Self-supervised learning for glitch classification

The dataset

- Contains 21 sets of images for glitches and 1 set of images for background noise.
- There are four versions for each image, the difference being the time window.
- It is divided between three subsets:
  - Train (6008 images)
  - Validation (1288 images)
  - Test (1287 images)

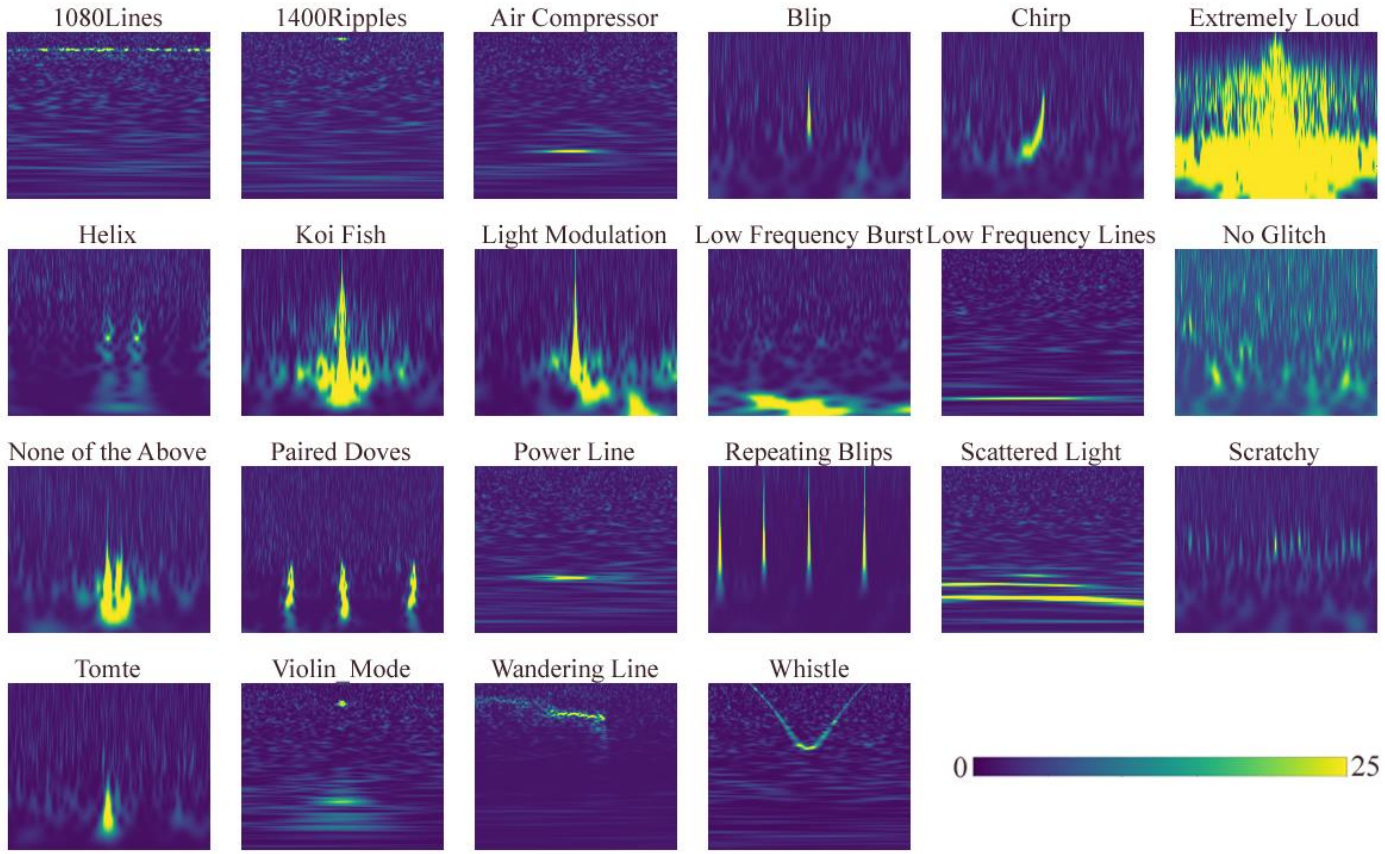

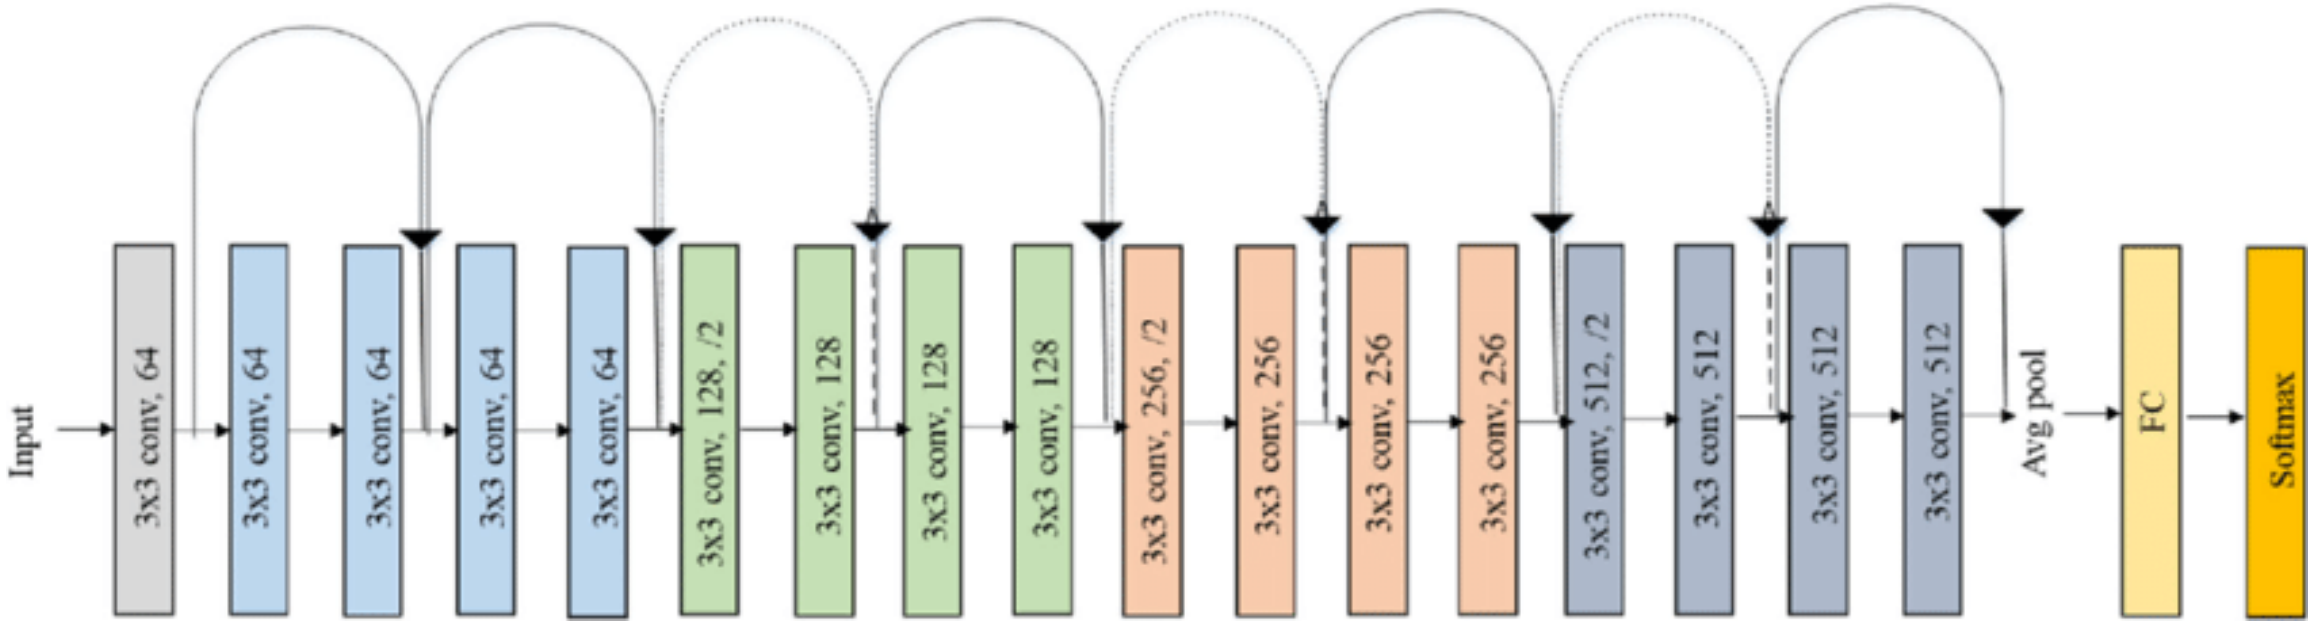

Resnet 18’s architecture

Image taken from: Ramzan, Farheen & Khan, Muhammad Usman & Rehmat, Asim & Iqbal, Sajid & Saba, Tanzila & Rehman, Amjad & Mehmood, Zahid. (2019). A Deep Learning Approach for Automated Diagnosis and Multi-Class Classification of Alzheimer’s Disease Stages Using Resting-State fMRI and Residual Neural Networks. Journal of Medical Systems. 44. 10.1007/s10916-019-1475-2.

The supervised approach

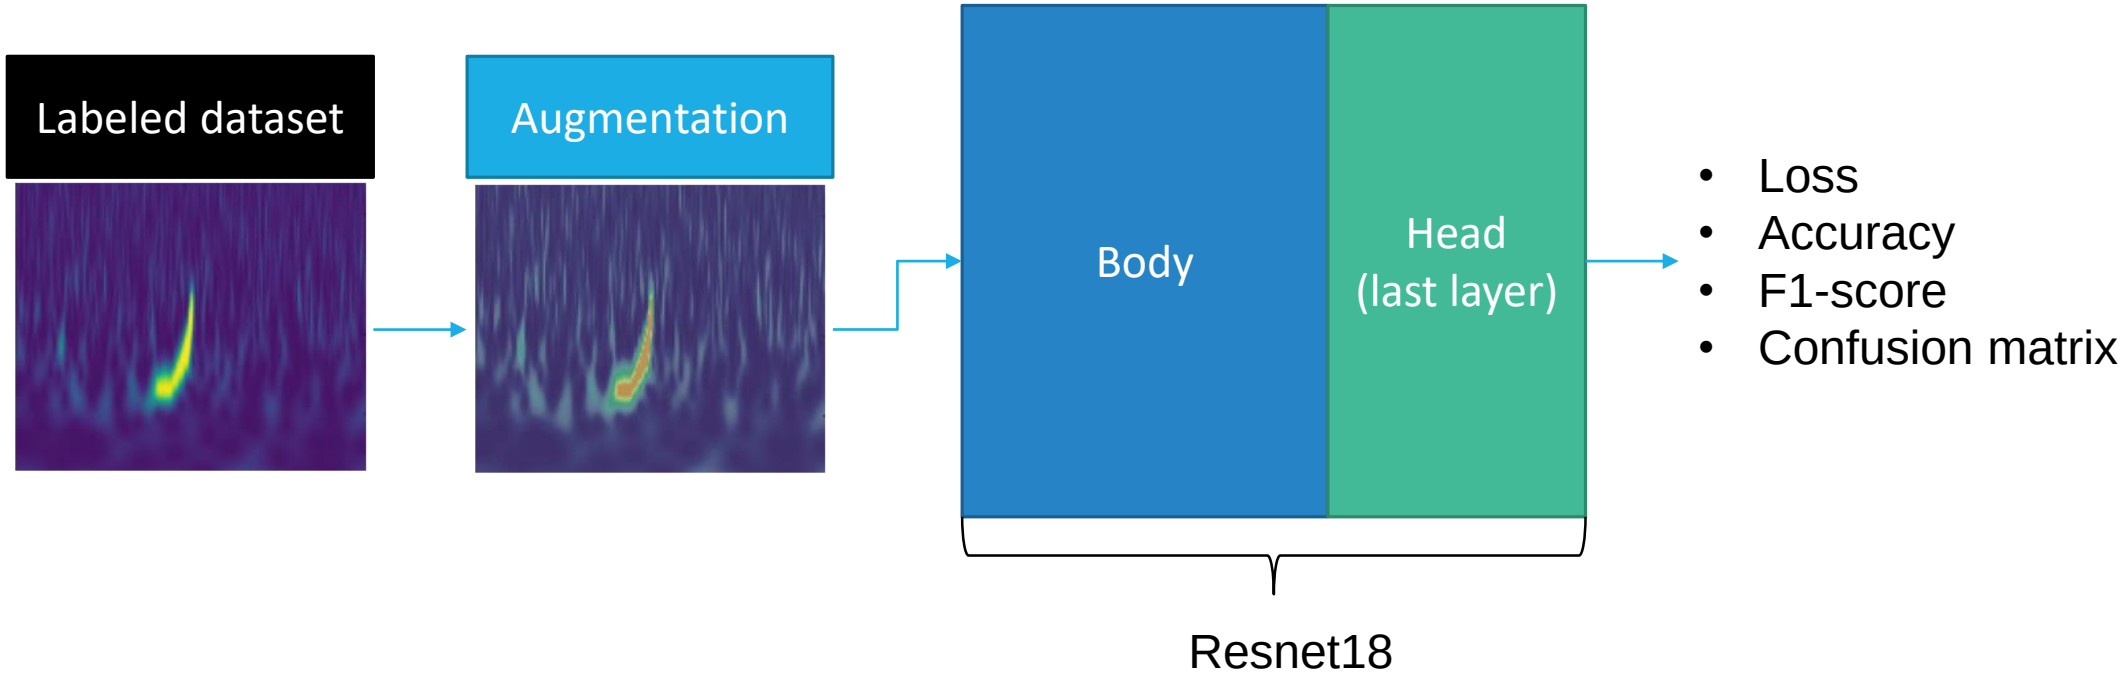

Self-supervised + transfer learning

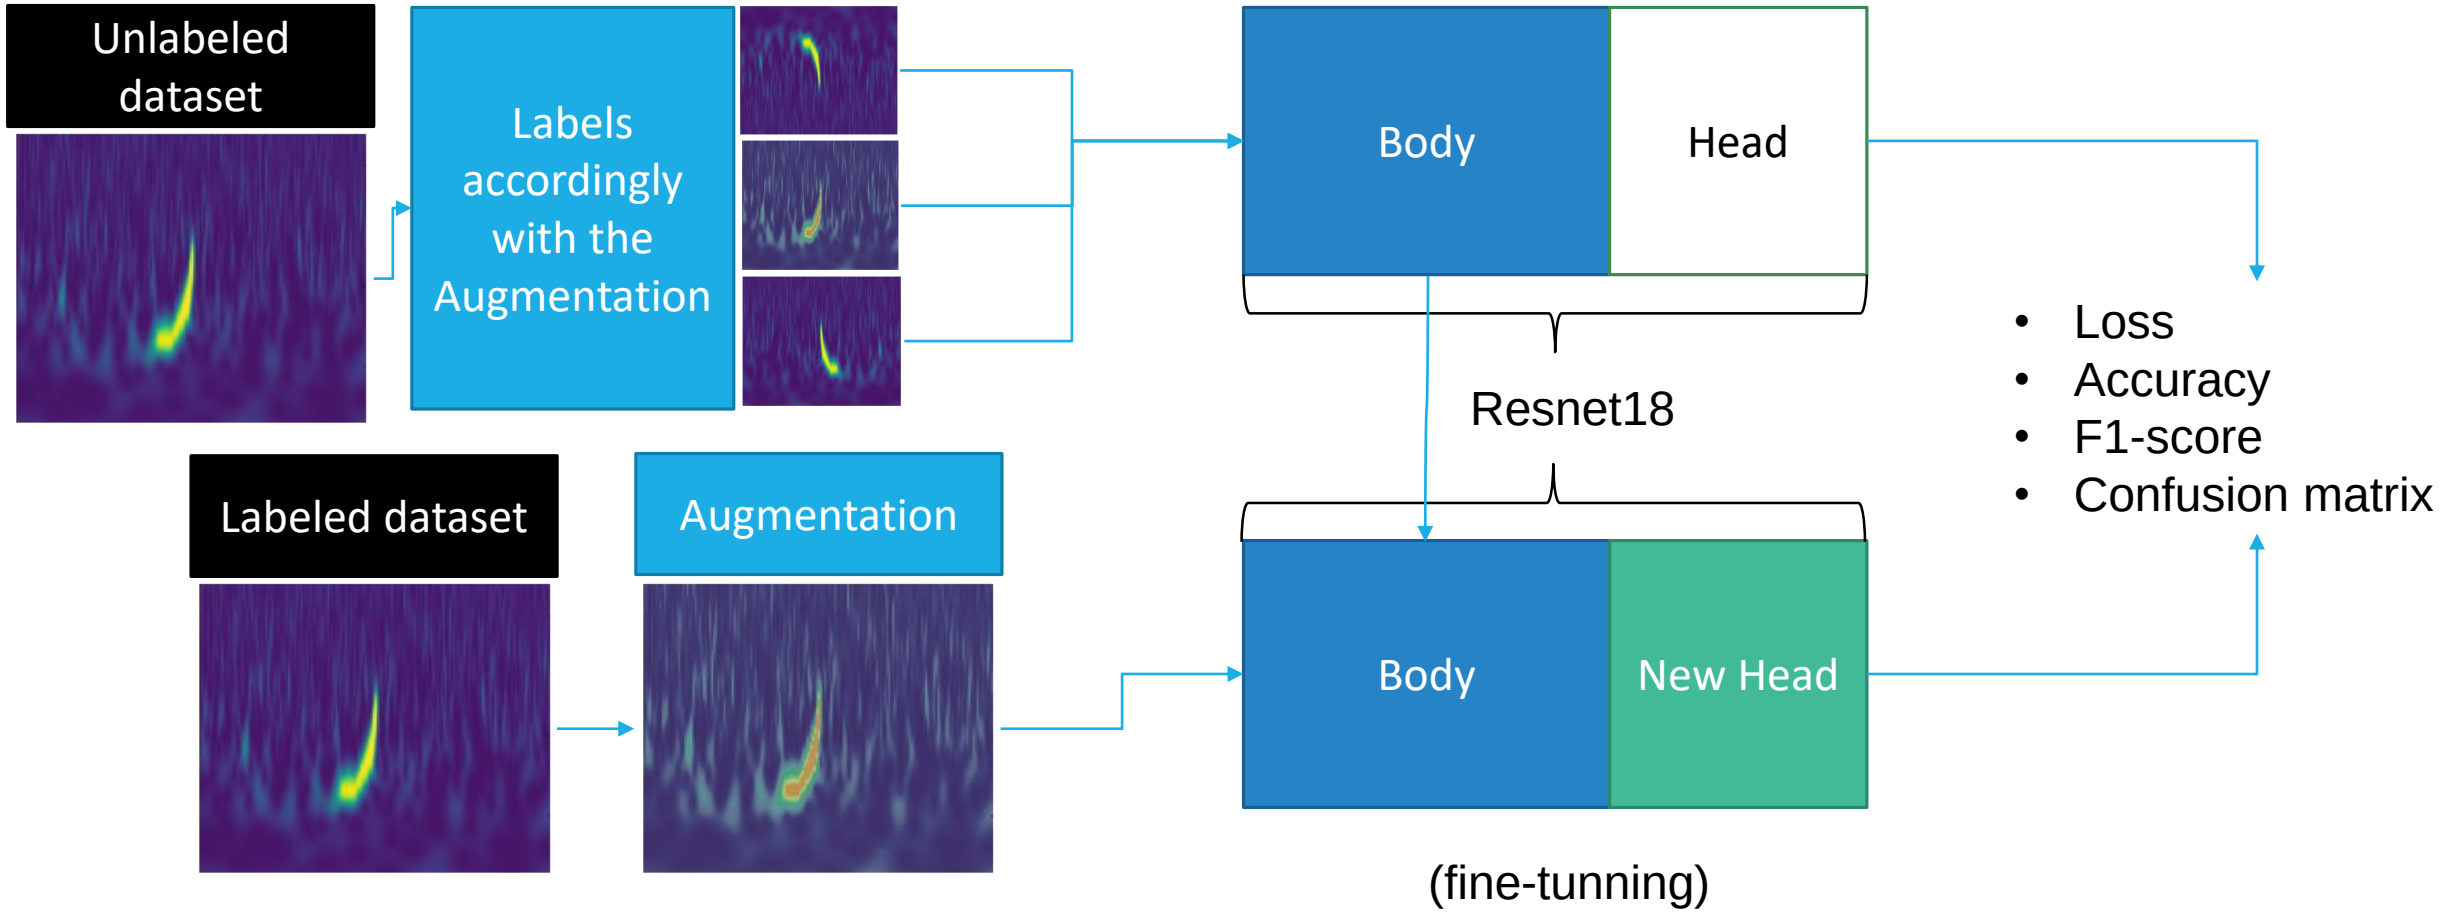

Supervised

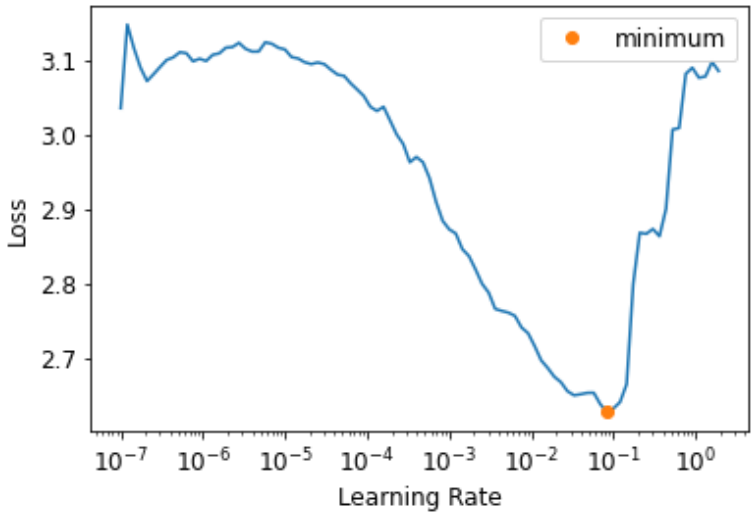

Self-supervised

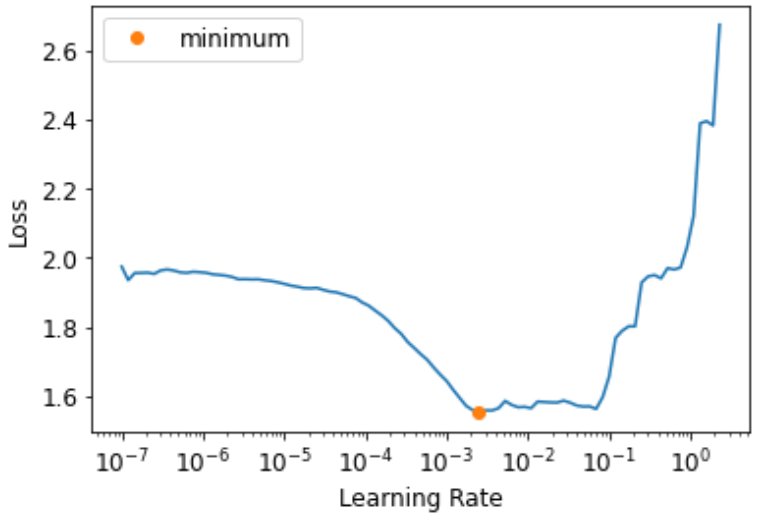

Self-supervised + transfer learning

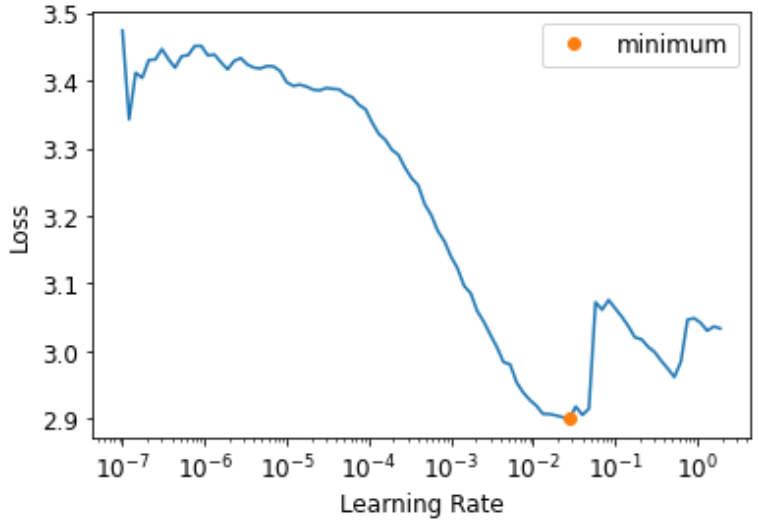

Finding the learning-rate

Supervised

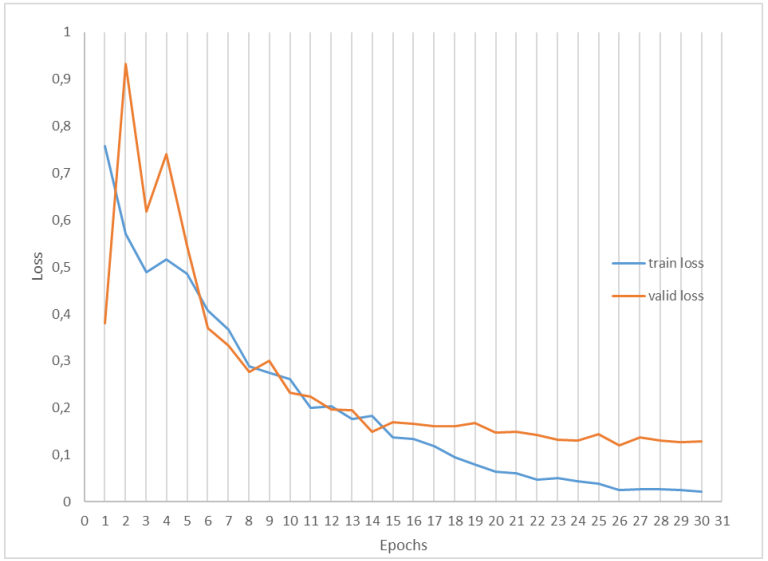

Self-supervised

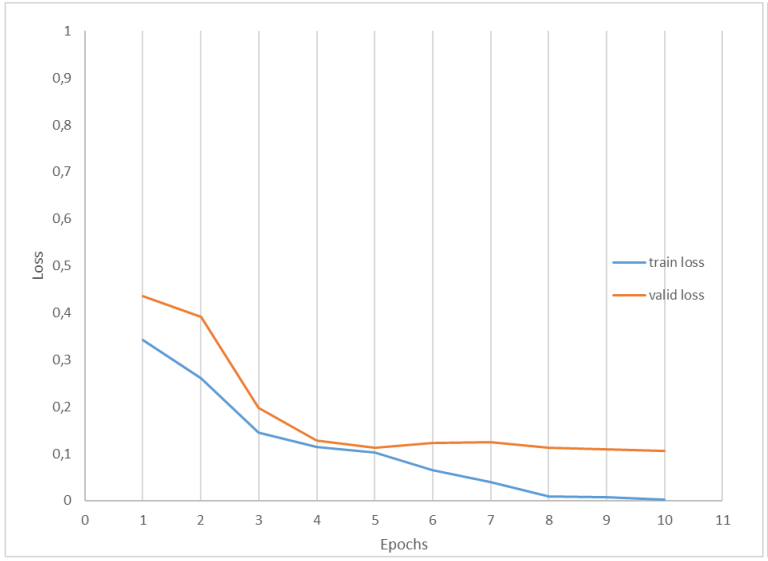

Self-supervised + transfer learning

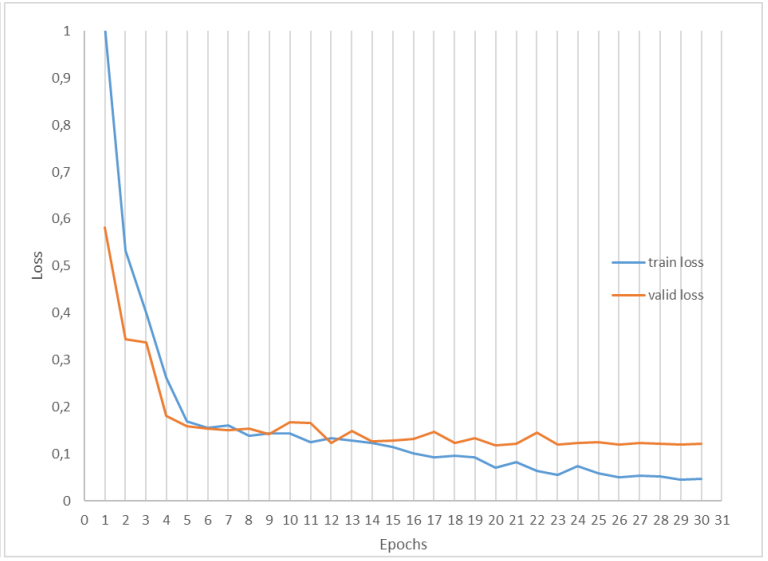

Train and valid loss across the epochs

Supervised

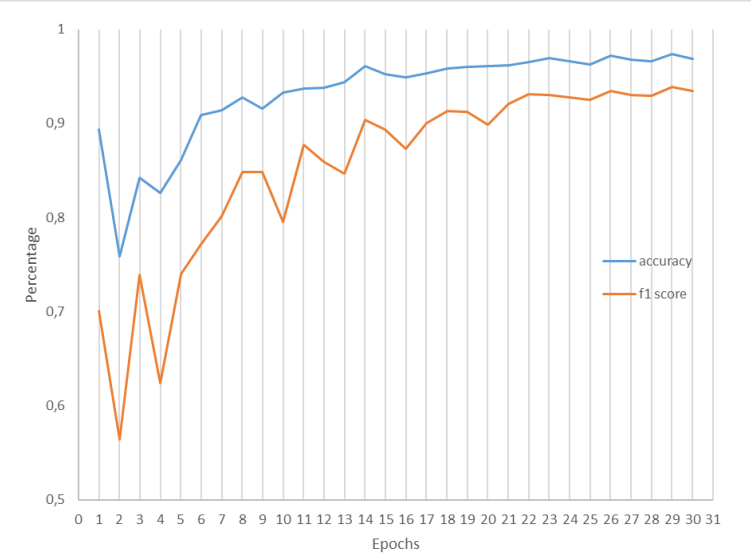

Self-supervised

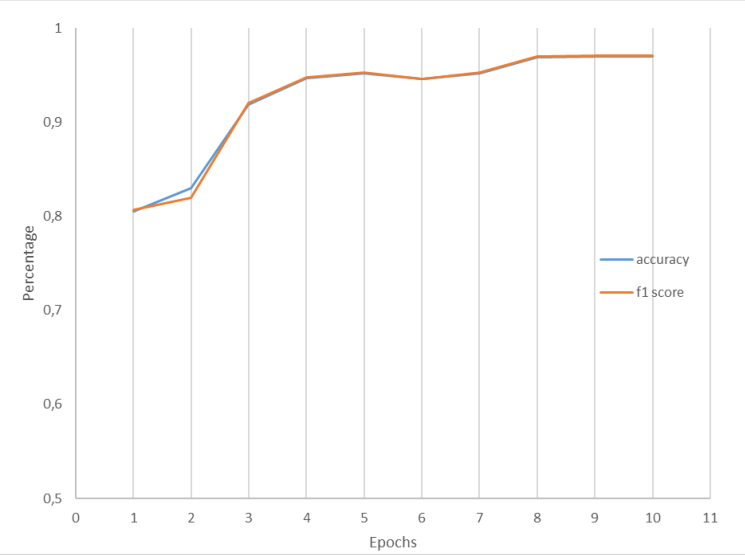

Self-supervised + transfer learning

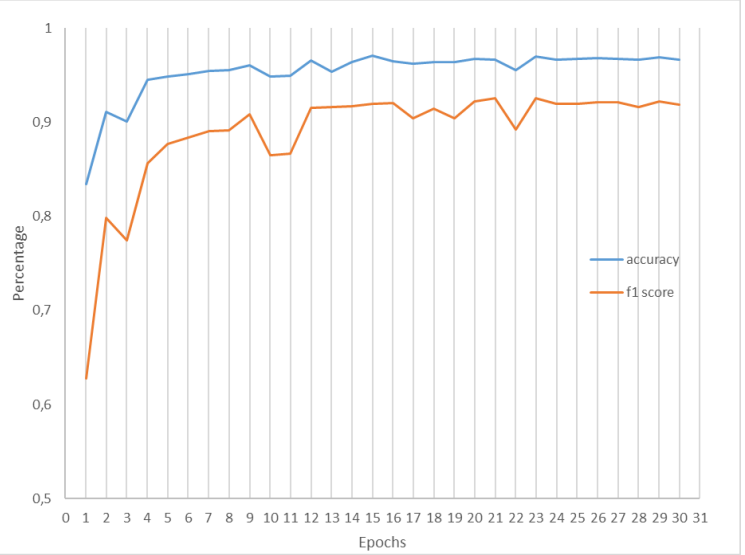

Accuracy and macro-averaged F1-score

Confusion matrices: supervised

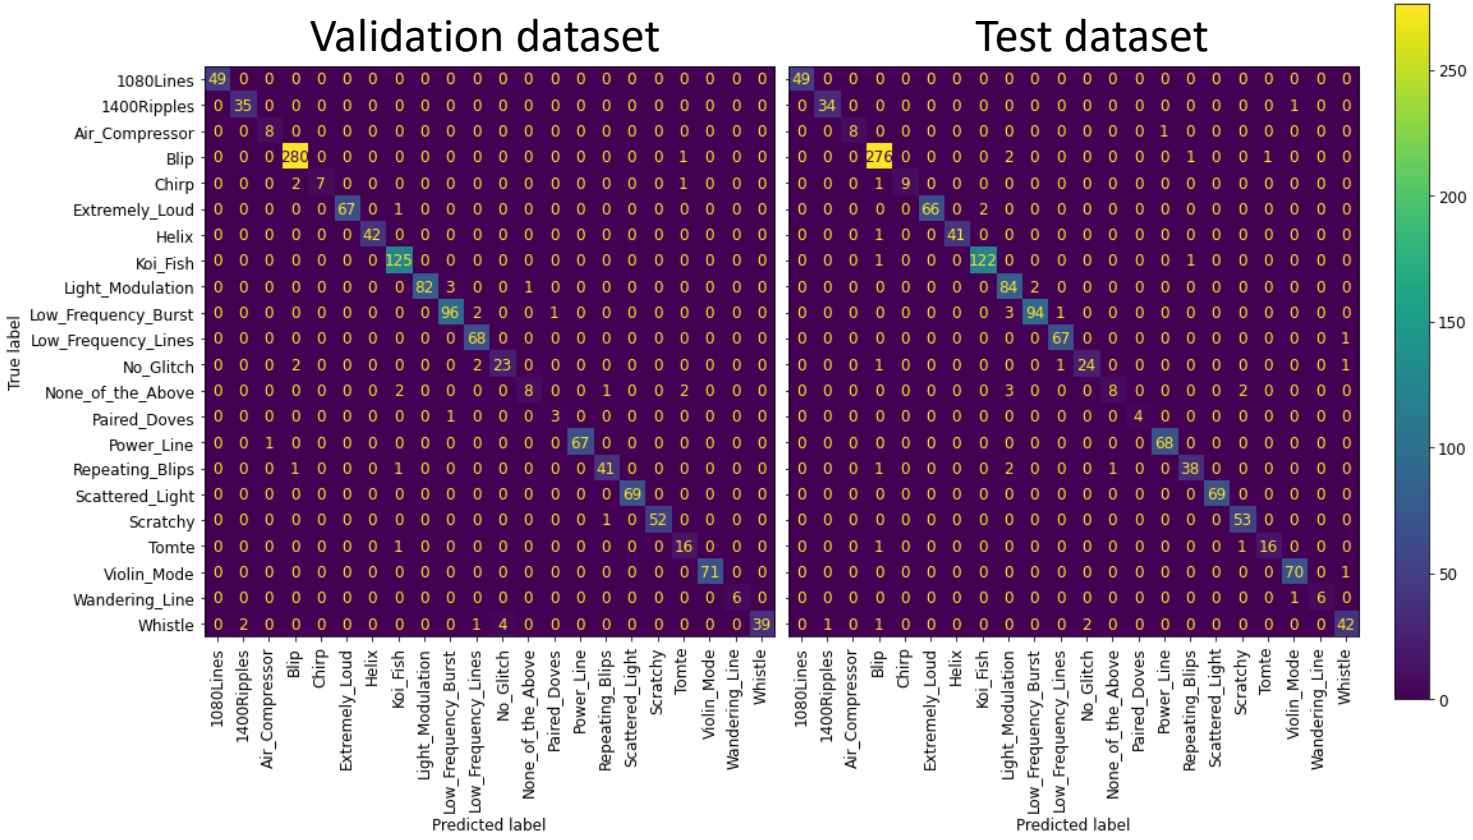

Validation:

- Accuracy: 97.36%
- Macro-averaged F1-score: 93.88%

Test:

- Accuracy: 96.97%
- Macro-averaged F1-score: 95.19%

Confusion matrices: self-supervised

Test:

- Accuracy: 96.97%
- Macro-averaged F1-score: 96.99%

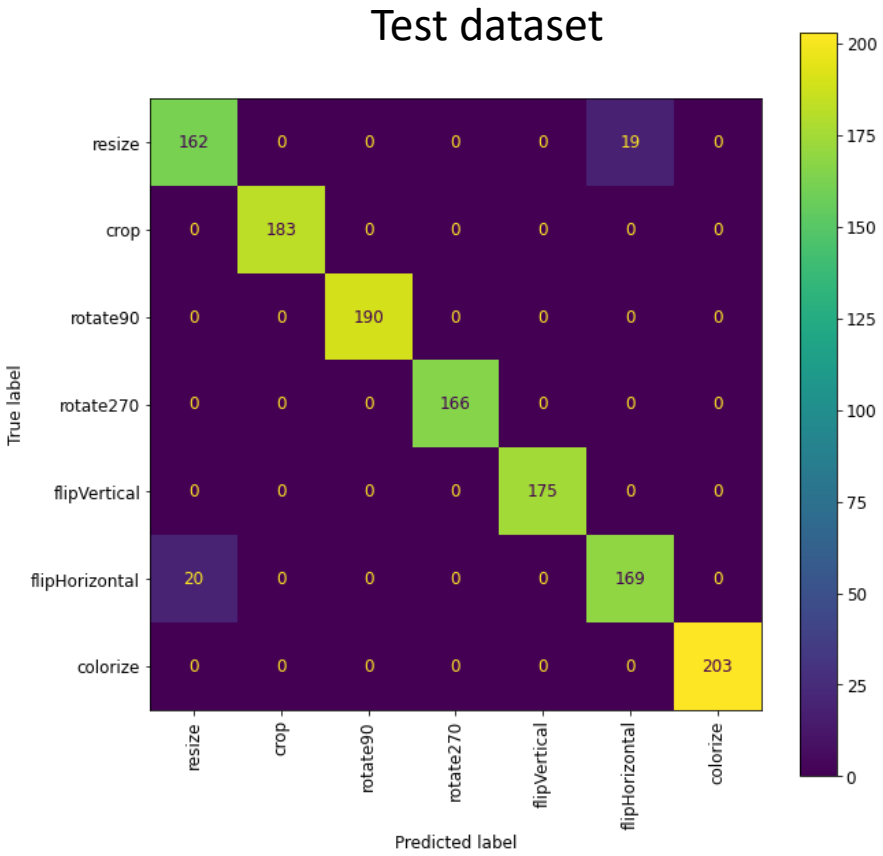

Confusion matrices: SSL+ transfer learning

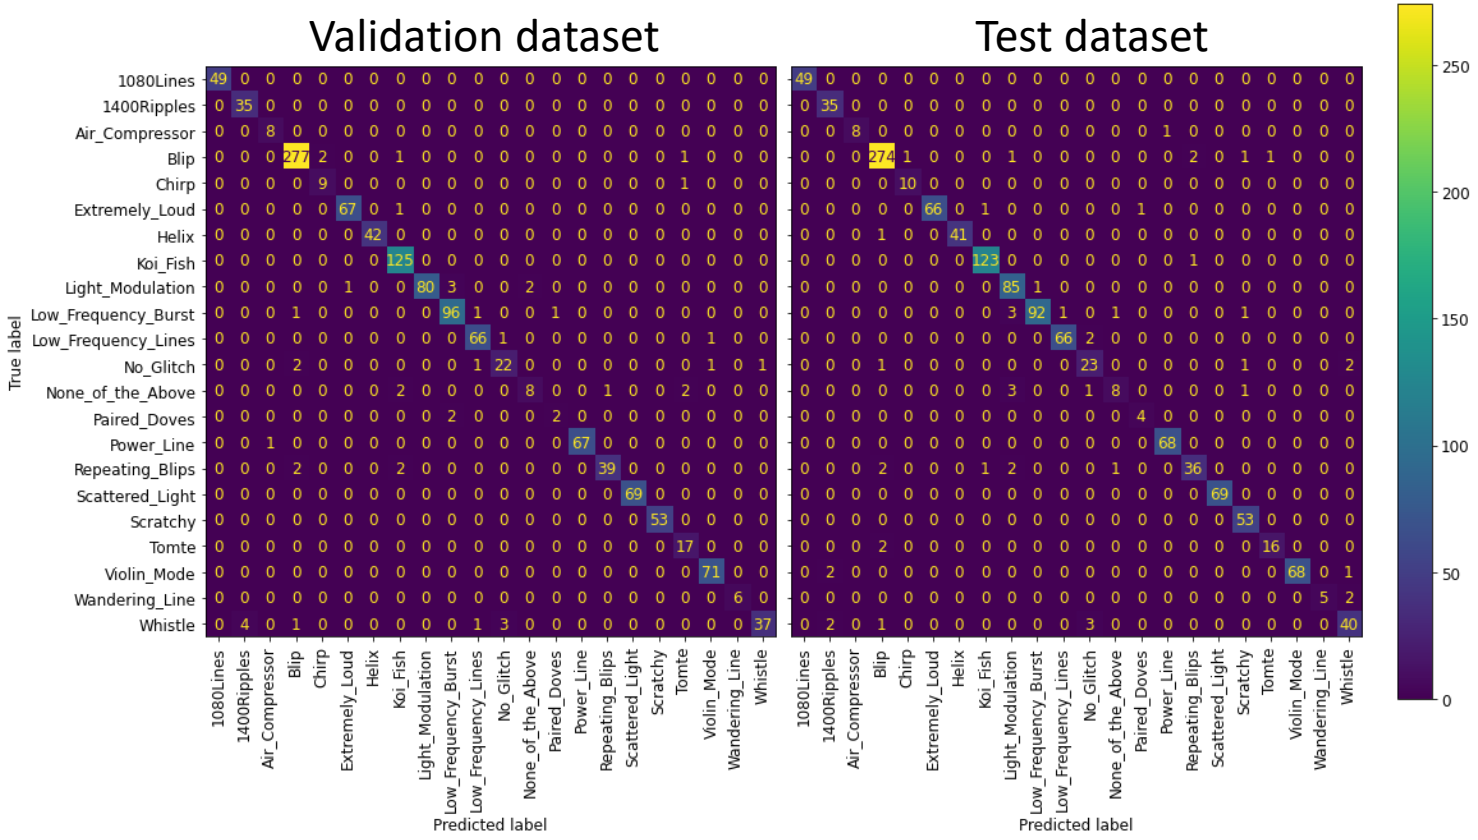

Validation:

- Accuracy: 96.66%
- Macro-averaged F1-score: 92.58%

Test:

- Accuracy: 96.27%
- Macro-averaged F1-score: 93.27%

Signal

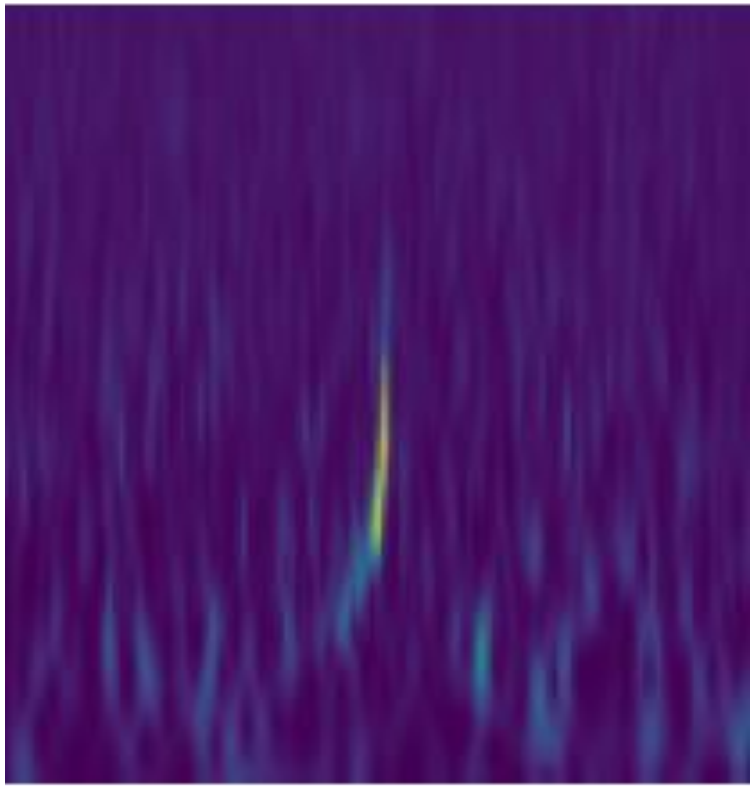

Chirp

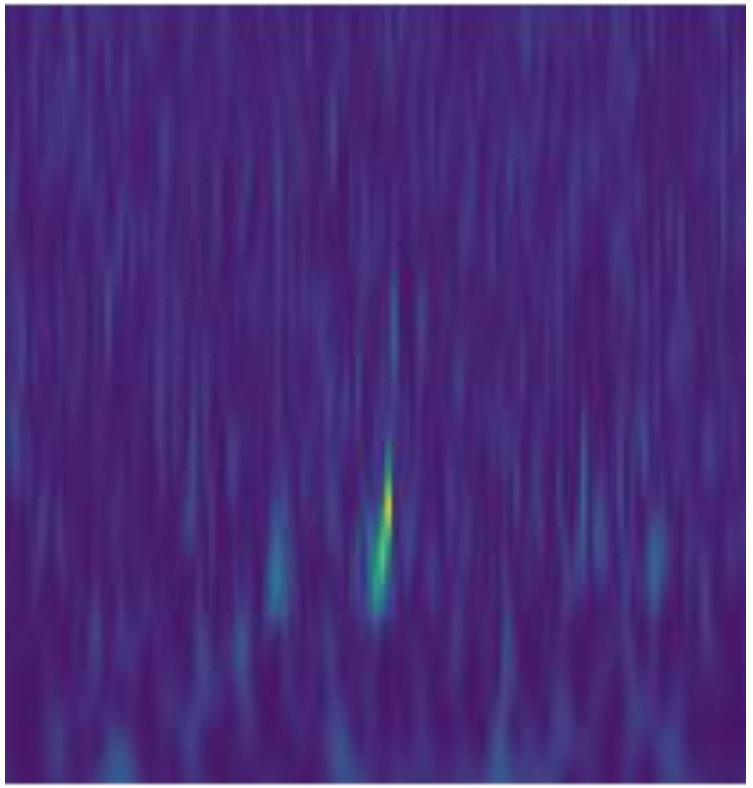

Blip

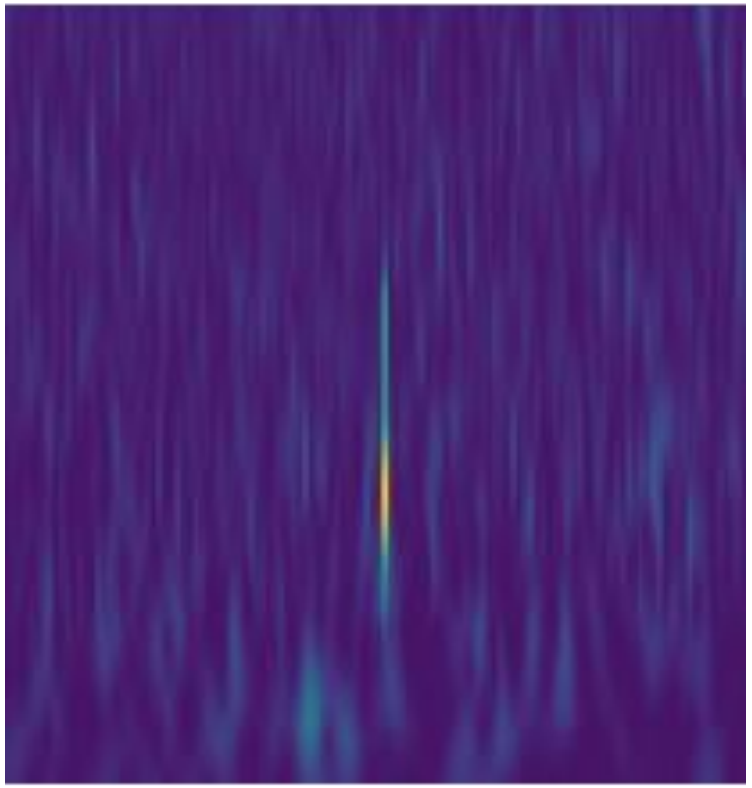

Can it detect as signals?

SINCE CHIRPS ARE SUPPOSED TO EMULATE A SIGNAL, CAN IT BE DETECTED BY THE NEURAL NETWORK?

Can it detect as signals?

There are 11 images of events and 11 images of just background noise, and it didn't go so well…

It classified the events:

- 54.5% as Blips (6/11)
- 18.1% as Chirps (2/11)

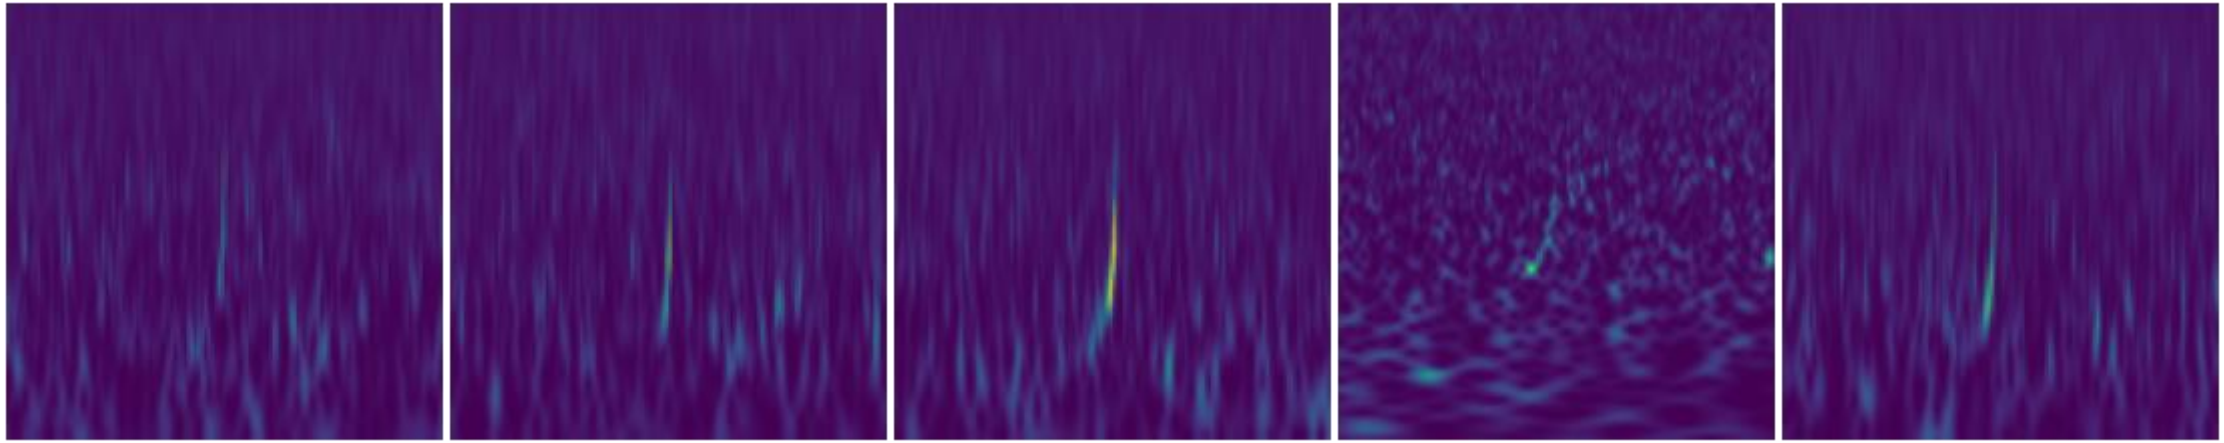

Answer:	Blip	Blip	Chirp	Low_Frequency_Lines	Chirp
Probability:	0.794	0.440	0.982	0.443	0.999