A neural network (NN) approach to cloud detection in NWP

Chris Burrows and Tony McNally

Many thanks to Marcin Chrust, Alan Geer and Mat Chantry

chris.burrows@ecmwf.int

ECMWF

Outline

- Motivation(s)
- Methods
- Initial results
- Outlook

Motivation 1 – the ML landscape

- Machine learning (ML) is a very hot topic.
- ECMWF is investigating ML for various applications (we have a "ML Roadmap"!).
- Tools for training neural networks are now widely available and are easy to use.

Motivation 2 – ML algorithms can be super-fast

- Operational NWP forecasts must be delivered on time – computational cost savings are always welcome (including IR cloud detection).

Motivation 3 – learning what cloud information is in the spectra

- Currently we use observations and clear-sky model equivalents to determine which channels are cloud-affected.
- However, is there sufficient information in the observations **alone** to determine this? If so, a neural network (NN) could do this classification given a large enough training set. This may help to avoid poor cloud detection in the presence of localised background errors.

Important notes

- This study uses **IASI** brightness temperatures.
- The approach is applied to **individual spectra** in order to dynamically identify which channels are cloud-affected and hence should not be assimilated.
- It is **not** an image-recognition/convolutional approach to identify spatial cloud patterns.

- We want to know if the observed spectra **alone** contain enough information to be able to identify which channels are affected by cloud – open science question.

Approach – train a NN to replicate the McNally and Watts cloud detection flags (and AVHRR clustering checks), but using **only** observed IASI values as inputs

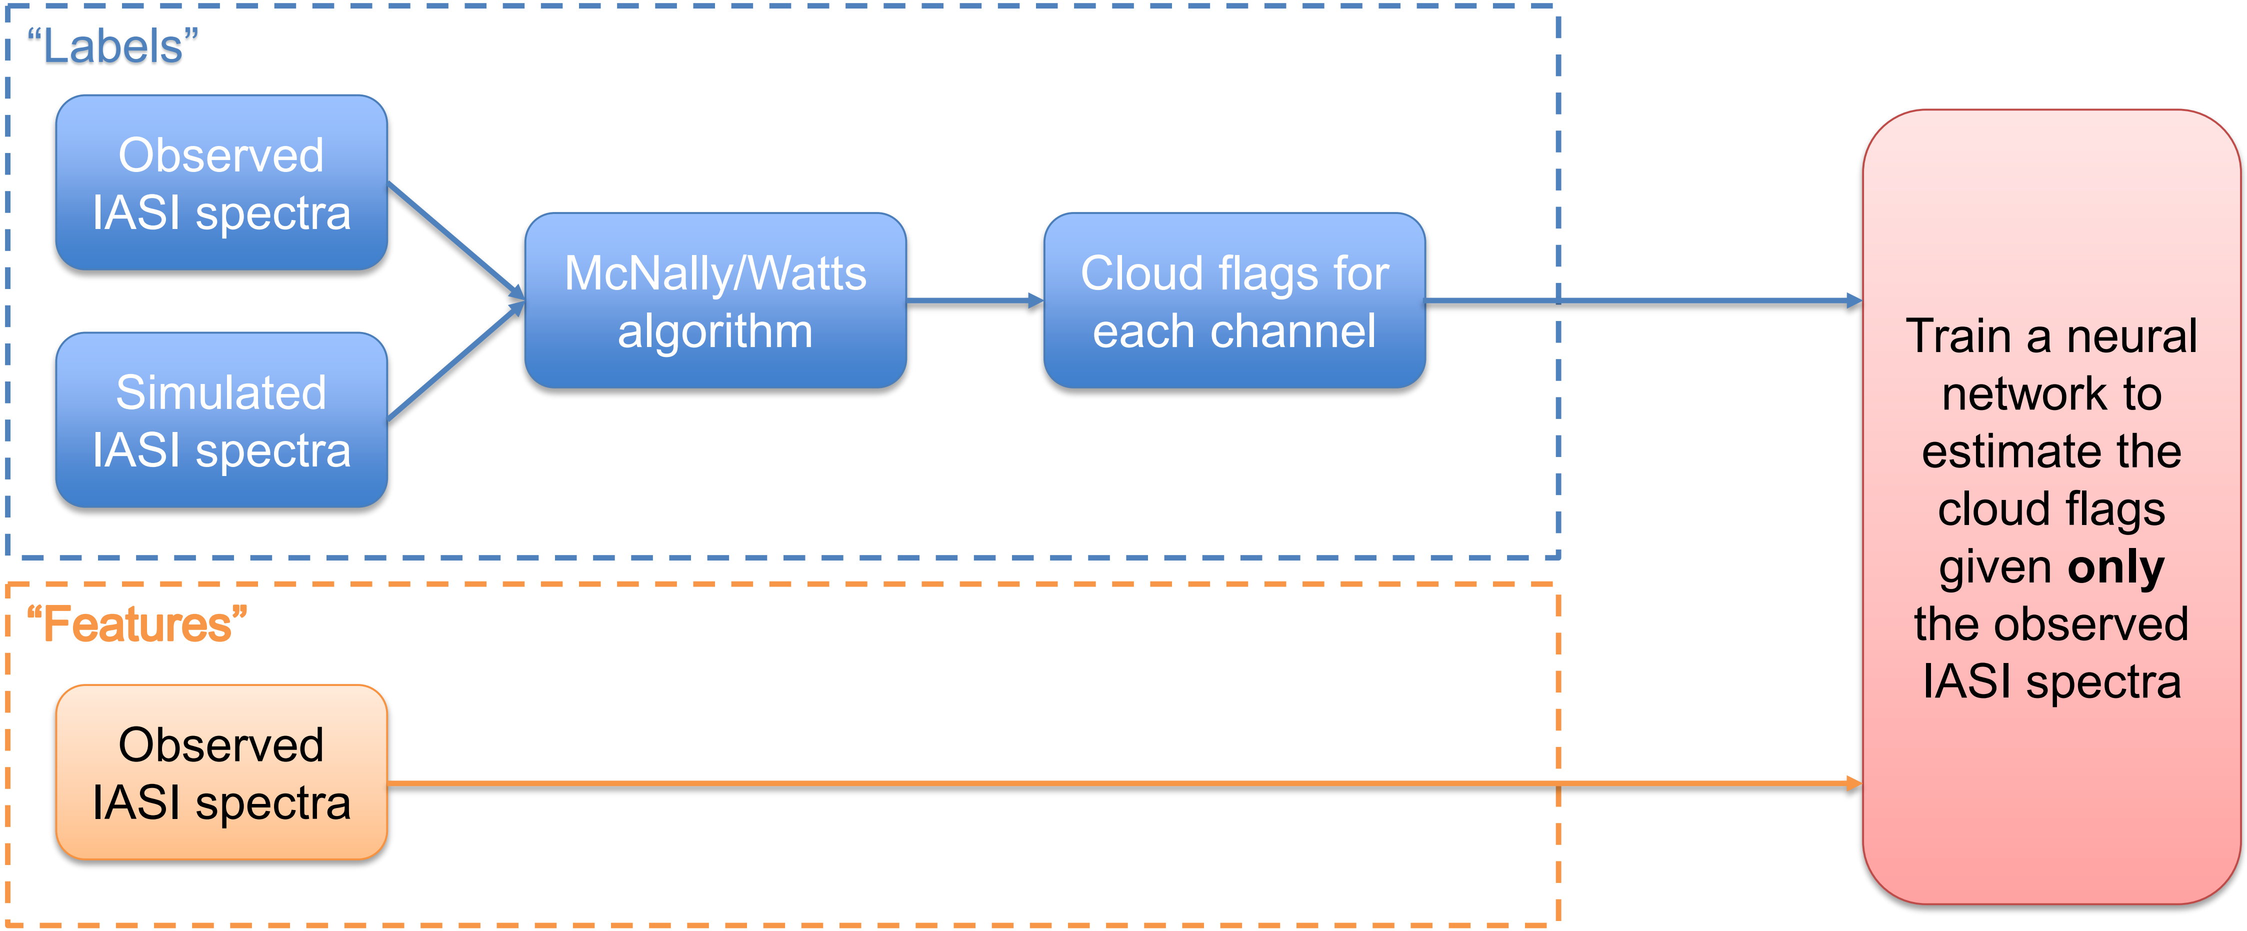

McNally A, Watts P. 2003. A cloud detection algorithm for high-spectral-resolution infrared sounders. Q. J. R. Meteorol. Soc. 129: 3411– 3423, doi: 10.1256/qj.02.208

For the ML aficionados...

- Supervised classification problem.
- Input brightness temperatures normalised: (BT-260)/60
- Input data: 1 week Metop-C IASI from 420 channels.
- 25% of data reserved for hold-out validation.
- 420 input neurons (one per channel)
- Two dense hidden layers (dimension 420); activation function: Rectified linear unit (ReLU).
- Output cloud flag layer: 420 neurons; activation function: sigmoid ($0<x<1$).
- Loss function: binary cross-entropy.
- Epochs: 250 (subjectively chosen to avoid over-fitting).
- Minimiser: “adam”.
- Thresholding of sigmoid output chosen to preserve ratio of clear/cloudy flags.

Histograms of observations minus simulations for a few channels

- The ‘all’ (blue) line shows both clear and cloudy obs (note the **cold tails**).
- The ‘ops’ (orange) line shows the cloud-free sample using the operational cloud detection scheme.
- The ‘ml’ (green) line shows the cloud-free sample from the neural network.
- **We want machine-learning statistics, ‘ml’ (green) to match the operational statistics, ‘ops’ (orange).**
- Generally, the agreement is very good and importantly, the cold tails are mostly removed.
- The window channel (861.5 cm^{-1}) shows the worst agreement, with significant broadening.

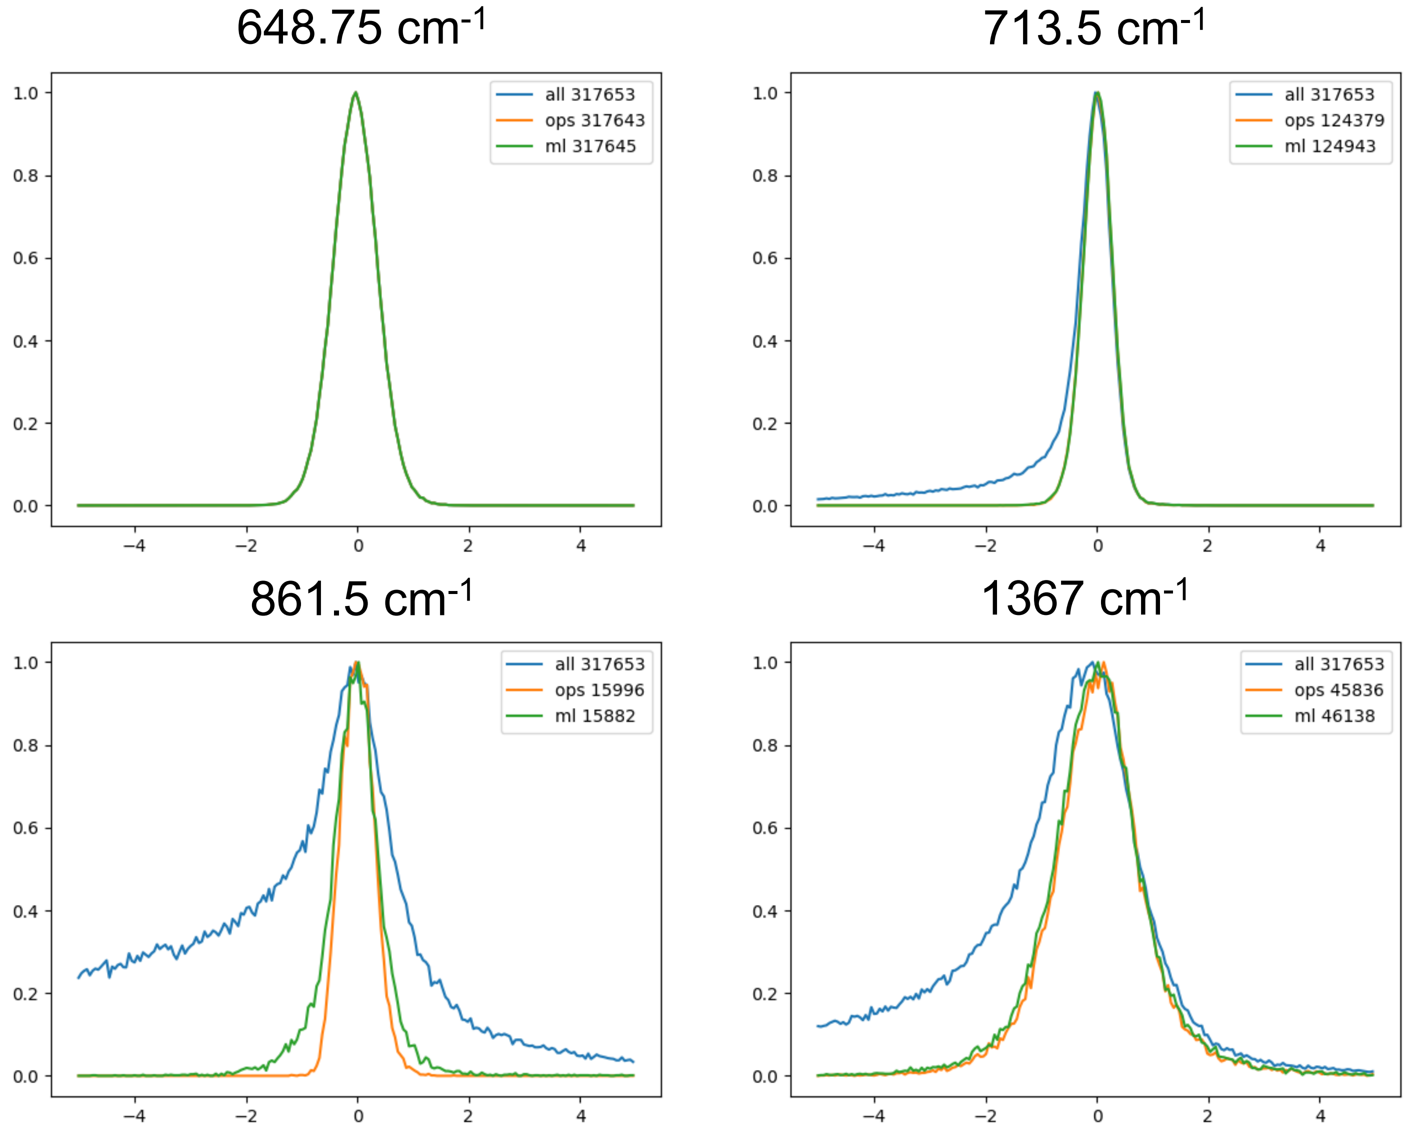

Initial fit looks good, even with a limited training set
Yellow is cloudy, blue is clear

Higher-peaking | Lower-peaking

Operations

Neural network

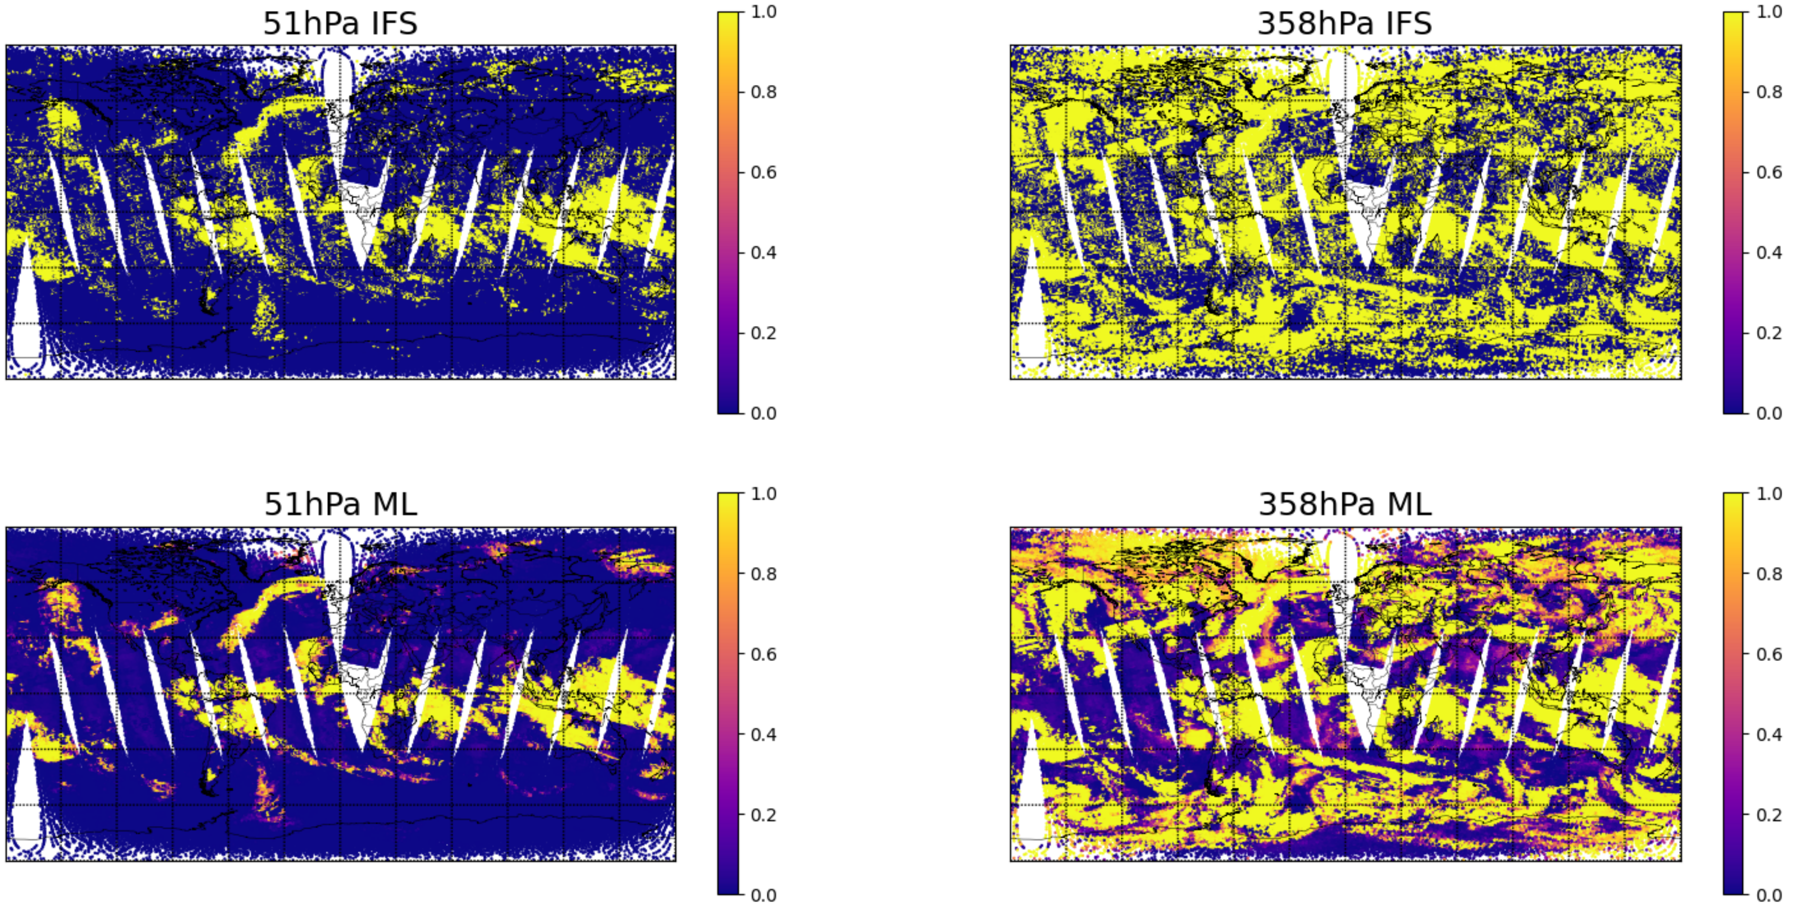

Preliminary assimilation experiments

- ECMWF has developed the "FNN" package (https://arxiv.org/pdf/2210.13817.pdf, https://github.com/cerea-daml/fnn/) which takes sequential NNs from Keras/TensorFlow and allows them to be run efficiently within Fortran code (as an aside, TLs and adjoints are calculated implicitly).
- This has allowed the existing cloud detection to be replaced by the NN version for testing.

- The number of used observations is reduced for all channels, particularly the lower-peaking channels
- The thresholding of the NN's output layer can have a big effect on this.

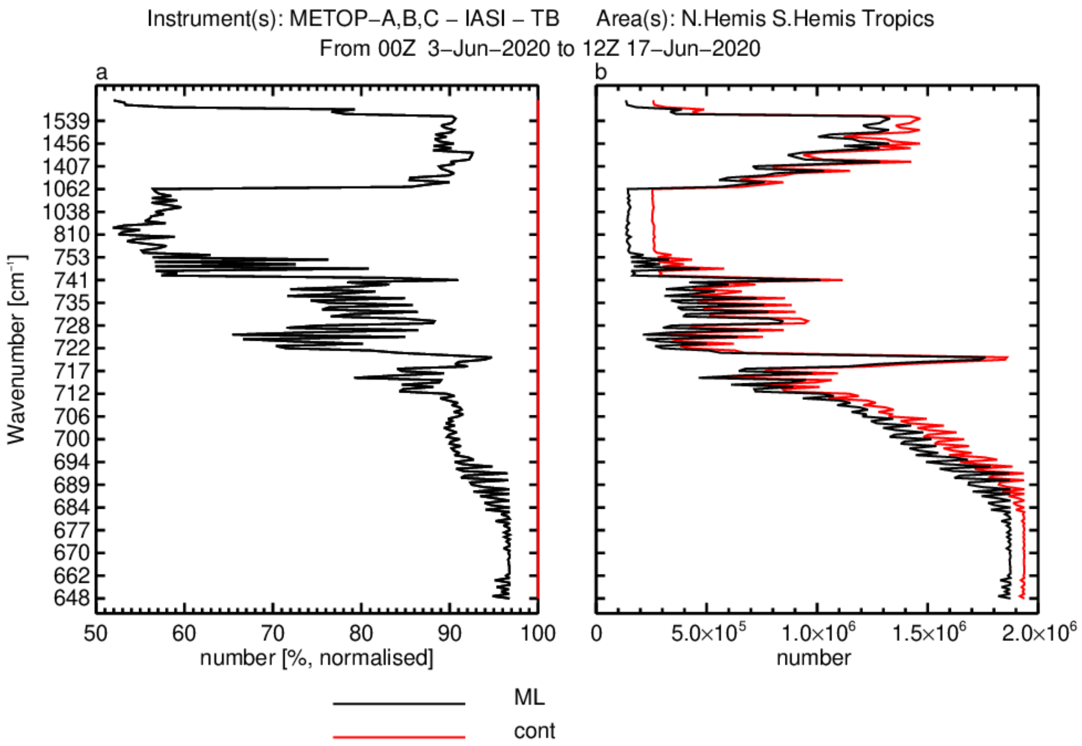

Short range forecast fit to independent observations – neutral is the goal!

- For this short verification period, the impact is neutral to slightly negative.
- IASI window channel numbers are reduced quite considerably, so this degradation could be due to using fewer observations rather than missing cloud-affected obs – more analysis is required.
- On the optimistic side, if the cloud detection was working badly, the degradation would be **very** apparent. So we're not doing too badly, but there is room for improvement.

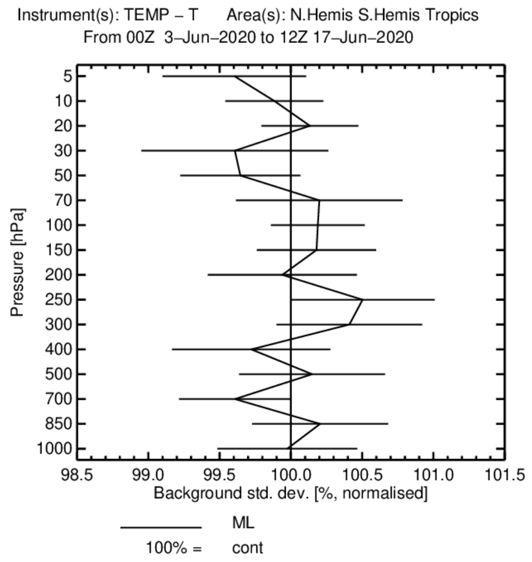

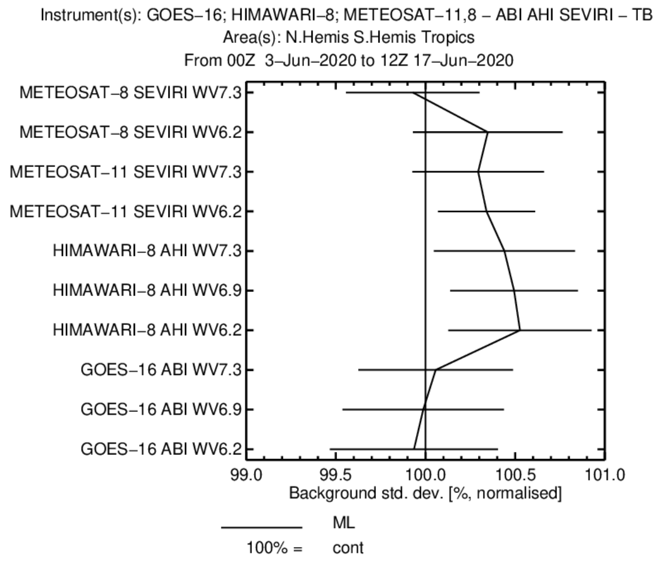

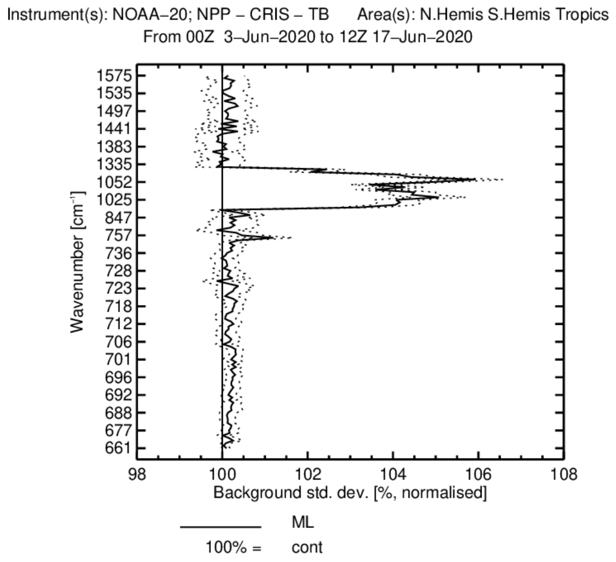

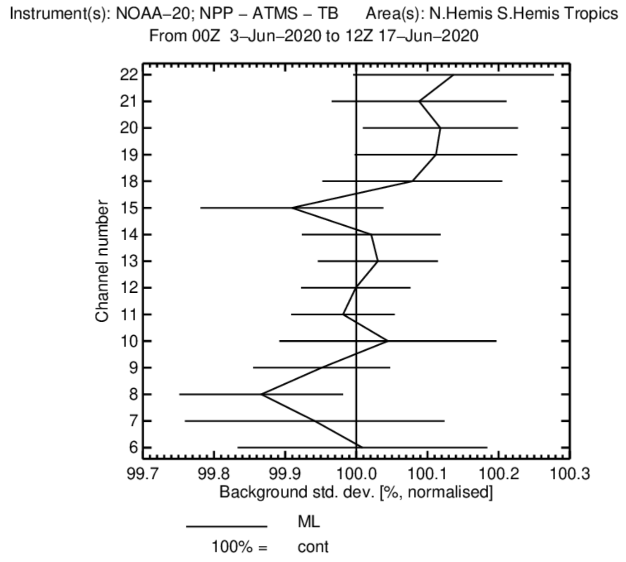

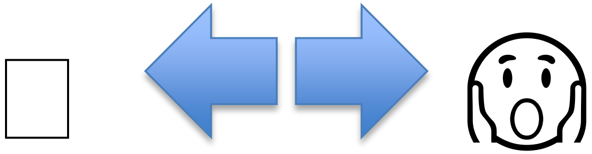

Outlook

- Open science question: is there enough information in the observed spectra to replicate the cloud detection? Tentative indication is “yes”!
- Further refinement of the NN is needed to demonstrate this.
- Assuming it is possible, operational cloud detection could be performed as soon as observations are received (i.e. outside the **“time critical path”**) – no need to wait for background fields to become available. Or, more likely, we can use the NN as an additional pre-screening check).
- Comparisons of computational cost have not been performed, but a day of IASI spectra can be cloud-screened on modest hardware in about 3 seconds.
- The NWP-SAF develops the **CADS** package. Including the neural network in the package is being considered by the CADS development team. We would welcome feedback!

Questions