# EoR CodeX Meeting Summary
# Görlitz, Germany

June 23rd to 25th, 2025


**ORGANIZERS**

Nichole Barry - nichole.barry@unsw.edu.au
Carolin Höfer - hofer@astro.rug.nl

**PARTICIPANTS**

Jacob Burba - jacob.burba@manchester.ac.uk
Kariuki Chege - chege@astro.rug.nl
Joel Dunstan - joel.g.dunstan@curtin.edu.au
Sonia Ghosh (online) - soniaghosh@astro.rug.nl
Wilden Hidayat - wilhid@gmail.com
Nivedita Mahesh (online) - nmahesh@caltech.edu
Florent Mertens - florent.mertens@obspm.fr
Satyapan Munshi - munshi@astro.rug.nl
Dev Null - dev.null@curtin.edu.au
Ridhima Nunhokee - ridhima.nunhokee@curtin.edu.au
André Offringa - offringa@astron.nl
Shijiao Yin - yinshijiao@shao.ac.cn
Zhenghao Shu - zhenghao@shao.ac.cn

# Contents

# Foreword

The Epoch of Reionization Code Exchange (EoR CodeX) is an attempt to bring together software developers, statistical analysts, and future users to encourage growth and scientific collaboration between different subgroups. For our first-ever EoR CodeX, we decided to focus on the synthesis between MWA and LOFAR.

Our first meeting was held the week after the SKAO Science Meeting 2025 in the historic town of Görlitz, Germany. While initially scheduled to take place at the Obermühle Görlitz, an old flour mill on the river converted into a brewery, Carolin quickly realized that the download speeds in the quiet part of a quaint village were not sufficient. She quickly enabled a shift of location to the Horshel, right in the middle of the town and next to the roaring lion on the clocktower, which provided ambiance throughout the workshop.

For our first day, Kariuki and Carolin guided the group through a Nextflow tutorial of a standard LOFAR EoR analysis. This was the first tutorial of its kind (as far as we are aware) of LOFAR EoR data reduction after pre-processing. The group spent the day going through the tutorial on their individual computers and, naturally, a good-natured race-to-the-finish began. This was inevitably won by Dev Null, who rented an AWS server and destroyed anyone's hopes of chugging across the finish line with a turtle of a laptop. Honorable mention goes to Florent Mertens and Jacob Burba.

On the second day, Dev Null and Joel Dunstan gave tutorials on the MWA Rust stack pipeline and the MWA python stack pipeline, respectively. These are independent pipelines with different analysis philosophies and different code bases. Dev started us off with some MWA context, going into the details on the fun quirks of MWA data. Then, despite a supercomputer maintenance hiccup that threatened to derail our access to tutorial data, Dev successfully guided us through power spectrum generation with CHIPS using a hyperdrive output.

In the afternoon, Nichole gave a brief explanation of model simulations for calibration and subtraction using WODEN. Joel then explained the foundations, documentation, and basics of running PyFHD, which debuted for the tutorial as a complete Version 1. The group had no trouble running through the small bandwidth selection provided for the tutorial, taking less than a minute for most.

Our last day was spent spontaneously going through other analyses that could potentially be helpful, including contributions from Wilden, Jacob, André, and Florent. We discussed the wider goal of combining analyses from different instrumentation, and some initial progress was made. We finished the workshop by having a lovely dinner at our initial location, the Obermühle Görlitz.

Links to tutorials, video recordings, and relevant documentation are provided in the rest of the Meeting Summary, along with initial updates of progress and ideas. We hope to see you at the next meeting!

Nichole and Carolin

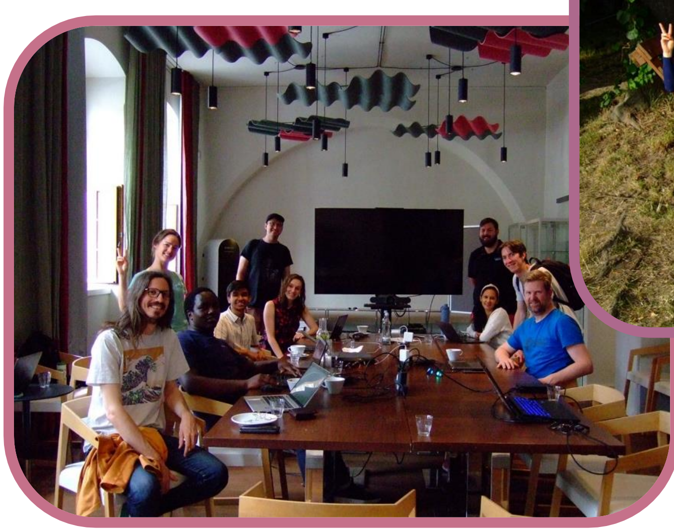

# Pipeline Tutorials

## Standard LOFAR EoR Analysis

Carolin Höfer and Kariuki Chege gave a Nextflow tutorial on the LOFAR EoR analysis pipeline. This implementation demonstrates direction-independent validation for fast time-varying gains, bandpass calibration, direction-dependent calibration, imaging, power spectrum estimation, collection of data quality statistics, and plotting. The recording and slides for the tutorial are available online.

Kariuki built this Nextflow pipeline, which relies heavily on DP3, AOFlagger, and pspipe. The demo uses a 4-minute dataset from the 3C196 field – the real data set is 6 hours long and is available upon request.

This was a wonderful first-attempt on such a detailed tutorial for the LOFAR analysis. Kariuki will be implementing some updates, including merging/splitting the subbands for correct spectral regularisation in the first and last calibration steps, as well as more documentation and examples.

## MWA Rust Pipeline

Dev Null gave a tutorial on running CHIPS with hyperdrive outputs, which is one of the two main pipelines used to reduce MWA data. Dev also gave an extended introduction to the MWA and the expected data quality. The recording for the tutorial is available online. The tutorial demo, MWA-context slides, and the presentation on CHIPS are also available.

Ridhima recently led a publication using this pipeline (Nunhokee et al., 2025). She performed extensive systematic mitigation, and that code is available for those keen to dive deeper into MWA data.

## MWA Python Pipeline

Joel Dunstan gave a tutorial on PyFHD, the python translation of FHD, which is one of the two main pipelines used to reduce MWA data. The recording for the tutorial is available online.

PyFHD can perform calibration, gridding, and image generation. It has been developed with MWA Phase I data in mind, but should be agnostic to the instrument as long as you have the proper beam models and model data. The code is publicly available and there is extensive documentation. The documentation contains installation, tutorials, a contribution guide, an IDL translation guide, change logs, API docs, testing/coverage results, and profiling.

The code for PyFHD now has a DOI attached to it, and there will be a JOSS paper soon to come. The test data used by GitHub Actions also has its own DOI – the main purpose of which is to allow for an external test data download (sneaky...). PyFHD is now past version 1 and is ready for robust use. PyFHD is available via pip and through dockerhub. Joel recommends `uv` or `poetry` for all your python versioning needs.

Joel will no longer be maintaining PyFHD, but he is always available by email or via the Pyfhd channel on the EoRAnalysis slack (shared slack for HERA and the MWA). The mantle will pass to

Bryna Hazelton, who will be investigating the excess power in the HEALPix generation of PyFHD and will be translating deconvolution.

The citation advice for PyFHD will change soon once a JOSS paper is published. Please see the PyFHD README for updates.

# Other Analysis Summaries and Progress

## GRANDWin

Wilden presented the newly named Gain-based RFI Analysis using Normalized Deviation with Winsorization, or GRANDWin. Even using all of the RFI techniques that exist on uncalibrated data, there still exists some low-level RFI near the FM band. To further remove this persistent RFI, Wilden is using his winsorization algorithm to find potential RFI outliers in the real component of the calibration gain solutions. This selects and removes visibilities associated with outlier gain solutions.

Initial results look promising, and everyone was impressed with how much the power spectra at FM bands improved. No one had actually seen a power spectrum at those frequencies! The initial code-base is currently available and will be updated as Wilden progresses towards a publication.

## pspipe

Florent gave a detailed walkthrough of the pspipe tutorial, which generates power spectra from calibrated visibilities using python. All code is publicly available, and can be easily installed with `pip`. The documentation is currently a work in progress.

Florent took the calibrated and ionospherically corrected visibilities from Dev's MWA Rust Pipeline tutorial and built a MWA pspipe tutorial in less than a day, immediately causing awe for his supreme efficiency. Florent then compared the power spectra from pspipe for this observation with the output from the CHIPS tutorial.

There were some notable differences that necessitate followup. Firstly, the horizon reflection around the coarse band harmonics is symmetric in pspipe, whereas it only occurs on the side with higher modes in CHIPS. This difference has also been see in comparisons with eppsilon (part of the MWA Python pipeline) and CHIPS. Secondly, pspipe has far more negative bins than CHIPS, which is either very good or very bad, depending. It was generally agreed that more data would be required – one observation has very high thermal noise, which could contribute to a seemingly large excess of negative bins.

The thermal noise estimate between CHIPS and pspipe is very similar, indicating a good agreement on normalization. That being said, Florent mentioned that the primary beam approximation in the power spectra normalization could be improved.

If you want to use pspipe, please also see the basic tutorial and the LOFAR tutorial. If you use pspipe, please cite Mertens et al. (2020) and Mertens et al. (2025).

## DP3

André attempted to use DP3 to perform direction-dependent calibration on MWA data, and walked us through his initial findings.

First of all, a measurement set with beam information is required. A measurement set made with Birli should automatically include this, but measurement sets made with other software (i.e. pyuvdata, casa) will require an extra fixing script. André wrote a script many eons ago called `fixmwams`, and it is located in the MWA-shared cotter repository. Fortunately, there is a docker container. A snippet of Dev's commands to make DP3-supported measurement sets is as follows:

```
# step1: casa to convert uvfits to ms
docker run --rm -it -v$PWD:$PWD -w$PWD --entrypoint=casa d3vnull0/casa -c
    "importuvfits('${uvfits}', '${ms}')"

# step2: fixmwams
docker run -it --rm -v$PWD:$PWD -w$PWD --entrypoint=fixmwams mwatelescope/cotter ${ms}
    ${metafits}

# This is the same as above if you use singularity / apptainer
singularity exec -B$PWD -B${outdir:-$PWD} -W$PWD --cleanenv docker://d3vnull0/casa casa
    -c "importuvfits('${uvfits}', '${ms}')"
singularity exec -B$PWD -B${outdir:-$PWD} -W$PWD --cleanenv
    docker://mwatelescope/cotter fixmwams ${ms} ${metafits}
```

André used a command line cluster/model editor from the lofartools repository, which can 1) cluster a model in direction automatically given the desired number of clusters, or 2) manually dived the model in clusters. André tried:

```
model_in="images/1069761080-sources-pb.txt"
editmodel -t 1 -set-cluster target -skymodel models/target.txt -near 00h00m00.0s
    -27d00m00s 20deg ${model_in}
editmodel -t 0.4 -set-cluster south -skymodel models/south.txt -near 00h10m03.077s
    -59d56m19.53s 7deg ${model_in}
editmodel -skymodel models/two-directions.txt models/target.txt models/south.txt
#editmodel -save-clusters models/two-directions-clusters.ann -kvis
    models/two-directions-kvis.ann models/two-directions.txt
render -t images/1069761080-MFS-image-pb.fits -o models/two-directions.fits -r
    models/two-directions.txt
```

For imaging, he used the command:

```
wsclean \
  -save-source-list \
  -size 10000 10000 -niter 1000000 \
  -auto-mask 4 -auto-threshold 0.5 \
  -mgain 0.8 -scale 30asec \
  -multiscale \
  -apply-primary-beam -reuse-primary-beam \
  -name images/1069761080 \
  -pb-grid-size 1250 \
  -channels-out 8 -join-channels -fit-spectral-pol 2 \
  -deconvolution-channels 2 \
```

```
-parallel-deconvolution 1024 \
-horizon-mask 0 \
hyp_1069761080_ssins_30l_src8k_300it_8s_80kHz.ms/
```

André split up the MWA image into two parts for DDcal. Unfortunately, this did not seem to yield logical results, but André notes that further work must be done to figure out the proper MWA configuration parameters. The infrastructure is now there to find the proper number of directions.

The documentation and source code for DP3 is available online, and is maintained and developed by a team at ASTRON.

## BayesEoR

Jacob showed us his work on BayesEoR, which is a package that estimates the power spectrum using a Bayesian approach. It performs it's estimation using model visibilities of the foregrounds and the EoR signal, calculating the posterior distribution over the power spectrum.

BayesEoR is publicly available, with test data and specific documentation available for learning how to run the code. BayesEoR doesn't explicitly require GPUs, but you will have a hard time if you try to do anything scientific without them. Please see the full documentation for more information.

BayesEoR is still in development. See Jacob's recent paper, Burba et al. (2023), for an understanding of the difficulties faced in sky foreground modeling and how that affects the result. Four papers describe the approach taken with BayesEoR (Burba et al., 2023; Sims et al., 2019; Sims & Pober, 2019; Sims et al., 2016) – please cite them if you use the package.

## hydra-pspec

Jacob also described hydra-pspec, a power spectra estimator that uses Gibbs sampling. It can filter foregrounds and in-paint flagged data. It is publicly available, and documentation will come shortly. In the meantime, there is currently a tutorial via SRCNet which walks through analysis with hydra-pspec.

Jacob mentioned that Bella Nasirudin is working on analyzing real HERA data with the pipeline, which will be a huge step forward. Please cite Burba et al. (2024) and Kennedy et al. (2023) if you use hydra-pspec.

## WODEN

Nichole gave a brief introduction to WODEN, which is a very precise simulator and now the official model generation code of PyFHD. Jack Line is the developer of WODEN. WODEN can generate instrumentally precise model simulations for the MWA, EDA2, SKA, LOFAR, and HERA, and it hypothetically can generate simulations for any instrument added in Everybeam or pyuvbeam. There is extensive documentation, including a variety of worked examples.

There are a variety of docker images available for ease-of-use. Jack also has a variety of catalogs available for the MWA which may be useful. If you use WODEN, please cite Line (2022).

## AOFlagger

Even though we didn't do a personalized walk-though, André has provided details on AOFlagger for those interested to pursue this in their own time. Documentation and source code is available online.

André notes that to flag on the standard deviation ("coherently averaged"), there are multiple options:

- Use `rfigui` to open a measurement set and select in the file->open->integration box to average over stddev of baselines.
- Use AOFlagger to do this; for this a Lua file is required which is set up to flag on the stddevs. André can provide such a file, but it is also reasonably well described on the Lua pages of AOFlagger's manual.
- To integrate this in Birli, Birli could itself calculate the standard deviations and use AOFlagger's API to flag these standard deviations. This is a bit more work but avoids dependencies on measurement sets or casacore. The advantage of using AOFlagger like this is that it uses pattern searching to catch RFI, to be more sensitive towards RFI that, in time-frequency space, also morphologically looks like RFI.

## MWA Flagging

Unfortunately, we ran out of time to hear from Dev about the MWA flagging scheme. Dev has an in-depth tutorial that they built for a radio school which goes into great detail about the various flagging required for MWA data. The demo is available online and Dev welcomes any questions via email.

# Future Interests

**Nichole Barry**
NB wants to try to get a representative beam into FHD/PyFHD without breaking the bank, which will allow analysing sims/data through FHD/PyFHD pipes. She also wants to use the DI LOFAR step in the tutorial to compare to the auto correlation calibration in FHD/pyFHD. NB will be putting FHD/PyFHD outputs through pspipe, which should be trivial. Her lofty goal is to use the LOFAR tutorial to analyse some MWA data and compare the noise estimation in the power spectrum for various pointings.

**Jacob Burba**
JB will be working on analyzing MWA validation simulations from Line et al. (2024) with NB and DN and then move to an analysis of real MWA data (likely EoR0) with both BayesEoR and hydra-pspec. HERA's first public data release (H1C IDR3.2) should be coming out soon (JB pushing on this from the SRCNet side). JB wants to try putting HERA data through pspipe and will work with FM on this.

**Kariuki Chege**
KC is interested in looking at the different flagging strategies mentioned, which might be helpful in the post-processing flagging of the 3C196 field data. KC mentions that a way forward might be looking into adopting the SSINS and/or EAVILS strategy – this is already in pspipe but needs a little work to work for his current data format. Additionally, KC will be updating the singularity container with all the LOFAR tools including the latest updates from André to ease calibration of MWA data with LOFAR tools, and he is happy to get involved in this as well.

**Joel Dunstan**
JD wants to get an idea of how to run LOFAR data through PyFHD. JD also wants to add any issues to the PyFHD Github that people run into.

**Sonia Ghosh**
SG wants to understand the synergies between LOFAR and MWA pipelines, how the MWA observes, how the data is structured, and the various calibration/power spectra strategies.

**Wilden Hidayat**
WH is interested to try the direction-dependent calibration strategy of LOFAR DDECal on the MWA ultra-low (75-105 MHz) drift scan data. WH notes that the work required involves converting files and getting the proper MWA sky model into the pipeline. WH is also interested in integrating pspipe in the drift scan data analysis.

**Carolin Höfer**
CH wants to use a MWA test data set and run through the LOFAR pipeline for cross-validation of pipelines. Specifically, she wants to calibrate MWA data with LOFAR software.

**Nivedita Mahesh**
In general, NM wants to understand the potential of utlizing LOFAR/MWA pipelines on OVRO-LWA data. Specifically, NM wants to try pspec on the OVRO-LWA simulated data. In addition, she wants to implement the OVRO-LWA beam in Everybeam.

**Florent Mertens**
FM is interested to try a simple `ml_gpr` configuration on the calibrated MWA dataset(s)! FM is also happy to learn more about the RFI strategies and metrics used in MWA and apply them to e.g. NenuFAR datasets, and is interested to follow the development of hydra-pspec.

**Satyapan Munshi**
SM is interested to in DDECal on MWA data, specifically the identification of optimal approaches for both DI/DD calibration, impact of frequency smoothness constraints, and baseline cuts. SM is also interested to try to implement GPR on pspipe gridded data cubes from MWA data.

**Dev Null**
DN wants to get a better AOFlagger strategy for implementation in Birli. DN also wants to infill missing flagged MWA data for better power spectra. They also want to get better I/O performance from the Rust casacore implementation.

**Ridhima Nunhokee**
RN wants to push more MWA data through pspipe and compare the results with CHIPS for the EoR0 and EoR1 fields. She also want to incorporate spectral regularisation for direction-independent calibration to MWA calibration solutions. RN is potentially thinking of incorporating the LOFAR flagging strategies to MWA, including using per baseline delay-transform-based flagging and applying AOFlagger before and after pre-processing and after calibration. She would be happy to contribute to applying BayesEoR on MWA data.

**André Offringa**
AO wants to apply MWA power spectra software to LOFAR data, and get DP3 DD calibration working on MWA data. Other goals include an improved AOFlagger strategy for MWA (and LOFAR).

**Shijiao Yin**
SY hopes to learn more about the LOFAR and MWA pipelines, and to understand the steps involved for generating the beam, implementing gridding, and power spectra estimation.

**Zhenghao Shu**
ZS want to see if the pipelines can be applied on simulation data, and if that works, try to develop an end-to-end pipeline from sky model building to EoR power spectrum reconstruction.

# References


Burba, J., Sims, P. H., & Pober, J. C. (2023). All-sky modelling requirements for Bayesian 21â cm power spectrum estimation with bayeseor. *Monthly Notices of the Royal Astronomical Society*, *520*(3), 4443–4455. https://doi.org/10.1093/mnras/stad401

Burba, J., Bull, P., Wilensky, M. J., Kennedy, F., Garsden, H., & Glasscock, K. A. (2024). Sensitivity of Bayesian 21â cm power spectrum estimation to foreground model errors. *Monthly Notices of the Royal Astronomical Society*, *535*(1), 793–806. https://doi.org/10.1093/mnras/stae2334

Kennedy, F., Bull, P., Wilensky, M. J., Burba, J., & Choudhuri, S. (2023). Statistical Recovery of 21 cm Visibilities and Their Power Spectra with Gaussian-constrained Realizations and Gibbs Sampling. *The Astrophysical Journal Supplement Series*, *266*(2), 23. https://doi.org/10.3847/1538-4365/acc324

Line, J. L. B., Trott, C. M., Cook, J. H., Greig, B., Barry, N., & Jordan, C. H. (2024). Verifying the Australian MWA EoR pipeline I: 21-cm sky model and correlated measurement density. *Publications of the Astronomical Society of Australia*, 1–13. https://doi.org/10.1017/pasa.2024.31

Line, J. L. B. (2022). 'WODEN': A CUDA-enabled package to simulate low-frequency radio interferometric data. *Journal of Open Source Software*, *7*(69), 3676. https://doi.org/10.21105/joss.03676

Mertens, F. G., Mevius, M., Koopmans, L. V. E., Offringa, A. R., Mellema, G., Zaroubi, S., Brentjens, M. A., Gan, H., Gehlot, B. K., Pandey, V. N., Sardarabadi, A. M., Vedantham, H. K., Yatawatta, S., Asad, K. M. B., Ciardi, B., Chapman, E., Gazagnes, S., Ghara, R., Ghosh, A., … Silva, M. B. (2020). Improved upper limits on the 21 cm signal power spectrum of neutral hydrogen at z â 9.1 from LOFAR [Publisher: Oxford Academic]. *Monthly Notices of the Royal Astronomical Society*, *493*(2), 1662–1685. https://doi.org/10.1093/mnras/staa327

Mertens, F. G., Mevius, M., Koopmans, L. V. E., Offringa, A. R., Zaroubi, S., Acharya, A., Brackenhoff, S. A., Ceccotti, E., Chapman, E., Chege, K., Ciardi, B., Ghara, R., Ghosh, S., Giri, S. K., Hothi, I., HÃ¶fer, C., Iliev, I. T., JeliÄ, V., Ma, Q., … Yatawatta, S. (2025). Deeper multi-redshift upper limits on the epoch of reionisation 21 cm signal power spectrum from LOFAR between z = 8.3 and z = 10.1. *Astronomy & Astrophysics*, *698*, A186. https://doi.org/10.1051/0004-6361/202554158

Nunhokee, C. D., Null, D., Trott, C. M., Barry, N., Qin, Y., Wayth, R. B., Line, J. L. B., Jordan, C. H., Pindor, B., Cook, J. H., Bowman, J., Chokshi, A., Ducharme, J., Elder, K., Guo, Q., Hazelton, B., Hidayat, W., Ito, T., Jacobs, D., … Zheng, Q. (2025). Limits on the 21 cm power spectrum at z=6.5-7.0 from mwa observations. https://arxiv.org/abs/2505.09097

Sims, P. H., Lentati, L., Pober, J. C., Carilli, C., Hobson, M. P., Alexander, P., & Sutter, P. M. (2019). Bayesian power spectrum estimation at the Epoch of Reionization. *Monthly Notices of the Royal Astronomical Society*, *484*(3), 4152–4166. https://doi.org/10.1093/mnras/stz153

Sims, P. H., Lentati, L., Alexander, P., & Carilli, C. L. (2016). Contamination of the Epoch of Reionization power spectrum in the presence of foregrounds. *Monthly Notices of the Royal Astronomical Society*, *462*(3), 3069–3093. https://doi.org/10.1093/mnras/stw1768

Sims, P. H., & Pober, J. C. (2019). Joint estimation of the Epoch of Reionization power spectrum and foregrounds. *Monthly Notices of the Royal Astronomical Society*, *488*(2), 2904–2916. https://doi.org/10.1093/mnras/stz1888

# Acknowledments

Nichole Barry would like to acknowledge Carolin Höfer for bearing the logistical burden of the meeting, including the last minute change of venue and the countless phone calls in German. Nichole Barry would also like to acknowledge the Scientia program at the University of New South Wales, which generously funded this endeavor as a part of her professional development.

Credit for photo page: Wilden Hidayat and André Offringa.

Credit for opening page: Old Bridge and St. Peter's Church, Gorlitz, Silesia, Germany., ca. 1890. Photograph. https://www.loc.gov/item/2002720701/