



First upper limits on the 21 cm signal power spectrum from cosmic dawn from one night of observations with NenuFAR

S. Munshi, F G Mertens, L V E Koopmans, A R Offringa, B. Semelin, D. Aubert, R. Barkana, A. Bracco, S A Brackenhoff, B. Cecconi, et al.

► To cite this version:

S. Munshi, F G Mertens, L V E Koopmans, A R Offringa, B. Semelin, et al.. First upper limits on the 21 cm signal power spectrum from cosmic dawn from one night of observations with NenuFAR. *Astronomy & Astrophysics* - A&A, 2024, 681, pp.A62. <10.1051/0004-6361/202348329>. <insu-04480297>

HAL Id: insu-04480297

<https://insu.hal.science/insu-04480297v1>

Submitted on 27 Feb 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 CC BY 4.0 - Attribution - International License

First upper limits on the 21 cm signal power spectrum from cosmic dawn from one night of observations with NenuFAR

S. Munshi¹, F. G. Mertens^{2,1}, L. V. E. Koopmans¹, A. R. Offringa^{3,1}, B. Semelin², D. Aubert⁴, R. Barkana^{5,6,7}, A. Bracco⁸, S. A. Brackenhoff¹, B. Cecconi^{9,10}, E. Ceccotti¹, S. Corbel^{10,14}, A. Fialkov^{11,12}, B. K. Gehlot¹, R. Ghara¹³, J. N. Girard⁹, J. M. Grießmeier^{15,10}, C. Höfer¹, I. Hothi^{2,8}, R. Mériot², M. Mevius³, P. Ocvirk⁴, A. K. Shaw¹³, G. Theureau^{15,10,16}, S. Yatawatta³, P. Zarka^{9,10}, and S. Zaroubi^{13,1}

¹ Kapteyn Astronomical Institute, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands
e-mail: munshi@astro.rug.nl

² LERMA, Observatoire de Paris, Université PSL, CNRS, Sorbonne Université, 75014 Paris, France

³ ASTRON, PO Box 2, 7990 AA Dwingeloo, The Netherlands

⁴ Université de Strasbourg, CNRS, Observatoire astronomique de Strasbourg (ObAS), 11 rue de l'Université, 67000 Strasbourg, France

⁵ School of Physics and Astronomy, Tel Aviv University, Tel Aviv 69978, Israel

⁶ Institute for Advanced Study, 1 Einstein Drive, Princeton, New Jersey 08540, USA

⁷ Department of Astronomy and Astrophysics, University of California, Santa Cruz, CA 95064, USA

⁸ Laboratoire de Physique de l'Ecole Normale Supérieure, ENS, Université PSL, CNRS, Sorbonne Université, Université de Paris, 75005 Paris, France

⁹ LESIA, Observatoire de Paris, CNRS, PSL, Sorbonne Université, Université Paris Cité, 92195 Meudon, France

¹⁰ ORN, Observatoire de Paris, CNRS, PSL, Université d'Orléans, 18330 Nançay, France

¹¹ Kavli Institute for Cosmology, Madingley Road, Cambridge CB3 0HA, UK

¹² Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

¹³ ARCO (Astrophysics Research Center), Department of Natural Sciences, The Open University of Israel, 1 University Road, PO Box 808, Ra'anana 4353701, Israel

¹⁴ AIM, CEA, CNRS, Université Paris Cité, Université Paris-Saclay, 91191 Gif-sur-Yvette, France

¹⁵ LPC2E, Université d'Orléans, CNRS, 45071 Orléans, France

¹⁶ LUTH, Observatoire de Paris, Université PSL, Université de Paris Cité, CNRS, 92190 Meudon, France

Received 19 October 2023 / Accepted 3 November 2023

ABSTRACT

The redshifted 21 cm signal from neutral hydrogen is a direct probe of the physics of the early universe and has been an important science driver of many present and upcoming radio interferometers. In this study we use a single night of observations with the New Extension in Nançay Upgrading LOFAR (NenuFAR) to place upper limits on the 21 cm power spectrum from cosmic dawn at a redshift of $z = 20.3$. NenuFAR is a new low-frequency radio interferometer, operating in the 10–85 MHz frequency range, currently under construction at the Nançay Radio Observatory in France. It is a phased array instrument with a very dense uv coverage at short baselines, making it one of the most sensitive instruments for 21 cm cosmology analyses at these frequencies. Our analysis adopts the foreground subtraction approach, in which sky sources are modeled and subtracted through calibration and residual foregrounds are subsequently removed using Gaussian process regression. The final power spectra are constructed from the gridded residual data cubes in the uv plane. Signal injection tests are performed at each step of the analysis pipeline, the relevant pipeline settings are optimized to ensure minimal signal loss, and any signal suppression is accounted for through a bias correction on our final upper limits. We obtain a best 2σ upper limit of $2.4 \times 10^7 \text{ mK}^2$ at $z = 20.3$ and $k = 0.041 \text{ h cMpc}^{-1}$. We see a strong excess power in the data, making our upper limits two orders of magnitude higher than the thermal noise limit. We investigate the origin and nature of this excess power and discuss further improvements to the analysis pipeline that can potentially mitigate it and consequently allow us to reach thermal noise sensitivity when multiple nights of observations are processed in the future.

Key words. methods: data analysis – techniques: interferometric – dark ages, reionization, first stars

1. Introduction

The redshifted 21 cm signal arising out of the hyperfine transition of neutral hydrogen (HI) is a very sensitive probe of the astrophysical processes active during the early stages of cosmic evolution (Madau et al. 1997; Shaver et al. 1999). Cosmic dawn is a particularly interesting epoch in the early universe for which the 21 cm signal can potentially give us rich insights. This is the period when the gas aggregated in dark-matter halos

became dense enough to undergo gravitational collapse and form the first luminous objects in the universe. The two main processes believed to dominate during cosmic dawn are Lyman- α coupling and X-ray heating. The Lyman- α radiation emitted by the first sources couples the spin temperature of HI to the kinetic temperature of the intergalactic medium (IGM) through the Wouthuysen-Field effect (Wouthuysen 1952; Field 1958). Due to large spatial fluctuations in the distribution of early galaxies, this Lyman- α coupling is believed to have occurred

inhomogeneously, leading to observable 21 cm signal intensity fluctuations (Barkana & Loeb 2005). At later stages, X-ray radiation by stellar remnants heats the IGM (Venkatesan et al. 2001; Chen & Miralda-Escudé 2004; Pritchard & Furlanetto 2007), and this in turn increases the spin temperature of the 21 cm transition, which is coupled to the IGM temperature. Since ionization can still be largely neglected at the high redshifts of interest, these physical processes guide the magnitude and spatial distribution of the brightness temperature of the 21 cm signal against the cosmic microwave background radiation (CMBR) during this era. Radio-wave observations by interferometric arrays targeting the redshifts corresponding to cosmic dawn can directly probe the fluctuations in the three-dimensional brightness temperature distribution of the 21 cm signal during this epoch. Thus, these observations have the immense potential of providing us with crucial constraints on the properties of the IGM and early star formation histories (e.g., Furlanetto et al. 2006; Pritchard & Loeb 2012; Fialkov & Barkana 2014; Mebane et al. 2020; Reis et al. 2020; Gessey-Jones et al. 2022; Hibbard et al. 2022; Munoz et al. 2022; Monsalve et al. 2019; Adams et al. 2023; Bevins et al. 2024).

The prospect of probing cosmic dawn and the subsequent epoch of reionization (EoR) has led to many interferometric experiments devoting significant observing time to 21 cm cosmology programs. Most of these experiments have focused on the EoR, and over the years increasingly stringent upper limits on the redshifted 21 cm signal power spectrum from the EoR have been set by the GMRT¹ (Paciga et al. 2013), PAPER² (Kolopanis et al. 2019), MWA³ (Barry et al. 2019; Li et al. 2019; Trott et al. 2020), LOFAR⁴ (Patil et al. 2017; Mertens et al. 2020), and HERA⁵ (Abdurashidova et al. 2022). Many current and upcoming interferometric experiments are also trying to detect or place limits on the 21 cm signal from cosmic dawn, with some of them having already set upper limits on the 21 cm power spectrum at $z > 15$. The first attempt at measuring the 21 cm power spectrum during cosmic dawn was by Ewall-Wice et al. (2016), who used 6 h of MWA data to set upper limits of $\Delta_{21}^2 < 8.3 \times 10^7 \text{ mK}^2$ at $z = 15.35$ and $k = 0.21 \text{ h Mpc}^{-1}$, and $\Delta_{21}^2 < 2.7 \times 10^8 \text{ mK}^2$ at $z = 17.05$ and $k = 0.22 \text{ h Mpc}^{-1}$. Later, Yoshiura et al. (2021) used 5.5 h of MWA data to set an upper limit of $\Delta_{21}^2 < 6.3 \times 10^6 \text{ mK}^2$ at $k = 0.14 \text{ h Mpc}^{-1}$ and $z = 15.2$. The OVRO-LWA⁶ was used first by Eastwood et al. (2019) to provide upper limits of $\Delta_{21}^2 < 10^8 \text{ mK}^2$ at $z \approx 18.4$ and $k \approx 0.1 \text{ Mpc}^{-1}$ using 28 h of data and later by Garsden et al. (2021), who reported an upper limit of $\Delta_{21}^2 < 2 \times 10^{12} \text{ mK}^2$ at $z = 28$ and $k = 0.3 \text{ h Mpc}^{-1}$ using 4 h of data. Gehlot et al. (2019) set 2σ upper limits for the 21 cm power spectrum at $\Delta_{21}^2 < 2 \times 10^8 \text{ mK}^2$ at $z = 19.8\text{--}25.2$ and $k = 0.038 \text{ h cMpc}^{-1}$ using 14 h of data obtained using the LOFAR LBA⁷ system. The ACE program (Gehlot et al. 2020), which uses the LOFAR AARTFAAC⁸, placed 21 cm power spectrum upper limits of $\Delta_{21}^2 < 5 \times 10^7 \text{ mK}^2$ at $z = 17.9\text{--}18.6$ and $k = 0.144 \text{ h cMpc}^{-1}$ using 2 h of data. Second-generation experiments such as HERA (DeBoer et al. 2017) and SKA-Low (Koopmans et al. 2015) are

expected to have higher sensitivities in these redshifts, and SKA-Low has the potential to detect the 21 cm signal from cosmic dawn in tomographic images.

In a complementary approach, several single dipole experiments have attempted to detect the sky-averaged (global) 21 cm signal at high redshifts. Among these, the EDGES⁹ (Bowman et al. 2008) and SARAS¹⁰ (Singh et al. 2017; Subrahmanyam et al. 2021) are the only ones currently in the sensitivity range of a detection. In 2018, the EDGES collaboration reported an absorption trough in the global 21 cm signal at the redshifts corresponding to cosmic dawn (Bowman et al. 2018a). However, the feature detected by EDGES is unusually deep and flat (Hills et al. 2018; Bowman et al. 2018b), and it becomes necessary to introduce “exotic” nonstandard models of the 21 cm signal to explain it if it is cosmological. One such approach to modeling the EDGES signal involves supercooling of gas by scattering off cold dark-matter particles (e.g., Barkana 2018; Fialkov et al. 2018; Berlin et al. 2018; Muñoz & Loeb 2018; Liu et al. 2019). Another alternative class of models suggests that the background against which we observe the 21 cm signal could have a component in addition to the CMBR, thus explaining the unusual depth of the trough (e.g., Feng & Holder 2018; Ewall-Wice et al. 2018; Dowell & Taylor 2018; Fialkov & Barkana 2019). These models, in turn, predict stronger fluctuations in the 21 cm signal from cosmic dawn, implying that the 21 cm power spectrum could potentially be easier to detect. However, as of now, no independent confirmation of the EDGES result has been obtained by other experiments, and a recent measurement by the SARAS team claims to disprove the EDGES detection with a 95% confidence level (Singh et al. 2022).

In this work we present the first upper limits on the 21 cm signal power spectrum from cosmic dawn using NenuFAR¹¹ (Zarka et al. 2012), a radio interferometer nearing completion at the Nançay Radio Observatory in France. NenuFAR is a low-frequency phased array instrument that will have 1976 receiving antennas, with a very dense core and a large collecting area, making it extremely sensitive to the large scales necessary for constraining the 21 cm signal (Joseph 2012). It is one of the most sensitive instruments for cosmic dawn observations and is currently the most sensitive instrument below 50 MHz. In this study we performed an end-to-end analysis of a single night of observations made with an incomplete array composed of 79 active interferometric elements. We used an analysis method that adopts many of the strategies that have been developed and improved over the years during the evolution of the LOFAR EoR analysis pipeline (Chapman et al. 2012, 2013; Patil et al. 2014, 2016, 2017; Asad et al. 2015, 2016, 2018; Mertens et al. 2018; Gehlot et al. 2018, 2019; Offringa et al. 2019b; Hothi et al. 2021; Mevius et al. 2022; Gan et al. 2023). We demonstrate the performance of the NenuFAR calibration pipeline on a single night of observations made with the current incomplete array and present its first upper limits on the 21 cm signal power spectrum. This pilot analysis also reveals the limitations of the current processing pipeline and helps shape our strategy for overcoming them in the future.

The paper is organized as follows: Sect. 2 introduces the NenuFAR instrument and its Cosmic Dawn Key Science Program (CD KSP). In Sect. 3 we describe the data acquisition process and the preprocessing steps. The different steps

¹ Giant Metrewave Radio Telescope.

² Precision Array to Probe EoR.

³ Murchison Widefield Array.

⁴ Low-Frequency Array.

⁵ Hydrogen Epoch of Reionization Array.

⁶ Owen’s Valley Radio Observatory – Long Wavelength Array.

⁷ Low Band Antenna.

⁸ Amsterdam ASTRON Radio Transients Facility And Analysis Center.

⁹ Experiment to Detect the Global EoR Signature.

¹⁰ Shaped Antenna measurement of the background RAdio Spectrum.

¹¹ New Extension in Nançay Upgrading LOFAR.

in the data calibration procedure are described in Sect. 4. In Sect. 5 we describe an additional step of post-calibration flagging. Section 6 describes the power spectrum estimation, and Sect. 7 describes the first application of the machine learning (ML) Gaussian process regression (GPR) foreground removal technique (Mertens et al. 2024) to our data. Section 8 summarizes our results and presents the robustness tests that were performed to estimate the level of signal suppression due to the analysis pipeline. In Sect. 9 we discuss the limitations and future improvements. A short summary of the main conclusions from this analysis is presented in Sect. 10. Throughout this paper, a flat Λ cold dark matter cosmology is used, with the cosmological parameters ($H_0 = 67.7$, $\Omega_m = 0.307$) consistent with results from *Planck* observations (Planck Collaboration XIII 2016).

2. NenuFAR Cosmic Dawn Key Science Program

The NenuFAR¹² CD KSP¹³ (KP01, P.I.: L. V. E. Koopmans, B. Semelin and F. G. Mertens) is one of the programs that started in the early scientific phase of NenuFAR, which aims to detect or place stringent limits on the 21 cm power spectrum from cosmic dawn in the redshift range of $z = 15\text{--}31$ (Mertens et al. 2021). The project has already been allocated 1267 h of observations, which started in September 2020 and are still going on with an array that is progressively more sensitive as more core and remote stations are being added. With NenuFAR's sensitivity, detection of the 21 cm signal predicted by EDGES-inspired exotic models is possible within 100 h of integration. Moreover, with 1000 h of integration, NenuFAR is expected to reach a thermal noise sensitivity comparable to the levels of the signal predicted by more standard models (Mesinger et al. 2011; Murray et al. 2020; Semelin et al. 2023).

2.1. The instrument

NenuFAR (Zarka et al. 2015, 2020) is a large low-frequency radio interferometer. The primary receiving antennas for NenuFAR are crossed inverted V-shaped dual polarization dipoles, similar to the antenna elements employed by the LWA (Hicks et al. 2012), but equipped with an original custom-made low noise amplifier (Charrier & the CODALEMA Collaboration 2007). In total, 19 such antennas, arranged in a hexagonal array, form a single mini-array (MA), whose station beam is analog steerable in different directions in the sky. The hexagonal MA has a six-fold symmetry, resulting in considerable grating lobes in an individual MA primary beam. In order to reduce the level of these grating lobes, the MAs are rotated with respect to each other by angles that are multiples of 10 degrees, with only 6 such non-redundant orientations being necessary due to their hexagonal symmetry. At completion, NenuFAR will have 96 MAs distributed in a 400 m diameter core and 8 remote MAs located at distances up to a few kilometers from the center. Figure 1 presents a schematic representation of the NenuFAR configuration used for obtaining the data used in this analysis. The dense core of NenuFAR with a large number of short baselines makes it extremely sensitive to large scales and especially well suited to constraining the 21 cm signal. The longer baselines involving the remote stations, combined with pointing observations and Earth rotation synthesis, yield a good uv coverage out to several kilometers (3.5–5 km at completion, currently 1.4 km)




Fig. 1. Schematic representation of the NenuFAR array configuration used in the observation.

and thus enable efficient point source foreground modeling and subtraction.

2.2. Observations and data selection

The primary target field for the NenuFAR CD KSP is the north celestial pole (NCP). The NCP is a particularly favorable field since it is visible at night throughout the year, allowing long observations on winter nights. The NCP is in a fixed direction, thus requiring no beam or phase tracking with time. Moreover, extensive sky models of the NCP, made using deep integrations with LOFAR, are readily available as a byproduct of the LOFAR EoR Key Science Project (KSP) and these aid the analysis and calibration efforts. The NCP observations with NenuFAR are preceded or succeeded by a 30 min observation of a calibrator – Cygnus A (Cyg A) or Cassiopeia A (Cas A), depending on which one of the two sources is at a high enough altitude to yield good quality observations.

Phase I observations of the NCP with NenuFAR started in September 2019 with the primary goal of understanding the instrument and monitoring its performance. Deep integrations of the NCP field commenced with the start of Phase II observations in September 2020 and as of May 2022, 1080 h of NCP data have been gathered, with the newer observations being carried out with more core and remote stations. In July 2021 the number of core MAs was increased from 56 to 80 and in September 2021, the third remote MA was installed. The night used in this analysis was selected on the basis of considerations related to low ionospheric activity among the observations that were available with three remote MAs active as of May 2022 when this pilot survey was initiated. Calibration solutions for the calibrator (Cyg A or Cas A) observations were obtained for all these nights and the temporal fluctuations in these solutions were used as a proxy for ionospheric activity. The trend in the temporal fluctuations as a function of observation nights was verified to be roughly consistent with the trend in the total electron content values for these nights obtained independently from GPS measurements using RMextract¹⁴ (Mevisius 2018). The night of 12–13 December 2021 had mild ionospheric activity and coincidentally relatively low level of flagged data due to

¹² <https://nenufar.obs-nancay.fr/en/homepage-en/>

¹³ <https://vm-weblerna.obspm.fr/nenufar-cosmic-dawn/>

¹⁴ <https://github.com/lofar-astron/RMextract/tree/master>

Table 1. Observation specifications for the raw data (L0) used in the analysis.

Parameter	Value
Core MAs	76
Remote MAs	3
Spectral window	61.1–72.2 MHz
Bandwidth	11.1 MHz
Spectral resolution	3.05 kHz
Sub-bands	59
Channels per sub-band	64
Maximum baseline	1.46 km
Target	NCP
Calibrator	Cas A
Target start time (UTC)	17:33 12/12/2021
Target duration	11.4 h
Calibrator start time (UTC)	17:00 12/12/2021
Calibrator duration	30 min
Time resolution	1 s

radio-frequency interference (RFI). This night was selected for the current analysis.

The 11.4 h NCP observation was preceded by a 30 min observation of Cas A, which was used for bandpass calibration. Though the observations are recorded in 4 spectral windows covering a wide frequency range of 37–85 MHz, we selected a single spectral window ranging from 61 to 72 MHz for this analysis on the basis of superior RFI statistics (Appendix A). The main observation specifications are summarized in Table 1. Figure 2 shows the baseline coverage of the observation.

3. Preprocessing

Observational data obtained with NenuFAR undergo several routine steps of preliminary processing, which prepare it for further stages of precise calibration and foreground subtraction. These steps are described in this section.

3.1. Data acquisition

The NenuFAR correlator is a replica of the COBALT-2 correlator used by LOFAR (Broekema et al. 2018) and is called NenuFAR Imager Correlation Kluster Elaborated from LOFAR’s (NICKEL). During a NenuFAR observation, the two linear polarization components of the signal from all MAs reach the receiver where they are sampled at a frequency of 200 MHz, followed by channelization into 512 sub-bands of 195.3 kHz width each using a polyphase filter bank. The data for each second is digitally phased to the target direction and cross correlated across the different polarizations to yield full polarization visibilities in 384 sub-bands with a total bandwidth of 75 MHz. Each sub-band is further channelized into 64 channels with a 3.05 kHz spectral resolution. The raw correlated visibilities are finally exported into a measurement set (MS) format.

3.2. RFI mitigation, bandpass calibration, data averaging, and compression

The preprocessing of NenuFAR data is performed using the `nenucal-cd` package¹⁵, which in turn uses the various tasks for

¹⁵ <https://gitlab.com/flomertens/nenucal-cd>




Fig. 2. Baseline coverage in the uv plane for the NCP observation. The baseline tracks for the core-remote baselines are indicated with different colors for the baselines involving the three remote MAs. For visual clarity, only the U, V points are plotted here, not the $-U, -V$ points. The full coverage including the $-U, -V$ points is much more symmetric.

calibration, flagging and averaging included in the default preprocessing pipeline (DP³; van Diepen et al. 2018) used for analysis of LOFAR data. NenuFAR visibility data during preprocessing is available at three data levels: L0, L1, and L2 in the order of decreasing time resolution, frequency resolution, and data volume. The raw visibilities output by the correlator (the L0 data) have a time resolution of 1 s and frequency resolution of 3.05 kHz with 64 channels per sub-band. To capture intermittent and narrowband RFI, the software AOFlogger (Offringa et al. 2012) is used directly at this highest time and frequency resolution to flag the data affected by strong RFI. To avoid the edge effects of the polyphase filter used in channelization, two channels at both ends of the sub-band are flagged, and the remaining data are averaged into 4 s time resolution and 15.26 kHz frequency resolution with 12 frequency channels per sub-band. This is followed by data compression with Dysco (Offringa 2016) to produce the L1 data, leading to the reduction in data volume by a factor of ≈ 80 compared to the L0 data. Though this compression introduces some noise into the data, it has been verified through simulations that the compression noise is significantly lower than the expected 21 cm signal. Moreover, the noise is uncorrelated across separate observations taken on different days, and hence averaging the data across different days decreases this compression noise just like thermal noise. The L1 data of the targeted observation of Cas A were calibrated using the software DDECa1 (described in Sect. 4) against a sky model of Cas A consisting of point sources and Gaussians created previously using LOFAR LBA observations. The parameters used in the bandpass calibration with DDECa1 are listed in Table 2. The bad calibration solutions were flagged based on a threshold-clipping algorithm and the remaining solutions were averaged in time to produce one solution per MA, per channel, and per polarization. These bandpass solutions were then applied to the L1 NCP data. The bandpass calibration sets the approximate amplitude scale of the target visibilities and more importantly, accounts for the cable reflections that can produce rapid frequency fluctuations in the data. The results of bandpass calibration for a few example MAs are shown in Appendix B.1. The bandpass calibration is followed by another step of RFI flagging with AOFlogger and averaging to a frequency resolution of 61.03 kHz with three channels per sub-band, to get the L2 data product. These data are now suitable for the next steps of calibration.

Table 2. Calibration parameters used in DDECa1 at different stages of the analysis.

Parameter	Bandpass calibration	A-team subtraction	3C subtraction	NCP subtraction
Input data level	L1	L2	L2	L3
Output data level	L2	L2	L3	L4
Mode	Diagonal	Diagonal	Diagonal	Diagonal
Time interval	4 s	2 min	4 min	8 min
Frequency interval	15.3 kHz	195.3 kHz	195.3 kHz	195.3 kHz
Directions	Cas A	Four A-team sources and NCP	Seven 3C sources and NCP	7 Clusters
Smoothness constraint	None	2 MHz	2 MHz	6 MHz
uv range	$>15\lambda$	$>20\lambda$	$>40\lambda$	$>40\lambda$

Notes. L1 data have a 15.26 kHz frequency resolution while L2, L3 and L4 data have 61.03 kHz frequency resolution. The time resolution for L1, L2, L3 and L4 data levels is 4 s.

4. Calibration

The dataset after the preprocessing was divided into 13 time segments (twelve 52 min segments and one 56 min segment), which were processed in parallel, in 13 computational nodes of the DAWN cluster (Pandey et al. 2020). This decreased the computation time and enabled us to conduct multiple calibration runs that allowed us to optimize the calibration settings for this dataset.

Calibration of NenuFAR data is done using DDECa1, a calibration software that is part of the DP³. DDECa1 employs a directional solving algorithm similar to the scalar algorithm described by Smirnov & Tasse (2015). In addition to the fulljones mode in which all four elements of the gain matrix are solved for, DDECa1 can also be used in the diagonal mode in which the off-diagonal elements are set to zero. A detailed description of the fulljones algorithm of DDECa1 is presented by Gan et al. (2023). DDECa1 also has a feature to solve for spectrally smooth solutions by applying constraints to the solutions, using a Gaussian kernel of a chosen full width at half maximum (FWHM). The calibration parameters that were used in DDECa1 at different stages of the data processing are listed in Table 2.

The frequency smoothing kernel prevents rapid nonphysical frequency fluctuations in the calibration solutions and ensures that the gain solutions do not over-fit the thermal noise and sources that are not part of the incomplete sky model used in calibration, including the 21 cm signal (Mouri Sardarabadi & Koopmans 2019; Mevius et al. 2022). The smoothing kernel FWHM was chosen by performing multiple signal injection tests with different kernel widths to find a compromise between maximizing the accuracy of source subtraction and minimizing the signal loss (described in Sect. 8.1). At every stage of calibration, we applied a baseline cut of at least 20λ and only used longer baselines for the calibration. This was done because the shorter baselines are especially sensitive to the diffuse Galactic emission, which is not present in the sky model used for calibration. If these baselines were used in calibration, this additional diffuse emission would affect the gain solutions and could reappear elsewhere in the final image or on different spatial and frequency scales (Patil et al. 2016). A similar calibration strategy is used by the LOFAR EoR KSP, but the difference is that LOFAR uses a much higher baseline cut (250 λ) and computes the power spectrum using only the shorter baselines (between 5 and 250 λ) that are not used in calibration. However, this is not feasible with NenuFAR for all the calibration steps since there are not enough long baselines to get reliable calibration solutions for all MAs, and the baselines used for generating power spectra are also used in one of the calibration steps. This creates a risk of signal suppression during this calibration step and this was assessed during signal injection tests (described in Sect. 8.1). Figure 3 summarizes all the different steps in the processing pipeline used in this analysis, except for the signal suppression tests.

4.1. Bright source subtraction and direction-independent correction

The sources Cas A, Cyg A, Taurus A (Tau A), and Virgo A (Vir A), collectively referred to as the A-team, are the brightest astronomical radio sources in the northern celestial hemisphere at low radio frequencies. In addition to these, there are several other radio sources in the northern sky with a high enough flux to have a significant impact on the NCP field through their point spread function (PSF) sidelobes. The spectra of most of these sources are steep, which means that they become extremely bright at the low frequencies that we are interested in. The result is that even though some of these sources are far away from the NCP and are attenuated by the primary beam, the PSF sidelobes of these sources leave a considerable imprint on the NCP field. Naturally, precise subtraction of these bright sources and mitigation of their PSF sidelobes is crucial for our analysis.

4.1.1. A-team subtraction and direction-independent correction

The NenuFAR primary beam model is not yet a part of DDECa1. Since the four A-team sources are well away (at least 30 degrees) from the NCP, we do not expect the MA gains in the direction of these sources to be the same as those in the direction of the NCP, mainly because of the primary beam. Hence, as these sources move through the stationary primary beam, the MA gain solutions in the direction of these sources will carry an imprint of the primary beam pattern. It is, therefore, necessary to perform direction-dependent (DD) calibration in order to obtain separate time and frequency-dependent gain calibration solutions for the NCP field and the individual A-team sources.

We used intrinsic sky models for all the A-team sources, obtained previously using the LOFAR LBA and consisting of point sources (delta-functions) and Gaussians. Not all the A-team sources are above the horizon for all 13 time segments into which we have divided the dataset. So for a given data segment, if a particular A-team source is below the horizon for most of the time, we excluded it from the sky model. This is illustrated in Fig. 4. For the NCP, we used an intrinsic model obtained from the global sky model¹⁶ that was created using data from The GMRT Sky Survey (TGSS; Intema et al. 2017), NRAO VLA Sky Survey (NVSS; Condon et al. 1998), Westerbork Northern Sky Survey (WENSS; Rengelink et al. 1997), and VLA Low-frequency Sky Survey (VLSS; Cohen et al. 2007). For each data segment, the primary beam¹⁷, averaged over all stations, was

¹⁶ <https://lcs165.lofar.eu/>

¹⁷ A preliminary model of the NenuFAR primary beam is available from the package nenupy: <https://github.com/AlanLoh/nenupy/>.

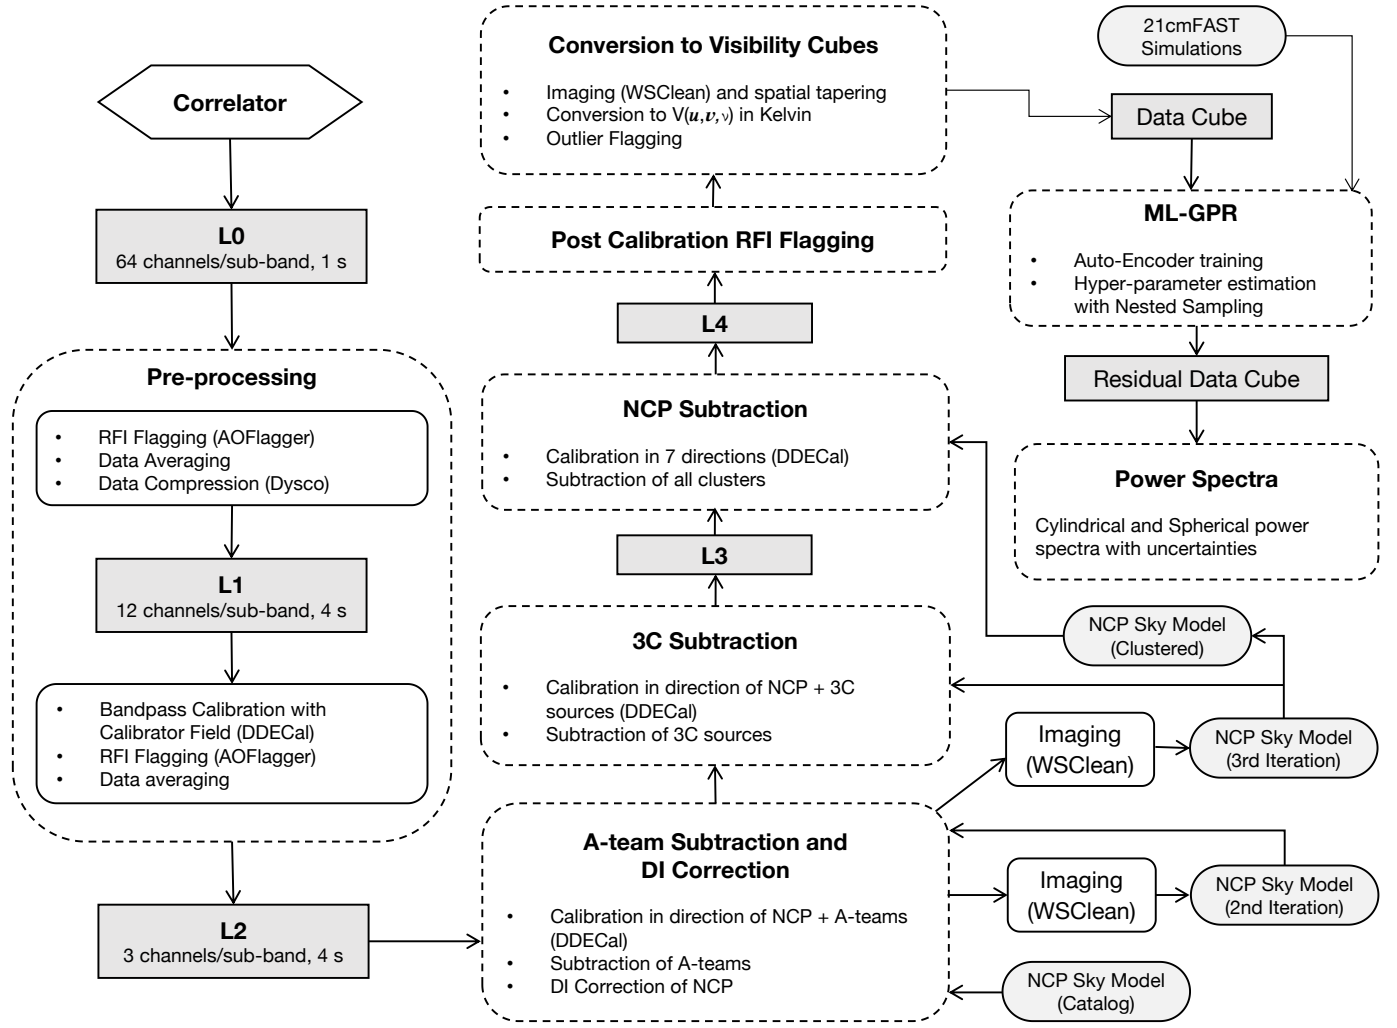


Fig. 3. Data processing pipeline with the different steps of preprocessing, calibration, foreground subtraction, and power spectrum estimation. The calibration settings used in the different steps in this pipeline are summarized in Table 2. Note that the signal injection tests performed on the calibration and foreground removal steps are not described in this flowchart (see Fig. 12).

computed at three equally spaced time points within the segment (corresponding to 1/4, 1/2 and 3/4 of the time duration of the segment) and averaged together to give the station and time averaged beam. The intrinsic sky model components were then attenuated by the beam values at the positions of the components. These “apparent” sky models (as seen by NenuFAR) for the different A-team sources and the NCP were subsequently put in separate sky-model-component clusters to get the final apparent sky model for each data segment. Each such cluster forms a direction for which a separate calibration solution is solved. It should be noted that this kind of attenuation by a simulated beam model accounts for the rotation of the primary beam against the sky on timescales longer than the duration of the time segments (since we used a single sky model for a segment), and up to the accuracy of the preliminary simulated beam model that was used to obtain the apparent sky model used in calibration. The inaccuracy of this preliminary beam model and the spatially averaged effect of the rotation of the beam on timescales as low as the gain solution intervals will be accounted for by the gain calibration solutions. However, since the NenuFAR primary beam model is currently not integrated into DDECal, the differences between the actual sky and the apparent sky model within a solution time

interval and inside one direction cannot be accounted for and will lead to errors in the calibration.

DDECal was used to calibrate the L2 data against the apparent sky model for each time segment. It returns separate calibration solutions for the NCP and each of the A-team sources per MA, polarization, time, and frequency step. Next, the visibilities corresponding to each of the A-team sky models were predicted and the calculated gain solutions were applied to the visibilities. These corrupted predicted visibilities were then subtracted from the data. This removes the A-team from the data if both the source model and the calibration are perfect. The gain solutions in the direction of the NCP were then applied to the data. We note that this is a single calibration solution for the entire NCP field, so this is similar to a direction-independent (DI) correction step if only the NCP field is concerned and this sets the absolute flux scale of the data. This completes the A-team subtraction and DI gain correction on the NCP field using a catalog sky model. The final parameters used for this calibration step were selected after multiple calibration trial runs (described in Appendix C). We find that including all four A-team sources when they are above the horizon, flagging MAs 6, 18, 62, 64, 72, and using diagonal mode gives us the optimal results for this dataset.




Fig. 4. Altitude- and beam-attenuated flux values for the four A-team sources, Cyg A, Cas A, Tau A, and Vir A, that are subtracted in the analysis. The vertical dotted lines demarcate the 13 different segments into which the data are divided. The solid lines indicate the altitude of the A-team and are to be read using the scale on the left-hand side axis. The dashed lines connect the intrinsic flux values for the A-team attenuated by the averaged preliminary primary beam model value for each data segment and are to be read using the axis on the right-hand side. The four horizontal bands on top indicate the data segments in which each A-team was included in the sky model and hence subtracted through the method described in Sect. 4.1.1.

The next step is to make a sky model of the NCP field using this calibrated NenuFAR data, which will serve as an updated version of the preliminary catalog sky model. This is similar to a self-calibration loop in which the DI gains and the sky model are iteratively updated. First, the A-team subtracted and DI-corrected data were used to make a deconvolved multifrequency synthesis (MFS) image with WSClean (Offringa et al. 2014). The parameters used in making the cleaned image are given in Table 3 under the column “High Res Clean”. The image was made using multi-scale clean (Offringa & Smirnov 2017) in which both point sources and Gaussians are part of the model. The clean model is given as an output by WSClean in the form of a component list, and this provides an updated sky model of the NCP field as observed by NenuFAR. This is an apparent sky model of the NCP because the NenuFAR primary beam, averaged over the observation, is not corrected for in the image. We then used this updated sky model to calibrate the L2 dataset, using the same calibration settings, which gave the best results for the A-team subtraction and DI correction using the catalog sky model. The only difference here is that for each data segment, in place of the intrinsic NCP sky model attenuated by the averaged primary beam, we used this apparent sky model of the NCP. This A-team subtraction and DI correction run provides slightly improved results compared to when the catalog NCP sky model is used, with the final residual image from WSClean having 1% lower rms. The corresponding clean model from a WSClean run on the DI corrected data now provides a third iteration model of the NCP field and has a few more sources than the previous sky model. The upper panel of Table 4 lists the three different sky models used at different stages: the initial catalog sky model from the Global Sky Model (“Initial”), the updated model after A-team subtraction and DI correction with the catalog model (“2nd Iteration”) and the third iteration model after A-team subtraction and DI correction with the 2nd Iteration model (“3rd Iteration”).

Table 3. Imaging parameters used for producing the high-resolution clean images (High Res Clean) and the dirty image cubes for the power spectrum (Dirty Cube).

Parameter	High Res Clean	Dirty Cube
Weighting	Briggs (Robust = −0.1)	Natural
Pixel size	3′	6′
Image size	600 pixels	400 pixels
uv range	>20 λ	5–100 λ
Channels	12 (joined)	177 (not joined)
Algorithm	Multiscale (Hogbom)	...
Auto mask ^(a)	3	...
Auto threshold ^(b)	1	...
Mgain ^(c)	0.6	...
Spectral pol ^(d)	2 terms	...

Notes. The values of parameters not applicable for making dirty images are specified as “...”. ^(a)Threshold at which a mask is constructed from the obtained components. ^(b)Cleaning threshold based on the local rms. ^(c)Cleaning gain for major iterations. ^(d)Number of terms in the polynomial fit over frequency to each clean component.

Table 4. NCP sky model used at different analysis stages.

Sky model	Direction	Components	Total flux (Jy)	Max flux (Jy)
Initial ^(*)	NCP	2787 ^(*)	6379 ^(*)	83 ^(*)
2nd iteration	NCP	1608	726	40
3rd iteration	NCP	1622	733	39
Clustered	Cluster 1	191	221	5
	Cluster 2	182	144	7
	Cluster 3	163	127	39
	Cluster 4	186	88	12
	Cluster 5	107	58	4
	Cluster 6	124	36	8
	Cluster 7	95	33	3

Notes. ^(*)The “Initial” model is an intrinsic sky model, while all other models are apparent sky models.

To assess the level of residuals at each stage of calibration and subtraction, the visibilities at each stage were imaged using WSClean and the MFS images of the data at different stages are presented in Fig. 5. The left panels of Fig. 5 show dirty wide-field images, which were made with the same parameters reported in the “High Res Clean” column of Table 3, except with an image size of 2400 pixels and without the last five parameters, which are needed only for cleaning. The wide-field image in the top-left panel clearly shows the level of contamination from the A-team sources (particularly Cas A and Cyg A), which are well subtracted in the A-team subtraction step (mid-left panel). The narrow-field cleaned images (right panels of Fig. 5) show how the PSF sidelobes of the A-team running through the NCP field, are also well subtracted. To assess the impact of the PSF sidelobes of wide-field sources on the uv plane, the standard deviation of the data in system equivalent flux density (SEFD) units was calculated from the frequency-channel-differenced noise for each uv cell in the gridded data cube (the gridding procedure is described in Sects. 6.1 and 6.2). This was repeated at different stages of source subtraction and the results are presented in




Fig. 6. Data standard deviation in SEFD units calculated from the channel-differenced noise at each uv cell of the Stokes- I data cube. The three panels correspond to the data after preprocessing (left), after A-team subtraction (middle), and after 3C subtraction (right). The dotted and dashed circles correspond to the 20λ and 40λ limits, respectively, between which the power spectrum is finally computed.

Fig. 6. It should be noted that only the central 16 degrees of the field was used in constructing these data cubes as well as the final power spectra. This figure illustrates the impact of the PSF sidelobes of the A-team sources, which appear as straight lines in the spatial frequency domain¹⁸ and are well subtracted using this method. However, the subtraction of Cyg A is not perfect and the PSF sidelobes of the Cyg A residuals can still be seen after the A-team subtraction step (middle and right panels of Fig. 6).

4.1.2. Subtraction of 3C sources

In addition to the A-team, there are several other compact sources in the northern sky that have fluxes of hundreds of janskys at the low frequencies of interest in our observations. In general, most of these sources will be highly attenuated by the primary beam and hence their PSF sidelobes will not be bright enough to contaminate the NCP field. However, in our data, we find that the PSF sidelobes of seven sources¹⁹ do have a strong effect on the NCP field. This is because all these sources pass through the grating lobes of the MA beam, and the primary attenuation that their fluxes get is due to the dipole beam and not due to the array factor of the MA, which is very high at the grating lobes. From Fig. 6, it is clear that the PSF sidelobes of the 3C sources have a strong signature on the uv plane. To subtract these sources, we performed a DD calibration in eight directions: the seven 3C sources and the NCP. The time solution interval of 4 min was selected to be small enough to allow the gain solutions to model the grating lobes and large enough such that there is sufficient signal-to-noise ratio (S/N) to obtain reliable calibration solutions. The sky model used for each of these sources is a single delta function at the location of the source, obtained from catalogs. The sources were included in the sky model for all segments where they are above 10 degrees altitude. This results in 3C 438 being excluded for the last seven segments, 3C 452 for the last six segments, 3C 48 for the last two segments, and 3C 84 for the last segment. The three other sources were included in the calibration for all segments. The third iteration NCP sky model made from the A-team subtracted and DI-corrected data was used as the model in the direction

of the NCP field. Once the calibration solutions were obtained, the visibilities for the 3C sources were predicted, the calculated gains in the respective direction were applied and the corrupted predicted visibilities were subtracted from the data. This largely removes these seven sources and their PSF sidelobes from the data. The impact of this step on the image and the uv plane are shown in Figs. 5 and 6, respectively. The impact of the PSF sidelobes is particularly prominent in Fig. 6 and the sidelobes are seen to be well subtracted using this method (right panel). The bright source subtracted and DI-corrected data will hereafter be referred to as the L3 data.

4.2. NCP subtraction

The sources in the NCP field can now be subtracted from the L3 data. Here it is important to account for the contaminating effects of ionospheric phase shifts and a non-axisymmetric central lobe of the primary beam of NenuFAR, both of which are expected to introduce DD effects on the data within the field of view (FOV). Currently, it is not possible to solve for small timescale ionospheric effects because the solution time interval needed to have a high enough S/N to yield reliable calibration solutions is much longer than the timescale on which these effects occur on short baselines (a few seconds to 1 min, following Vedantham & Koopmans 2015). Ideally one would want to obtain separate gain matrices for all the different components in the NCP sky model. However, this is infeasible because the flux for each component is not sufficient to yield reliable calibration solutions and it would also be extremely computationally expensive. Additionally, solving for a large number of directions would also increase the degrees of freedom, thus increasing the risk of absorption of the noise and the 21 cm signal. Therefore, we used an alternative approach, employed by the LOFAR EoR KSP, in which the NCP sky model is divided into several clusters and a separate calibration solution is obtained for each cluster under the assumption that DD effects do not vary strongly over each cluster (Kazemi et al. 2013). For this purpose, we used the third iteration sky model made from the data after DI correction (Table 4). The sources outside the primary beam beyond 13.4 degrees from the phase center were first removed from the sky model since they have a low flux density. The remaining sources were divided into seven clusters using a k -means-based clustering algorithm. The number seven was chosen to maximize

¹⁸ These lines are perpendicular to the ripples of the PSF sidelobes of sources located far away from the phase center.

¹⁹ 3C 123, 3C 134, 3C 111, 3C 84, 3C 48, 3C 452, and 3C 438.

the number of clusters and retain a sufficiently large S/N for the faintest cluster. The flux in the faintest cluster is 33 Jy, which is at a 6σ level when compared to the expected thermal noise in the visibilities for the chosen 8 min and 183.1 kHz solution interval. The details of the clustered sky model are presented in the bottom half of Table 4 and the clusters are shown in the bottom-left panel of Fig. 5. The L3 data were then calibrated against this clustered sky model, with separate solutions calculated for the direction of each cluster. Once the calibration solutions were obtained, the visibilities corresponding to each cluster were predicted, then corrupted by the calculated gains and finally subtracted from the data. We performed multiple runs of the DD calibration in conjunction with signal injection simulations using a fiducial simulated 21 cm signal (described in Sect. 8.1) to converge to the final calibration settings, which gives the best compromise between optimal results in calibration solutions and images and minimum signal suppression. The data after NCP subtraction will hereafter be referred to as the L4 data.

In the bottom-right panel of Fig. 5 we can see that most sources are subtracted well until we reach the confusion noise level (of 0.9 Jy beam^{-1} ²⁰) beyond which we do not have sky model components and hence they are not subtracted. The rms of the fluxes in the image pixels in the central 5 degrees radius is $0.15 \text{ Jy beam}^{-1}$, with a maximum and minimum flux of $1.05 \text{ Jy beam}^{-1}$ and $-0.45 \text{ Jy beam}^{-1}$, respectively. The remaining poorly subtracted sources at the edge of the clusters are likely due to the fact that the DD effects are not accounted for on scales smaller than the cluster sizes. As a result, the average gain solution for each cluster does not correctly represent the gains in the direction of these sources and is a complex flux-weighted average over the cluster.

5. Post-calibration RFI flagging

The data after the different stages of calibration and point source subtraction still have significant residual power at high delay (Fourier dual of frequency) modes. Strong features are seen in the delay and fringe rate (Fourier dual of time) power spectra of multiple baselines (Fig. 7), which are likely due to near-field RFI.

To identify local sources of RFI, we constructed a near-field image (Paciga et al. 2011) from the L4 data. This image was created by coherently summing all visibilities after their amplitudes had been set to unity and assuming that phase differences are only due to sources on the ground. Thus, the amplitude in the image does not necessarily correspond to the strength of the local RFI source. It only allows us to pinpoint its location. The near field image reveals that the buildings within the NenuFAR core, housing the electronic containers, are producing significant RFI (left panel of Fig. 8). We performed a comprehensive study of the impact of such near-field sources of RFI on the power spectrum. The details of this study will be published in a separate paper. To mitigate this near-field RFI, we adopted a simple approach of examining the delay power spectra of individual baselines and selecting baselines that are most severely affected by the RFI source. The RFI source at the two buildings within the array is seen to have a periodic fluctuation in intensity with a periodicity of 18 min, leading to a well-defined signature in the delay-fringe rate space (Fig. 7). The origin of this periodicity is still unknown. Any baseline showing a similar feature in the delay-fringe rate power spectra can be easily identified and flagged. This can potentially allow us to filter it out from the data instead of manu-

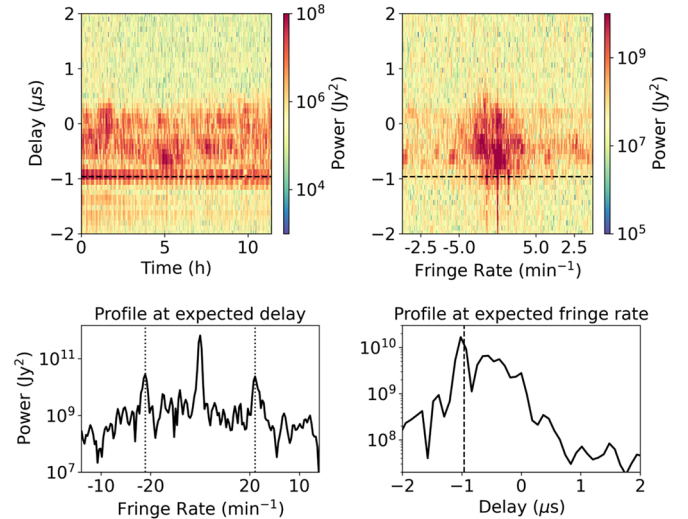


Fig. 7. Delay and fringe rate power spectra of an example baseline where a strong RFI feature is seen. Delay vs. time plot (top-left) and delay vs. fringe rate plot (top-right) for the Stokes-*I* data after NCP subtraction. The dashed horizontal lines indicate the expected delay corresponding to the local RFI source at the building. Bottom left: power as a function of fringe rate at the expected delay for the local RFI source. The vertical dotted lines correspond to a fringe rate of 18 min^{-1} . Bottom right: power as a function of delay at the expected fringe rate of 18 min^{-1} . The vertical dashed line is the expected delay.

ally flagging entire baselines, but we defer this to future analysis. We examined the Stokes-*I* delay power spectra for all baselines manually and identified those that show a strong periodic feature. For each such baseline, a feature is always also present in the Stokes-*V* delay spectra, where it is usually more prominent because the power from sky sources is negligible. The histogram of the MAs contributing to the flagged baselines reveals that baselines involving MAs close to the building are more strongly affected by the RFI (middle panel of Fig. 8). In addition to this periodic RFI signature, the delay spectra of many baselines also show other unusually strong features beyond the horizon delays. The histogram of the MAs contributing to these baselines shows a peak near the northeast of the array (right panel of Fig. 8) in a region where we see a strong RFI source in the near-field image. These baselines were flagged as well. In this process, we flagged 7.5% of the data. After this additional flagging step, the data are deemed sufficiently clean, calibrated and sky model subtracted, and can be used for residual foreground subtraction using ML-GPR and power spectrum estimation.

6. Power spectra estimation

The estimation of power spectra was done using the power spectrum pipeline `pspipe`²¹, which is also used to generate the power spectrum for the LOFAR EoR KSP.

6.1. Image cubes

The data need to be gridded in a *uvw* grid in order to construct the power spectra. For this purpose, `WSClean` was used to construct dirty image cubes from the visibilities. `WSClean` uses a *w*-stacking algorithm while making the image, which accounts for wide-field effects due to the *w*-term. The imaging parameters

²⁰ This is estimated using Eq. (6) of van Haarlem et al. (2013).

²¹ <https://gitlab.com/flomertens/pspipe>




Fig. 8. Local RFI sources at the NenuFAR site and their impact on the MAs. Left: normalized near-field image of the data. The locations of the buildings housing the electronic containers are indicated with black triangles, and the MAs are indicated with gray circles. Center: histogram of the MAs contributing to the baselines that exhibit a periodic amplitude fluctuation in time. Right: histogram of the MAs contributing to all other baselines, which have been flagged based on unusual features seen in the delay power spectra. The size and color of the circles in the middle and right panels indicate the number of baselines involving that MA that have been flagged.

are specified in the column named “Dirty Cube” in Table 3. Separate image cubes were made for alternating odd and even time samples and these were used at a later stage in the estimation of the noise level in the data. A Hann filter with an FOV of 16 degrees was applied to the dirty image cubes in the image plane in order to suppress primary-beam effects on sources far away from the phase center as well as aliasing artifacts. These dirty image cubes were subsequently used to obtain the power spectra.

6.2. Conversion to visibility cubes

The gridded image cubes (I^D) produced as described above were first converted from units of Jy/PSF to units of Kelvin using the relation (Offringa et al. 2019a; Mertens et al. 2020)

$$T(l, m, \nu) = \frac{10^{-26} c^2}{2k_B \nu^2 \delta_l \delta_m} \mathcal{F}_{u,v}^{-1}[\mathcal{F}_{l,m}[I^D] \oslash \mathcal{F}_{l,m}[I^{\text{PSF}}]], \quad (1)$$

where k_B is the Boltzmann constant, δ_l, δ_m are the pixel sizes in the l and m directions (in units of radians), respectively, $\mathcal{F}_{l,m}$ denotes a Fourier transform and $\mathcal{F}_{u,v}^{-1}$ is its inverse, I^{PSF} is the point spread function and \oslash is the element-wise division operator. These image cubes in the (l, m, ν) space were then Fourier transformed in the spatial direction to get the data cube in the uv -space: $\tilde{T}(u, v, \nu)$. A uv range of $15\text{--}50\lambda$ was chosen for further analysis, and the remaining data were nulled for all frequencies. This avoids baselines shorter than 15λ , which could have strong contaminating effects from mutual coupling, and those longer than 50λ , which have a higher thermal noise contribution. We note that though we constructed the final spherical power spectra in the $20\text{--}40\lambda$ baseline range, here we tried to retain as many baselines as possible since it allows better foreground modeling and removal through ML-GPR. A final outlier rejection was done on the gridded visibility cubes to flag potentially remaining low-level RFI using a simple threshold-clipping method. The uv cells were flagged based on outliers in the uv weights, Stokes- V variance, and channel-differenced Stokes- I

variance. About 15% of the uv cells were flagged in this procedure.

6.3. Cylindrically and spherically averaged power spectra

The data cube $\tilde{T}(u, v, \nu)$ at any stage of calibration and foreground subtraction can now be Fourier transformed along the frequency axis (after applying a Blackman-Harris filter to suppress aliasing) to obtain $\tilde{T}(u, v, \eta)$, where η is the Fourier dual of frequency, commonly referred to as the delay. The corresponding power spectrum as a function of the wave numbers k_l, k_m, k_{\parallel} is given by

$$P(k_l, k_m, k_{\parallel}) = \frac{\Omega_{lm} B_{\text{bw}} D_M^2(z) \Delta D}{\langle A_{\text{pb}}^2(l, m) A_w^2(l, m) \rangle \langle B_w^2(\nu) \rangle} |\tilde{T}(u, v, \eta)|^2. \quad (2)$$

The wave modes are related to u, v, η as (Morales & Hewitt 2004; McQuinn et al. 2006)

$$k_l = \frac{2\pi u}{D_M(z)}, \quad k_m = \frac{2\pi v}{D_M(z)}, \quad k_{\parallel} = \frac{2\pi H_0 \nu_{21} E(z)}{c(1+z)^2} \eta, \quad (3)$$

where $\nu_{21} = 1420\text{ MHz}$, H_0 is the Hubble constant, $E(z)$ is the dimensionless Hubble parameter with $E(z) = H(z)/H_0$ where $H(z)$ is the Hubble parameter. Ω_{lm} denotes the angular extent of the image cube and B_{bw} denotes the frequency bandwidth of the data cube. $D_M(z)$ and ΔD are conversion factors to go to comoving distance units from angular and frequency units, respectively. The denominator in Eq. (2) accounts for the limitation of the angular extent due to the primary beam $A_{\text{pb}}(l, m)$ and the spatial tapering function $A_w(l, m)$ and also the limitation in the frequency extent due to the frequency filter $B_w(\nu)$ used before the Fourier transform along the frequency direction. $\langle \dots \rangle$ denotes an average over the respective domains. The power values $P(k_l, k_m, k_{\parallel})$ are next averaged in cylindrical and spherical shells to yield the cylindrical (two-dimensional) power spectrum $P(k_{\perp}, k_{\parallel})$ and the spherical power spectrum $P(k)$, respectively, where $k_{\perp}^2 = k_l^2 + k_m^2$ and $k^2 = k_{\perp}^2 + k_{\parallel}^2$.

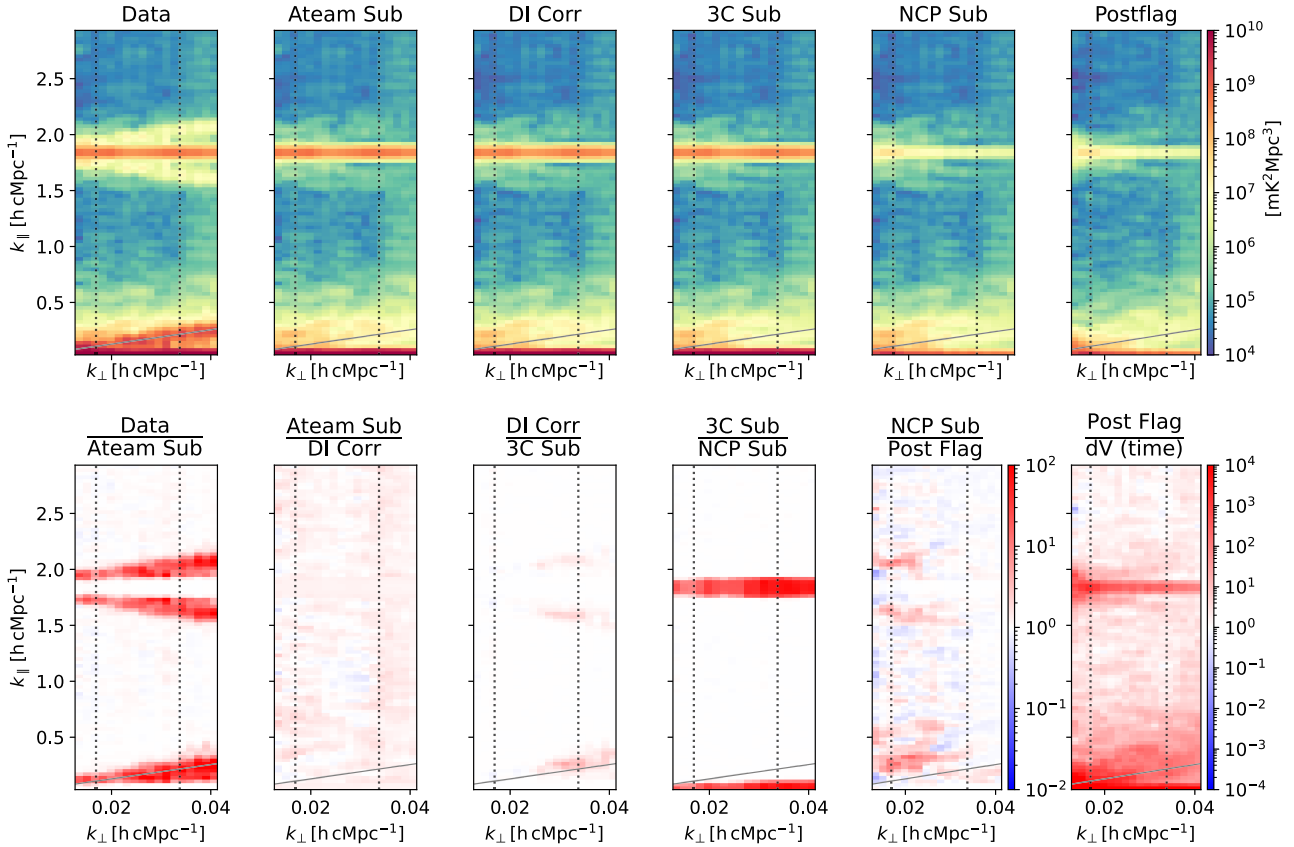


Fig. 9. Cylindrical power spectra at different stages of calibration and source subtraction. Top row: cylindrical power spectra after preprocessing (“Data”), after A-team subtraction (“Ateam Sub”), after DI correction (“DI Corr”), after 3C subtraction (“3C Sub”), after NCP Subtraction (“NCP Sub”), and after post-calibration RFI flagging (“Postflag”). Bottom row: ratio of successive power spectra in the top row. This shows how much power has been subtracted and from which part of the k space at the different calibration stages. The rightmost plot in the bottom row is the ratio of the power spectrum of the data after post-calibration RFI flagging to the noise power spectrum obtained from time-differenced Stokes-V data. Note that the color bars for all plots in the same row are the same, except for the plot on the extreme right in the bottom row.

We computed the Stokes- I cylindrical power spectra at all the stages of calibration as a diagnostic to compare the power levels after each step. The cylindrical power spectra for each calibration step and the ratio of the power spectra in successive steps are presented in Fig. 9. The gray lines correspond to the horizon delays and the foreground wedge is clearly visible to have several orders of magnitude higher power than the “EoR window”. The bright horizontal feature at $k_{\parallel} \approx 1.84 \text{ h cMpc}^{-1}$ is the result of the flagging of sets of 2 channels at the ends of each sub-band during preprocessing, in order to avoid the edge effect of the polyphase filter used in forming the sub-bands. This feature at such a high k mode does not affect our power spectrum analysis significantly, since it is focused on k modes typically much smaller than this. It is evident that a significant amount of power near the wedge’s horizon is removed during the A-team and 3C subtraction steps, while the NCP subtraction removes power at the lowest k_{\parallel} modes corresponding to the low delay values of the NCP field at the phase center. Post-calibration flagging decreases the power by a factor of more than 10 in the high k_{\parallel} -modes beyond the horizon line ($0.2\text{--}0.6 \text{ h cMpc}^{-1}$). In some k -modes this factor is less than unity because the thermal noise level is higher after flagging due to the lower volume of remaining data. It should be noted that the data prior to the DI correction stage is not absolutely calibrated. The noise power spectrum is estimated by taking the difference of the Stokes-V dirty image cubes of the even and odd time samples and then forming the power spectra as usual, accounting for the extra fac-

tor of 2 increase in variance in the process. The panel on the bottom right in Fig. 9 shows the power spectrum of the data after post-calibration RFI flagging divided by the noise power. We see that within the wedge, there is still more than three orders of magnitude of power beyond the thermal noise limit, likely due to Galactic diffuse emission and confusion noise due to extragalactic sources. The power far beyond the wedge could be due to a variety of factors such as residual RFI and polarization leakage. Well away from the wedge, at $k_{\parallel} > 1 \text{ h cMpc}^{-1}$, the residual power approaches the thermal noise limit.

7. Residual foreground removal

The data, after the different steps of calibration and compact source subtraction, are still dominated by residual foregrounds, such as the diffuse Galactic emission and extragalactic point sources, which have a flux density at or below the confusion noise limit. However, the fact that foregrounds have a larger frequency coherence scale than thermal noise or the 21 cm signal can be utilized to model and subtract the foregrounds from the data. One approach for subtracting foregrounds from the data is through GPR (Mertens et al. 2018).

7.1. Gaussian process regression

In GPR, the data are modeled as a sum of Gaussian processes describing the foreground, thermal noise, and the 21 cm signal

components. Each Gaussian process is characterized by a certain frequency covariance function and zero mean. The covariance function is parametrized by a set of adjustable hyperparameters that control properties such as the variance, coherence scale, and the shape of the covariance function. Using a Bayesian approach, the maximum a posteriori values of the hyperparameters are derived from their posterior probability distribution conditioned on the observed data. The expectation value of the foreground component at each data point is subtracted from the data to yield the residual foreground subtracted data.

To limit the computational requirements, we applied GPR to the gridded visibility data cube $\tilde{T}(u, v, \nu)$ before power spectrum generation. Performing GPR along the frequency direction in the uv space allows us to easily take into account the baseline dependence of the frequency coherence scale of the different components, such as the foreground wedge and the thermal noise. Following the same approach as taken by [Mertens et al. \(2018\)](#), the data \mathbf{d} is modeled as a sum of different components, namely the foregrounds \mathbf{f}_{fg} , the 21 cm signal \mathbf{f}_{21} and the noise \mathbf{n} as functions of frequency ν ,

$$\mathbf{d} = \mathbf{f}_{\text{fg}}(\nu) + \mathbf{f}_{21}(\nu) + \mathbf{n}(\nu). \quad (4)$$

The different GP components should in principle be possible to separate by virtue of their having different spectral behavior. This spectral behavior is specified by the covariance of the components between frequency channels, with the total covariance matrix of the data being a sum of the individual GP covariances:

$$\mathbf{K}(\nu_p, \nu_q) = \mathbf{K}_{\text{fg}}(\nu_p, \nu_q) + \mathbf{K}_{21}(\nu_p, \nu_q) + \mathbf{K}_n(\nu_p, \nu_q). \quad (5)$$

Here \mathbf{K} is the total covariance matrix of the data whose entries are a function of two frequencies ν_p and ν_q , and is given by the sum of the foregrounds covariance matrix \mathbf{K}_{fg} , the 21 cm covariance matrix \mathbf{K}_{21} and the noise covariance matrix \mathbf{K}_n .

In an earlier approach to implementing GPR, which was employed by [Gehlot et al. \(2019, 2020\)](#), and [Mertens et al. \(2020\)](#), both the foreground and 21 cm covariances have a specific functional form along with hyperparameters that guide their respective variance and frequency coherence scale. However, one concern with this method is that there is a risk of signal loss if the 21 cm covariance function is not a good match to the frequency covariance of the 21 cm signal in actual data ([Kern & Liu 2021](#)). This becomes particularly important when the objective is to use the 21 cm power spectrum upper limits to rule out astrophysical models that have a 21 cm power spectrum that is not well described by the covariance function adopted for the 21 cm signal in GPR. Addressing these concerns, a novel approach to GPR-based foreground subtraction called ML-GPR has been developed by [Mertens et al. \(2024\)](#), which employs ML methods to build a covariance model of the 21 cm signal directly from simulations. The main steps in the ML-GPR approach are summarized in the following section.

7.2. Training

In ML-GPR, a variational auto-encoder (VAE) kernel is trained to build a low dimensional representation of the 21 cm signal covariance from a large number of simulations of the 21 cm signal. We used the 21cmFAST ([Mesinger et al. 2011](#); [Murray et al. 2020](#)) code to produce simulations of the 21 cm signal at our redshift of interest ($z = 20$), with a comoving box size of 600 Mpc and a resolution of 2 Mpc per pixel. We followed the




Fig. 10. Normalized spherically averaged power spectra of the 1000 21cmFAST simulations at $z = 20$ used as the training set for the 21 cm VAE kernel.

parameterization introduced by [Park et al. \(2019\)](#) with the following range of astrophysical parameters:

- $f_{*,10}$ the normalization of the fraction of Galactic gas in stars at high z , evaluated for halos of mass $10^{10} M_{\odot}$: $\log_{10}(f_{*,10}) = [-3, 0]$.
- α_* the power-law scaling of f_* with halo mass: $\alpha_* = [-0.5, 1]$.
- $f_{\text{esc},10}$ the normalization of the ionizing UV escape fraction of high z galaxies, evaluated for halos of mass $10^{10} M_{\odot}$: $\log_{10}(f_{\text{esc},10}) = [-3, 0]$.
- α_{esc} the power-law scaling of f_{esc} with halo mass: $\alpha_{\text{esc}} = [-0.5, 1]$.
- t_* the star formation timescale taken as a fraction of the Hubble time: $t_* = [0, 1]$.
- M_{turn} the turnover halo mass below which the abundance of active star forming galaxies is exponentially suppressed: $\log_{10}(M_{\text{turn}}/M_{\odot}) = [8, 10]$.
- E_0 the minimum X-ray photon energy capable of escaping the galaxy, in keV: $E_0 = [0.1, 1.5]$.
- L_X/SFR the normalization of the soft-band X-ray luminosity per unit star formation computed over the band $E_0 - 2$ keV: $\log_{10}(L_X/\text{SFR}) = [38, 42]$.

Generating meaningful simulations of cosmic dawn made it necessary to perform IGM spin temperature fluctuations. Inhomogeneous recombination was also turned on in the simulations. Latin hypercube sampling was used to sample 1000 sets of parameters in this eight-dimensional space and 21cmFAST was used to perform 1000 simulations and obtain the corresponding brightness temperature cubes. We note that this relatively sparsely sampled parameter space is sufficient since the VAE, being a generative model, is able to interpolate between the training samples. This has been shown by [Mertens et al. \(2024\)](#). Additionally, since we used a VAE instead of an auto-encoder (AE), the added regularization allowed us to avoid overfitting the sparse sample ([Kingma & Welling 2013](#)). Next, the spherically averaged power spectra for all 21 cm brightness temperature cubes were computed and normalized in the k range $0.03 - 2.0 h \text{ cMpc}^{-1}$ to have a variance of unity. This was done because we aim to use the VAE to learn only the shape of the 21 cm power spectrum and keep its variance as a separate free parameter, thus allowing it to account for boosted signals predicted by exotic models. Though this provides more freedom to the models, in the case of a detection it would be necessary to check if the converged model is physically plausible. The normalized spherically averaged power spectra of all simulations used in the training set can be seen in Fig. 10, showing the large variety of power spectrum shapes.

These power spectra were then used to train a VAE with a two-dimensional latent space, meaning that we want to capture the shape of the 21 cm signal using two parameters. It should be noted that due to the 21 cm signal power spectrum being isotropic (to first order, ignoring peculiar velocities), the spherical power spectrum contains all information about the signal under the assumption of Gaussianity. So it is sufficient to train the VAE on the spherical power spectra, rather than the covariance matrix itself. The VAE has two components, an encoder that maps the normalized power spectra to the latent space, and a decoder that can be used to recover the normalized power spectra corresponding to any point in the two-dimensional latent space. Both the encoder and decoder were trained, and the optimization was performed by minimizing the reconstruction loss between the training power spectrum used as input to the encoder, and the recovered power spectrum given as output by the decoder. We divided the simulated power spectra into a training set of 950 power spectra and a validation set of 50 power spectra. The reconstruction loss, defined as the mean squared error (MSE) between the output and the input power spectra, stabilized after 500 out of a total of 4000 iterations. When comparing the behavior of the reconstruction loss for the training and validation sets, no over-fitting was observed. After training, we also checked the ratio between the input and output. A median value of 1 and rms of 0.1 was observed for both the training and validation sets. This rms is well below what is typically expected in terms of measurement errors with the first-generation detection experiments. The power spectrum obtained from the trained decoder at any given point in the latent space can now be used to calculate the frequency-frequency covariance matrix, thus effectively capturing the covariance of the 21 cm signal from the simulations into two latent space quantities. We also tested training on a higher dimension of the latent space but did not find any improvement in the reconstruction. This is likely due to the sparse sampling of the eight-dimensional parameter space using 1000 simulations and a denser sample with a higher latent space dimension could possibly capture subtler changes in the power spectra. However, we defer that to future analyses with higher sensitivities, where such small effects on the power spectrum will be more important.

7.3. Covariance model

The trained VAE kernel serves as the 21 cm covariance (\mathbf{K}_{21}) with three parameters: the two latent space dimensions x_1 and x_2 and a scaling factor for the 21 cm signal variance σ_{21}^2 . For the foregrounds, we used an analytical covariance model, which is a good description of the spectral structure we see in the data based on multiple trials using different combinations of covariance functions. The form of the functions used in our covariance model can be described using the Matern class functions,

$$\kappa_{\text{Matern}}(\nu_p, \nu_q) = \sigma^2 \frac{2^{1-\eta}}{\Gamma(\eta)} \left(\frac{\sqrt{2\eta}r}{l} \right)^\eta K_\eta \left(\frac{\sqrt{2\eta}r}{l} \right). \quad (6)$$

Here, σ^2 is the variance, l is the frequency coherence scale, $r = |\nu_p - \nu_q|$ is the frequency separation, Γ is the Gamma function and K_η is the modified Bessel function of the second kind. Different values of η correspond to different functional forms that are special cases of the Matern class functions. The GP covariance has a foreground, 21 cm, and noise components (Eqs. (4) and (5)). However, we find that it is necessary to use two components to model the foregrounds: an “intrinsic” and a

“mode-mixing” component. In addition, an “excess” component was used to account for the excess power seen in the data. The different components used in the GP covariance model are:

Intrinsic foregrounds – $\mathbf{K}_{\text{int}}(l_{\text{int}}, \sigma_{\text{int}}^2)$. Diffuse Galactic emission and extragalactic point sources below the confusion limit within the FOV constitute the intrinsic foregrounds after sky model subtraction. These foregrounds are expected to have a very large frequency coherence scale due to the smooth spectrum of the synchrotron emission mechanism. We modeled the covariance of the intrinsic foregrounds using a radial basis function (RBF; obtained by setting $\eta = \infty$ in Eq. (6)), which yields very smooth models along frequency (Mertens et al. 2018). We used a uniform prior $\mathcal{U}(20, 40)$ MHz on l_{int} to capture the very large frequency coherence scale features due to intrinsic foregrounds. We note that a wider prior does not affect the converged value of l_{int} and the specific range $\mathcal{U}(20, 40)$ was chosen to have a narrow prior enclosing the converged value, which sped up the ML-GPR runs for the 100 signal injection tests that were performed on ML-GPR (described in Sect. 8.3).

Mode-mixing foregrounds – $\mathbf{K}_{\text{mix}}(l_{\text{mix}}, \sigma_{\text{mix}}^2)$. Interferometers are chromatic instruments and flat-spectrum sources far away from the phase center till the horizon occupy a region in the k_\perp, k_\parallel space known as the foreground wedge (Morales et al. 2012; Vedantham et al. 2012). Apart from this, additional frequency modulations can be imparted on the foreground data by the chromatic primary beam, especially near nulls and sidelobes, the instrumental bandpass, and other systematic effects. To account for these effects, we used a mode-mixing foreground component. The covariance model used for this component is an RBF but with a $\mathcal{U}(0.1, 0.5)$ MHz prior on l_{mix} accounting for the smaller frequency scale fluctuations due to the mode mixing. An RBF was chosen here since it has a rapid fall in power at high delay due to the smooth models it yields. This also makes it easier to separate it from the 21 cm signal without signal loss and this has been tested through simulations and signal injection tests.

Excess – $\mathbf{K}_{\text{ex}}(l_{\text{ex}}, \sigma_{\text{ex}}^2)$. We find that the data cannot be adequately described by just a foreground and a 21 cm component, and there is additional power in the data with a small coherence scale that is difficult to differentiate from the 21 cm signal. This could be caused by small-scale frequency fluctuations introduced into the data by instrumental effects and RFI. Suboptimal calibration and polarization leakage could also be contributing factors to this additional power. This “excess power” is seen to be well described by an exponential covariance model. An exponential function is obtained by setting $\eta = 0.5$ in Eq. (6). We used a $\mathcal{U}(0.2, 2)$ MHz prior on l_{ex} for this component. We note that even though this prior range is similar to the prior range for l in \mathbf{K}_{mix} , an exponential kernel does not have a sharp drop in power at large k_\parallel like the RBF, making it considerably more difficult to separate from the 21 cm signal component. Hence, this excess component was not subtracted from the data. This avoids a potential signal loss due to the absorption of the 21 cm signal into this component.

Noise – \mathbf{K}_n . The noise covariance is calculated from the time-differenced Stokes-V image cubes as a proxy for the thermal noise. We find that using a fixed noise covariance is sufficient, and multiplying the noise covariance by a scaling factor does not affect our results since such a scaling factor converges to a value very close to 1.

Table 5. Components of the covariance model used in ML-GPR along with the priors and converged values of the parameters.

Component	Covariance	Parameter	Description	Prior bounds	Estimated value
Intrinsic foregrounds	Radial basis function	σ_{int}^2	Variance	$[-1, 1]$	-0.1 ± 0.02
		l_{int}	Lengthscale	$[20, 40]$	32.9 ± 1.4
Mode-mixing foregrounds	Radial basis function	σ_{mix}^2	Variance	$[-2, 0]$	-1.32 ± 0.01
		l_{mix}	Lengthscale	$[0.1, 0.5]$	0.275 ± 0.001
21 cm signal	Trained ML Kernel	x_1	Latent space dimension	$[-4, 4]$	–
		x_2	Latent space dimension	$[-4, 4]$	–
		σ_{21}^2	Variance	$[-7, -1]$	< -4.97
Excess power	Exponential function	σ_{ex}^2	Variance	$[-3, -1]$	-1.88 ± 0.02
		l_{ex}	Lengthscale	$[0.2, 2]$	0.56 ± 0.03

Notes. All σ^2 values are in logarithmic scale and are expressed as a fraction of the variance of the data input into GPR. σ_{21}^2 reaches the lower bound and hence only the upper limit is shown. The priors are all uniform priors. All l values are in units of MHz.

7.4. Application to data

The data used as input to ML-GPR are the gridded visibility cubes described in Sect. 6.2. The optimal parameters for our covariance model were obtained using Monte Carlo-based algorithms, Markov chain Monte Carlo (MCMC) or nested sampling, to generate samples from the posterior distribution. Monte Carlo-based methods offer an advantage over simple gradient-based optimization techniques since the former yield both the optimal parameters and the uncertainties associated with them, which can be propagated down to obtain the corresponding uncertainties on the final power spectrum. We get very similar results with both MCMC and nested sampling as the optimization method, but nested sampling, while being more computationally expensive, yields a more complete sampling of the parameter space within the prior bounds and also provides an evidence value. Hence, we used nested sampling with 100 live points to obtain the optimal set of parameters. The prior ranges and converged values for the parameters in our covariance model are listed in Table 5. Figure 11 shows a corner plot of the posterior probability distribution of the parameters. The parameters x_1 and x_2 , which describe the 21 cm signal shape, do not converge, and the variance σ_{21}^2 hits the lower bound of the prior range, as we would expect for data in which the thermal noise level is well above the 21 cm signal. All other parameters for the foreground and excess components converge to well-constrained l and σ^2 values, within the prior bounds. Finally, multiple realizations of foreground cubes were sampled from the posterior probability distribution of the parameters and subtracted from the data to yield residual data cubes, which were then used to calculate the residual power spectrum (Fig. 15) along with its uncertainties, representing the spread in the distribution of samples in the parameter space.

8. Robustness tests and residual power spectra

To ensure that the different calibration and foreground subtraction steps do not result in signal loss (e.g., Patil et al. 2016), it was necessary to perform robustness tests on these steps. The test consists of injecting an artificial 21 cm signal into the data, passing it through the calibration step, and comparing the recovered signal with the injected signal. This is a similar procedure to that followed by Mevius et al. (2022). The injected signal power is well above the expected 21 cm signal, but as long as it is small compared to the foregrounds and in the linear perturbation

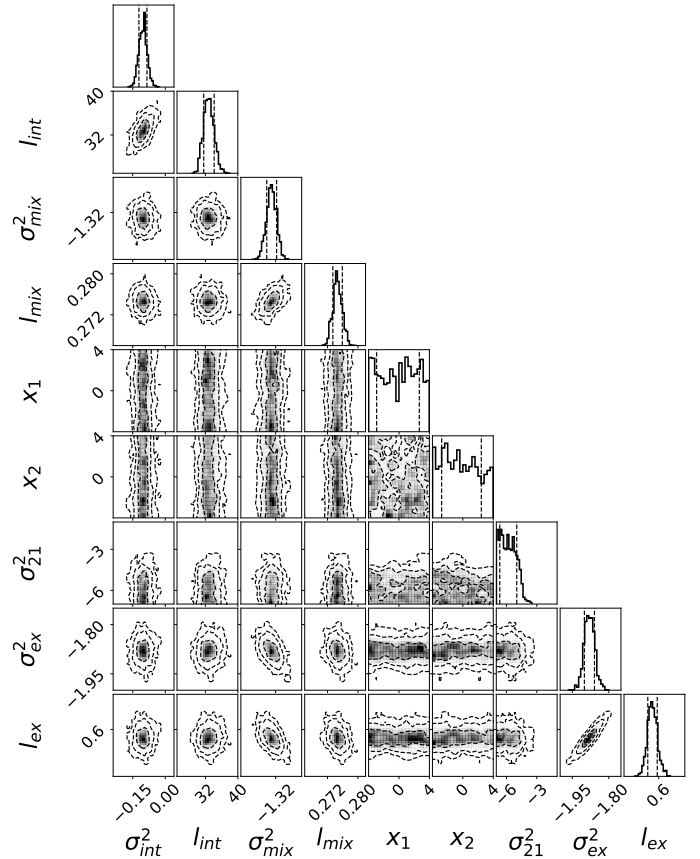


Fig. 11. Corner plot showing the posterior probability distributions for the different parameters used in ML-GPR. The dashed contours correspond to the 68%, 95%, and 99.7% confidence levels. The vertical dashed lines in the one-dimensional histogram enclose the central 68% of the probability.

regime, the suppression factor is expected to be the same as the actual faint signal (Mouri Sardarabadi & Koopmans 2019). We verified this by performing injection tests with different signal strengths. We performed robustness tests on all the steps involving DD calibration and also on ML-GPR. A flowchart describing the workflow of the signal injection simulations is given in Fig. 12.

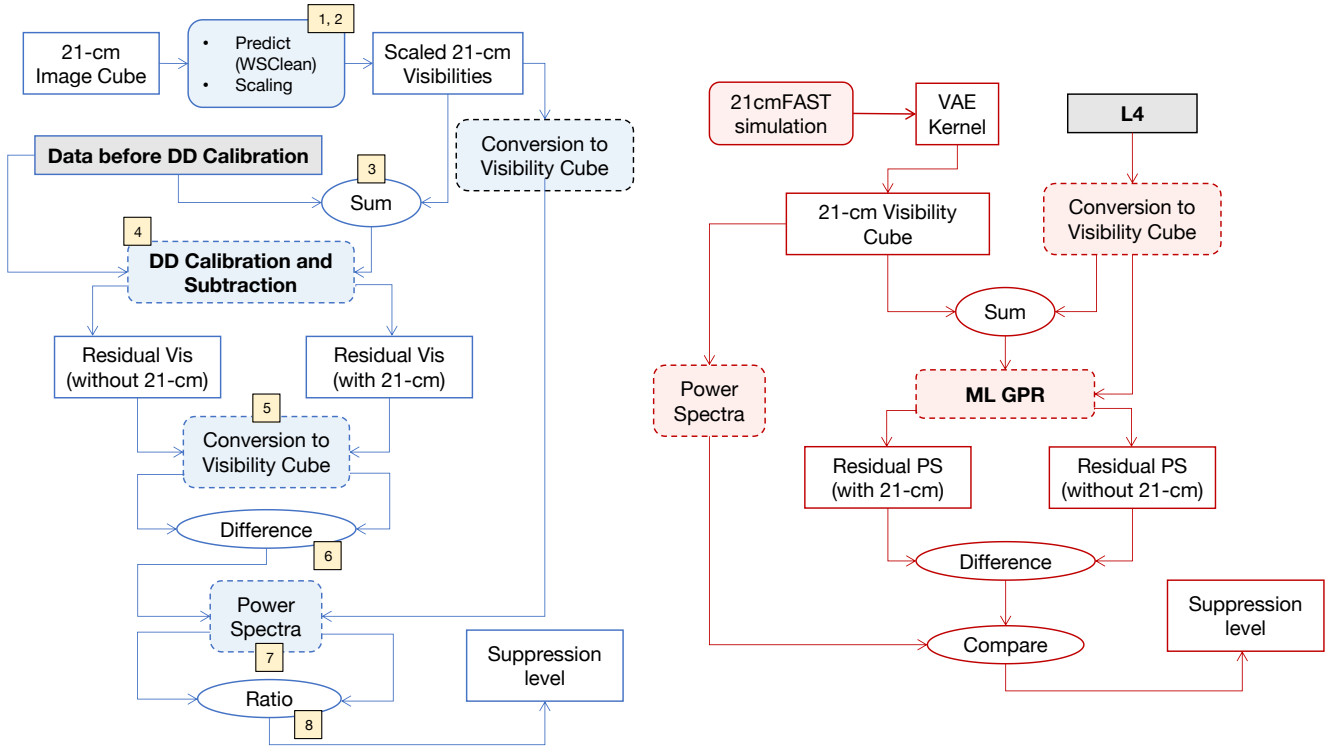


Fig. 12. Flowchart for the signal injection tests performed on a DD calibration and subtraction step (outlined in blue), and on ML-GPR (outlined in red). The numbers refer to the sequence of steps described in the text in Sect. 8.1.

8.1. Robustness test on DD subtraction steps

DD calibration of radio interferometric data hinges on the fact that for arrays with a large number of interferometric elements, there are enough excess degrees of freedom in the system of calibration equations to allow us to obtain separate calibration solutions for several directions. However, having a large number of directions also increases the number of free parameters in the calibration model and the risk of over-fitting features other than the DD effects increases. In other words, the MA gains for the different directions can conspire to create a PSF that, convolved with the sky model in that direction, can mimic a part of the 21 cm signal or any other component present in the data and absent in the sky model. This can lead to suppression of the signal and any un-modeled component in the data, in particular on the larger angular and frequency scales (Patil et al. 2016). In our analysis, the steps of A-team subtraction, 3C subtraction, and NCP subtraction involve a step of DD calibration followed by model visibility prediction, corruption, and subtraction. Hence, it is necessary to check that these steps do not result in significant signal suppression. Therefore, we performed a signal injection test on all these three calibration steps separately. Below we list the steps followed to perform the signal injection test during the calibration step:

1. The mock signal used in this analysis was obtained from the 21 cm brightness temperature cube produced by Jelić et al. (2008) using the 21cmFAST code, which was also used by Mevius et al. (2022). We note that the signal injection test in the linear perturbation regime is independent of the injected signal (Mouri Sardarabadi & Koopmans 2019), and to be self-consistent with earlier work, we chose to use the same cube. This image cube was used to predict the corresponding visibilities employing the predict task in WSClean, which effectively

performs a Fourier transform of the data to generate visibilities at the observed uvw coordinates for each frequency.

2. The 21 cm simulated visibilities were multiplied by a factor of 1000 to boost up the signal to be close to the thermal noise level of the data for the lowest k modes. At the same time, it was made sure that the injected signal is still more than two orders of magnitude fainter than the foregrounds so that the signal itself has a negligible impact on the amplitude of the gain solutions.

3. These scaled simulated 21 cm signal visibilities were added to the data visibilities just before the calibration step being tested.

4. The calibration and subtraction were then performed with exactly the same parameter settings as was done for the data without the injected signal.

5. The DD subtracted visibilities were then gridded in the manner described in Sects. 6.1 and 6.2 to obtain the data cube in the uvw space.

6. The data cube without the injected signal was subtracted from the data cube with the injected signal to obtain the residual 21 cm signal data cube.

7. The cylindrical and spherical power spectra for this residual data cube were obtained in the manner described in Sect. 6.3.

8. The ratio of these power spectra over the power spectra of the injected 21 cm signal gives the factor by which the signal power is suppressed at the different wave modes due to this particular DD calibration and subtraction step.

The workflow of the signal injection tests on DD calibration and subtraction is shown on the left of Fig. 12. The numbers indicate the sequence of steps.

The signal injection test was repeated for different spectral gain smoothing kernel widths and baseline cuts. Smoothness constraints on the gain solutions can decrease signal suppression since it decreases the degrees of freedom of the calibration




Fig. 13. Results of the robustness tests performed on different steps of the calibration pipeline. The results of the tests on the three DD calibration steps for different smoothing kernel widths and baseline cuts are presented in the first three columns from the left. The top row shows the suppression factor in the spherical power spectra for the different steps. The lines in red correspond to the settings that were used for the final analysis, and the plots in the bottom row show the suppression factor in the cylindrical power spectra corresponding to these settings. The vertical dotted lines in the bottom row correspond to the 20λ and 40λ cuts and the region between these two lines is used for constructing the final spherical PS. The fourth column from the left shows the cumulative suppression of a signal injected into the L2 data. The top panel shows the cumulative signal suppression after A-team subtraction and after NCP subtraction for the injected 21 cm signal scaled by a factor of 1000 (Signal 1: circles) and 2000 (Signal 2: inverted triangles). The inverted triangle markers have been shifted horizontally for visual clarity. The rightmost column presents the spherical (top) and cylindrical (bottom) power spectra of the injected signal and their comparison with the noise power spectra.

algorithm by imposing constraints on the parameters that are being solved for, thereby decreasing the chance of over-fitting (Barry et al. 2016; Ewall-Wice et al. 2016; Patil et al. 2016). Applying a baseline cut involves the use of only the baselines longer than a limit for calculating the calibration solutions. The station-based gain solutions are then applied to the predicted model visibilities and subtracted from all baselines. This restricts any absorption of the signal, and hence signal suppression, to only the baseline range higher than the cut that was used in calibration. Now only the baselines smaller than the cut can be used for power spectrum estimation and there is a much lower risk of signal suppression in these baselines. However, using a baseline cut decreases the accuracy of the gain solutions for a given station, since a small number of baselines that are longer than the cut are available for a given station for calculating the gains. This problem is particularly important for NenuFAR since having a baseline cut larger than the core diameter (400 m) would mean that for a given core station, only three baselines involving the remote stations are available for obtaining the gain solution (for the current NenuFAR configuration with three remote MAs). This decreases the accuracy of the model subtraction and introduces an excess variance on baselines lower than the cut. Also, the baseline-based gains (combination of the gains of two stations) are only constrained by the sky model on long baselines and not so much on short baselines. Hence, any inaccuracy of the sky model and the gains on the longer baselines can lead to

spurious signals when applied to the model and subtracted from the short baselines. This can also contribute to excess variance.

The results of the signal injection tests done on the different DD calibration steps are summarized in Fig. 13. As expected, we see that there is negligible signal suppression on the baselines that are not used in the calibration step. Also, using a wider gain smoothing kernel reduces the signal suppression in the baselines that are used in calibration. The final values of the gain smoothing kernel width and baseline cut used for the three different DD calibration and subtraction tests were decided by comparing the results of the signal injection tests with the level of residual power after subtraction.

For the A-team subtraction step, using a 40λ baseline cut is not feasible. This is because it is the first step in our calibration pipeline, and nonoptimal A-team subtraction due to the lower number of baselines used in calibration adversely affects all the following steps of DI correction and sky model construction. The smoothing kernel of 2 MHz for A-team subtraction is necessary since a 4 MHz smoothing is not able to account for the small-scale frequency fluctuations of the primary beam at the location of the A-team sources. Using a 4 MHz kernel does decrease the signal suppression factor but results in high residuals at the location of the A-team, which again affects all the successive calibration steps. So we used a frequency smoothing kernel of 2 MHz and a baseline cut of 20λ for the A-team subtraction step, thus favoring accurate A-team subtraction and allowing for a $<10\%$

signal loss for which a bias correction was performed when calculating the power spectra (Sect. 8.4).

For 3C subtraction and NCP subtraction, we find that using a baseline cut of 40λ does result in higher residual power after subtraction compared to when we used a 20λ cut. However, we get negligible signal suppression for a 40λ cut and the increase in residual power is less than the inverse of the signal suppression factor for a 20λ cut. Thus, we used a 40λ cut for both the 3C subtraction and NCP subtraction steps. A frequency smoothing kernel of 2 MHz was used for the 3C subtraction step since it leads to a more accurate subtraction of the 3C sources, which are far from the phase center. For the NCP subtraction step, we used a 6 MHz smoothing kernel since it leads to better source subtraction than a 2 MHz kernel. This is because the phase center is at the NCP and a 6 MHz kernel is sufficient to account for the frequency behavior of the main lobe of the primary beam. Using a 6 MHz kernel also increases the S/N of the gain solutions and we find that the noise in the gain amplitude and phase solutions is smaller when a 6 MHz kernel is used for NCP subtraction. (Appendix B.2). The S/N becomes an important issue for the NCP subtraction since the total flux is divided into multiple clusters and it is necessary to make sure that the faintest cluster has enough S/N to yield reliable calibration solutions.

8.2. Robustness test on the full calibration pipeline

Once the results of the robustness tests for each DD subtraction step were assessed and the final settings to use in each processing and analysis step were selected, we performed a signal injection test on the entire calibration pipeline. This involved adding the scaled 21 cm simulated visibilities to the L2 data and passing it through all the steps of A-team subtraction and DI correction, 3C subtraction, and NCP subtraction one by one. The residual data cube without the injected signal was subtracted from the residual data cube with the injected signal to obtain the residual 21 cm signal data cube. The ratio of the power spectrum of this data cube with that of the injected signal then yields the suppression factor for the entire calibration pipeline. We note that the signal was injected before DI correction, so the calculated DI gains get applied to the injected signal in the course of passing the data through the pipeline. So for all steps after DI correction, to calculate the suppression factor, it is necessary to divide the power spectrum of the residual signal by the power spectrum of the injected signal with these DI gains applied.

The fourth column in Fig. 13 shows the results of injecting the signal to the L2 data and passing it through the pipeline with the final settings. We find that in addition to the A-team subtraction step, the DI correction step introduces a small additional suppression (2% of the signal). In the top panel, the additional suppression after the A-team subtraction step till the NCP subtraction step comes almost solely from the DI correction step. This additional factor is due to the fact that the calculated DI gains in the direction of the NCP are slightly higher when there is an injected signal. We repeated this signal injection test with twice the magnitude of the injected signal (Signal 2 in the top row, fourth and fifth columns of Fig. 9). We find that the signal suppression percentage in all DD calibration steps and in the DI correction step remains the same. This is in agreement with the results of Mouri Sardarabadi & Koopmans (2019) who conclude that in the linear regime where the signal is much weaker than the foregrounds, the suppression factor due to calibration is independent of the injected signal strength. In the bottom row of Fig. 13, we can see the factor by which the signal will be suppressed in the k_\perp, k_\parallel space due to the calibration pipeline. It should be

noted that we only used the baseline range 20λ – 40λ (the region between the vertical dotted lines) to construct the final spherical power spectra.

8.3. Robustness test on ML-GPR

Since we applied ML-GPR to the gridded visibility cubes, an additional signal injection test was done on the gridded data cube after NCP subtraction. This allows us to test the robustness of the ML-GPR step using a large variety of simulated 21 cm signals. This was not necessary for the calibration steps since the exact shape of the signal does not affect the signal injection results during calibration, as long as the injected signal is small compared to the foregrounds (Mouri Sardarabadi & Koopmans 2019). The decoder of the trained VAE kernel (described in Sect. 7) was used to generate power spectra corresponding to uniformly spaced points in the two-dimensional latent space. We chose five points between the values -2 and 2 for both x_1 and x_2 , making a total of 25 different shapes of the 21 cm power spectrum. For each such signal, the power was scaled to be equal to 0.5, 1, 4, and 20 times the thermal noise power, thus covering a large range of intensities of the 21 cm signal for which we want to test the performance of ML-GPR. Each of these 100 mock 21 cm signals was separately injected into the data. To do this, one realization of the visibility cube was predicted per signal and added to the complex gridded data cube before ML-GPR. The data cube with the injected signal was next used to perform GPR with the same priors as was used for the actual data. The residual power spectrum of the data without an injected signal was then subtracted from the residual power spectrum obtained by applying ML-GPR to data with an injected signal. The workflow used in the signal injection tests for ML-GPR is shown on the right of Fig. 12. The recovered 21 cm power spectrum from the injection test (Δ_{rec}^2) can be compared to the power spectrum of the injected signal (Δ_{inj}^2) to see if the signal is absorbed in the procedure. We used two metrics to quantify the results of the signal injection simulations, the z -score and bias, defined as follows:

$$z\text{-score}(k) = \frac{\Delta_{\text{rec}}^2(k) - \Delta_{\text{inj}}^2(k)}{\sigma_{\Delta_{\text{rec}}^2}(k)},$$

$$\text{Bias}(k) = \frac{\Delta_{\text{rec}}^2(k)}{\Delta_{\text{inj}}^2(k)}. \quad (7)$$

The z -score specifies the number of standard deviations by which the recovered signal power spectrum over or underestimates the input. A negative z -score indicates suppression of the signal and a value below -2 would thus indicate a suppression of the signal beyond the 2-sigma upper limits at that particular k bin. The z -scores for all signal injection tests are shown in the top panel of Fig. 14. The different k bins are indicated with different colored points. We find that none of the signals are suppressed beyond the 1-sigma error bars, and thus the 2-sigma upper limits are comfortably valid. The z -score for brighter signals has a large positive value since $\sigma_{\Delta_{\text{rec}}^2}(k)$ have very low values compared to the signal strength for stronger signals. This is because the contribution to the error bar from the uncertainty in the converged hyperparameter values is lower for stronger signals since the signals are estimated with higher precision. The bias tells us the factor by which the recovered signal power is higher than the injected signal. The bias at each k bin for all signal injection tests is illustrated in the bottom panel of Fig. 14. The signals for which the bias values are less than unity are all within the




Fig. 14. Results of the robustness tests performed on the ML-GPR step. Top: z-score for all injected signals at the different k bins. The vertical dashed lines demarcate the different strengths of the injected signals. The horizontal gray bands indicate the 1-sigma and 2-sigma levels. Bottom: Bias for all injected signals at the different k bins. The 25 cells in each panel correspond to the 25 different shapes of the injected signal. The different rows correspond to the different strengths of the injected signal and the different columns show the results in different k bins. The color scale indicates the value of the bias in each case.

1 sigma limits and correspond to the points in the top panel that lie below zero. There is a positive bias, particularly for signals with high x_2 values. This suggests that there is some degeneracy between the foregrounds, excess, and the 21 cm components when modeled as Gaussian processes, but it only results in the absorption of the foregrounds into the excess or the 21 cm signal component, which are not subtracted from the data. This is acceptable as long as we set upper limits on the magnitude of the signal, and do not claim a detection.

8.4. Residual power spectra

The residual foreground subtracted data cube from ML-GPR was used to compute the spherical power spectra in logarithmically spaced bins in k space between $k = 0.035 \text{ h cMpc}^{-1}$ and $k = 0.5 \text{ h cMpc}^{-1}$. To compute the spherical power spectra, we used only the data in the uv range of 20λ – 40λ , in order to avoid the signal suppression in calibration due to the 3C subtraction and NCP subtraction steps. The spherical and cylindrical power spectra before and after the entire processing are shown in Fig. 15. We find that the residual power spectra are around two orders of magnitude higher than the thermal noise level at the lowest k bins, contributed mainly by the strong excess component that is left in the data after GPR. We note that the thermal noise power spectrum for “Residual” is slightly lower than that

Table 6. 2σ upper limits on the 21 cm power spectrum $\Delta_{ul}^2(k)$ at different k bins. $\Delta_{21}^2(k)$ is the residual power, $\Delta_{21, \text{err}}^2(k)$ is the error on the residual power, and $\Delta_N^2(k)$ is the thermal noise power.

k bin (h cMpc^{-1})	$\Delta_{21}^2(k)$ (K^2)	$\Delta_{21, \text{err}}^2(k)$ (K^2)	$\Delta_{ul}^2(k)$ (K^2)	$\Delta_N^2(k)$ (K^2)
0.036–0.055	17.81	2.90	23.60	0.12
0.055–0.086	71.30	11.11	93.52	0.55
0.086–0.134	261.63	30.96	323.54	2.01
0.134–0.207	456.33	49.54	555.41	8.10
0.207–0.322	731.75	57.18	846.12	31.29
0.322–0.5	1157.37	85.06	1327.49	104.66

for “Data” since the post-calibration RFI flagging removes 7.5% of the data.

The noise bias at each k bin was calculated by passing the time-differenced Stokes- V image cube (see Sect. 6.1) through the power spectrum pipeline, and accounting for the factor of 2. This was subtracted from the residual power spectrum to obtain the noise bias subtracted dimensionless power spectrum. In Sect. 8.2, we found that there is a maximum of 11% suppression in the signal as a cumulative effect of the entire calibration pipeline, coming mainly from the A-team subtraction and DI correction steps in which the 20λ to 40λ baselines were used (Fig. 13). To account for this, we performed a bias correction, which amounts to applying the inverse of the calculated average cumulative signal suppression factors at each k bin to the residual power spectra. We place the upper limits at a level of 2σ above these values at the different k modes. The power spectrum values with the 2σ error bars after noise bias subtraction and after bias correction are shown in Fig. 15 in the rightmost panel. The 2σ upper limits are listed in Table 6. The best 2σ upper limit is $2.4 \times 10^7 \text{ mK}^2$ at $k = 0.041 \text{ h cMpc}^{-1}$.

9. Discussion and next steps

In this section, we discuss the results and the limitations of the current analysis and present the plans for future improvements to the analysis pipeline.

9.1. Nature of the “excess”

The residual data after foreground subtraction is seen to be dominated by excess power up to two orders of magnitude beyond the thermal noise limit. This excess power is not possible to subtract from the data using GPR since its spectral behavior is similar to what we expect from the 21 cm signal.

In order to understand the behavior of the excess as a function of time, we utilized the fact that the visibility data were divided into 13 time segments of 52 min duration each (with a 56 min last segment). The gridded data cube for each of these segments was obtained following Sects. 6.1 and 6.2. A separate GPR was performed on each cube, with the length scales fixed to the values obtained for the full data, and the variances left as free parameters, which were optimized in GPR. Fixing the length scales prevents overfitting of the data, and optimizing the variances allows the time dependence of the strength of the GPR components to be modeled over timescales longer than a time segment. The maximum a posteriori solutions of the foreground cubes were subtracted from the data cube of each time segment to give the residual data cubes, which were used to construct the




Fig. 15. Spherical and cylindrical power spectra at some key stages of the analysis pipeline. The left-most panel shows the spherical power spectra after preprocessing (“Data”), after sky model subtraction (“Skymodel Sub”), and after GPR (“GPR Residual”), along with the thermal noise power spectrum (“Thermal Noise”). For the cylindrical power spectra (second and third panels), the ratio with respect to the thermal noise power spectrum is shown. The white dashed lines indicate the horizon limit. The spherical power spectra after noise bias subtraction and suppression factor correction are shown in the rightmost panel.

respective power spectra, shown in the left panel of Fig. 16. The excess is stronger in the first six segments, and the last seven segments have a similar power level. We note that the last segment has the lowest power because it is 56 min long and has a lower thermal noise power level than the other 52 min segments.

Possible contributors to this excess are the residuals from Cyg A, which dominate the uv plane after sky model subtraction as seen in Fig. 6. Also, Cyg A goes to lower elevations as time progresses and the beam gain in its direction has very low values after the first 6 segments (Fig. 4). The residuals from Cyg A subtraction are expected to be directly related to its apparent brightness and this is confirmed by wide-field images of the data (for example, the images on the left panels of Fig. 5) where the residuals at its location are seen to go down with time. Thus, Cyg A residuals are expected to be a major contributor to the excess. It is important to note that this excess power is a result of the variance on the NCP field caused by the residual PSF sidelobes of Cyg A. Therefore, even if the peak flux at the position of the Cyg A is reduced to be below the thermal noise limit, it does not necessarily remove the effect of its PSF sidelobes on the NCP field where the power spectrum is computed. This can only be improved in the future with careful time and frequency modeling of the beam, ionosphere, and any other DD effects, during calibration.

To determine if the excess power is coherent with time within a single night, the power spectrum was constructed with data from an increasing number of time segments included. This was done by computing the weighted average (using the uv weights) of the residual data cubes for those segments and constructing the power spectra. It should be noted that the weighted average of all the time segments retrieves the exact full data cube since the gridding is linear. The residual power is seen to go down as we integrate more data (right panel of Fig. 16). We find that the power spectrum indicated by “ $N = 13$ ” has slightly lower power than that of “full data”. This is because, in the case of $N = 13$, GPR has been applied to the individual data cubes with the variances of the individual GPR components allowed to vary in time instead of having a single variance representing the strength of the GPR components for the duration of the entire




Fig. 16. Spherical power spectra of the data after applying GPR on individual time segments. The left panel shows the power spectra of the GPR residuals for individual time segments of the data and the right panel shows the power spectra with increasing volume of data, in the order of observation.

observation as is the case for full data, thus allowing slightly better foreground modeling and subtraction. The ratio between the power of the first segment and that of the full data is seen to be more than a factor of 20, as opposed to a factor of 13 that we expect from using 13 times as much data if the excess is incoherent between time segments. This is likely due to the first few segments having more power from the Cyg A residuals, as discussed above. The fact that the power goes down by a factor of 13 or more as we integrate 13 times as much data suggests that the excess power, which dominates the residuals,

is incoherent. Thus, as we integrate more nights of observations, it is possible that this excess could go down significantly with time if a part of the excess is incoherent across different nights. This needs further investigation by processing multiple nights of data. However, even if this is confirmed, the final power spectra would remain well above what can be expected from thermal noise alone, which also averages down in a similar manner.

9.2. Impact of wide-field sources

The NenuFAR primary beam has very strong grating lobes roughly 60 degrees away from the phase center. Wide-field images of the data reveal a large number of strong radio sources falling in these grating lobes, which remain in the data after sky model subtraction. In the 3C subtraction step, seven of the brightest among these sources were subtracted through DD calibration. However, the remaining sources have too little S/N to obtain reliable calibration solutions. Though these sources are far away from the phase center, the PSF sidelobes of these sources will still pass through the NCP field and add a significant contribution to the excess power. To understand the impact of these sources far away from the phase center, we computed wide-field MFS dirty images from the data for each of the time segments into which the data had been divided. Each image was divided into four annular regions around the NCP, labeled “Main Lobe”, “Nulls”, “Grating Lobes”, and “Edge”, based on the amplitude of the radial profile of the NenuFAR primary beam (bottom panel in Fig. 17). For the pixels within the n th annulus, the cross-coherence between the i th and j th time segments is given by

$$C_{ij}(n) = \frac{\langle I_i(l, m) I_j(l, m) \rangle}{\sqrt{\langle I_i^2(l, m) \rangle \langle I_j^2(l, m) \rangle}}, \quad (8)$$

where $\langle \dots \rangle$ indicates an average over all l, m lying within the n th annulus. Figure 17 shows the cross-coherence at four different calibration and foreground subtraction stages. It should be noted that the cross-coherence of a segment with itself is unity by definition.

Before A-team subtraction, at the “data” level, the nulls annulus is dominated by flux from Cas A, and the grating lobes annulus is dominated by Cyg A. Cas A is high in the sky throughout the observation, but Cyg A goes below the horizon halfway through the observation, resulting in lower cross-coherence values for the later time segments. After A-team subtraction, the grating lobes annulus is dominated by the strong 3C sources. The 3C subtraction step removes seven of the brightest sources, but the remaining sources have a significant cross-coherence in the grating lobes annulus. There is also higher cross-coherence for the separation of five time segments, corresponding to a time difference of 4.3 h. This is possibly because of the six-fold symmetry of the hexagonal MAs, which results in the grating lobes in the sky that overlap exactly once every 4 h. In the main lobe, the cross-coherence values are close to unity until before NCP subtraction. NCP subtraction decreases the cross-coherence in the main lobe and the modulation seen in the cross-coherence with increasing time difference is because there are large residuals from sources at the edge of the main lobe. As the elongated main lobe rotates against the sky, it completes close to a 180 degree rotation in the duration of the observation as a result of which the time segments separated by 13 segments again have high cross-coherence values. We note that this modulation due to the rotation of the non-axisymmetric main lobe of the beam is also seen

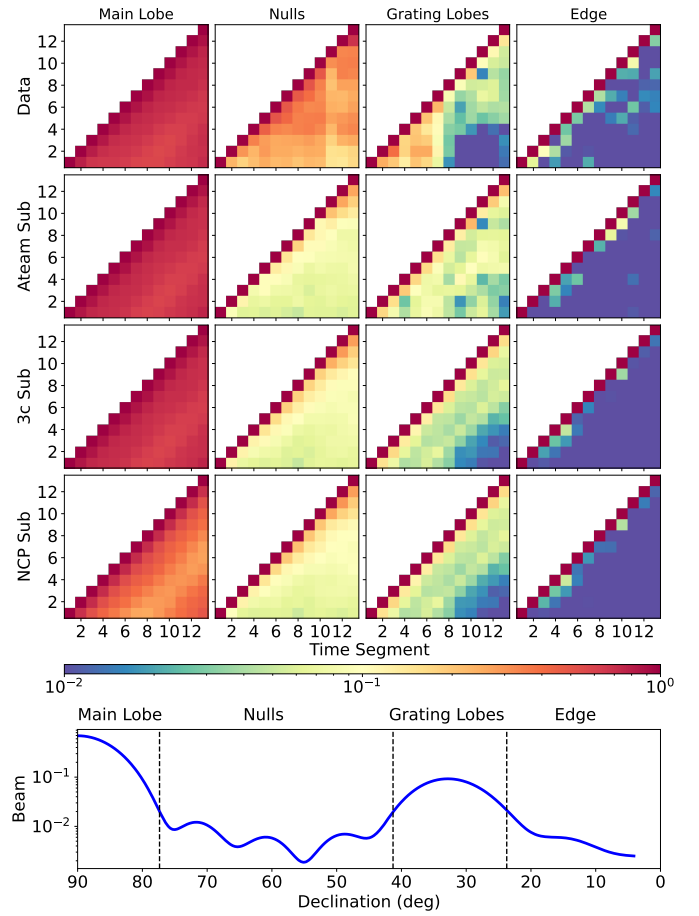


Fig. 17. Cross-coherence of the data between different time segments. The top panel shows the cross-coherence computed in four different regions of the primary beam (different columns), for the data at four different stages of sky model subtraction (different rows). The bottom panel shows what the four regions correspond to in the radial profile of the time-averaged primary beam.

before NCP subtraction, but the effect is lower since the field has a more uniform distribution of flux before NCP subtraction. The edge region is dominated by noise, so there is negligible cross-coherence between time segments.

These cross-coherence results show that the timescale of coherence is heavily dependent on the distance from the phase center because of the primary beam. For the main lobe, this timescale is nearly the duration of the full night, while for the nulls, the coherence drops to 20% or less for consecutive segments. For the grating lobes, the coherence is about 10% for the separation of between one and six segments and falls to less than 2% for longer timescales. If the edge region is a reasonable estimate of the noise, we can state that the coherence of the nulls and grating lobes (and, of course, the main lobe) are still dominated by real sources and not too much diluted by noise. These sources behave incoherently in their flux or position from segment to segment due to the primary beam fixed to the Earth. Thus, the power from these sources, coupled with the non-smooth spectral nature of the primary beam at the nulls and grating lobes, could have a large contribution to the excess power at high k_{\parallel} seen in the residual data, which averages down with time like incoherent noise (as seen in Sect. 9.1), on timescales that exceed their own coherence time (as seen in Fig. 17).

9.3. Limitations and future improvement

This first analysis on a single night of NenuFAR observations reveals the limitations of the current analysis pipeline. Here, we describe these limitations and discuss future improvements in the processing pipeline necessary to overcome them.

Near-field RFI sources. Near-field imaging (i.e., imaging of sources on the ground plane in the vicinity of the array) of the data after calibration and sky model subtraction revealed that two prominent sources of local RFI contaminate the data. The approach adopted in this analysis to mitigate its impact was by identifying the baselines where these RFI sources have a strong signature in the delay power spectra and flagging those baselines completely. However, this is an overly simplistic approach, and all un-flagged baselines still have some power due to these local RFI sources. It was seen that flagging the affected baselines significantly reduced excess power at high delay modes. So, local RFI in the unflagged baselines could make up much of the excess power we see in the data. A more targeted approach in mitigating the RFI sources could be by filtering them out from delay or fringe rate power spectra of all baselines, utilizing the fact that the RFI sources have a localized signature in the delay-fringe rate space. We defer that to a future analysis. We note that the impact of the near-field RFI sources is expected to be maximum for NCP observations since the NCP is the only point in the northern sky that is fixed with respect to the Earth, and the local stationary RFI sources add up coherently at the NCP even as the sky rotates with time. This is less the case for other fields, not including the NCP, and the contaminating effects of these local RFI sources on observations of these fields could be significantly less.

Impact of diffuse (polarized) foregrounds. NenuFAR, with its dense core consisting of many small baselines, is very sensitive to diffuse Galactic foregrounds. These (polarized) foregrounds are not present in the sky model used in calibration, and the power from the diffuse foregrounds can be absorbed in the station gains, leading to calibration errors. We used a 20λ baseline cut in calibration for the A-team subtraction step and a 40λ cut for the other calibration steps. This avoids the power due to diffuse foregrounds at spatial modes measured on baselines shorter than 20λ and 40λ , respectively. However, using a baseline cut results in excess power on the excluded baselines. Thus, both 3C and NCP subtraction result in an excess noise on the baselines smaller than 40λ , which are used in power spectrum generation. The errors in calibration due to diffuse foregrounds can be improved by including diffuse foreground models. However, the ideal way to model the diffuse foregrounds, perhaps using shapelets decomposition (Yatawatta 2011; Gehlot et al. 2022), must still be tested. A future step could also be to build a broadband spectral model of diffuse foregrounds by combining AARTFAAC High Band Antenna and NenuFAR observations.

Beam model in calibration. The beam model of NenuFAR is currently not integrated into the calibration algorithm. As a result, we performed calibration against an apparent sky model obtained by multiplying the intrinsic sky model with the average beam for each time segment of data. The effect of the beam on the apparent flux of sources on timescales more than the calibration solution interval is accounted for by the calibration gain solutions. However, the variation of the beam on time and frequency scales shorter than a solution interval is currently not accounted for. The high residuals due to Cyg A (seen in the middle-left panel of Fig. 5 and the middle and right panels of

Fig. 6), which are expected to be a major contributor to the excess variance (Sect. 9.1), is partly a consequence of this. The solution interval used in the A-team subtraction step is not small enough to model the primary beam accurately at the grating lobes in which Cyg A lies. This could be addressed to some extent when the primary beam model is integrated into the calibration software, and the level of improvement will depend on the accuracy of the beam model.

For the NCP subtraction step, including a beam model in the calibration step will allow us to use longer solution intervals since this will partly mitigate the requirement of the gain solutions to represent the rotating non-axisymmetric beam. This will lower the expected thermal noise for each solution interval and thus make it possible to use smaller clusters with enough flux to be significantly beyond the thermal noise level. Thus, it will be possible to account for DD effects on smaller spatial scales than currently possible. Including a beam model in the calibration software will also mean that the gain solutions do not need to account for the frequency behavior of the primary beam. Therefore, even with a large value for the smoothing scale, it will be possible to fit and subtract the sources well, thus avoiding increasing the risk of overfitting by decreasing the width of the frequency smoothing kernel.

Another advantage of having the beam in the calibration step is using a wider sky model of the NCP field. This is because the spatial gradient of the source fluxes due to attenuation by the beam will now be partly accounted for by the beam model, thus allowing us to use wider clusters with higher S/N. We see in the bottom-right panel of Fig. 5 and the bottom-left panel of Fig. 17 that after NCP subtraction, there are residual sources near the edge of the main lobe of the primary beam. Using an extended sky model, along with the beam model, will allow us to subtract these sources better.

Finally, including a beam model will allow us to subtract many sources in the grating lobes of the NenuFAR primary beam. Currently, these sources cannot be subtracted because to model the beam, we needed to calibrate and absorb the beam in the gain solutions. Performing such a calibration run is impossible since the sources do not have enough S/N to yield reliable calibration solutions within the time and frequency solution interval needed to represent the beam faithfully. Once a beam model is available, the visibilities for these sources with the primary beam attenuation can be predicted and subtracted directly from the data. The accuracy of this subtraction will depend on the accuracy of the simulated beam model. We note that all sources that are not subtracted and the residuals from the subtracted sources will continue to have the effect of the rotating beam and will still contribute to the excess power.

Confusion noise limit. The NCP subtraction step is limited by confusion noise due to the maximum angular resolution of NenuFAR, which is given by the baseline formed by two remote MAs. Once more remote MAs are available, a 5 km maximum baseline in the future will improve the angular resolution by a factor of ≈ 3.4 and allow us to construct a deeper sky model due to a lower confusion limit (by a factor of ≈ 6.7). This will allow us to use smaller clusters in the NCP subtraction step since smaller clusters will have a high enough S/N to give reliable calibration solutions.

Polarization leakage. In this analysis, we have used diagonal mode in calibration with DDEC1 and Stokes- I and Stokes- V data for all steps. However, we expect some polarization leakage in NenuFAR since it is a wide-field phased array

instrument. Additionally, due to its dense core, NenuFAR is particularly sensitive to diffuse foregrounds, which are expected to have a polarized component. The leakage of Faraday rotated polarized foregrounds into Stokes- I can result in the foregrounds assuming a non-smooth spectral behavior. This can be tested with simulations once the full Stokes beam model is available, and the inclusion of such a beam model in calibration can potentially decrease instrumental polarization.

Ionospheric effects. An important effect that can impact the power spectrum is ionospheric scintillation noise (Vedantham & Koopmans 2015, 2016; Gehlot et al. 2018). A turbulent ionosphere introduces phase shifts on the incoming electromagnetic waves, which can cause the sources to move around in the sky with time, thus affecting the accuracy of their subtraction through calibration. In this analysis, we have selected a relatively quiet ionospheric night, and the calibration solution time intervals used at each stage are longer than the expected timescales of ionospheric effects, making it impossible to solve for these effects. A detailed simulation of the ionosphere is necessary to quantify its impact on NenuFAR data and will be done in the future.

In summary, a few specific effects could be the dominant contributors to the excess power we see in the data. Targeted improvements in the processing pipeline can account for, mitigate, and possibly eliminate these effects.

10. Summary

This paper describes the first end-to-end analysis of the spectral window of 61–72 MHz from an 11.4 h observation of the NCP field with NenuFAR. We draw the following main conclusions from this analysis:

- We set the best 2σ upper limits on the 21 cm power spectrum at $2.4 \times 10^7 \text{ mK}^2$ at $k = 0.041 \text{ h cMpc}^{-1}$ and $z = 20.3$. These upper limits are two orders of magnitude higher than the thermal noise power spectrum.
- Localized RFI sources strongly impact the data and are a major contributor to the excess power. The RFI is mitigated to an extent by baseline flagging, and a more targeted approach to filtering these local RFI sources could significantly mitigate the excess power.
- The excess power seems largely incoherent with time within a single night and averages down like noise as we include more data to construct the power spectra. Processing of multiple nights is necessary to understand whether coherent averaging of multiple nights of data will decrease this excess power.
- Many sources lying in the nulls and grating lobes of the NenuFAR primary beam are left in the data after sky model subtraction. Coupled with the frequency fluctuations of the primary beam in these regions, these sources could strongly contribute to the excess power seen in the data. Including a simulated primary beam model in the calibration step would allow us to subtract these sources to some extent and partly mitigate this effect.

We will continue to investigate the origin of this excess power and implement improvements in the analysis pipeline to mitigate it. This will allow us to exploit the full sensitivity of NenuFAR. In parallel, we plan to process multiple nights of observations and obtain upper limits closer to the thermal noise and expected 21 cm signal levels, which will allow us to constrain the underlying astrophysical models.

Acknowledgements. SM, LVEK, and BKG acknowledge the financial support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 884760, "CoDEX"). FGM acknowledges the support of the PSL Fellowship. AB acknowledges support from the European Research Council through the Advanced Grant MIST (FP7/2017–2022, No. 742719). EC would like to acknowledge support from the Centre for Data Science and Systems Complexity (DSSC), Faculty of Science and Engineering at the University of Groningen. AF is supported by the Royal Society University Research Fellowship. The post-doctoral contract of IH was funded by Sorbonne Université in the framework of the Initiative Physique des Infinis (IDEX SUPER). We are grateful to Bradley Greig for his valuable input and guidance in setting up 21cmFAST simulations for cosmic dawn. This paper is based on data obtained using the NenuFAR radiotelescope. NenuFAR has benefitted from the following funding sources : CNRS-INSU, Observatoire de Paris, Station de Radioastronomie de Nançay, Observatoire des Sciences de l'Univers de la Région Centre, Région Centre-Val de Loire, Université d'Orléans, DIM-ACAV and DIM-ACAV+ de la Région Ile de France, Agence Nationale de la Recherche. We acknowledge the Nançay Data Center resources used for data reduction and storage.

References

- Abdurashidova, Z., Aguirre, J. E., Alexander, P., et al. 2022, *ApJ*, **925**, 221
- Adams, T., Aguirre, J. E., Alexander, P., et al. 2023, *ApJ*, **945**, 124
- Asad, K., Koopmans, L., Jelić, V., et al. 2015, *MNRAS*, **451**, 3709
- Asad, K., Koopmans, L., Jelić, V., et al. 2016, *MNRAS*, **462**, 4482
- Asad, K., Koopmans, L., Jelić, V., et al. 2018, *MNRAS*, **476**, 3051
- Barkana, R. 2018, *Nature*, **555**, 71
- Barkana, R., & Loeb, A. 2005, *ApJ*, **624**, L65
- Barry, N., Hazelton, B., Sullivan, I., Morales, M., & Pober, J. 2016, *MNRAS*, **461**, 3135
- Barry, N., Wilensky, M., Trott, C., et al. 2019, *ApJ*, **884**, 1
- Berlin, A., Hooper, D., Krnjaic, G., & McDermott, S. D. 2018, *Phys. Rev. Lett.*, **121**, 011102
- Bevins, H. T., Heimersheim, S., Abril-Cabezas, I., et al. 2024, *MNRAS*, **527**, 813
- Bowman, J. D., Rogers, A. E., & Hewitt, J. N. 2008, *ApJ*, **676**, 1
- Bowman, J. D., Rogers, A. E., Monsalve, R. A., Mozdzen, T. J., & Mahesh, N. 2018a, *Nature*, **555**, 67
- Bowman, J. D., Rogers, A. E., Monsalve, R. A., Mozdzen, T. J., & Mahesh, N. 2018b, *Nature*, **564**, E35
- Broekema, P. C., Mol, J. J. D., Nijboer, R., et al. 2018, *Astron. Comput.*, **23**, 180
- Chapman, E., Abdalla, F. B., Harker, G., et al. 2012, *MNRAS*, **423**, 2518
- Chapman, E., Abdalla, F. B., Bobin, J., et al. 2013, *MNRAS*, **429**, 165
- Charrier, D., & the CODALEMA Collaboration 2007, in *2007 IEEE Antennas and Propagation Society International Symposium (IEEE)*, 4485
- Chen, X., & Miralda-Escudé, J. 2004, *ApJ*, **602**, 1
- Cohen, A., Lane, W., Cotton, W., et al. 2007, *AJ*, **134**, 1245
- Condon, J. J., Cotton, W., Greisen, E., et al. 1998, *AJ*, **115**, 1693
- DeBoer, D. R., Parsons, A. R., Aguirre, J. E., et al. 2017, *PASP*, **129**, 045001
- Dowell, J., & Taylor, G. B. 2018, *ApJ*, **858**, L9
- Eastwood, M. W., Anderson, M. M., Monroe, R. M., et al. 2019, *AJ*, **158**, 84
- Ewall-Wice, A., Dillon, J. S., Hewitt, J. N., et al. 2016, *MNRAS*, **460**, 4320
- Ewall-Wice, A., Chang, T.-C., Lazio, J., et al. 2018, *ApJ*, **868**, 63
- Feng, C., & Holder, G. 2018, *ApJ*, **858**, L17
- Fialkov, A., & Barkana, R. 2014, *MNRAS*, **445**, 213
- Fialkov, A., & Barkana, R. 2019, *MNRAS*, **486**, 1763
- Fialkov, A., Barkana, R., & Cohen, A. 2018, *Phys. Rev. Lett.*, **121**, 011101
- Field, G. B. 1958, *Proc. IRE*, **46**, 240
- Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, *Phys. Rep.*, **433**, 181
- Gan, H., Mertens, F., Koopmans, L., et al. 2023, *A&A*, **669**, A20
- Garsden, H., Greenhill, L., Bernardi, G., et al. 2021, *MNRAS*, **506**, 5802
- Gehlot, B. K., Koopmans, L. V., de Bruyn, A., et al. 2018, *MNRAS*, **478**, 1484
- Gehlot, B., Mertens, F., Koopmans, L., et al. 2019, *MNRAS*, **488**, 4271
- Gehlot, B., Mertens, F., Koopmans, L., et al. 2020, *MNRAS*, **499**, 4158
- Gehlot, B., Koopmans, L., Offringa, A., et al. 2022, *A&A*, **662**, A97
- Gessey-Jones, T., Sartorio, N. S., Fialkov, A., et al. 2022, *MNRAS*, **516**, 841
- Hibbard, J. J., Mirocha, J., Rapetti, D., et al. 2022, *ApJ*, **929**, 151
- Hicks, B. C., Paravastu-Dalal, N., Stewart, K. P., et al. 2012, *PASP*, **124**, 1090
- Hills, R., Kulkarni, G., Meerburg, P. D., & Puchwein, E. 2018, *Nature*, **564**, E32
- Hothi, I., Chapman, E., Pritchard, J. R., et al. 2021, *MNRAS*, **500**, 2264
- Intema, H. T., Jagannathan, P., Mooley, K. P., & Frail, D. A. 2017, *A&A*, **598**, A78
- Jelić, V., Zaroubi, S., Labropoulos, P., et al. 2008, *MNRAS*, **389**, 1319
- Joseph, R. 2012, Bachelor's Thesis, Faculty of Science and Engineering
- Kazemi, S., Yatawatta, S., & Zaroubi, S. 2013, *MNRAS*, **430**, 1457
- Kern, N. S., & Liu, A. 2021, *MNRAS*, **501**, 1463

- Kingma, D. P., & Welling, M. 2013, arXiv e-prints [arXiv:[1312.6114](#)]
- Koopmans, L., Pritchard, J., Mellema, G., et al. 2015, arXiv e-prints [arXiv:[1505.07568](#)]
- Kolopanis, M., Jacobs, D. C., Cheng, C., et al. 2019, *ApJ*, **883**, 133
- Li, W., Poher, J., Barry, N., et al. 2019, *ApJ*, **887**, 141
- Liu, H., Outmezguine, N. J., Redigolo, D., & Volansky, T. 2019, *Phys. Rev. D*, **100**, 123011
- Madau, P., Meiksin, A., & Rees, M. J. 1997, *ApJ*, **475**, 429
- McQuinn, M., Zahn, O., Zaldarriaga, M., Hernquist, L., & Furlanetto, S. R. 2006, *ApJ*, **653**, 815
- Mebane, R. H., Mirocha, J., & Furlanetto, S. R. 2020, *MNRAS*, **493**, 1217
- Mertens, F., Ghosh, A., & Koopmans, L. 2018, *MNRAS*, **478**, 3640
- Mertens, F. G., Mevius, M., Koopmans, L. V., et al. 2020, *MNRAS*, **493**, 1662
- Mertens, F. G., Semelin, B., Koopmans, L. V. E., et al. 2021, in *SF2A-2021: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics*, eds. A. Siebert, K. Baillié, E. Lagadec, et al., 211
- Mertens, F. G., Bobin, J., & Carucci, I. P. 2024, *MNRAS*, **527**, 3517
- Mesinger, A., Furlanetto, S., & Cen, R. 2011, *MNRAS*, **411**, 955
- Mevius, M. 2018, RMextract: Ionospheric Faraday Rotation calculator, Astrophysics Source Code Library [record ascl:[1806.024](#)]
- Mevius, M., Mertens, F., Koopmans, L., et al. 2022, *MNRAS*, **509**, 3693
- Monsalve, R. A., Fialkov, A., Bowman, J. D., et al. 2019, *ApJ*, **875**, 67
- Morales, M. F., & Hewitt, J. 2004, *ApJ*, **615**, 7
- Morales, M. F., Hazelton, B., Sullivan, I., & Beardsley, A. 2012, *ApJ*, **752**, 137
- Mouri Sardarabadi, A., & Koopmans, L. 2019, *MNRAS*, **483**, 5480
- Muñoz, J. B., & Loeb, A. 2018, *Nature*, **557**, 684
- Munoz, J. B., Qin, Y., Mesinger, A., et al. 2022, *MNRAS*, **511**, 3657
- Murray, S. G., Greig, B., Mesinger, A., et al. 2020, *J. Open Source Software*, **5**, 2582
- Offringa, A. 2016, *A&A*, **595**, A99
- Offringa, A., & Smirnov, O. 2017, *MNRAS*, **471**, 301
- Offringa, A., Van De Gronde, J., & Roerdink, J. 2012, *A&A*, **539**, A95
- Offringa, A., McKinley, B., Hurley-Walker, N., et al. 2014, *MNRAS*, **444**, 606
- Offringa, A., Mertens, F., Van der Tol, S., et al. 2019a, *A&A*, **631**, A12
- Offringa, A., Mertens, F., & Koopmans, L. 2019b, *MNRAS*, **484**, 2866
- Paciga, G., Chang, T.-C., Gupta, Y., et al. 2011, *MNRAS*, **413**, 1174
- Paciga, G., Albert, J. G., Bandura, K., et al. 2013, *MNRAS*, **433**, 639
- Pandey, V., Koopmans, L., Tiesinga, E., Albers, W., & Koers, H. 2020, *Astronomical Data Analysis Software and Systems XXIX*, 527, 473
- Park, J., Mesinger, A., Greig, B., & Gillet, N. 2019, *MNRAS*, **484**, 933
- Patil, A. H., Zaroubi, S., Chapman, E., et al. 2014, *MNRAS*, **443**, 1113
- Patil, A. H., Yatawatta, S., Zaroubi, S., et al. 2016, *MNRAS*, **463**, 4317
- Patil, A., Yatawatta, S., Koopmans, L., et al. 2017, *ApJ*, **838**, 65
- Planck Collaboration XIII. 2016, *A&A*, **594**, A13
- Pritchard, J. R., & Furlanetto, S. R. 2007, *MNRAS*, **376**, 1680
- Pritchard, J. R., & Loeb, A. 2012, *Rep. Prog. Phys.*, **75**, 086901a
- Reis, I., Fialkov, A., & Barkana, R. 2020, *MNRAS*, **499**, 5993
- Rengelink, R. B., Tang, Y., de Bruyn, A. G., et al. 1997, *A&AS*, **124**, 259
- Semelin, B., Mériot, R., Mertens, F., et al. 2023, *A&A*, **672**, A162
- Shaver, P., Windhorst, R., Madau, P., & De Bruyn, A. 1999, arXiv e-prints [arXiv:[astro-ph/9901320](#)]
- Singh, S., Subrahmanyan, R., Shankar, N. U., et al. 2017, *ApJ*, **845**, L12
- Singh, S., Nambissan, T., Subrahmanyan, R., et al. 2022, *Nat. Astron.*, **1**
- Smirnov, O., & Tasse, C. 2015, *MNRAS*, **449**, 2668
- Subrahmanyan, R., Somashekar, R., Shankar, N. U., et al. 2021, *Exp. Astron.*, **51**, 193
- Trott, C. M., Jordan, C., Midgley, S., et al. 2020, *MNRAS*, **493**, 4711
- van Diepen, G., Dijkema, T. J., & Offringa, A. 2018, Astrophysics Source Code Library [record ascl:[1804.003](#)]
- van Haarlem, M. P., Wise, M. W., Gunst, A., et al. 2013, *A&A*, **556**, A2
- Vedantham, H., & Koopmans, L. 2015, *MNRAS*, **453**, 925
- Vedantham, H., & Koopmans, L. 2016, *MNRAS*, **458**, 3099
- Vedantham, H., Shankar, N. U., & Subrahmanyan, R. 2012, *ApJ*, **745**, 176
- Venkatesan, A., Giroux, M. L., & Shull, J. M. 2001, *ApJ*, **563**, 1
- Wouthuysen, S. 1952, *AJ*, **57**, 31
- Yatawatta, S. 2011, in *2011 XXXth URSI General Assembly and Scientific Symposium* (IEEE), 1
- Yoshiura, S., Pindor, B., Line, J., et al. 2021, *MNRAS*, **505**, 4775
- Zarka, P., Girard, J., Tagger, M., & Denis, L. 2012, in *SF2A-2012: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics*, 687
- Zarka, P., Tagger, M., Denis, L., et al. 2015, *2015 International Conference on Antenna Theory and Techniques (ICATT)* (IEEE), 1
- Zarka, P., Denis, L., Tagger, M., et al. 2020, in *URSI GASS 2020*

Appendix A: RFI statistics




Fig. A.1. RFI percentage at the L1 data level. The two panels show the flagging percentage as a function of time and frequency (left) and as a function of baseline (right).

RFI flagging was performed at different stages of the data analysis pipeline. In Fig. A.1, we present the percentage of flagged data by AOFlogger until the L1 data level as a function of time, frequency, and baseline.

Appendix B: Calibration solutions

B.1. Bandpass solutions

Bandpass calibration is a crucial step to account for the frequency ripples produced in the data due to cable reflections. In NenuFAR the different MAs have different cable lengths, and the cable reflections are clearly visible in the autocorrelations before bandpass calibration. This is illustrated in Fig. B.1 where the delay power spectra of the autocorrelations of the NCP data before (middle panel) and after (right panel) bandpass calibration are plotted along with the delay power spectra of the bandpass solutions obtained separately from the Cas A observation (left panel), for a few example MAs. The vertical dashed lines indicate the expected delays due to the cable reflections for the MAs calculated from their respective cable lengths. We see a peak near the expected delay in both the bandpass solutions from the Cas A observation and the autocorrelations of the NCP data before bandpass calibration. The peaks are not present in the autocorrelations after the bandpass solutions have been applied to the data.

In the peak-normalized delay power spectra of the autocorrelations before and after bandpass calibration (middle and right panels in Fig. B.1), we find that even though the peaks due to the cable reflections are accounted for by bandpass calibration, the relative power at high delay values increases after bandpass calibration. To understand the impact of this effect on the power spectra, we performed a test in which the bandpass solutions were smoothed by a Gaussian filter of 2 MHz width, separately on the real and imaginary parts, before applying them to the data. The delay spectra of the autocorrelations with the smoothed bandpass show the peaks at the expected delays corresponding to the cable lengths but have a lower noise floor. However, in the final power spectra, we found negligible differences in the power levels of the autocorrelations with the actual bandpass against the autocorrelations with the smoothed bandpass. This is because the thermal noise level is well above the noise floor of the bandpass gain solutions. For the NCP field, the total flux




Fig. B.1. Bandpass calibration results for some example MAs. Left column: Delay power spectra of bandpass solutions obtained from the Cas A observation. Middle column: Delay power spectra of autocorrelations for the NCP observation before bandpass calibration. Right column: Delay power spectra of autocorrelations for the NCP observation after bandpass calibration. The cable lengths for the respective MAs are indicated by vertical dashed lines.

is not high enough to give a dynamic range, which is needed for these effects to show up in the cylindrical and spherical power spectra.

B.2. DD gain errors

We find that for the NCP subtraction step, a 6 MHz smoothing kernel results in better subtraction of sources compared to a 2 MHz smoothing kernel. Here, we compare the noise in the gain solutions in these two cases. The difference in the gain solutions at successive time steps is used as a proxy for the error in the gain solutions under the assumption that the gains do not change substantially over two solution intervals. Fig. B.2 shows the errors in the amplitude and phase of the DD gain solutions




Fig. B.2. Gain calibration errors at the NCP subtraction step. The top and bottom panels show the errors in the amplitude and phase of the gain solutions, respectively. The dashed lines correspond to the solutions obtained using 2 MHz smoothing while the solid lines correspond to those obtained using 6 MHz smoothing. The thick solid lines show the beam factor and are to be read using the vertical axis on the right-hand side. The different colors refer to the gains in the direction of the different clusters.

obtained in this manner, by taking the standard deviation of the time-differenced gains across frequencies and stations. Clusters 1 to 7 are arranged in the decreasing order of flux and we see that fainter clusters have higher calibration errors. For each cluster, the calibration errors for a 6 MHz smoothing kernel are smaller compared to a 2 MHz smoothing kernel. The modulation of the errors in the gain solutions toward Clusters 6 and 7 with time is due to the fact that the central lobe of the NenuFAR primary beam is not axisymmetric and is elongated in one direction. As the beam rotates with respect to the sky with time, Cluster 7 overlaps with the beam for the first few hours and then moves out of the beam. So the errors are lower in the first few hours. The opposite is true for Cluster 6 and hence we see the reverse trend in the gain errors. The average primary beam for each data segment at the center of Cluster 6 and Cluster 7 are plotted in Fig. B.2 and the anticorrelation of the gain calibration errors with the primary beam can be seen clearly as a function of time.

Appendix C: A-team subtraction trials




Fig. C.1. Standard deviation of the real part of the visibilities along frequency and baselines, as a function of time for different calibration trials. At each successive trial, only one setting was changed in order to understand its impact on the calibration. Trial 1 is the result of a fulljones calibration with Cyg A and Cas A subtracted when they are at elevations above 10 degrees. The subsequent changes in the other trials are described as follows. Trial 2: MAs 62, 64, 72 flagged, Trial 3: A-team included till 0 degree elevation, Trial 4: Tau A included, Trial 5: MA 6 flagged, Trial 6: MA 18 flagged, Trial 7: Vir A included, Trial 8: Calibration done in the diagonal mode.

We varied the number of stations to be flagged, the number of A-team sources to be included, the elevation range for which an A-team is included, and the calibration mode (fulljones or diagonal) in these trial runs and converged to the optimized calibration parameters by monitoring the calibration solutions, data statistics, and images at the end of each trial. As an example, Fig. C.1 shows the standard deviation of the real parts of the XX and YY calibrated visibilities, integrated along all frequencies and baselines, for these different trial runs. We see that it improves progressively as we converge to the final calibration settings. Many more trials were performed with more MAs flagged (MAs 1, 9, 22, 30, 56, 59, and 63), but they did not lead to any further improvement in the solutions, statistics, or images; hence, these stations were not flagged in the final data used for the analysis.