

Deep learning introduction



netherlands
eScience center

ChatGPT

<https://poe.com/>



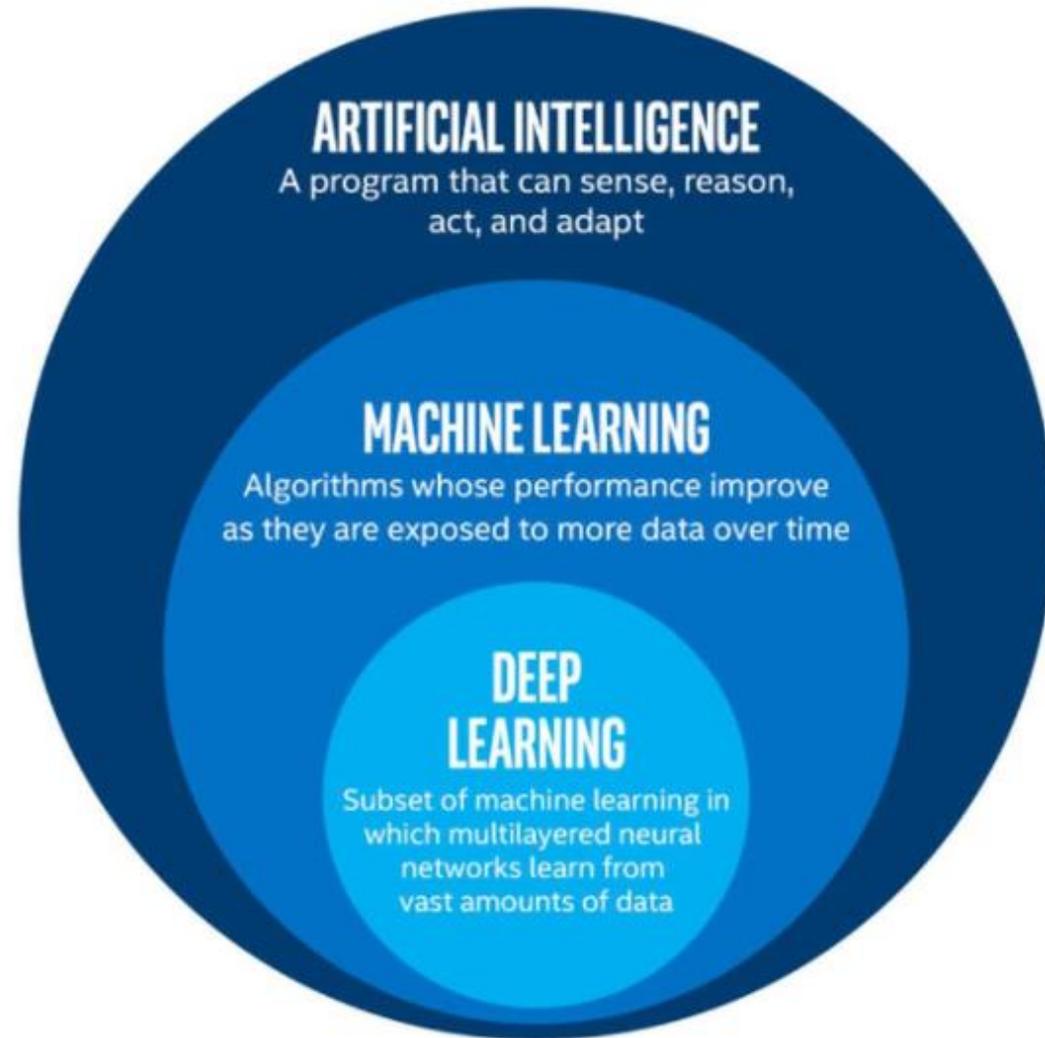
Rationale

Provide you with just enough intuition about deep neural networks to start implementing them.

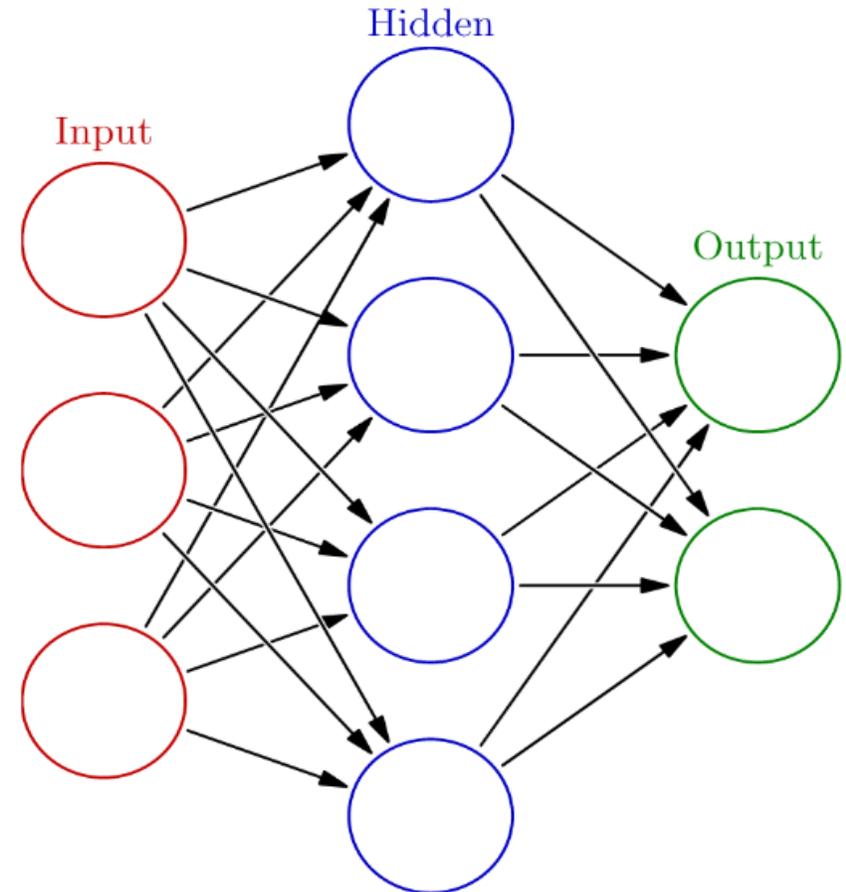
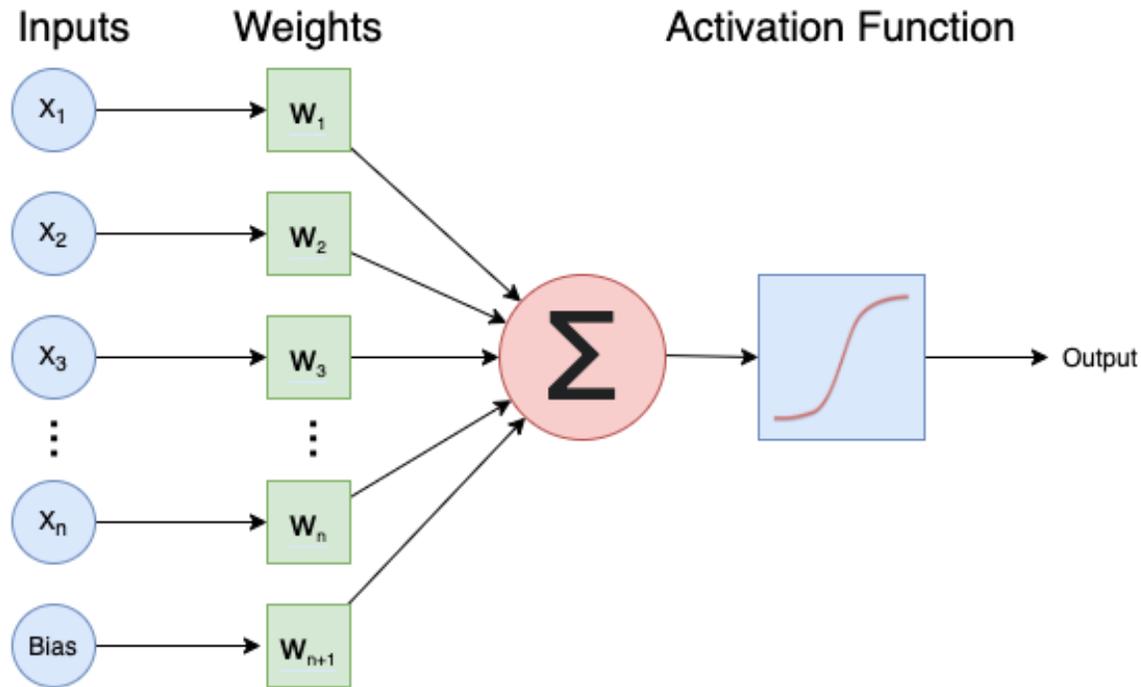
Then learn by doing.



AI Landscape



Neural network



Networks as a team of detectives

- Each neuron has a certain expertise, focusing on particular patterns in the evidence (the input data)
- Together they make a decision on who is guilty.



GPT-4 explaining what neurons do in GPT

Bills et al. (2023) trained GPT-4 to give explanations of what neurons in another GPT-4 network do:

<https://openaipublic.blob.core.windows.net/neuron-explainer/paper/index.html#sec-intro>



Calculate the output for one neuron

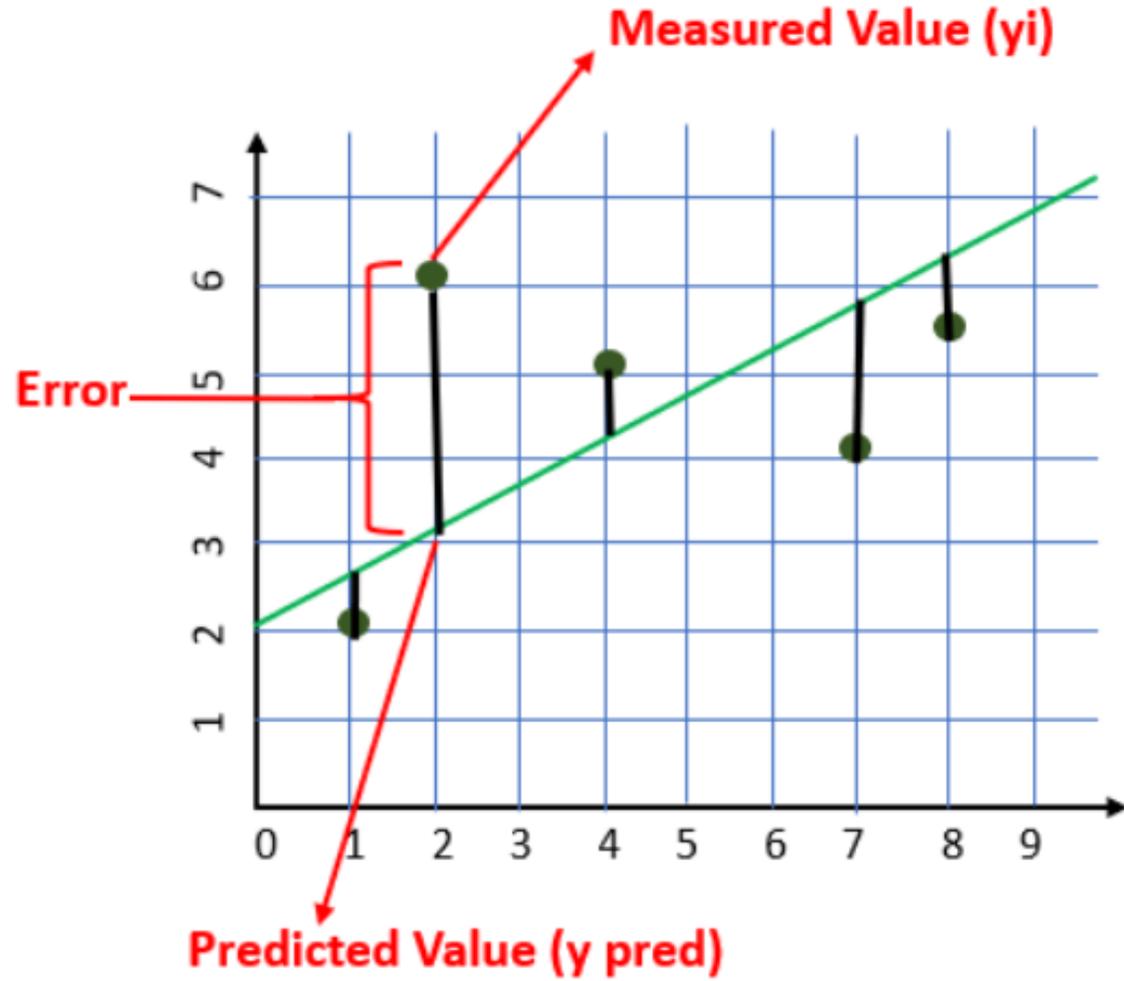
Suppose we have

- Input: $X = (0, 0.5, 1)$
- Weights: $W = (-1, -0.5, 0.5)$
- Bias: $b = 1$
- Activation function *relu*: $f(x) = \max(x, 0)$

What is the output of the neuron?

Note: You can use whatever you like: brain only, pen&paper, Python, Excel...



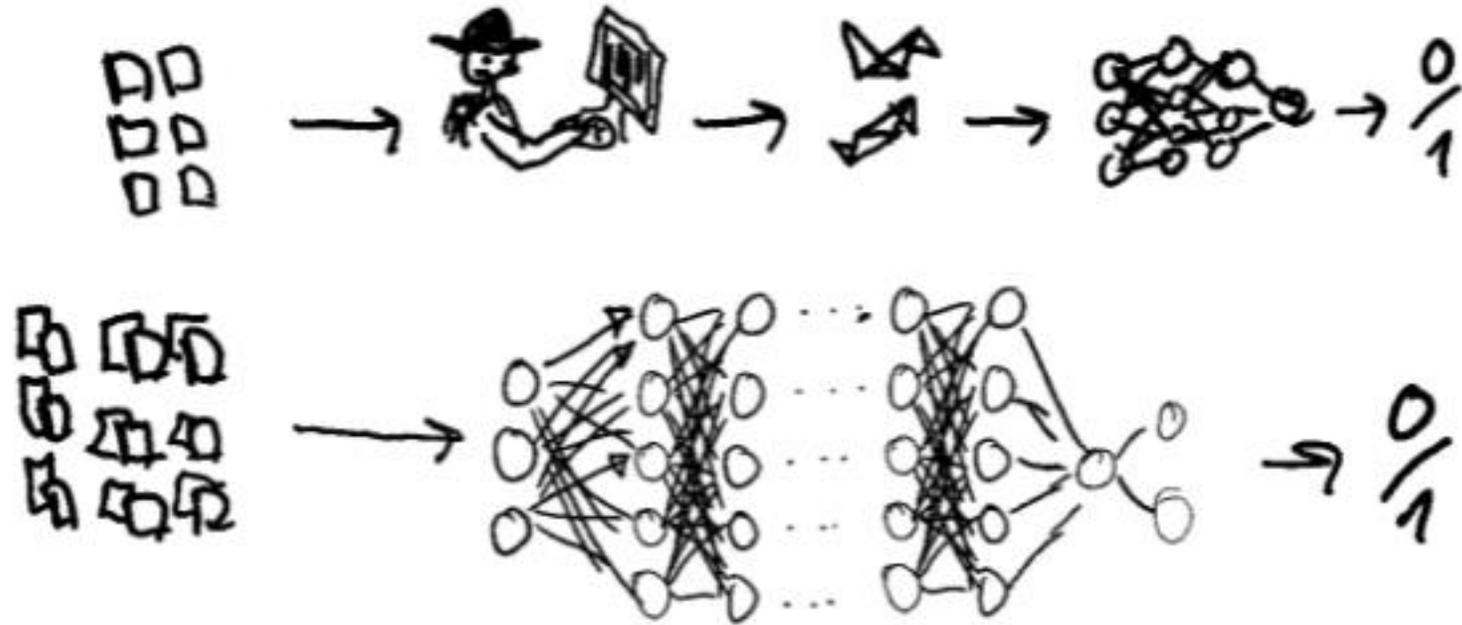


Why deep neural networks? Intuition.

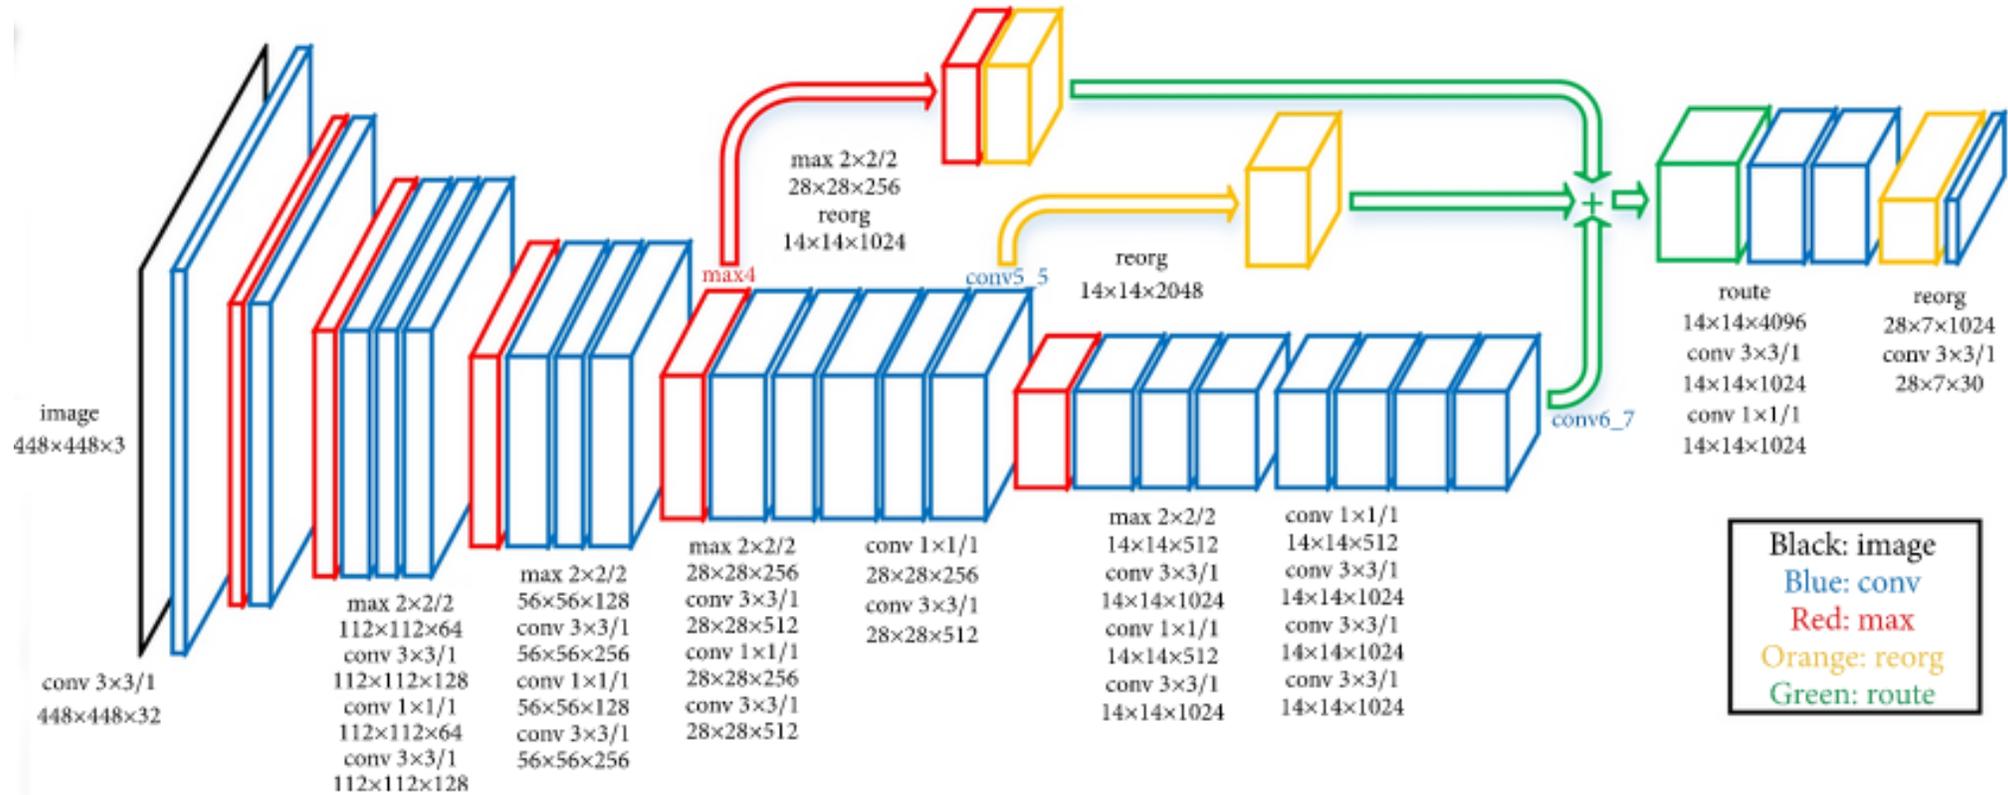
- No hidden layer, simple dataset
- No hidden layer, complex dataset
- Hidden layer, complex dataset



Deep vs shallow neural networks



Example deep network



What sort of problems can deep learning solve?

- Pattern/object recognition



DOG



CAT

What sort of problems can deep learning solve?

- Segmenting images (or any data)



What sort of problems can deep learning solve?

- Translating between one set of data and another, for example natural language translation.



What sort of problems can deep learning solve?

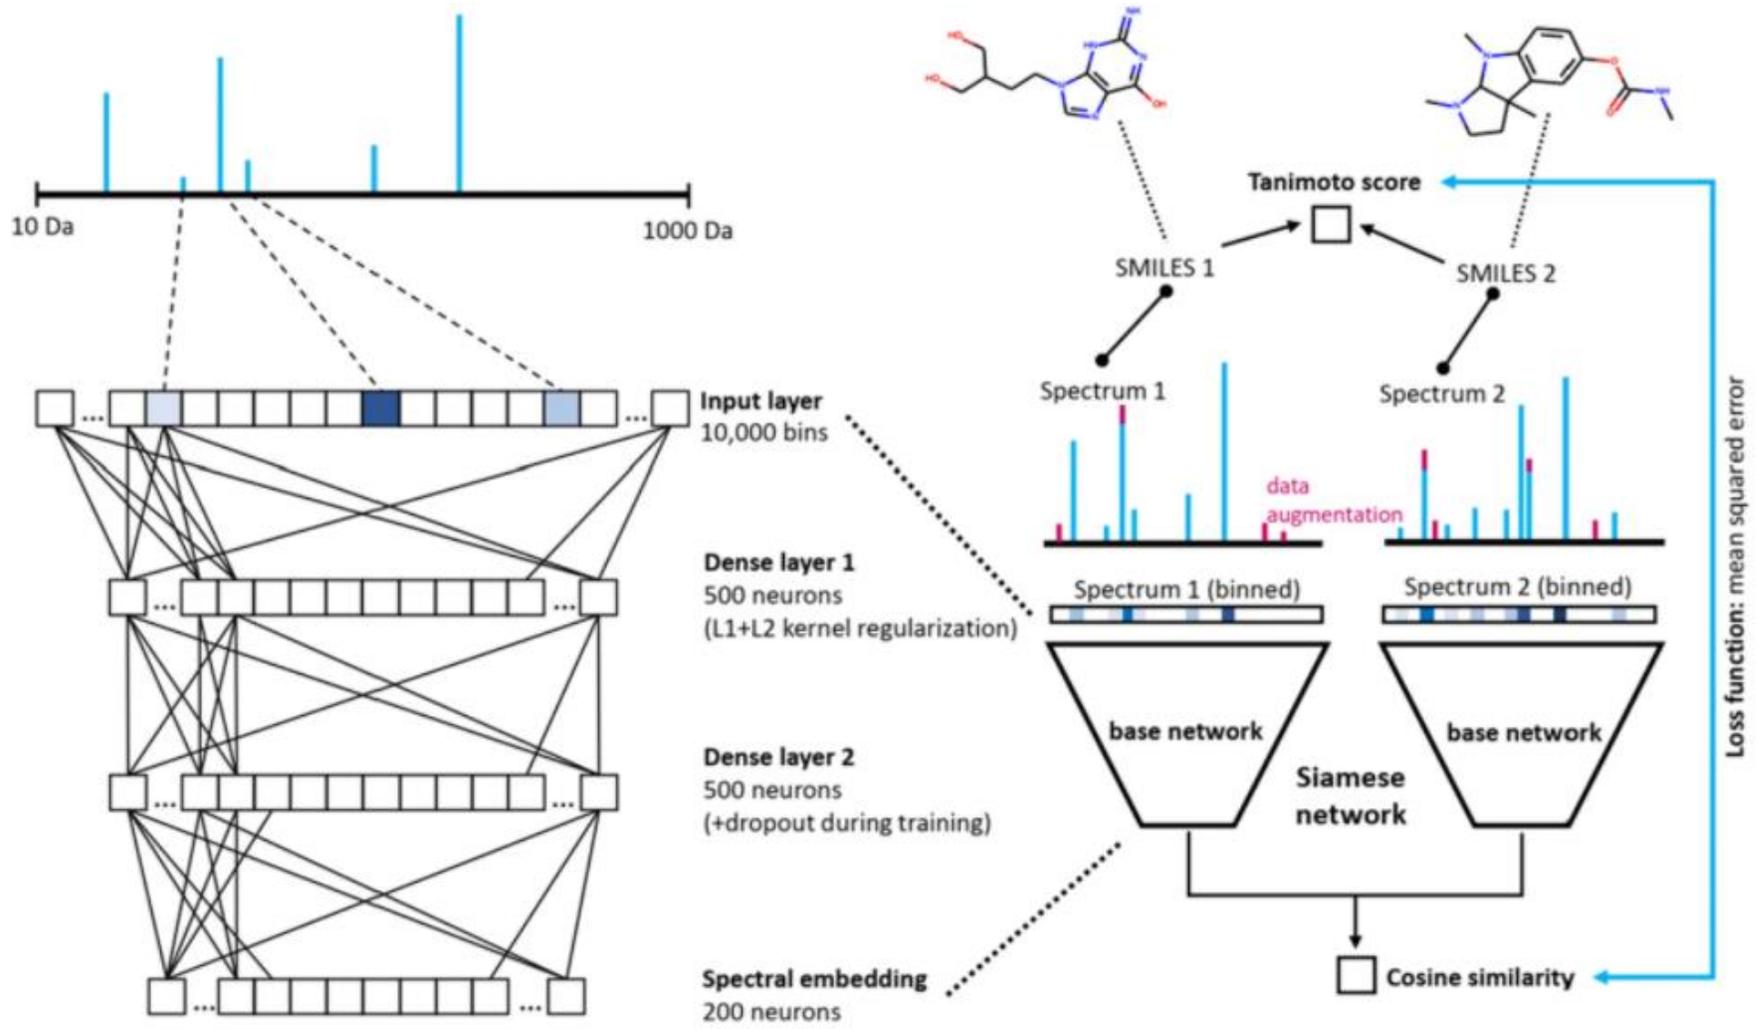
- Generating new data that looks similar to the training data, often used to create synthetic datasets, art or even “deepfake” videos.



What are examples that you have of deep learning applied to research?



Example of Deep Learning applied to Research: MS2DeepScore: a novel deep learning similarity measure to compare tandem mass spectra



What sort of problems deep learning *cannot* solve?

- Where only small amounts of training data is available.
- Tasks requiring an explanation of how the answer was arrived at.
- Being asked to classify things which are nothing like their training data.

Deep Learning Problems Exercise

Which of the following would you apply Deep Learning to?

1. Recognising whether or not a picture contains a bird.
2. Calculating the median and interquartile range of a dataset.
3. Identifying MRI images of a rare disease when only one or two example images available for training.
4. Identifying people in pictures after being trained only on cats and dogs.
5. Translating English into French.



Deep learning workflow

1. Formulate/outline the problem
2. Identify inputs and outputs
3. Prepare data
4. Choose a pre-trained model or build a new architecture from scratch
5. Choose a loss function and optimizer
6. Train the model
7. Perform a Prediction
8. Measure Performance
9. Tune hyperparameters
10. Share model



Deep learning workflow exercise

Think about a problem you'd like to use Deep Learning to solve. As a group, pick 1 problem.

1. What do you want a Deep Learning system to be able to tell you?
2. What data inputs and outputs will you have?
3. Do you think you'll need to train the network or will a pre-trained network be suitable?
4. What data do you have to train with? What preparation will your data need? Consider both the data you are going to predict/classify from and the data you'll use to train the network.

Discuss your answers in your breakout group. Write them (concisely) down in the collaborative document.



Deep Learning Frameworks

- Tensorflow
 - Google
 - General Tensor Network Calculations
- PyTorch
 - Facebook
 - Very Fast
- Keras
 - Designed for Ease of Use
 - On top of Tensorflow
 - Good Documentation
- And many others (Theano, Caffe, etc.)

What we are using

- Keras – Deep Learning Framework
- Sklearn – Machine Learning Framework
- Seaborn – Visualization Library

What have we learned?

- What deep learning is
- How neural networks work
- When to apply deep learning
- Applying deep learning with the ‘deep learning workflow’



You are pioneers



Check versions

```
from tensorflow import keras  
print(keras.__version__)
```

```
import seaborn as sns  
print(sns.__version__)
```

```
import sklearn  
print(sklearn.__version__)
```



Next steps

- Practice!
 - For example try a Kaggle competition: <https://www.kaggle.com/competitions>
- Read/watch blogs and videos
 - For example: [2-minute papers](#)
- Fast.ai course: making neural nets uncool again
 - <https://course.fast.ai/>
- Try a quick and dirty approach first, then optimize
- Get access to a GPU
 - If experiments take seconds instead of 20 minutes, you can iterate over hypothesis -> experiment -> conclusion much much quicker

