# Applied ML

AutoEncoders

Troels C. Petersen (NBI)

"Statistics is merely a quantisation of common sense - Machine Learning is a sharpening of it!"

Encoders

An **encoder** is a network that can **take a signal and produce a code**!

Typically, the code is a description of the signal, which could be images, sound, etc. The code is usually "smaller" than the signal, not unlike a summary. An image classifier is an example of an encoder.

Hollerith 80-column punch card (1950s)

AutoEncoders

An **AutoEncoder** (AE) is a **coupled pair of encoder and decoder**. The encoder maps signals into code, and the decoder reconstructs the original signal from those codes.

The pair is trained to have the most accurate reconstruction: If you give a signal **x** to an encoder E to get **y** = E(**x**), then the decoder D should ensure that D(**y**) is close to **x**.

One application is unsupervised feature learning, where it tries to construct a useful feature set from a set of unlabelled images. We could use the code produced as a source of features.

PCA as an autoencoder

A PCA is a linear AE. One can project a higher dimension down on (fewer) principle components (encoding) and then "reconstruct" the original data from the latent space, by choosing the low dimensional points in the original dimensionality.

PCA as an autoencoder

A PCA is a linear AE. One can project a higher dimension down on (fewer) principle components (encoding) and then "reconstruct" the original data from the latent space, by choosing the low dimensional points in the original dimensionality.

5

Usage of AE

AEs are used for many things, such as:

- Unsupervised learning (e.g. clustering) on images, sound, graphs, etc.
- Compression (with loss!) of e.g images
- De-noising and inpainting images
- Anomaly detection
- Training on large dataset with few labels

Most AEs are CNN based and produced for images. However, the AE concept is more general, and applies to anything, that can be passed through an NN.

Usage of AE

AEs are used for many things, such as:
- Unsupervised learning (e.g. clustering) on images, sound, graphs, etc.
- Compression (with loss!) of e.g images
- De-noising and impainting images
- Anomaly detection
- Training on large dataset with few labels

Most AEs are CNN based and produced for images. However, the AE concept is more general, and applies to anything, that can be passed through an NN.

The most central hyperparameters to consider are:
- Size of the latent space (code)
- Architecture of NN (layers and nodes)
- Loss function

As we shall see, these HPs to some extend determine what type of AE you're making.

AE Unsupervised

Since one might have no knowledge of the content of images (or sound, etc.), training an AE is inherently unsupervised. The result is simply a latent space that is a good **representation** of the images.

However, this can be used to cluster "less simple data", as one can apply both dimensionality reduction and/or clustering to the latent space.

This enables one to analyse very complex data in an unsupervised manner. (e.g. "How many zebra calls exists?")

AE Compression

The latent space is also called the "code", and is a good (but not perfect) representation of the image.

This very "code" is an efficient way of compressing images - or anything else that can be fed into an NN - down to a fixed size.

The approach has proven to be rather performant, competing with JPEG2000.

An Autoencoder–based Learned Image Compressor:

We propose a lossy image compression system using the deep-learning autoencoder structure… []. Our aim is to produce reconstructed images with good subjective quality subject to the 0.15 bits-per-pixel constraint.

cs > arXiv:1902.07385

9

De-noising AutoEncoders

AutoEncoders can be used for de-noising for example images.

The method consists of first introducing noise (partially ruining the input), and then train an autoencoder to **produce the original image from the noisy input**. Once this is learned, noisy images can be de-noised using this network.
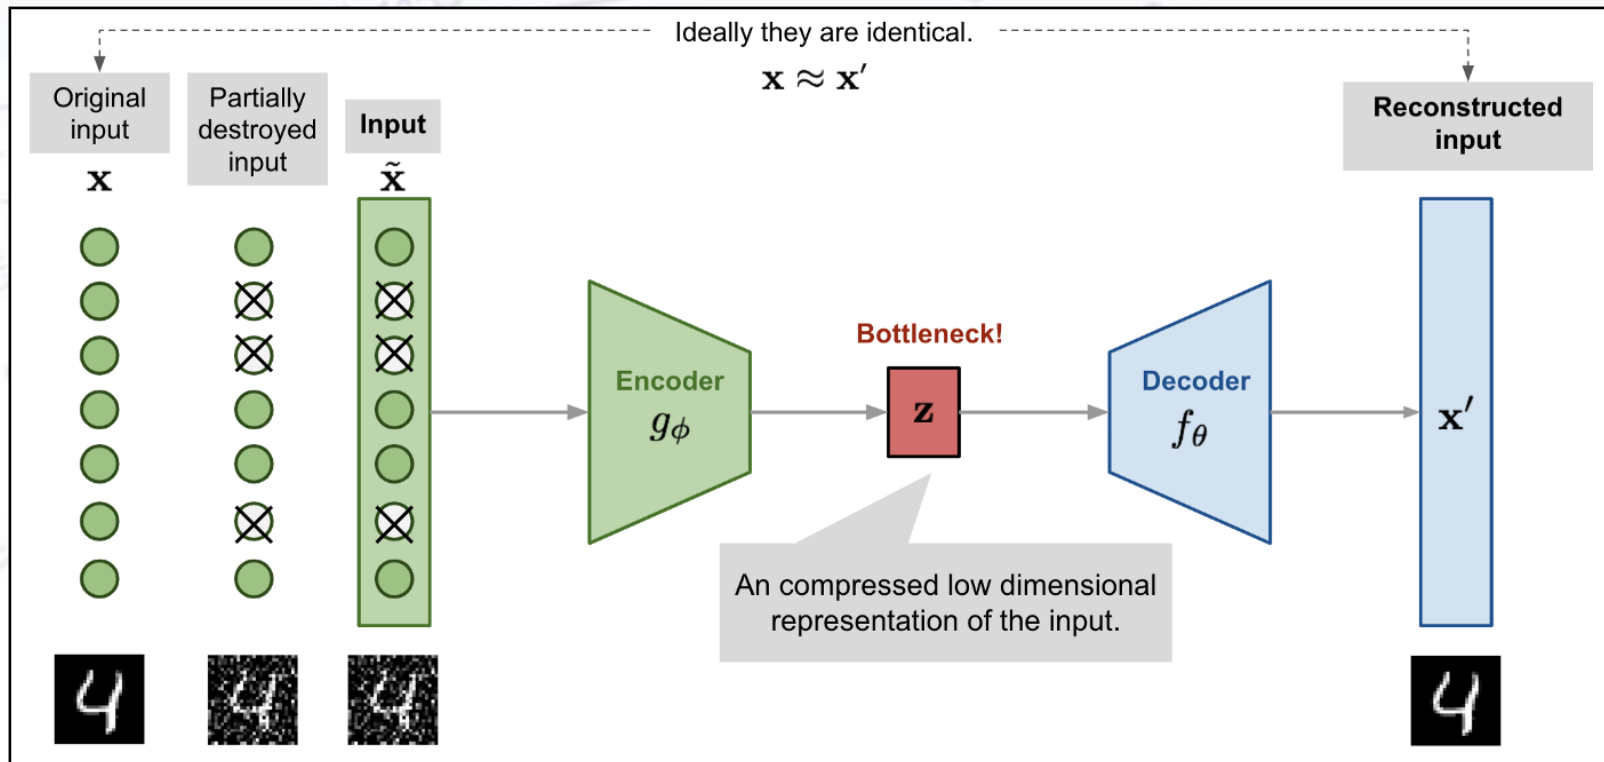
De-noising AutoEncoders

Another way of making this work is perceptual loss. Here one does not (only) compare the input and output images, but also (or only) the subsequent layers in the CNN. These layers learns how images in general works, and thus aren't as affected by noise.

This idea can be taken further to reconstructing larger parts of an image. This is called "inpainting".

Tutorial on De-noising AutoEncoder

Inpainting with AEs

AEs can also be used for "inpainting", which means replacing damaged/lost parts of an image.

Inpainting with AEs

AEs can also be used for "inpainting", which means replacing damaged/lost parts of an image.

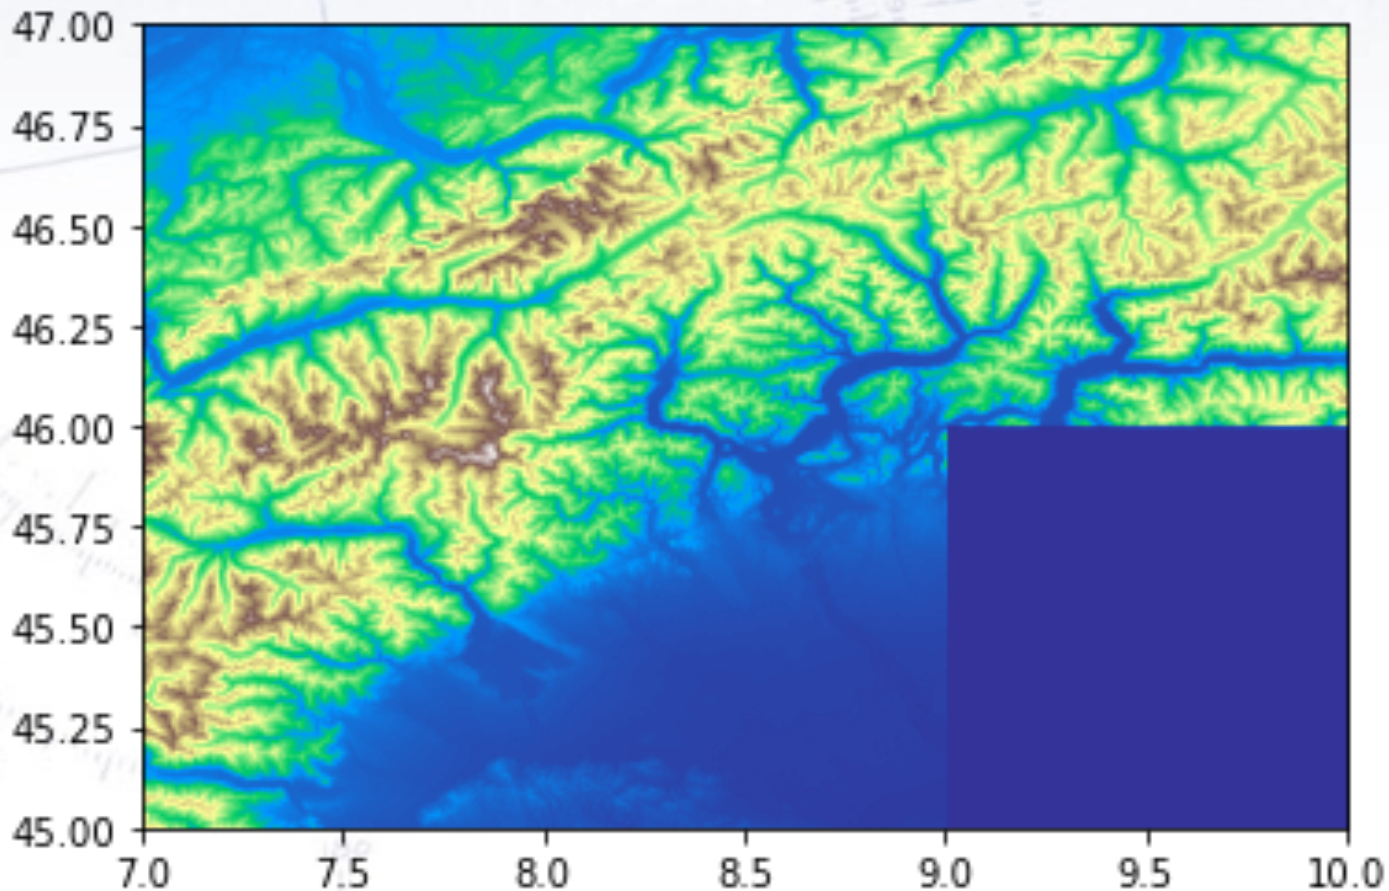Original Image   Free-Form Input   Our Result   Original Image   Free-Form Input   Our Result

Inpainting with AEs

AEs can also be used for "inpainting", which means replacing damaged/lost parts of an image.

This can be used for estimating the volume of glaciers, given an altitude map.

AE Anomaly Detection

By learning to replicate the most salient features in the training data under some of the constraints described previously, the model is encouraged to learn to precisely reproduce the most frequently observed characteristics.

When facing anomalies, the model should worsen its reconstruction performance.

AE Anomaly Detection

By learning to replicate the most salient features in the training data under some of the constraints described previously, the model is encouraged to learn to precisely reproduce the most frequently observed characteristics.

When facing anomalies, the model should worsen its reconstruction performance.

In most cases, **only data with normal instances are used to train the autoencoder**. In others, the frequency of anomalies is small compared to the observation set so that its contribution to the learned representation could be ignored. After training, the autoencoder will accurately reconstruct "normal" data, while failing to do so with unfamiliar anomalous data.

Reconstruction error (the error between the original data and its low dimensional reconstruction) is used as an anomaly score to detect anomalies.

Note: *It has been observed that sometimes the autoencoder "generalizes" so well that it can also reconstruct anomalies well, leading to the miss detection of anomalies.*

AE Training Assistance

Another AE use is "training assistance". Image that you have a large data sample, but only few labelled cases (i.e. few cases where you know the result). The few labelled cases might not be enough to train a full CNN (which can have millions of parameters!).

AE Training Assistance

Another AE use is "training assistance". Image that you have a large data sample, but only few labelled cases (i.e. few cases where you know the result). The few labelled cases might not be enough to train a full CNN (which can have millions of parameters!).

But if you train an AE on the large un-labelled dataset, then once this network is in place, you can continue from the latent space.

Now using the small labelled data set, you train a classifier (or regressor).

Insolubles in ice cores

…an example case

The data

The data consists of two parts - a training set (control) and a test set (ice core):
- Training: Obtained from (controlled) samples, that was carefully selected.
- Testing: Obtained from melted ice core filtering process.

All the images are gray scale and scaled to be the same size: 128 x 128
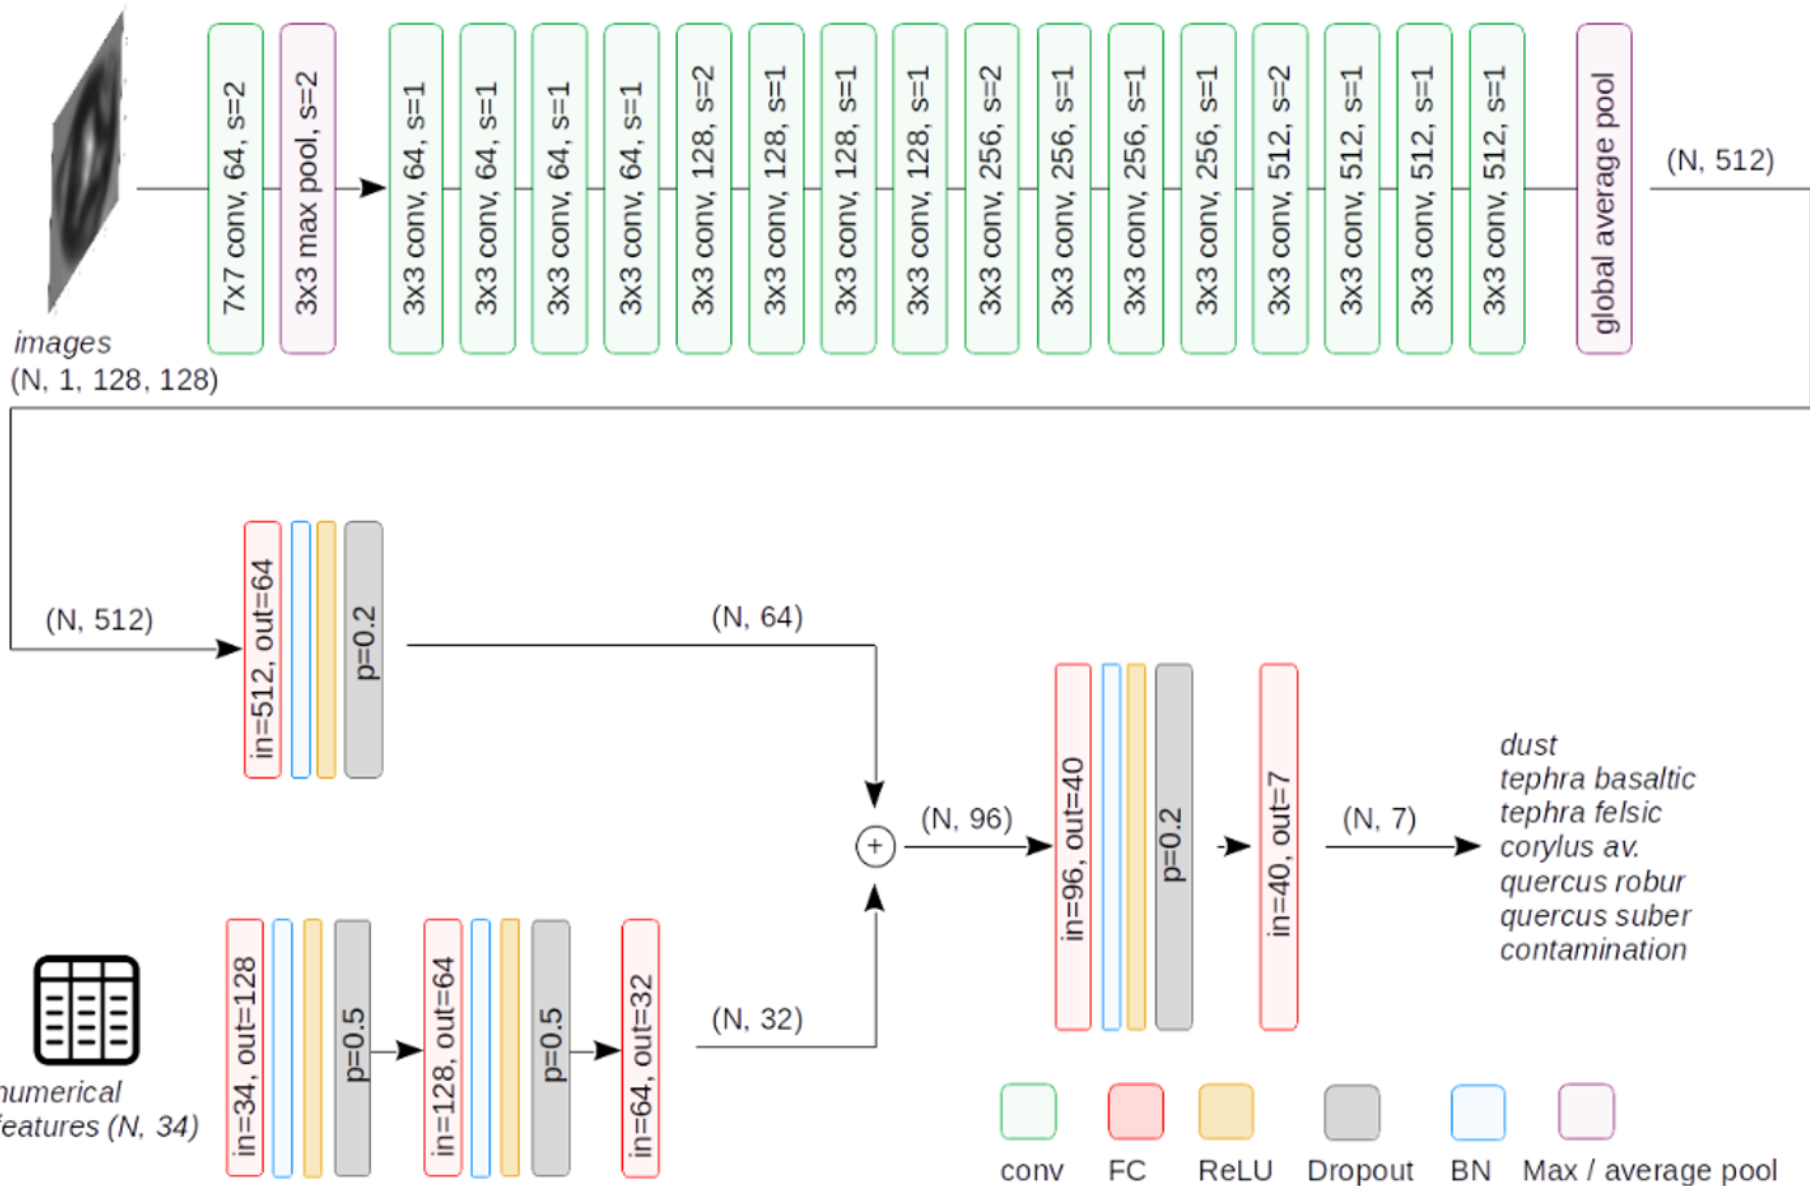In addition, we have 34 numerical features from the microscope imaging.

The initial goals are:
- Be able to classify the training images (supervised learning).
- Apply the classification to the testing images (predictions).

But how can we then know, if there are **other things** in the ice cores?
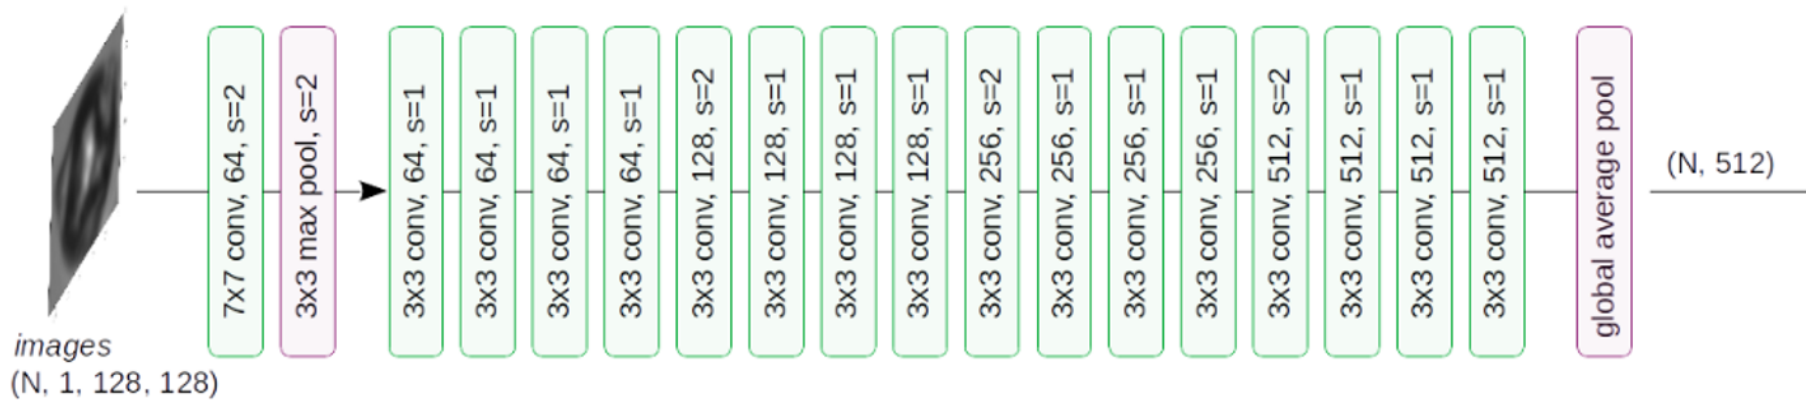
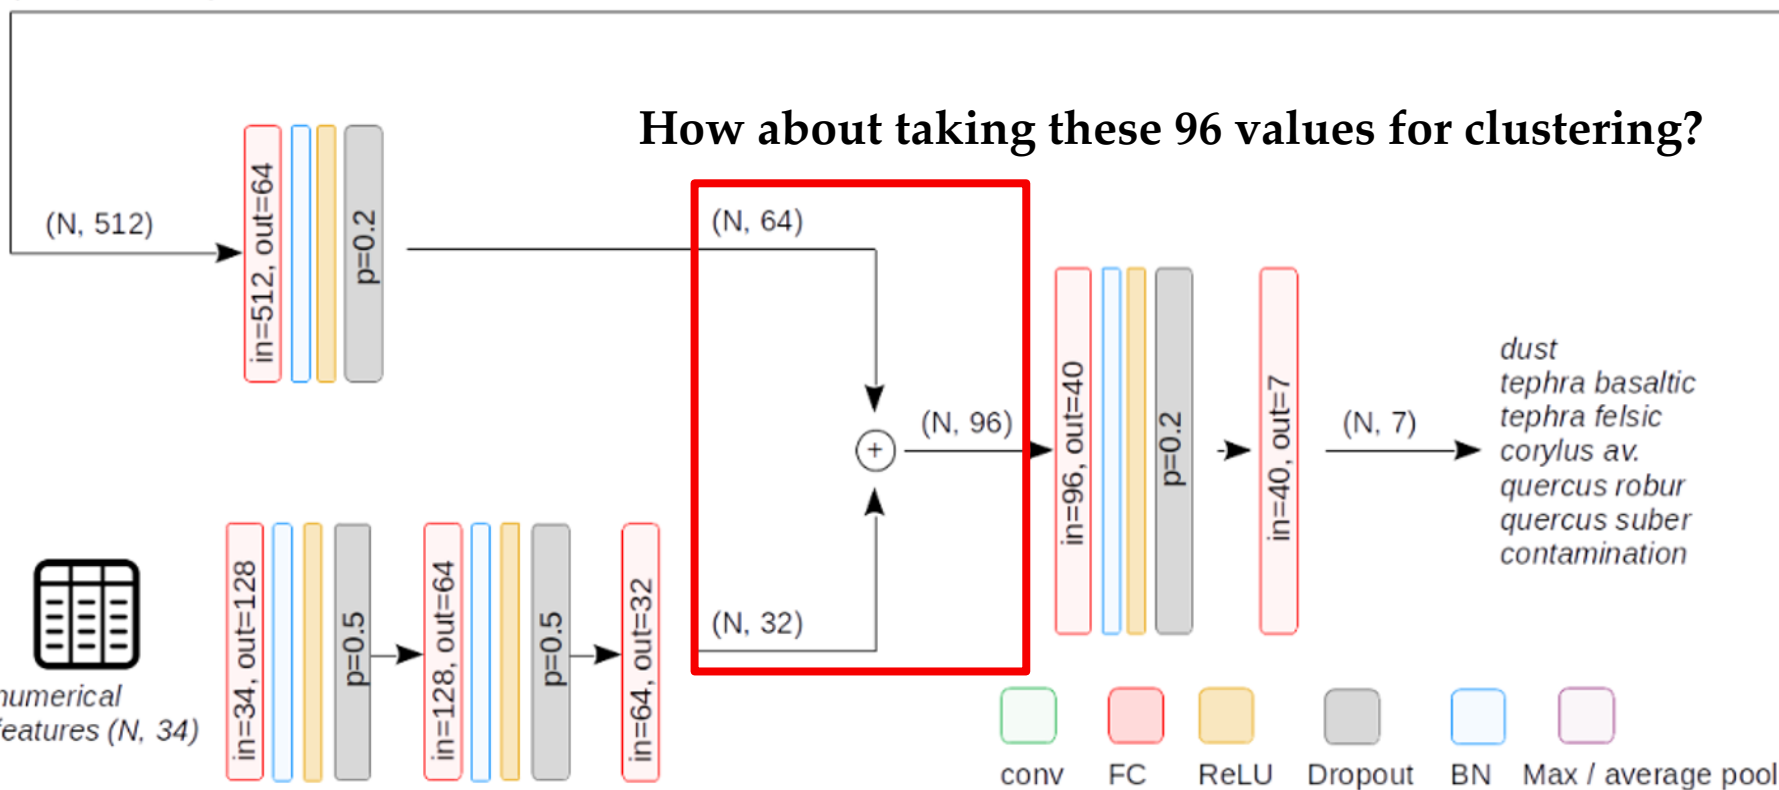# Classification CNN architecture

Performance

The CNN performance is rather good:
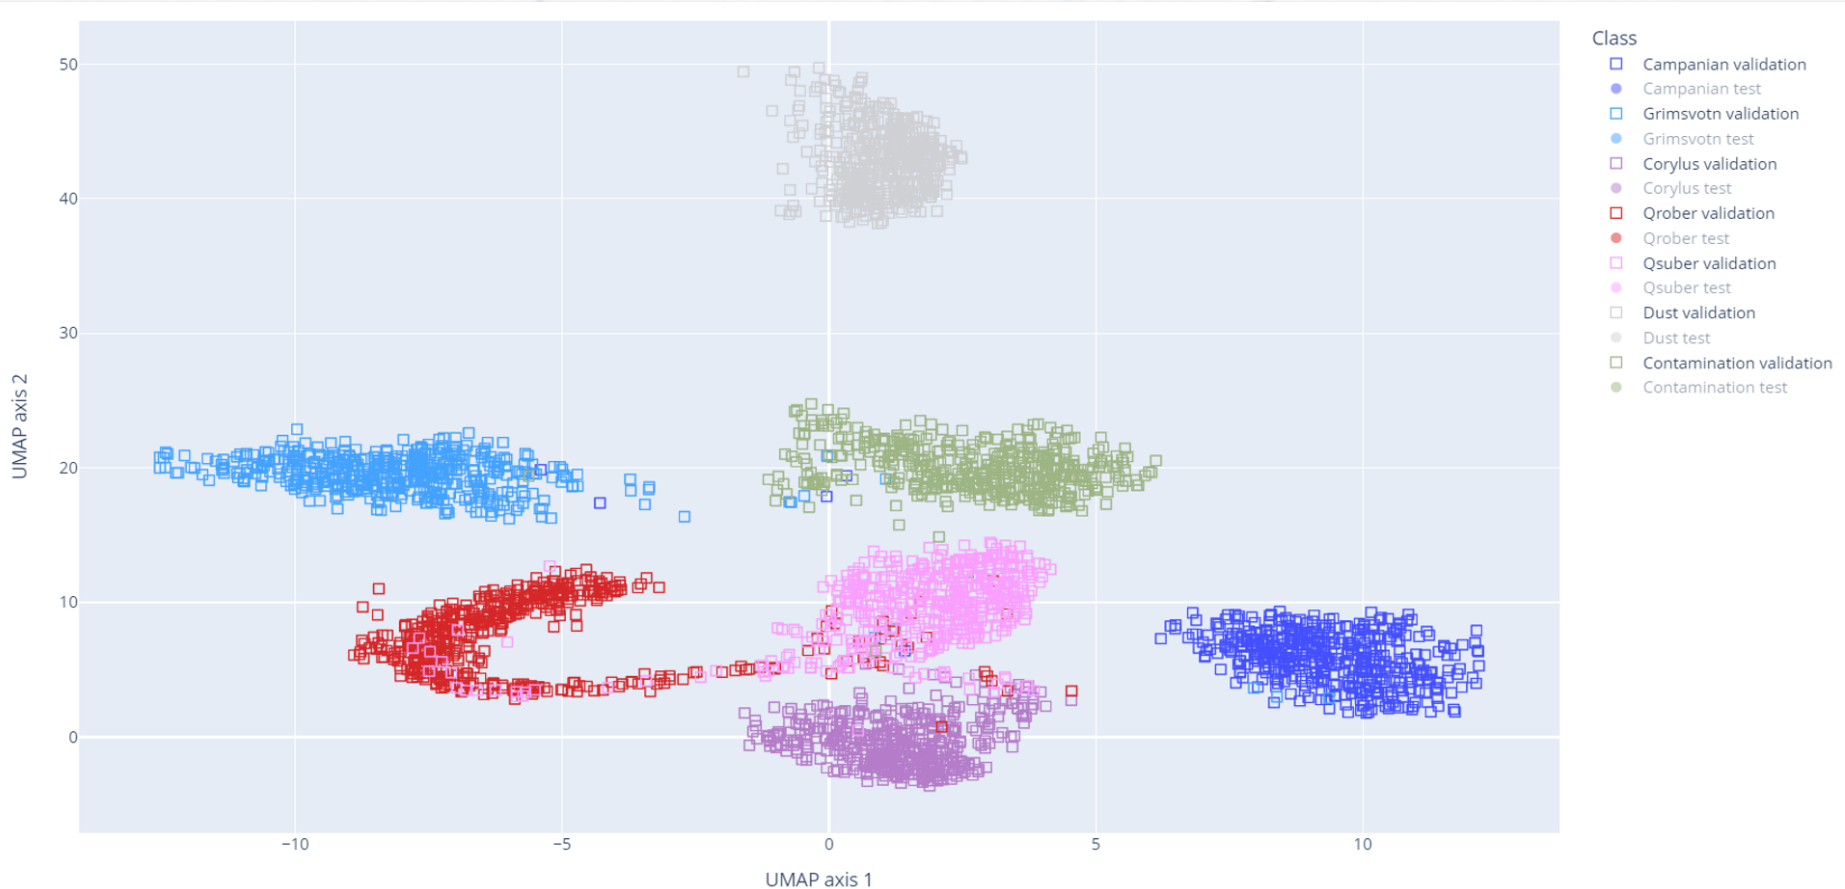
Classification CNN architecture

How about taking these 96 values for clustering?

UMAP clustering

Clearly, this clusters the sample quite well.

As was seen before, the two types of pollen - naturally - have some overlap.

UMAP on testing set…

Applying the UMAP to the testing set… Well, how would you even do that?

UMAP on testing set…

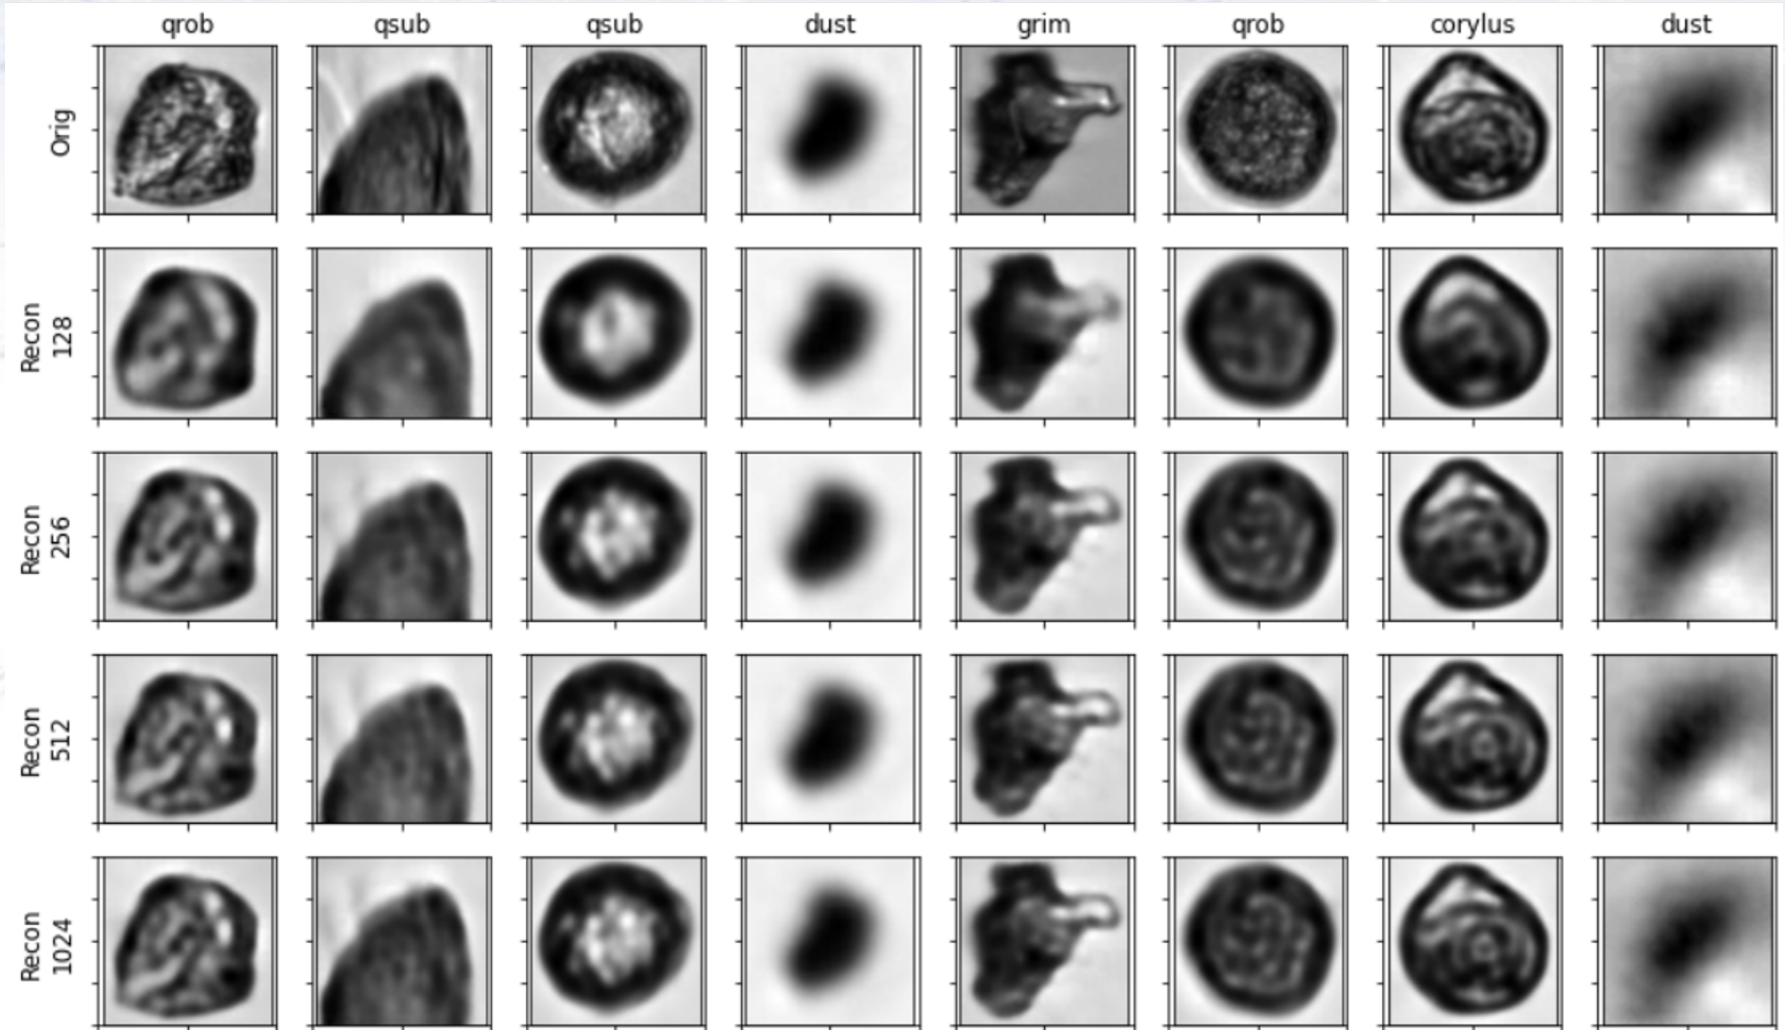Applying the UMAP to the testing set… Well, how would you even do that?

The result using ParametricUMAP is as follows (dust and cont. dominates).

Autoencoding the images

What size should the latent space be?
You should consider what quality of the (output) images you want.

The unsupervised result...

Not using the training sample (which might bias the classification), we have managed to produce the following distribution.
Dust is clearly separated, as is the pollen (though not internally) and the two tephra types. Contaminations overlap mainly with the tephra.