AI for LOFAR calibration

By: Jurjen de Jong

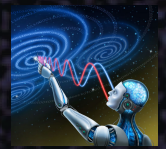

Content

1. LOFAR high-resolution imaging
2. Calibration
3. Machine learning for calibration validation
4. Future ideas with GenAI

LOw-Frequency ARray (LOFAR)

- ★ Pan-European radio telescope
- ★ Degree-scale field-of-view
- ★ High sensitivity
- ★ 10 - 250 MHz

LOFAR Two-Metre Sky Survey (LoTSS)

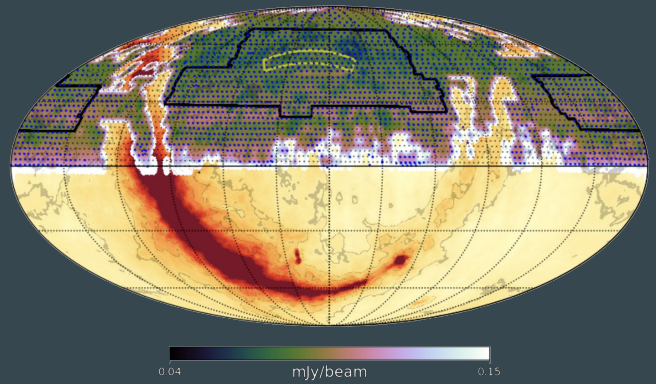

- ★ 2013 - Present
  (LoTSS; Shimwell et al. 2017; 2019; 2022)
- ★ 6” angular resolution
- ★ 80 μJy/beam sensitivity
- ★ 150 MHz

International LOFAR → 0.3” resolution

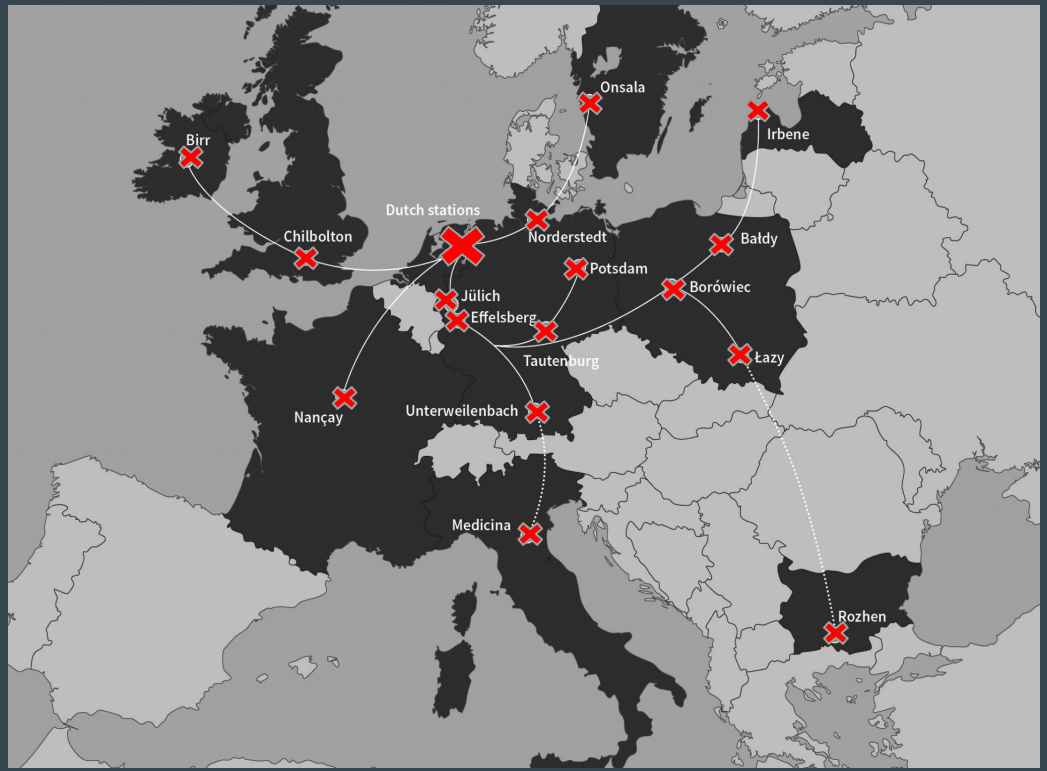

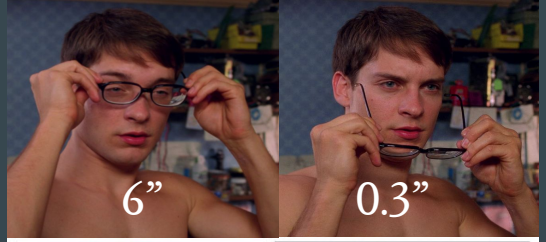

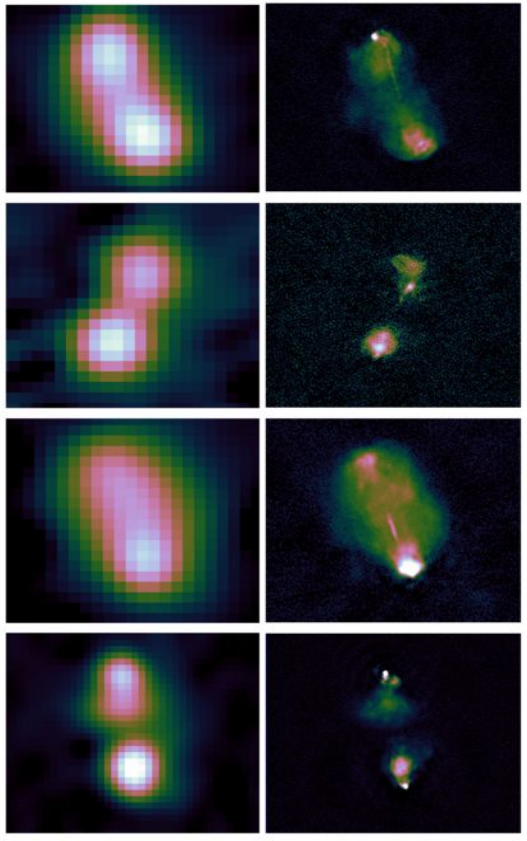

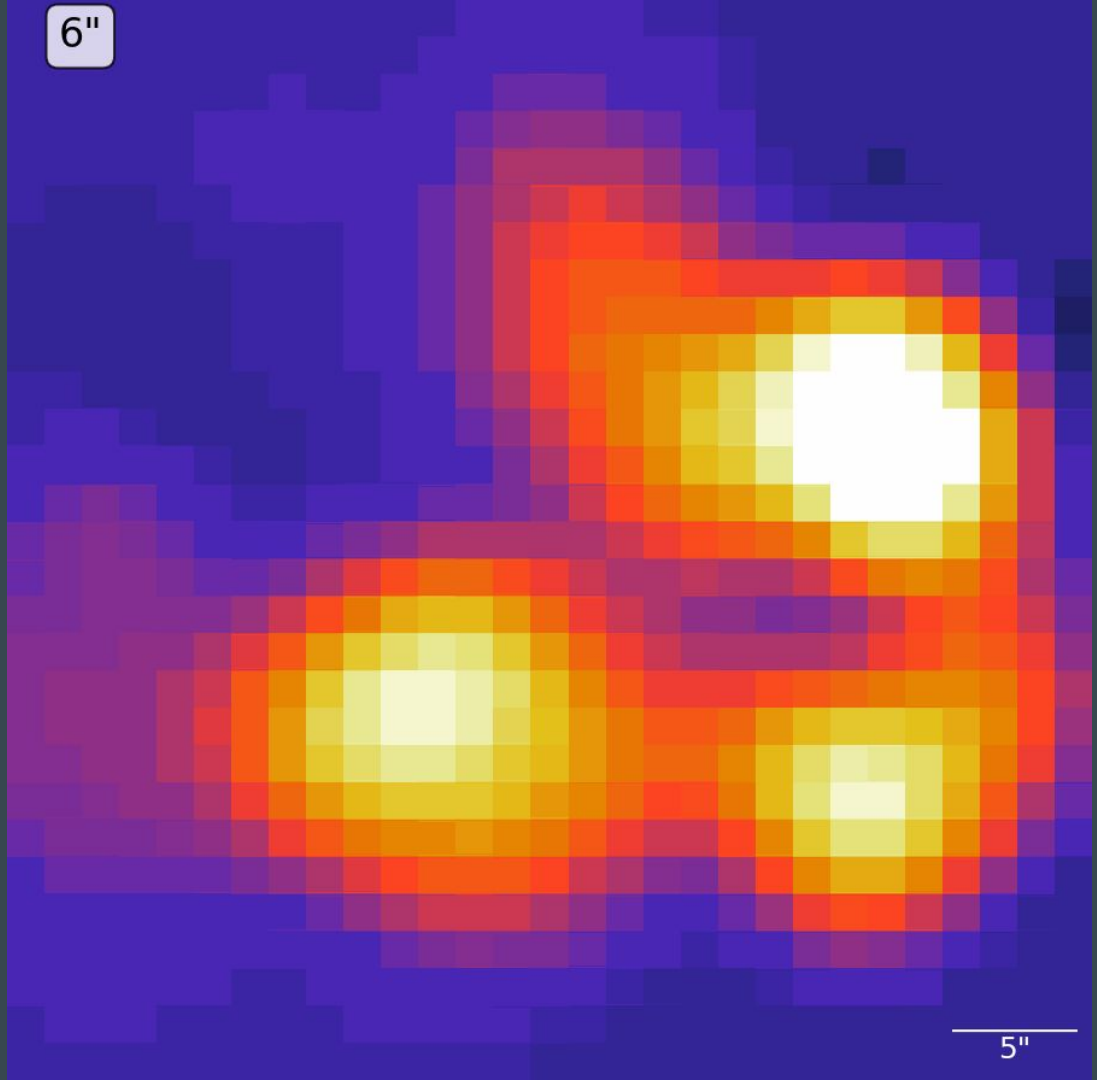
6"
5"

Wide-field high-resolution imaging (8 billion pixels)

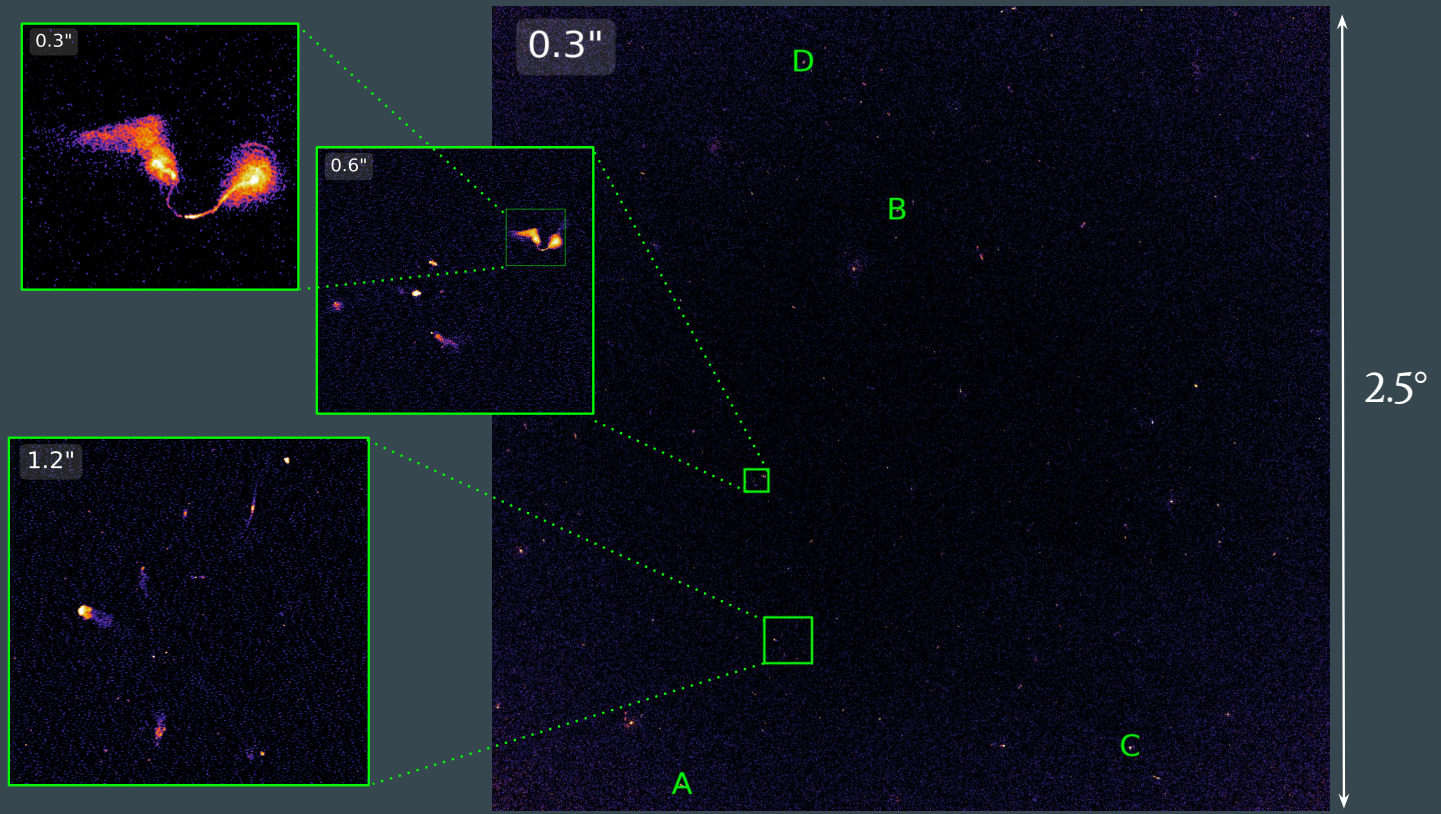

Pipeline for the International LOFAR Telescope (PILOT)

(van der Wild et al. subm.)

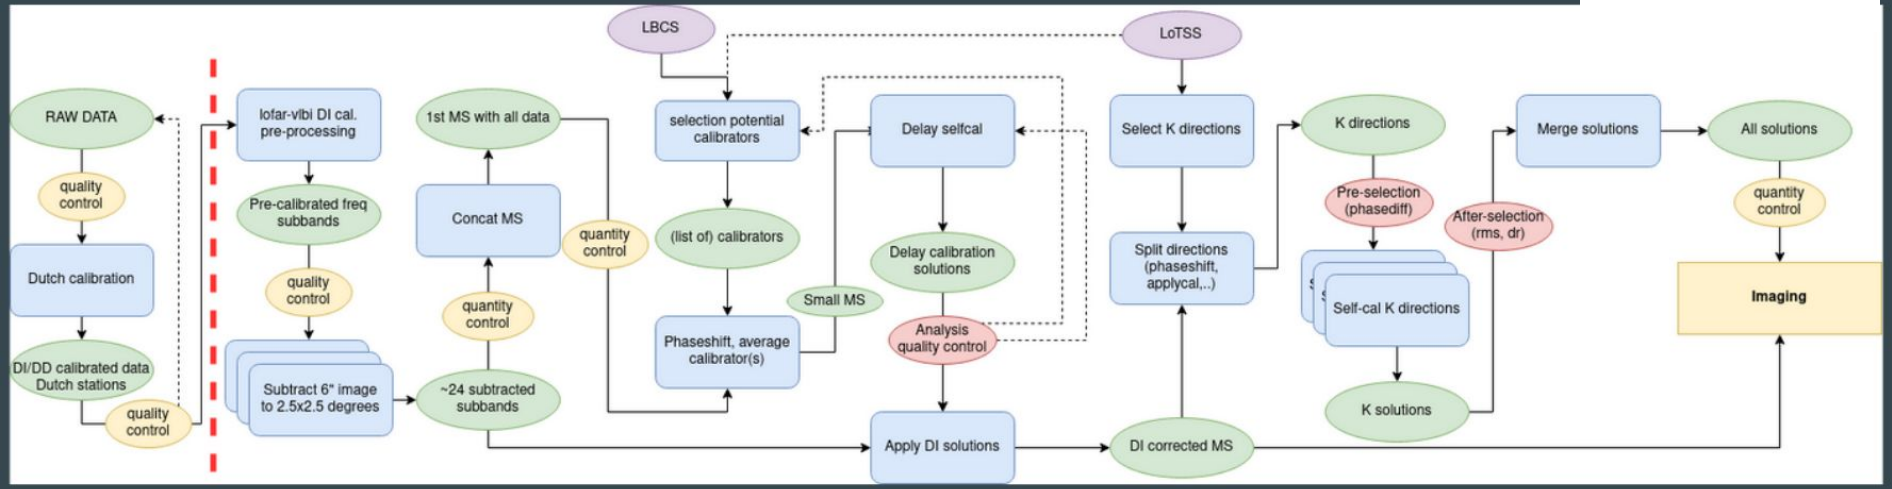

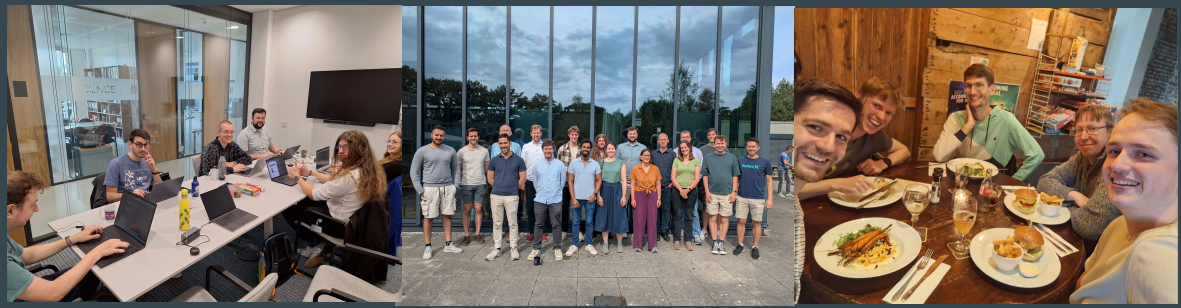

Computational cost (de Jong et al. 2025b)

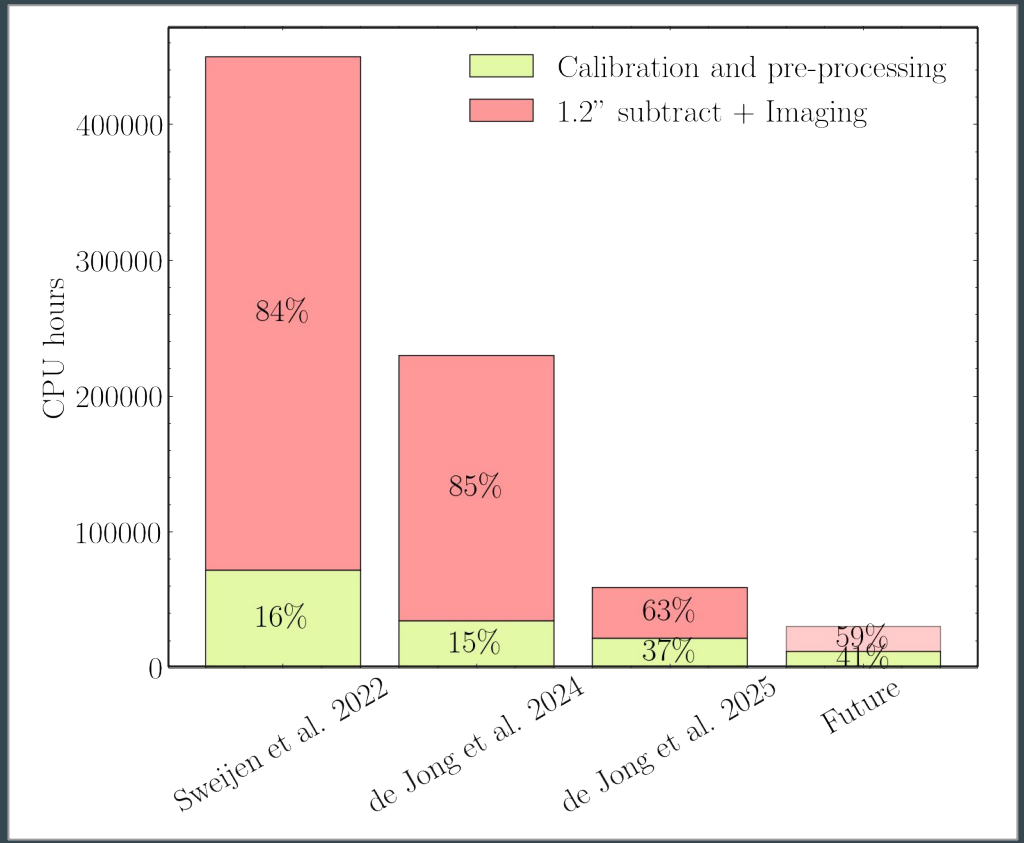

Direction-independent calibration

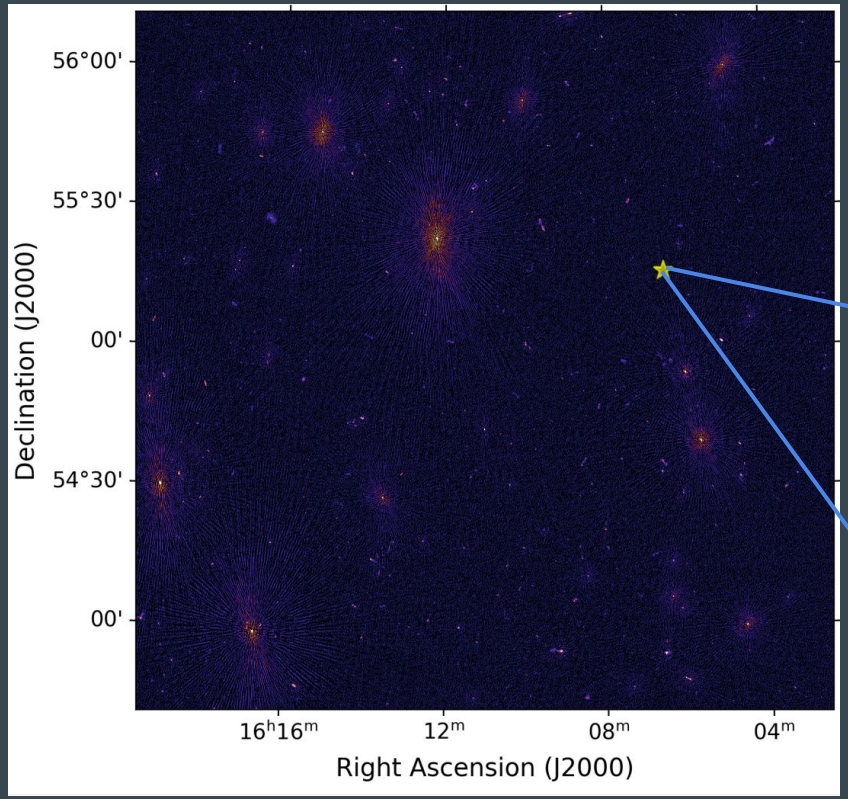

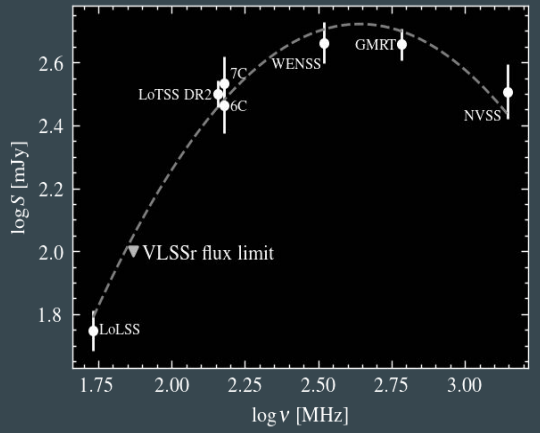

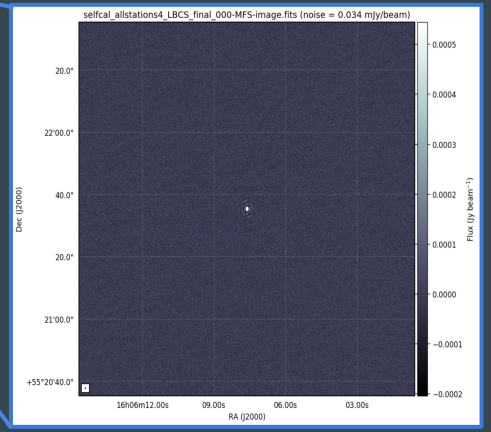

Direction-dependent self-calibration

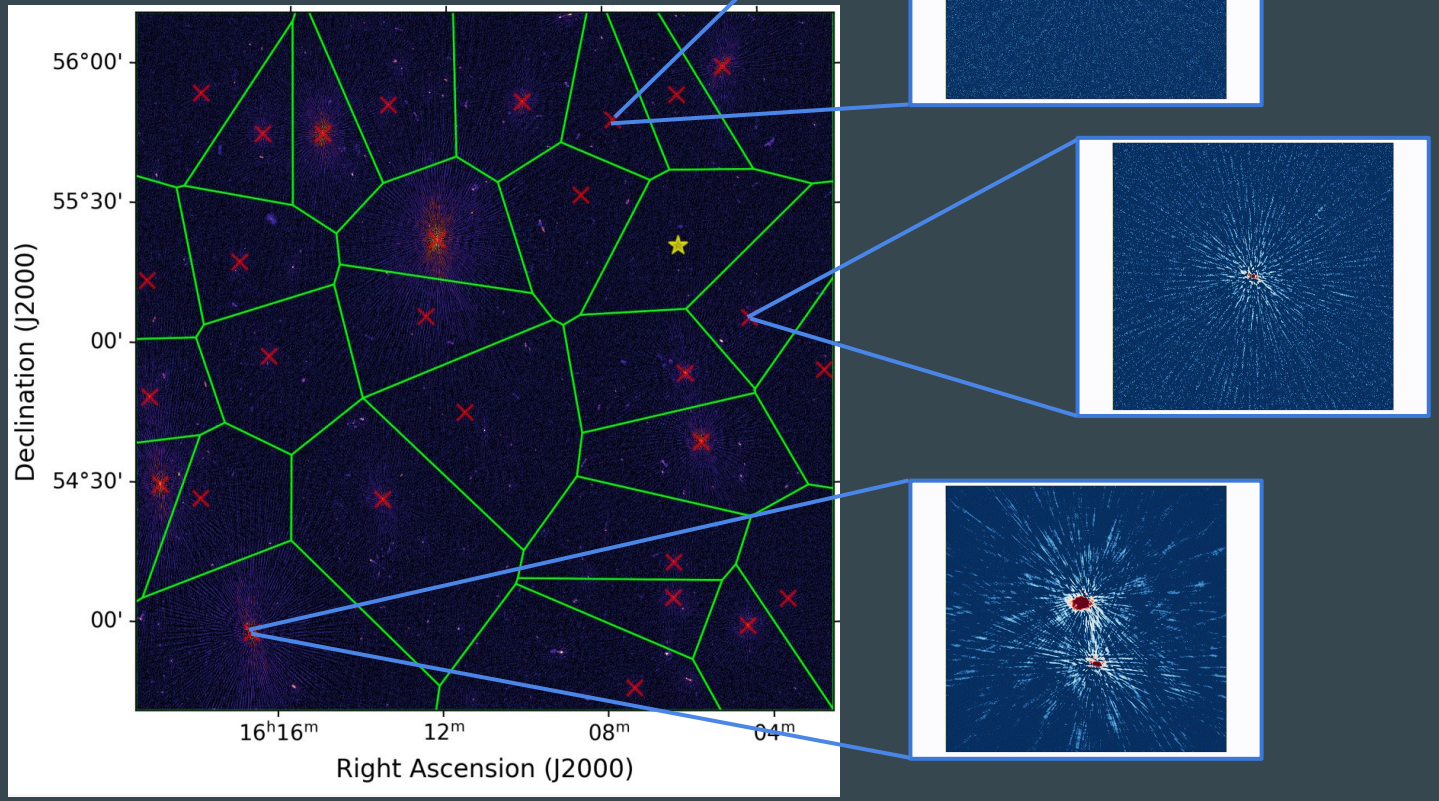

When are we happy?

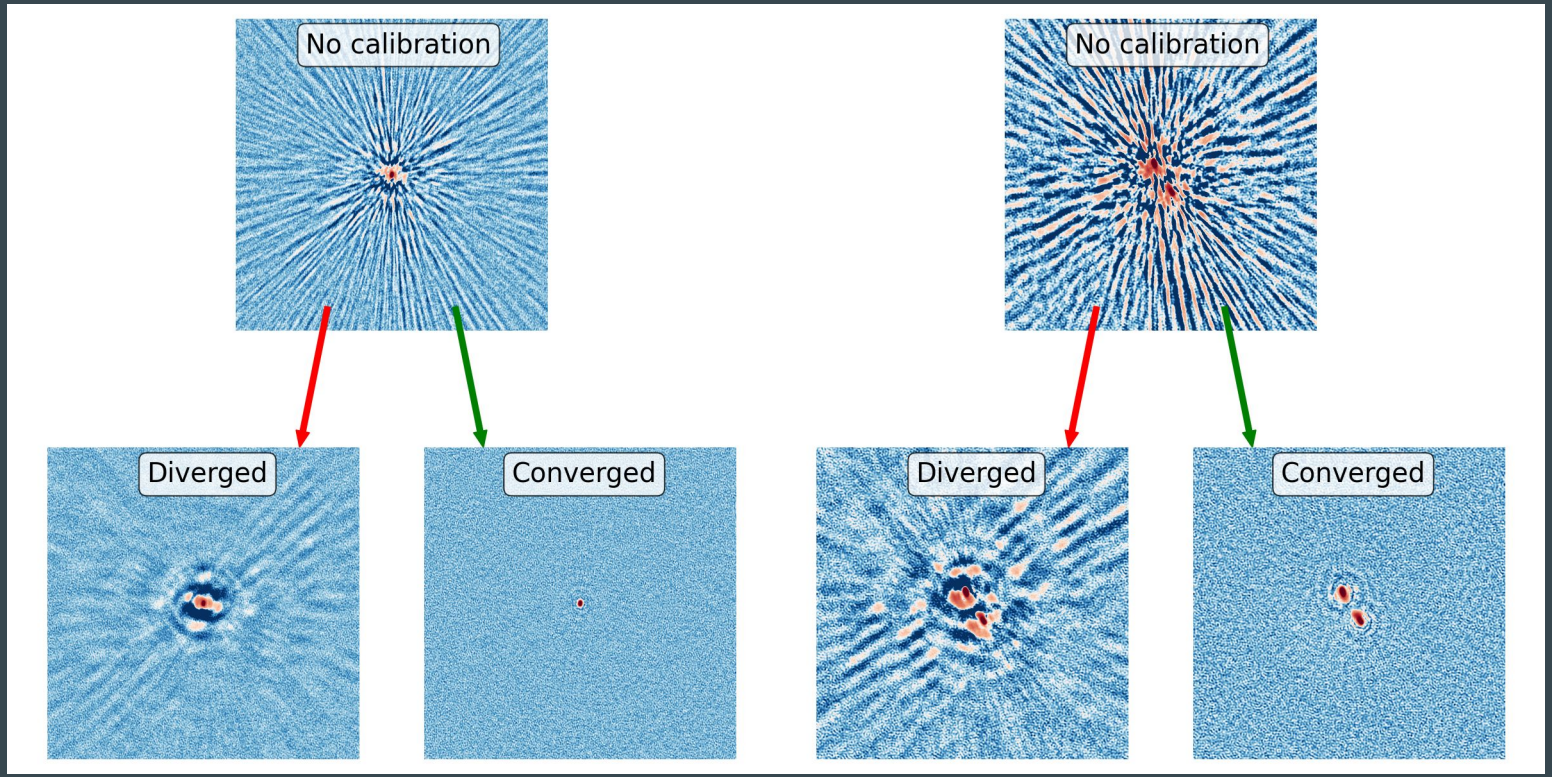

Self-calibration validation

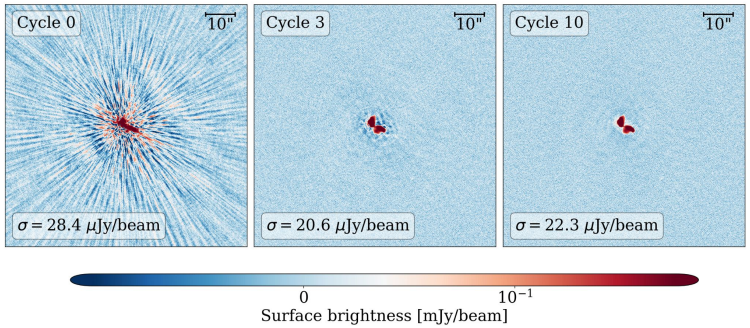

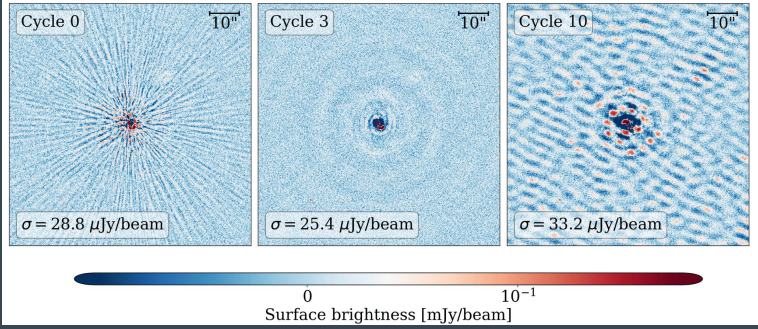

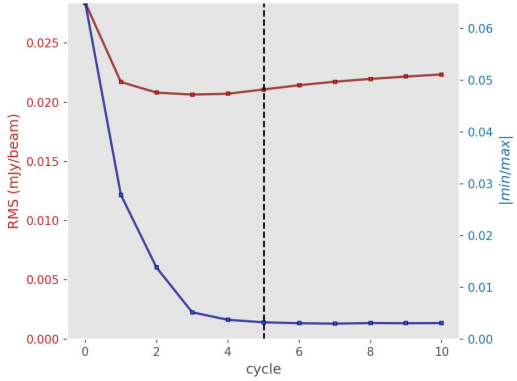

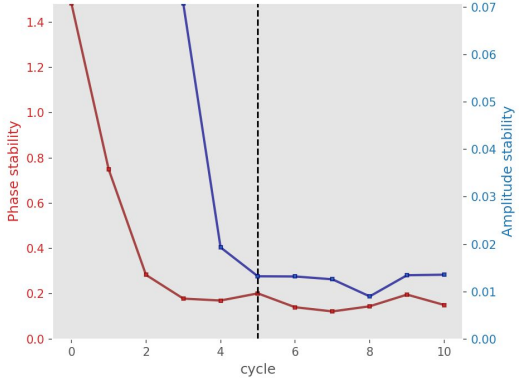

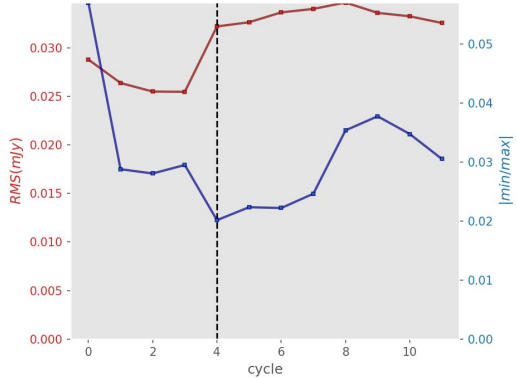

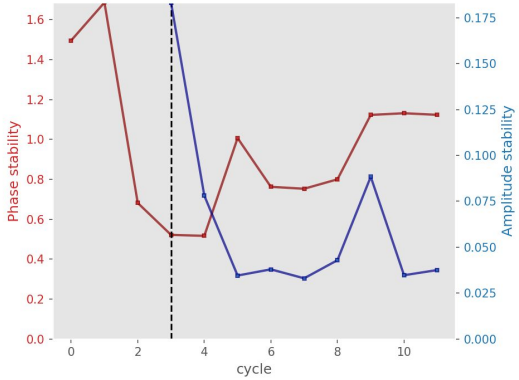

Validation with machine learning

- 3,000 labelled images
- Different calibration cycles
- Transformer model
- DINOv2

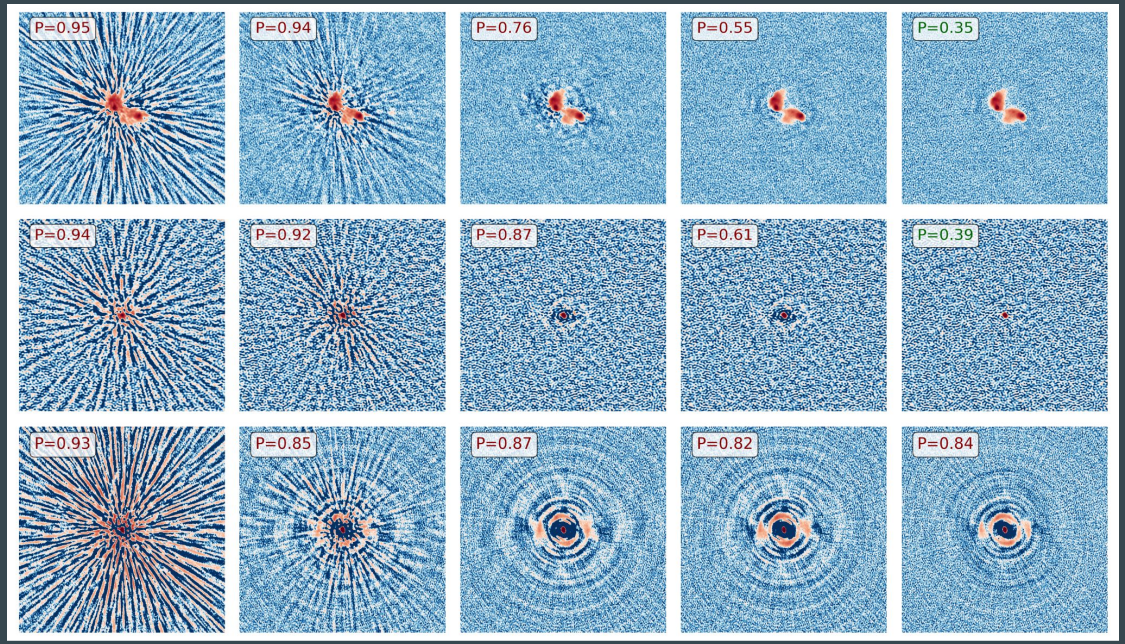

Validation with machine learning

- 3,000 labelled images
- Different calibration cycles
- Transformer model
- DINOv2

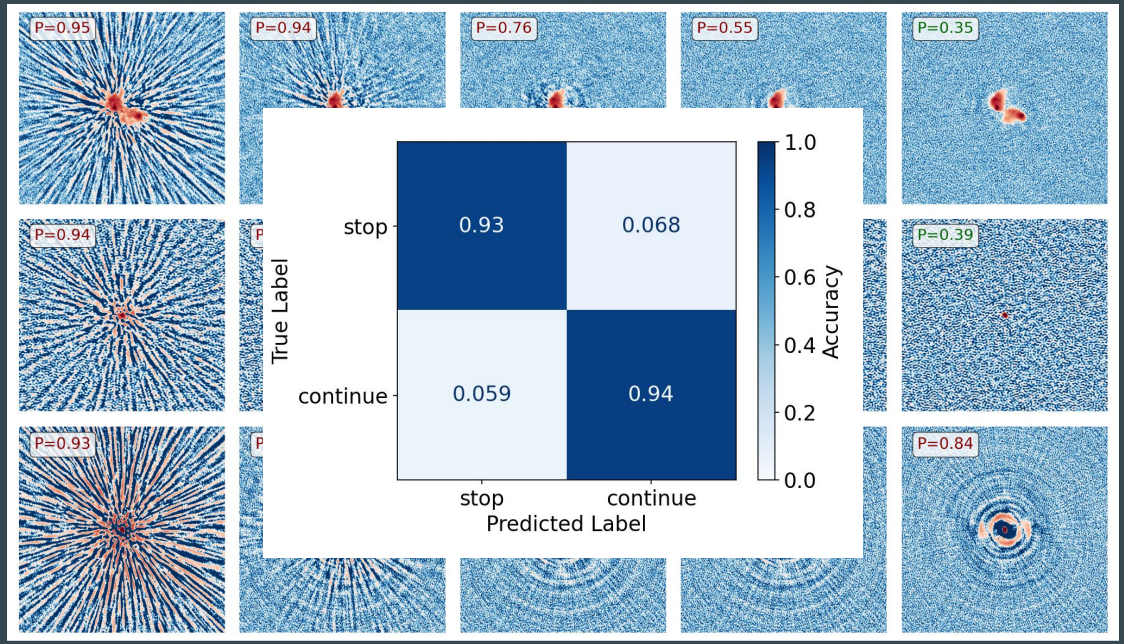

Transformer models

- “**Attention** is all you need” paper (Vaswani et al. 2017) – Google
  ChatGPT: “A model that looks at *everything at once* and learns *which parts matter most* using attention”

-

- More efficient processing and parallelism (perfect for GPUs)

- Images → Vision Transformers (ViT)

Transformers (for language)

1. Semantical meaning is determined by its context:

*I sat down at the **bank** and watched the water flow*

vs

*I went to the **bank** to deposit some money*

2. First bank is likely a **riverbank** , second a **financial bank**

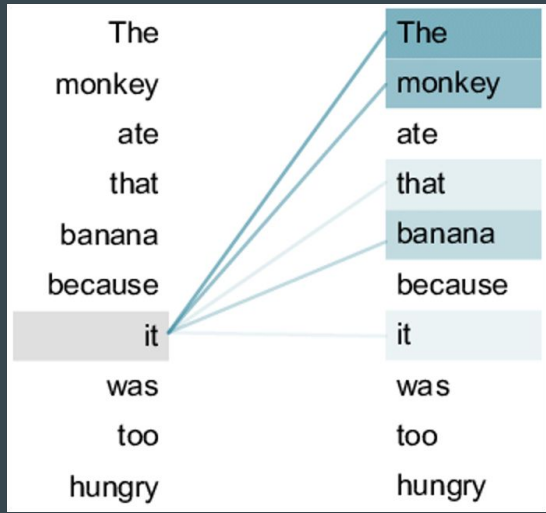

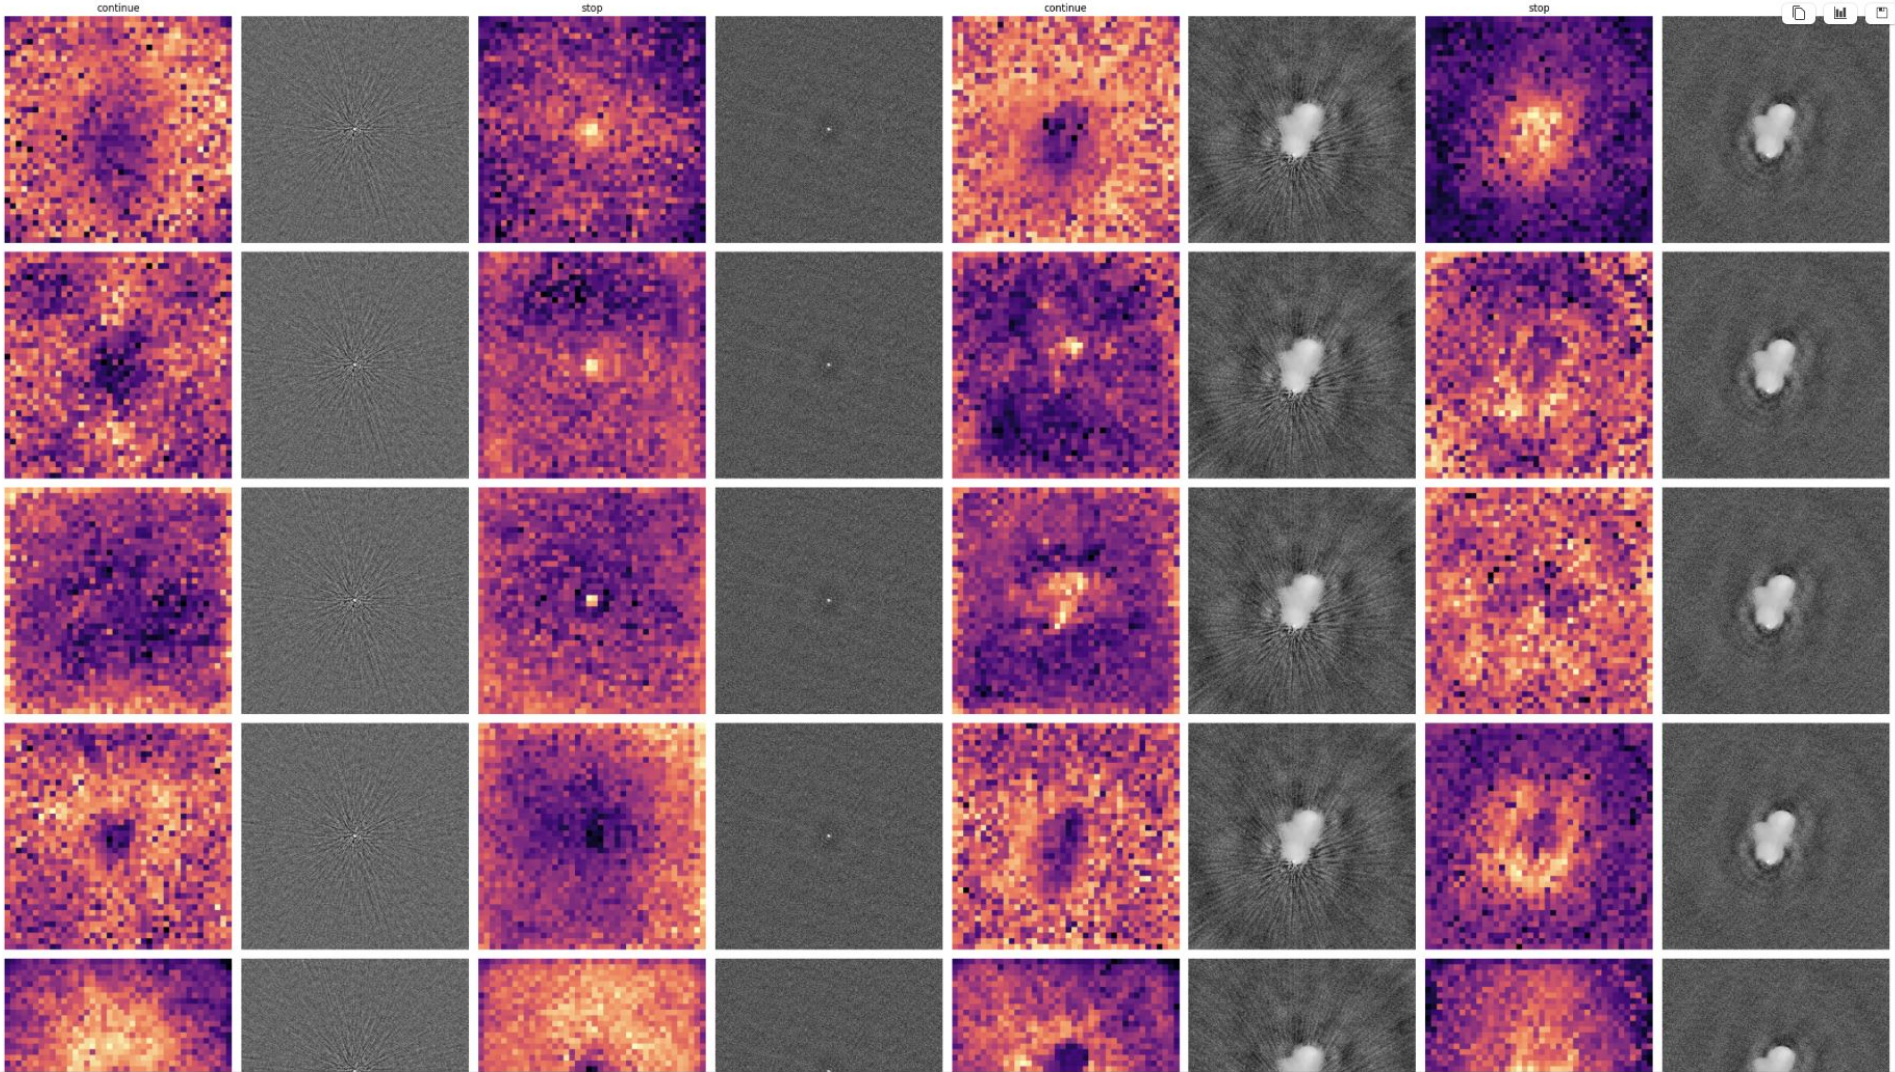
continue
stop
continue
stop

Visual Transformer: DINOv2 (Meta; Oquab et al. 2023)

1. Pre-trained ViT with 1.1B params, trained on 142M images
   Combination of self-curated and pre-curated cleaned datasets (e.g. ImageNet, Google Landmarks)

2. Added two-layered perceptron (ReLU + Sigmoid output)

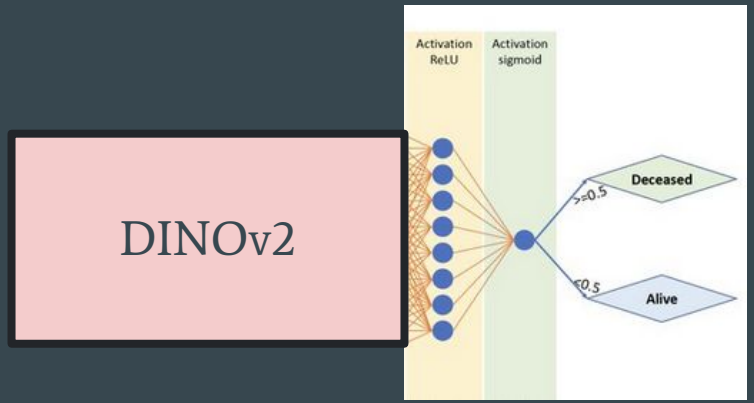

Preventing Overfitting

1. **Data augmentation**
   a. Mirroring, rotation, cropping, etc.
2. **Dropout**
   a. Ignore certain outputs with some probability *p*
   b. This prevents the model from focussing too much on some features
3. **Label smoothing** : turn hard labels into soft labels
   a. e.g. turn stop/True/1 to 0.9, and continue/False/0 to 0.1
   b. This prevents overconfidence
   c. Useful for noisy labeling

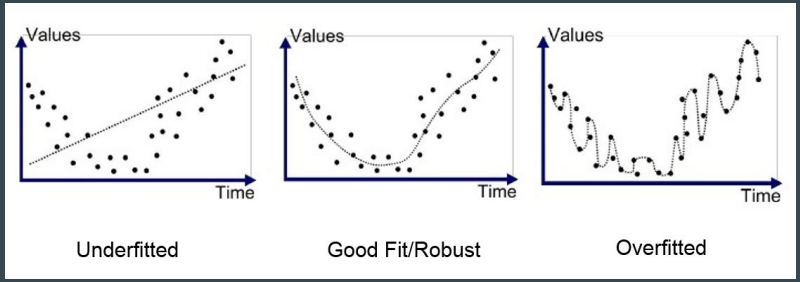

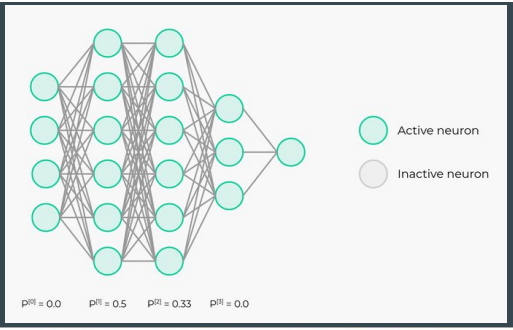

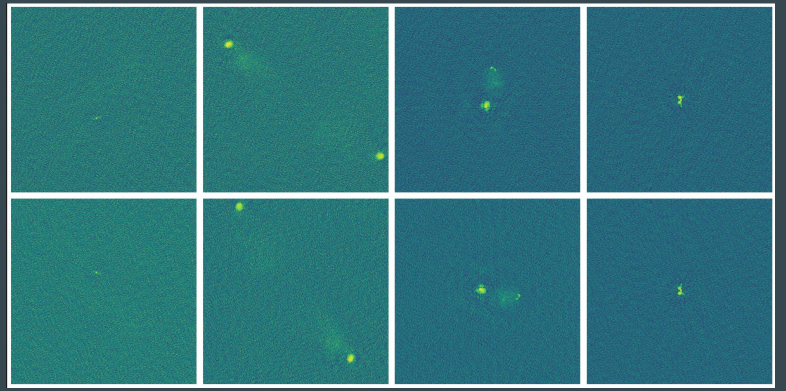

Potential issues

- Not ‘general’ enough
- Hasn’t ‘seen’ unexpected image artefacts
- Complex extended sources

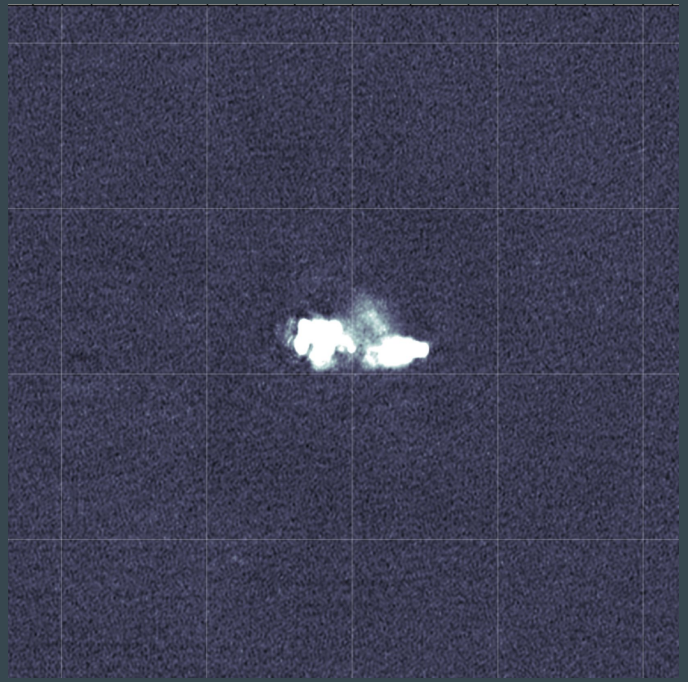

Future work?

GenAI

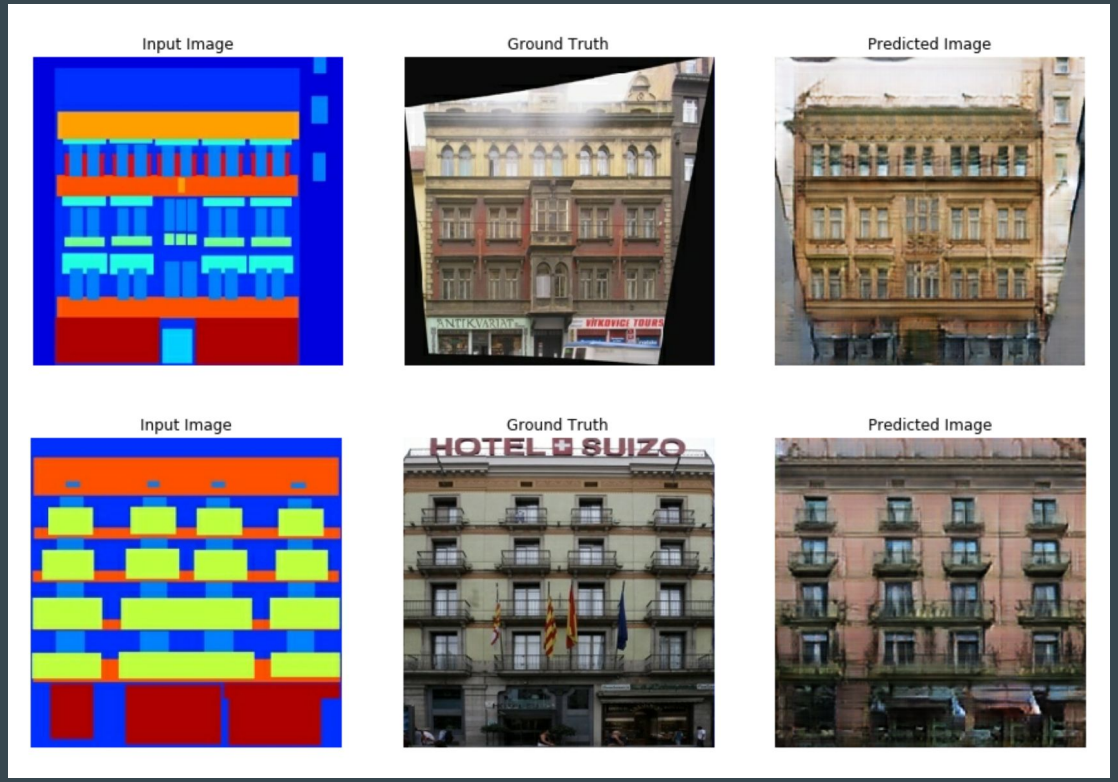
Input Image
Ground Truth
Predicted Image
Input Image
Ground Truth
HOTEL SUIZO
Predicted Image

Generative Adversarial Network (GANs)

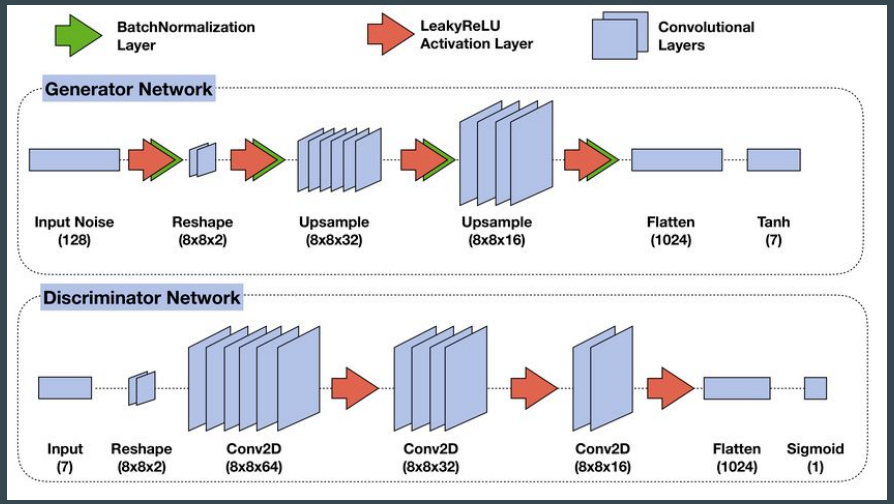

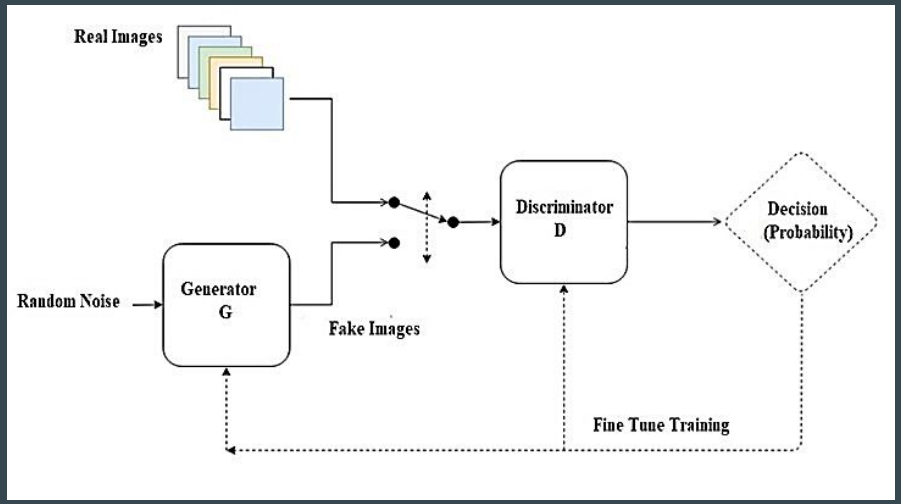

Diffusion models (guided)

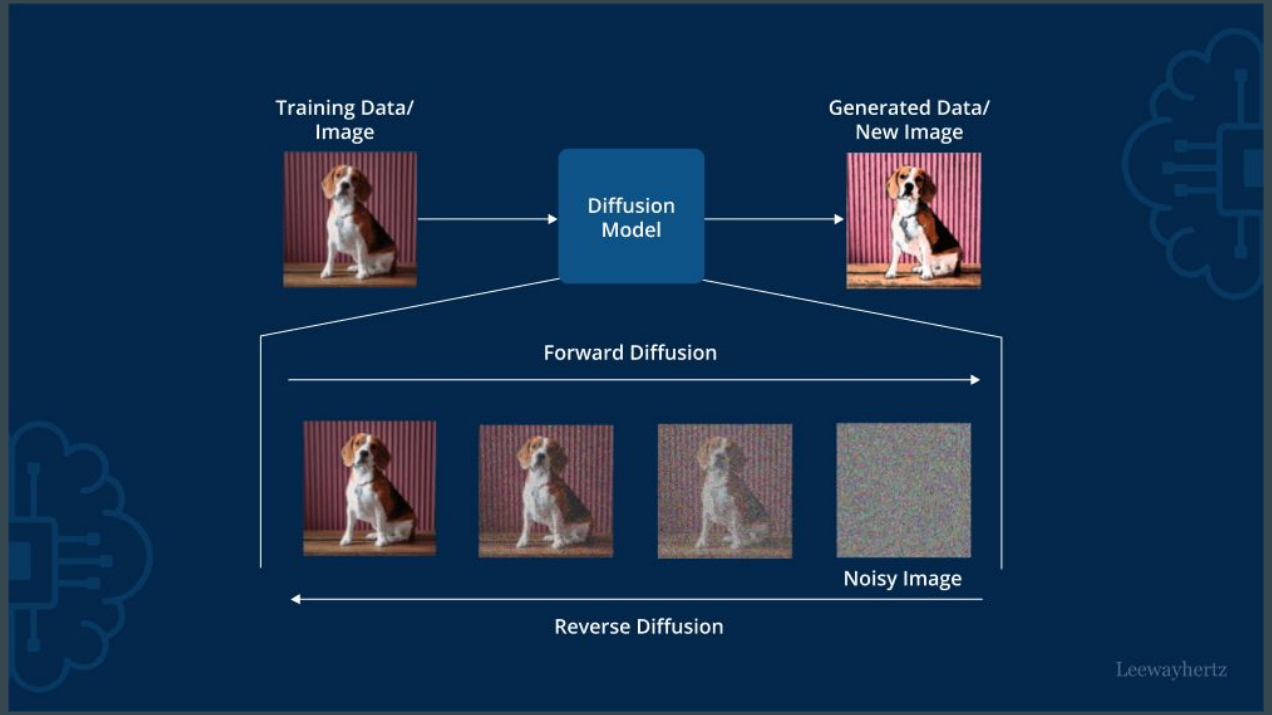

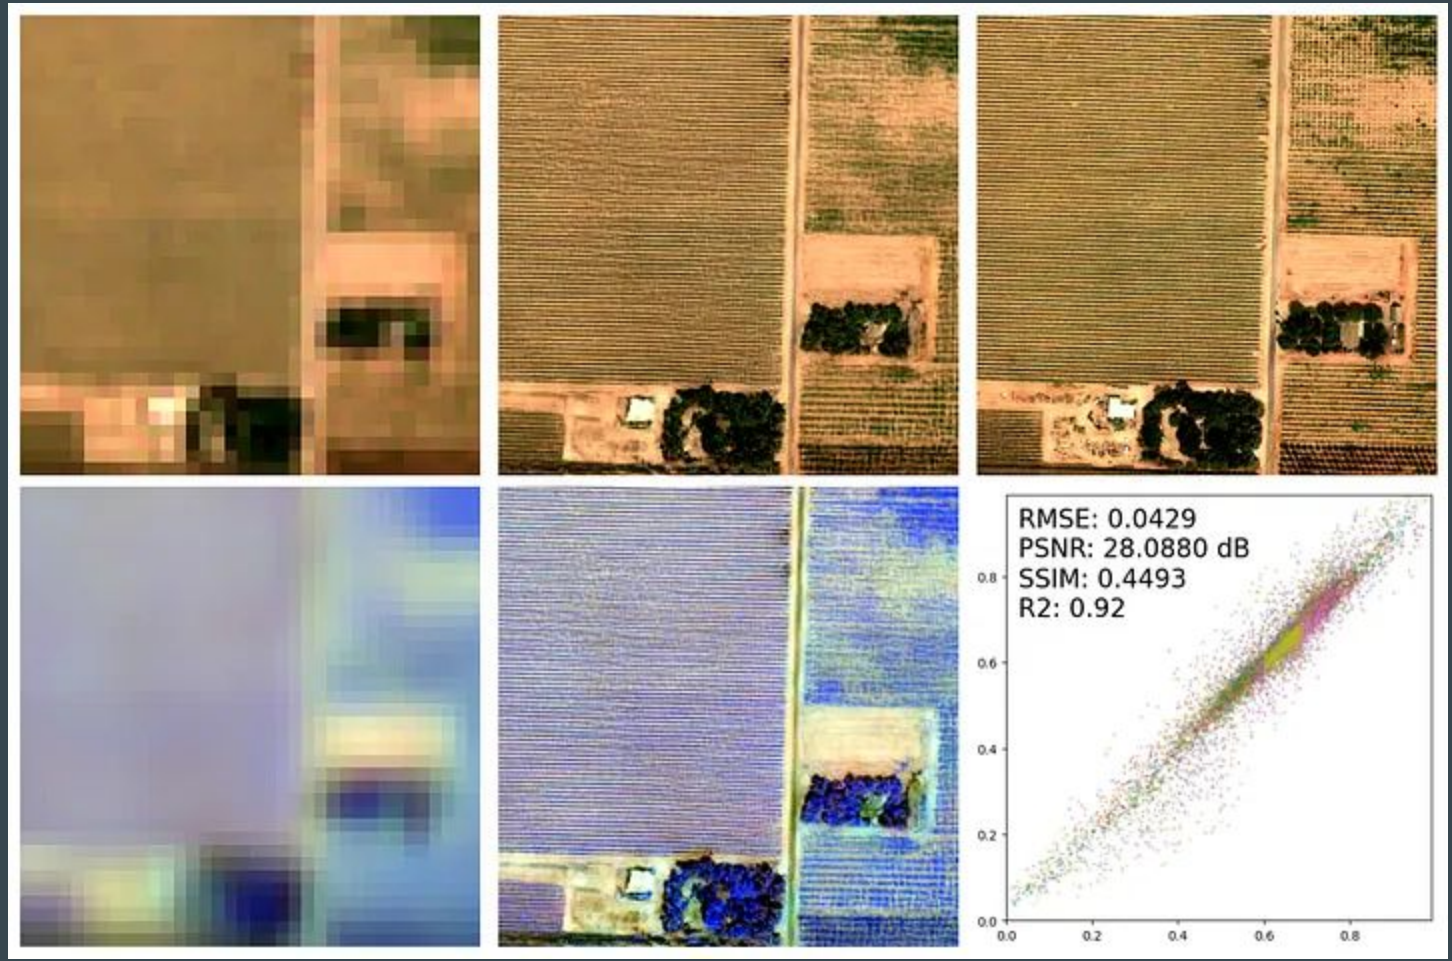
RMSE: 0.0429
PSNR: 28.0880 dB
SSIM: 0.4493
R2: 0.92
0.0
0.2
0.4
0.6
0.8
0.0
0.2
0.4
0.6
0.8

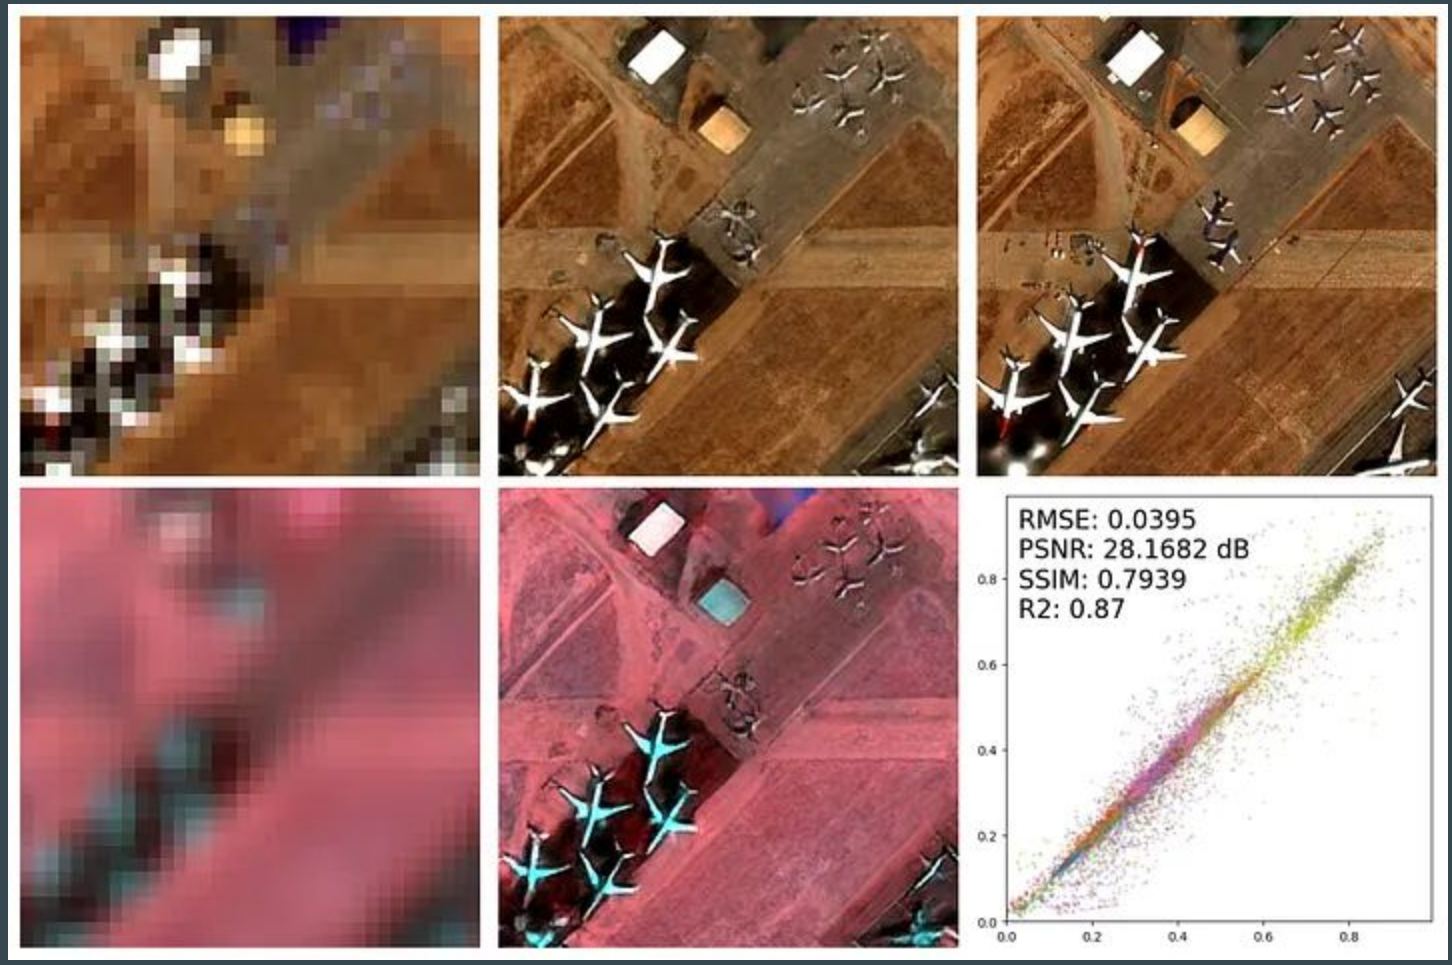
RMSE: 0.0395
PSNR: 28.1682 dB
SSIM: 0.7939
R2: 0.87
0.0
0.2
0.4
0.6
0.8

Current 2.5x2.5 degrees images

Lockman Hole

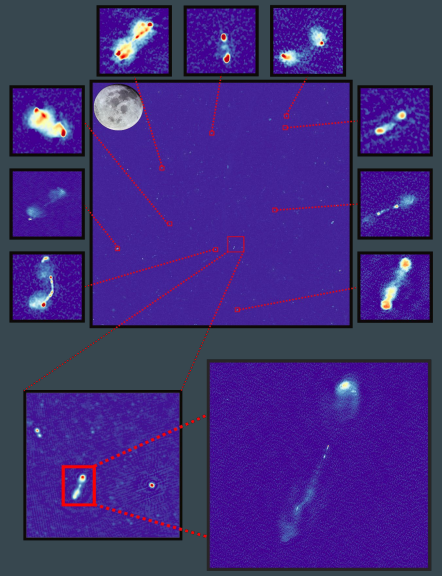

Sweijen et al. 2022

ELAIS-N1

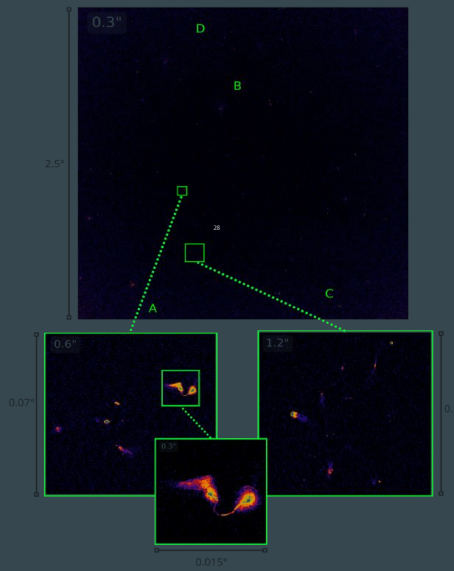

de Jong et al. 2024

Boötes

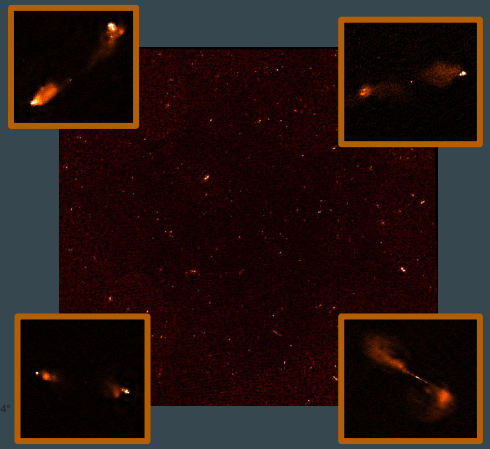

Escott et al. (subm.)

Euclid DFN

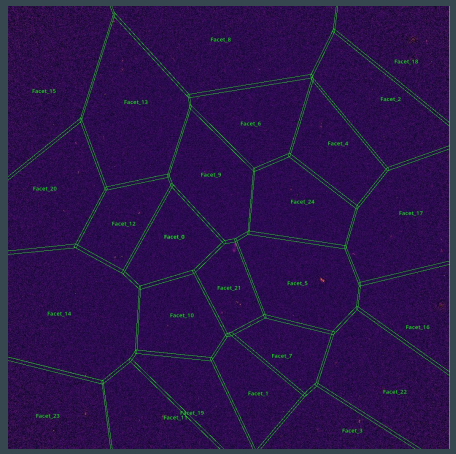

Bondi et al. (in prep.)

GenAI for calibration

1. LoTSS at 6” → 95% of the sky

   LOFAR-VLBI at 0.3” → <<1% of the sky
2. Predicting the sky at high-resolution
3. Use predicted sky as model data for calibration
4. Computational and/or quality improvements?

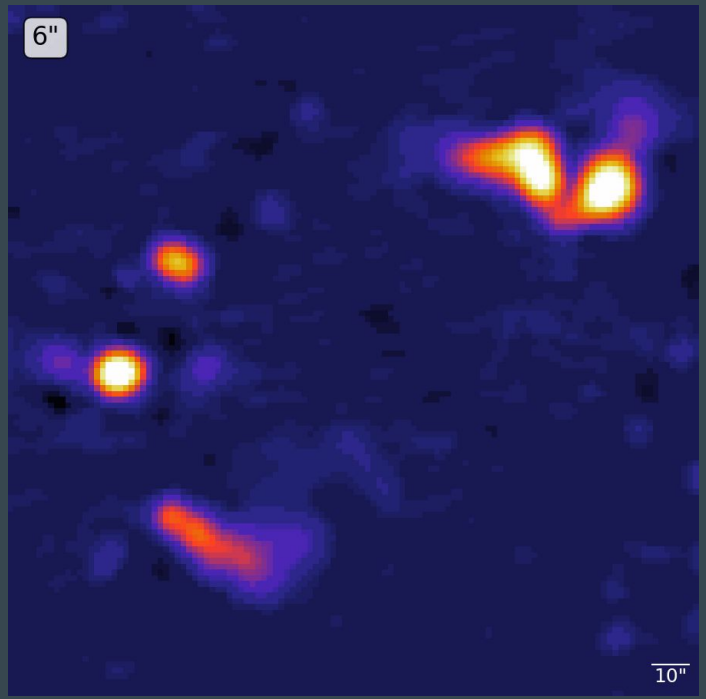

Ethan Woest

Predicting model data

1. Take calibrated and uncalibrated model data images
2. Learn to map from uncalibrated to calibrated
3. Computational improvements?

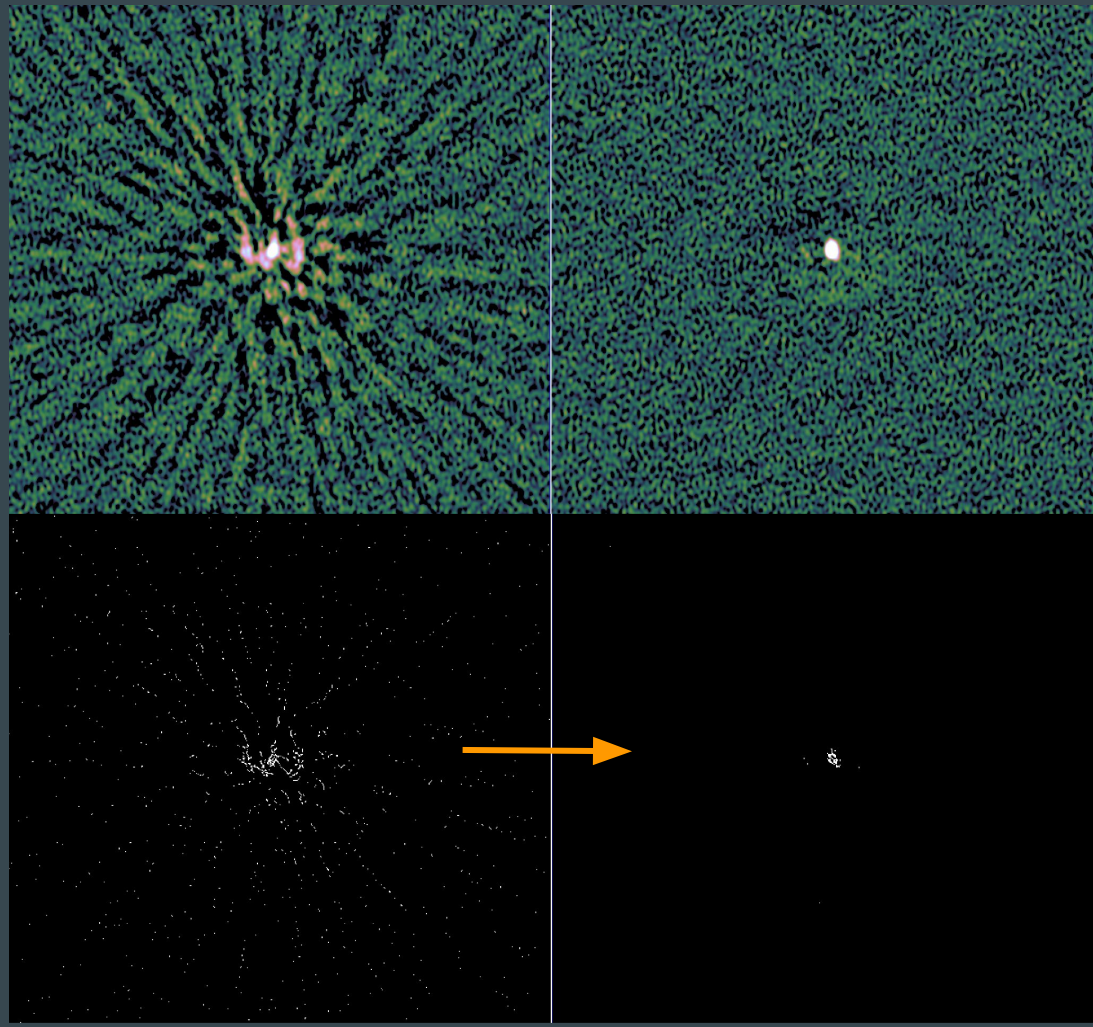

Conclusions

- High-resolution LOFAR imaging pipeline exists
- Replacing human validation for calibration with AI (Transformer DINOv2 model)
- Future work: GenAI for improving calibration

de Jong et al. 2025

Extra slides

Reinforcement learning for calibration parameters

- **Agent** – The learner or decision-maker
- **Action** – A choice the agent makes
- **Environment** – The ‘world’ the agent interacts with
- **State** – A snapshot of the environment
- **Reward** – Feedback signal for each action

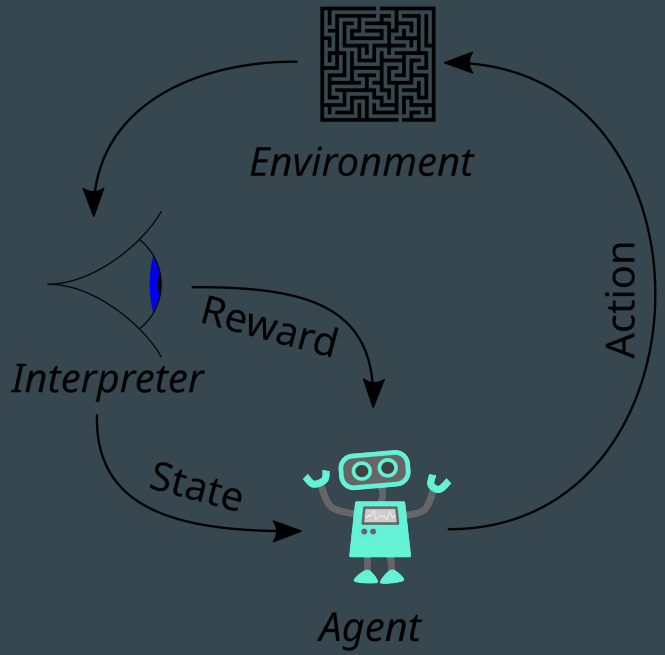

Reinforcement learning for calibration parameters

- **Agent** – Calibration algorithm
- **Action** – Adjust calibration parameter(s)
- **Environment** – Visibility or image space
- **State** – Current calibration solutions or image
- **Reward** – Calibration quality metric

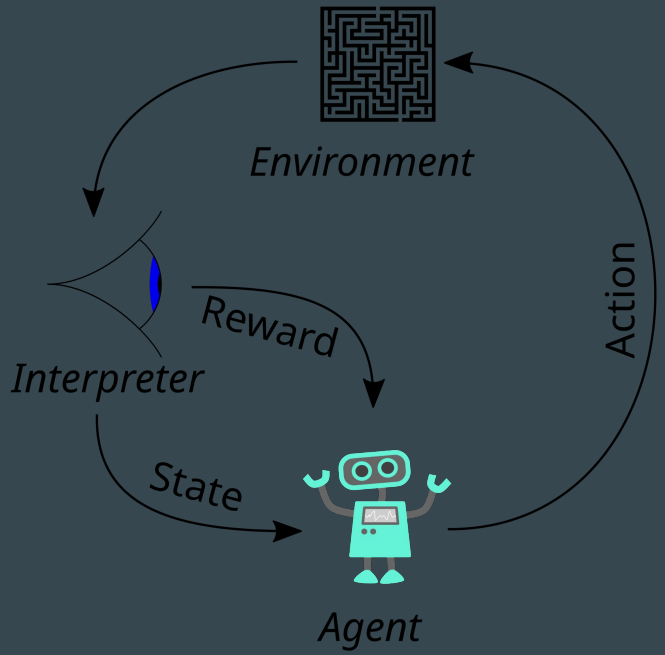

Yatawatta et al. (2019; 2021)

- Same image quality, but different calibration quality

- Reward: influence maps

- Grid-search (GS) vs. Reinforcement learning (RL):
  RL ~6 times more efficient in learning parameters than GS