SnT 2019

CTBT: SCIENCE AND TECHNOLOGY CONFERENCE

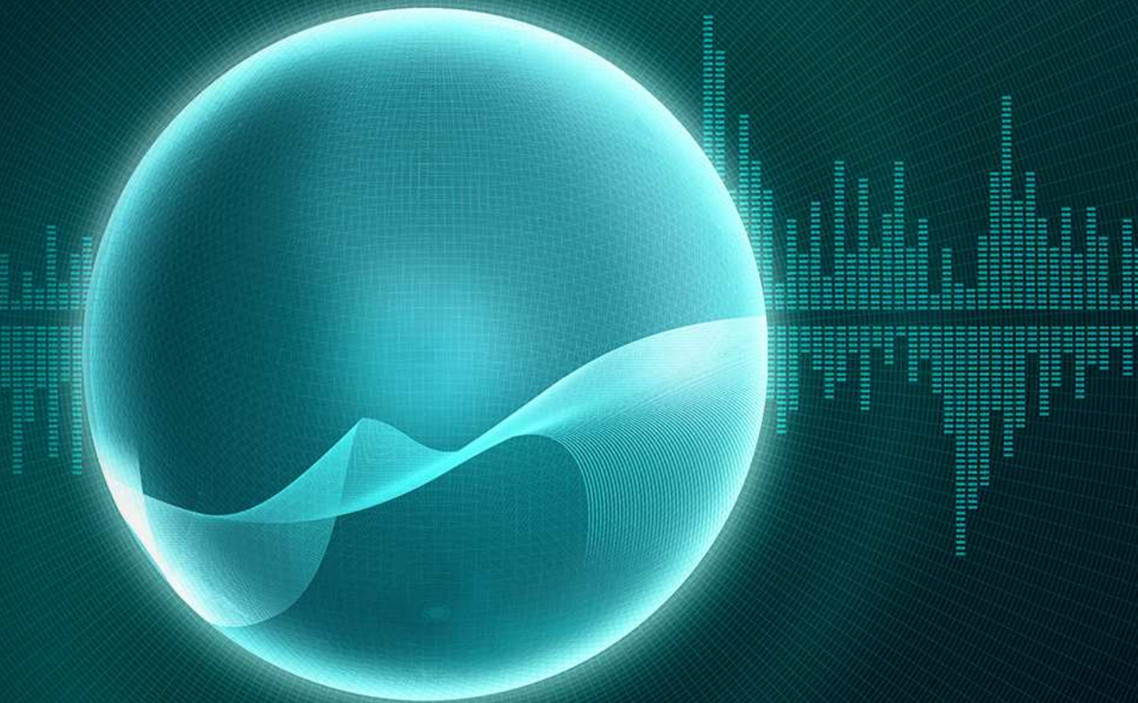

Tryggvi EDWALD, Chief, IDC/APS
tryggvi.edwald@ctbto.org

Machine Learning to Categorize Radionuclide Spectra

Contents

This quick talk outlines experiments with artificial Neural Networks (NNs) which I have applied to Radionuclide Spectra Review, since December.

The ultimate goal of these efforts is to produce tools to assist and support (but not to replace) the human analysts.

This, and similar work by my colleagues, was made Practically possible by IDC's acquisition of a work-station with GPUs that is specialized for numerical processing.

SnT 2019
CTBT: SCIENCE AND TECHNOLOGY CONFERENCE

CTBTO
PREPARATORY COMMISSION

Problem Statement

Classify radionuclide particulate spectra into categories 0-5 as a human analyst would, from the Raw Spectra alone (PS1)

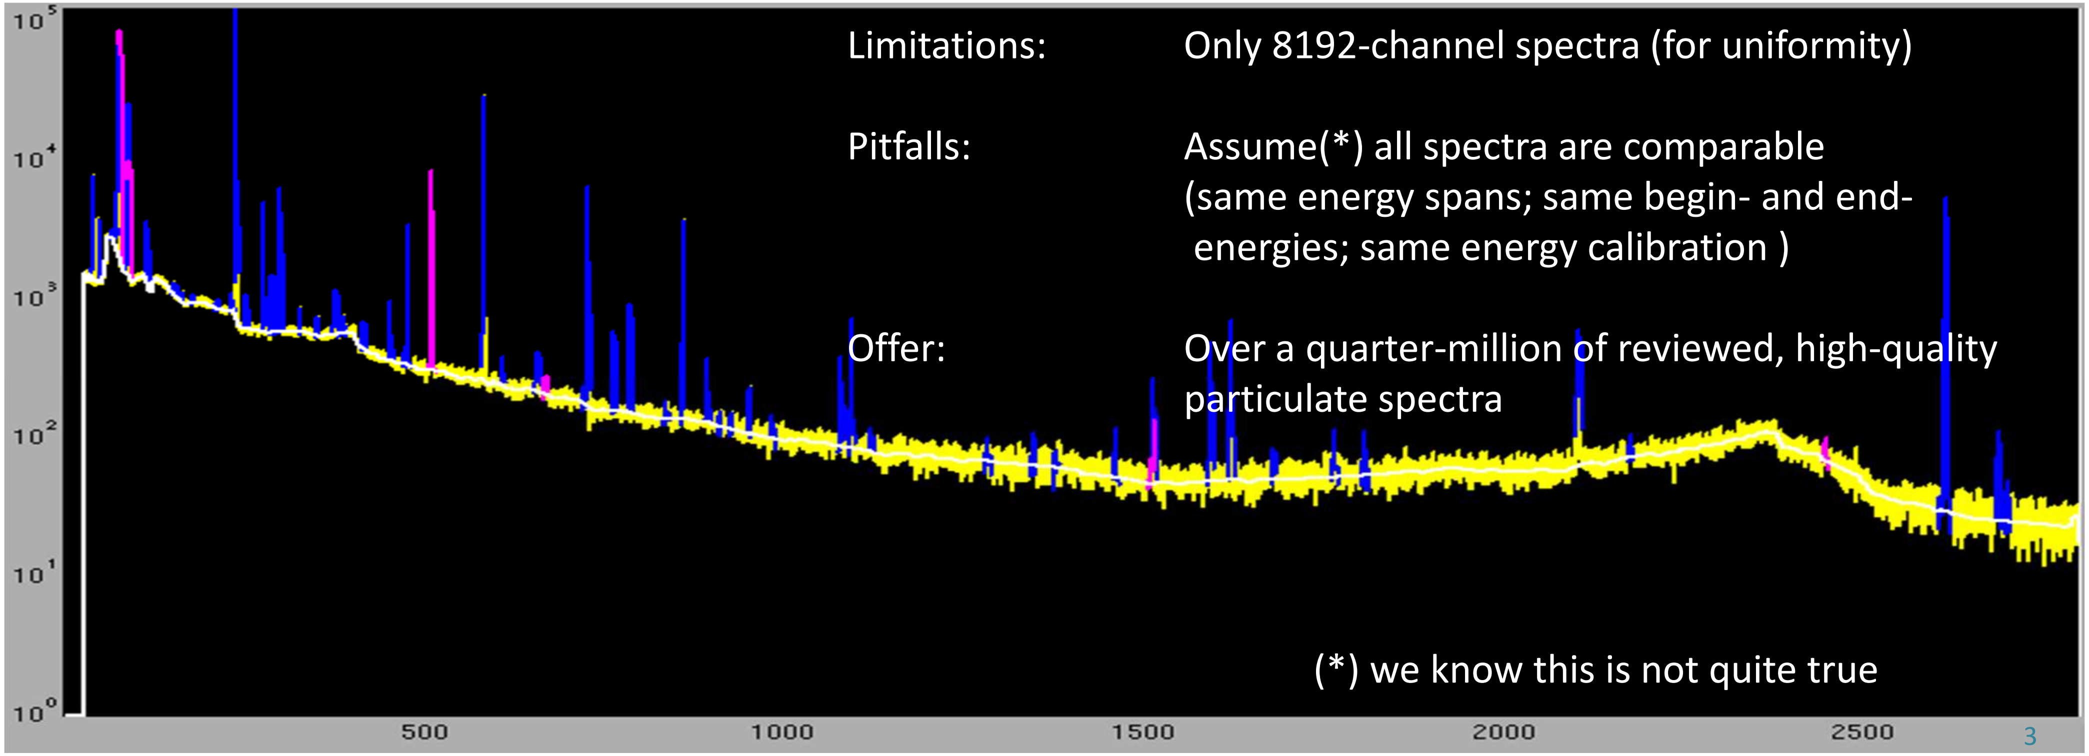

Limitations: Only 8192-channel spectra (for uniformity)

Pitfalls: Assume(*) all spectra are comparable (same energy spans; same begin- and end- energies; same energy calibration)

Offer: Over a quarter-million of reviewed, high-quality particulate spectra

(*) we know this is not quite true

Six-way Classification

Cat	~%	Definition	Concern / Problem
0	4.2	Spectra released w/o categorization (for any reason)	*Could look like a perfectly good spectra with non-0 categorization, as far as the NN can tell.*
1	77.9	Normal natural radionuclide measurement	
2	10.3	Abnormal natural radionuclide measurement	*Hard for the NN to determine what is "abnormal" as that is location-dependent (and wind-direction, and ..)*
3	4.4	Normally observed man-made products present	*This is also problematic due to the definition of "normal". This can be location –dependent, time-dependent, ..*
4	3.0	Abnormal man-made products present	*Hard to determine difference between cat 3 and cat 4*
5	0.2	Multiple abnormal man-made products and at least one fission product present	*This is really the interesting case, very rare, and liable to get lost in the learning*

Decisive Features

Coming at this problem from RN-ignorance, one might theorize:

- Peak shapes are of primary importance

  The shapes of peaks (high counts) in the spectra could be very important. Shapes can be easily recognized by 'Convolutions', in so-called CNNs, Convolutional NNs.

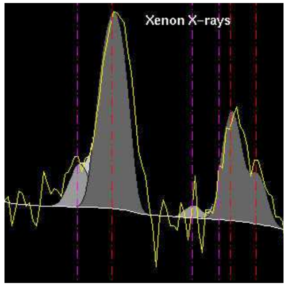

- Peak locations are of primary importance

  The Peak Locations might also be of prime importance, that is, the energy values assigned to the peak.

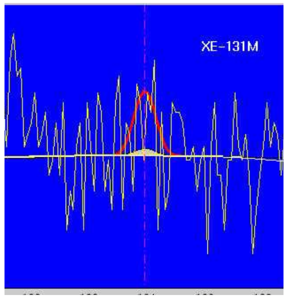

- Or both ..

Neural Networks -1

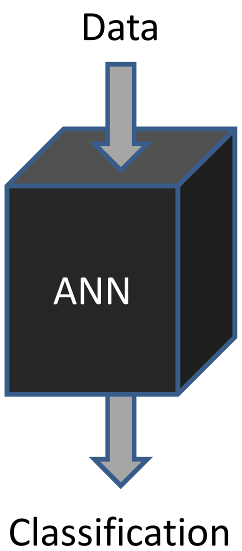

NNs are sometimes seen as 'Black Boxes', but are, for the purposes of my discussion, a mechanism to approximate a mathematical, nonlinear, function to separate (multidimensional) space into well-defined subspaces.
(Eg. the 8192-dimensional space of RN Spectra separated into six well-defined subspaces, corresponding to their Categories.)

As such, an ANN is *subject to under- and overfitting*, much like when fitting a curve to a set of (x,y) coordinates in X,Y space.
And they are *very heavily dependent on the quality* and volume of the input data (the training data).

Neural Networks -2

The NN takes a spectra, applies its function to it, and predicts the category. The NN then compares the predicted category to the human-bestowed category, and finds the difference. This difference, or Prediction error, is distributed back into the function, and we try again. And again. Gradually, the function converges to a form that predicts quite correctly.

Due to the size of the dataset and the input data, the calculations become massive matrix manipulations, and therefore difficult to follow. But there is nothing here that C.F. Gauss couldn't have improved on, back in 1850.

Neural Networks -3

Happily, one really doesn’t need to know the full details, to use NNs. Many open source toolkits are now available that essentially remove this need.

Much of the research in NNs is on speeding up convergence to a good predictor, and improving performance.

Recent work aims at explaining the results, eg. why a network has decided the category “5” for instance. NNs currently cannot “explain themselves” and this is a significant and serious weakness in general.

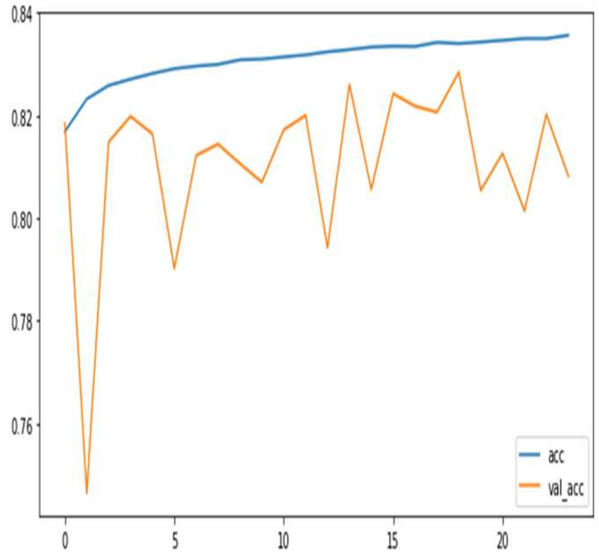

(On the right: tracking accuracy and validation accuracy during training of a NN.)

Model

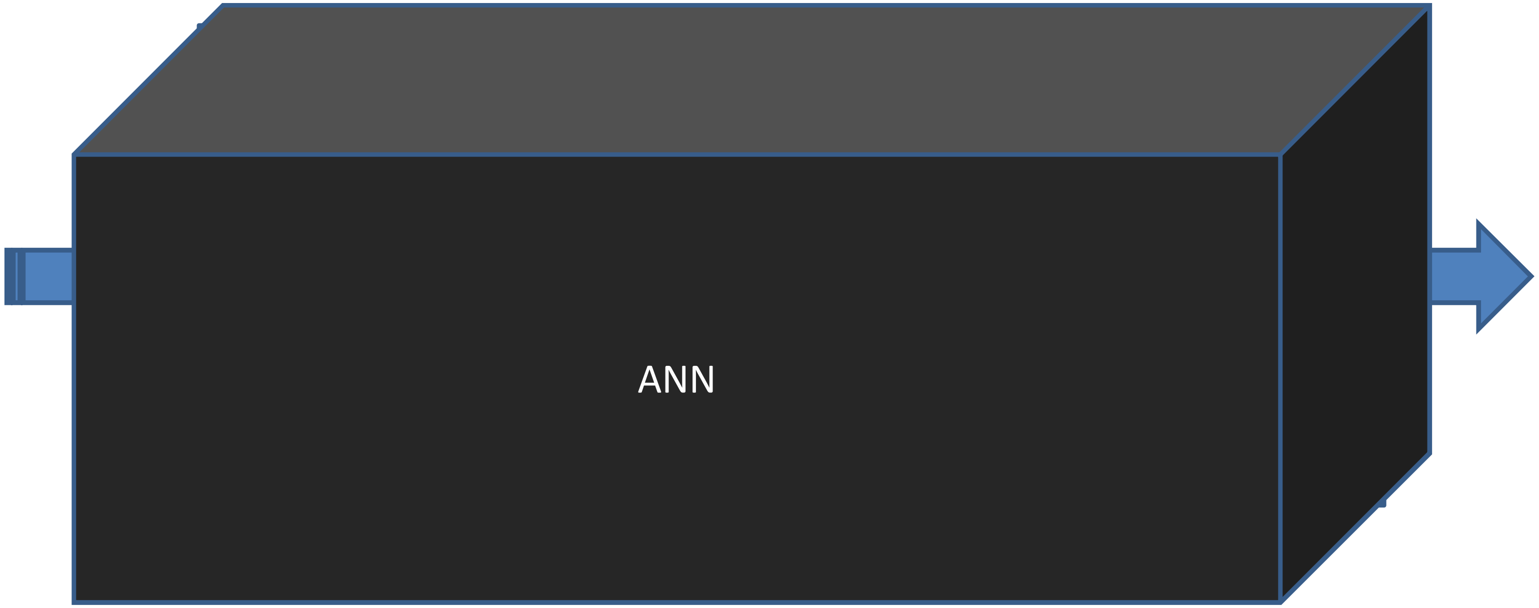

This particular model had the structure shown above, and a relatively modest 17,886,038 trainable parameters

Model

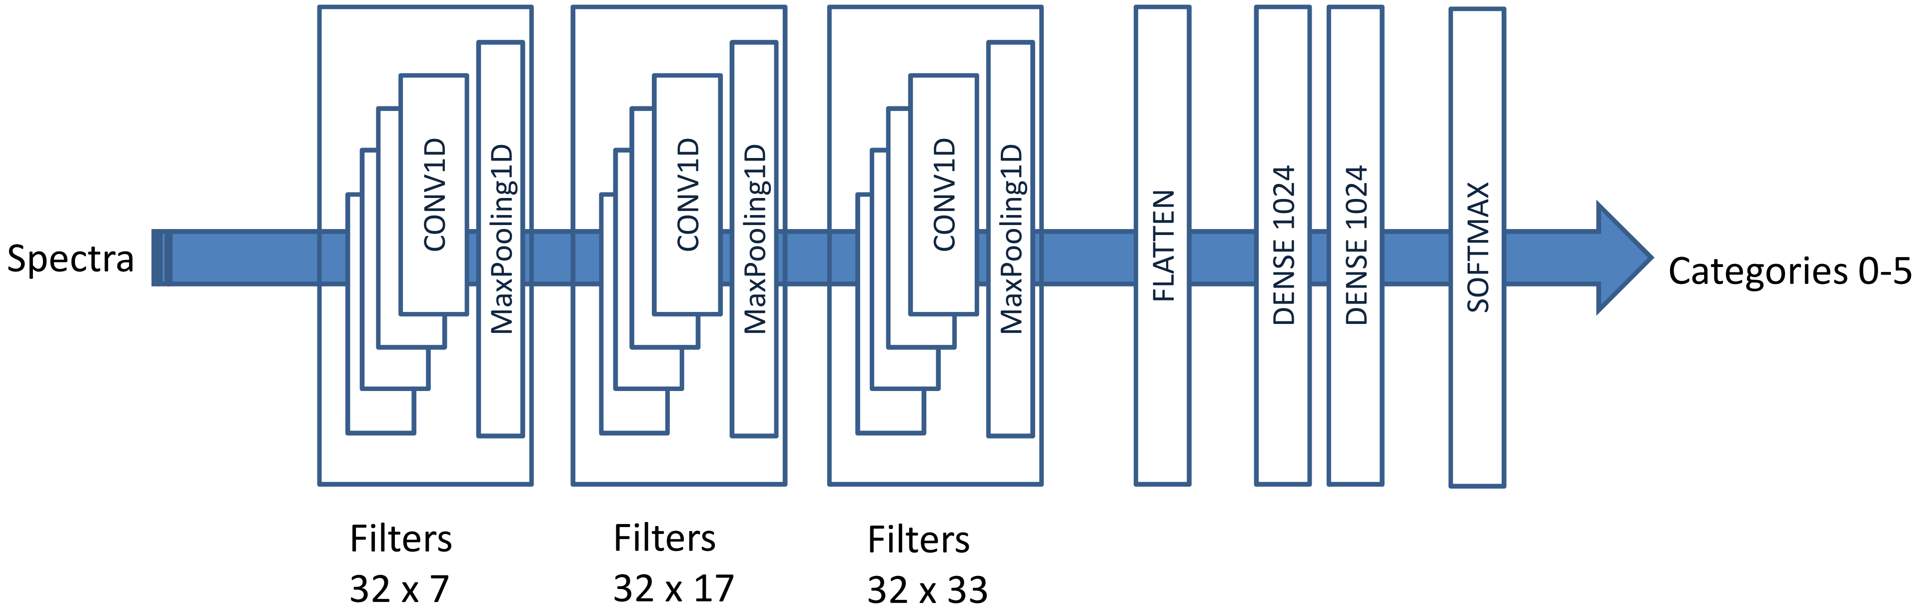

This particular model had the structure shown above, and a relatively modest 17,886,038 trainable parameters

Results / Predictions

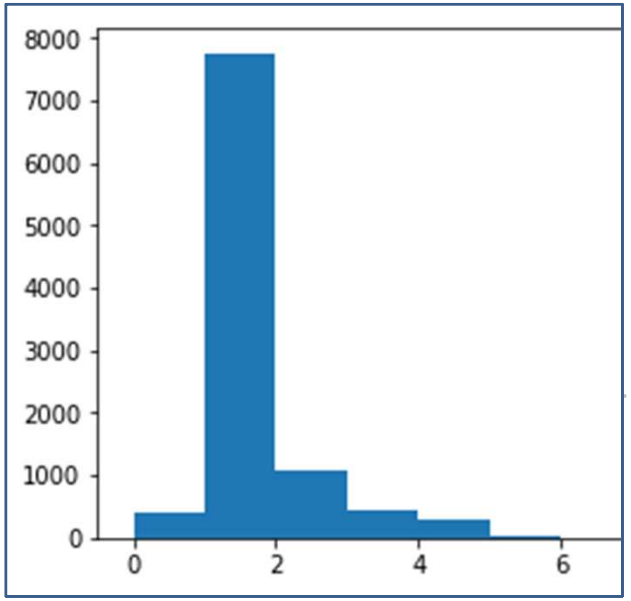

Predictions on unseen test data

This run: 30 epochs; batch size 128;

Accuracy on unseen test set 78.25%
(Test set 10K spectra, with same Cat-distribution as the main training set.)

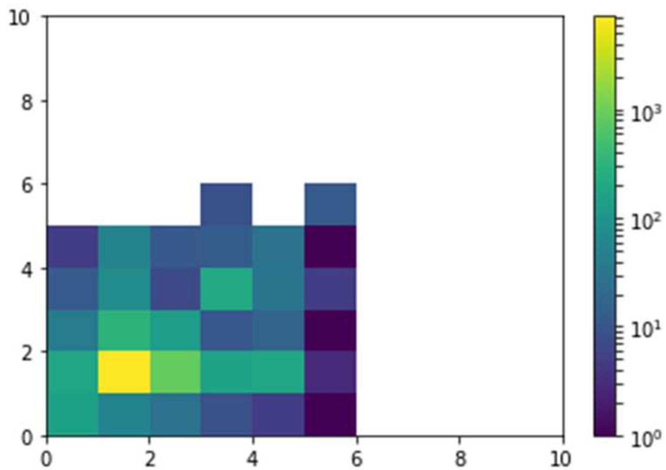

Confusion matrix

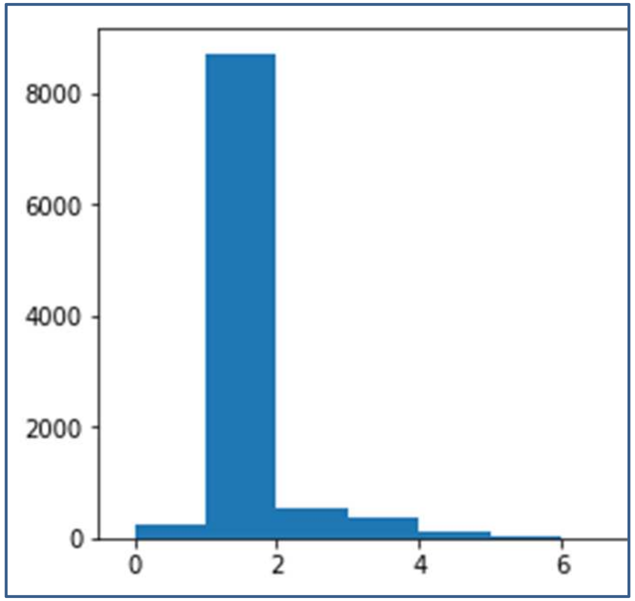

True distribution of same data

Binary Classification

Class	%	Definition	Concern / Problem
1	~ 80	Normal natural radionuclide measurement	
0	~ 20	Any other situation (categories 2-5)	

Use only the above two classification.

The idea being, we can still save analysts a lot of work by drawing attention only to the “unusual” spectra (~20% of the whole set).

Revised “Problem Statement”:

Classify radionuclide particulate spectra into two classes as a human analyst would, from the Raw Spectra alone (PS1’)

Model

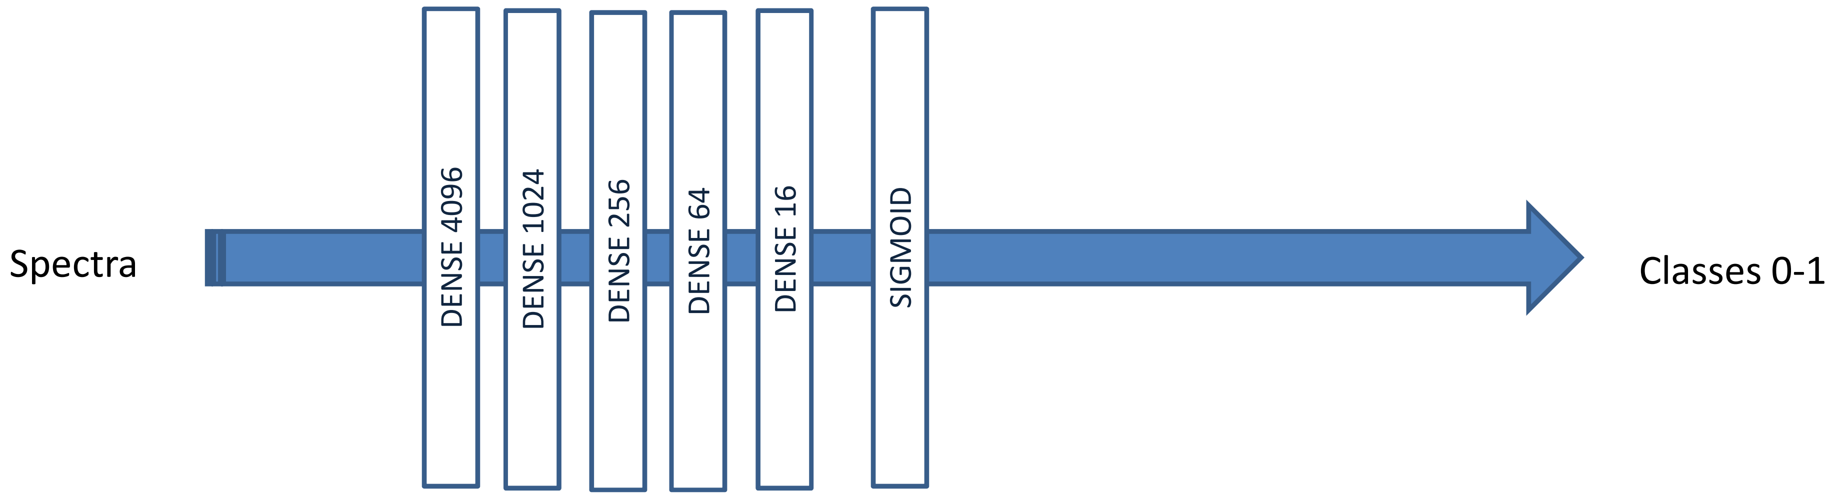

One particular Binary-Classifier model had the structure shown above, which looks pretty harmless, and Weighs in with 38.033.761 trainable parameters. A training session took around 40 minutes.

Predictions (Binary Model)

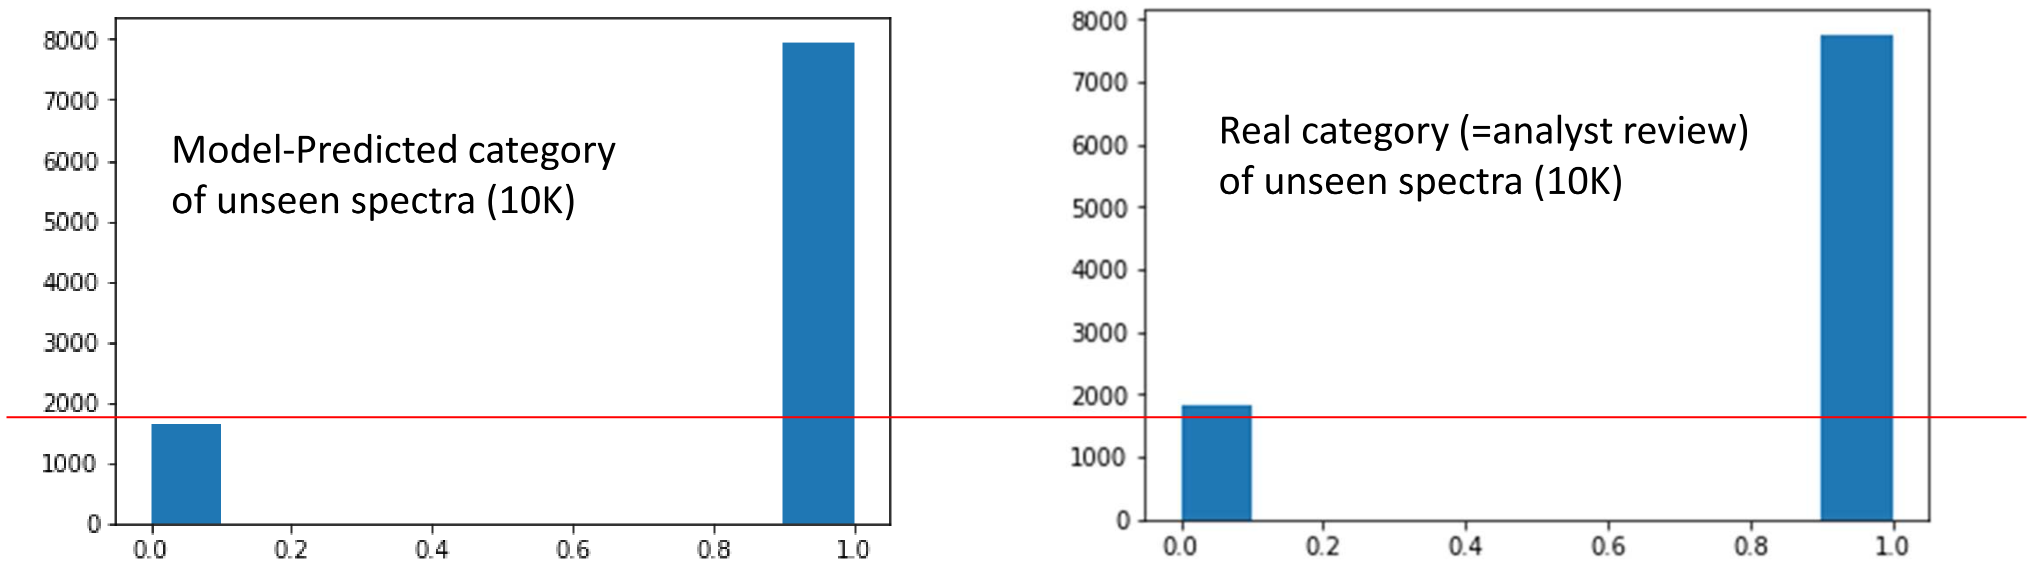

After training, the model was tested on 10K unseen spectra, and resulted in:

`Loss 0.4265 , Accuracy 82.62% on test data`

Not Yet Good Enough

This result doesn't look too bad, for a starting investigation. But let us not forget that the "model" below would also have around 80% correctness. (And no trainable parameters.)

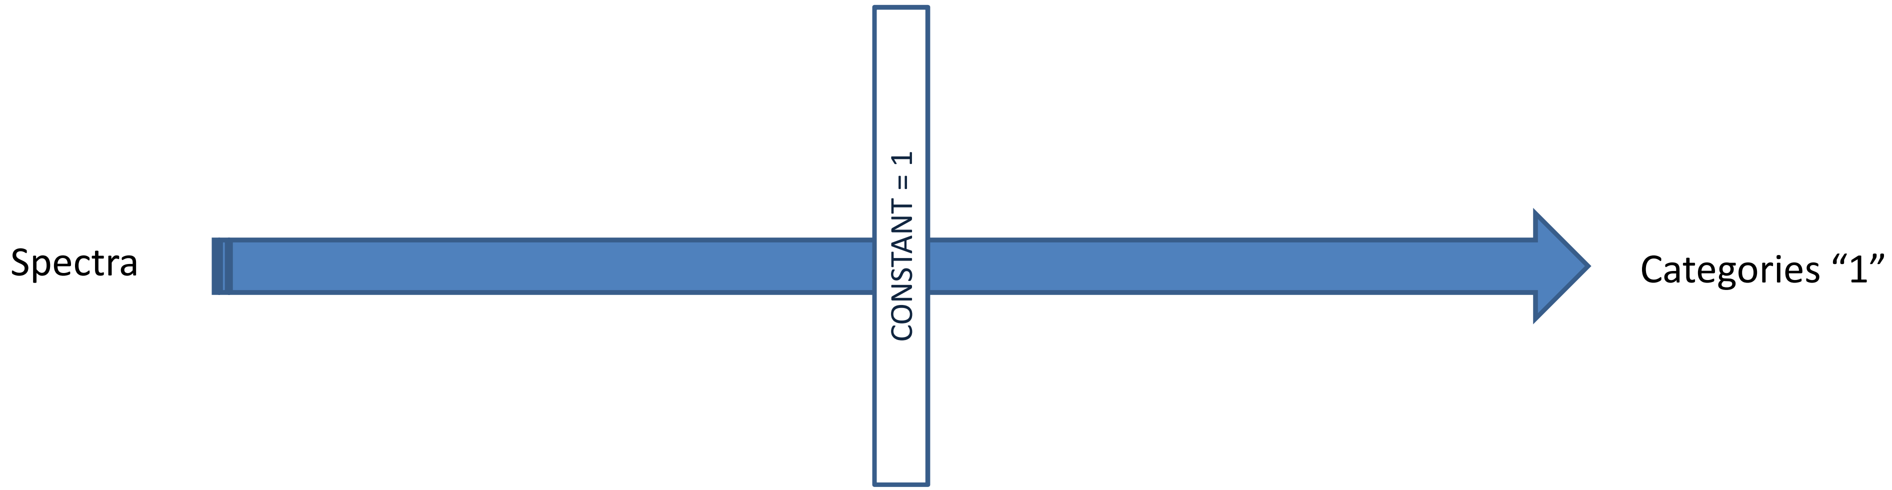

Possible Paths Forward – Key Takeaways -1

Calibration

Calibrate and equalize the spectra, to correct for variations in detectors (make the spectra fully comparable) The necessary parameters are available in the RN databases

Apply (Detector) history to spectra

Bring location (detector) history, eg. short- and long-term averages (possibly including Backgrounds) when considering what is "normal"

Possible Paths Forward – Key Takeaways -2

Using detector IDs as input data

Could add a 8193’rd input parameter: the detector ID, assuming that each detector remains relatively stable.

This might avoid the need for applying calibration to the training spectra.

This is easy to implement, but possibly invites a *“Clever Hans Effect”*.

Other "Next Steps"

Investigation has begun into using 2D classifications on beta-gamma spectra. There are not enough results yet, to discuss here.

Some similar situations exist as in the Particulate case, eg. *image feature location is of prime importance*. Perhaps it is not possible to learn anything from the raw 2D images, as there is a large variation in energy spans on both axes.

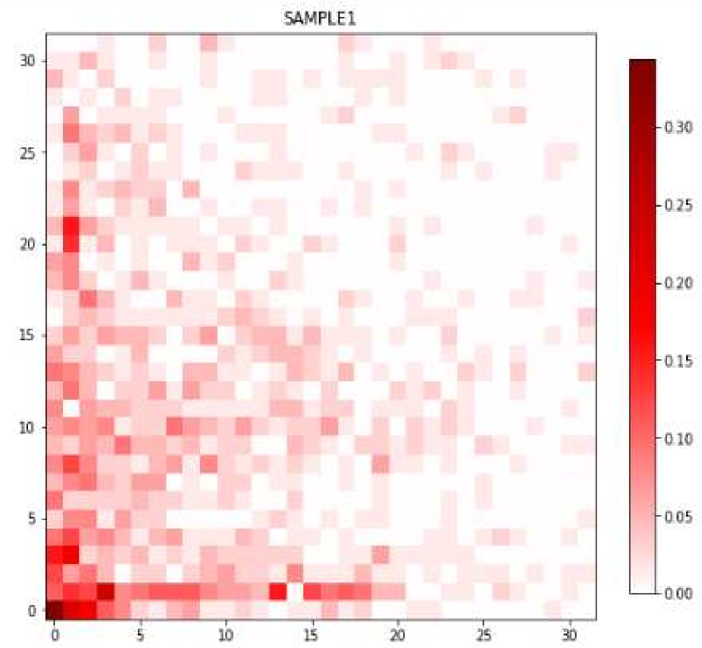

There will be a live run in the Gardehalle I, Thursday 27/6 at 11:30, where I will be delighted to discuss technical details with anyone even slightly interested.

(and anytime via tryggvi.edwald@ctbto.org)


THANK YOU