Microsoft

Hochschule für Wirtschaft und Recht Berlin
Berlin School of Economics and Law

Berliner Forum Digitalisierung
Hochschule und Praxis im Dialog

Ausblick und Faszination: AI über APIs einbinden

Anja Fiegler
Microsoft Deutschland

19. September 2019

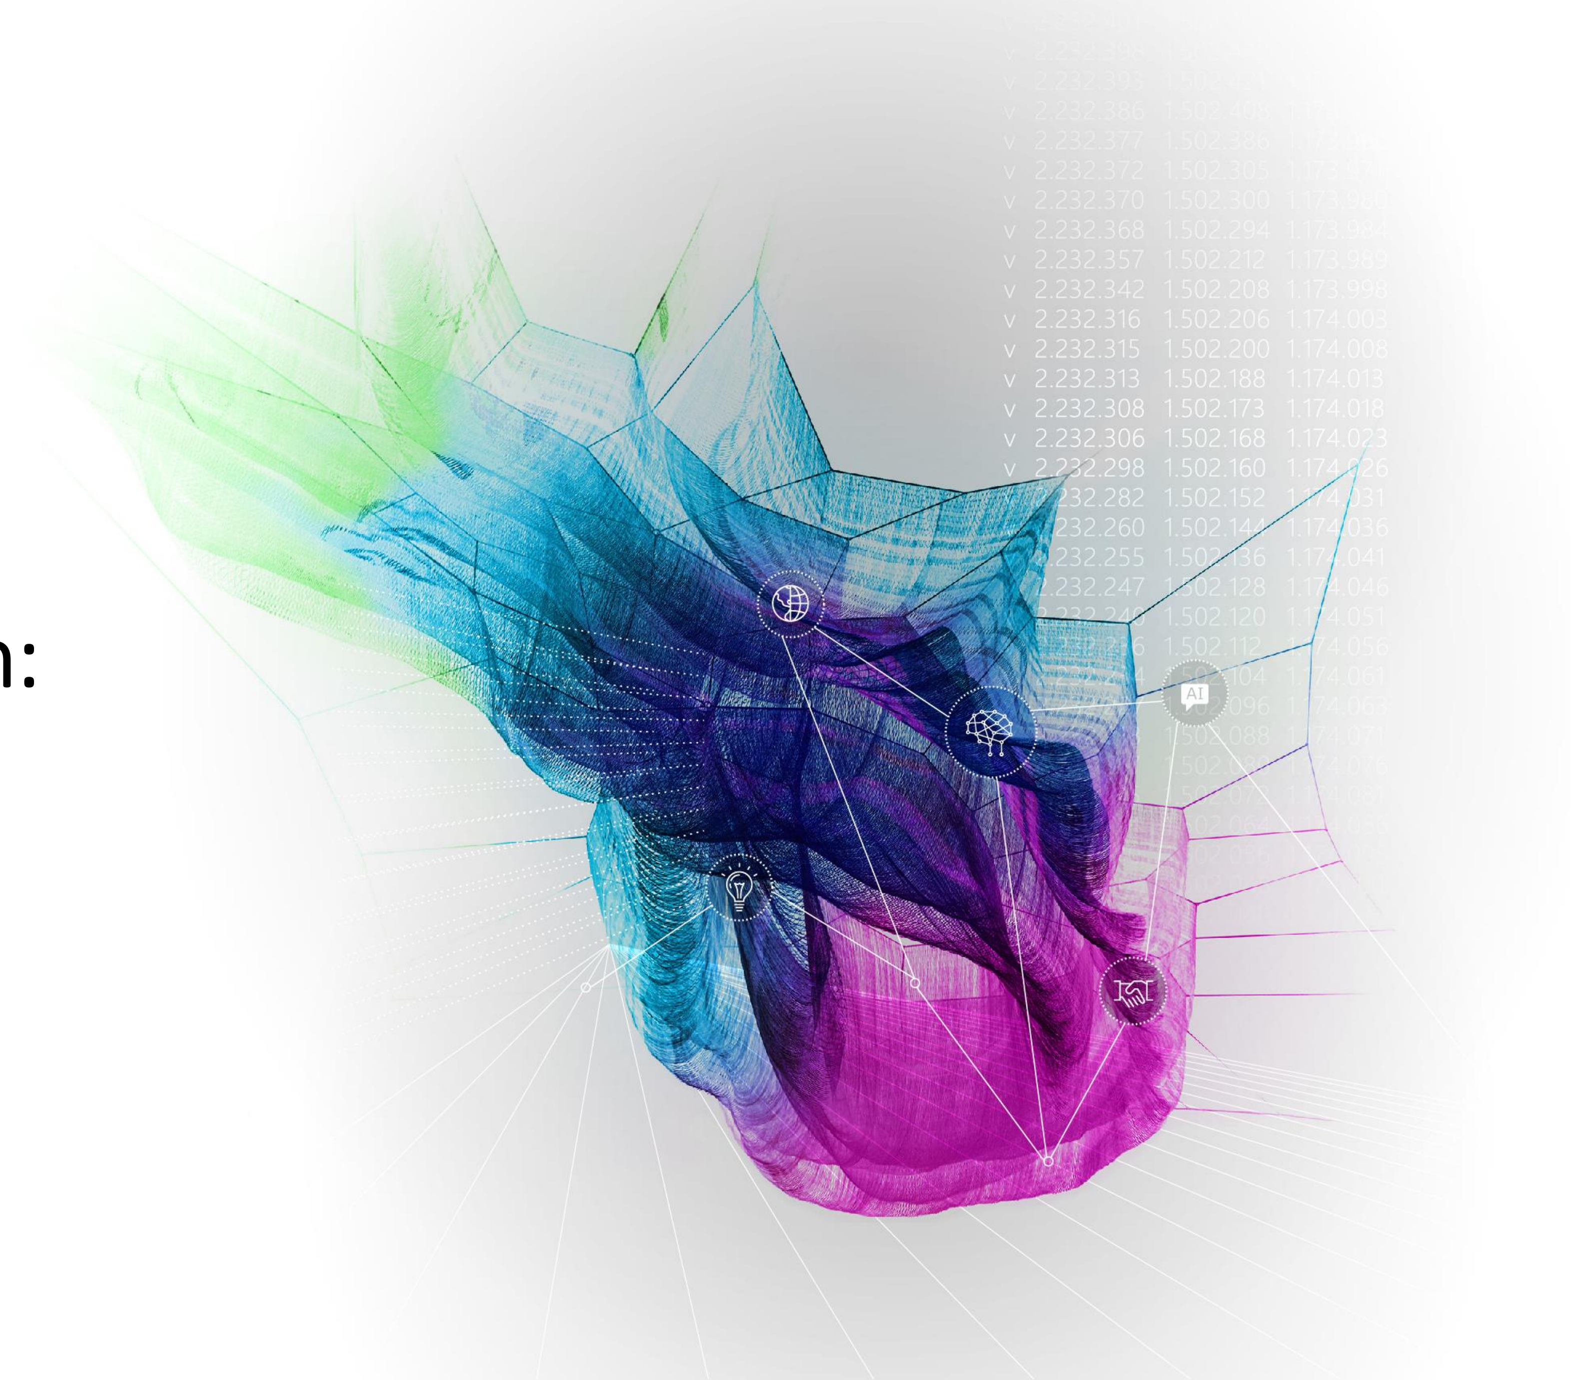

Microsoft

Agenda

1. Intro / Recap ML & AI
2. AI Ethics
3. AI Platforms & APIs
4. Four Important AI Trends
5. Cognitive Service – API Demo

What is AI?

Assigning human-like qualities to digital experiences

Perceives its environment

Mimics cognitive functions

Learns from example in volumes of data

Program that writes itself based on examples

Classifies, recommends, predicts, groups, segments

Strong AI - AGI

Combining weak AI with a consciousness or "mind"

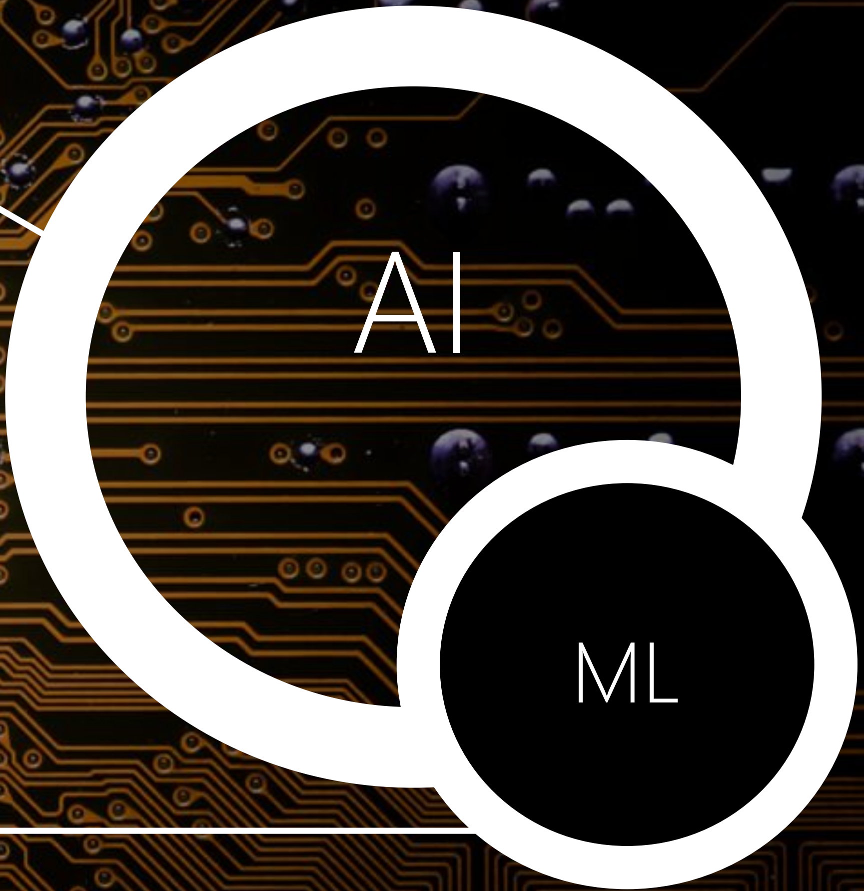

Weak AI

Separate cognitive functions, seeing, natural language, vision

Machine Learning Overview

Machine Learning Core Domains

The significant difference

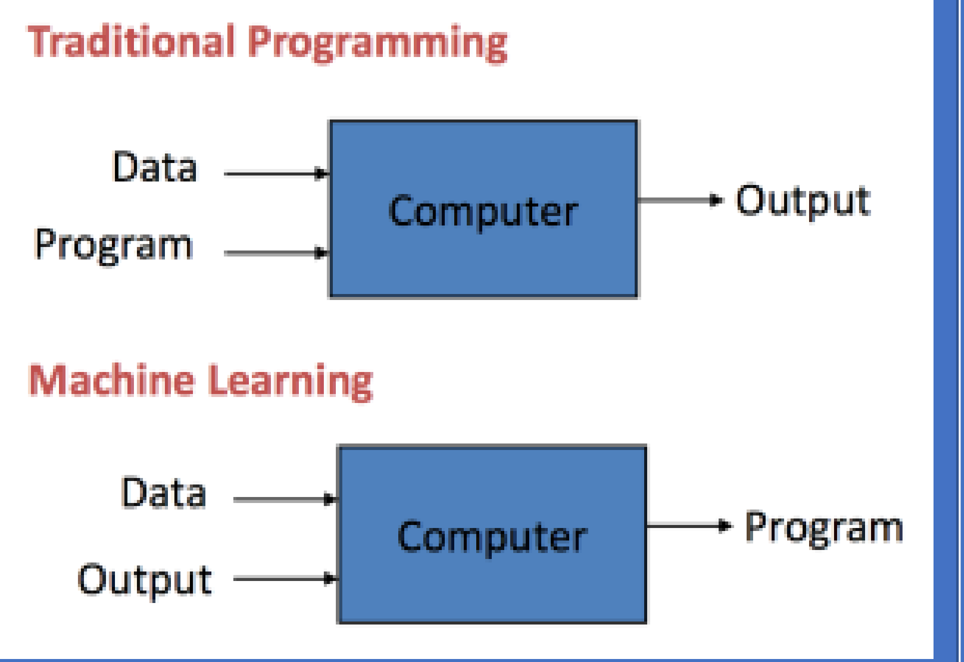

AutoML

Transfer Learning

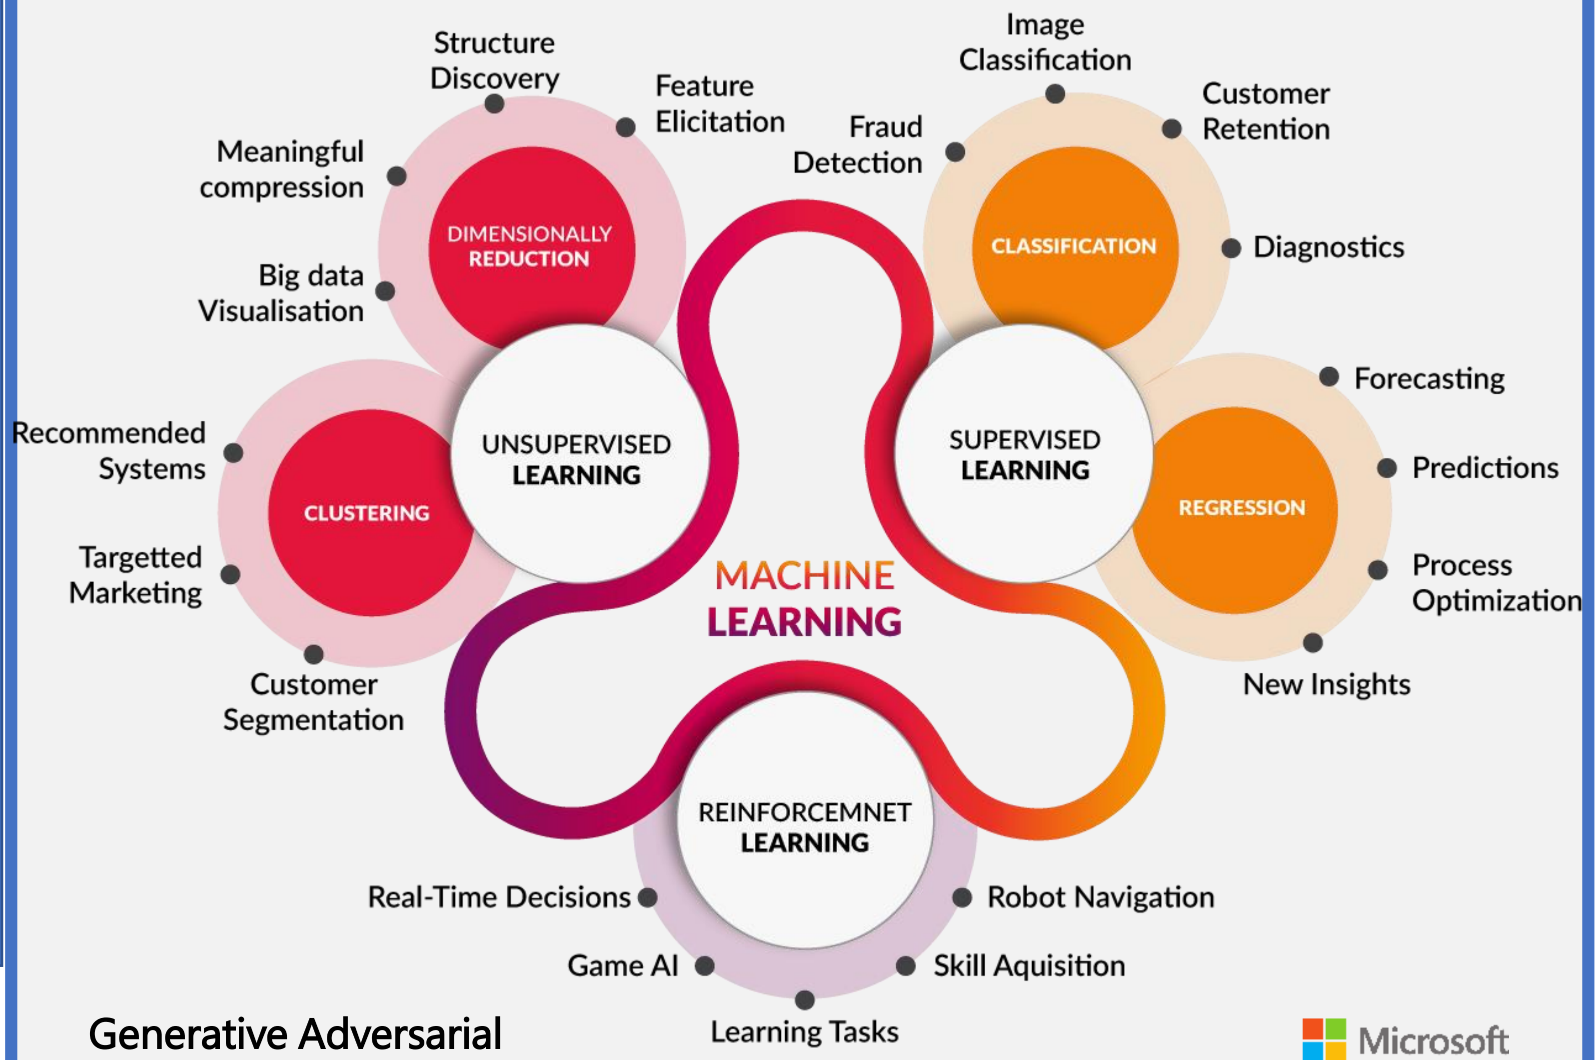

Generative Adversarial Network (GANs)

Microsoft

Machine Teaching

AI Progression Overview past 3 years

Breakthroughs human parity

94.9% on Switchboard test

69.9% with MT Research system

89.4% on Stanford CoQA test

96% on RESNET vision test

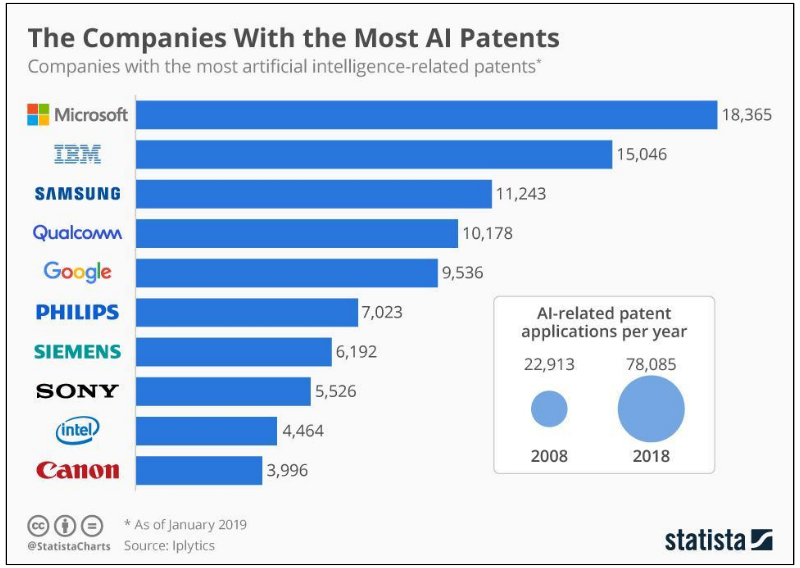

General Language Understanding Evaluation (GLUE)

	Rank	Name	Model	URL	Score
	1	Facebook AI	RoBERTa		88.5
	2	XLNet Team	XLNet-Large (ensemble)		88.4
+	3	Microsoft D365 AI & MSR AI	MT-DNN-ensemble		87.6
	4	GLUE Human Baselines	GLUE Human Baselines		87.1
+	5	王玮	ALICE large ensemble (Alibaba DAMO NLP)		86.3
	6	Stanford Hazy Research	Snorkel MeTaL		83.2

Conversational Q&A (CoQA) Test

Rank	Model	In-domain	Out-of-domain	Overall
	Human Performance *Stanford University* (Reddy & Chen et al. TACL '19)	89.4	87.4	88.8
1 Mar 29, 2019	Google SQuAD 2.0 + MMFT (ensemble) *MSRA + SDRG*	89.9	88.0	89.4
2 Mar 29, 2019	Google SQuAD 2.0 + MMFT (single model) *MSRA + SDRG*	88.5	86.0	87.8
2 Mar 29, 2019	ConvBERT (ensemble) *Joint Laboratory of HIT and iFLYTEK Research*	88.7	85.4	87.8
3 Mar 28, 2019	ConvBERT (single model) *Joint Laboratory of HIT and iFLYTEK Research*	87.7	84.6	86.8
3 Jan 25, 2019	BERT + MMFT + ADA (ensemble) *Microsoft Research Asia*	87.5	85.3	86.8

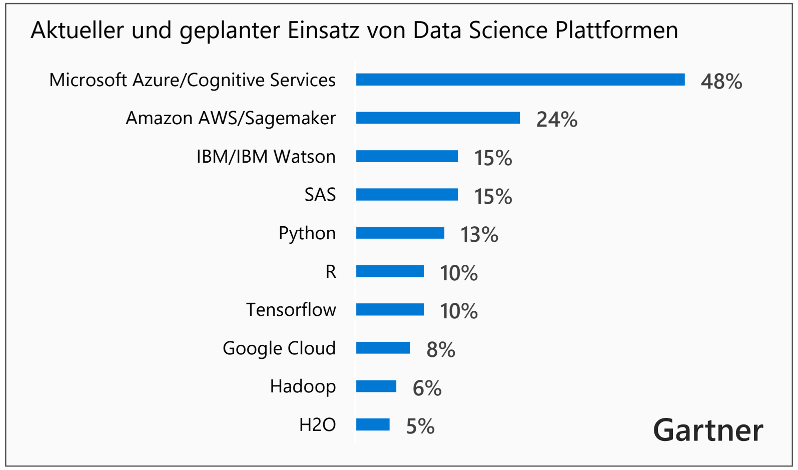

Enterprise-ready AI matters

"Organizations may be dabbling with other cloud providers, but when it comes to putting AI in production, that's when they come to Microsoft because they know it will work and has all the enterprise qualities they need." – Jim Hare, Gartner

Microsoft

AI – the most important aspect is …..

„We need to ask ourselves
not only what computers **can do**,
but what computers **should do** –
that time has come“

Satya Nadella
CEO Microsoft
Mai 2018


What do these people have in common?


They do not exist
www.thispersondoesnotexist.com

GPT Text Generation

Intro formulated by human

In a shocking finding, scientist discovered a herd of unicorns living in a remote, previously unexplored valley, in the Andes Mountains. Even more surprising to the researchers was the fact that the unicorns spoke perfect English.

Continued by OpenAI Modell GPT-2 (Feb. 2019)

The scientist named the population, after their distinctive horn, Ovid's Unicorn. These four-horned, silver-white unicorns were previously unknown to science.

Now, after almost two centuries, the mystery of what sparked this odd phenomenon is finally solved.

Dr. Jorge Pérez, an evolutionary biologist from the University of La Paz, and several companions, were exploring the Andes Mountains when they found a small valley, with no other animals or humans. Pérez noticed that the valley had what appeared to be a natural fountain, surrounded by two peaks of rock and silver snow.

Pérez and the others then ventured further into the valley. "By the time we reached the top of one peak, the water looked blue, with some crystals on top," said Pérez.

Pérez and his friends were astonished to see the unicorn herd. These creatures could be seen from the air without having to move too much to see them – they were so close they could touch their horns.

While examining these bizarre creatures the scientists discovered that the creatures also spoke some fairly regular English. Pérez stated, "We can see, for example, that they have a common 'language,' something like a dialect or dialectic."

Dr. Pérez believes that the unicorns may have originated in Argentina, where the animals were believed to be descendants of a lost race of people who lived there before the arrival of humans in those parts of South America.

While their origins are still unclear, some believe that perhaps the creatures were created when a human and a unicorn met each other in a time before human civilization. According to Pérez, "In South America, such incidents seem to be quite common."

However, Pérez also pointed out that it is likely that the only way of knowing for sure if unicorns are indeed the descendants of a lost alien race is through DNA. "But they seem to be able to communicate in English quite well, which I believe is a sign of evolution, or at least a change in social organization," said the scientist.

OpenAI

Beispiel Medizin: Erheblicher Nutzen durch Künstliche Intelligenz – Gleichzeitig Missbrauchsrisiko falls die Technologie in falsche Hände gerät

KI basierte Erkennung seltener Erbkrankheiten mittels Portraitfoto (DeepGestalt)

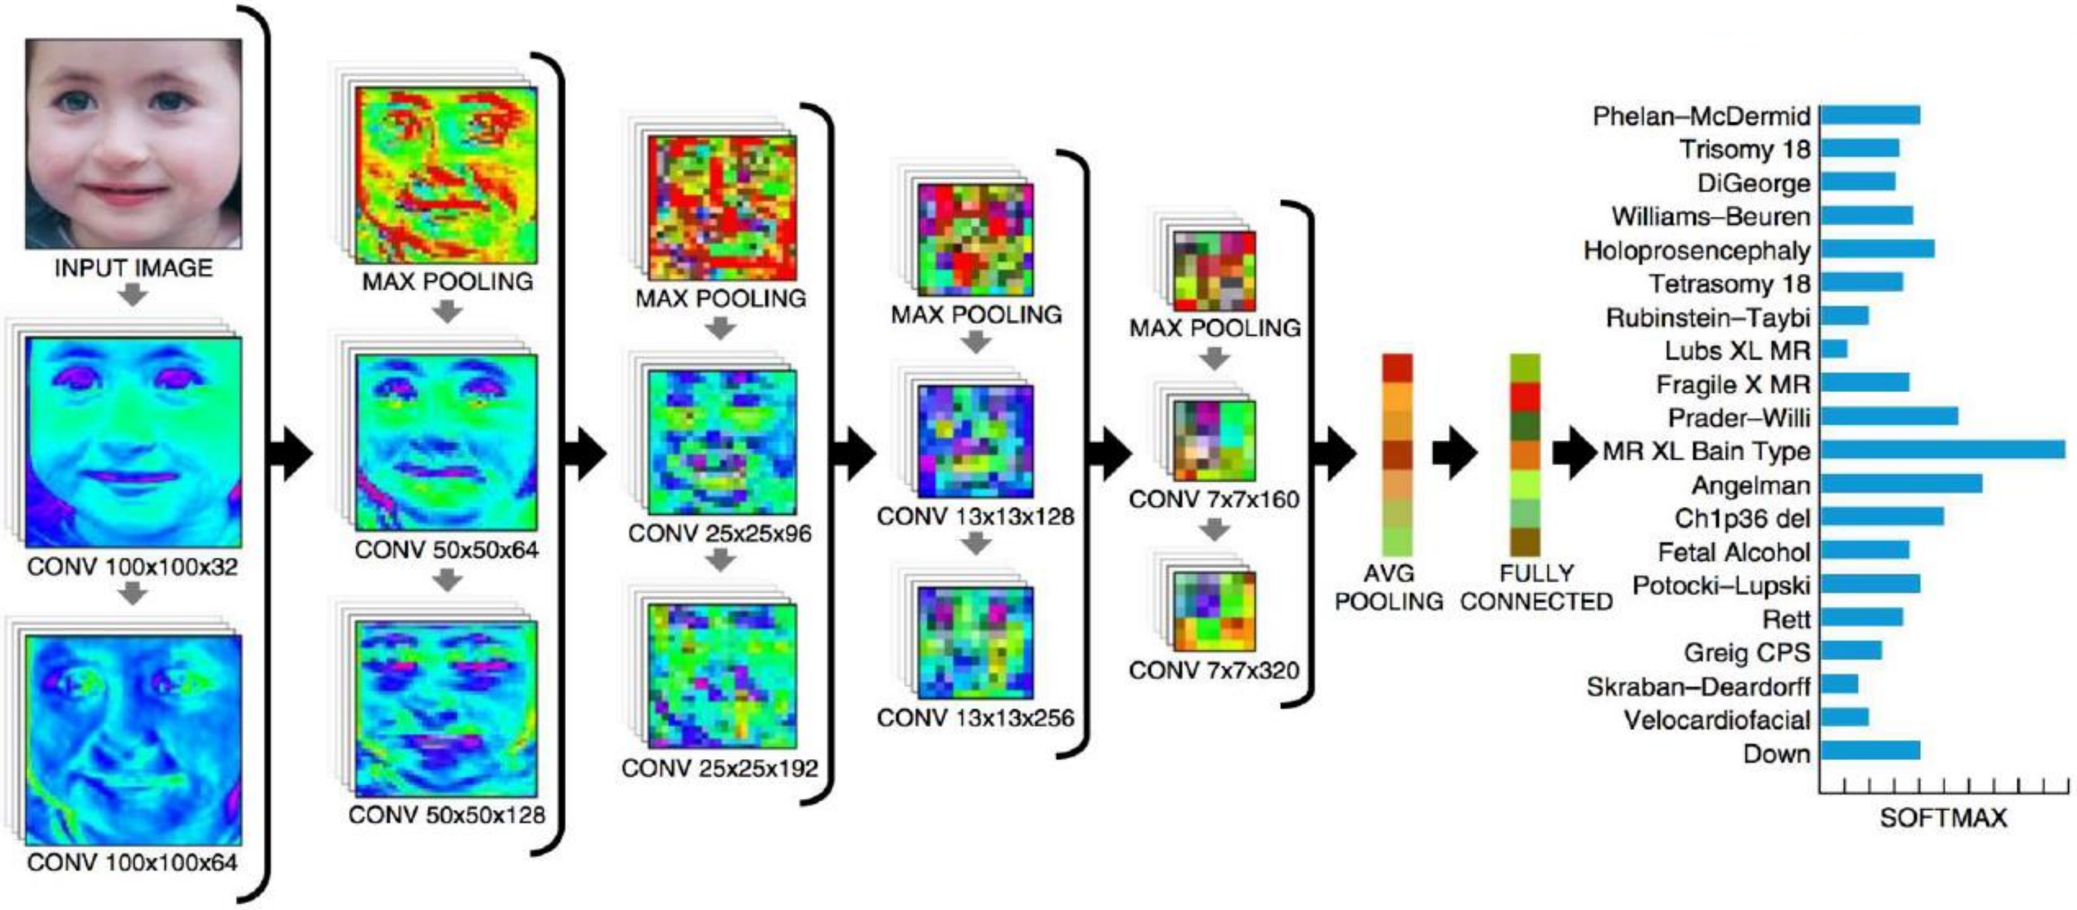

Quelle: https://deeplearning.mit.edu

Face2Gene

The ethics of AI

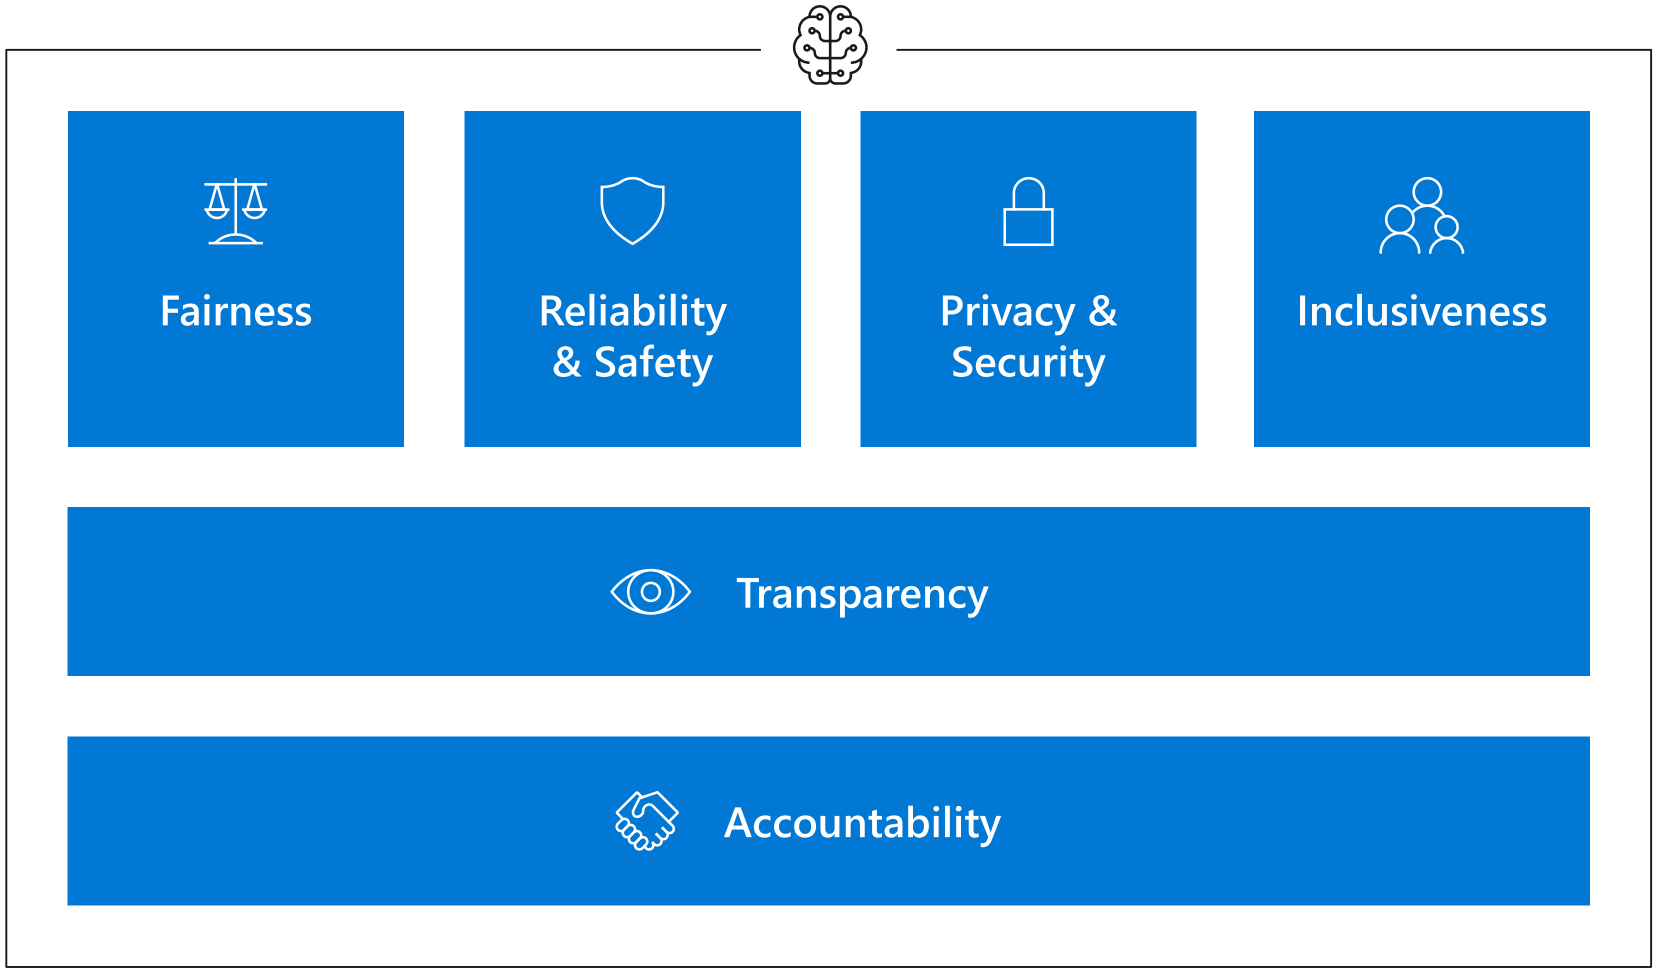

LIME

Local
Interpretable
Model-Agnostic
Explanations

Example 1:

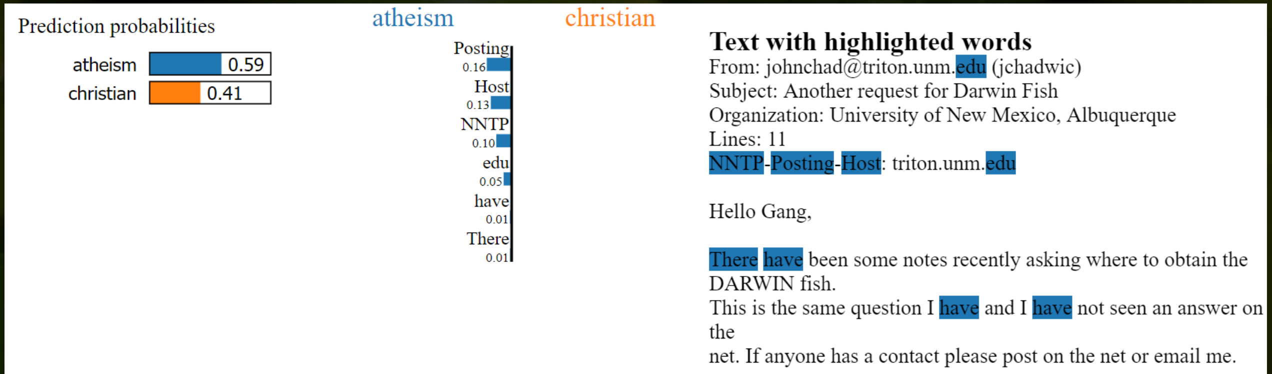

How It Works

1. Permute Data **
2. Calculate distance between permutations and original observations
3. Make predictions on new data using complex model
4. Pick *m* features best describing the complex model outcome from the permuted data **
5. Fit a simple model to the permuted data with *m* features and similarity scores as weights **
6. Feature weights from the simple model make explanations for the complex models local behavior

Example 2:

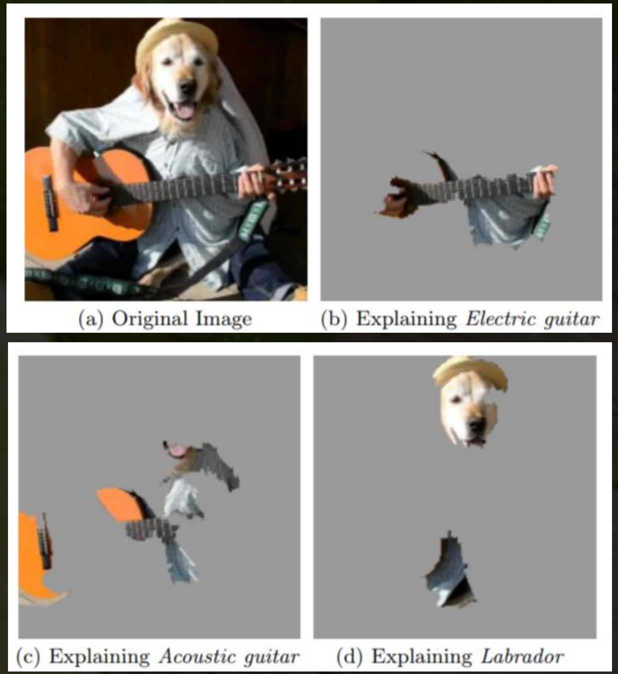

(a) Original Image (b) Explaining *Electric guitar* (c) Explaining *Acoustic guitar* (d) Explaining *Labrador*

Source: https://www.youtube.com/watch?v=CY3t11vuuOM

Kasia Kulma, PHD

https://github.com/marcotcr/lime

AI – Portfolios and the Link to APIs

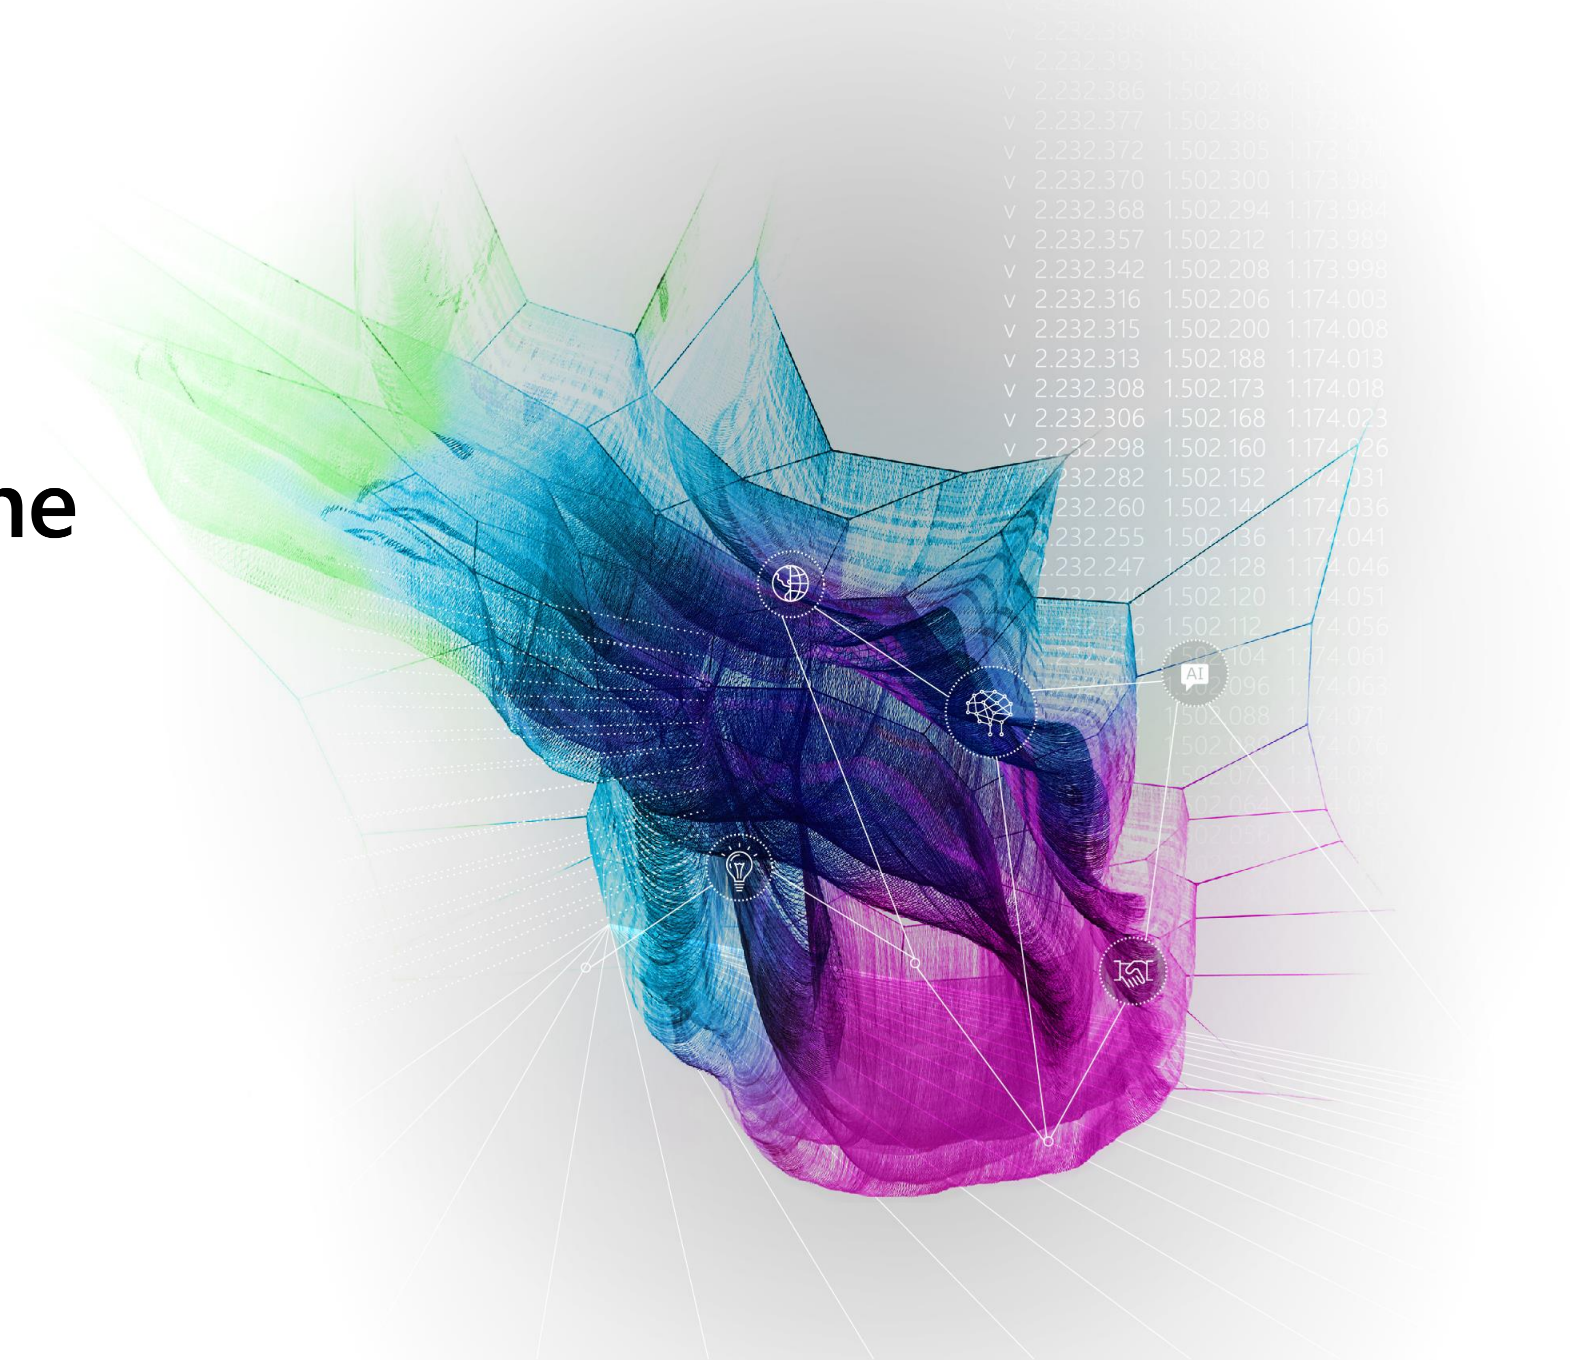

Microsoft AI Strategy Pilars

For all skill levels

Automated ML + Visual Interface + Code first

Industry leading MLOps

Integrated with Azure DevOps

Open

Any tool + any framework

Microsoft investiert eine Milliarde in KI-Startup OpenAI

Beide Unternehmen wollen künstliche Intelligenz der gesamten Menschheit zu Gute kommen lassen. Auf Azure soll eine neue Plattform für KI-Entwickler entstehen.

Lesezeit: 1 Min. heise online Jul 22, 2019 17

AI – Level of Consumption

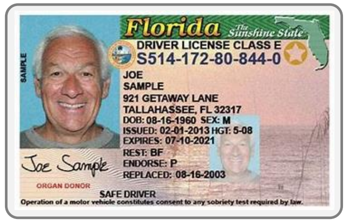

Extract identity

A que horas fecham hoje?

Comparable Software Level

- SaaS
- PaaS
- SDK
- Compiler

AI Consumption Model

how deep in the stack?

Effort: *Everyone* – *Expert* – *Developer* – *Pioneer*

- General Pre-Built AI Services
- Customizable AI Services
- Transfer Learning
- Pre-Trained Models
- AutoML
- Custom KI

Technology: *Specialized AI-Services* → *General AI-Technology*

AI Portfolio

Domain specific pretrained models
To simplify solution development

Vision · Speech · Language · Search

Familiar Data Science tools
To simplify model development

Visual Studio Code · Azure Notebooks · Jupyter · Command line

Popular frameworks
To build advanced deep learning solutions

PyTorch · TensorFlow · Scikit-Learn · ONNX

Productive services
To empower data science and development teams

Azure Databricks · Azure Machine Learning · Machine Learning VMs

Powerful infrastructure
To accelerate deep learning

CPU · GPU · FPGA

From the Intelligent Cloud to the Intelligent Edge

Microsoft Cognitive Services

Vision

- Video Indexer
- Computer Vision
- Face
- Emotion
- Content Moderator
- Custom Vision
- Forms Understanding

Speech

- Speaker Recognition
- Bing Speech
- Custom Speech
- Translator Speech
- Unified Speech
  - Speech to Text w. Custom Speech
  - Text to Speech w. Custom Voice
  - Speech Translation w. Custom Translator

Language

- Text Analytics
- Bing Spell Check
- Translator Text
- Language Understanding (LUIS)
- Custom Text Analytics

Knowledge

- QnA Maker
- Custom Decision

Search

- Bing Entity Search
- Bing Autosuggest
- Bing Search
  - Web Search
  - Image Search
  - News Search
  - Video Search
- Bing Statistics add-in
- Bing Visual Search
- Bing Custom Search

Labs

- Project Gesture
- Project Local Insights
- Project Academic Knowledge
- Project Entity Linking
- Project Knowledge Exploration
- Project Event Tracking
- Project Answer Search
- Project URL Preview
- Project Anomaly Finder
- Project Conversation Learner
- Project Personality Chat

Why Microsoft Cognitive Services?

Easy

Roll your own with REST APIs

Simple to add: just a few lines of code required

Flexible

Integrate into the language and platform of your choice

Breadth of offerings helps you find the right for your app

Bring your own data for your custom experience

Tested

Built by experts in their field from Microsoft Research, Bing, and Azure Machine Learning

Quality documentation, sample code, and community support

Pre-Built AI API Use – Deployment View

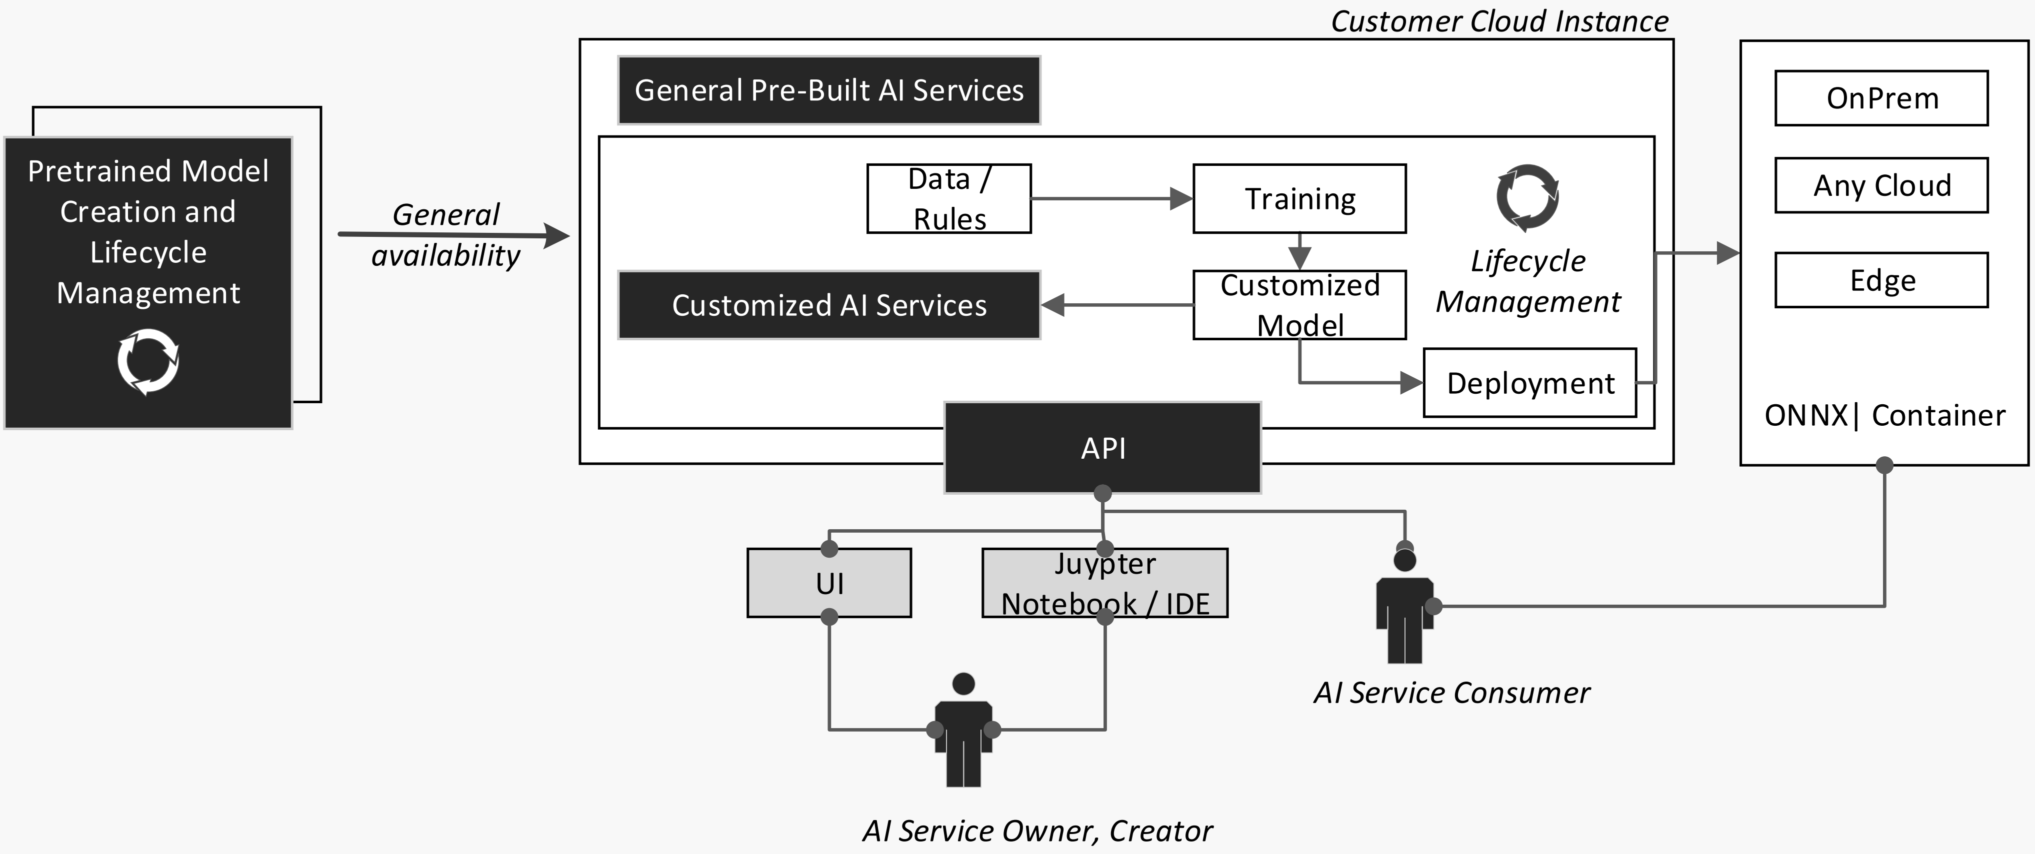

Democratizing AI: Time to Value

2018

Multiple Data Scientists
with custom algorithms

Months

2019

1 Domain
Expert with
Custom Vision
Cognitive
Service

Hours

COMPARABLE RESULTS

Microsoft

AI Core Trends

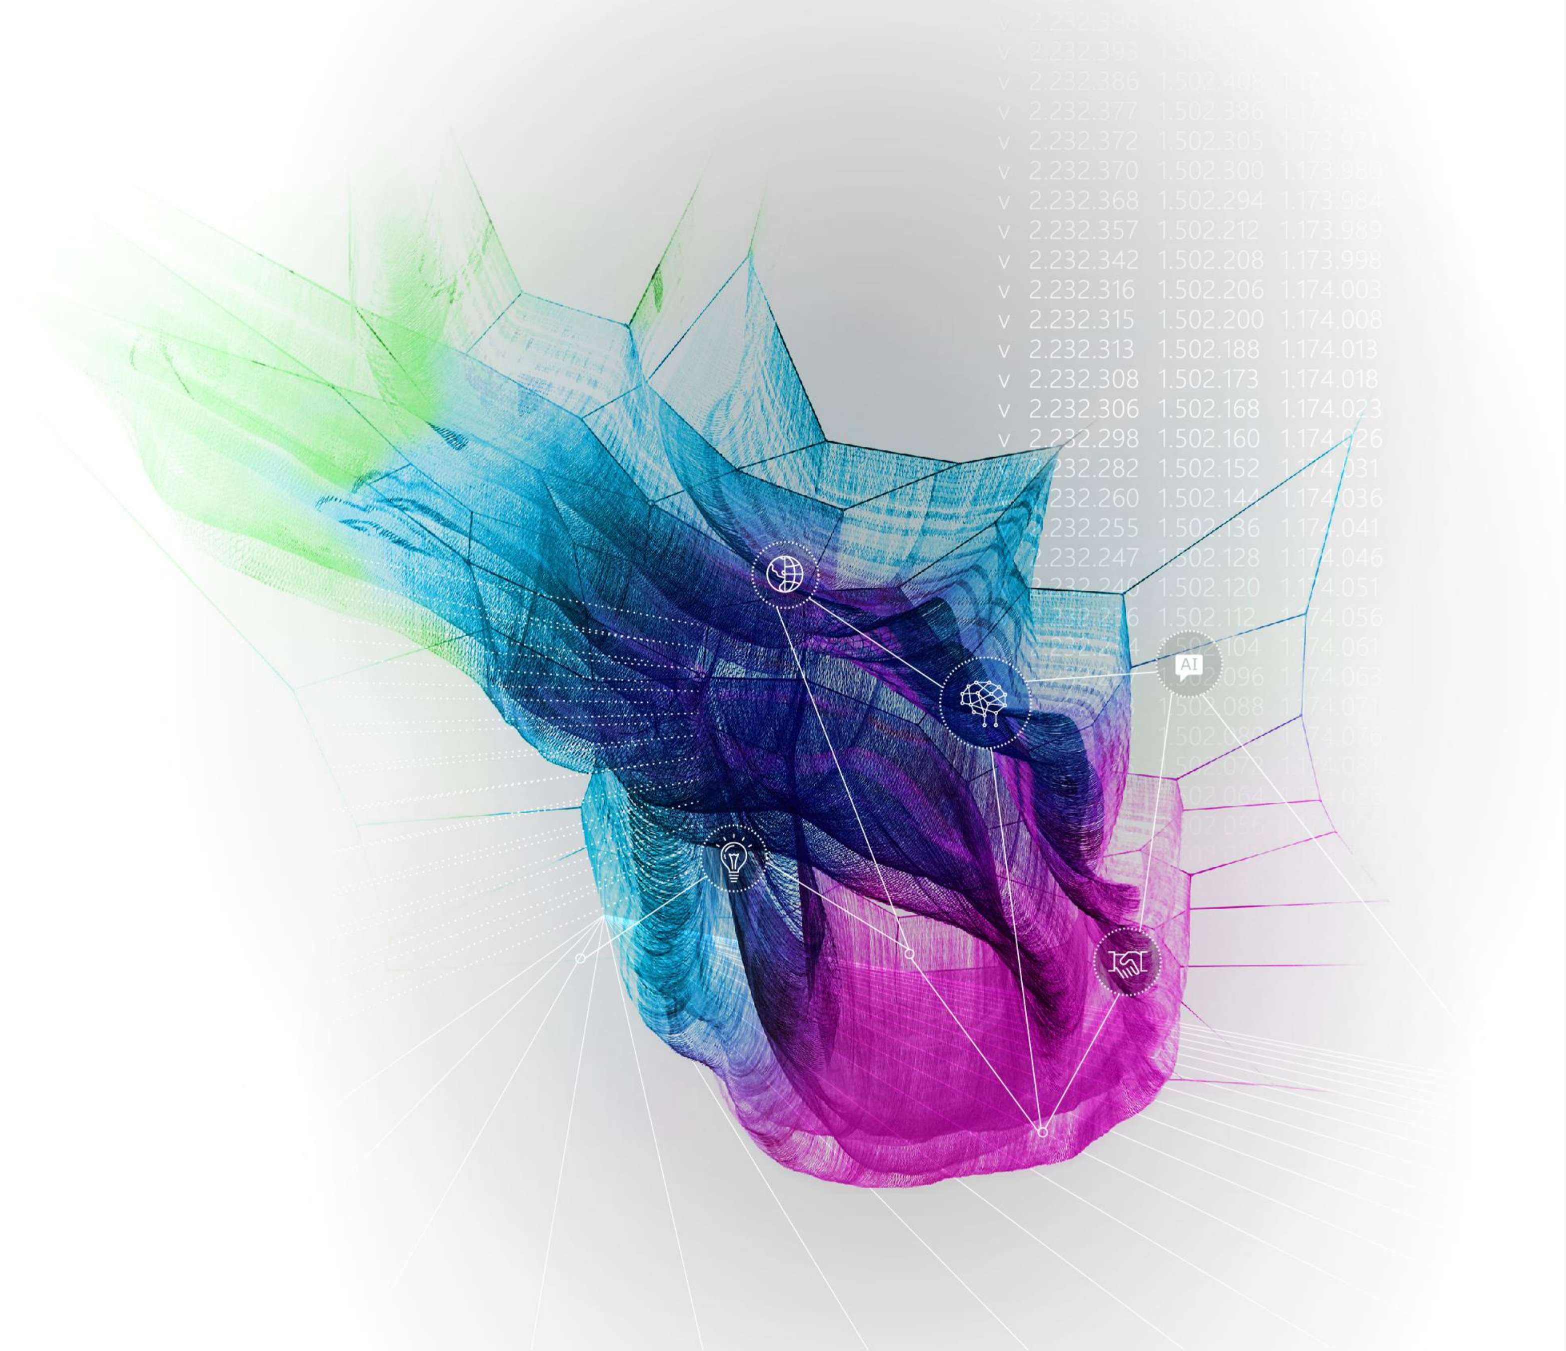

Machine Learning DevOps

End-to-end ML lifecycle management

CI/CD for ML pipelines

Feedback loop

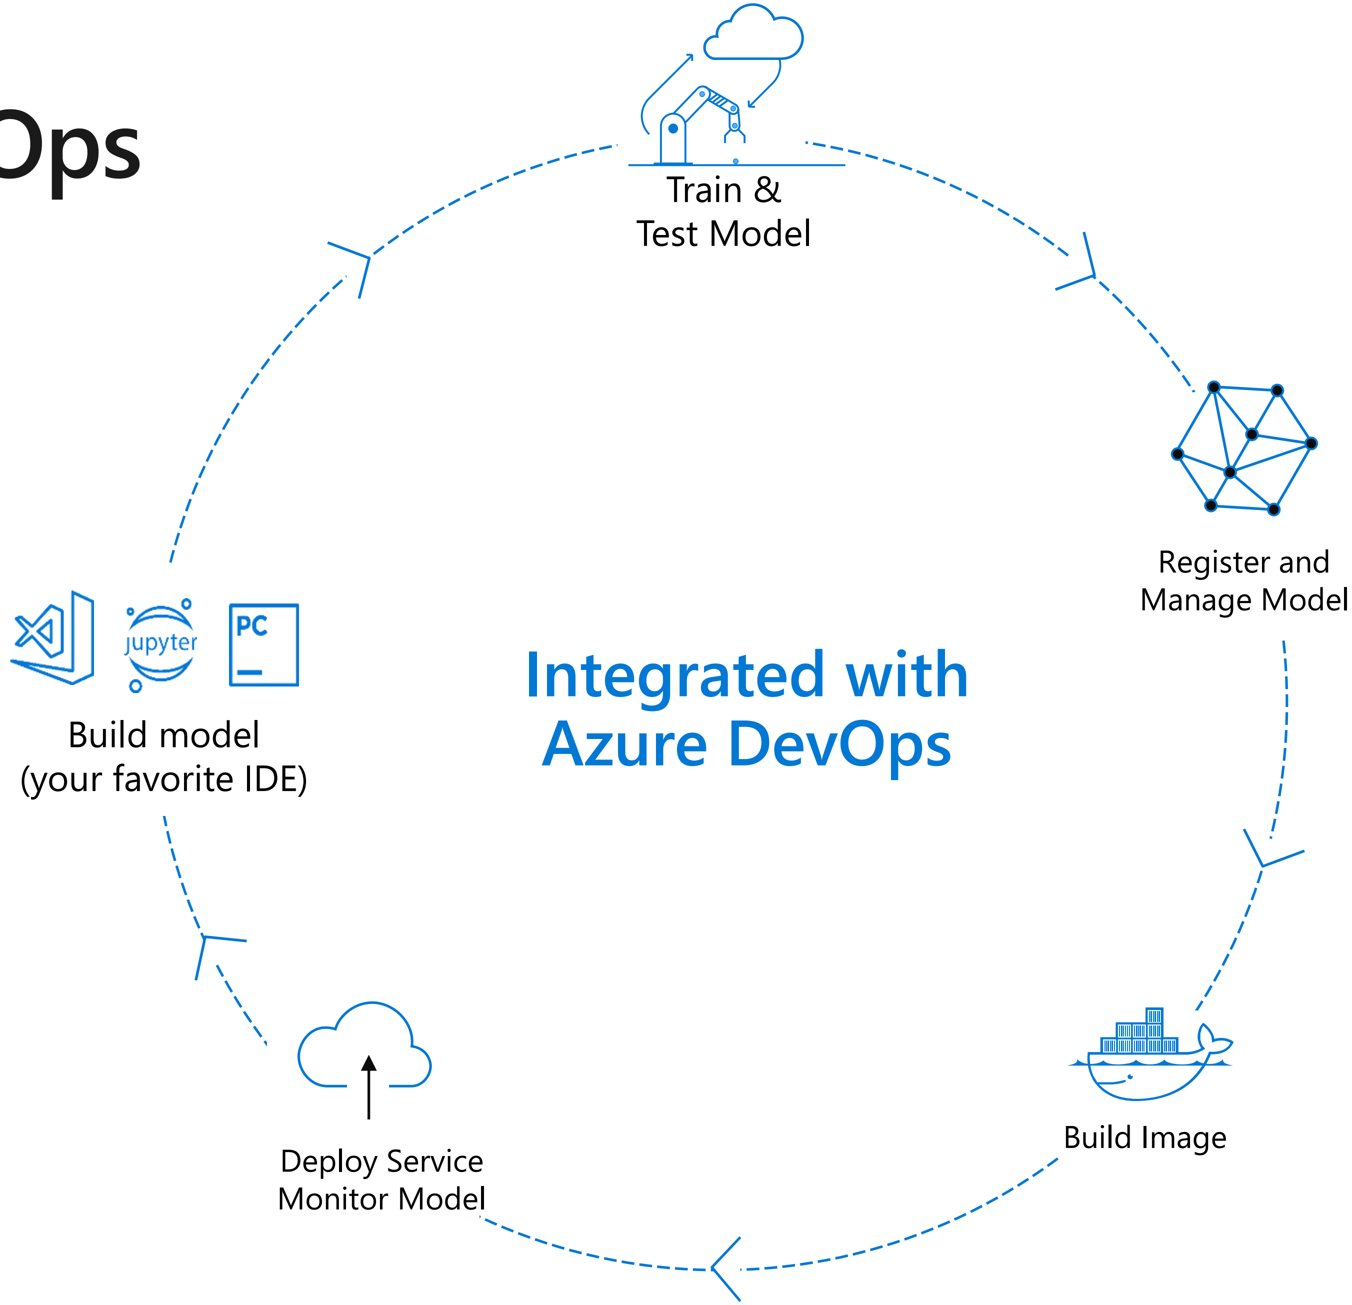

Home / Solutions / Confidential Compute

Azure confidential computing

Protect and secure your cloud data while it's in use

Start free

- Safeguard data from malicious and insider threats while it's in use
- Maintain control of data through its lifetime
- Protect and validate the integrity of code in the cloud
- Ensure that data and code is opaque to the cloud platform provider

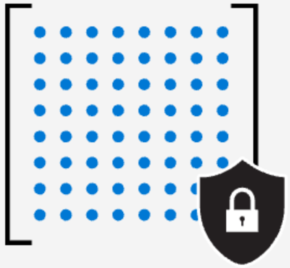

Take data security to the next level with confidential computing

Azure confidential computing protects the confidentiality and integrity of your data and code while it's processed in the public cloud. Cloud security is the cornerstone of our confidential cloud vision, which aims to remove Microsoft from the trusted computing base (TCB) of Azure.

Deep reinforcement learning – DRL (Machine Teaching)

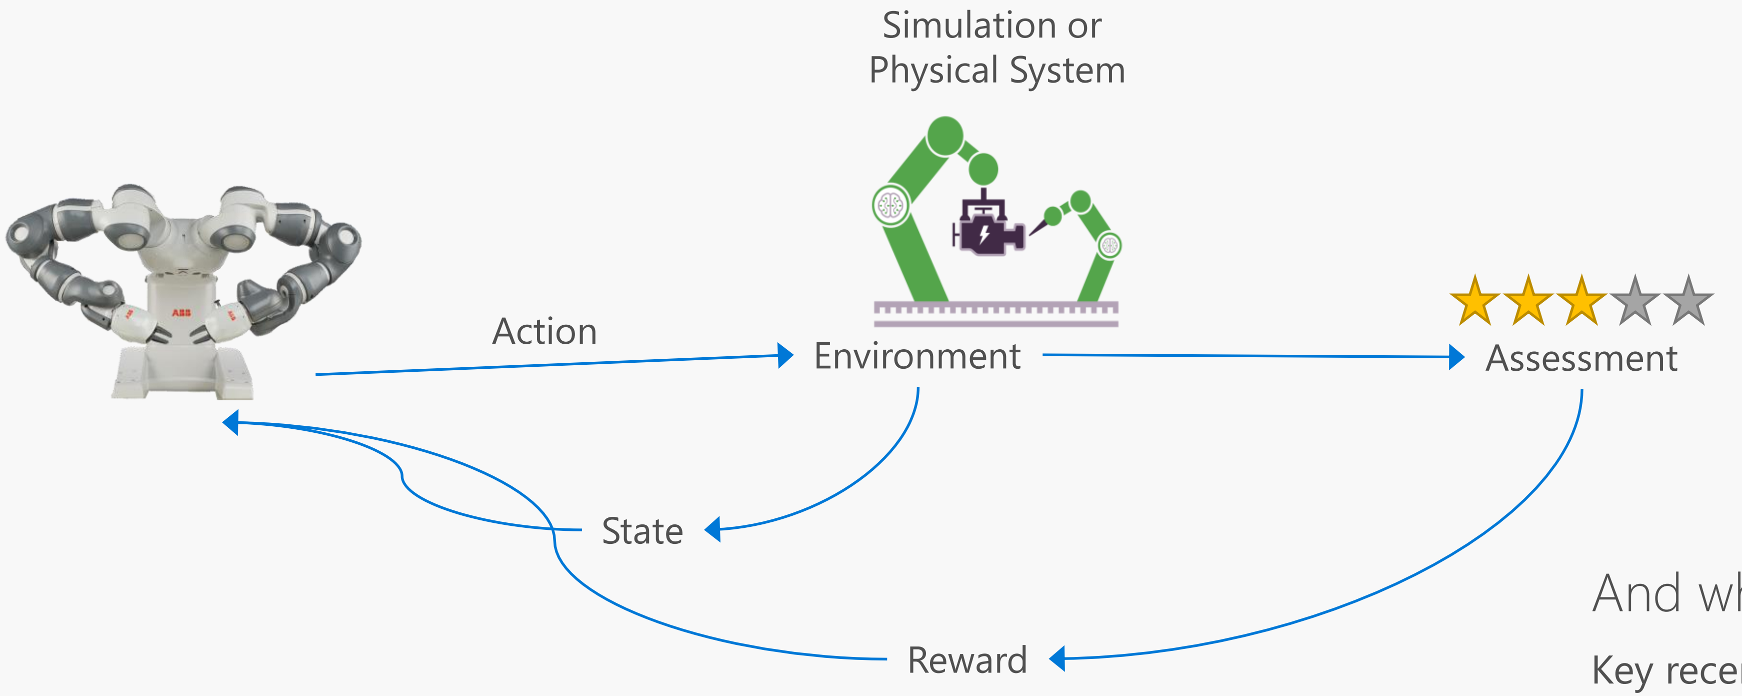

And why now?

Key recent advancements in reinforcement learning: neural networks, compute infrastructure and algorithms

DRL is well suited for complex and dynamic environment (unlike other ML techniques)

Proliferation of connected physical devices and sensors

Multi-party and Distributed Machine Learning

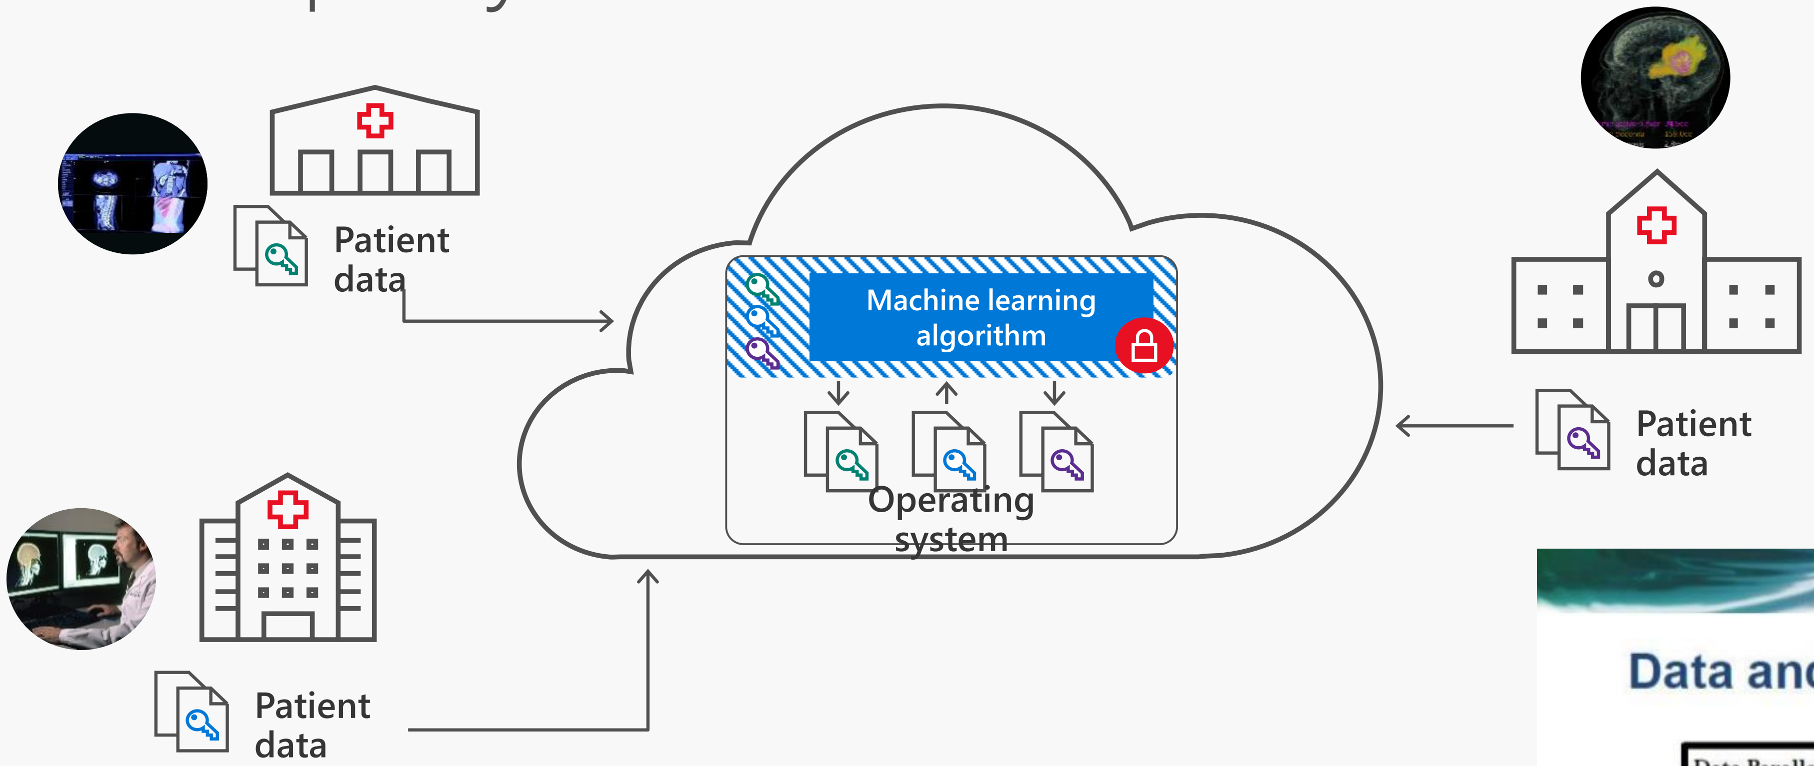

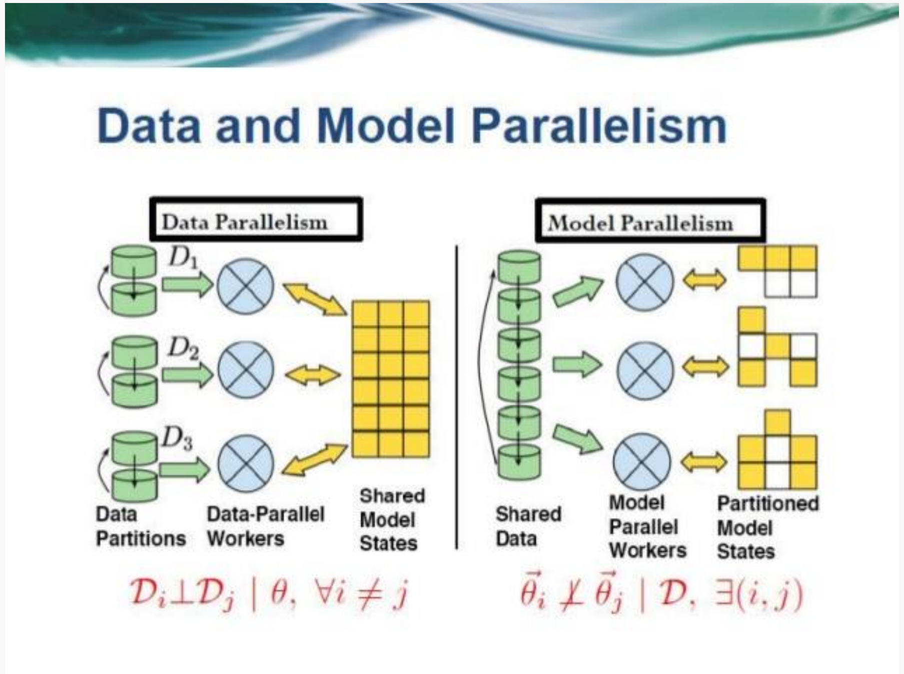

Microsoft

Cognitive Services & Demo

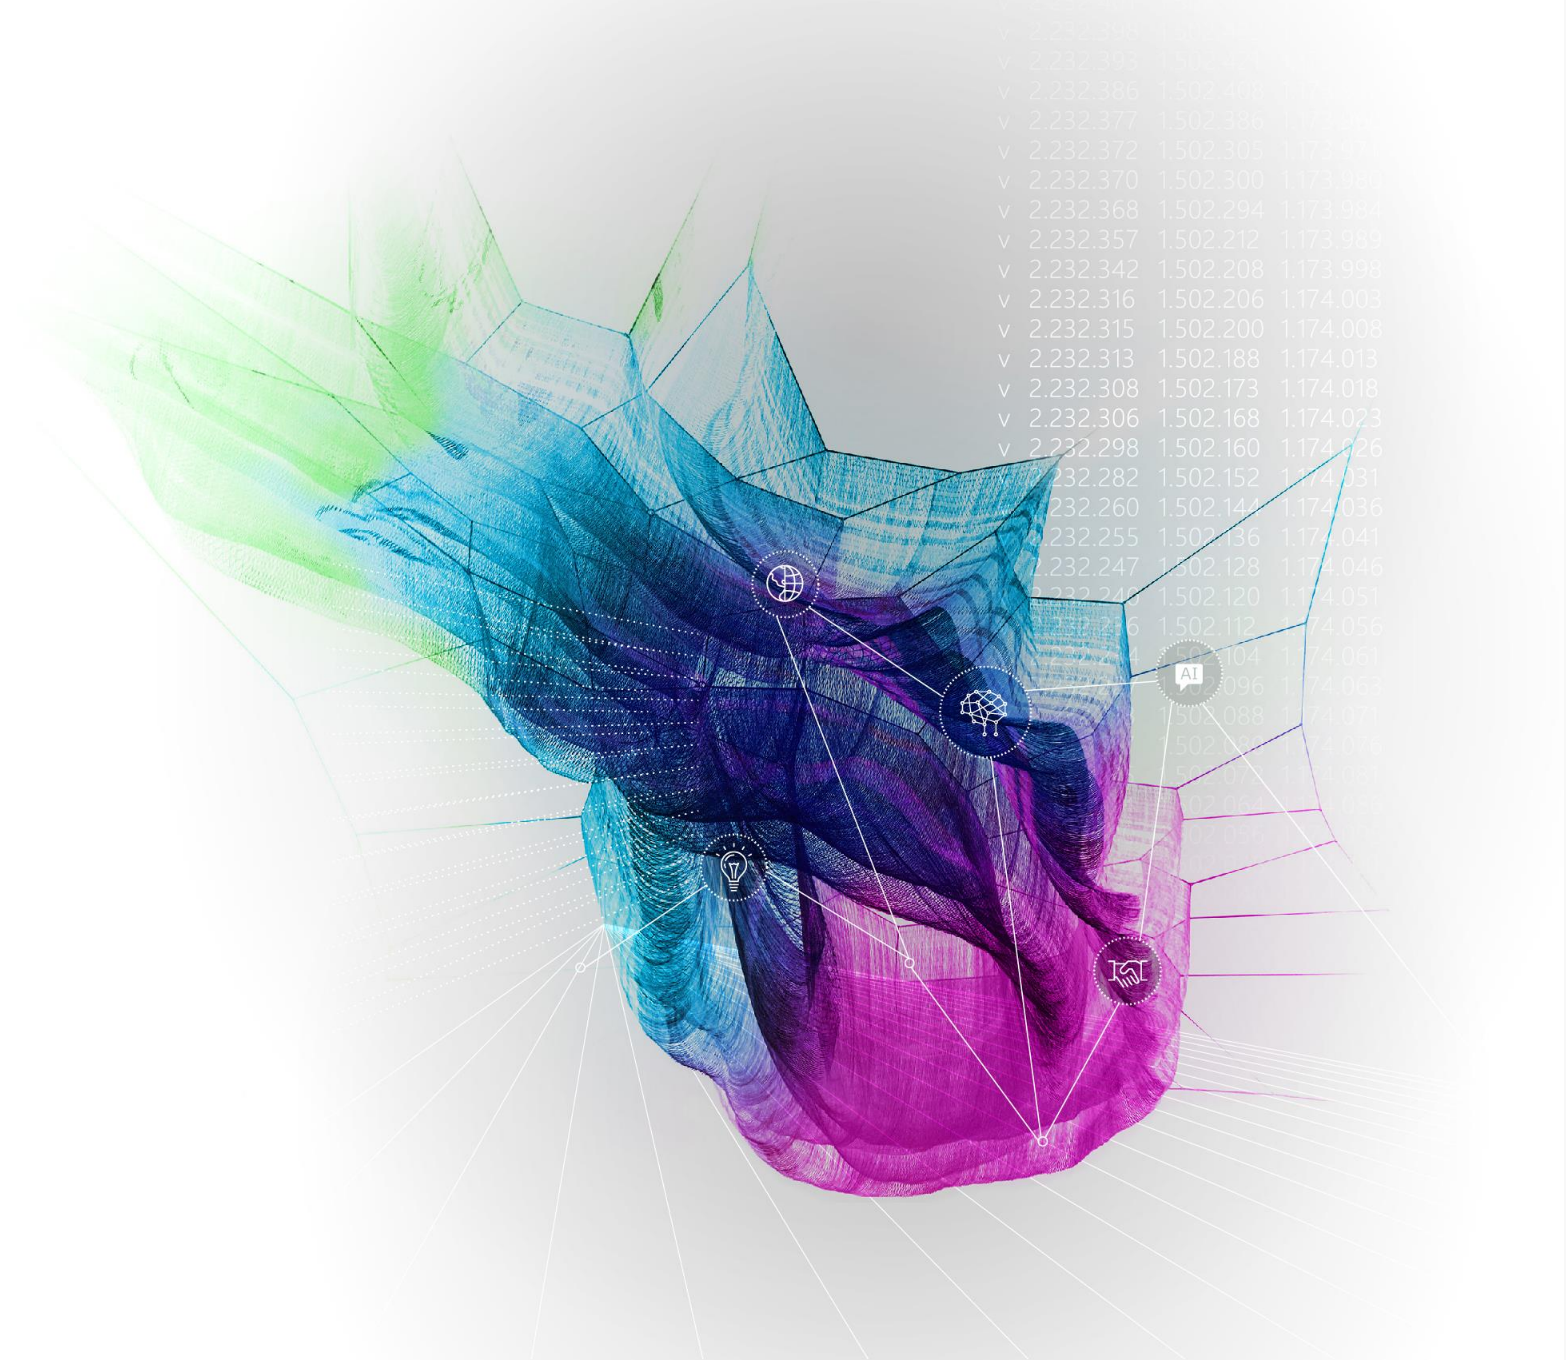

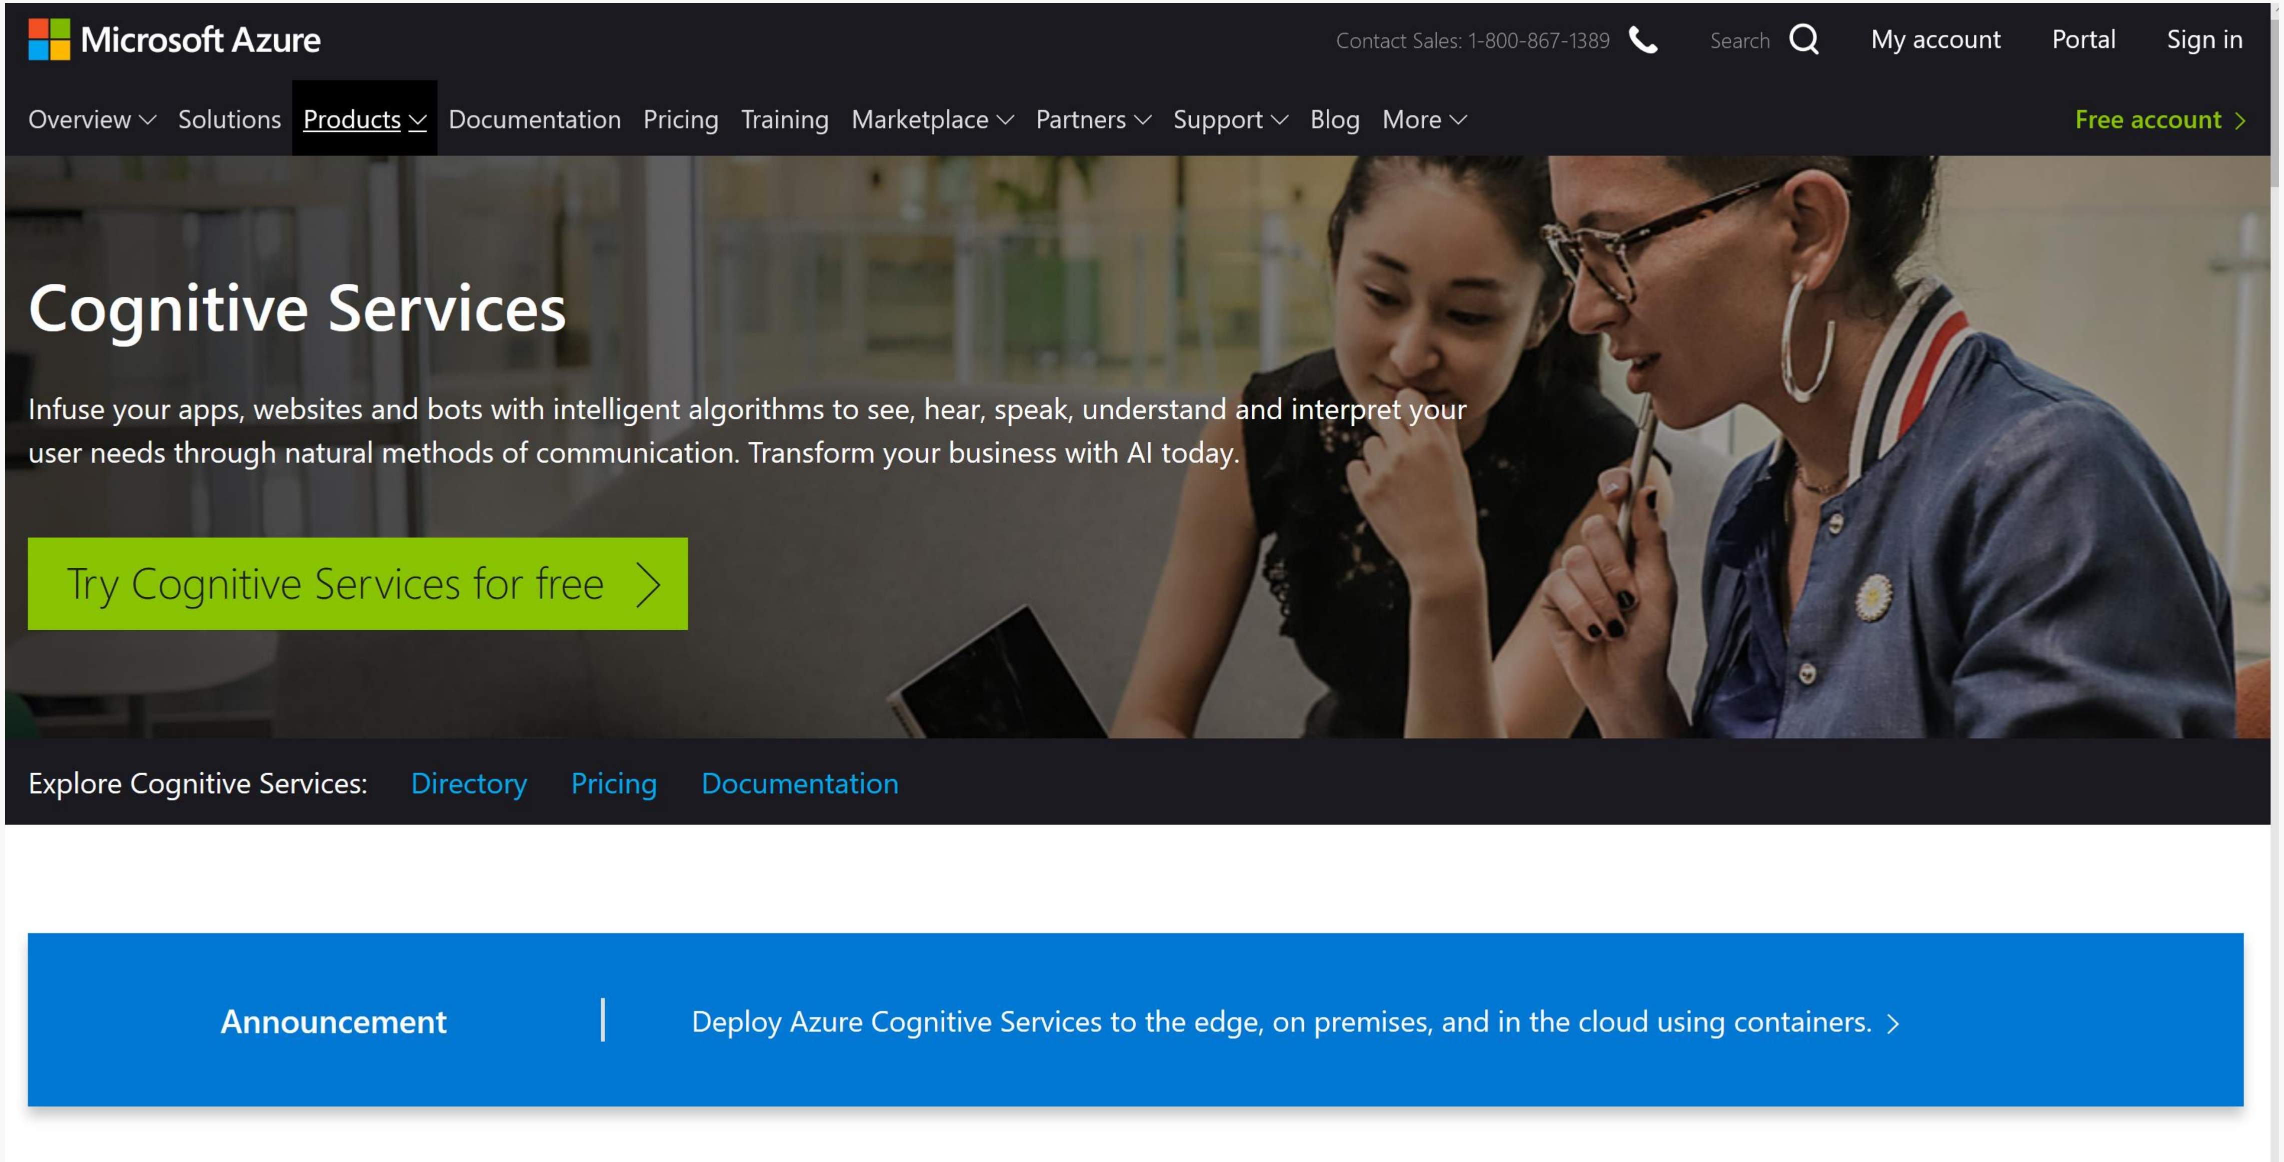

https://azure.microsoft.com/en-us/services/cognitive-services/

Computer Vision

FEATURE NAME:	VALUE
Description	{ "tags": ["train", "platform", "station", "building", "indoor", "subway", "track", "walking", "waiting", "pulling", "board", "people", "man", "luggage", "standing", "holding", "large", "woman", "yellow", "suitcase"], "captions": [{ "text": "people waiting at a train station", "confidence": 0.833099365 }] }
Tags	[{ "name": "train", "confidence": 0.9975446 }, { "name": "platform", "confidence": 0.995543063 }, { "name": "station", "confidence": 0.9798007 }, { "name": "indoor", "confidence": 0.927719653 }, { "name": "subway", "confidence": 0.838939846 }, { "name": "pulling", "confidence": 0.431715637 }]
Image format	"Jpeg"

Example: AI helping to fight Breast Cancer

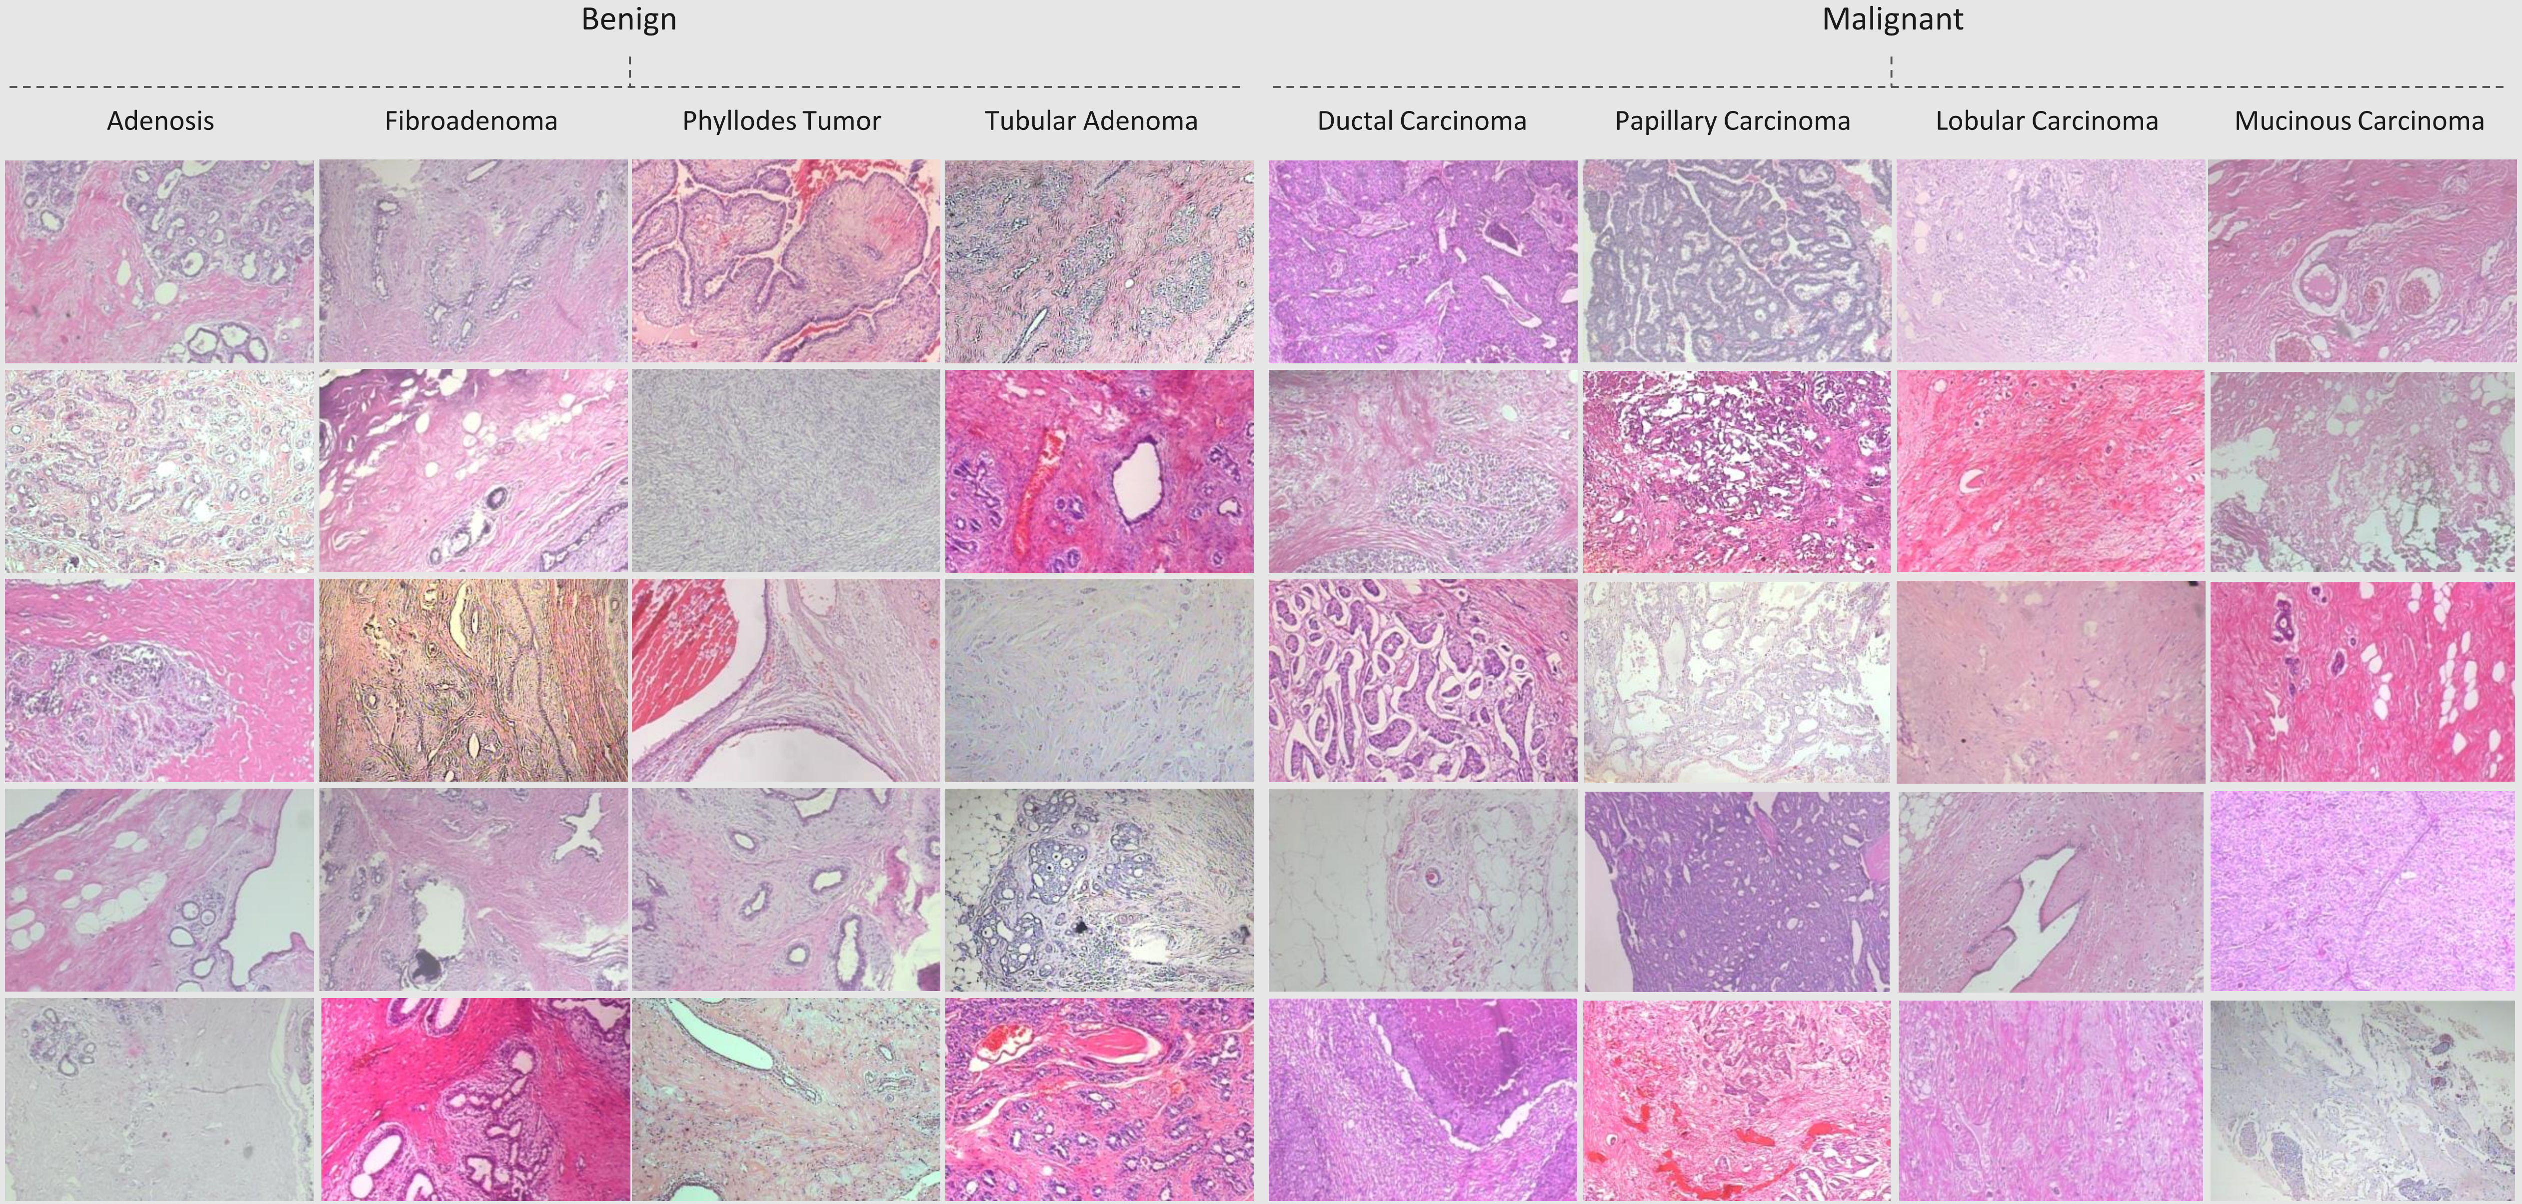

Source: [1] Spanhol, F., Oliveira, L. S., Petitjean, C., Heutte, L., A Dataset for Breast Cancer Histopathological Image Classification, IEEE Transactions on Biomedical Engineering (TBME), 63(7):1455-1462, 2016

Benchmark: Custom Convolutional Neural Network Analysis

Breast Cancer Histopathological Image Classification using Convolutional Neural Networks

Fabio A. Spanhol, Luiz S. Oliveira
Federal University of Parana
Department of Informatics (DInf)
Curitiba, PR - Brazil
Email: {faspanhol, lesoliveira}@inf.ufpr.br

Caroline Petitjean, and Laurent Heutte
University of Rouen
LITIS Lab
Saint Etienne du Rouvray, France
Email: {caroline.petitjean, laurent.heutte}@univ-rouen.fr

***Abstract*—The performance of most conventional classification systems relies on appropriate data representation and much of the efforts are dedicated to feature engineering, a difficult and time-consuming process that uses prior expert domain knowledge of the data to create useful features. On the other hand, deep learning can extract and organize the discriminative information from the data, not requiring the design of feature extractors by a domain expert. Convolutional Neural Networks (CNNs) are a particular type of deep, feedforward network that have gained attention from research community and industry, achieving empirical successes in tasks such as speech recognition, signal processing, object recognition, natural language processing and transfer learning. In this paper, we conduct some preliminary experiments using the deep learning approach to classify breast cancer histopathological images from BreaKHis, a publicly dataset available at http://web.inf.ufpr.br/vri/breast-cancer-database. We propose a method based on the extraction of image patches for training the CNN and the combination of these patches for final classification. This method aims to allow using the high-resolution histopathological images from BreaKHis as input to existing CNN, avoiding adaptations of the model that can lead to a more complex and computationally costly architecture. The CNN performance is better when compared to previously reported results obtained by other machine learning models trained with hand-crafted textural descriptors. Finally, we also investigate the combination of different CNNs using simple fusion rules, achieving some improvement in recognition rates.**

I. Introduction

NOWADAYS, cancer is a massive public health problem around the world. According to the International Agency for Research on Cancer (IARC), part of the World Health Organization (WHO), there were 8.2 million deaths caused by cancer in 2012 and 27 million of new cases of this disease are expected to occur until 2030 [1]. Among the cancer types, breast cancer (BC) is second most common for women, excluding skin cancer. Besides, the mortality of BC is very high when compared to other types of cancer. Even in face of recent advances in the comprehension of the molecular biology of BC progression and the discovery of new related molecular markers, the histopathological analysis remains the most widely used method for BC diagnosis [2]. Despite significant progress reached by diagnostic imaging technologies, the final BC diagnosis, including grading and staging, continues being done by pathologists applying visual inspection of histological samples under the microscope. Recent advances in image processing and machine learning techniques allow to build Computer-Aided Detection/Diagnosis (CAD/CADx) systems that can assist pathologists to be more productive, objective and consistent in diagnosis. Classification of histopathology images into distinct histopathology patterns, corresponding to the non-cancerous or cancerous condition of the analyzed tissue, is often the primordial goal in image analysis systems for cancer automatic aided diagnosis applications. The main challenge of such systems is dealing with the inherent complexity of histopathological images.

The automatic imaging processing for cancer diagnosis has been explored as a topic of research for more than 40 years [3] but is still challenging due to the complexity of the images to analyze. For example, Kowal *et al.* [4] compare and test different algorithms for nuclei segmentation, where the cases are classified as either benign or malignant on a dataset of 500 images, and report accuracies ranging from 96% to 100%. Filipczuk *et al.* [5] present a BC diagnosis system based on the analysis of cytological images of fine needle biopsies, to discriminate the images as either benign or malignant. Using four different classifiers trained with a 25-dimensional feature vector, they report a performance of 98% on 737 images. Similarly to [4] and [5], George *et al.* [6] propose a diagnosis system for BC based on the nuclei segmentation of cytological images. Using different machine learning models, such as neural networks and support vector machines, they report accuracy rates ranging from 76% to 94% on a dataset of 92 images. Zhang *et al.* [7] propose a cascade approach with rejection option. In the first level of the cascade, authors expect to solve the easy cases while the hard ones are sent to a second level where a more complex pattern classification system is used. They assess the proposed method on a database proposed by the Israel Institute of Technology, which is composed of 361 images and report results of 97% of reliability. In another work [8], the same authors assess an ensemble of one-class-classifiers on the same database achieving a recognition rate of 92%.

Most of these recent works related to BC classification are focused on Whole-Slide Imaging (WSI) [7], [8], [6], [4], [9]. However, the broad adoption of WSI and other forms of digital pathology still facing obstacles such as the high cost of implementing and operating the technology, insuffi-

CNN Architecture Overview

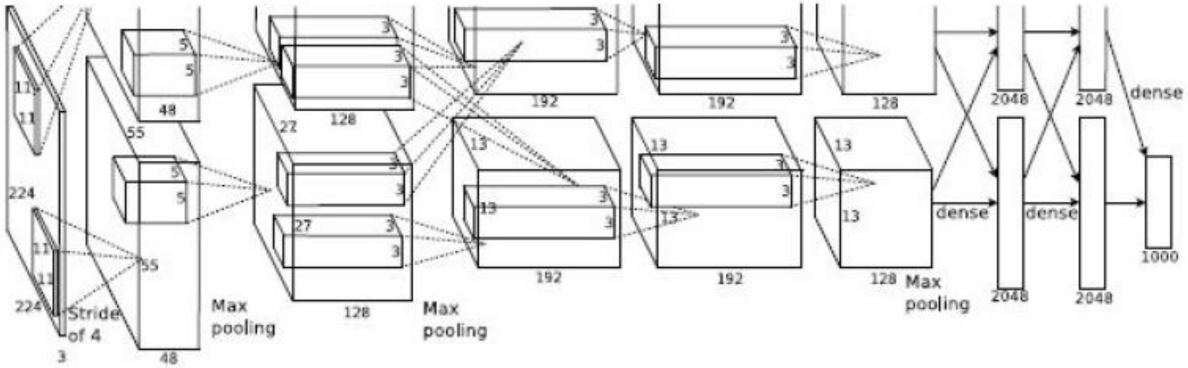

Accuracy of best performing network (based on AlexNet)

Accuracy at	Strategy	Magnification Factors			
		40×	100×	200×	400×
Image Level	1	79.9 ± 2.6	80.8 ± 3.7	**84.0 ± 3.2**	80.7 ± 1.8
	2	80.6 ± 2.1	81.0 ± 3.0	82.7 ± 1.9	**80.8 ± 3.1**
	3	81.8 ± 3.3	82.3 ± 4.9	82.4 ± 2.8	80.3 ± 4.0
	4	**89.6 ± 6.5**	**85.0 ± 4.8**	82.8 ± 2.1	80.2 ± 3.4

http://www.inf.ufpr.br/lesoliveira/download/IJCNN2016-BC.pdf

Custom Vision Model Performance

Default Training (<1 Minute)

- Training only the top fully connected of the neural net
- Basic hyperparameter tuning and data augmentation

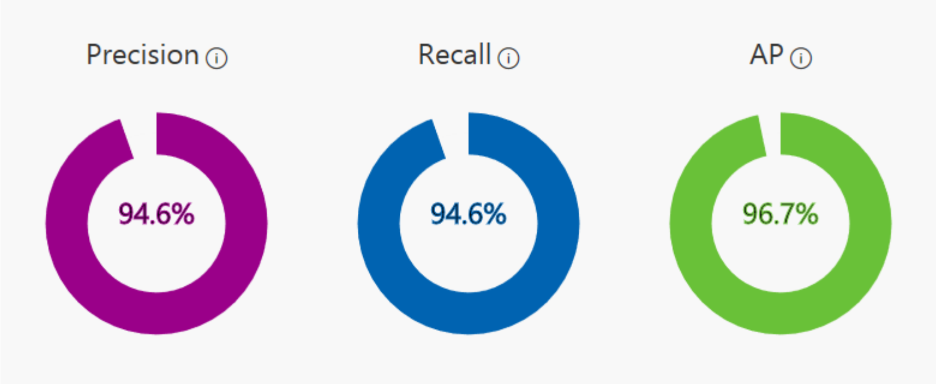

Performance Per Tag

Tag	Precision	Recall	A.P.	Image count
Malignant	95.4%	96.3%	97.4%	2157
Benign	92.9%	91.3%	94.9%	1150

Advanced Training (≈1h)

- Also fine tuning the last blocks of the base network
- Advanced hyperparameter tuning
- More data augmentation
- Different improvement strategies explored (depending on provided time budget)

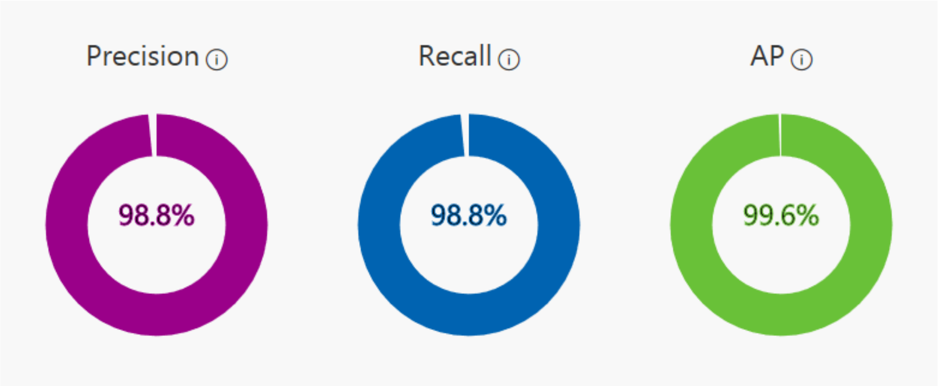

Performance Per Tag

Tag	Precision	Recall	A.P.	Image count
Malignant	99.1%	99.1%	99.8%	2157
Benign	98.3%	98.3%	99.1%	1150

Experiment: BC03, Performance metrics are based on cross-validation during training

Benchmarking on separate Test Set

Custom Vision (2h Advanced Training)

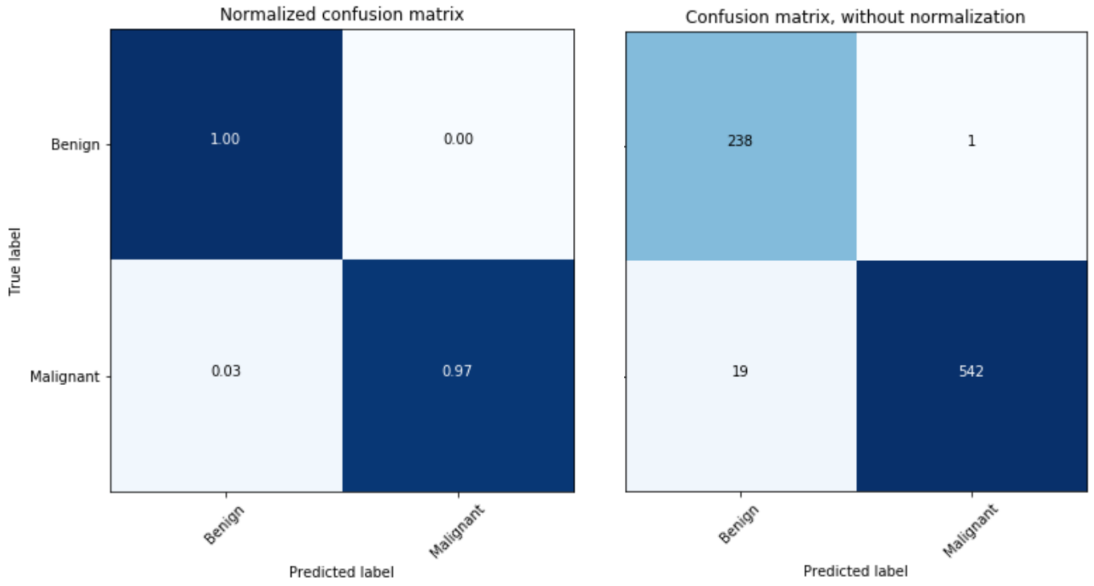

Classification report	Precision	Recall	F1
Benign	0.93	1.00	0.96
Malignant	1.00	0.97	0.98
Micro avg.	0.97	0.97	0.97
Macro avg.	0.96	0.98	0.97
Weighted avg.	0.98	0.97	0.98

Customized Transfer Learning Model (effort: 7 days)

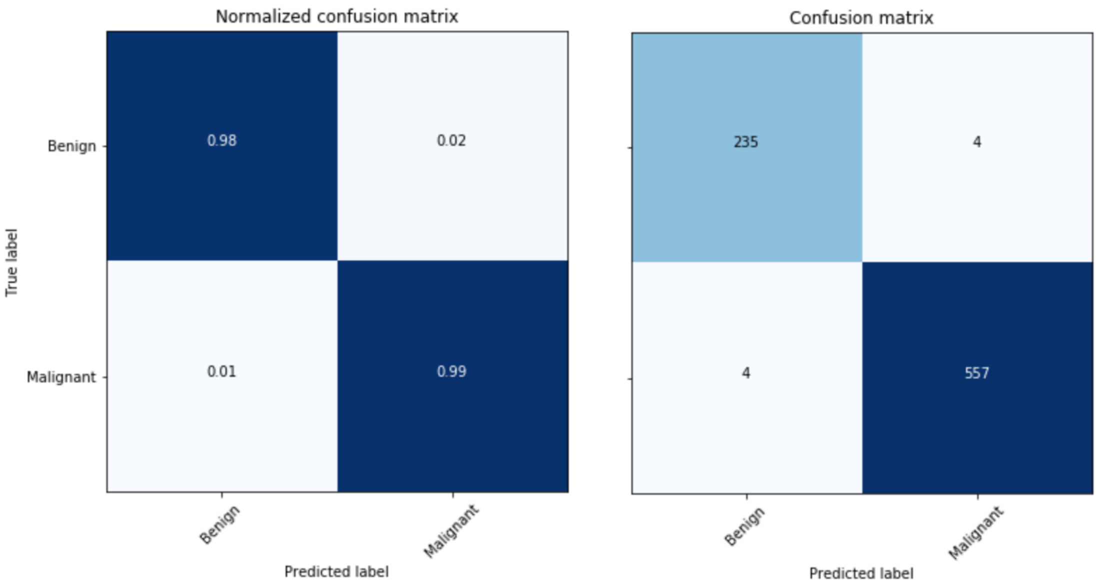

Classification report	Precision	Recall	F1
Benign	0.98	0.98	0.98
Malignant	0.99	0.99	0.99
Micro avg.	0.99	0.99	0.99
Macro avg.	0.99	0.99	0.99
Weighted avg.	0.99	0.99	0.99

Experiment: BC03, Performance metrics are based on separated test set (n=800), magnification factors: 40x, 100x, 200x, 400x

Classifying Malignant Tumor Categories

Finished training on **27.3.2019, 13:49:53** using **General** domain
Classification type: **Multiclass (Single tag per image)**

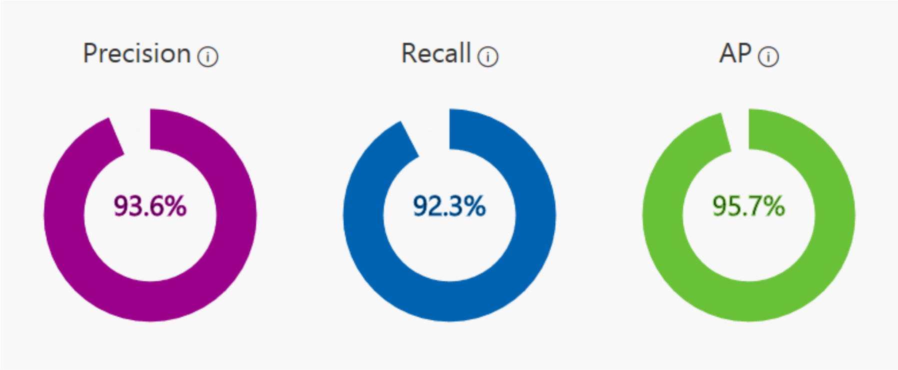

Performance Per Tag

Tag	Precision	Recall	A.P.	Image count
Lobular Carcinoma	95.8%	95.8%	98.4%	594
Ductal Carcinoma	94.9%	82.8%	92.8%	786
Papillary Carcinoma	92.1%	98.1%	97.2%	536
Mucinous Carcinoma	91.8%	95.4%	96.2%	761

Lobular Carcinoma

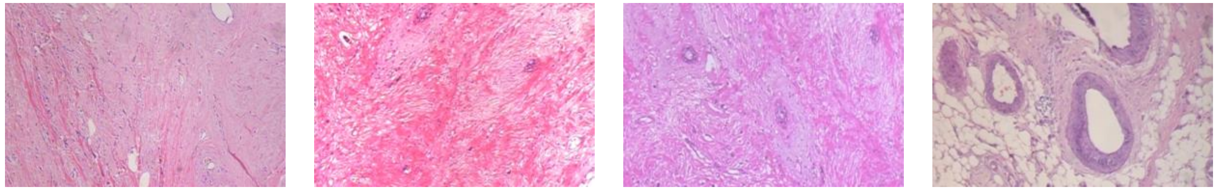

Ductal Carcinoma

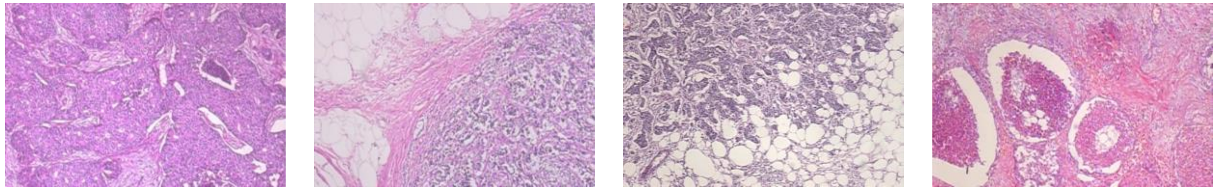

Papillary Carcinoma

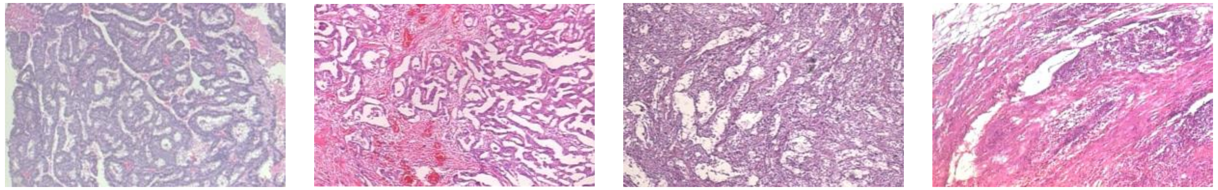

Mucinous Carcinoma

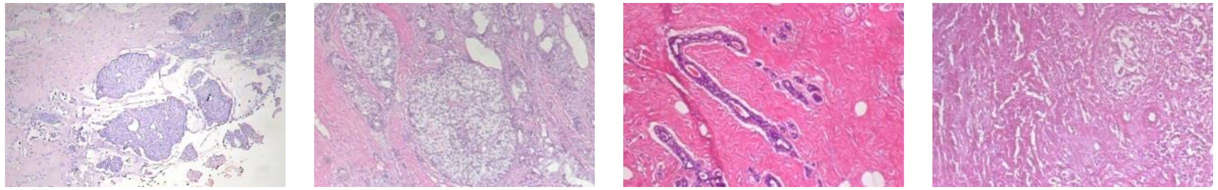

Experiment: BC04, Performance metrics are based on cross-validation during training

Data Sets

Breast – Cancer „BreaKHis_v1“

https://github.com/jhole89/classifying-cancer/tree/master/cnn_image_classifier

Bees vs. Ants = „hymenoptera_data“

https://www.kaggle.com/ajayrana/hymenoptera-data

Microsoft

AI

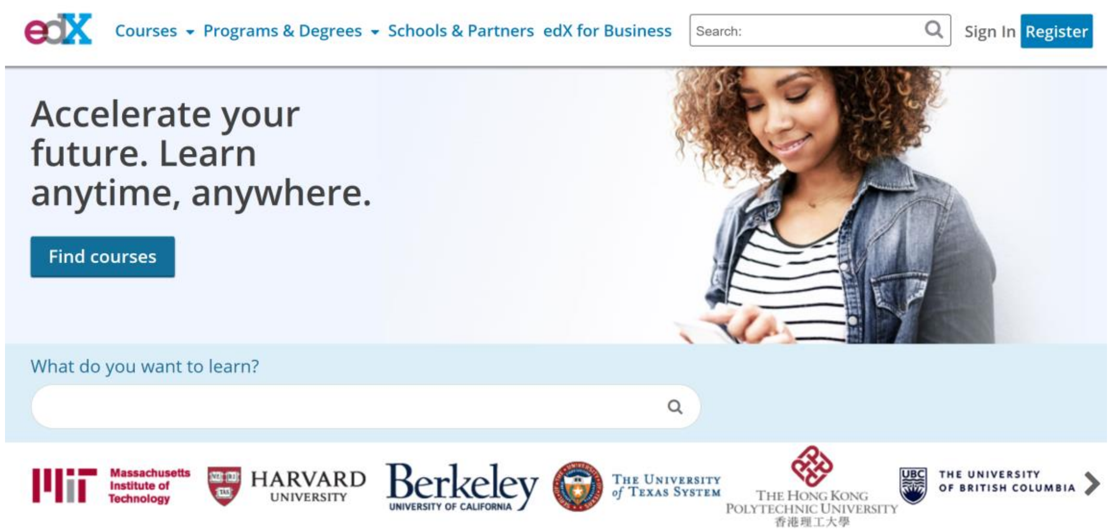

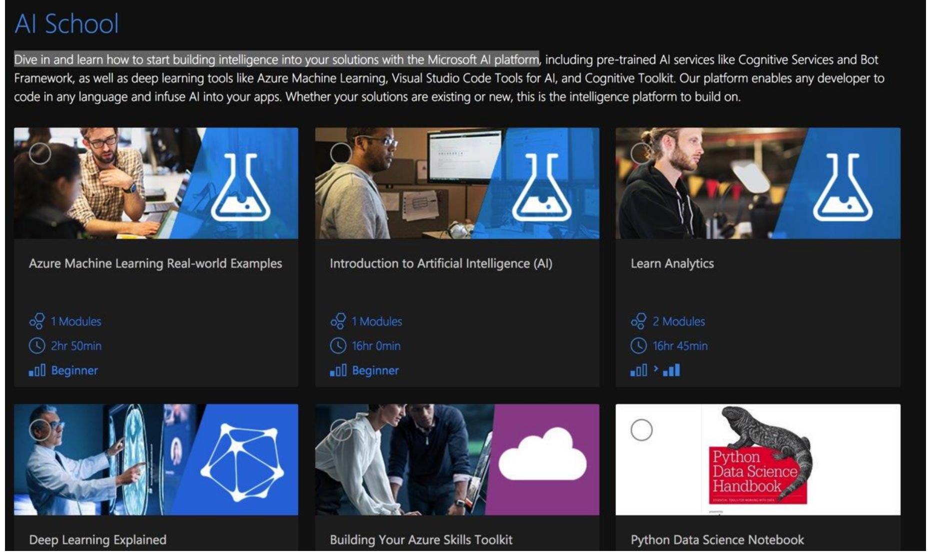

https://www.edx.org/

https://aischool.microsoft.com/

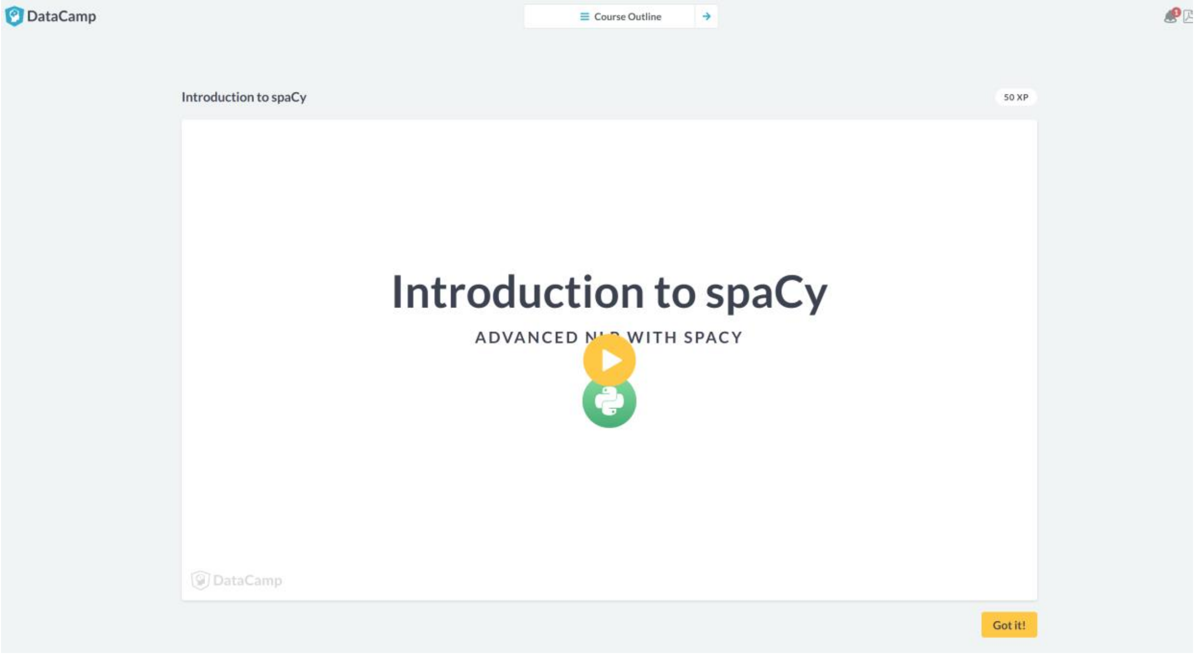

https://www.datacamp.com/

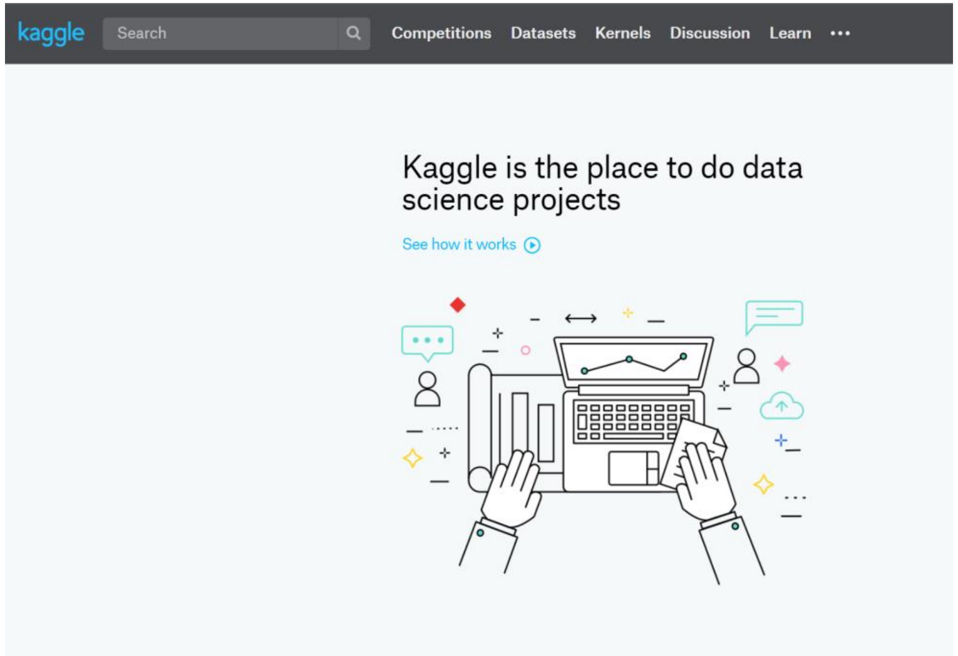

https://www.kaggle.com

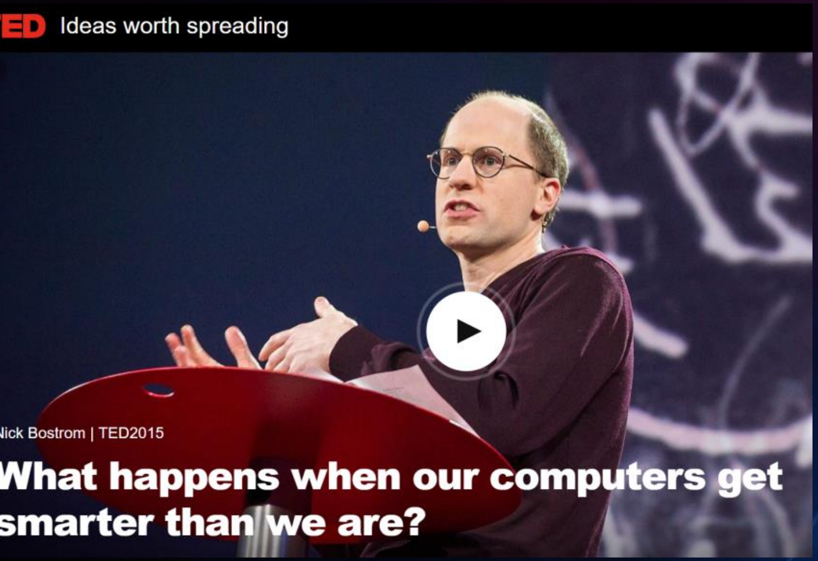
ED Ideas worth spreading
Nick Bostrom | TED2015
What happens when our computers get smarter than we are?

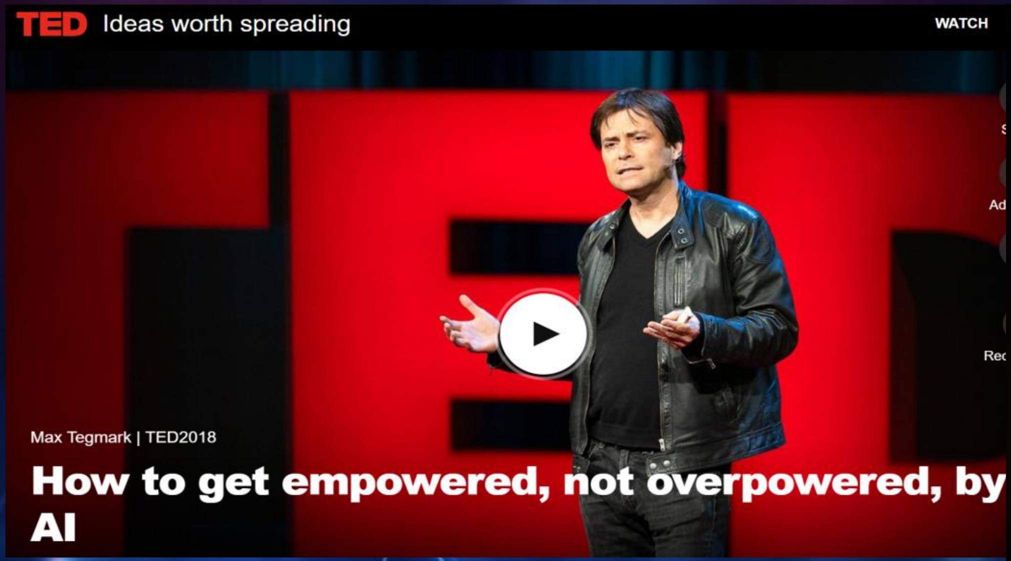
TED Ideas worth spreading
WATCH
Max Tegmark | TED2018
How to get empowered, not overpowered, by AI

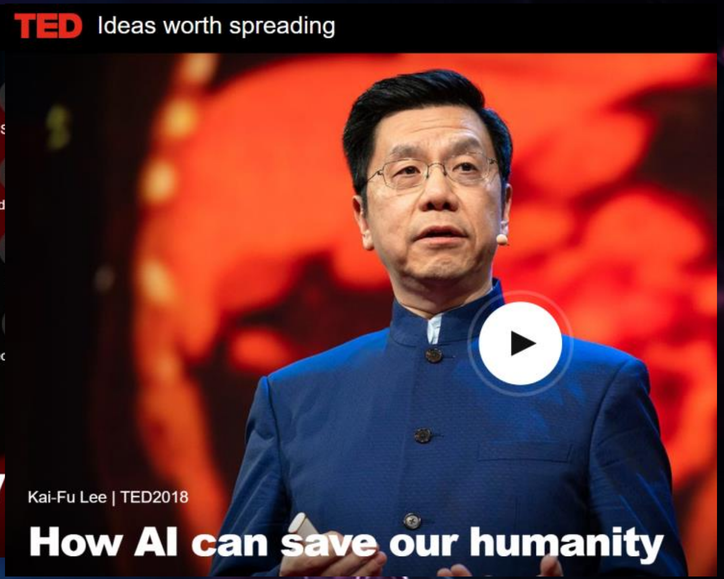
TED Ideas worth spreading
Kai-Fu Lee | TED2018
How AI can save our humanity

THANK YOU


BLACK MIRROR

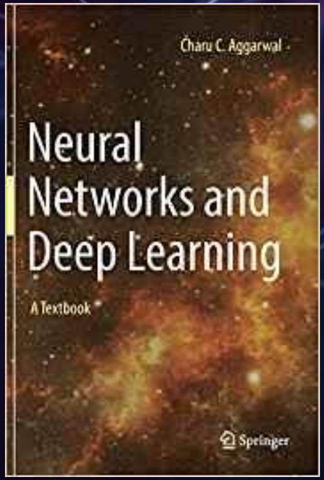
Charu C. Aggarwal
Neural Networks and Deep Learning
A Textbook
Springer

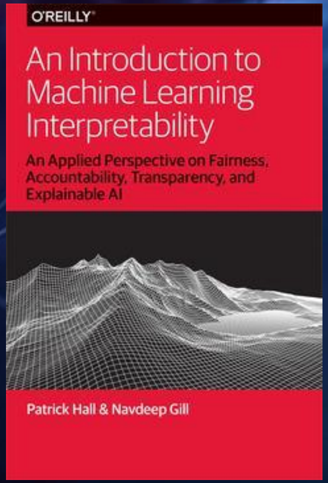
O'REILLY
An Introduction to Machine Learning Interpretability
An Applied Perspective on Fairness, Accountability, Transparency, and Explainable AI
Patrick Hall & Navdeep Gill

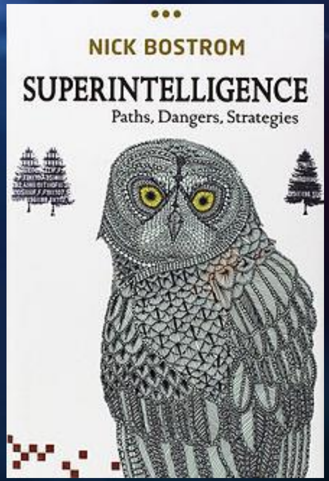
NICK BOSTROM
SUPERINTELLIGENCE
Paths, Dangers, Strategies

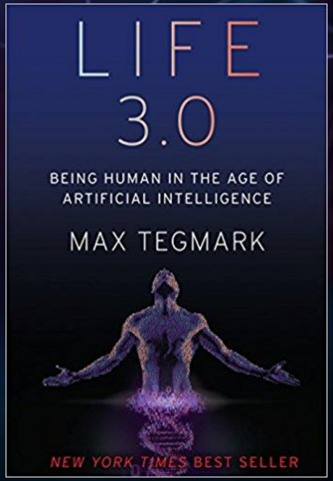
LIFE 3.0
BEING HUMAN IN THE AGE OF ARTIFICIAL INTELLIGENCE
MAX TEGMARK
NEW YORK TIMES BEST SELLER

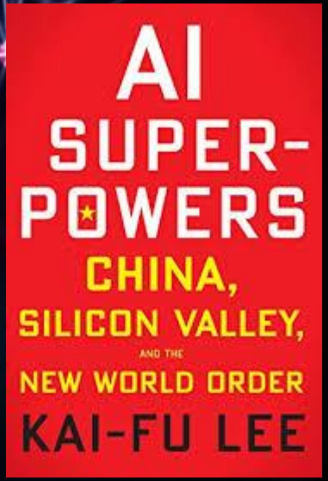
AI SUPER-POWERS
CHINA, SILICON VALLEY,
AND THE
NEW WORLD ORDER
KAI-FU LEE

Microsoft Azure