Introduction to AI Engineering

Code Sync, London, 2018

Garrett Smith
Founder, Guild AI

HI

guild.ai

AI

Artificial General Intelligence (AGI)

Photographer: Joe Mehling

The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves.

J. McCarthy et al.; Aug. 31, 1955

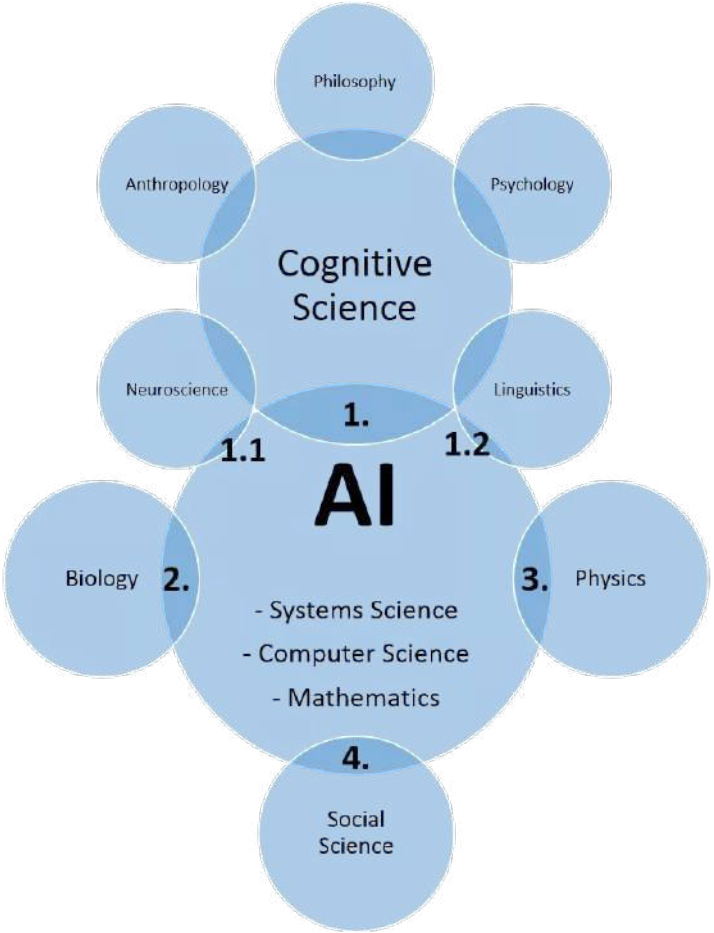

Source: *https://opentechai.blog/2017/11/10/navigating-in-the-ai-sciences-jungle/*

AI ≠ AGI

(for our purposes)

AI ≈ Intelligent Systems

“Hyper-automation”

Data Science

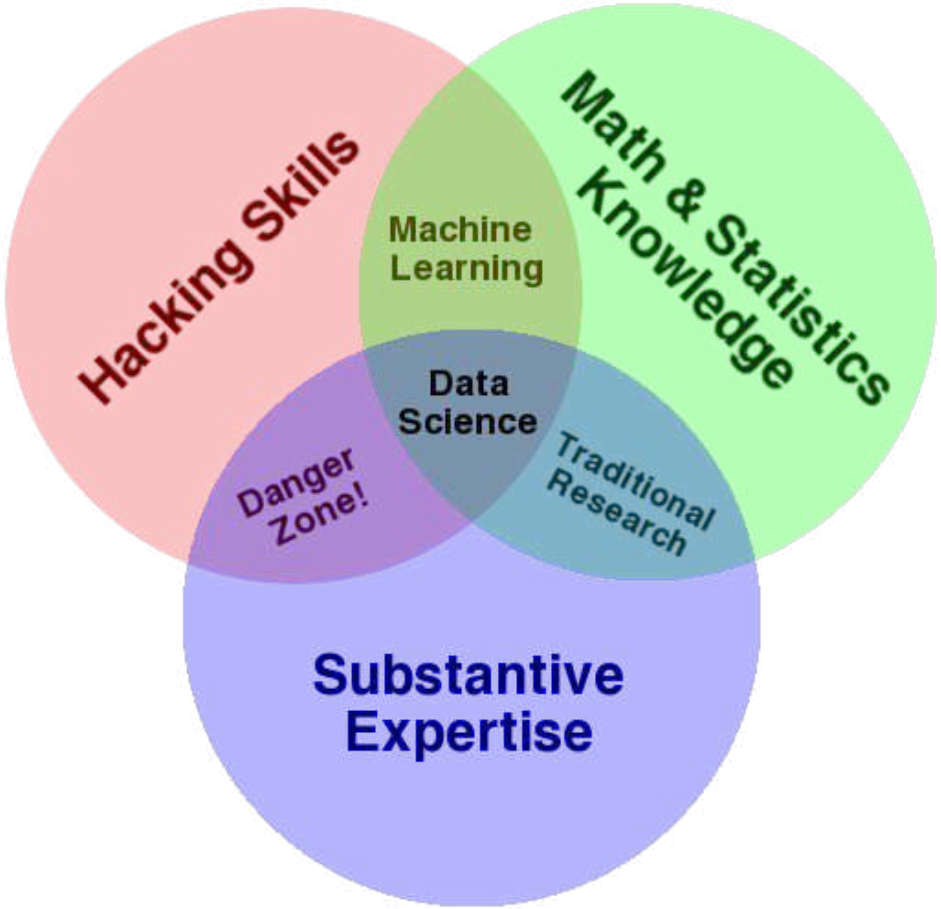

Source: *Conway, 2010*

Machine Learning

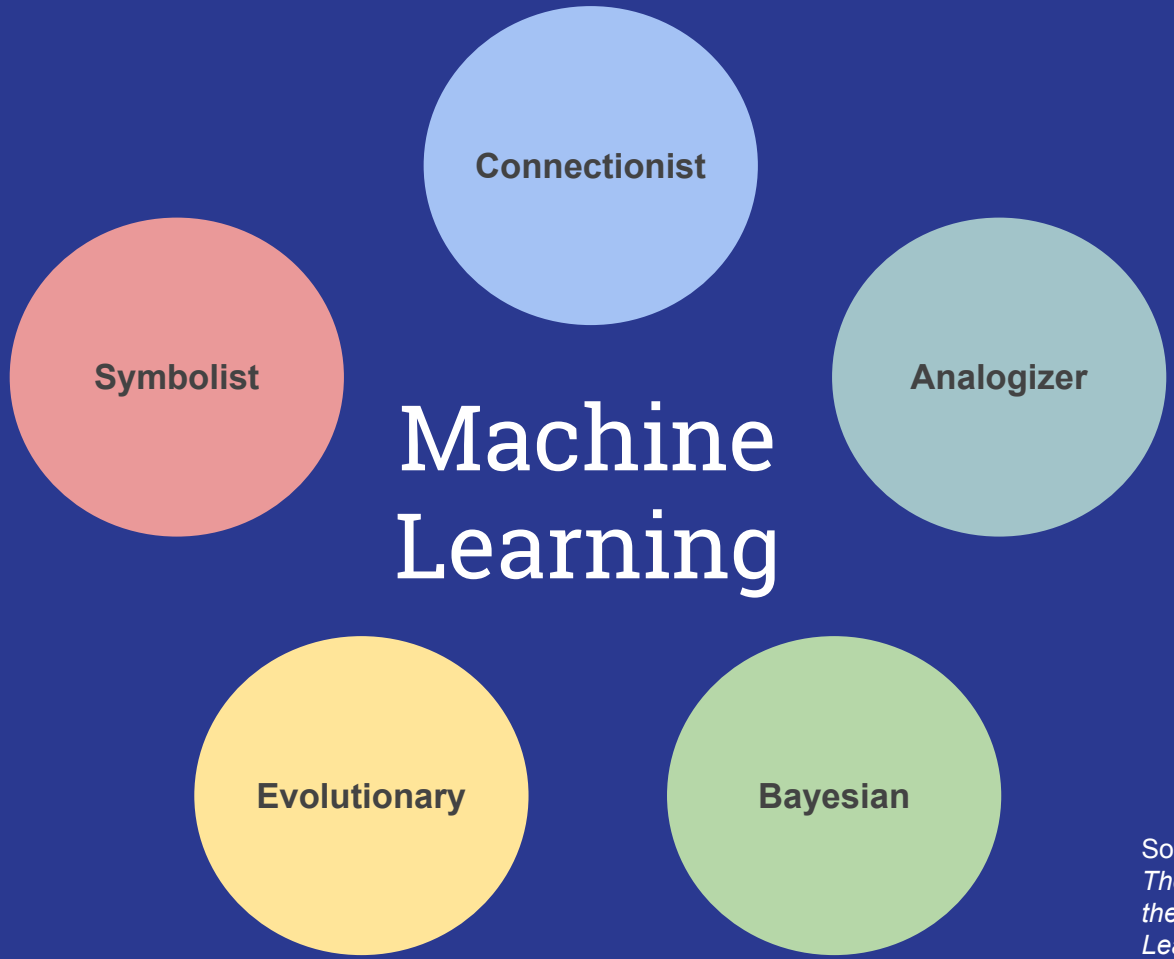

Source: Domingos, Pedro. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World.*

	Symbolist	Connectionist	Evolutionary	Bayesian	Analogizer
Representation	Logic	Neural Networks	Genetic Programs	Graphical Models	Support Vectors
Evaluation	Accuracy	Squared Error	Fitness	Posterior Probability	Margin
Optimization	Inverse Deduction	Gradient Descent	Generic Search	Probabilistic Inference	Constrained Optimization

Source: Domingos, Pedro. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World.*

arXiv - Artificial Intelligence (cs.AI)

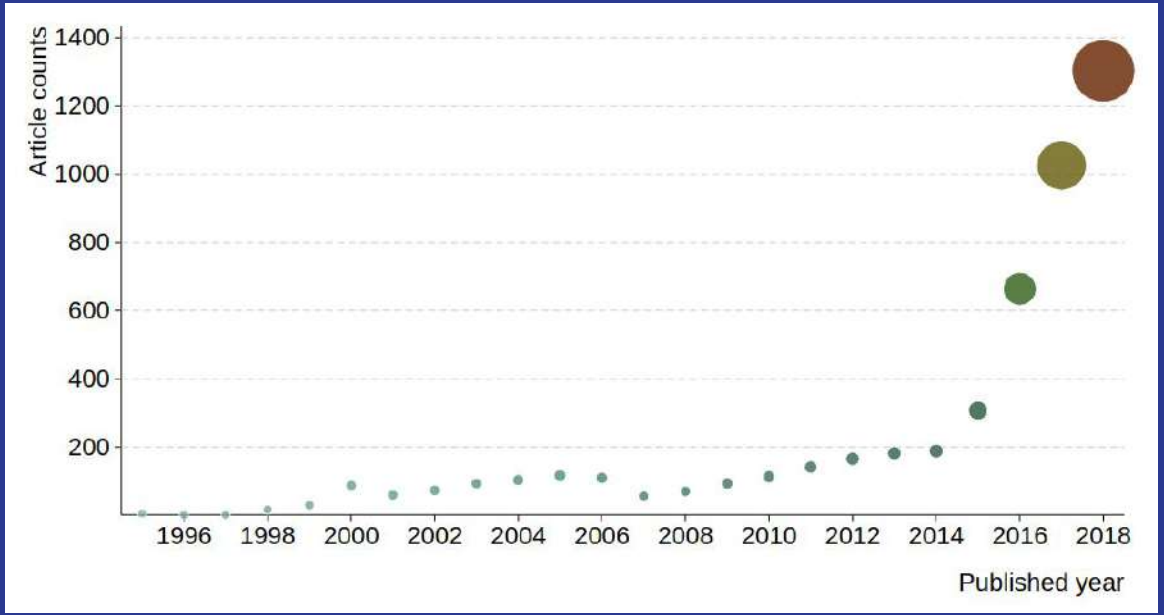

arXiv - Machine Learning (cs.LG)

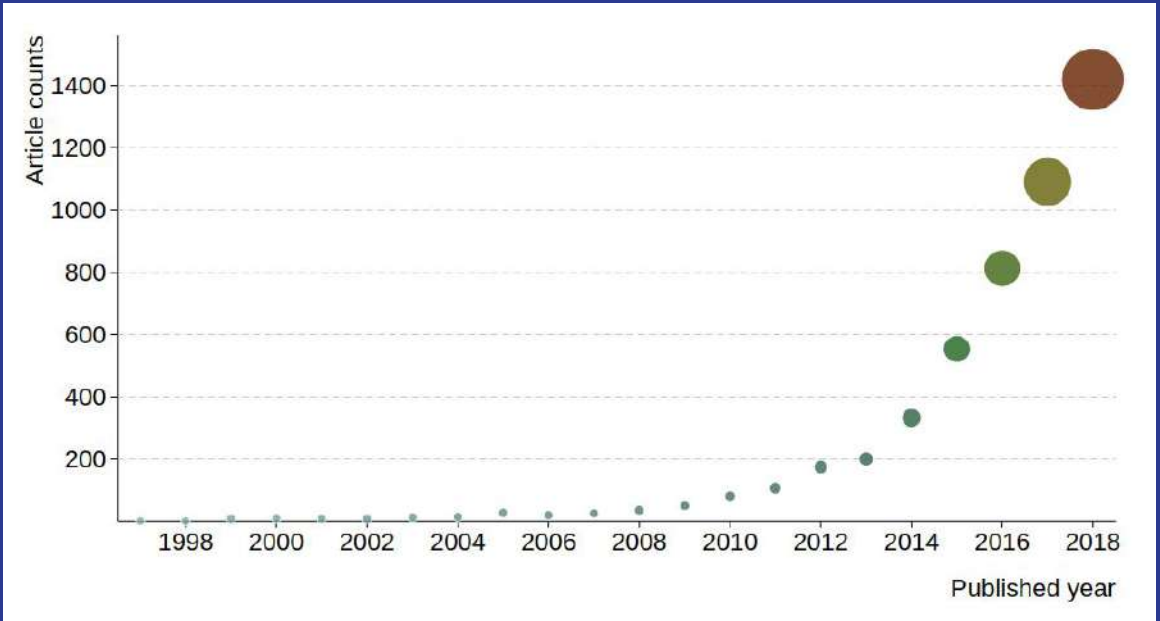

arXiv - Computer Vision and Pattern Recognition (cs.CV)

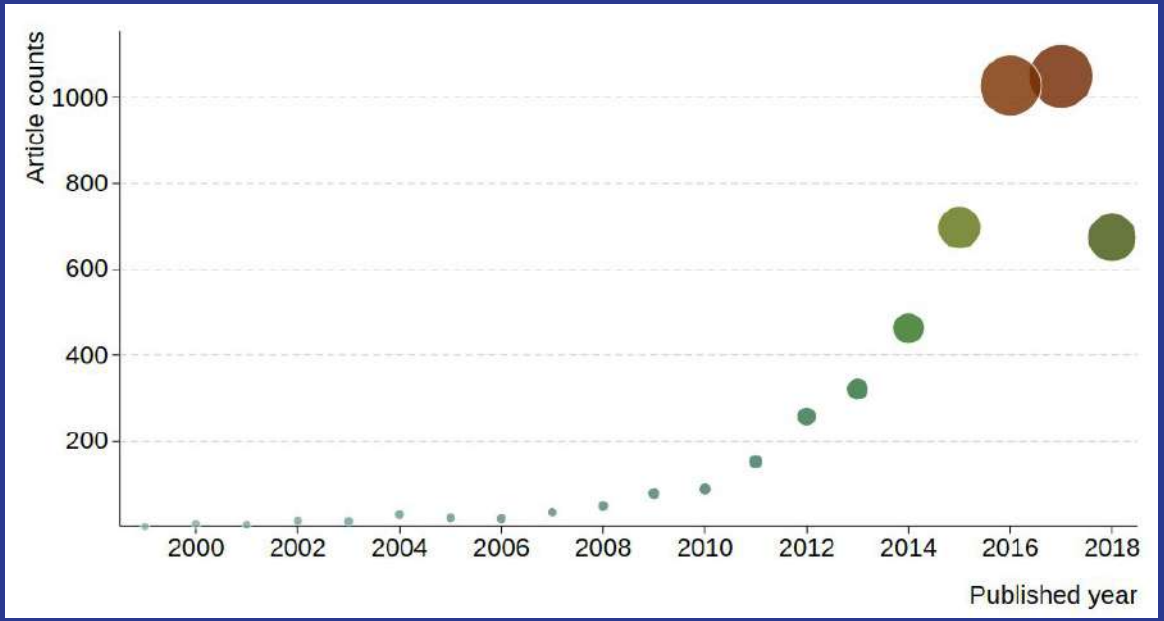

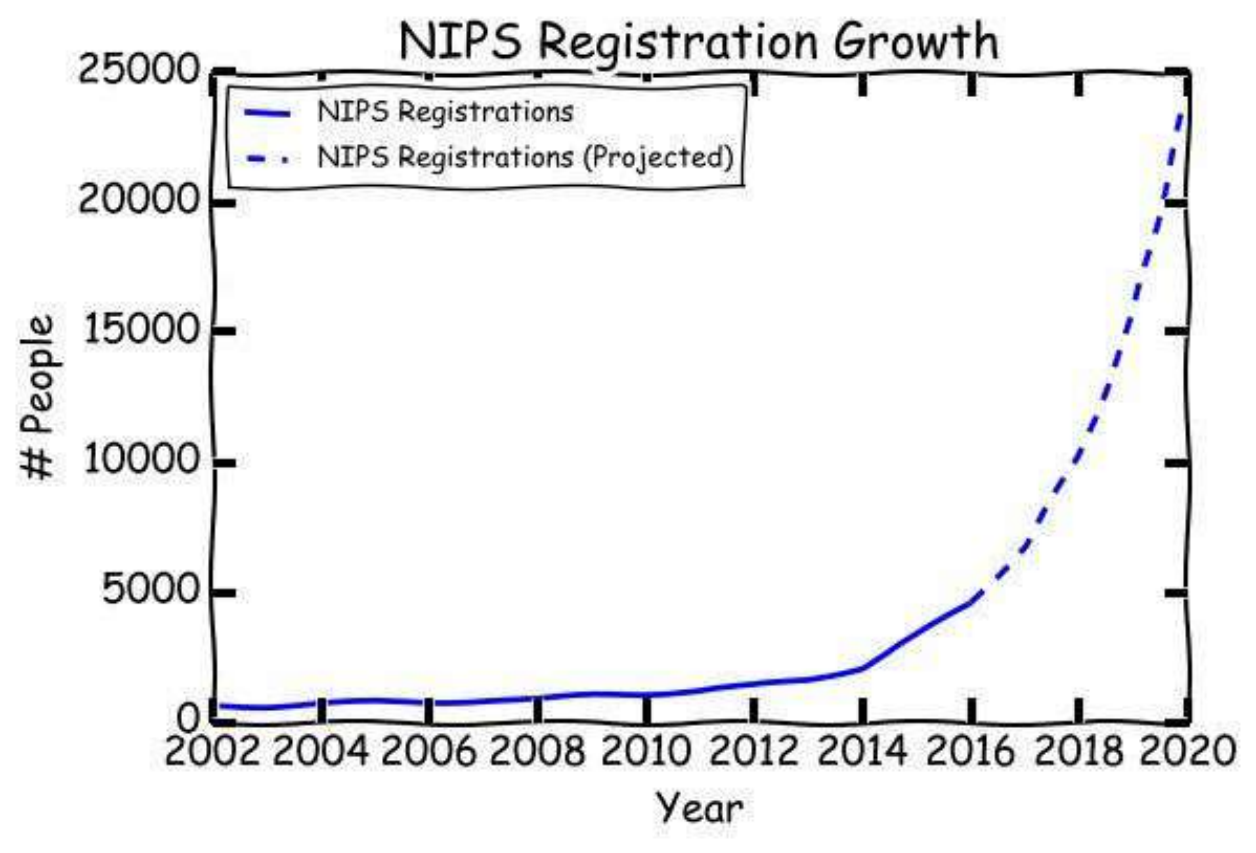

Source: *Ian Goodfellow*

Deep Learning

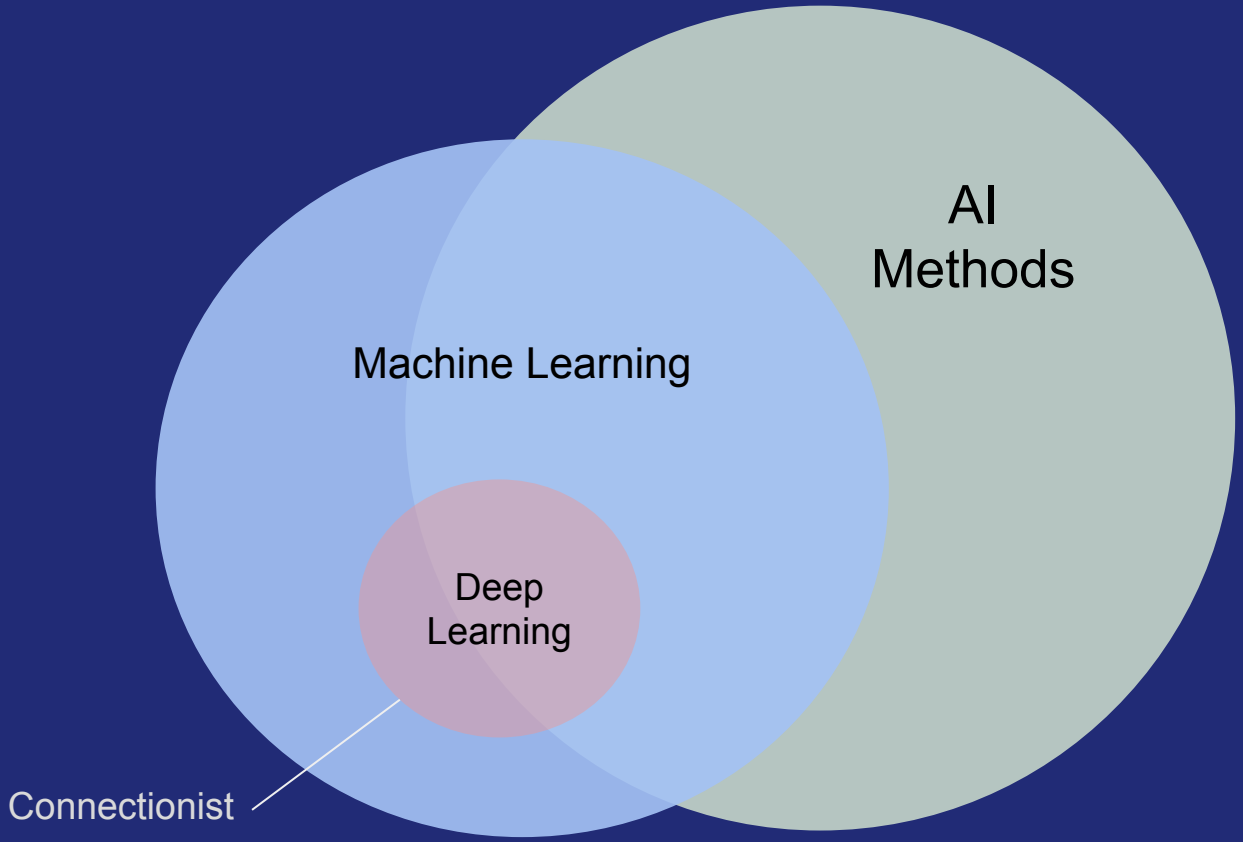
AI
Methods
Machine Learning
Deep
Learning
Connectionist

Original

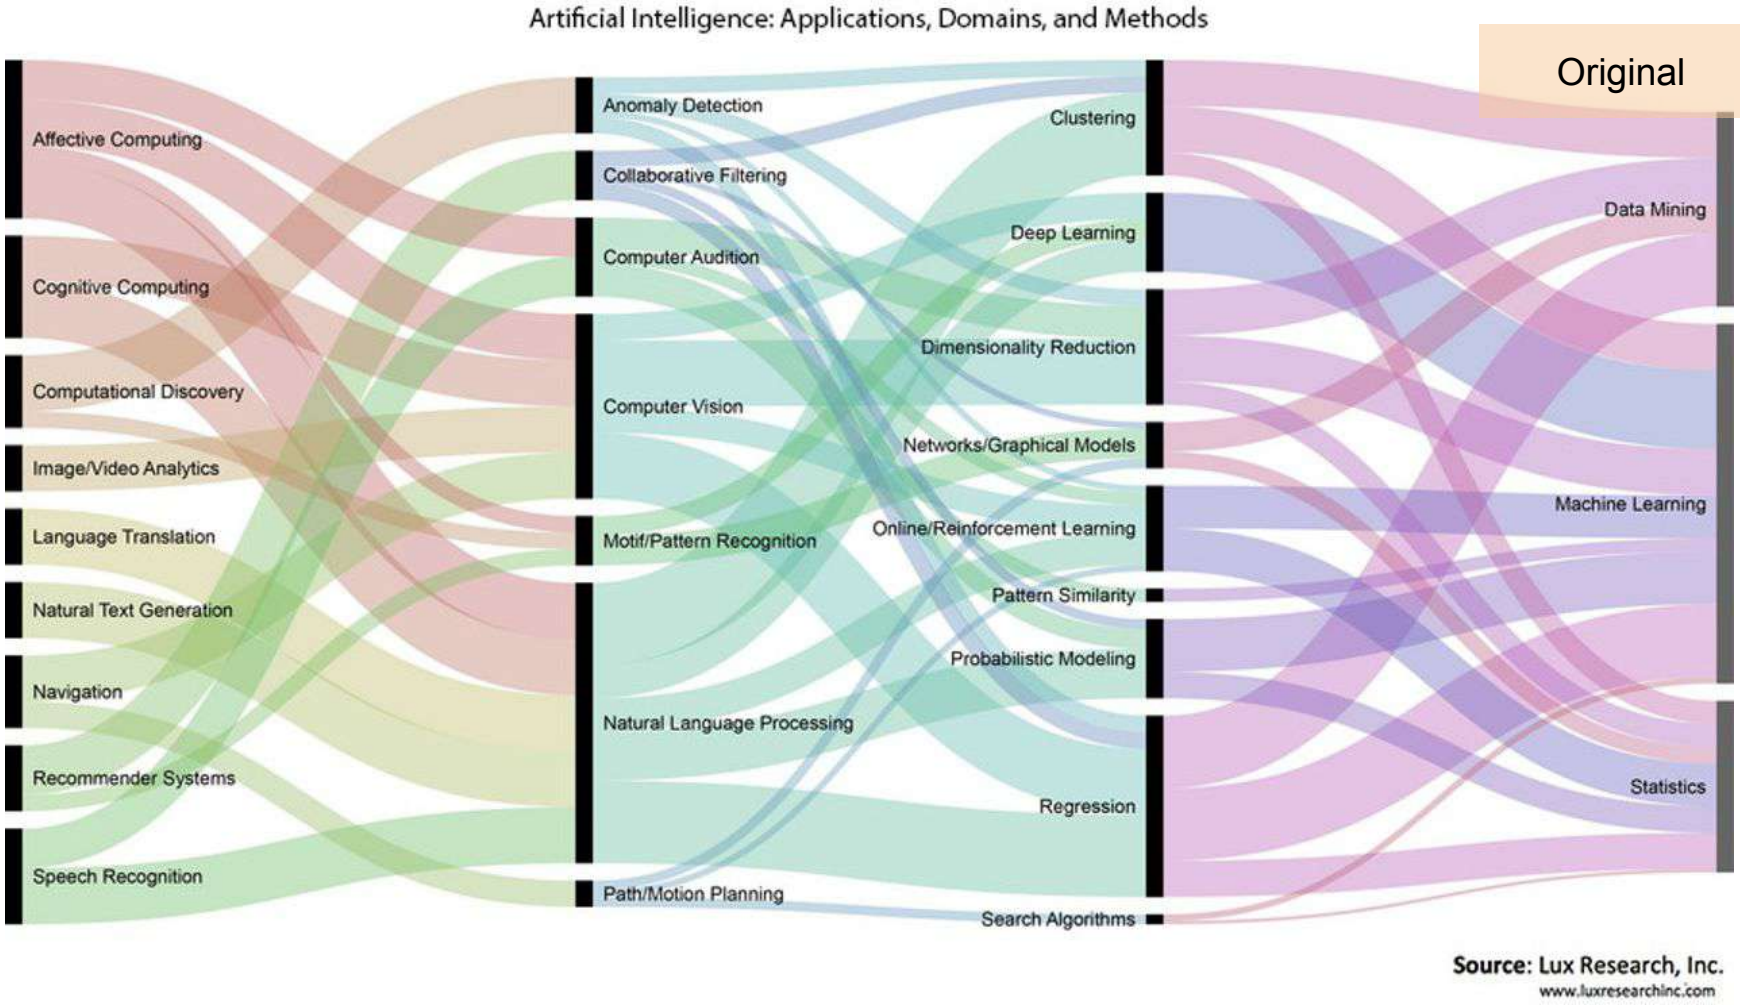
Artificial Intelligence: Applications, Domains, and Methods
Affective Computing
Cognitive Computing
Computational Discovery
Image/Video Analytics
Language Translation
Natural Text Generation
Navigation
Recommender Systems
Speech Recognition
Anomaly Detection
Collaborative Filtering
Computer Audition
Computer Vision
Motif/Pattern Recognition
Natural Language Processing
Path/Motion Planning
Clustering
Deep Learning
Dimensionality Reduction
Networks/Graphical Models
Online/Reinforcement Learning
Pattern Similarity
Probabilistic Modeling
Regression
Search Algorithms
Data Mining
Machine Learning
Statistics
Source: Lux Research, Inc.
www.luxresearchinc.com

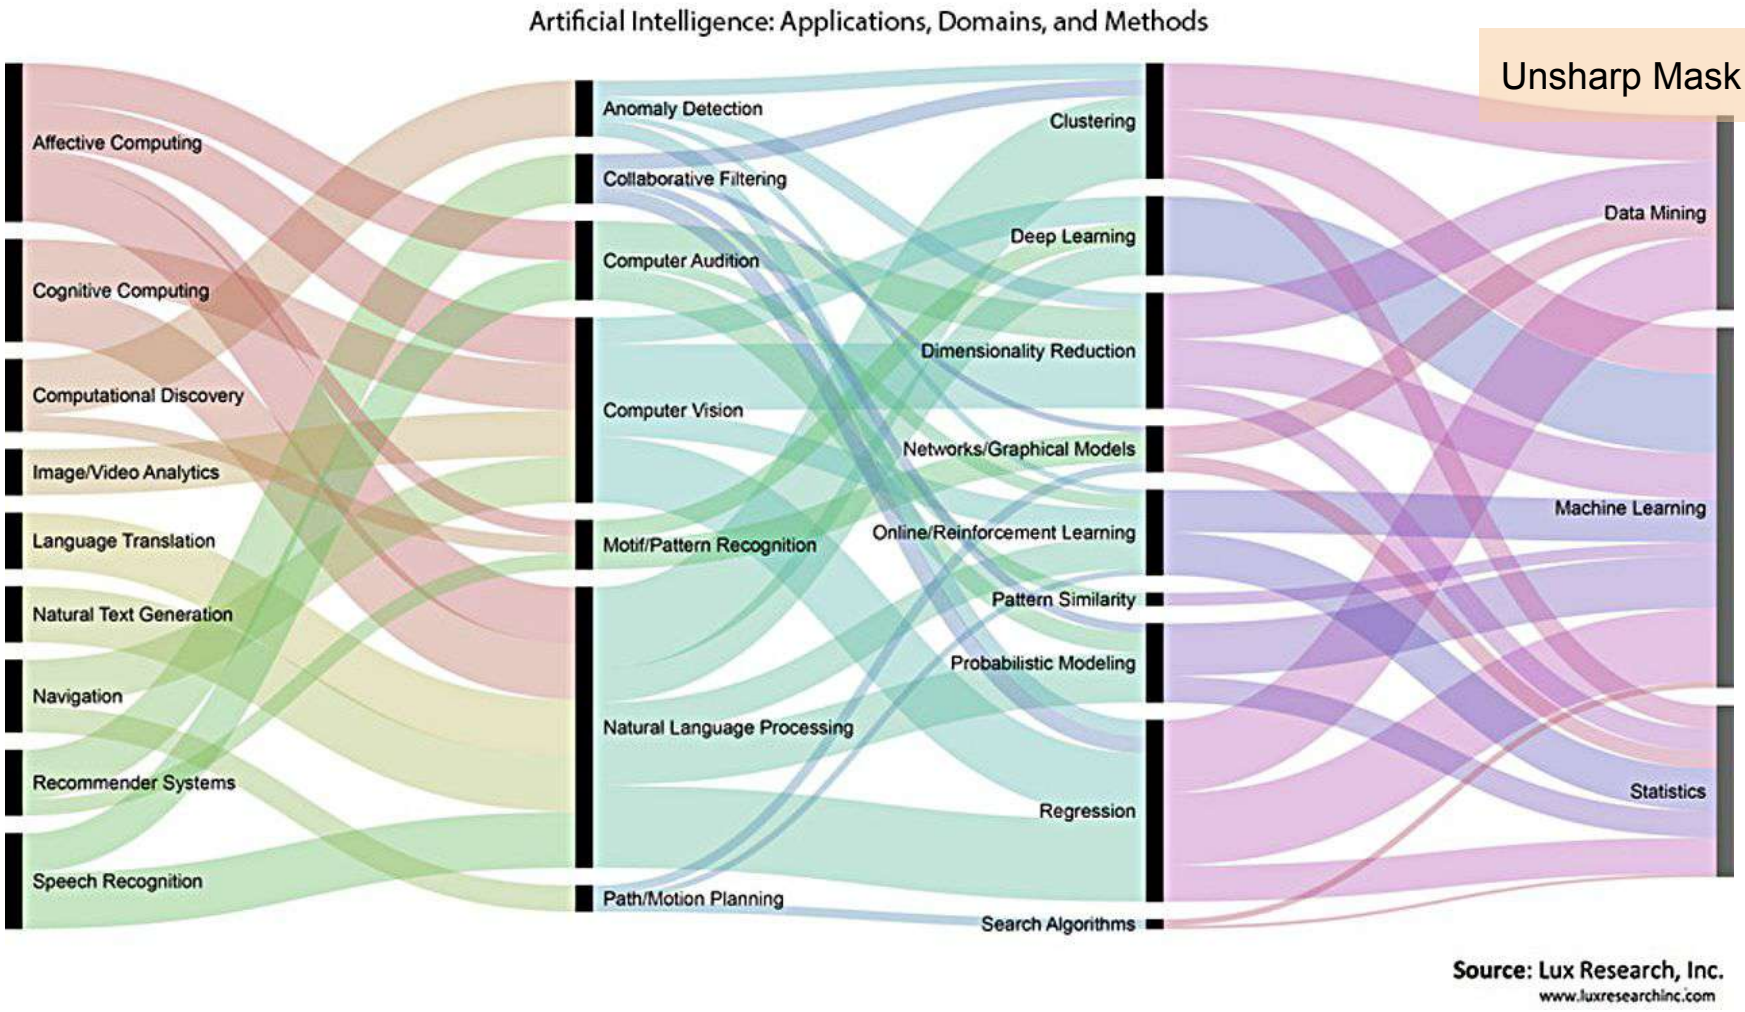
Artificial Intelligence: Applications, Domains, and Methods
Unsharp Mask
Affective Computing
Cognitive Computing
Computational Discovery
Image/Video Analytics
Language Translation
Natural Text Generation
Navigation
Recommender Systems
Speech Recognition
Anomaly Detection
Collaborative Filtering
Computer Audition
Computer Vision
Motif/Pattern Recognition
Natural Language Processing
Path/Motion Planning
Clustering
Deep Learning
Dimensionality Reduction
Networks/Graphical Models
Online/Reinforcement Learning
Pattern Similarity
Probabilistic Modeling
Regression
Search Algorithms
Data Mining
Machine Learning
Statistics
Source: Lux Research, Inc.
www.luxresearchinc.com

Neural Enhance

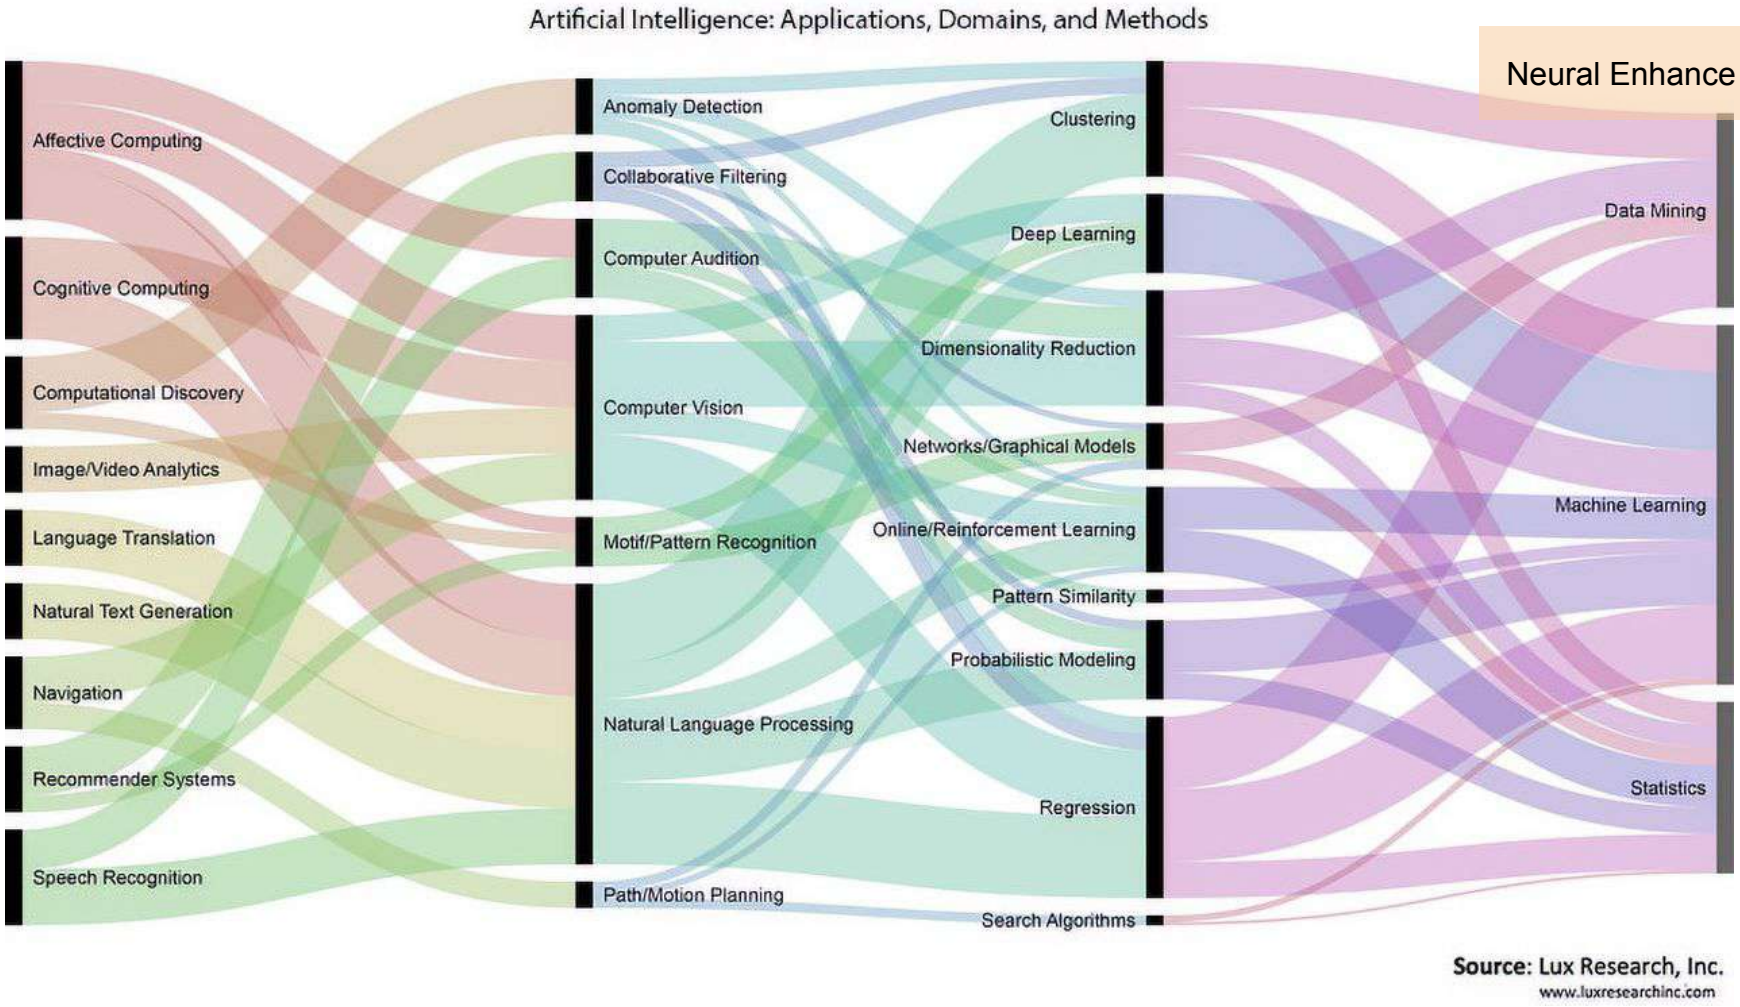

alexjc/neural-enhance on GitHub

README.rst

Neural Enhance

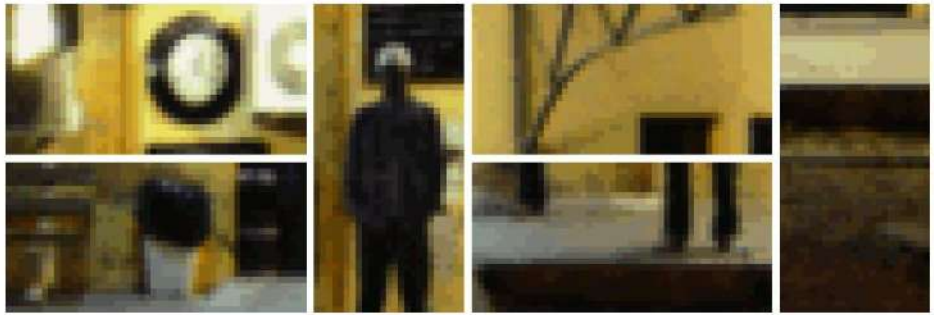

Example #1 — Old Station: view comparison in 24-bit HD, original photo CC-BY-SA @siv-athens.

As seen on TV! What if you could increase the resolution of your photos using technology from CSI laboratories? Thanks to deep learning and `#NeuralEnhance`, it's now possible to train a neural network to zoom in to your images at 2x or even 4x. You'll get even better results by increasing the number of neurons or training with a dataset similar to your low resolution image.


PICCADILLY
PICCADILLY

Methodology

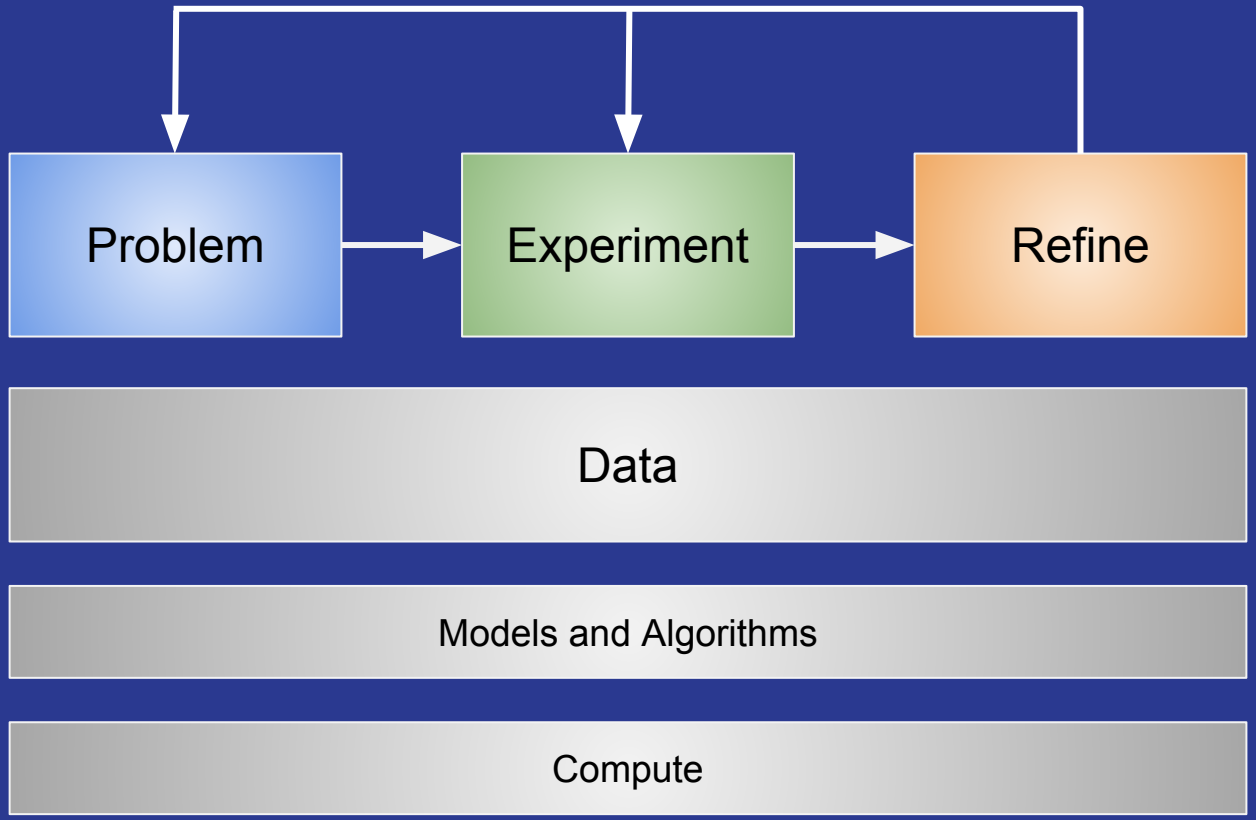
Problem
Experiment
Refine
Data
Models and Algorithms
Compute

Three roles

Research scientist

Research engineer

Software/system engineer

Content funnel

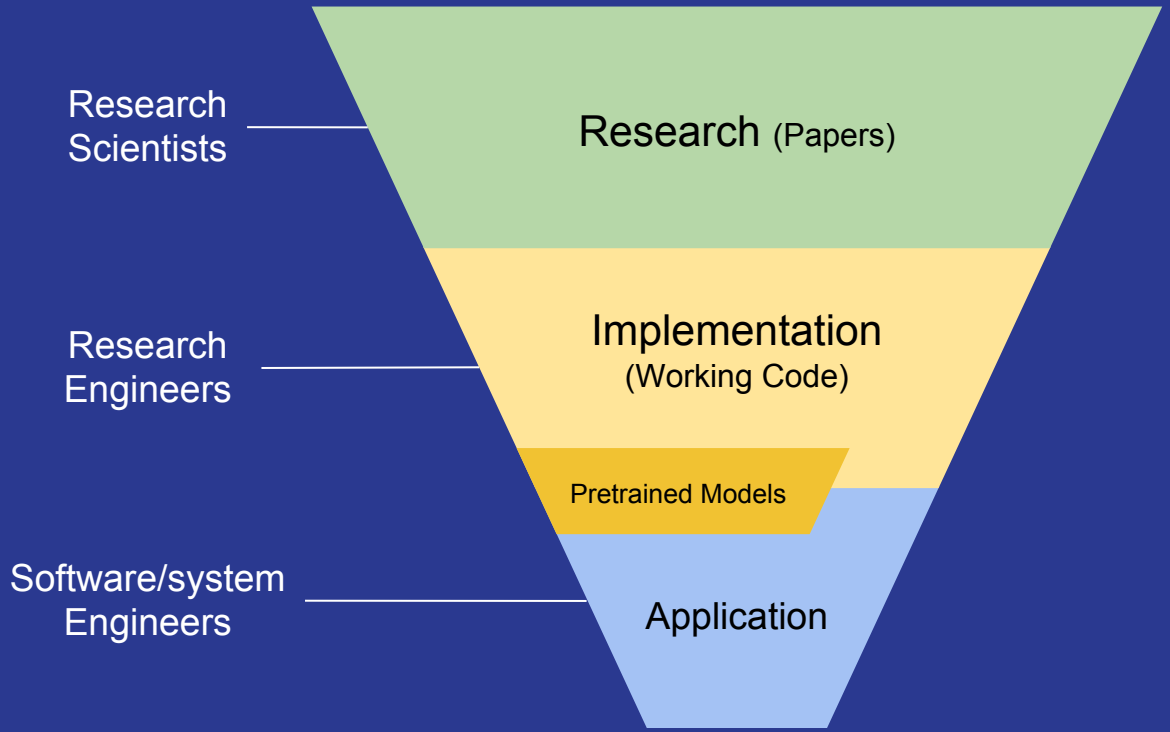

Best Paper - CVPR 2018

Taskonomy: Disentangling Task Transfer Learning

Amir R. Zamir[1,2] Alexander Sax[1*] William Shen[1*] Leonidas Guibas[1] Jitendra Malik[2] Silvio Savarese[1]

[1] Stanford University [2] University of California, Berkeley

http://taskonomy.vision/

Abstract

Do visual tasks have a relationship, or are they unrelated? For instance, could having surface normals simplify estimating the depth of an image? Intuition answers these questions positively, implying existence of a ***structure*** *among visual tasks. Knowing this structure has notable values; it is the concept underlying transfer learning and provides a principled way for identifying redundancies across tasks, e.g., to seamlessly reuse supervision among related tasks or solve many tasks in one system without piling up the complexity.*

We proposes a fully computational approach for modeling the structure of space of visual tasks. This is done via finding (first and higher-order) transfer learning dependencies across a dictionary of twenty six 2D, 2.5D, 3D, and semantic tasks in a latent space. The product is a computational taxonomic map for task transfer learning. We study the consequences of this structure, e.g. nontrivial emerged

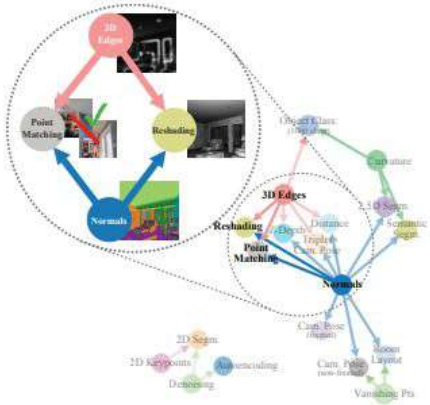

StanfordVL/taskonomy on GitHub

README.md

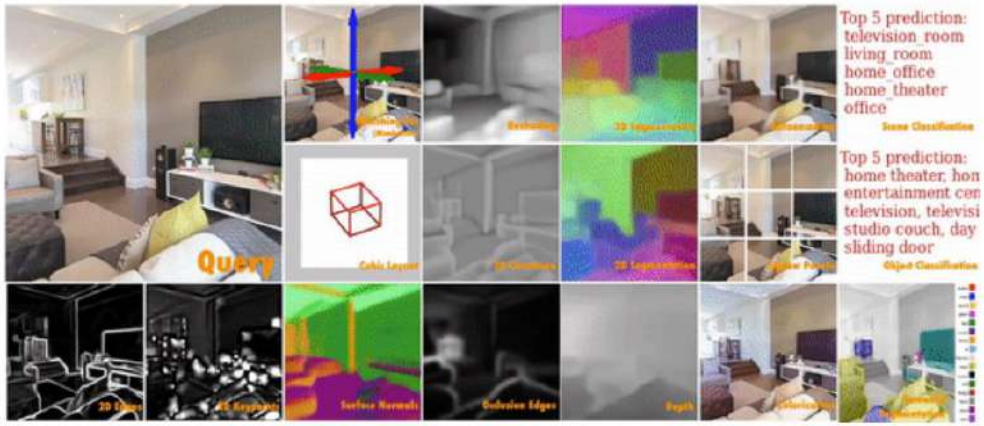

TASK BANK: A Unified Bank of 25 Visual Estimators

This repository shares a unified bank of pretrained models for **25 vision tasks spanning a wide range of 2D, 3D, and semantic problems**. Given a query image, the produced 25 estimations give a broad visual understanding useful for different

http://taskonomy.stanford.edu/tasks/

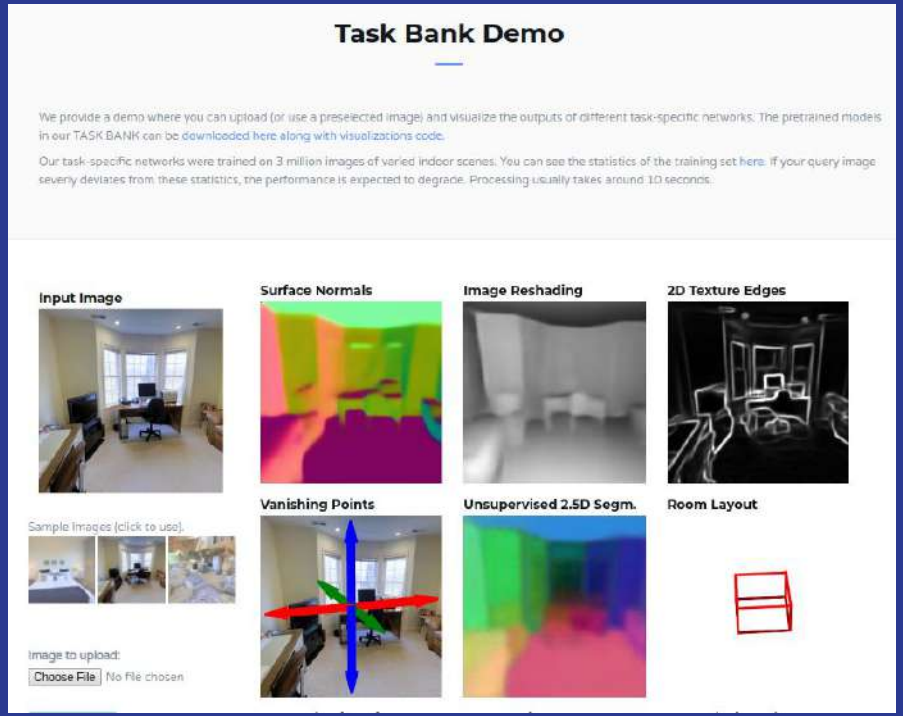

Languages and Libraries

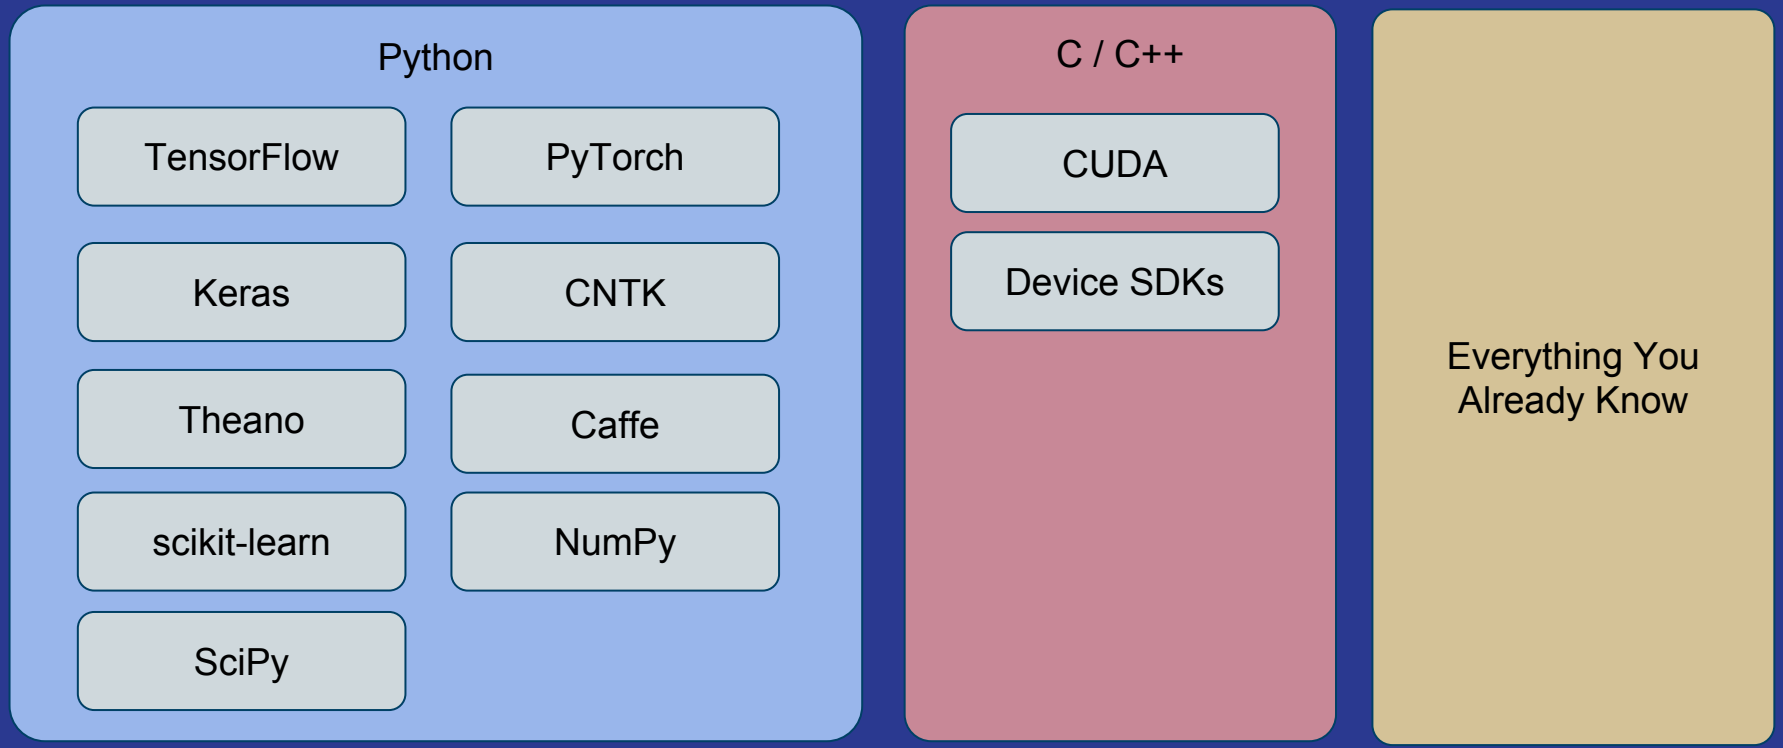

Systems

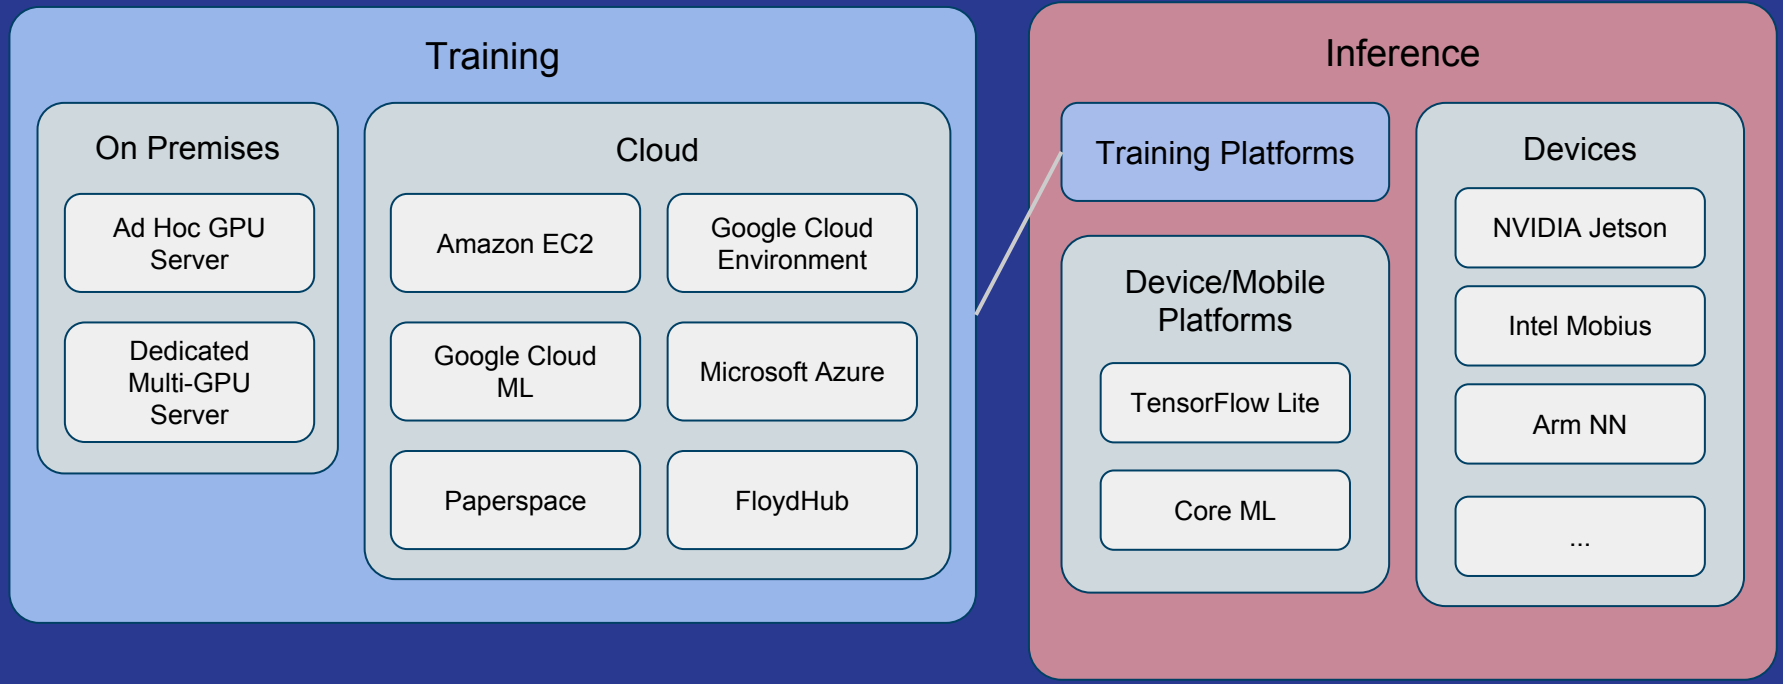
Training
On Premises
Ad Hoc GPU Server
Dedicated Multi-GPU Server
Cloud
Amazon EC2
Google Cloud Environment
Google Cloud ML
Microsoft Azure
Paperspace
FloydHub
Inference
Training Platforms
Device/Mobile Platforms
TensorFlow Lite
Core ML
Devices
NVIDIA Jetson
Intel Mobius
Arm NN
...

Hard to be an expert in everything

The Floating City Problem


A
B
C
D

“The Data, Stupid”

Individuals

Build something small to start

Find a mentor or community for help

Attend local area meetups

Build something else, this time with others

Follow your nose

Organizations

Lower expectations

Select a senior engineer to spearhead a project

Select a problem that has multiple known solutions

Ship something non-critical

Expect to throw-away, but to learn

Books

The Master Algorithm, Pedro Domingos

Deep Learning with Python, Francois Chollet

Deep Learning, Ian Goodfellow

Starting Projects

Object Detection / Segmentation

Natural Language Processing

Generative Networks

Speech Recognition

Build Stuff

Starting Projects

Object Detection / Segmentation

Natural Language Processing

Generative Networks

Speech Recognition

Do It

guild.ai

https://guild.ai