

Machine Learning & Artificial Intelligence for Physics

Part 1: Principles & Key Tools

Julia Gonski

17 July 2024

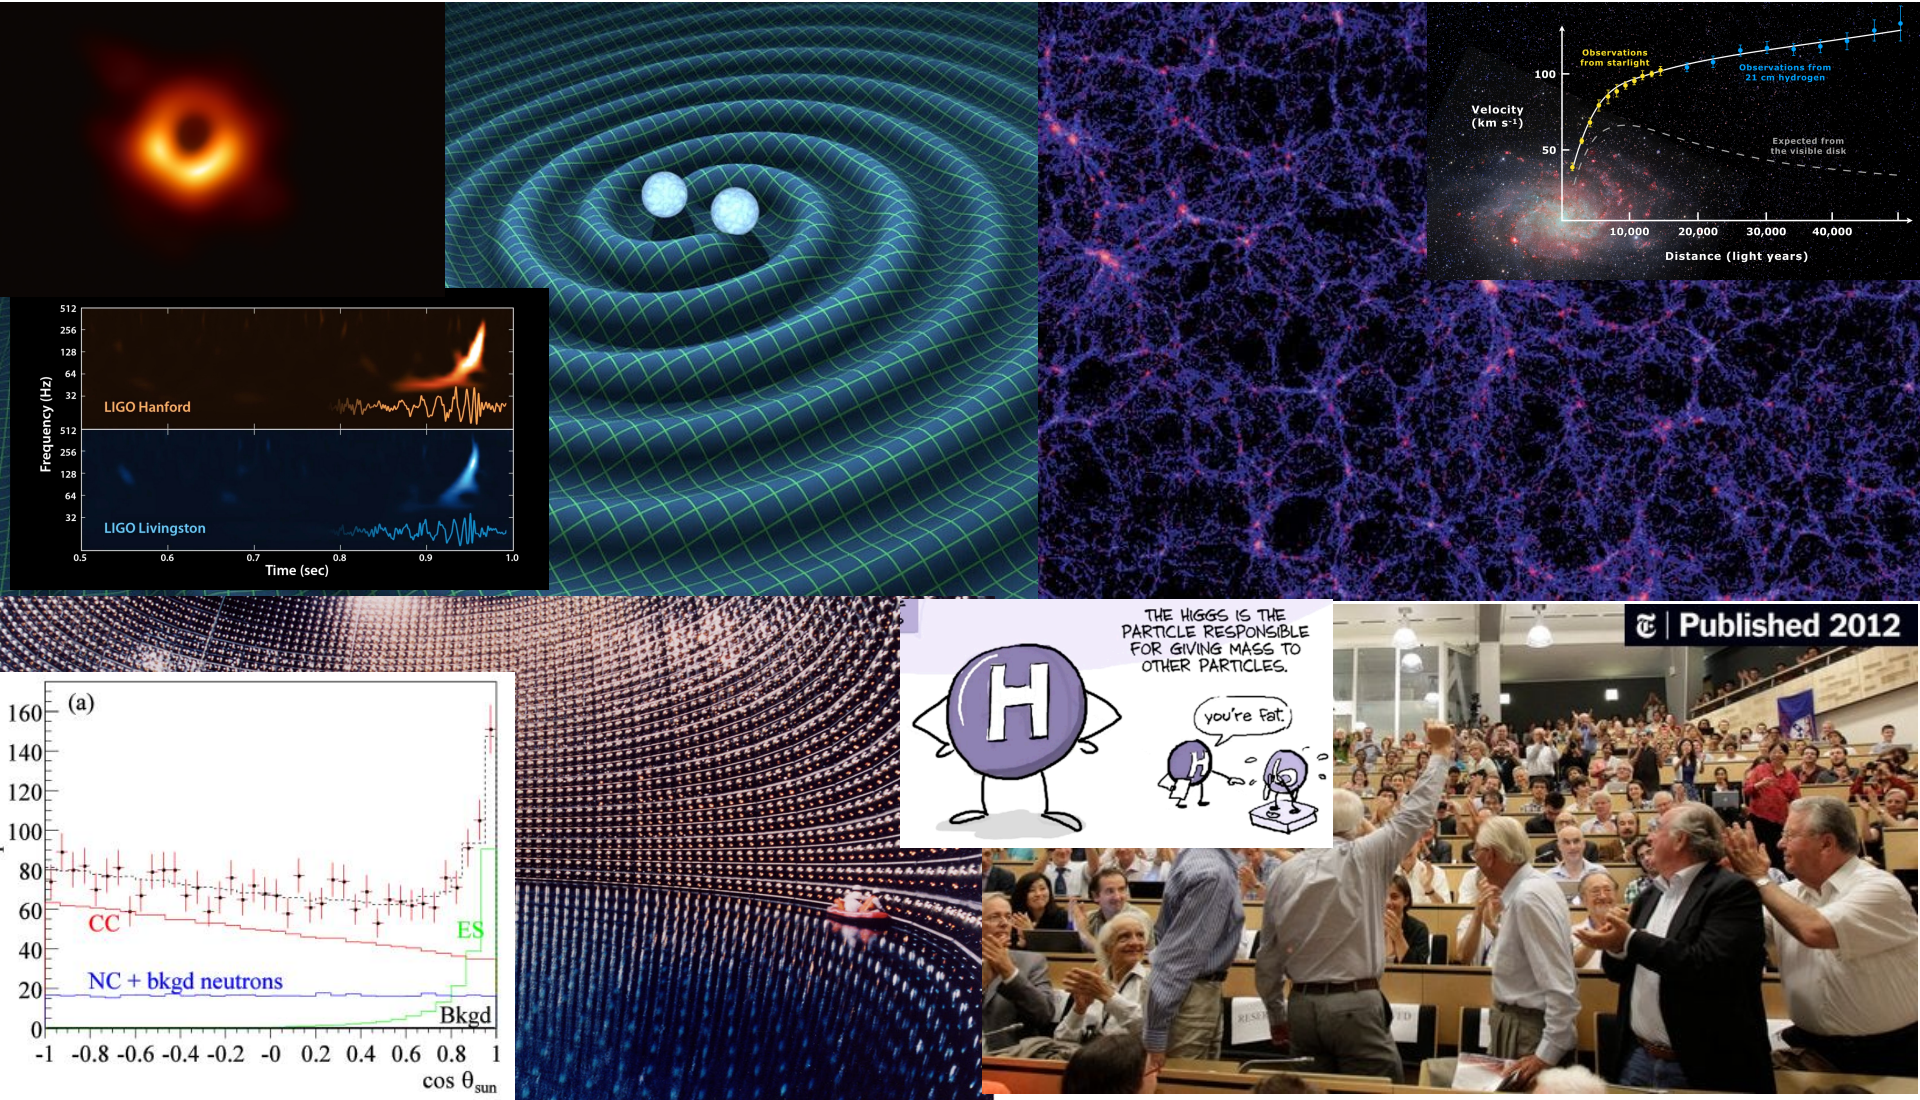
African School of Physics
Marrakech Morocco





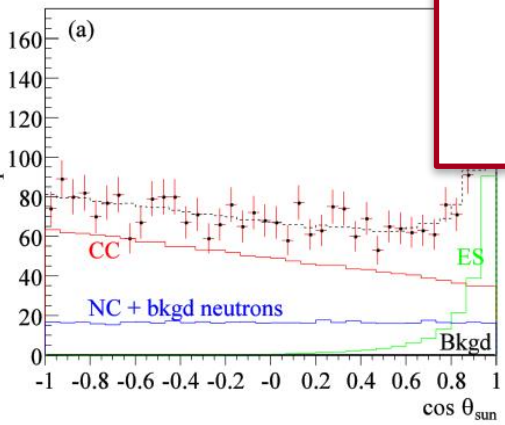
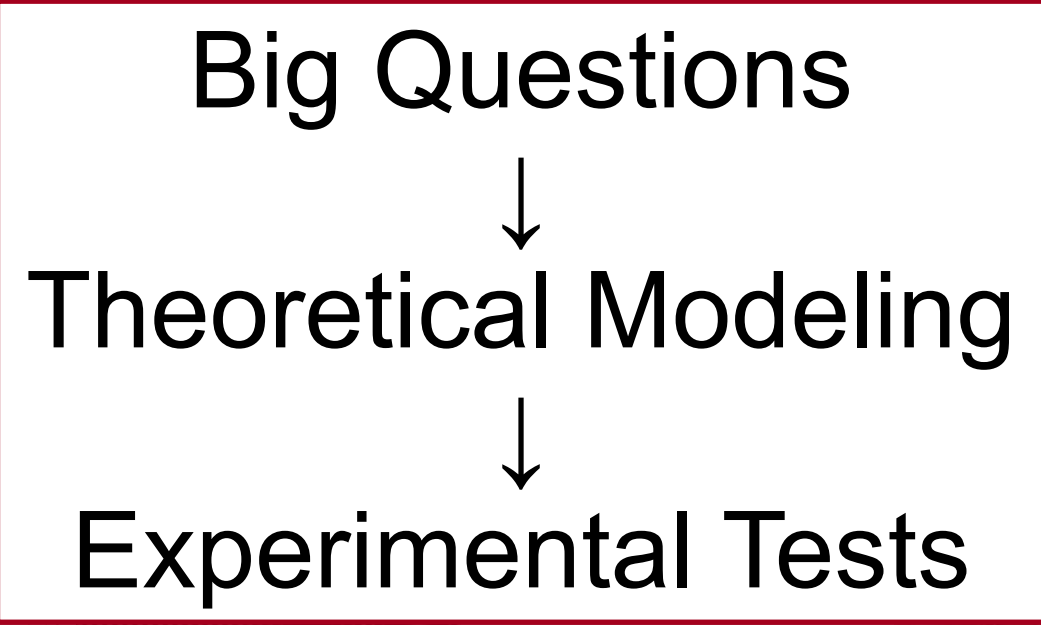



Outline

- Lecture 1: principles & key tools
 - What is AI/ML? How is it useful in physics research?
 - Basics of neural nets: architecture & development
- Lecture 2: applications & advanced models
 - Anomaly detection
 - Geometrical ML
 - Hardware acceleration

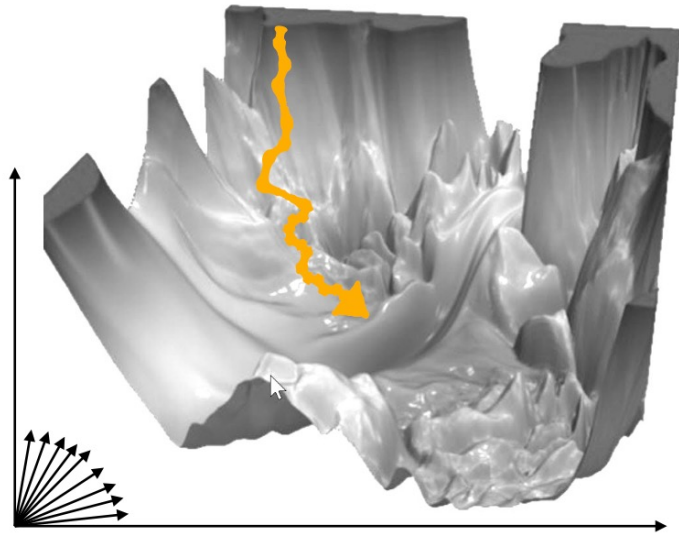
The Why & How of Physics Research



The Why & How of Physics Research



The AI Revolution



Billion parameter gradient descent



Unreal Engine Kite Demo (Epic Games 2015)

Prompt: a landscape from the Moon with the Earth setting on the horizon, realistic, detailed
Negative prompt: whimsical interpretation of the prompt
Parameters: Steps: 30, Sampler: Euler a, CFG scale: 7.0, Seed: 4252913504, Size: 512x384, Model hash: 82aac931



The “Machine” in Machine Learning

Classic Computing

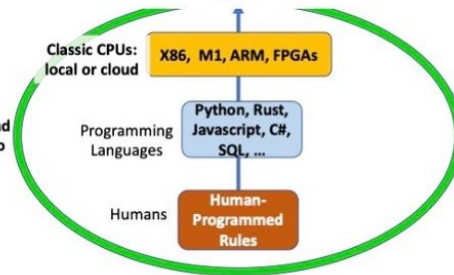
Running (Executing) Programs

Programs take input from humans, sensors or other programs. They run on the same type of classic CPUs as they were developed on.



Programming

Predefined Rules, Logic and Knowledge are coded into programs by humans




The “Machine” in Machine Learning

Classic Computing

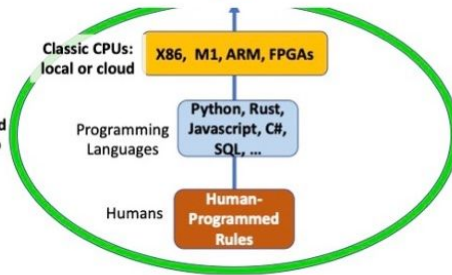
Running (Executing) Programs

Programs take input from humans, sensors or other programs. They run on the same type of classic CPUs as they were developed on.



Programming


Predefined Rules, Logic and Knowledge are coded into programs by humans



Machine Learning

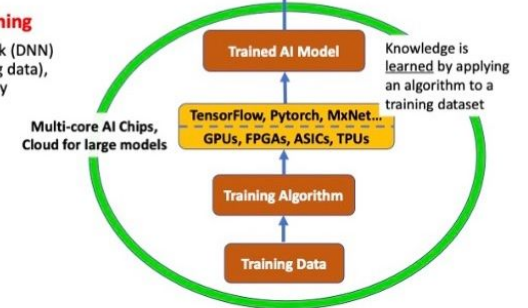
Inference

Once the system has been trained it can be deployed/run on lower performance chips then needed for training



Machine Learning - Training

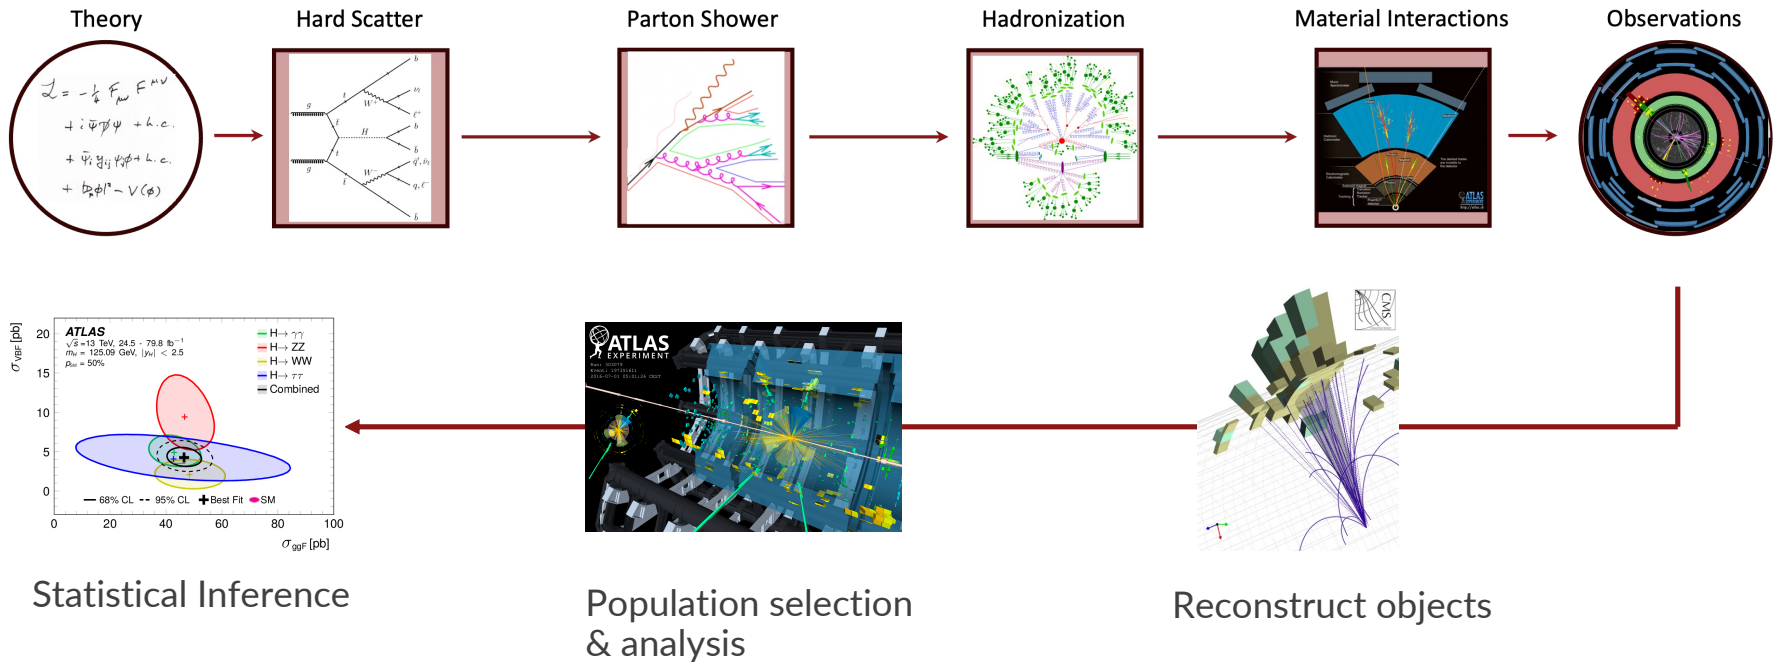
“Teaches” a deep neural network (DNN) to learn from examples (training data), rather than being explicitly programmed.



Fundamental Physics ↔ AI/ML

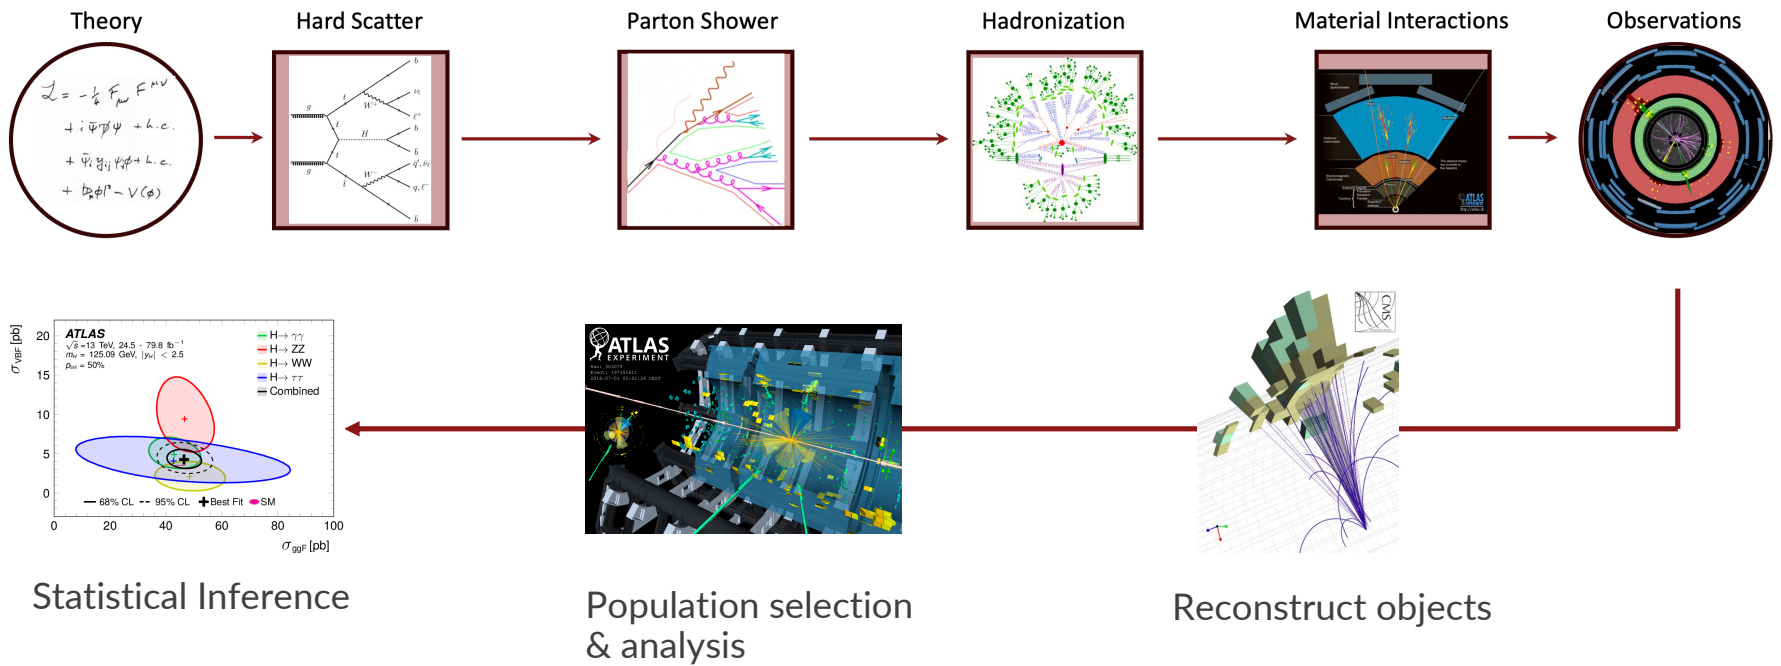
O(10) Physics parameters

O(10⁹) sensor elements



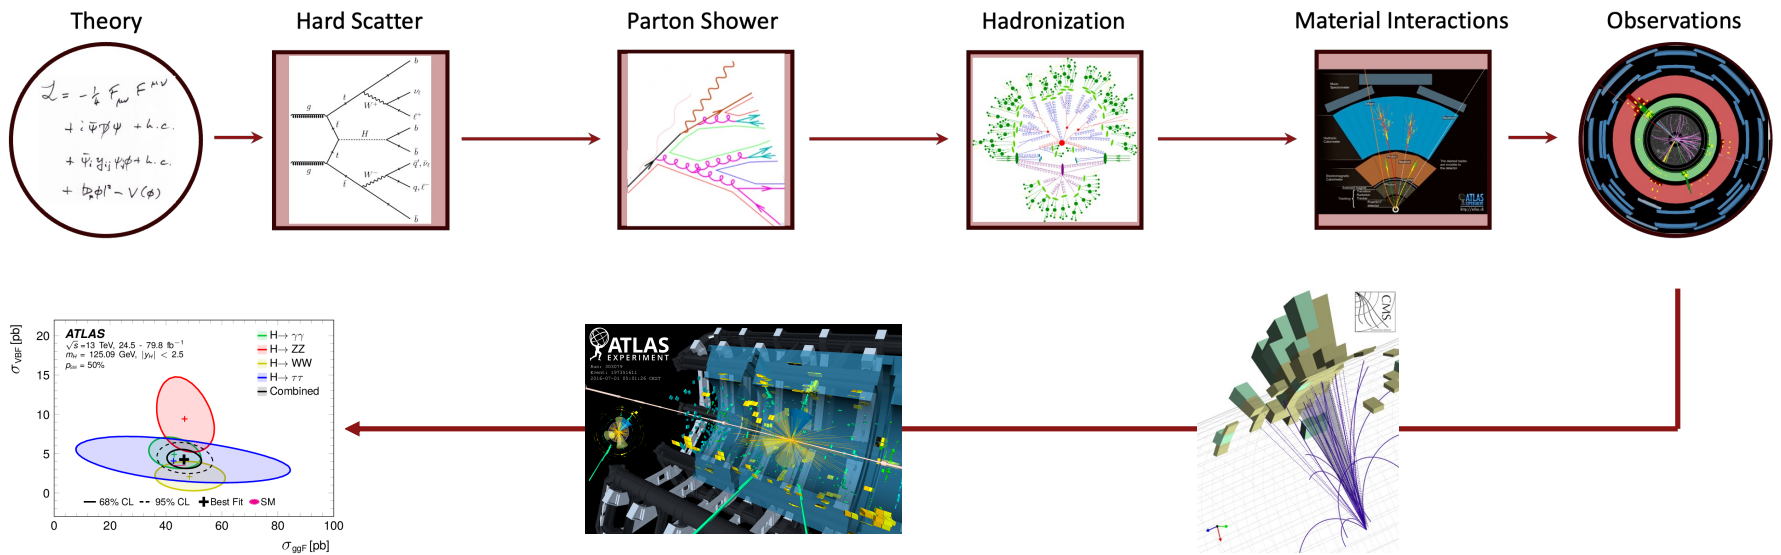
Fundamental Physics ↔ AI/ML

Generate plausible, high-dim. data from high-level concepts



Fundamental Physics ↔ AI/ML

Generate plausible, high-dim. data from high-level concepts



Statistical Inference

Population selection & analysis

Reconstruct objects

Extract high-level concepts from low-level, high-dim. data

Fundamental Physics ↔ AI/ML

G Generate plausible, high-dim data from high-level concepts

Prompt:
street style photo of a woman selling pho at a Vietnamese street market, sunset, shot on fujifilm





Extract high-level concepts from low-level, high-dim. data

Classification:
A photo of guacamole, a type of food



Extract high-level concepts from low-level, high-dim. data




Stati




ts



rations



Frontiers of Physics



The Energy Frontier

Origins of Mass

Matter/Anti-matter
Asymmetry

Dark matter

Origin of Universe

Unification of Forces

New Physics
Beyond the Standard Model

Dark energy


Neutrino Physics

Proton Decay



Cosmic Particles

The Intensity Frontier



The Cosmic Frontier



Atomic/Molecular/
Optical





Biophysics





+ theory, computing,
advanced
technology, & more!

How it All Began...



ML Building Blocks: Linear Regression

- Fitting a straight line to a linear function
- **Cost function** = returns a global error between the predicted values from a mapping function $h(x)$ (predictions) and all the target values (observations) of the training data set



$$h_{\theta}(x) = \theta_0 + \theta_1 x$$


1D \Rightarrow 2 parameters:

$$\theta = \begin{pmatrix} \theta_0 \\ \theta_1 \end{pmatrix}$$


Cost Functions

- Ex. linear regression: cost function = mean-squared error

$$J(\theta_0, \theta_1) = \frac{1}{2m} \sum_{i=1}^m \left(h_{\theta}(x^{(i)}) - y^{(i)} \right)^2$$




“Fit the data”



Find θ parameters
minimizing the cost J


$$\min_{\theta_0, \theta_1} J(\theta_0, \theta_1)$$



Cost Functions

- Ex. linear regression: cost function = mean-squared error

$$J(\theta_0, \theta_1) = \frac{1}{2m} \sum_{i=1}^m \left(h_{\theta}(x^{(i)}) - y^{(i)} \right)^2$$




“Fit the data”



Find θ parameters
minimizing the cost J

$$\min_{\theta_0, \theta_1} J(\theta_0, \theta_1)$$

→ 1-dimensional **gradient descent!**



Gradient Descent

- **Gradient descent:** iterative procedure to update the parameters in the direction of ‘descending gradient’, i.e. dictated by the slope of the cost function’s partial derivatives for each parameter
 - Until: N iterations or partial derivatives $\rightarrow 0$
 - Convergence to values for the parameters minimizing the cost function

n input features \rightarrow

m data samples \downarrow

	x_1	x_2	x_3	x_4
0	53.844621	-2.236084	2158.276717	-0.035128
1	45.108137	1.204087	1935.968922	0.038599
2	43.764648	-1.509527	2080.412520	-0.080645
3	44.739359	-2.604297	2048.357208	-0.040369
4	38.390222	-2.325613	1984.086668	0.243353
5	49.152215	-0.600885	1972.311885	-0.034590

(θ_0, θ_1)

$\frac{\partial}{\partial \theta_0} J(\theta_0, \theta_1)$

$\frac{\partial}{\partial \theta_1} J(\theta_0, \theta_1)$

$\theta'_0 = \theta_0 - \alpha \frac{\partial}{\partial \theta_0} J(\theta_0, \theta_1)$

$\theta'_1 = \theta_1 - \alpha \frac{\partial}{\partial \theta_1} J(\theta_0, \theta_1)$

Hyperparameter α = learning rate

Shallow Learning: The Neural Net


x_j : input nodes

w_j : weights

b : bias term

Σ : weighted sum

f : activation function



Shallow Learning: The Neural Net

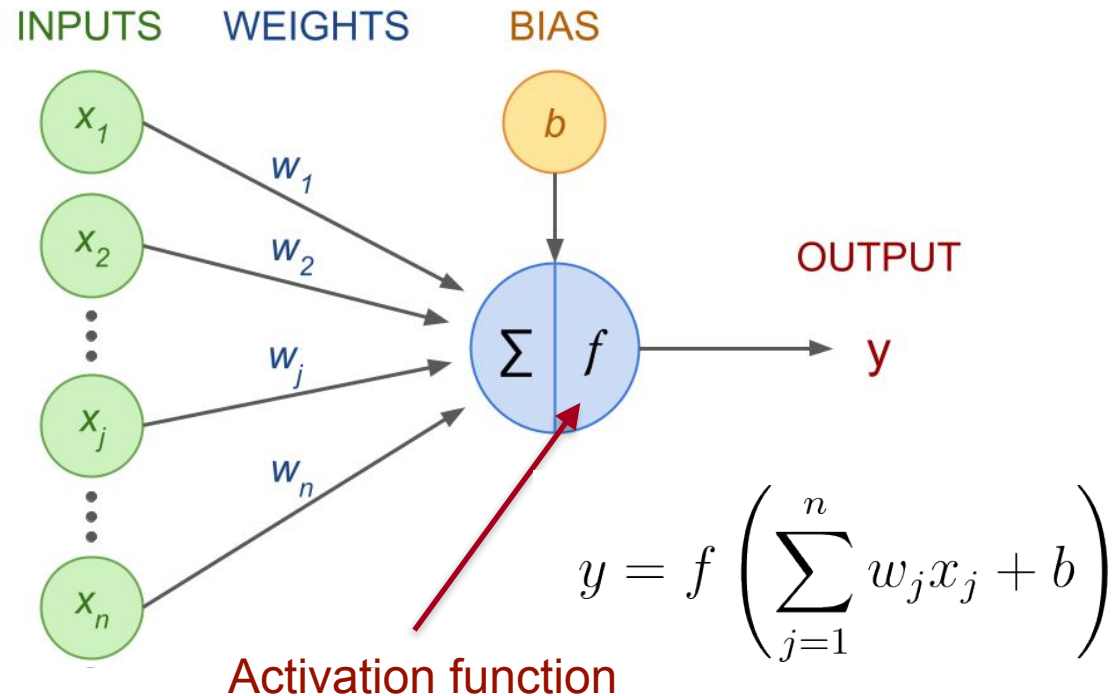
x_j : input nodes

w_j : weights

b : bias term

Σ : weighted sum

f : activation function



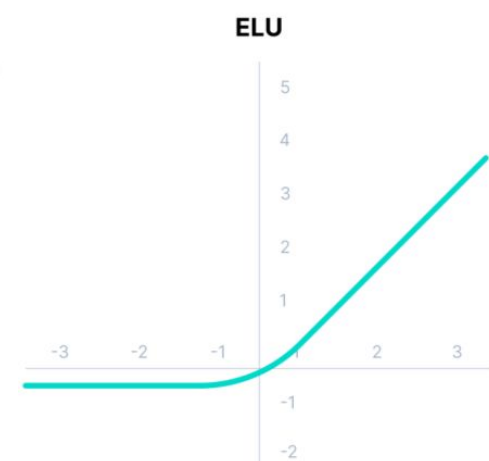


The Activation Function

The **Activation Function** is a mathematical operation deciding whether the neuron's input to the network is important or not.

Returning a non-zero values means the neuron is “activated”, or “fired.”

The purpose of the activation function is to **introduce non-linearity** into the output of a neural network.

- **Differentiable**, ideally continuously differentiable
- **Non-linear**: the only way for the network to learn



Shallow Learning: The Neural Net


x_j : input nodes

w_j : weights

b : bias term

Σ : weighted sum

f : activation function



- **Feed-forward propagation:** process of computing all activation units of a NN
 - At last layer (output), leads to predictions (which are compared to observations in training)

Shallow Learning: The Neural Net


x_j : input nodes

w_j : weights

b : bias term

Σ : weighted sum

f : activation function



- **Backpropagation:** uses the chain rule to compute how much each activation unit contributed to the overall error & adjust weights/biases to reduce overall error

Deep Learning: The Power of Scale

- More dataset complexity means you need a larger model, which means you need more input data, which means you need more computational power!






Image credit: [D. McCandless, T. Evans, P. Barton](#)

Putting it in Action


1. Select your training data
2. Choose an input modeling
3. Build & train your model
4. Gauge performance (ROC/AUC)
5. Optimize (hyperparameters)

1. Selecting the Data

Training set: dedicated to the fitting procedure (minimizing the cost function).


Validation set: assess the performance of the model & tune the model's hyperparameters.

Test set: final assessment done on the model → error rate is called **generalization error**.





2. Input Modeling

- What does your data *naturally* look like?
 - An image? → convolutional layers
 - A sequence? → recurrent architecture
 - A graph? → geometrical ML (graphs, transformers)

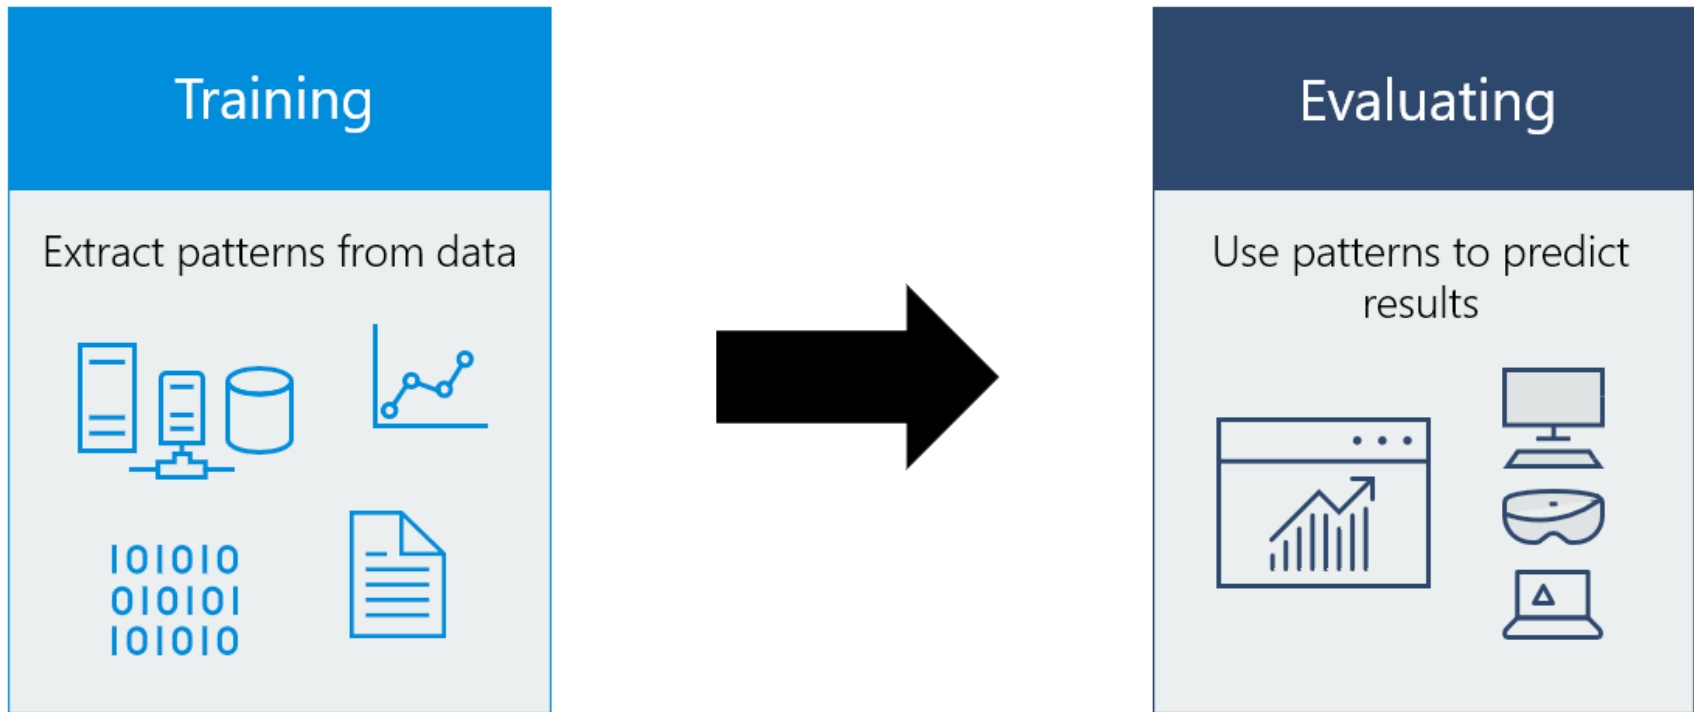


Sequential access

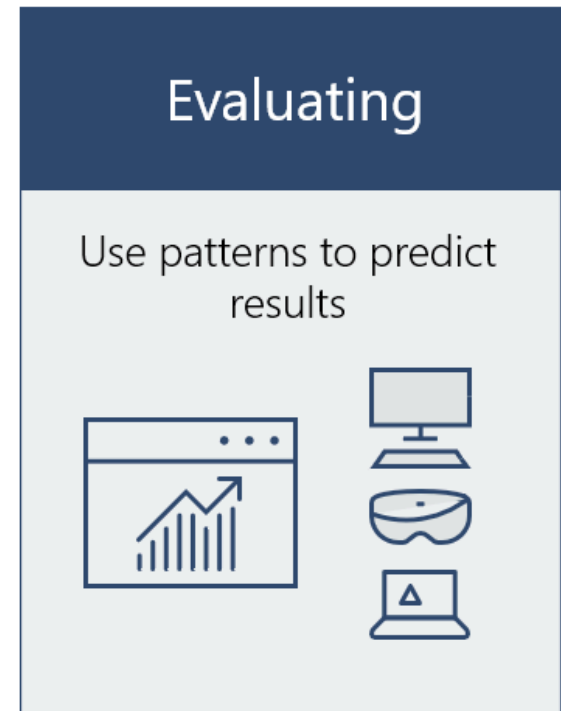



- What information can I give my ML tool to help it best learn from the data?

3. Building, Training, and Evaluating the Model



3. Building, Training, and Evaluating the Model



Epoch: iteration of training calculations

Example:

- **Classification:** predict the category of the input
- **Regression:** predict value of key quantity

3. Building, Training, and Evaluating the Model

- This is all very easy to implement!
- Pre-existing functionality through ML python packages such as keras or PyTorch

```
import tensorflow as tf
mnist = tf.keras.datasets.mnist

(x_train, y_train), (x_test, y_test) = mnist.load_data()
x_train, x_test = x_train / 255.0, x_test / 255.0

model = tf.keras.models.Sequential([
    tf.keras.layers.Flatten(),
    tf.keras.layers.Dense(512, activation=tf.nn.relu),
    tf.keras.layers.Dropout(0.2),
    tf.keras.layers.Dense(10, activation=tf.nn.softmax)
])
model.compile(optimizer='adam',
              loss='sparse_categorical_crossentropy',
              metrics=['accuracy'])

model.fit(x_train, y_train, epochs=5)
model.evaluate(x_test, y_test)
```

Input Data

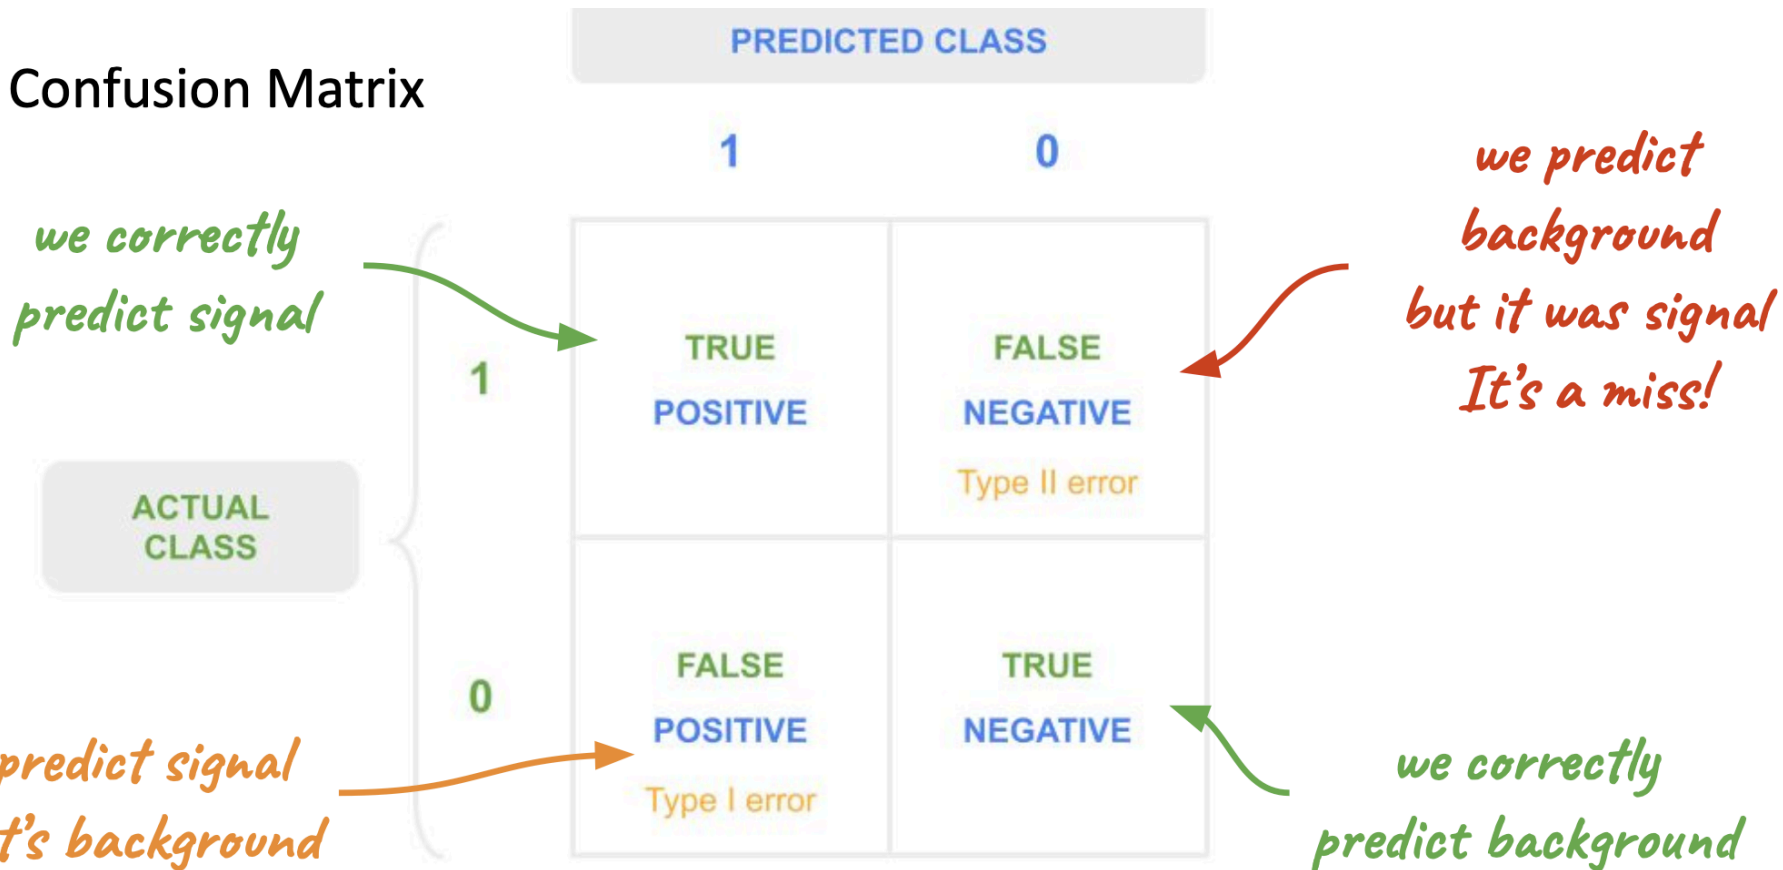
Sequential Model

Train

Evaluate

4. Quantifying Performance

The Confusion Matrix




4. Quantifying Performance

The **Receiver Operating Characteristic (ROC)** curve is a graphical display that plots the True Positive Rate (TPR) against the False Positive Rate (FPR) for each value of the decision threshold going over the classifier's output score range.

Reducing
threshold:

TPR ↗

FPR ↘



Increasing
threshold:

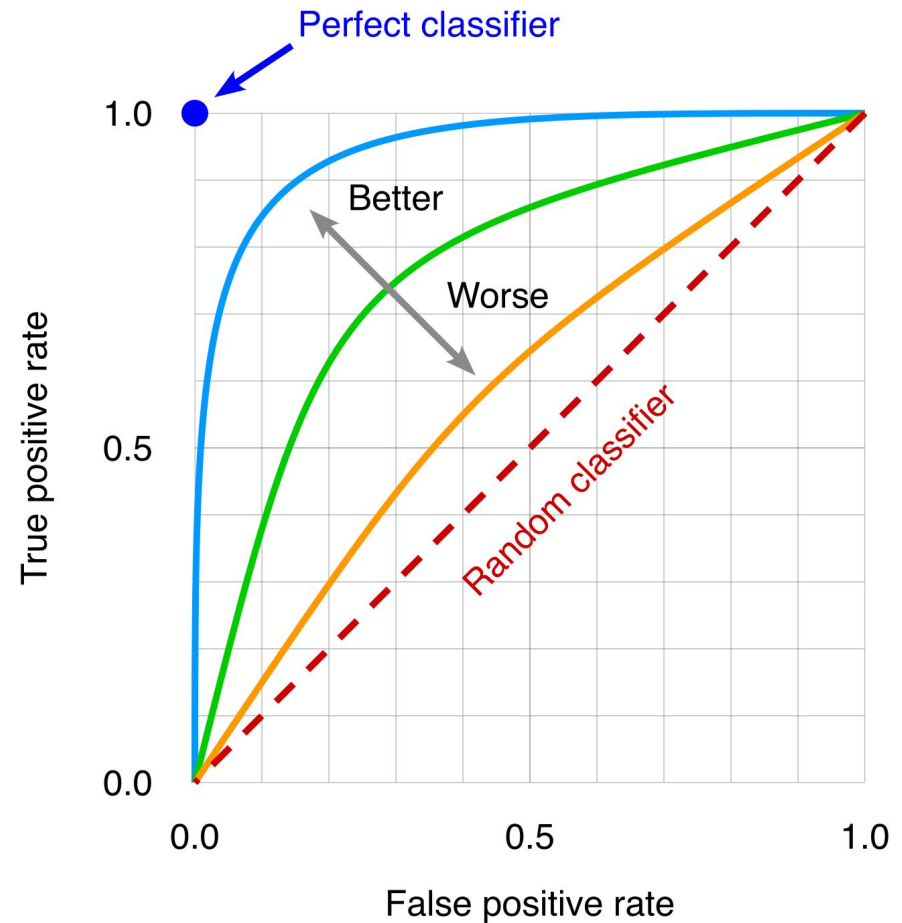
↘ TPR

↗ FPR

Receiver-Operating Characteristic (ROC)


The **Area Under Curve (AUC)** is the integral of the ROC curve, from $FPR = 0 \rightarrow 1$

A perfect classifier has $AUC = 1$.



5. Optimization

- Model parameters = basic structure (number of layers/nodes, weights/biases)
- **Hyperparameters**: choosable parameters of model that can significantly affect performance
 - Must be well-matched to complexity of training dataset and input modeling
 - Batch size (how many events feed into a single update)
 - Train-test split ratio
 - Learning rate: how large is the step in your gradient descent? (below)




Good 😊

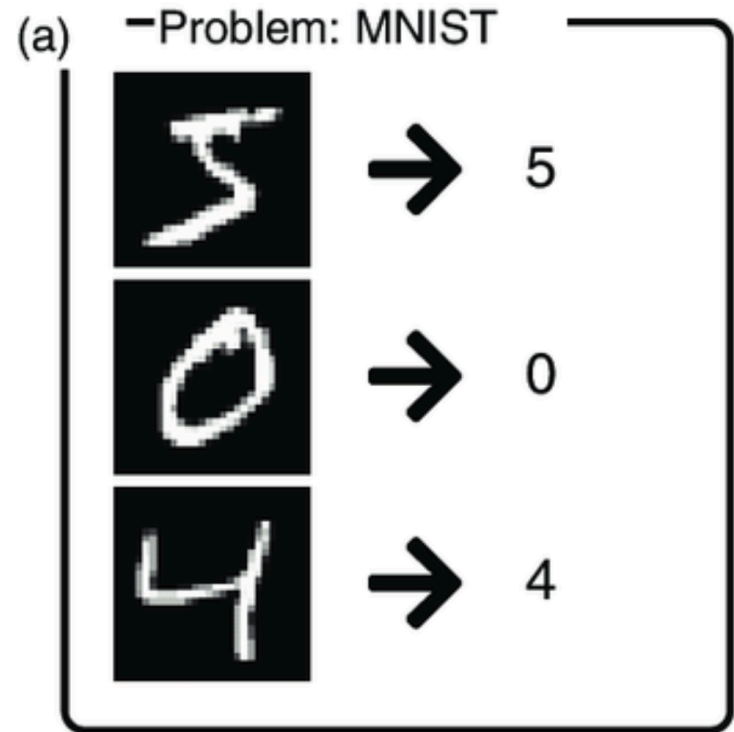
Bad 😐

Ugly 😱

Example: MNIST Numbers

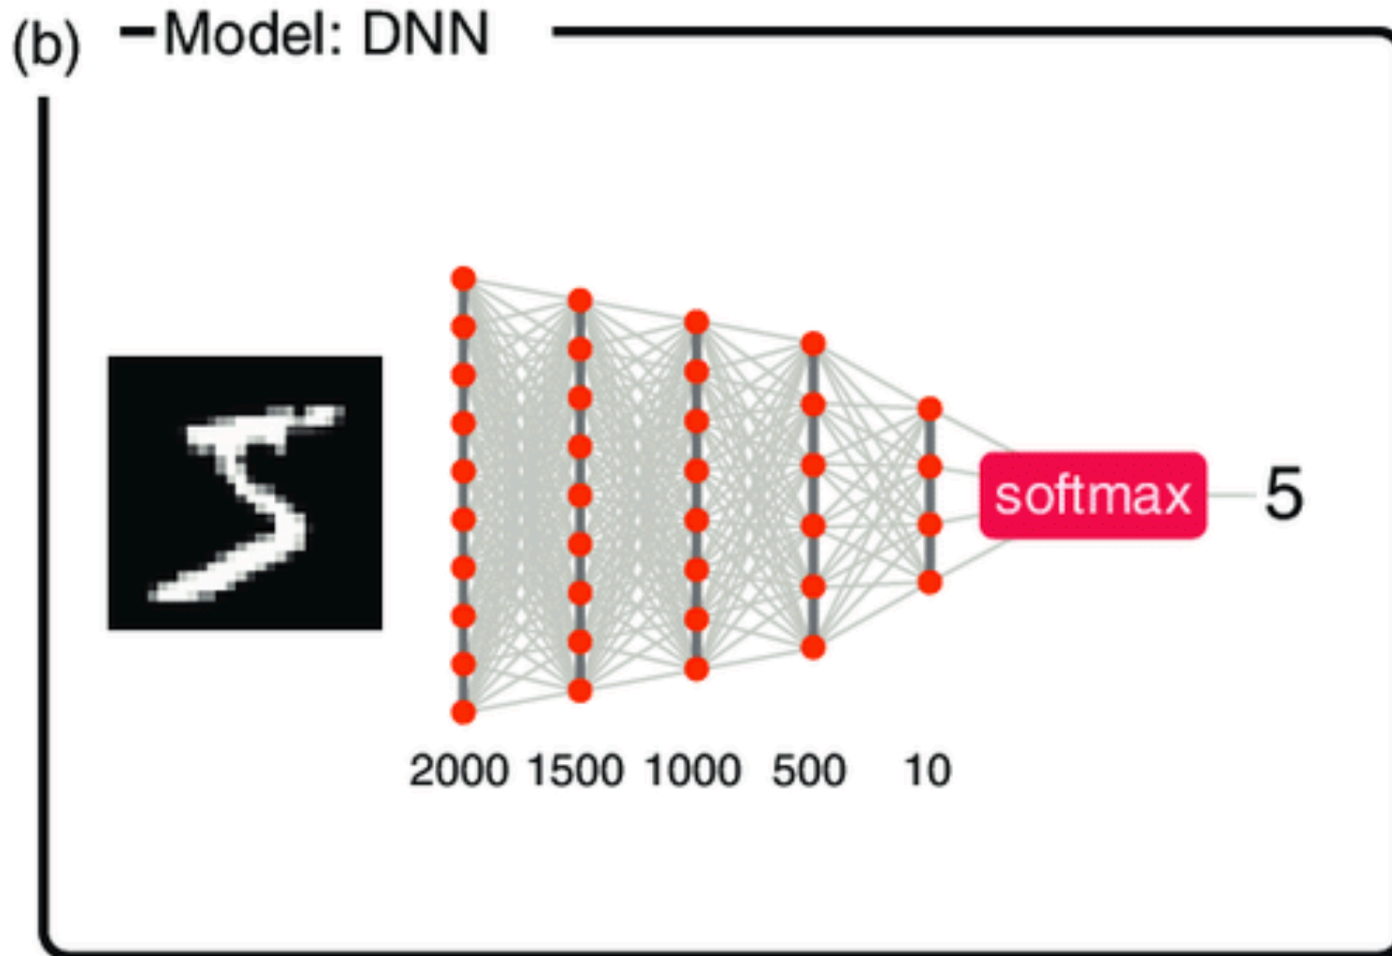


Example: MNIST Numbers



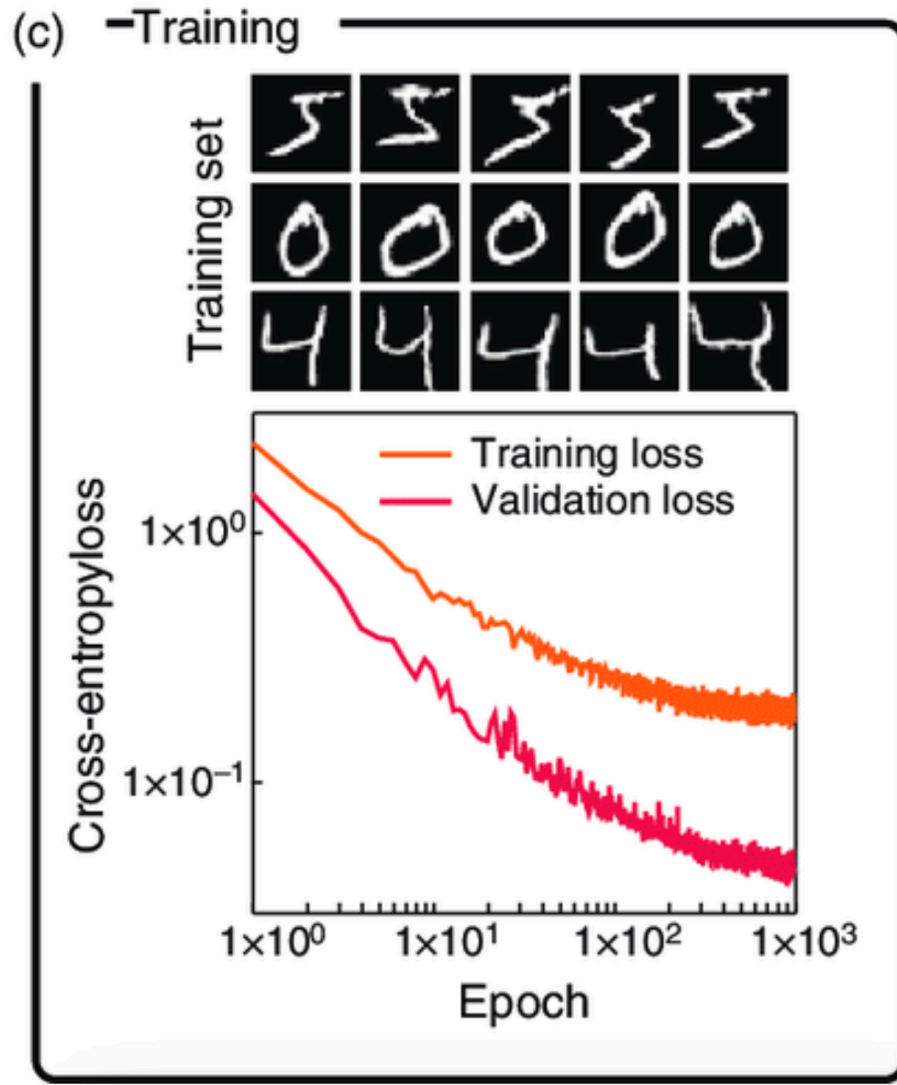
Classification

Example: MNIST Numbers




[Applied Physics Reviews](#)
8:011310 (2011)

Example: MNIST Numbers



Loss should decrease as epochs increase \rightarrow model is learning

Example: MNIST Numbers



Confusion matrix: model is very accurate!

Recap

Recap

- ML at the most basic level is a high-dimensional *non-linear* fitting procedure
 - Described by minimizing **loss function** using **gradient descent**
- Building an ML model requires training/test data, a choice of input modeling, and decisions on model architecture & hyperparameters
- ML performance can be assessed by studying the loss vs. epoch and the ROC/AUC (along with many other diagnostics)
- Developing ML methods on your own is facilitated by many preexisting python packages!

Conclusions

- **Why AI/ML?**

- Exciting and rapidly growing technology that will change our world in ways yet to be seen (and physics research is no different!)

- **What we discussed so far:**

- Motivation for AI/ML in physics
- Basic math of simple neural nets
- Optimizing & qualifying models

- *Coming up next*

- Applications to physics research fields
- Advanced models: anomaly detection, graphs
- Real-time ML and hardware systems

Backup

