2020 embedded VISION summit

An Introduction to Machine Learning:
How to Teach Machines to See

Facundo Parodi
Tryolabs
September 2020

tryo·labs

Who are we?


tryo·labs
tryo-labs
tryo-labs
tryo-labs
tryo-labs
tryo-labs
tryo-labs
IT'S NOT ME,
IT'S YOU

Outline

Machine Learning

- Types of Machine Learning Problems
- Steps to solve a Machine Learning Problem

Deep learning

- Artificial Neural Networks

Image Classification

- Convolutional Neural Networks

What is a cat?

What is a cat?

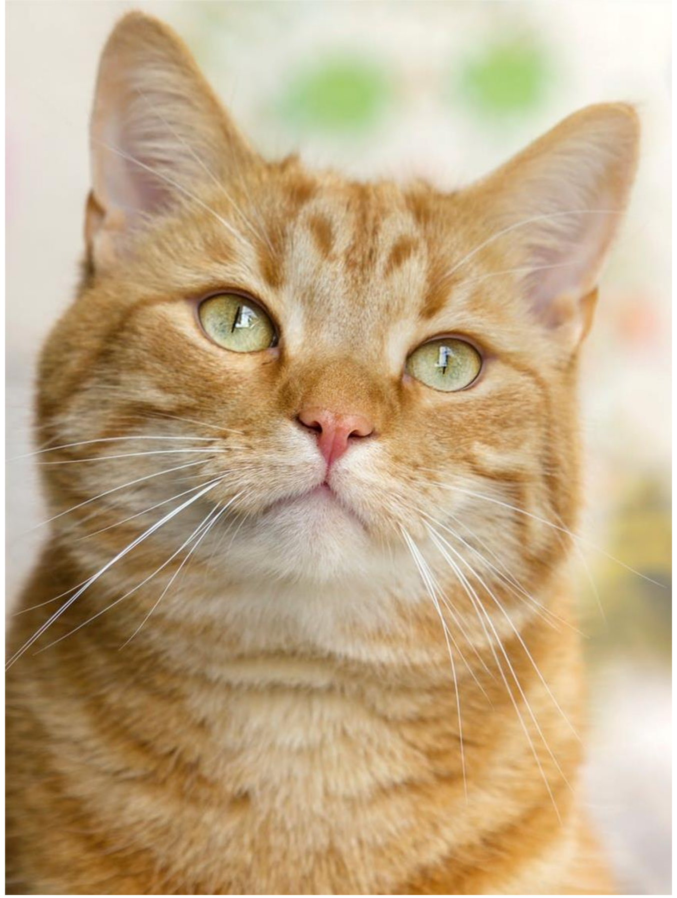

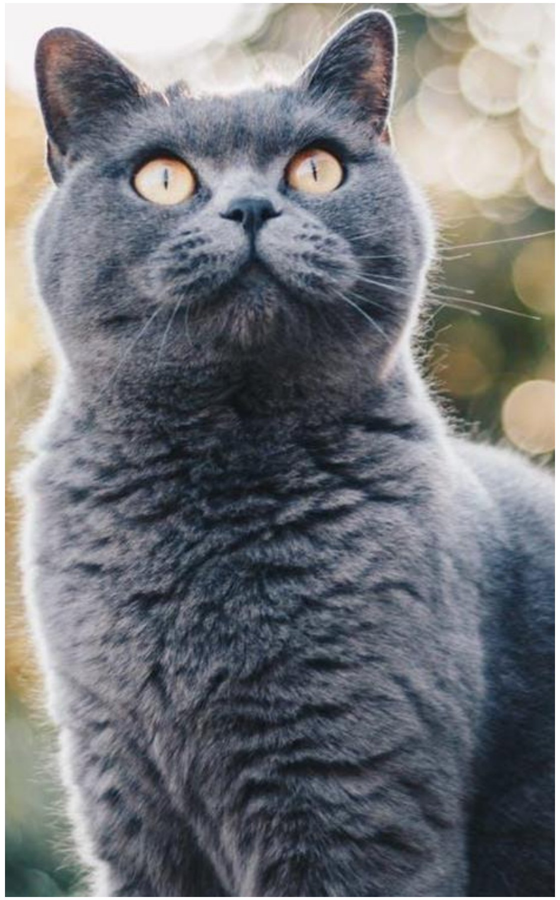

What is a cat?

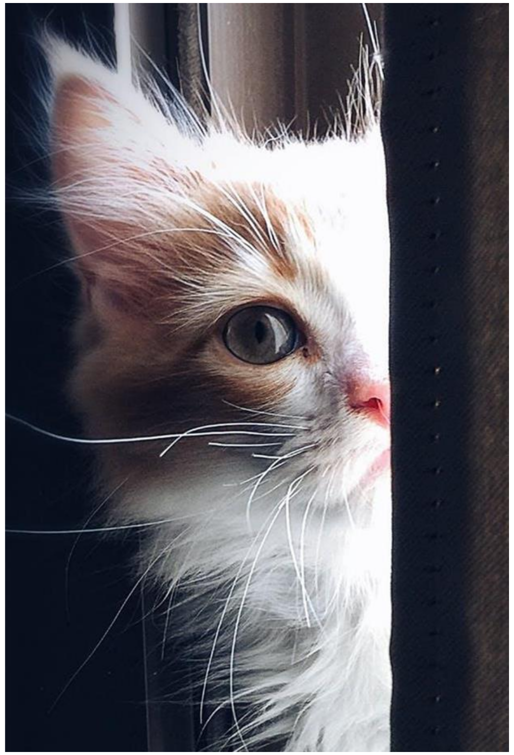
Occlusion

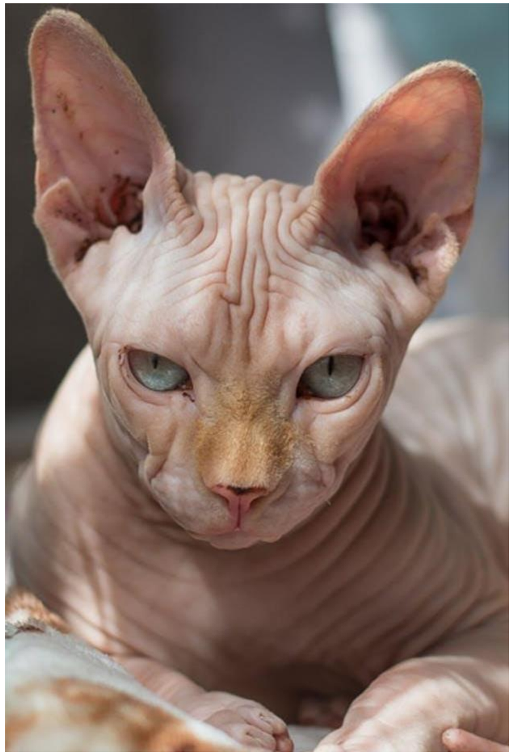
Diversity

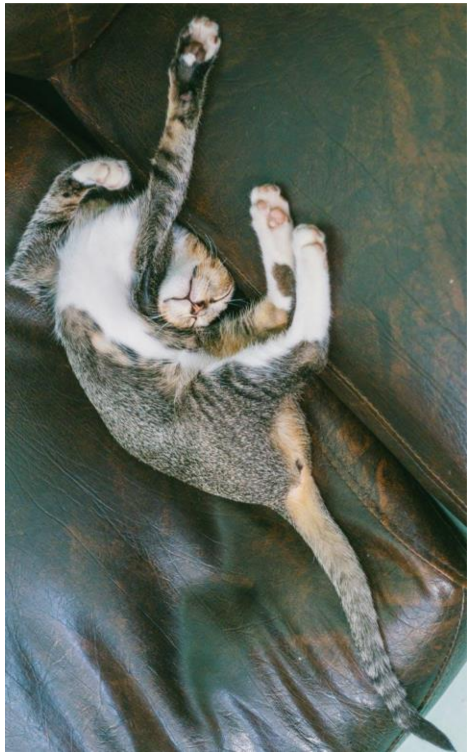
Deformation

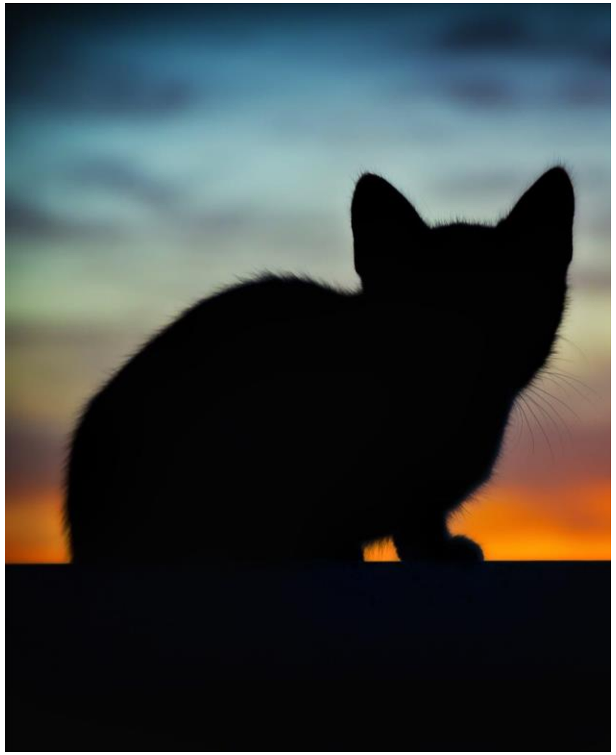
Lighting variations

Introduction to Machine Learning

tryo•labs

What is Machine Learning?

2020 embedded VISION summit

The subfield of computer science that "gives computers the ability to learn without being explicitly programmed".
(Arthur Samuel, 1959)

A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P if its performance at tasks in T, as measured by P, improves with experience E."
(Tom Mitchell, 1997)

Using data for **answering questions**

Training — **Predicting**

tryo·labs

The Big Data era

2020 embedded VISION summit

Accessibility

Data already available everywhere

Low storage costs: everyone has several GBs for "free"

Hardware more **powerful** and **cheaper** than ever before

Devices

Everyone has a powerful computer packed with sensors:

- GPS
- Cameras
- Microphones

Permanently connected to the Internet

Services

Cloud Computing:

- Online storage
- Infrastructure as a Service

User applications:

- YouTube
- Gmail
- Facebook
- Twitter

Types of Machine Learning problems

Supervised

Unsupervised

Reinforcement

Types of Machine Learning problems

Supervised

Unsupervised

Reinforcement

Learn through **examples** of which we know the desired output (what we want to predict).

Is this a cat or a dog?

Are these emails spam or not?

Predict the market value of houses, given the square meters, number of rooms, neighborhood, etc.

Types of Machine Learning problems

Supervised

Unsupervised

Reinforcement

Classification

Output is a **discrete** variable (e.g., cat/dog)

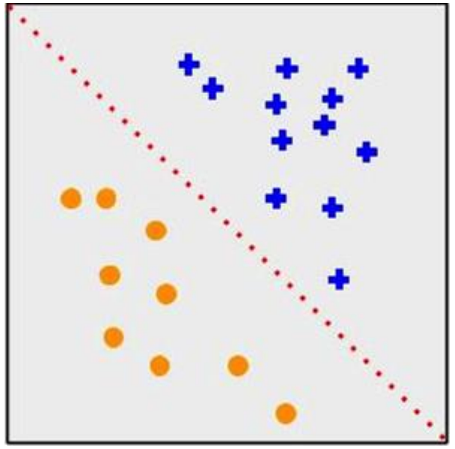

Regression

Output is **continuous** (e.g., price, temperature)

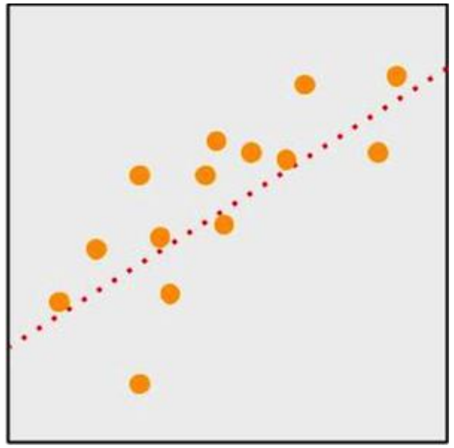

Types of Machine Learning problems

Supervised

Unsupervised

Reinforcement

There is ***no desired* output**. Learn something about the data. *Latent* relationships.

I have photos and want to put them in 20 groups.

I want to find anomalies in the credit card usage patterns of my customers.

Types of Machine Learning problems

Supervised

Unsupervised

Reinforcement

Useful for learning structure in the data **(clustering)**, hidden correlations, reduce dimensionality, etc.

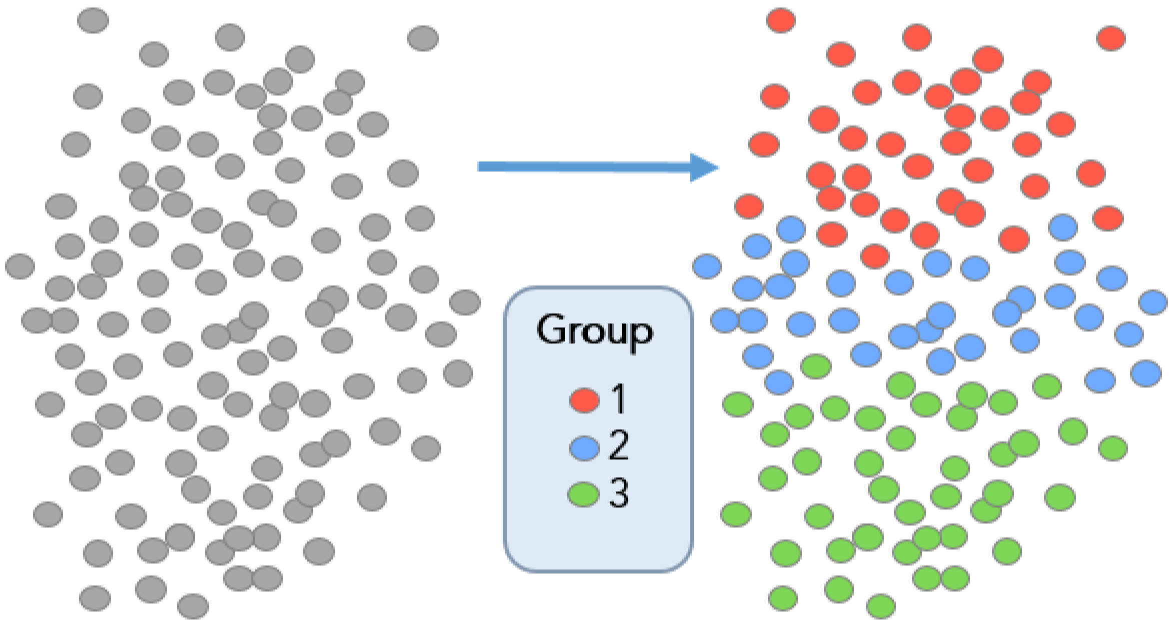

tryo·labs

Types of Machine Learning problems

Supervised

Unsupervised

Reinforcement

An agent **interacts** with an **environment** and watches the result of the interaction.

Environment gives feedback via a positive or negative **reward signal**.

Steps to solve a Machine Learning problem

Data Gathering	Data Preprocessing	Feature Engineering	Algorithm Selection & Training	Making Predictions
Collect data from various sources	Clean data to have homogeneity	Making your data more useful	Selecting the right machine learning model	Evaluate the model

Data gathering

Might depend on human work

- Manual labeling for supervised learning.
- Domain knowledge. Maybe even experts.

May come for free, or “sort of”

- E.g., Machine Translation.

The more the better: Some algorithms need large amounts of data to be useful (e.g., neural networks).

The **quantity** and **quality** of data dictate the model **accuracy**

Data processing

*Is there anything **wrong** with the data?*

- Missing values
- Outliers
- Bad encoding (for text)
- Wrongly-labeled examples
- Biased data

  Do I have many more samples from one class than the rest?

Need to fix/remove data?

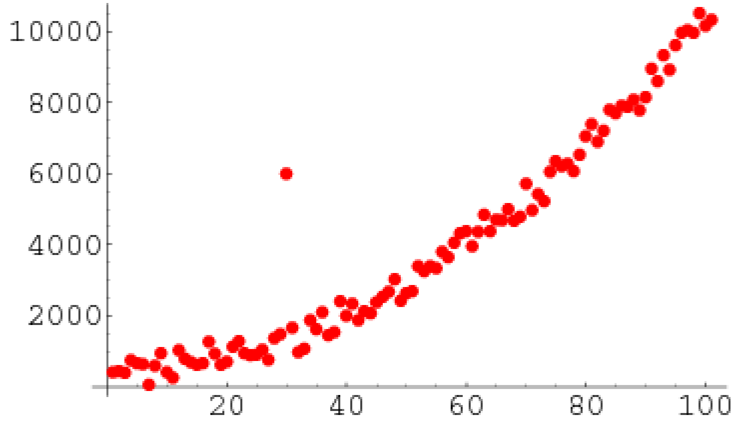

Feature engineering

Our inputs are represented by a **set of features**.

A feature is an individual measurable property of a phenomenon being observed

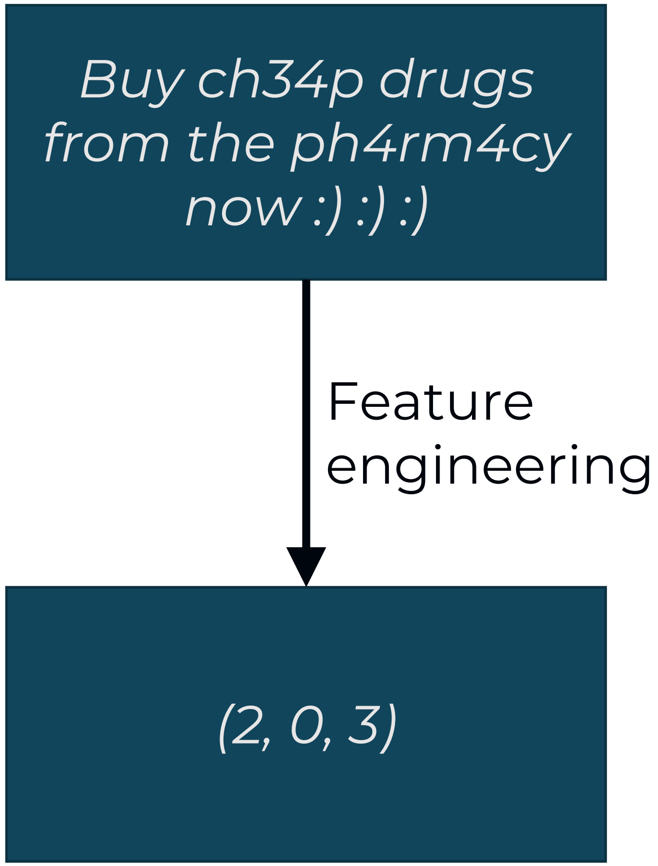

To classify spam email, features could be:

- Number of words that have been *ch4ng3d* like this.
- Language of the email (0=English, 1=Spanish)
- Number of emojis

Feature engineering

Extracts **more** information from **existing** data, doesn't add "new" data

- Making it more **useful**
- With good features, most algorithms can learn **faster**

It can be an art

- Requires thought and knowledge of the data

Two steps:

- Variable transformation (e.g., dates into weekdays, normalizing)
- Feature creation (e.g., n-grams for texts, if word is capitalized to detect names)

Algorithm selection & Training

Supervised

- Linear classifier
- Naive Bayes
- Support Vector Machines (SVM)
- Decision Tree
- Random Forests
- k-Nearest Neighbors
- **Neural Networks (Deep learning)**

Unsupervised

- PCA
- t-SNE
- k-means
- DBSCAN

Reinforcement

- SARSA–λ
- Q-Learning

Algorithm selection & Training

Goal of training: make the correct prediction as often as possible

- Incremental improvement:

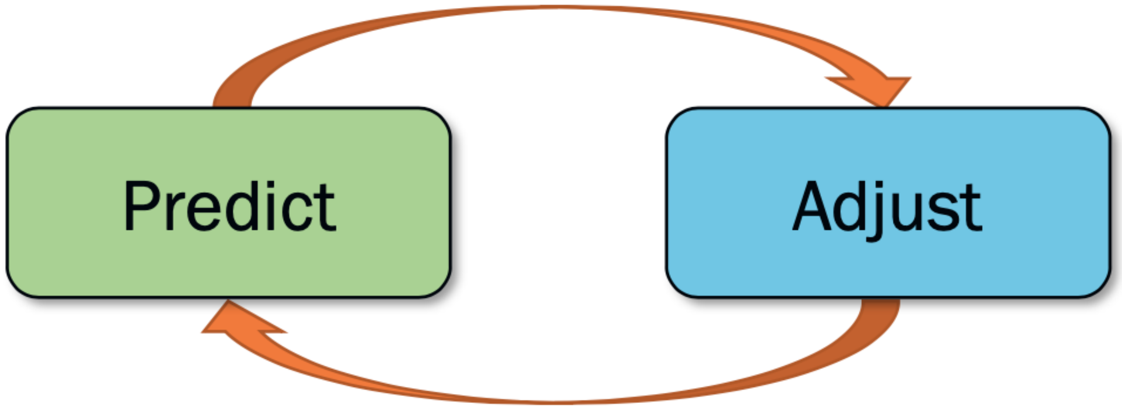

- Metrics for **evaluating** performance and comparing solutions.
- **Hyperparameter tuning**: more an art than a science

Making predictions

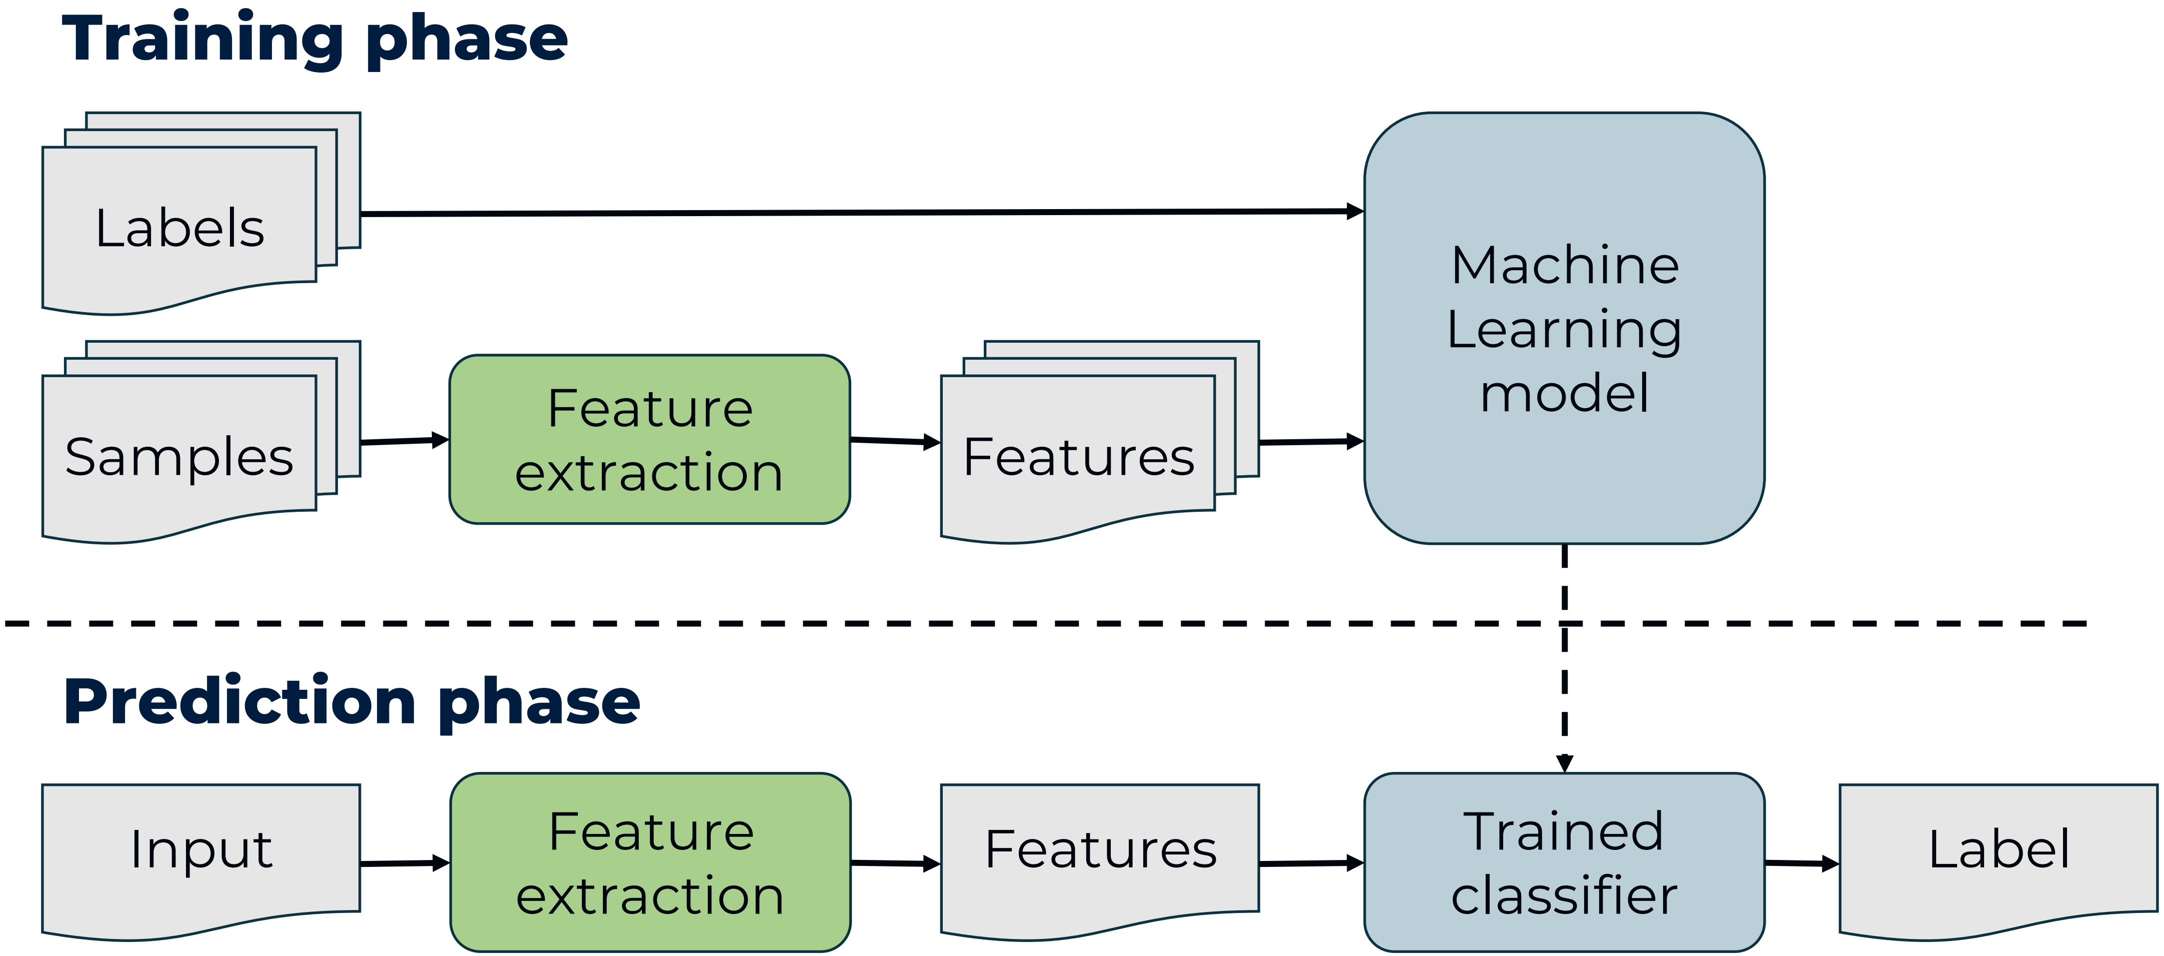

Summary

- Machine Learning is intelligent **use of data** to **answer questions**
- Enabled by an **exponential** increase in computing power and data availability
- Three big types of problems: **supervised**, **unsupervised**, **reinforcement**

Summary

5 steps to **every** machine learning solution:

1. Data Gathering
2. Data Preprocessing
3. Feature Engineering
4. Algorithm Selection & Training
5. Making Predictions

Deep Learning

"Any sufficiently advanced technology is indistinguishable from magic." (Arthur C. Clarke)

tryo·labs

Artificial Neural Networks

- First model of artificial neural network proposed in **1943**
- Analogy to the *human brain* greatly exaggerated
- Given some inputs (x), the network calculates some outputs (y) using a set of weights (w)

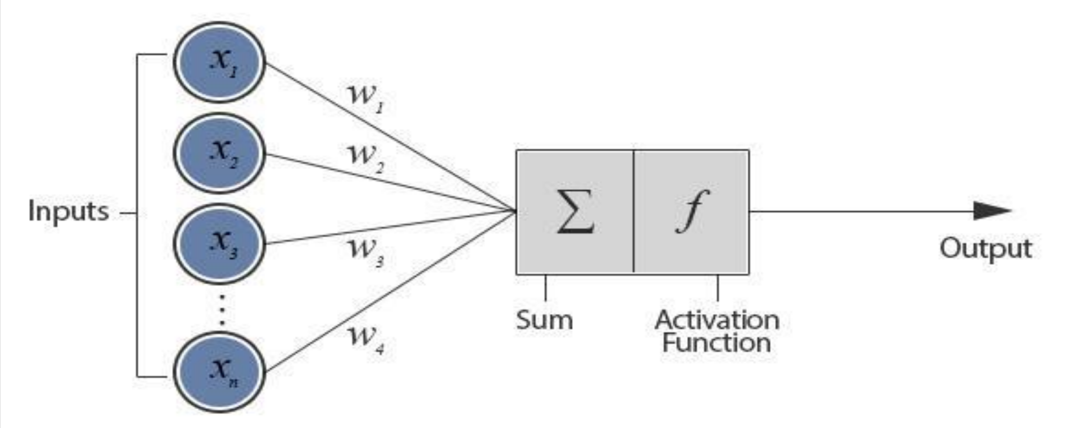

Perceptron *(Rosenblatt, 1957)*

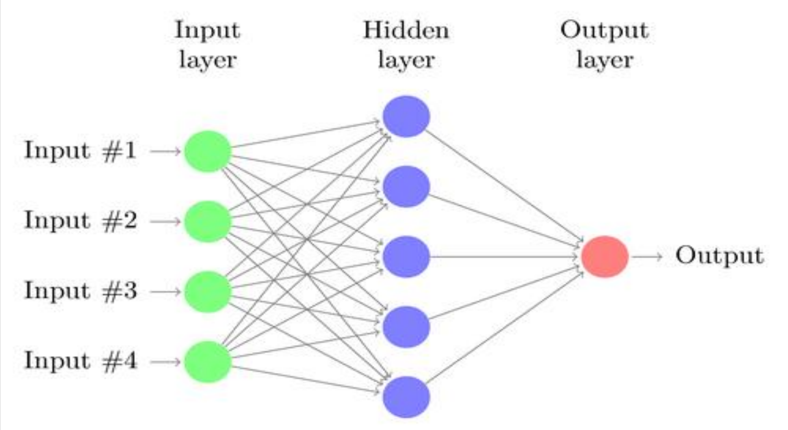

Two-layer Fully Connected Neural Network

Artificial Neural Networks

- Weights must be adjusted (**learned** from the data)
- Idea: define a function that tells us **how “close”** the network is to generating the desired output
- **Minimize** the loss ➔ **optimization** problem
- With a **continuous** and **differentiable** loss function, we can apply **gradient descent**

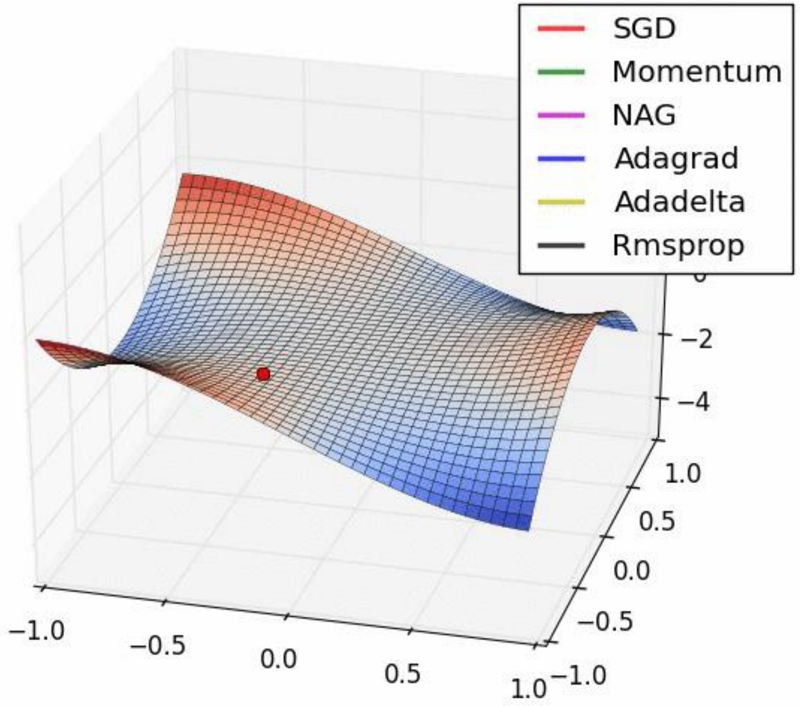

The rise, fall, rise, fall and rise of Neural Networks

- Perceptron gained popularity in the 60s
  - Belief that would lead to true AI
- XOR problem and AI Winter (1969 – 1986)
- Backpropagation to the rescue! (1986)
  - Training of multilayer neural nets
  - LeNet-5 (Yann LeCun et al., 1998)
- Unable to scale. Lack of good data and processing power

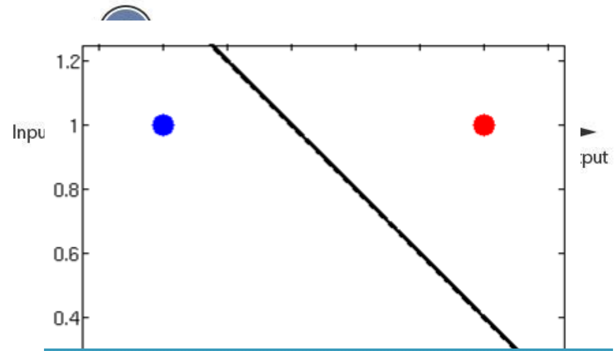

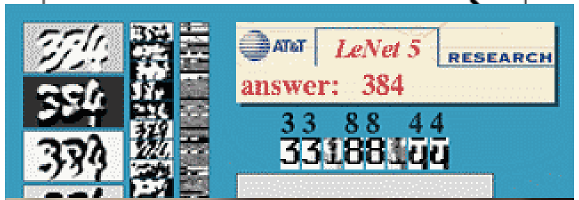

The rise, fall, rise, fall and rise of Neural Networks

- Regained popularity since ~2006.
  - Train each layer at a time
  - Rebranded field as **Deep Learning**
  - Old ideas rediscovered (e.g., Convolution)
- Breakthrough in 2012 with AlexNet (Krizhevsky et al.)
  - Use of GPUs
  - Convolution

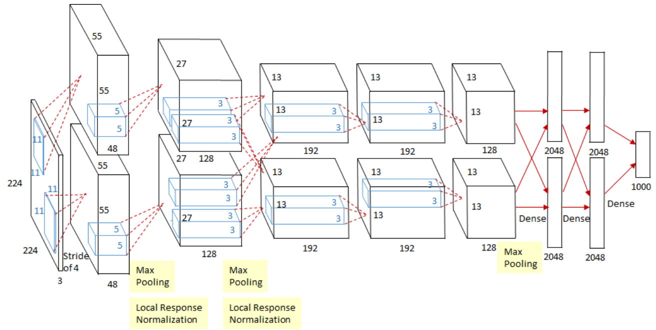

Image classification with Deep Neural Networks

tryo·labs

Image classification with Deep Neural Networks
The convolution operator

=

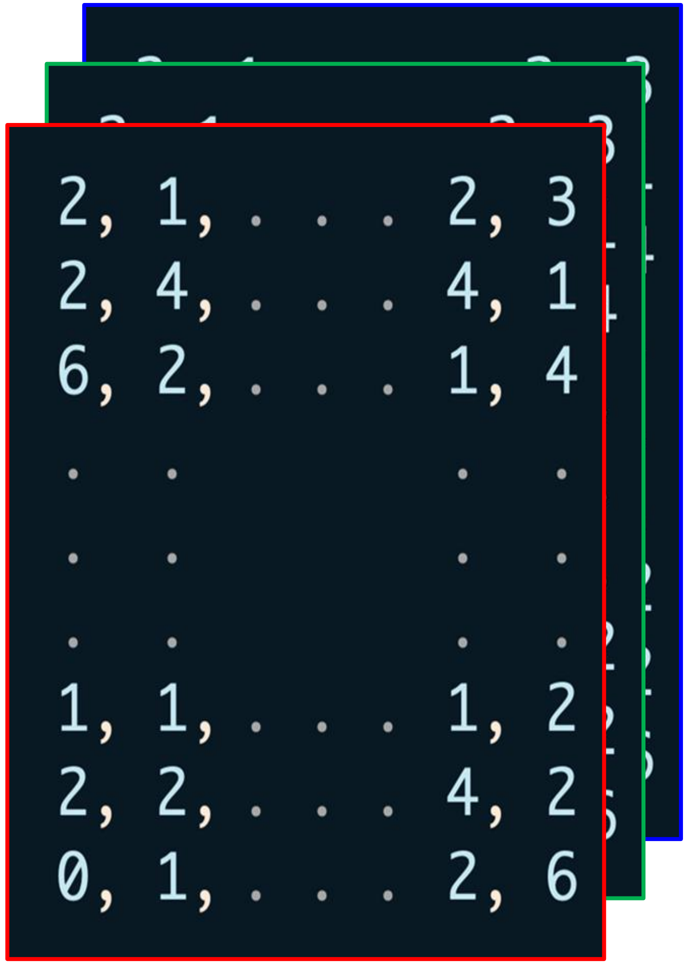

Image classification with Deep Neural Networks

The convolution operator

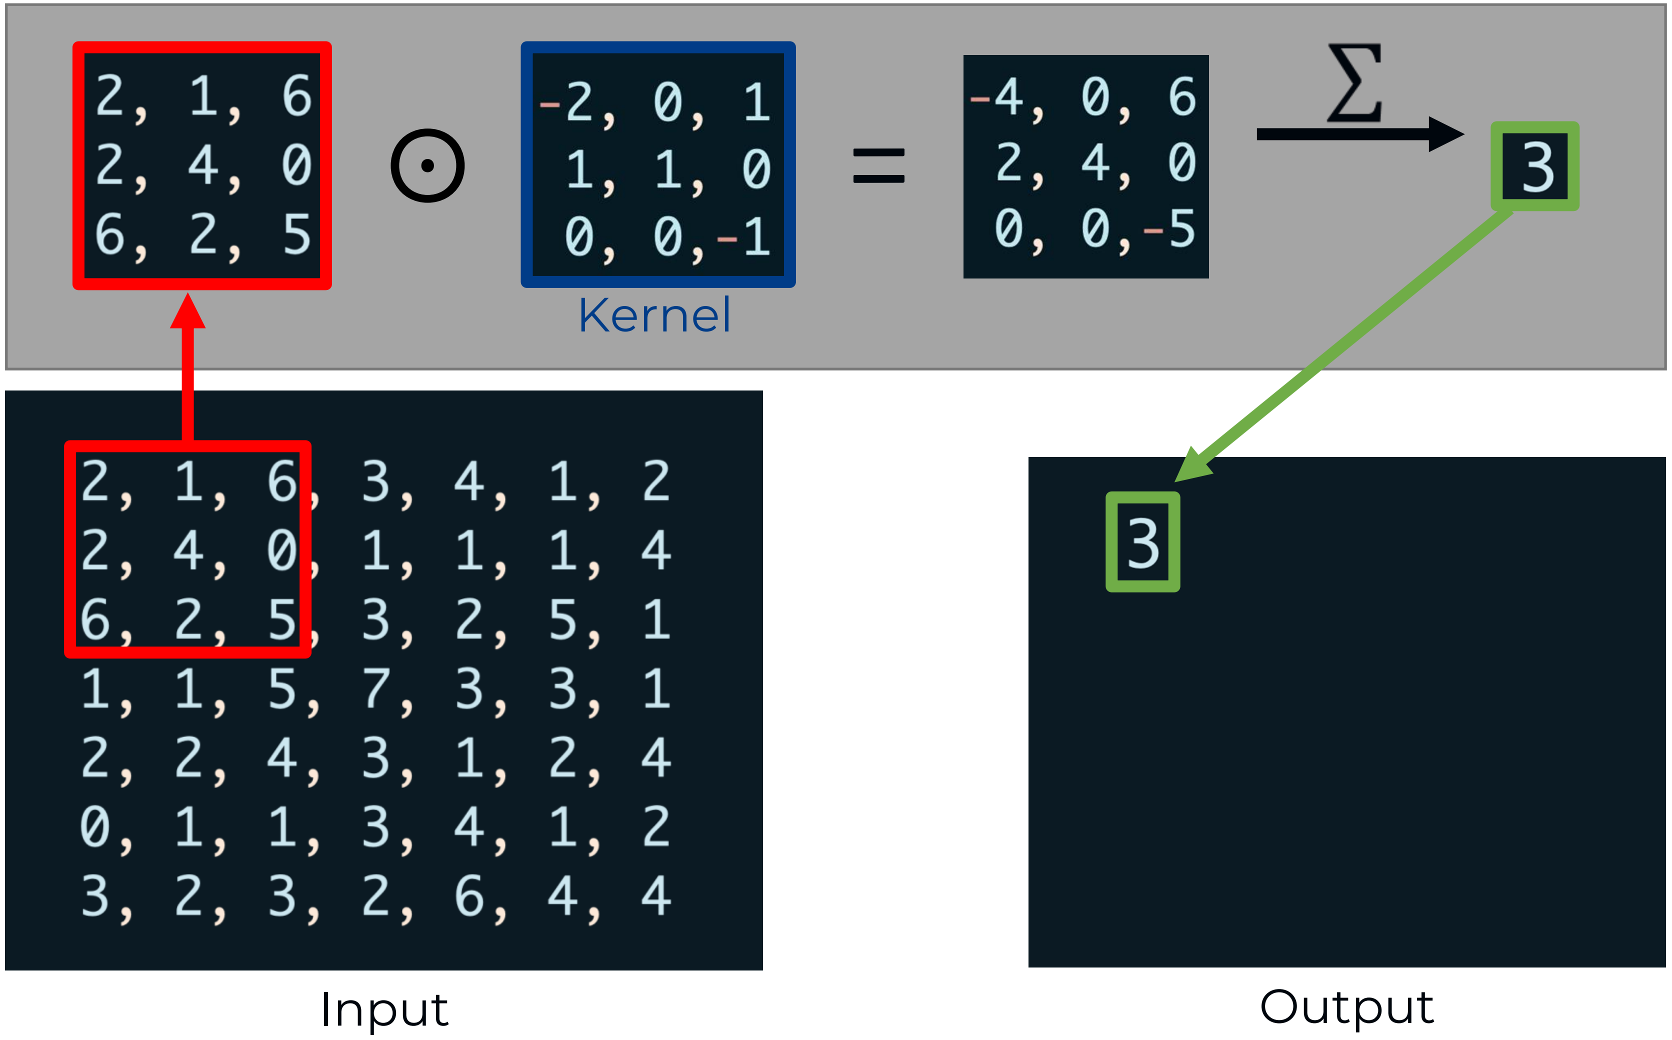

Image classification with Deep Neural Networks

The convolution operator

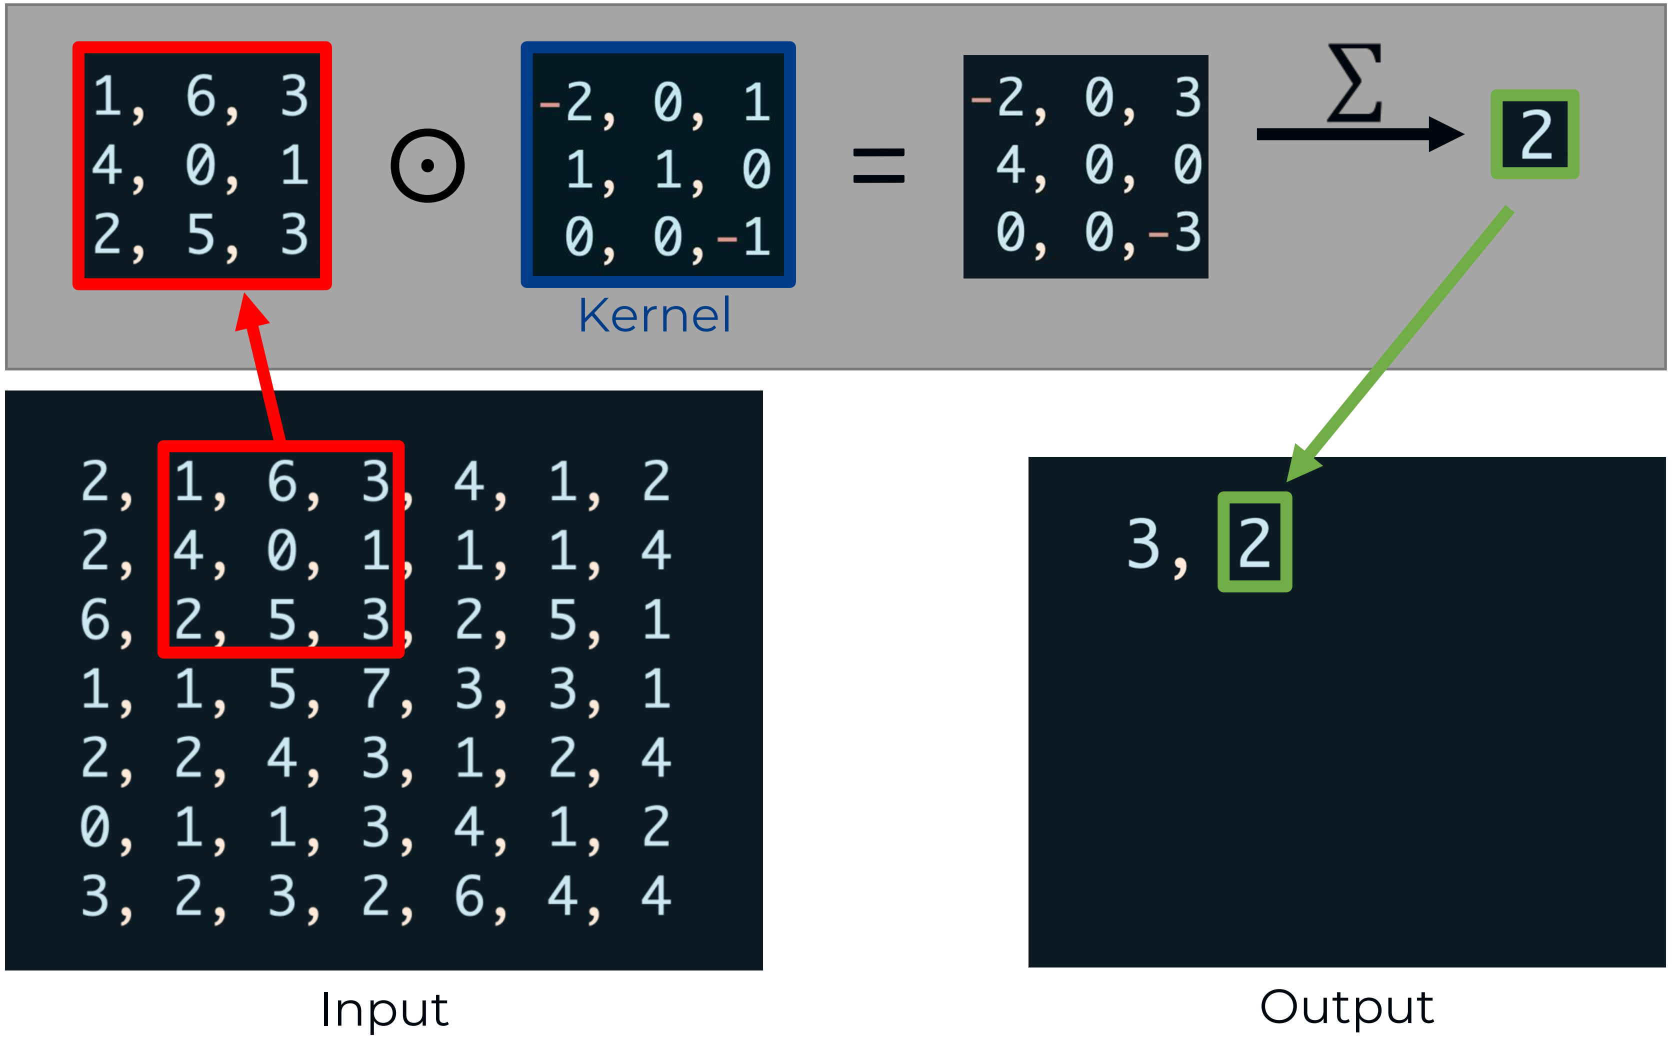

tryo·labs

The convolution operator

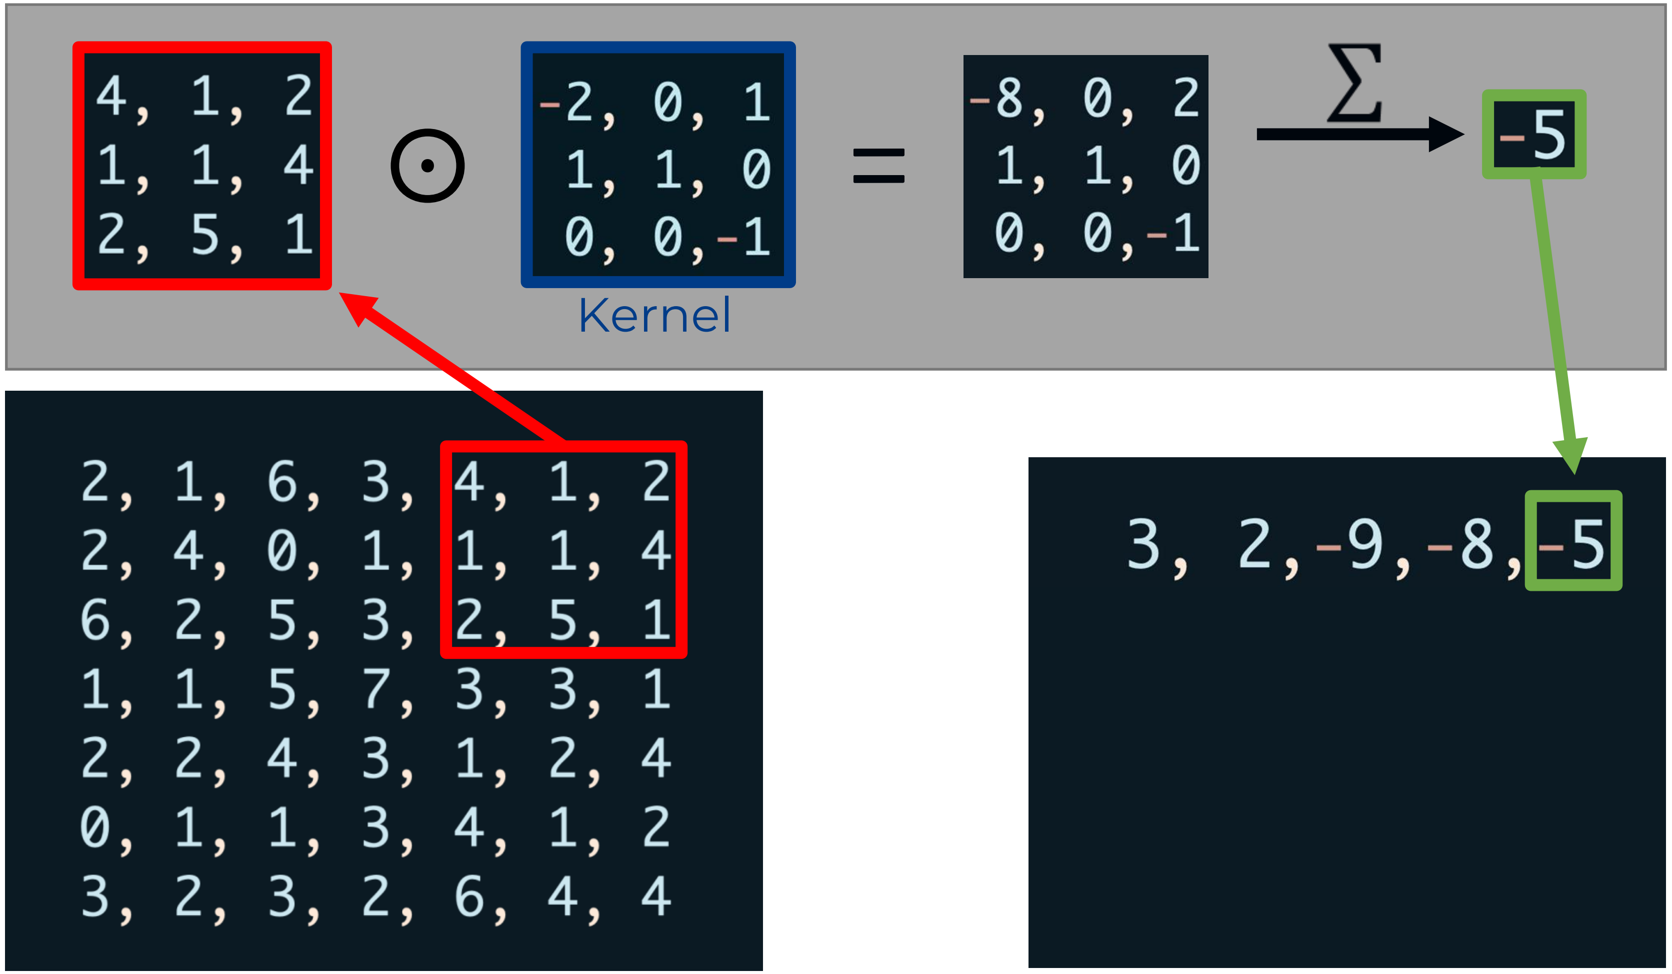

Input

Output

The convolution operator

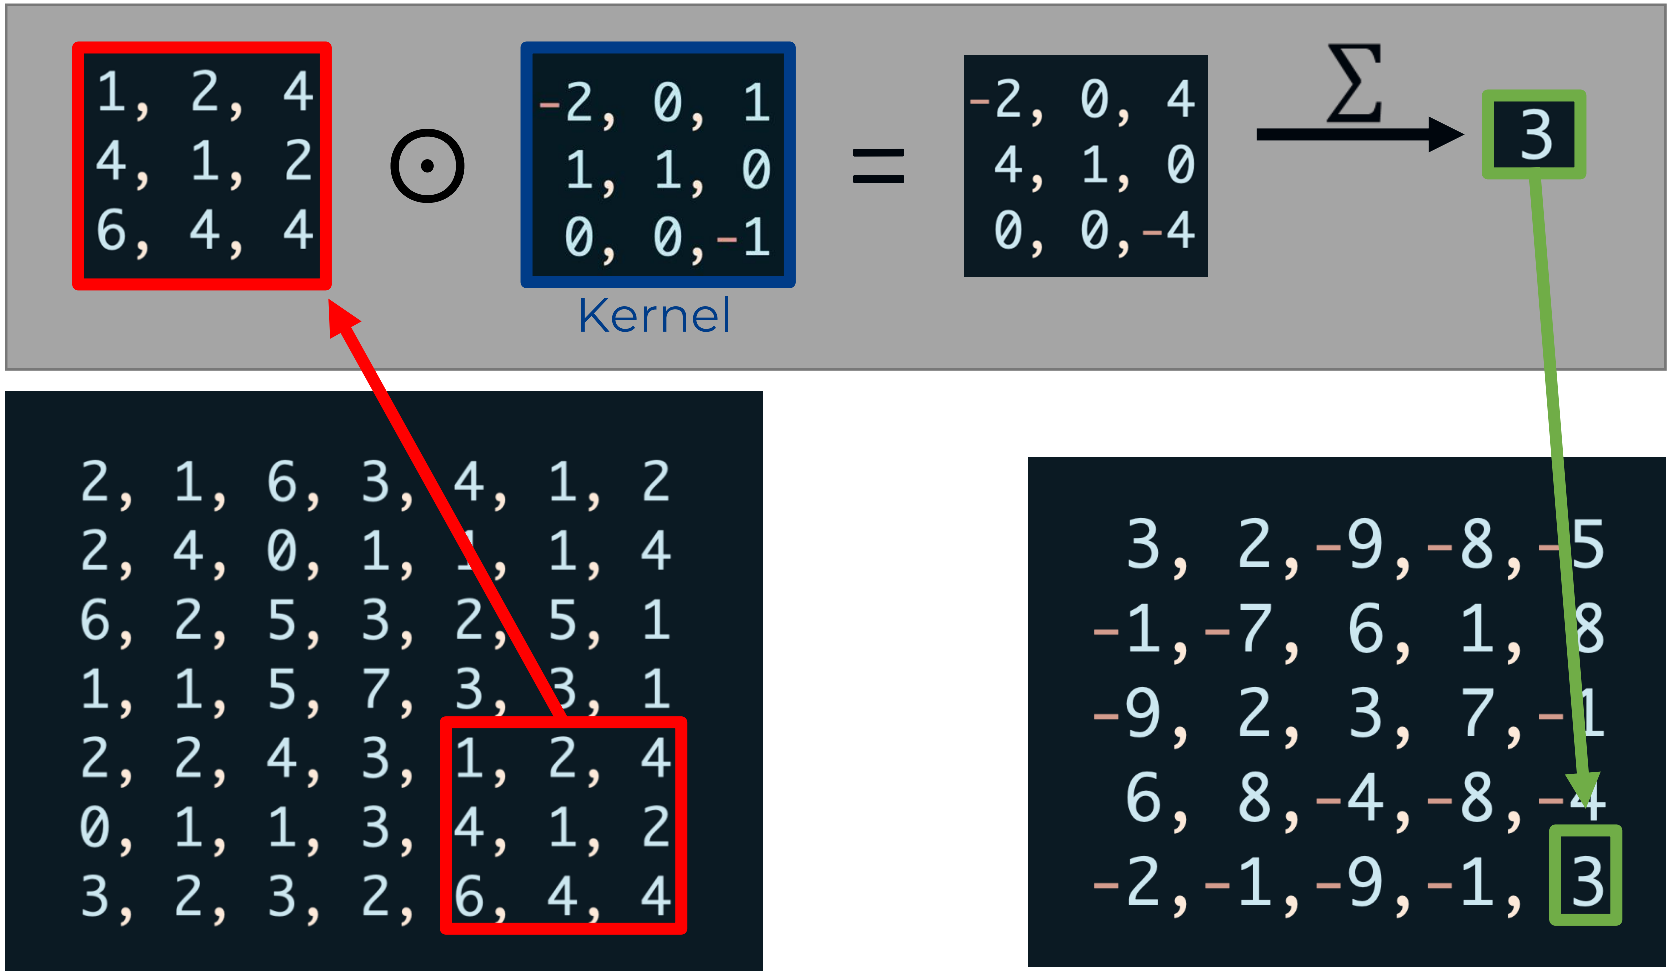

Input

Output

tryo·labs

The convolution operator

Input

```
-1, 0, 1
-1, 0, 1
-1, 0, 1
```

Kernel

Feature Map

The convolution operator

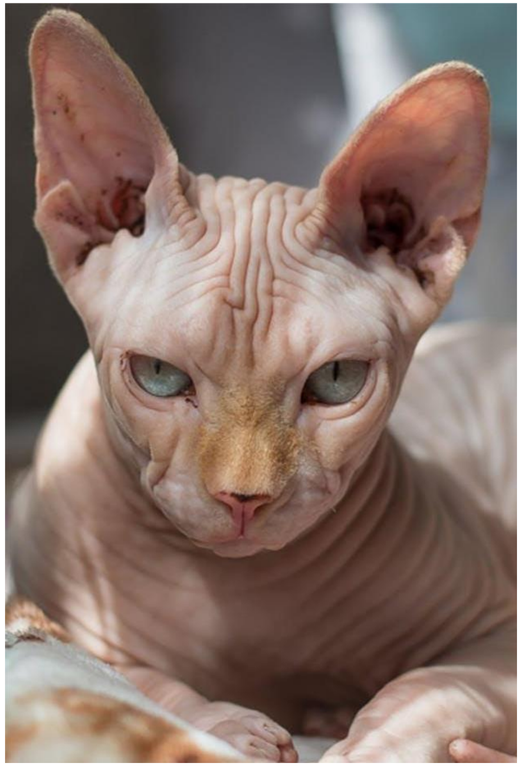

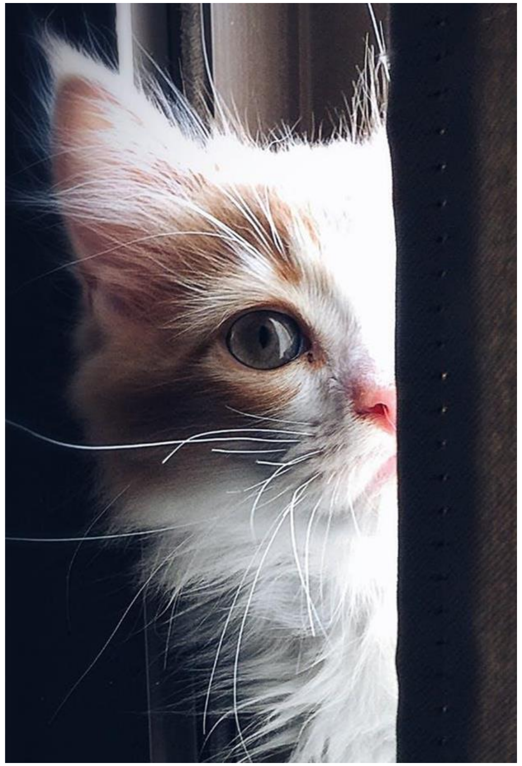

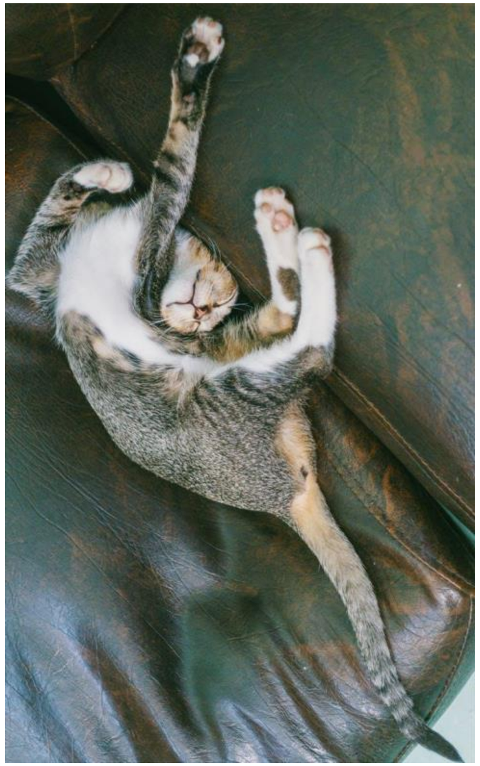

The convolution operator

- Takes spatial dependencies into account
- Used as a **feature extraction** tool
- **Differentiable** operation ➔ the kernels can be **learned**

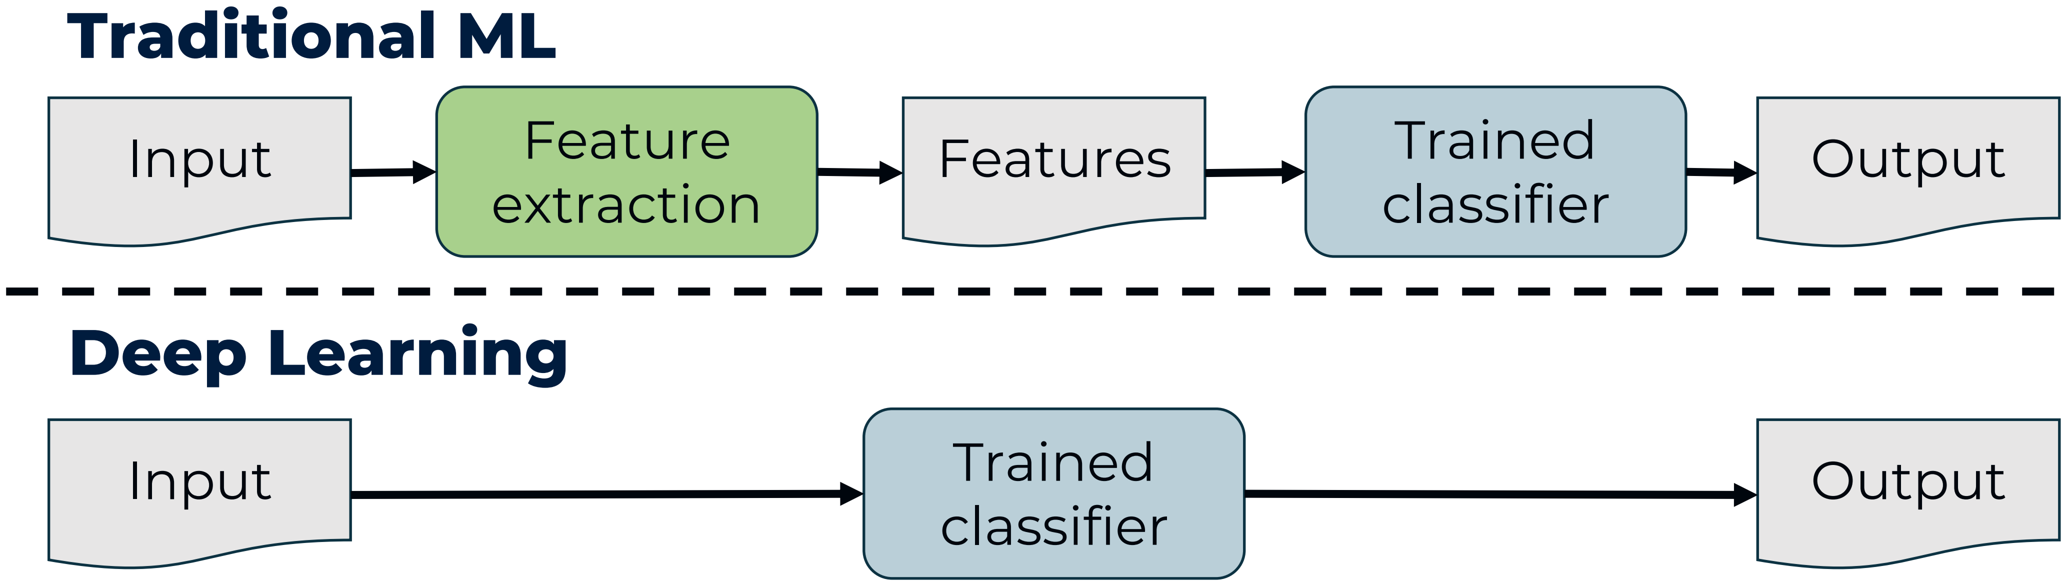

tryo·labs

Non-linear activation functions

Increment the network's **capacity**

- Convolution, matrix multiplication and summation are linear

Sigmoid

$$f(x) = \frac{1}{1 + e^{-x}}$$

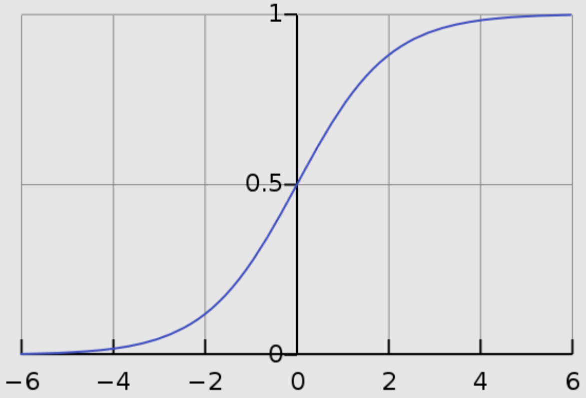

Hyperbolic tangent

$$tanh(x) = \frac{e^{2x-1}}{e^{2x+1}}$$

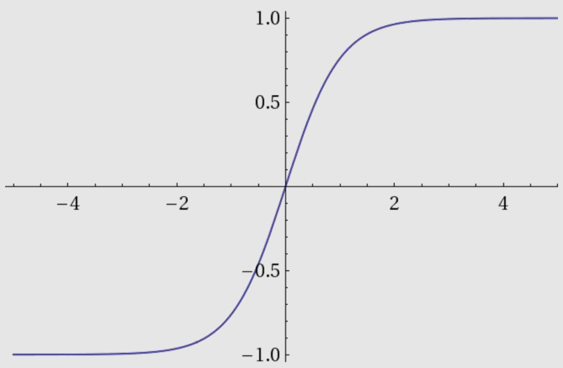

ReLU

$$f(x) = \max(0, x)$$

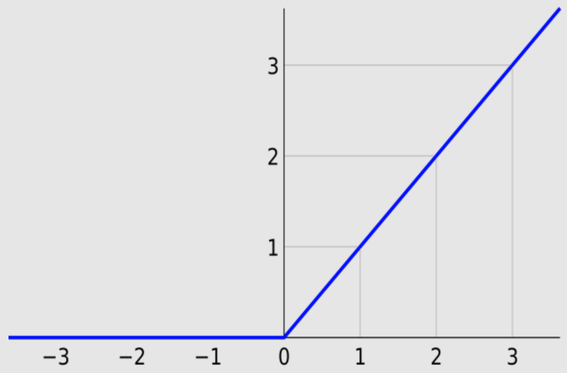

Image classification with Deep Neural Networks

Non-linear activation functions

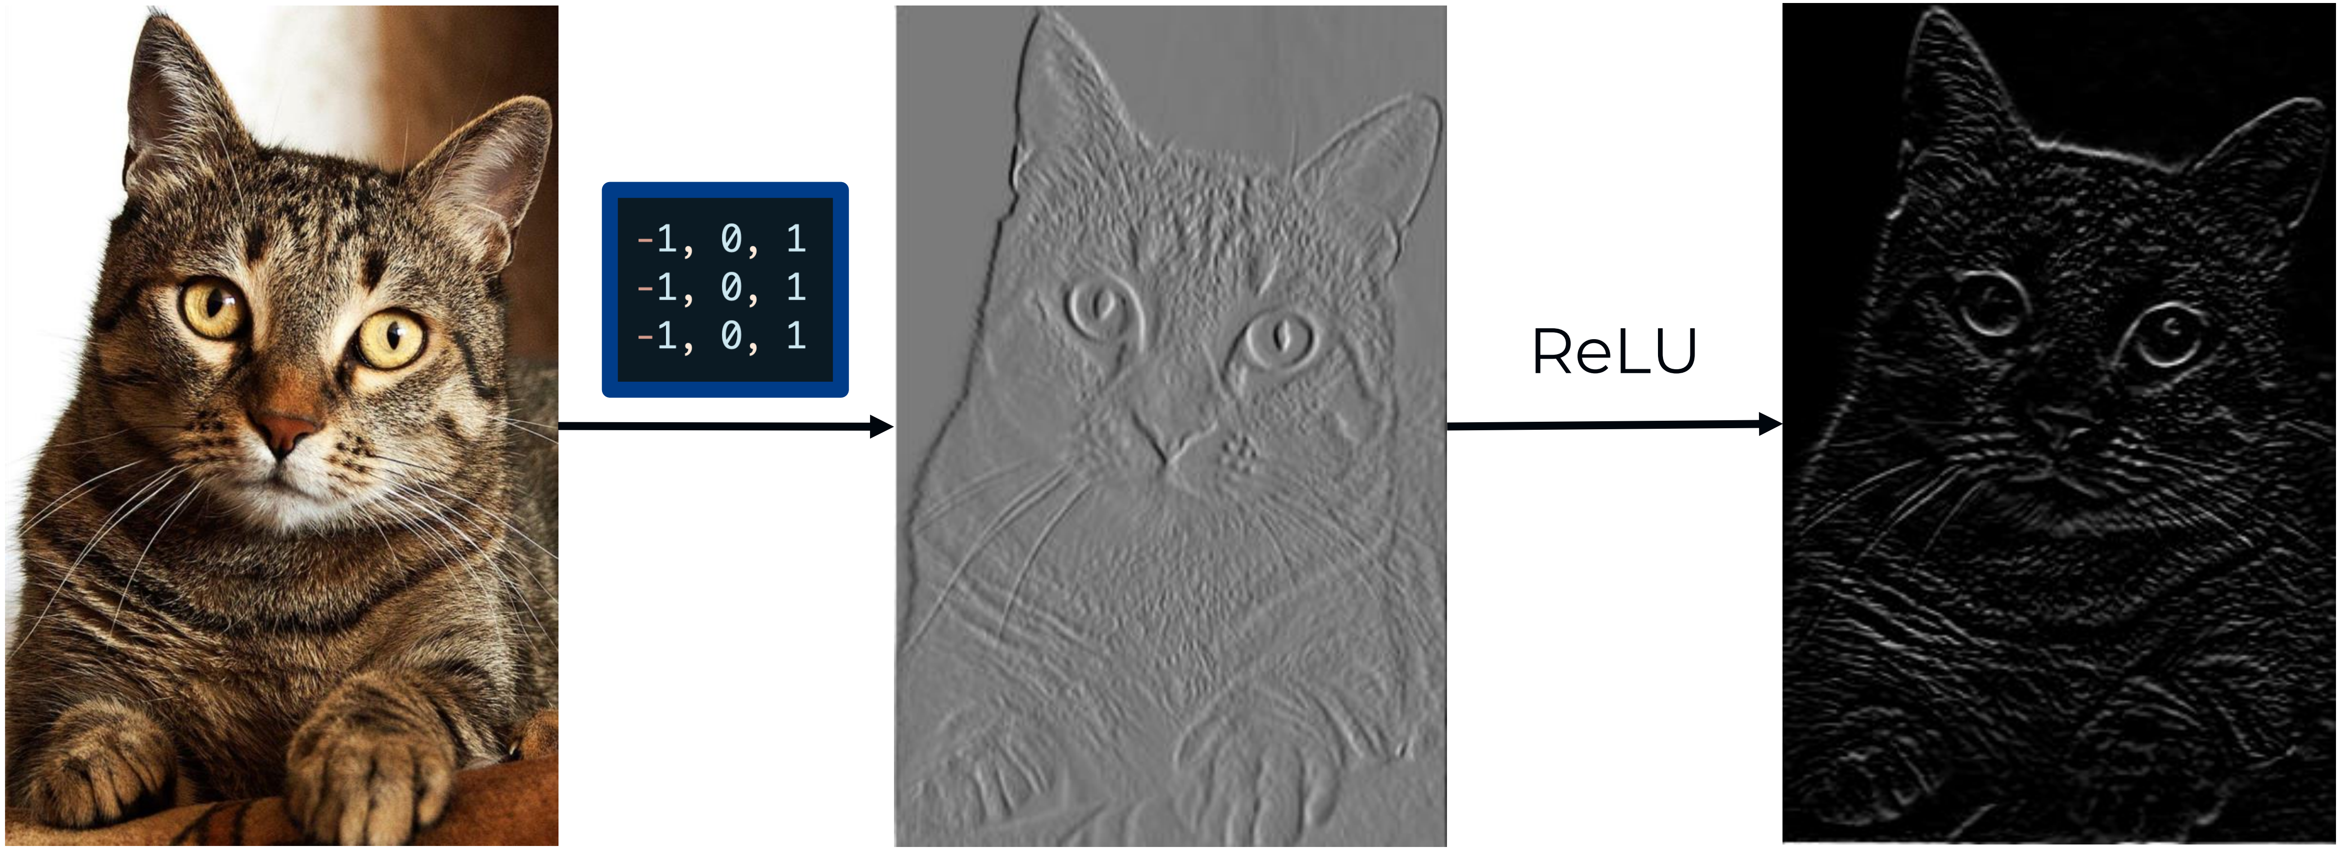

The pooling operation

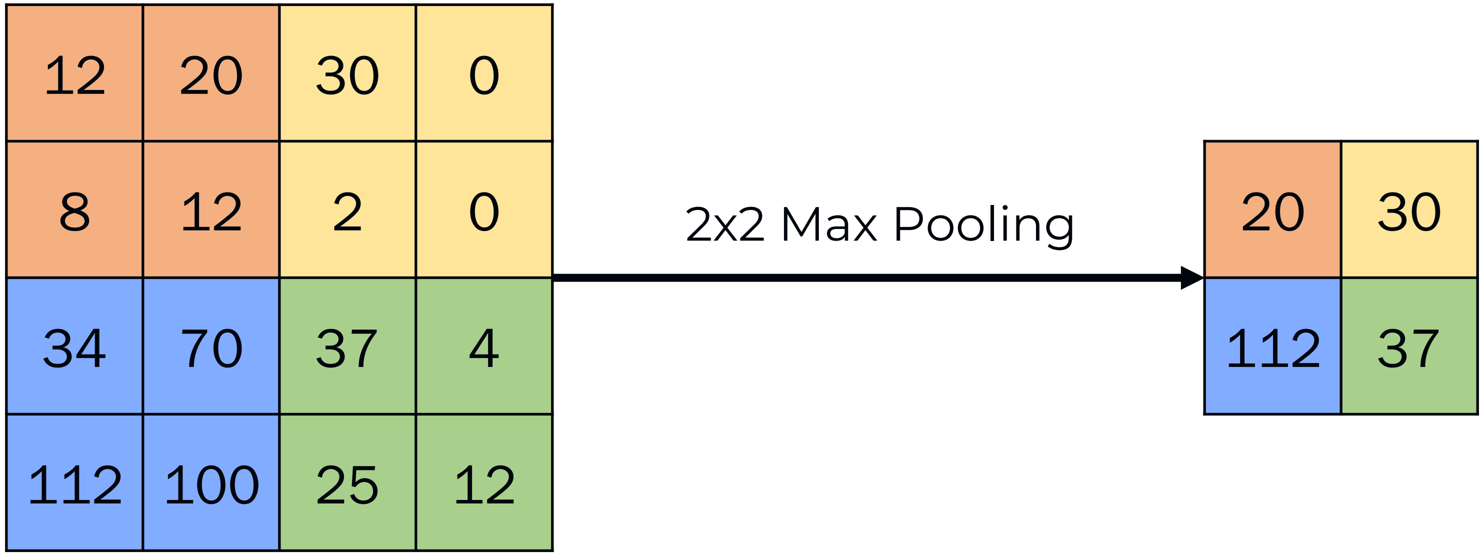

- Used to reduce dimensionality
- Most common: **Max pooling**
- Makes the network invariant to small **transformations**, **distortions** and **translations.**

Putting all together

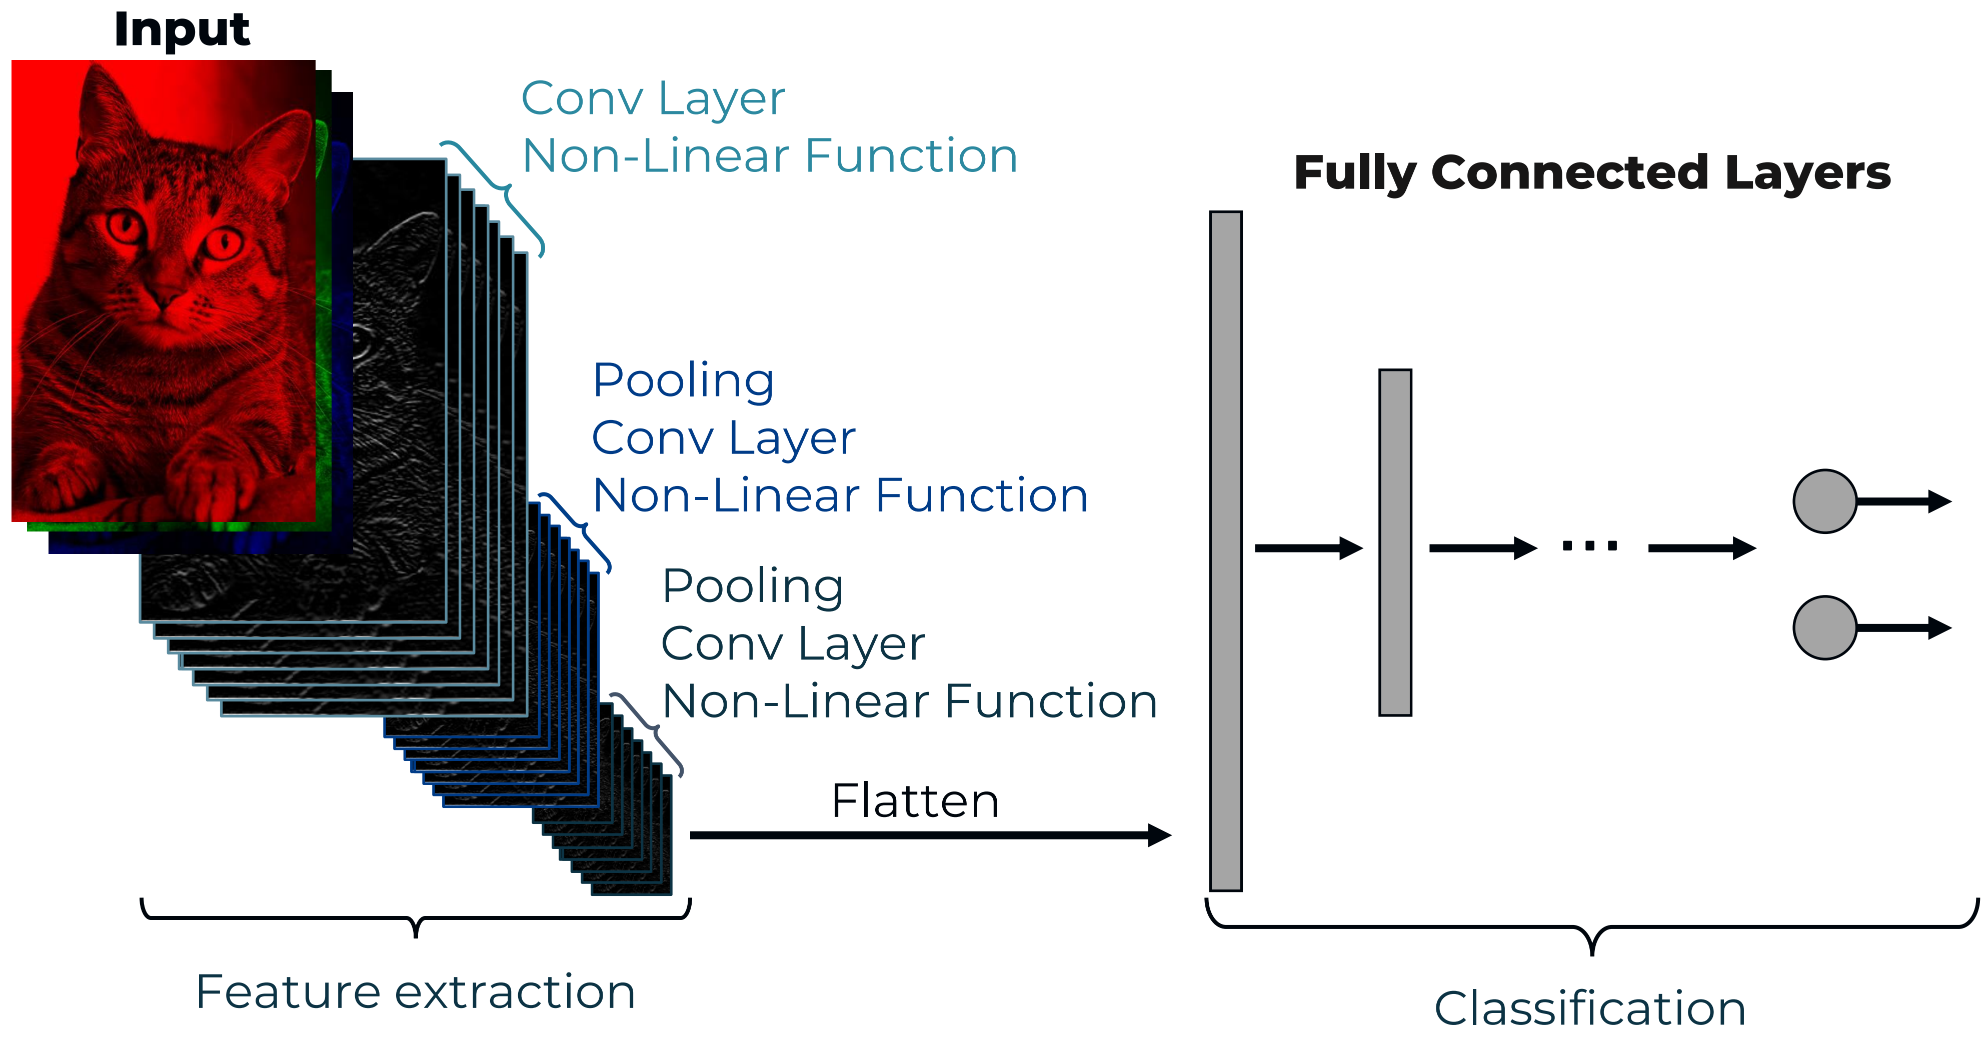

tryo·labs

Training Convolutional Neural Networks

Image classification is a **supervised problem**

- Gather images and **label** them with desired output
- Train the network with **backpropagation**!

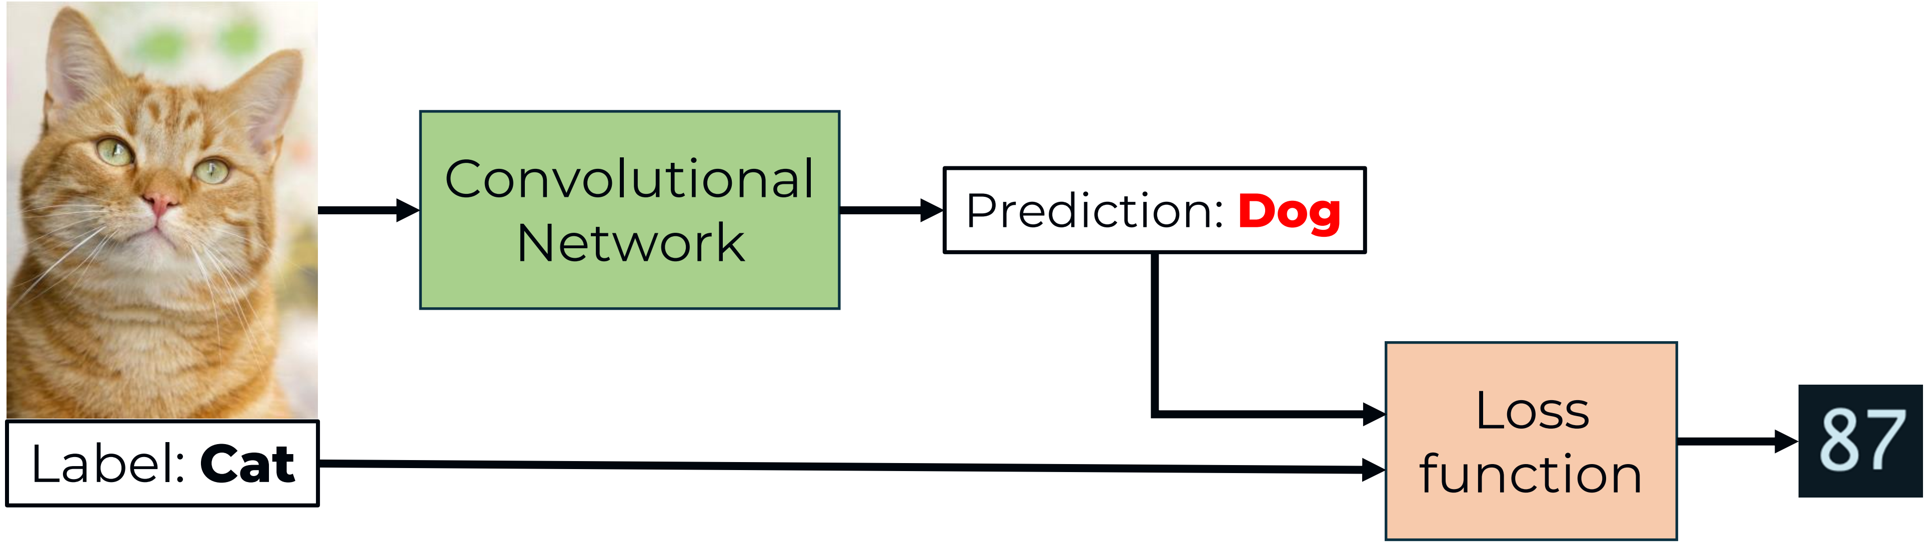

Training Convolutional Neural Networks

Image classification is a **supervised problem**

- Gather images and **label** them with desired output
- Train the network with **backpropagation**!

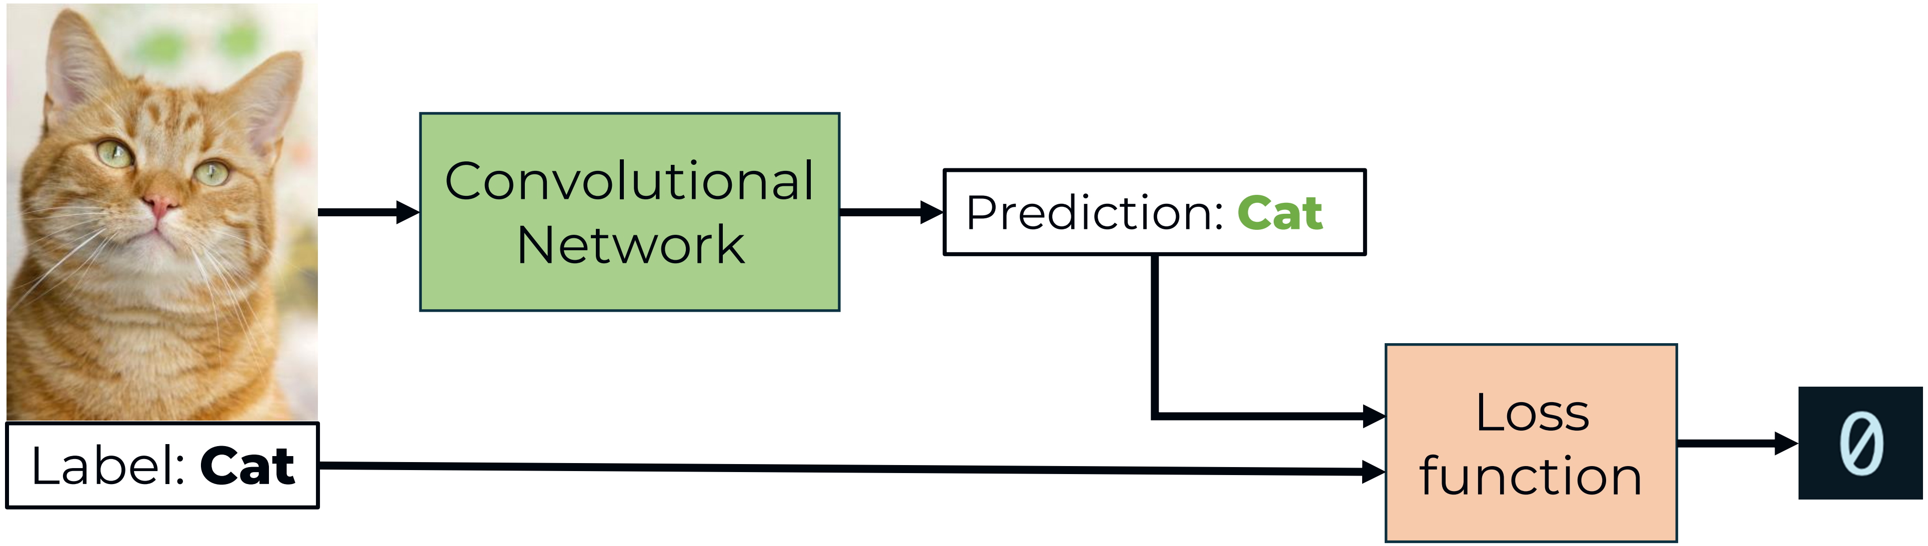

Image classification with Deep Neural Networks
Surpassing human performance

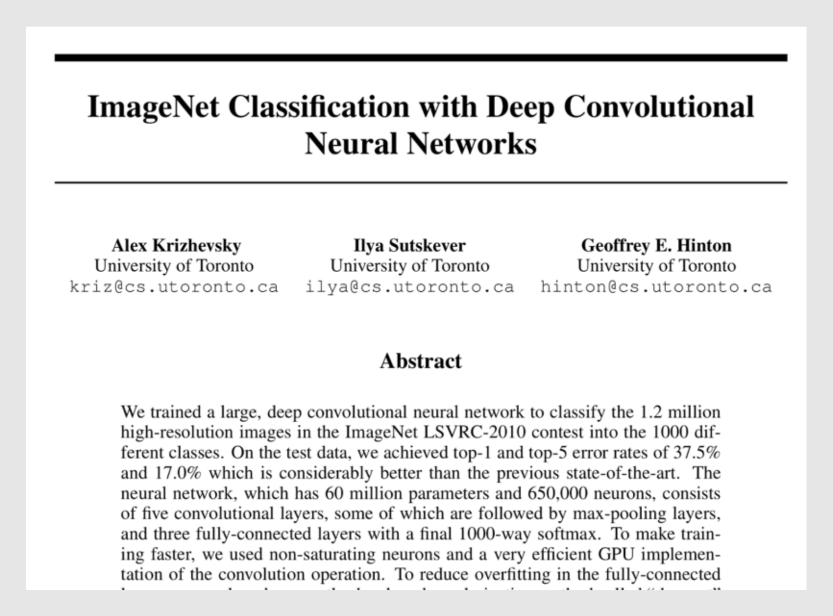

ImageNet Classification with Deep Convolutional Neural Networks

Alex Krizhevsky
University of Toronto
kriz@cs.utoronto.ca

Ilya Sutskever
University of Toronto
ilya@cs.utoronto.ca

Geoffrey E. Hinton
University of Toronto
hinton@cs.utoronto.ca

Abstract

We trained a large, deep convolutional neural network to classify the 1.2 million high-resolution images in the ImageNet LSVRC-2010 contest into the 1000 different classes. On the test data, we achieved top-1 and top-5 error rates of 37.5% and 17.0% which is considerably better than the previous state-of-the-art. The neural network, which has 60 million parameters and 650,000 neurons, consists of five convolutional layers, some of which are followed by max-pooling layers, and three fully-connected layers with a final 1000-way softmax. To make training faster, we used non-saturating neurons and a very efficient GPU implementation of the convolution operation. To reduce overfitting in the fully-connected

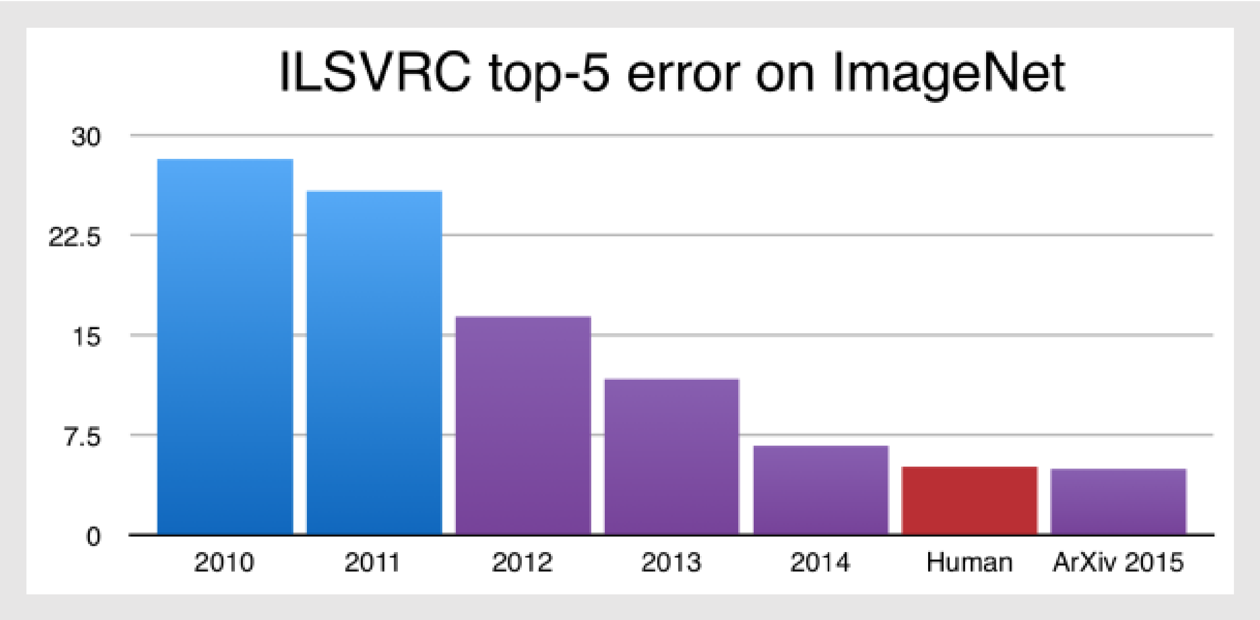

tryo·labs

Deep Learning in the wild

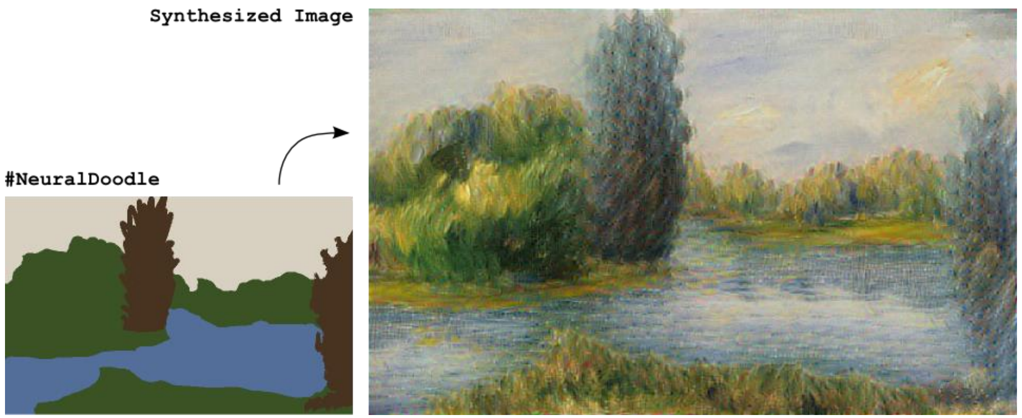

Deep Learning is here to stay

2020 embedded vision summit

Data

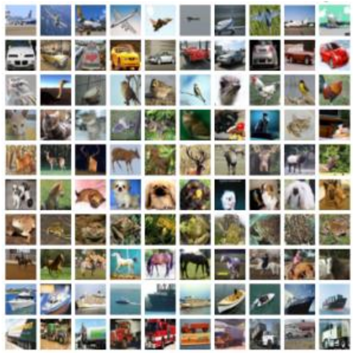

Architectures

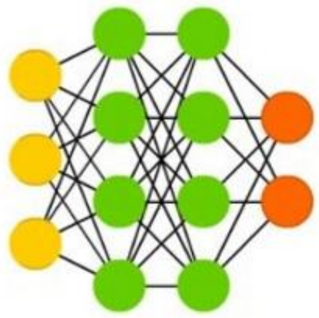

Frameworks

TensorFlow™
PyTorch
torch
theano
Microsoft CNTK

Power

amazon web services
Google Cloud Platform
NVIDIA

Players

Microsoft
Google
Facebook
Baidu 百度
DeepMind

tryo·labs

Conclusions

Conclusions

Machine learning algorithms learn from data to find hidden relations, to make predictions, to interact with the world, …

A machine learning algorithm is as good as its input data

Good model** + **Bad data** = **Bad Results

Deep learning is making significant breakthroughs in: speech recognition, language processing, computer vision, control systems, …

If you are not using or considering using Deep Learning to understand or solve vision problems, **you almost certainly should try it**

Resource

Our work

An Introductory Guide to Computer Vision
https://tryolabs.com/resources/introductory-guide-computer-vision/

Tryolabs Blog
https://www.tryolabs.com/blog

Norfair
https://github.com/tryolabs/norfair

Resource

To learn more...

Google Machine Learning Crash Course
https://developers.google.com/machine-learning/crash-course/

Stanford course **CS229**: Machine Learning
https://developers.google.com/machine-learning/crash-course/

Stanford course **CS231n**: Convolutional Neural Networks for Visual Recognition
http://cs231n.stanford.edu/

Thank you!