Neural Networks for Images

How to see like a human

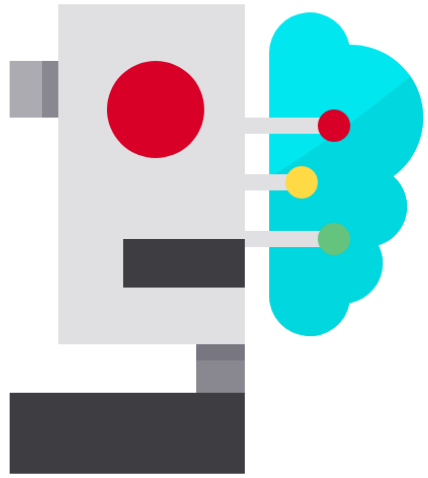

Yordan Darakchiev

iordan93@gmail.com

Machine Learning

- Making a program which performs a task *without* explicitly programming it
  - Like the way people learn

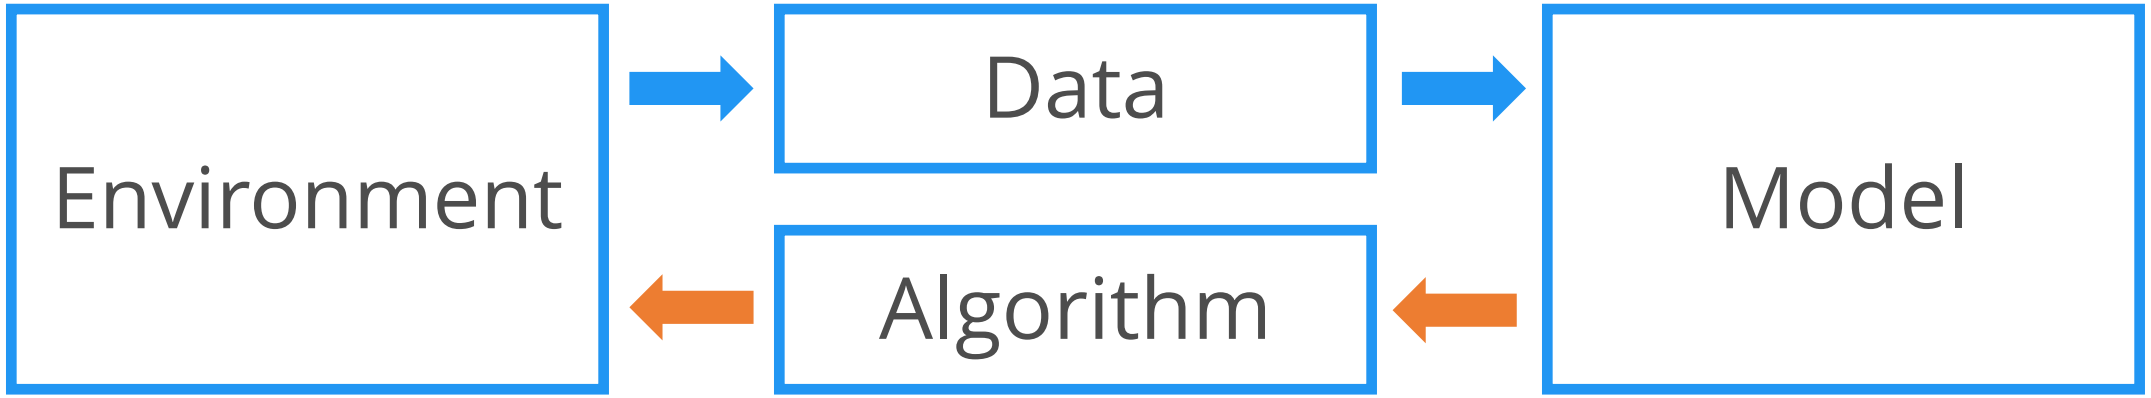

- Seeing new things
- Learning

Neural Network

Data

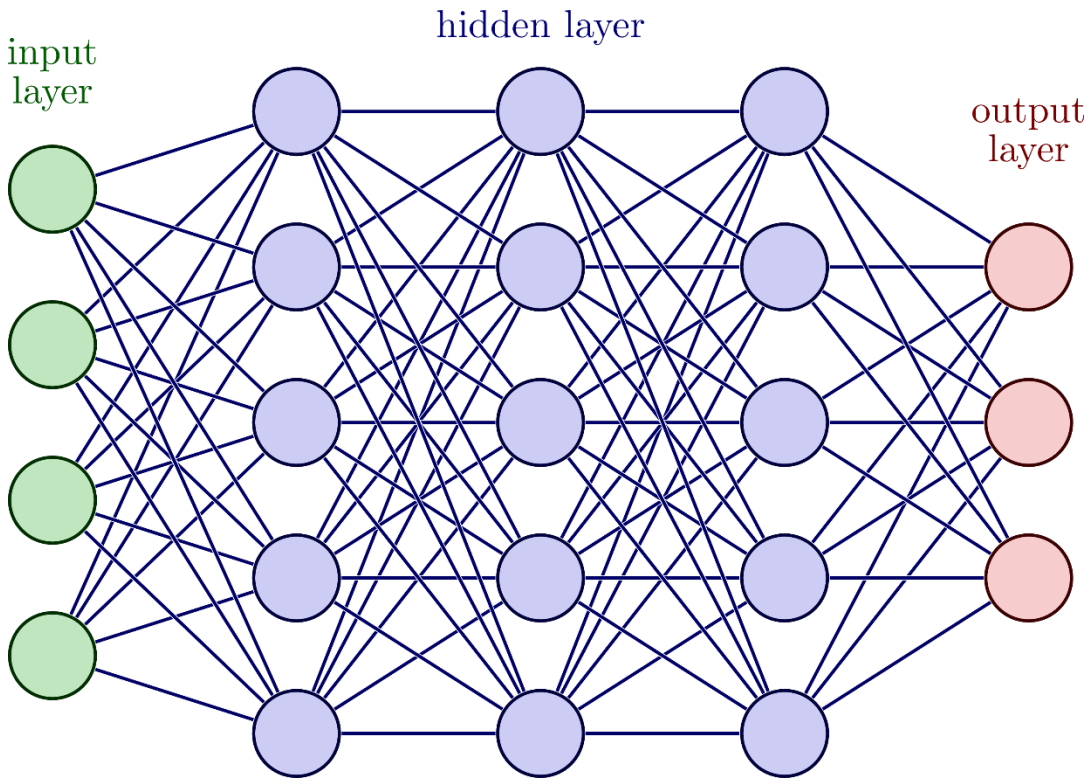

Result

Neural Network Learning

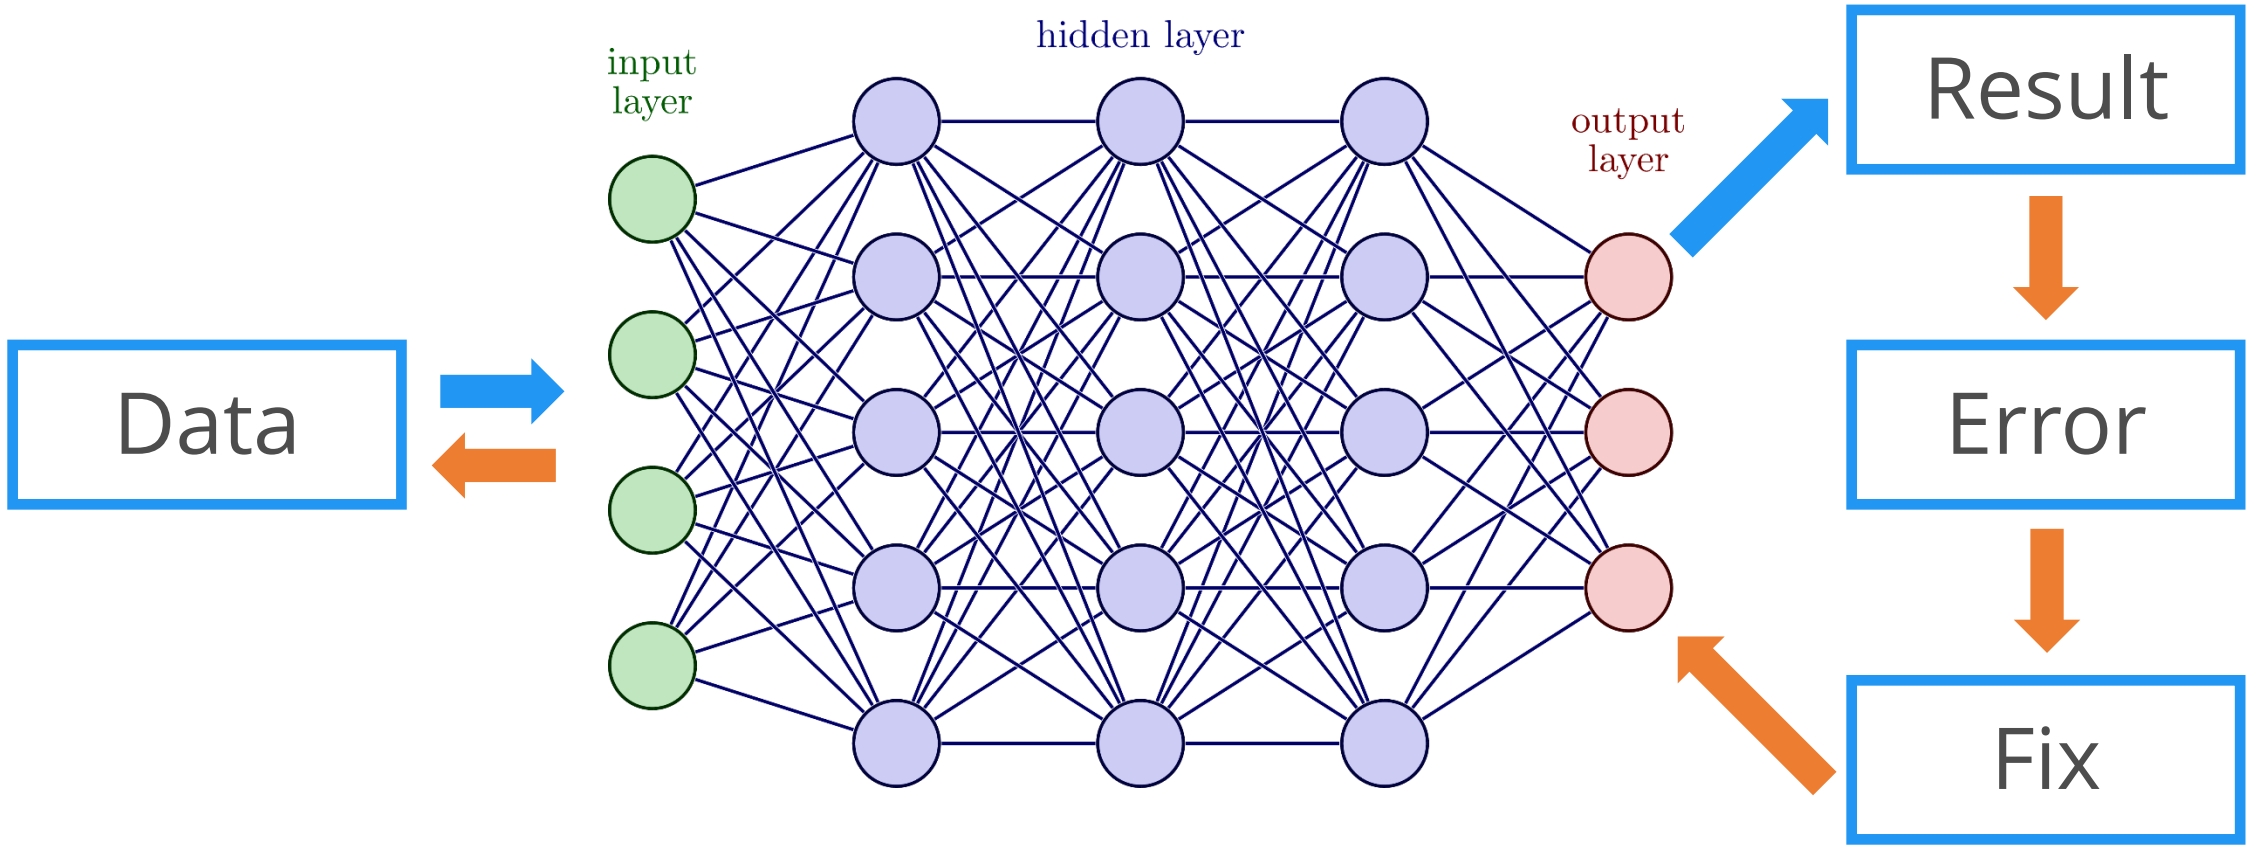

Convolution

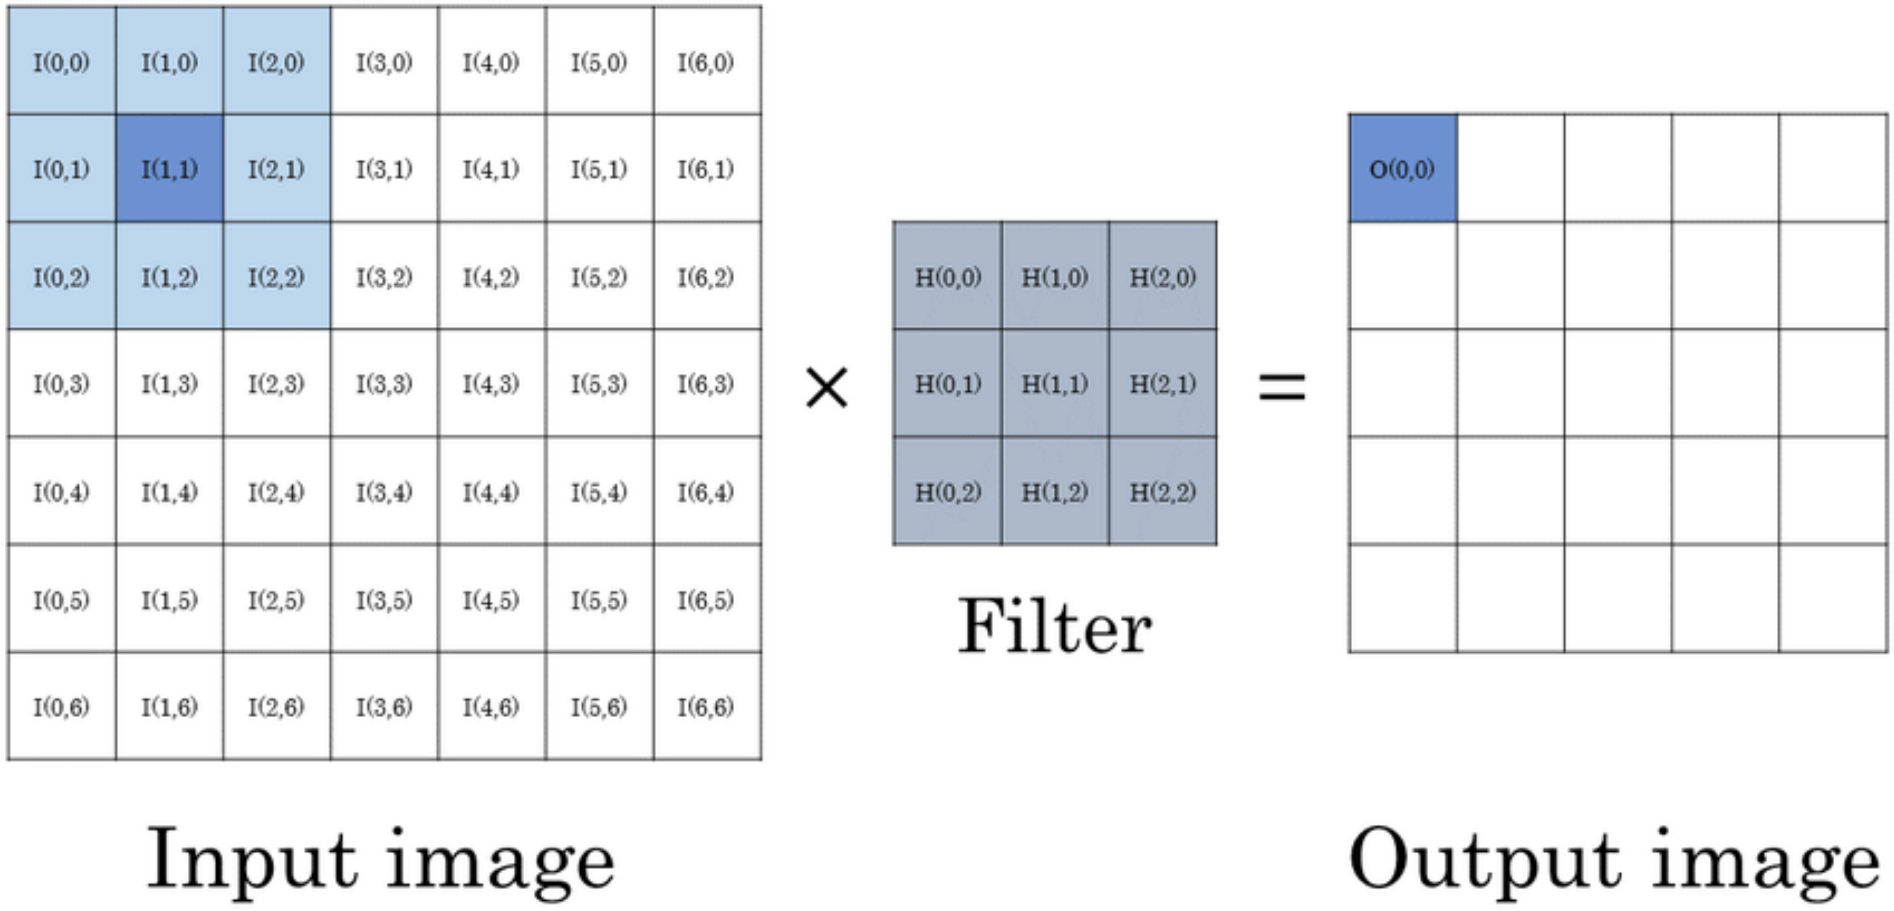

- How about *many* convolutions?

Convolution Layer

- Many channels
- Many filters (kernels)

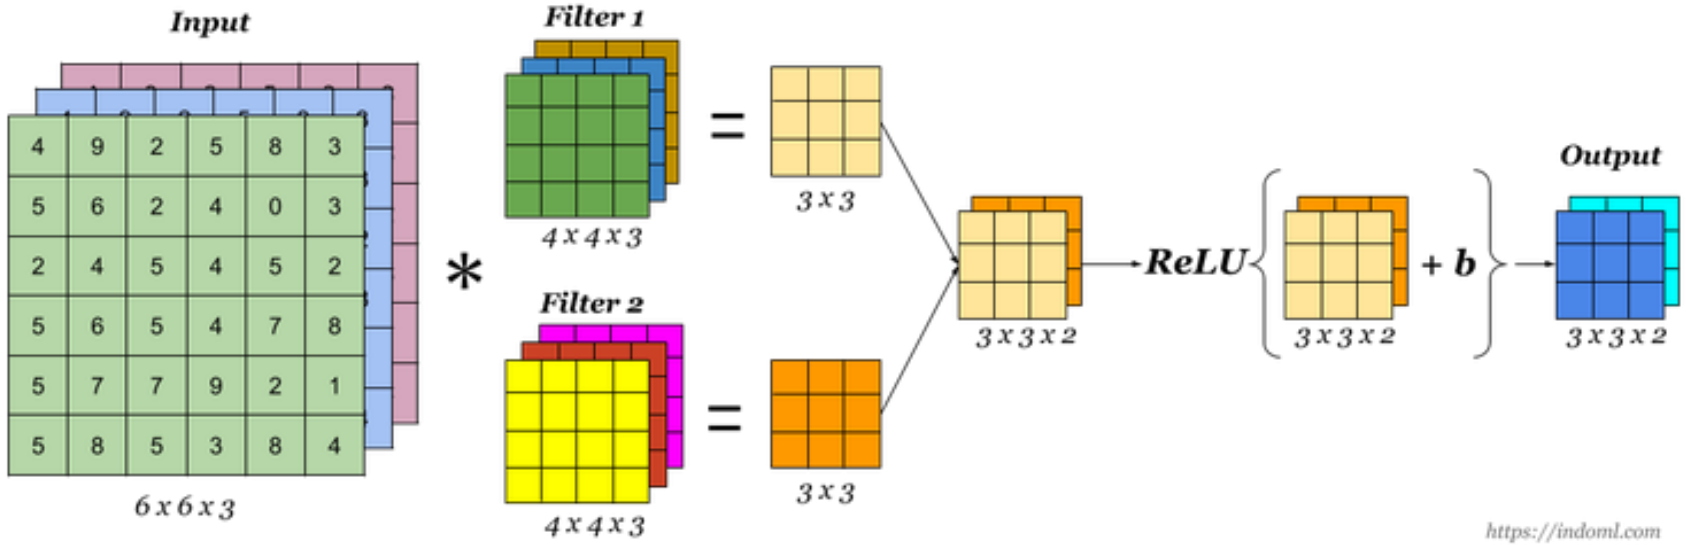

- Looks just like a "usual" NN!
  - But works perfectly on images!

A Complete Model

- Convolution layers
- Pooling layers
- "Standard" layers

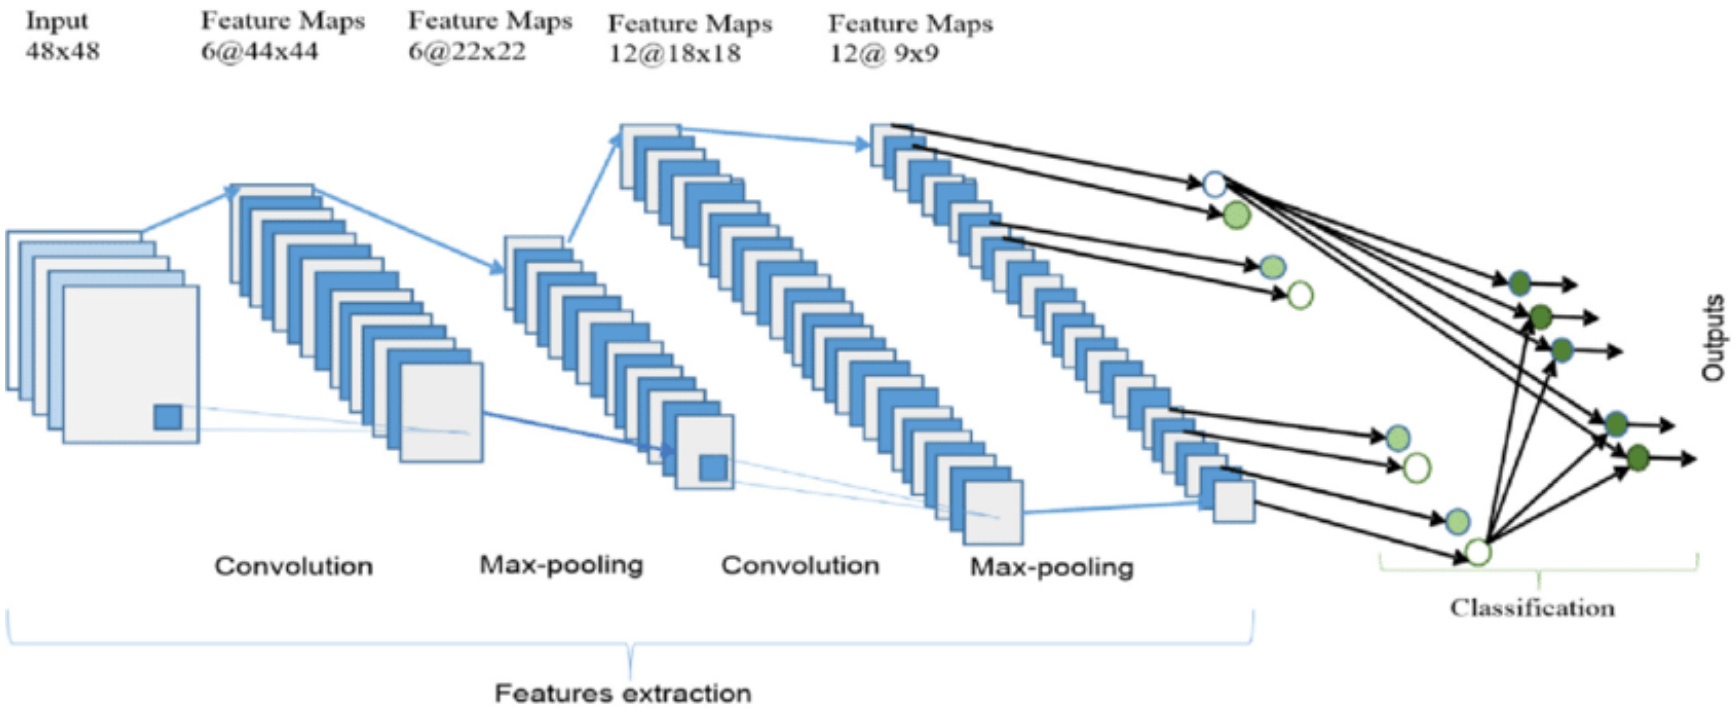

How Deep Can We Go?

- Inception ResNet V2
- Each rectangle is a layer

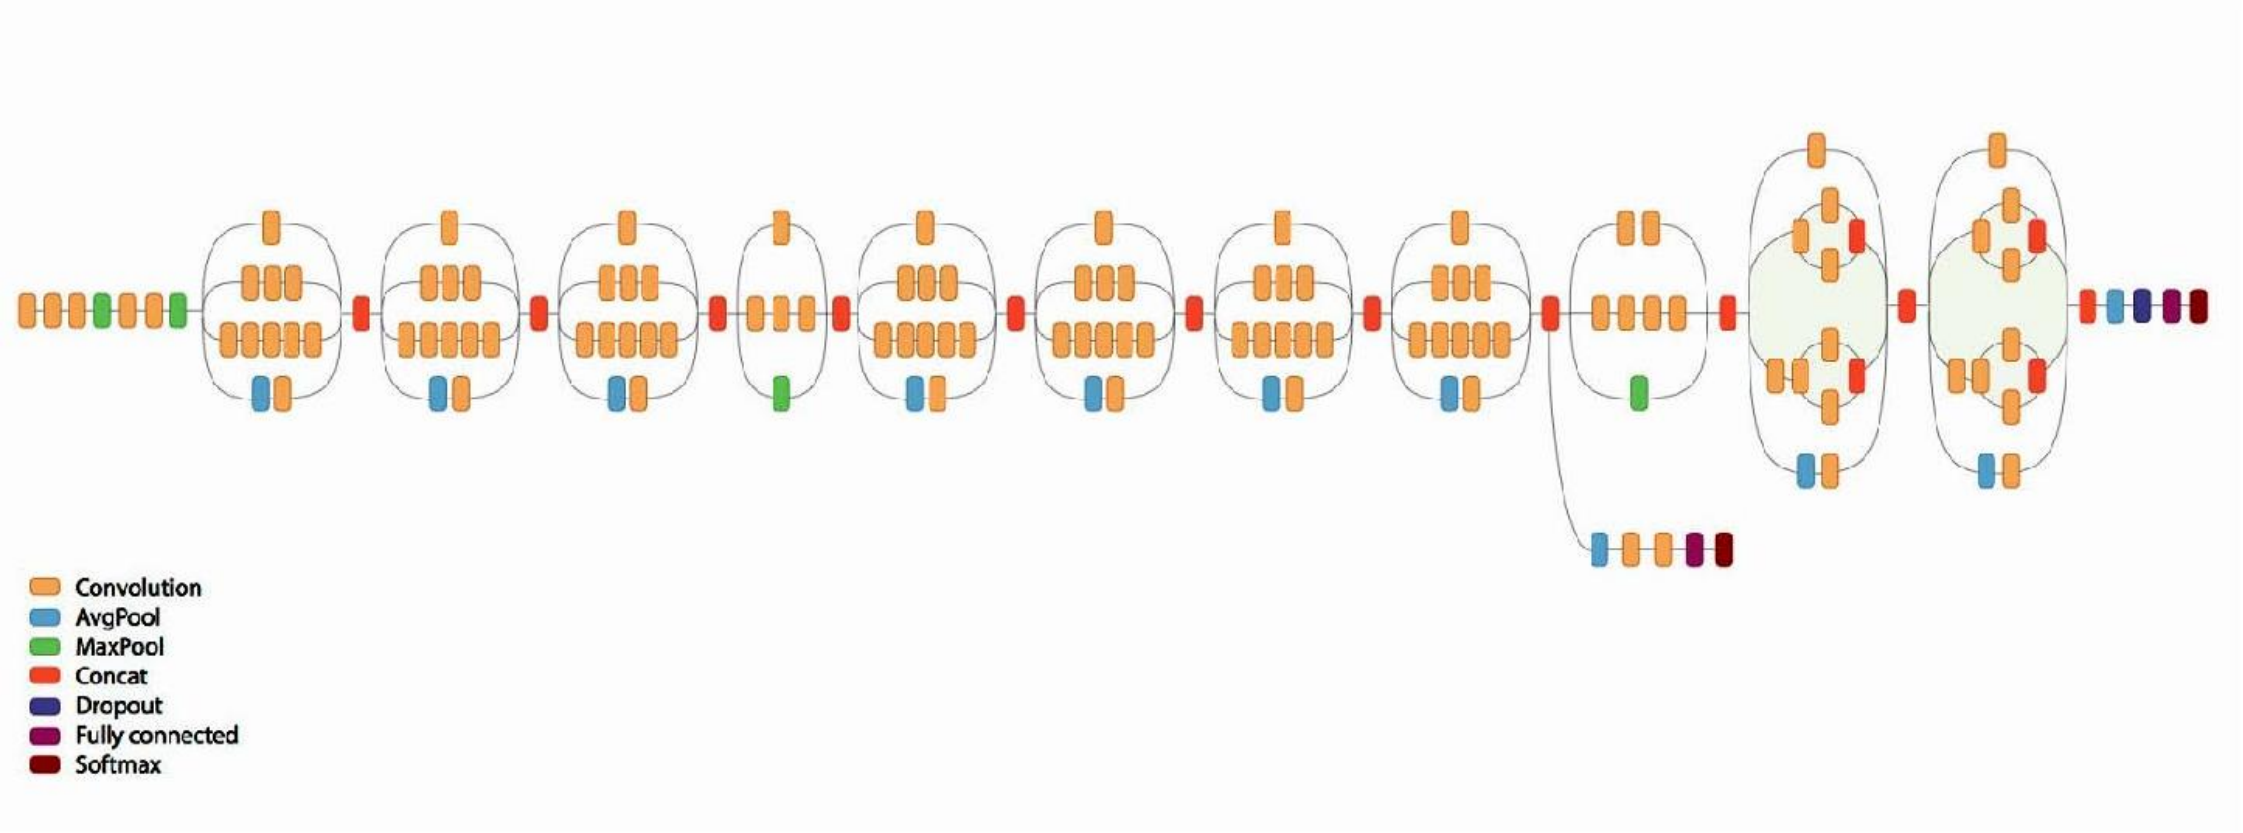

Fig. 1. Architecture of Inception Resnet V2 Network

How Deep Can We Go?

- Inception ResNet V2
- Each rectangle is a layer

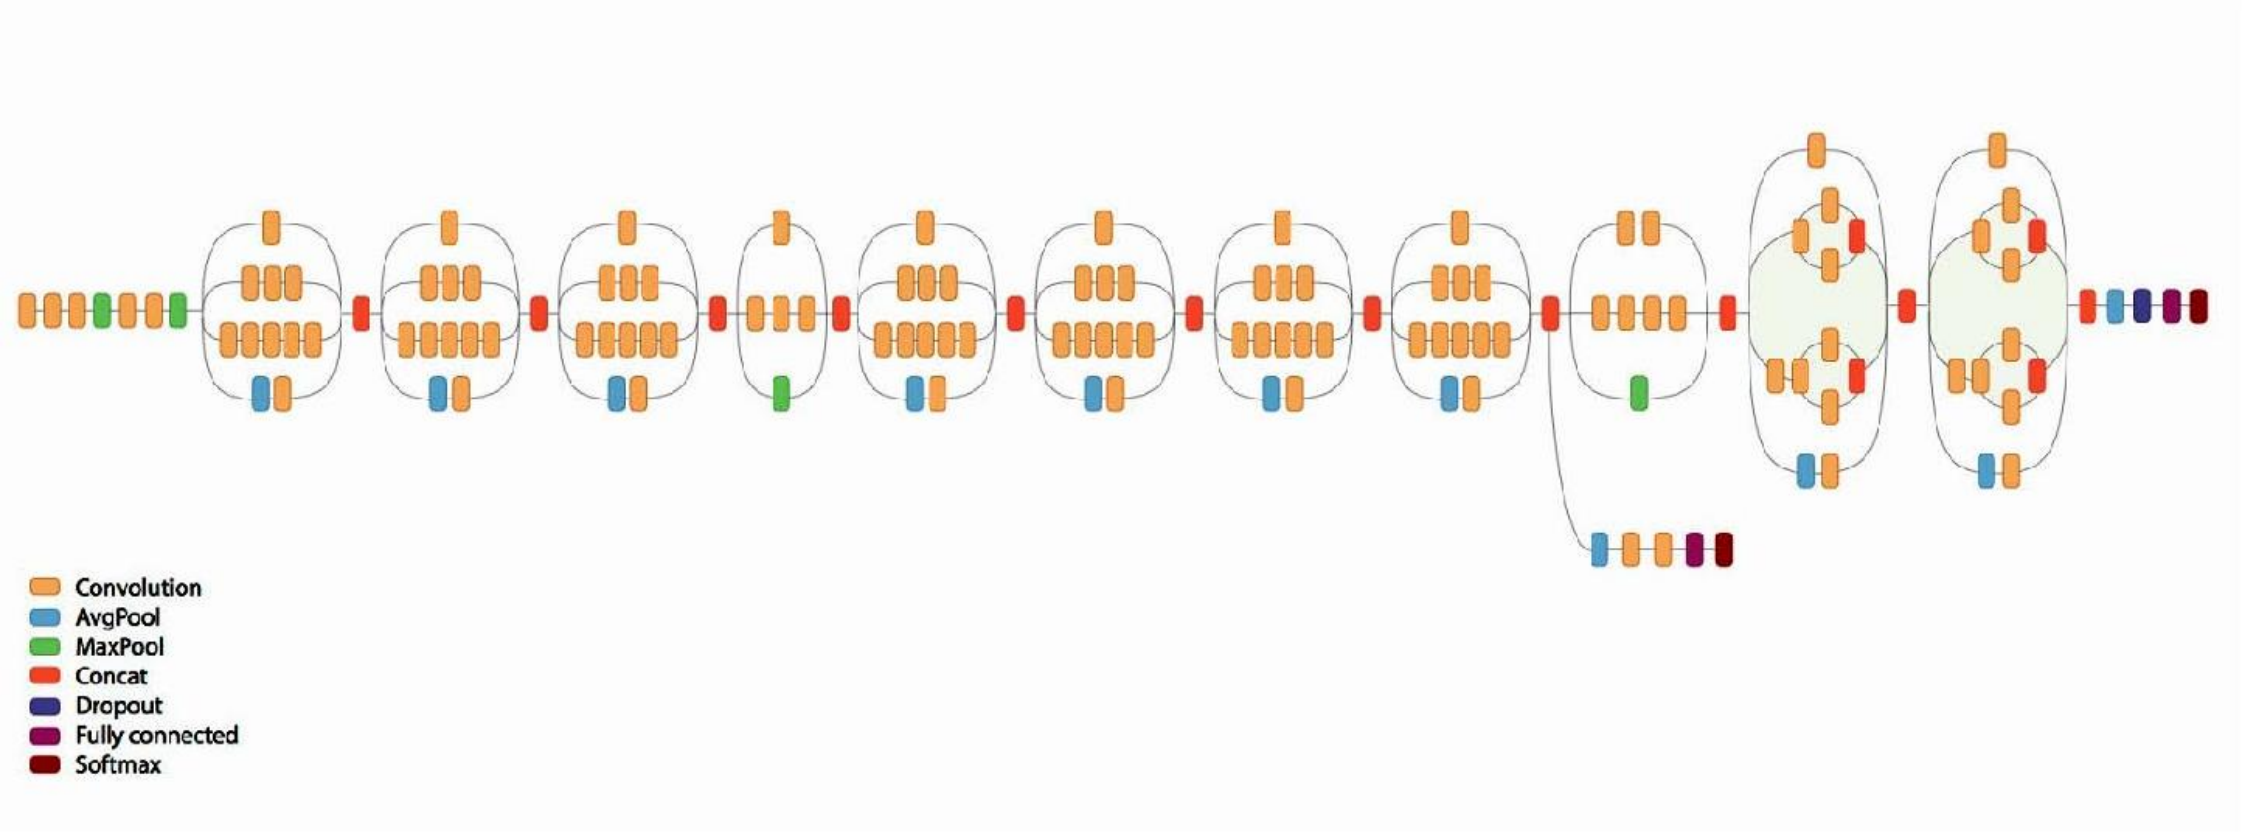

Fig. 1. Architecture of Inception Resnet V2 Network

The Gritty Details

- Activation function
- Hyperparameters
- Batch normalization
- Regularization
- Function optimization

- ... and other "-ations"

The Applications

- Image classification & friends

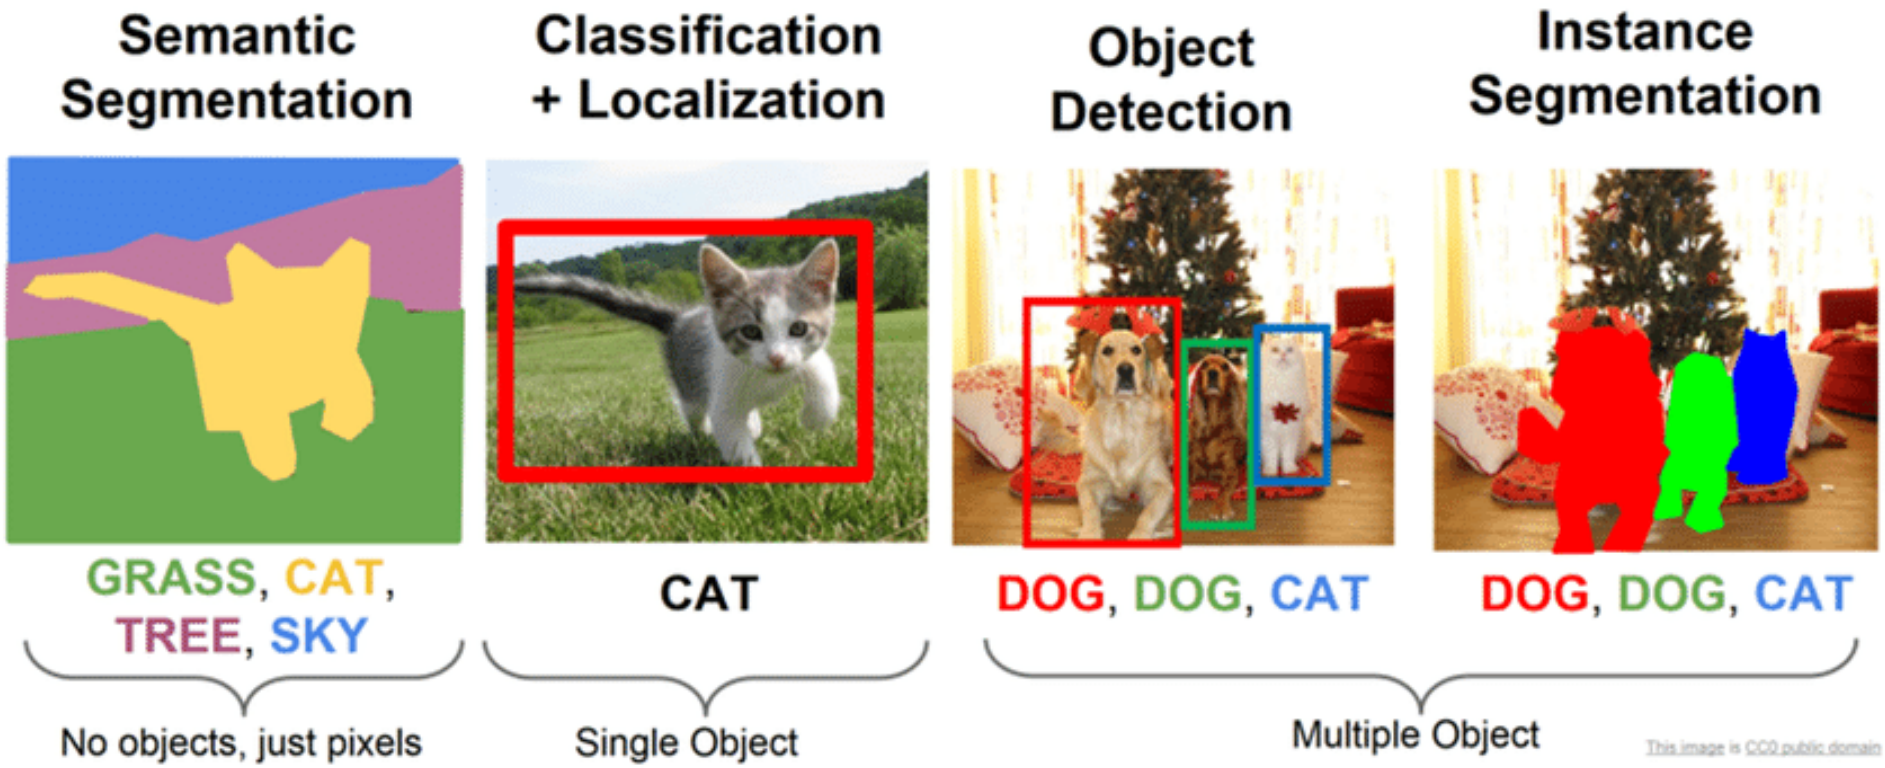

The Applications (2)

- Image captioning & generation

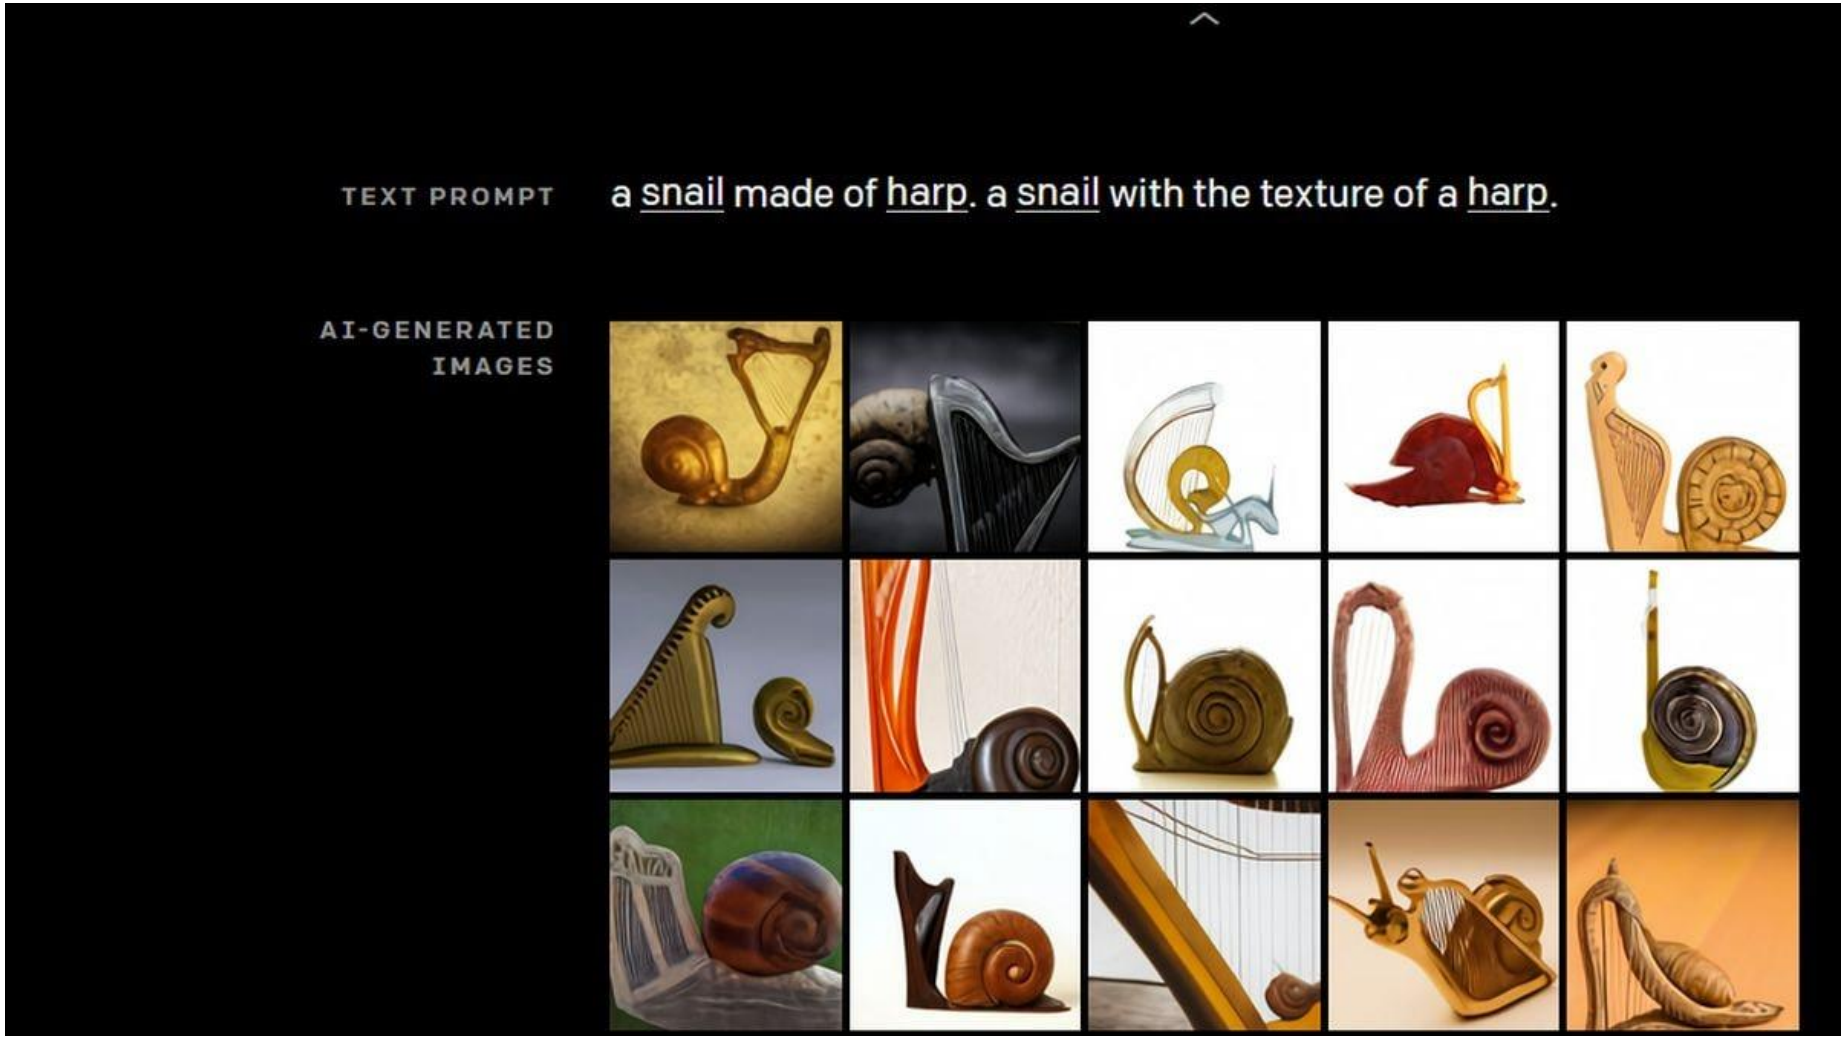

The Applications (3)

- Keypoint detection, pose estimation, activity recognition

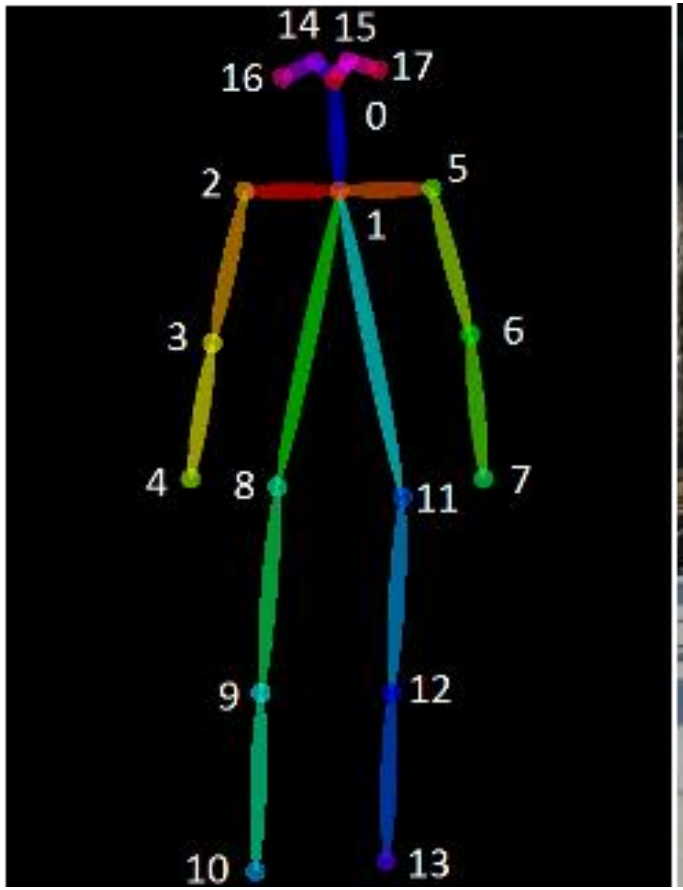

The Applications (4)

- Robotic control, autonomous vehicles

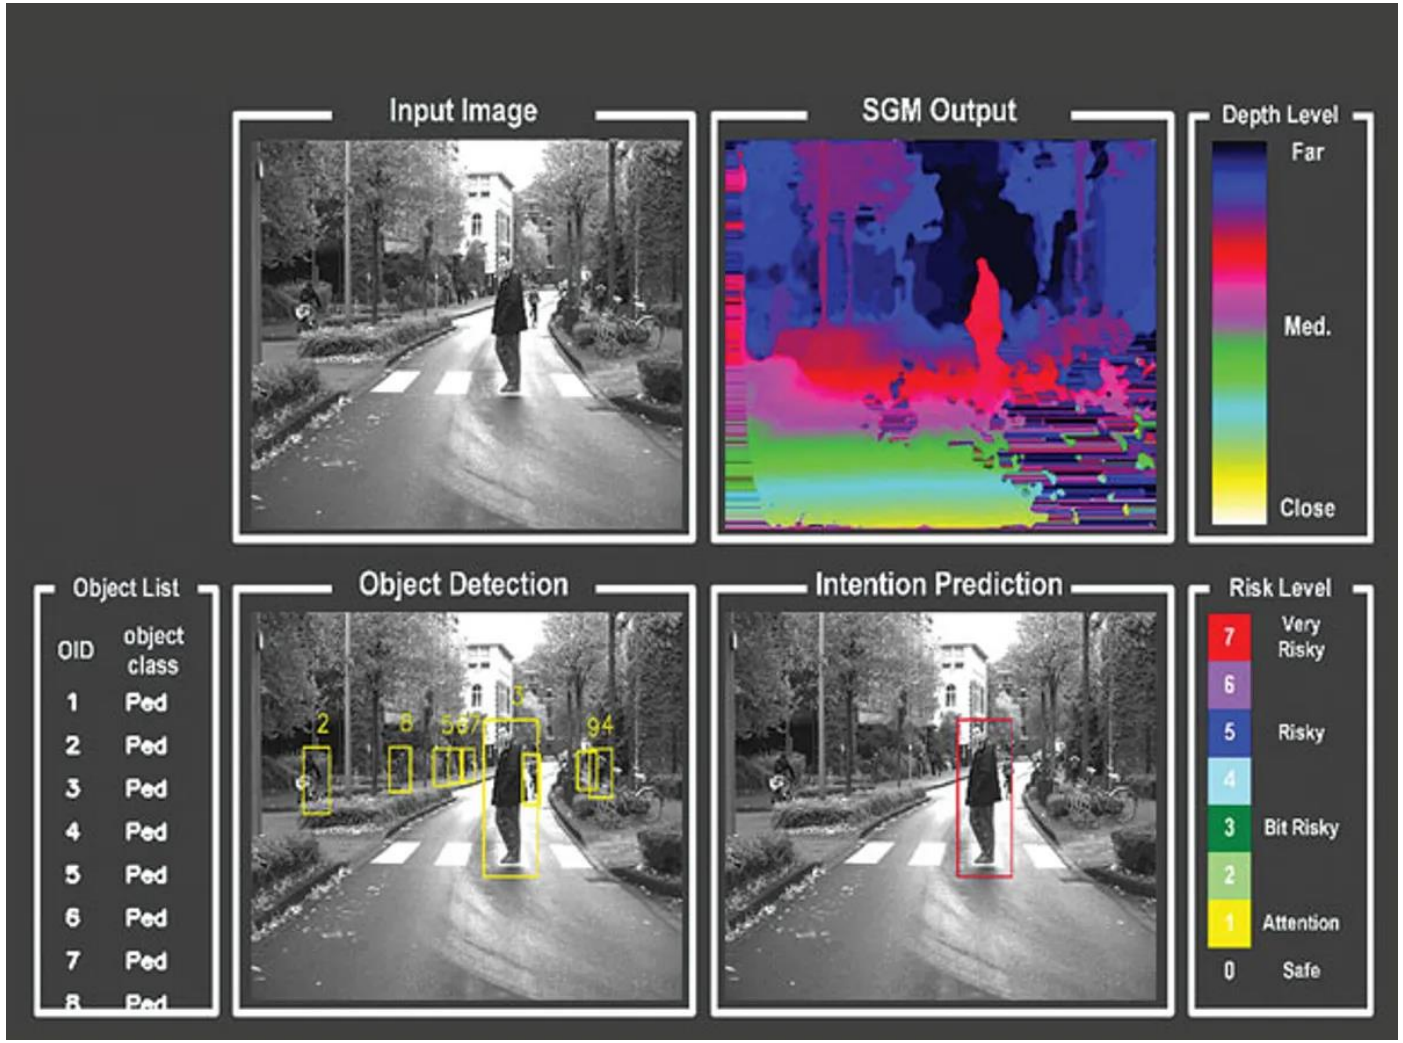

In Real Life

- Lots of data
- Data quality
- Finding biases
- Producing metrics
- Fine-tuning models

- Explainability
- Responsibility

Thank you!