WHAT IS MACHINE LEARNING GOOD AT?

GLAM DATA TYPES FOR MACHINE LEARNING

- Tabular Data
- Images (Computer Vision)
- Text (Natural Language Processing)

"CLASSICAL MACHINE LEARNING"

Machine learning on tabular data is sometimes know as "classical machine learning", and there are decades of algorithms that all have strengths and weaknesses.

Choosing the correct algorithm requires a lot of testing and validation, and they can often be bundled together as *ensemble* models.

DEEP LEARNING

Deep learning refers to a machine learning model where several layers of neural networks are connected in order to make a prediction.

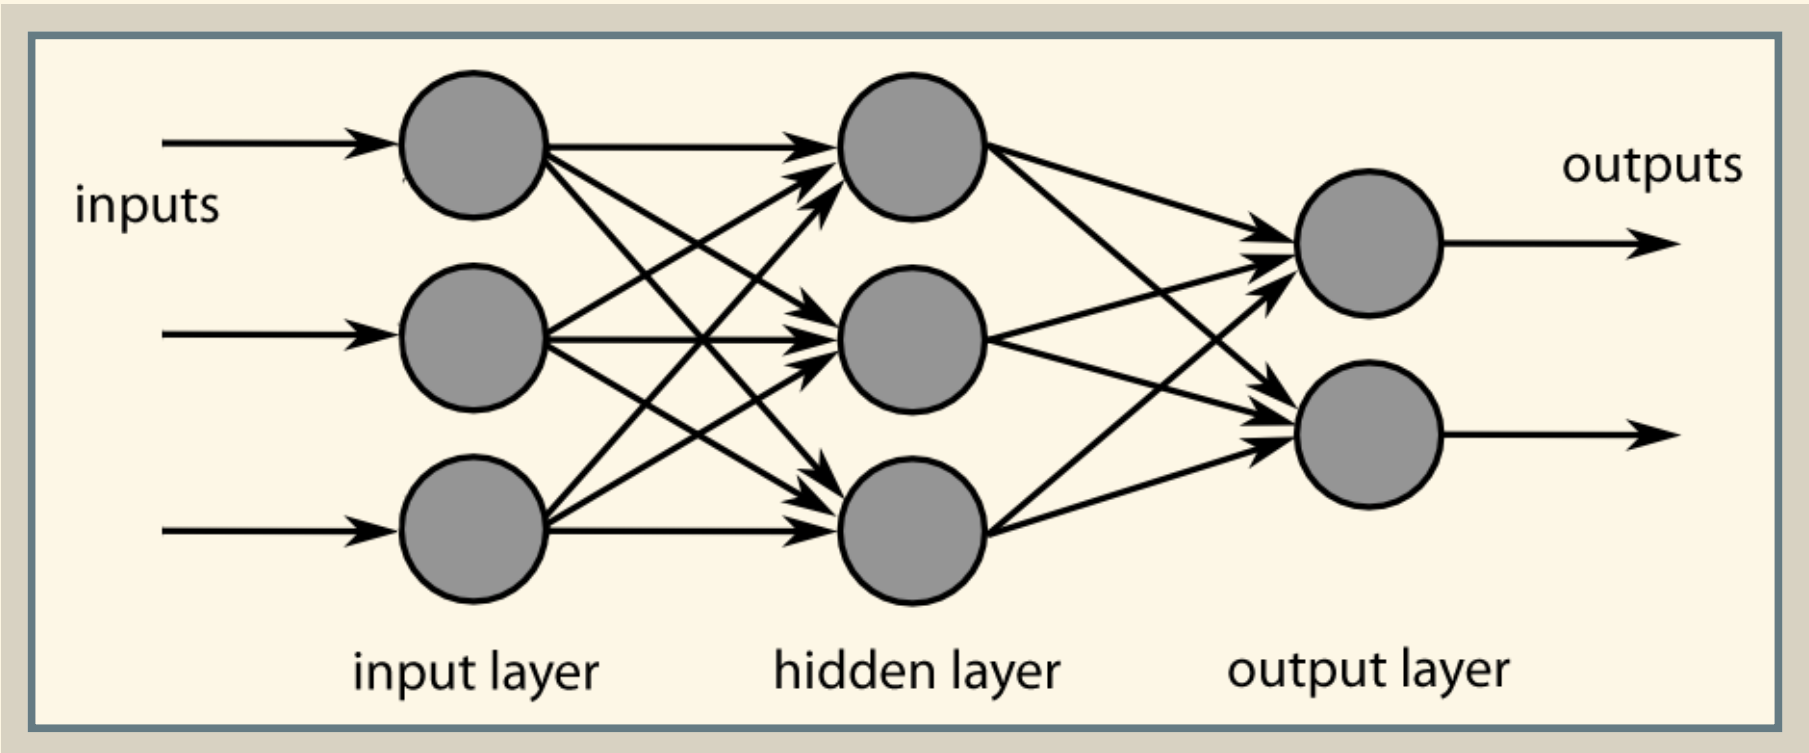

COMPUTER VISION

Machine learning on images is often refered to as **Computer Vision**.

Image data can almost be thought of as a form of tabular data (rows and columns of numerical pixel values that represent pixels).

DEEP LEARNING CNN

Most CV Deep Learning models use a model architecture called Convolutional Neural Networks (CNN). Convolutions use matrix math to pull features from neighboring pixels (such as edges).

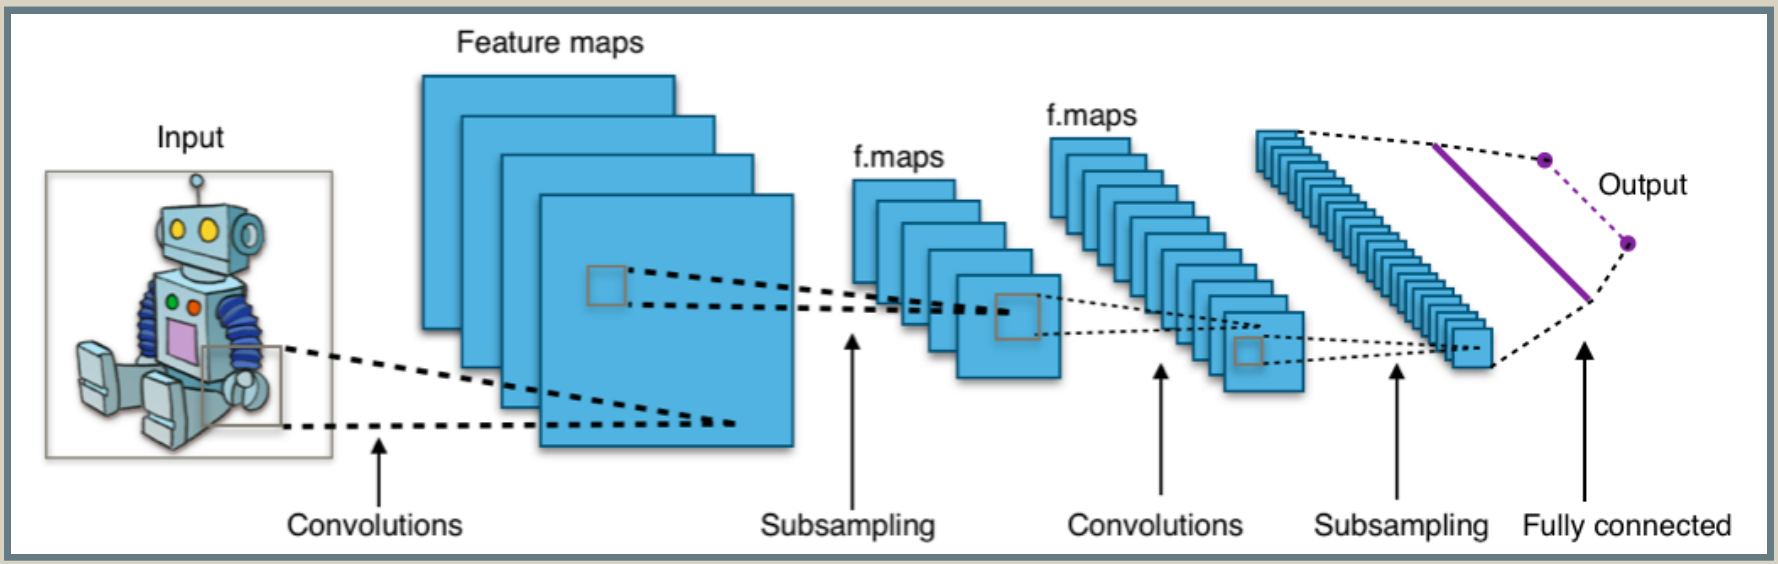

COMPUTER VISION TASKS

IMAGE CLASSIFICATION

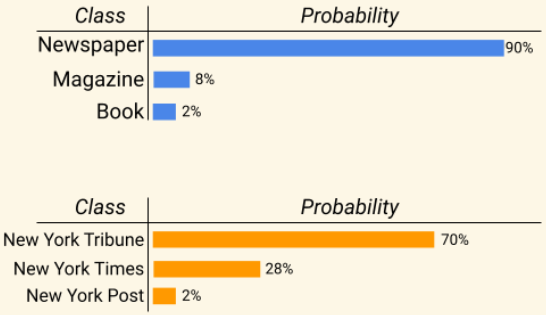

OBJECT DETECTION

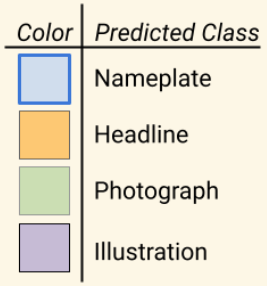

ACTIVITY

Now that you have learned the types of computer vision tasks where machine learning excels, what are some things you might try to do with this image?

Imagine you had overhead views of each tray of bird specimens. How might object detection help speed up digitisation?

NATURAL LANGUAGE PROCESSING

Machine learning on text data is often referred to as **Natural Language Processing (or NLP)**.

NLP TRADITIONAL

Remember, machine learning algorithms require numerical inputs, so how do we get numerical data from text?

Bag of Words

NLP DEEP LEARNING

Recurrent Neural Networks and Transformer models produce *word embeddings* for words based on their context.

NLP TASKS

- Sentiment Analysis
- Named Entity Recognition (NER)
- Language Translation
- Information Summarization

ACTIVITY

Which of these typical Natural Language Processing tasks do you think requires a complex language model, and which could be achieved with a simpler one?

- Sentiment analysis: Is a sentence positive or negative?
- Machine Translation: converting text in one language, to another
- Spam detection
- Entity extraction: identifying people, places, organisations in text
- Text generation: generating text, maybe a poem, in the style of an author
- Question answering: interpreting a written question and retrieving the answer from a database

OCR = CV + NLP

Computer Vision models to **classify** each character, and then NLP models to make sure the sequences of letters and words actually make sense.

GENERATIVE MODELS

Both CV and NLP models can be used to generate novel images and text.

UNSUPERVISED LEARNING WITH FEATURE VECTORS

You can take feature vectors generated by pre-built models, and calculate distances between them, and produce clusterings.