جامعة الملك عبدالله
للعلوم والتقنية
King Abdullah University of
Science and Technology

The Machine Learning (ML) Hub @ KAUST

Prof. Bernard Ghanem

The ML Hub Organizers

Prof. Panagiotis Kalnis
CS, ECRC

Prof. Marco Canini
CS, ECRC

Prof. Peter Richtarik
CS, AMCS, STAT, VCC

Prof. Bernard Ghanem
EE, CS, STAT, VCC

ml.kaust.edu.sa

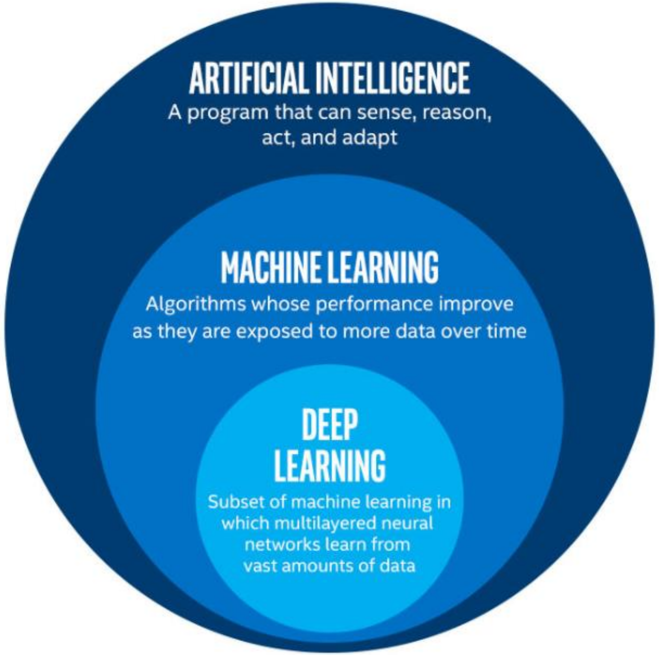
ARTIFICIAL INTELLIGENCE
A program that can sense, reason, act, and adapt
MACHINE LEARNING
Algorithms whose performance improve as they are exposed to more data over time
DEEP LEARNING
Subset of machine learning in which multilayered neural networks learn from vast amounts of data

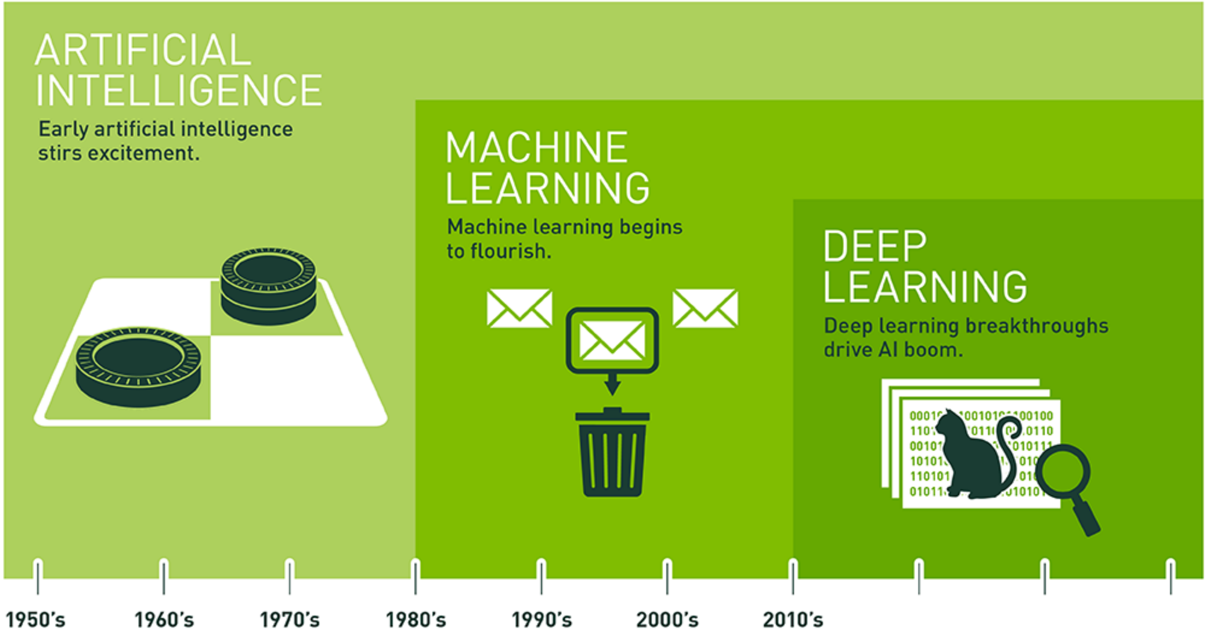
ARTIFICIAL INTELLIGENCE
Early artificial intelligence stirs excitement.
MACHINE LEARNING
Machine learning begins to flourish.
DEEP LEARNING
Deep learning breakthroughs drive AI boom.
1950's
1960's
1970's
1980's
1990's
2000's
2010's

Google CEO: A.I. is more important than fire or electricity

Catherine Clifford | 12:56 PM ET Thu, 1 Feb 2018

Beck Diefenbach | Reuters

Google CEO Sundar Pichai takes the stage during the presentation of new Google hardware in San Francisco on Oct. 4, 2016.

CXO

AI skills reign supreme in the fastest-growing jobs of the year

Six out of the 15 top emerging jobs in 2018 were related to artificial intelligence, according to LinkedIn.

By Macy Bayern | December 13, 2018, 6:07 AM PST

The big takeaways for tech leaders:

- *AI has a prominent presence on the emerging jobs list, making up six of the 15 fastest-growing jobs of the year. — LinkedIn, 2018*
- *The top five emerging jobs of 2018 include blockchain developer, machine learning engineer, application sales executive, machine learning specialist, and professional medical representative. — LinkedIn, 2018*

POLITICS FEBRUARY 11, 2019 / 8:06 AM / UPDATED 5 HOURS AGO

REUTERS

Trump administration unveils order to prioritize and promote AI

3 MIN READ

WASHINGTON (Reuters) - U.S. President Donald Trump on Monday signed an executive order asking federal government agencies to dedicate more resources and investment into research, promotion and training on artificial intelligence, known as AI.

AI Index Report 2018 (aiindex.org)

Our Mission is to ground the conversation about AI in data.

The AI Index is an effort to track, collate, distill, and visualize data relating to artificial intelligence. It aspires to be a comprehensive resource of data and analysis for policymakers, researchers, executives, journalists, and the general public to develop intuitions about the complex field of AI.

Impact of AI on Research

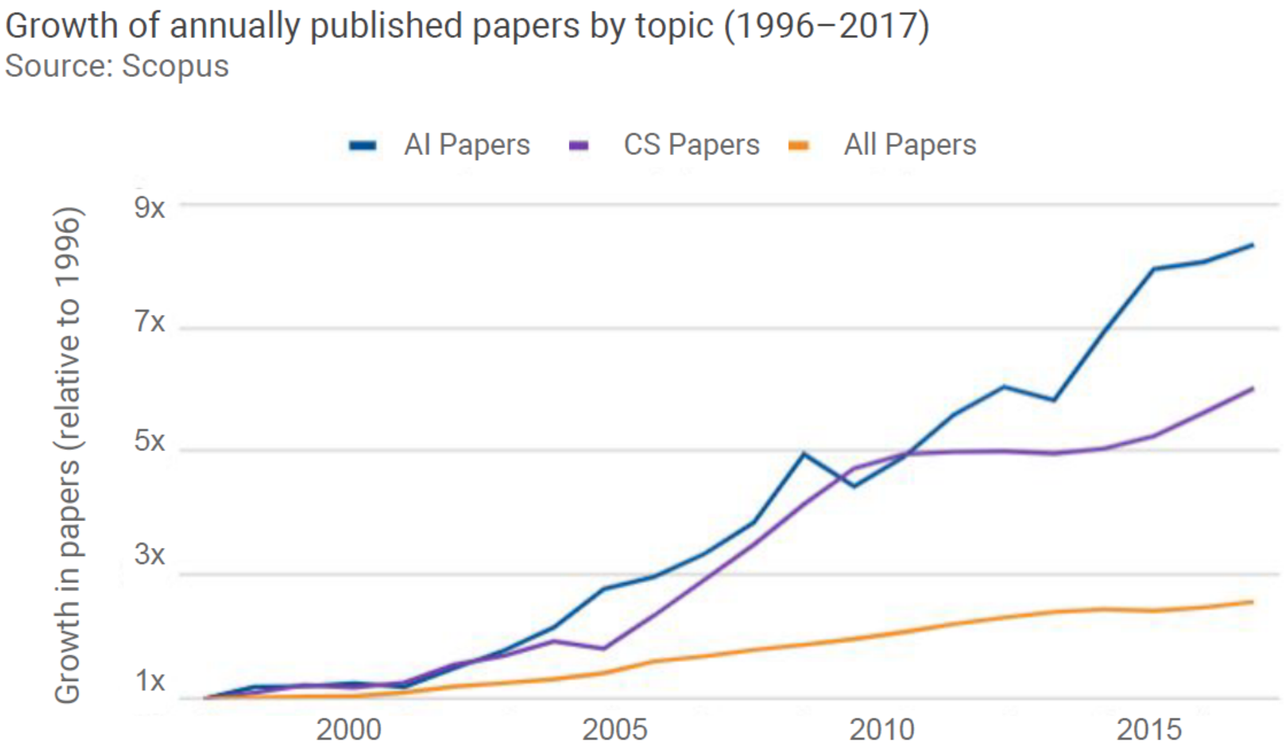

Impact of AI on Research

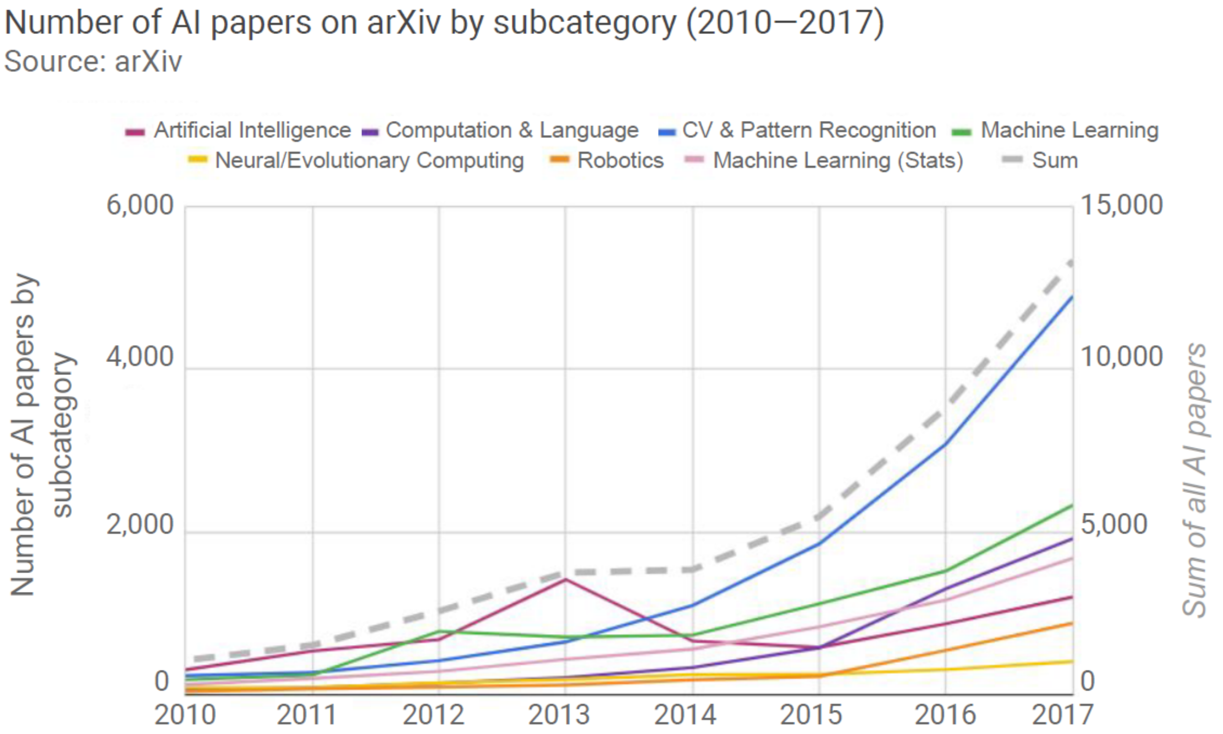

Impact of AI on Research (interdisciplinary)

Relative Activity Focus by Region and AI Research Sector in 2000 and 2017

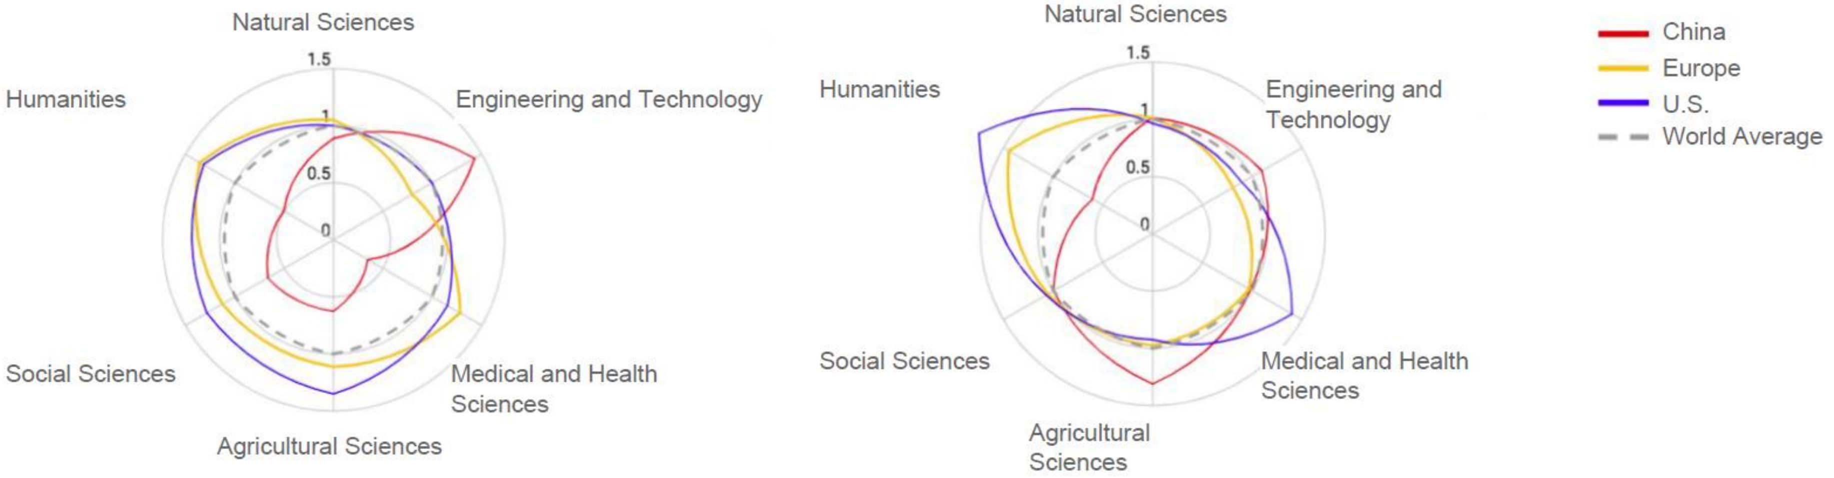

Impact of AI on Education/Training

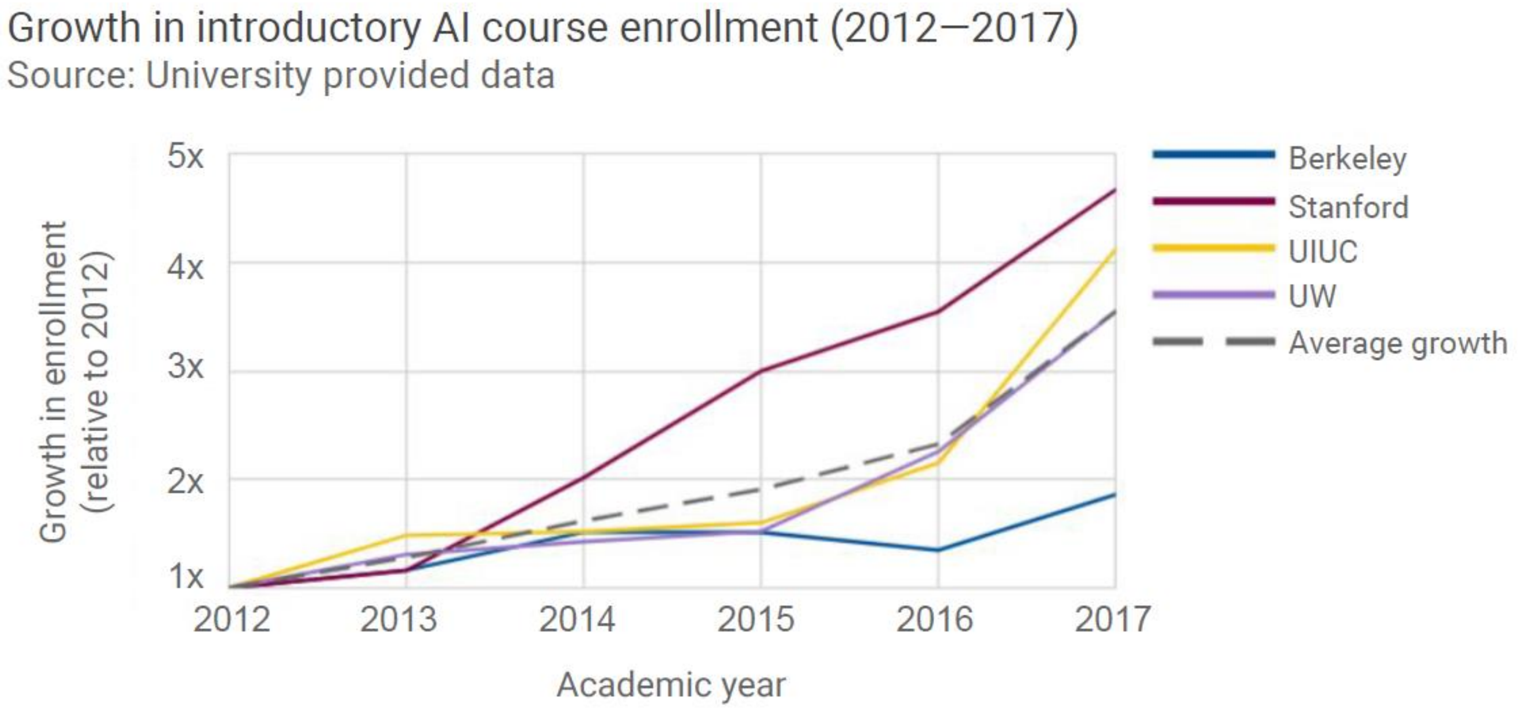

Impact of AI on Education/Training

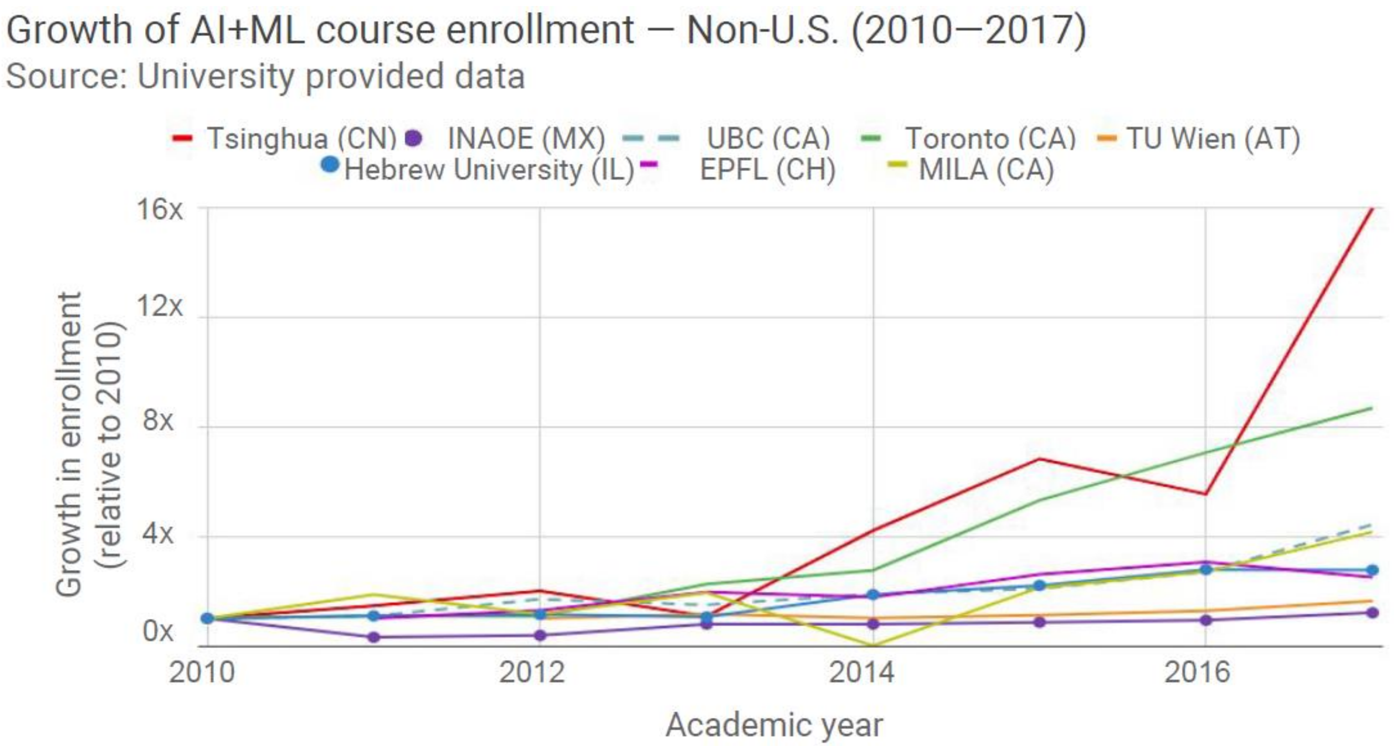

Impact of AI on Education/Training

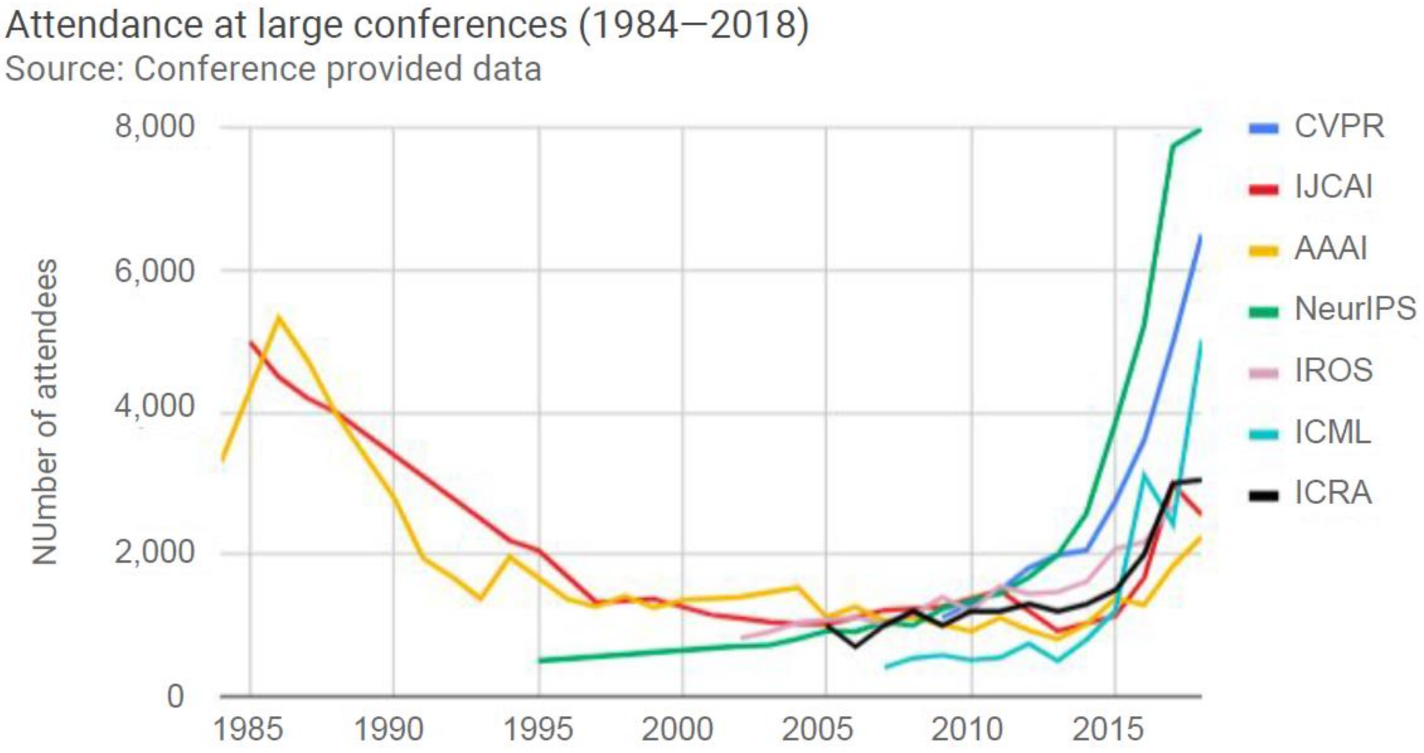

Impact of AI on Education/Training

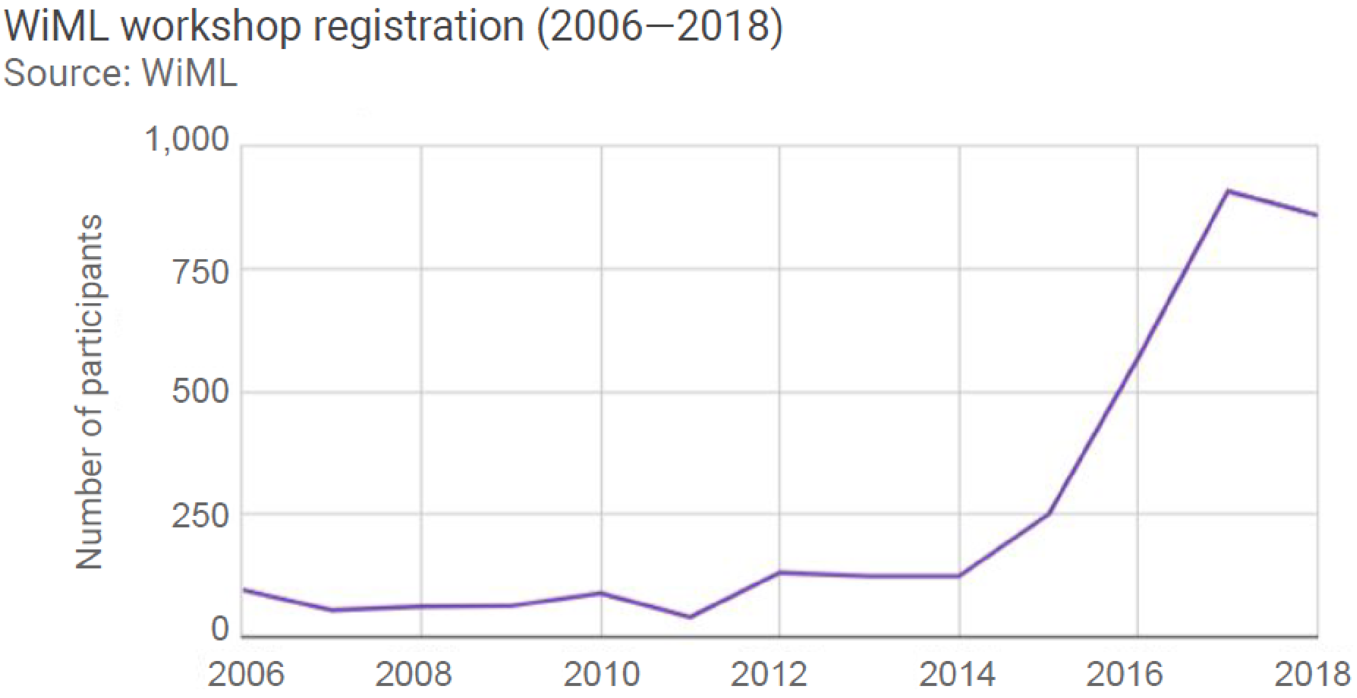

Impact of AI on IED/Jobs

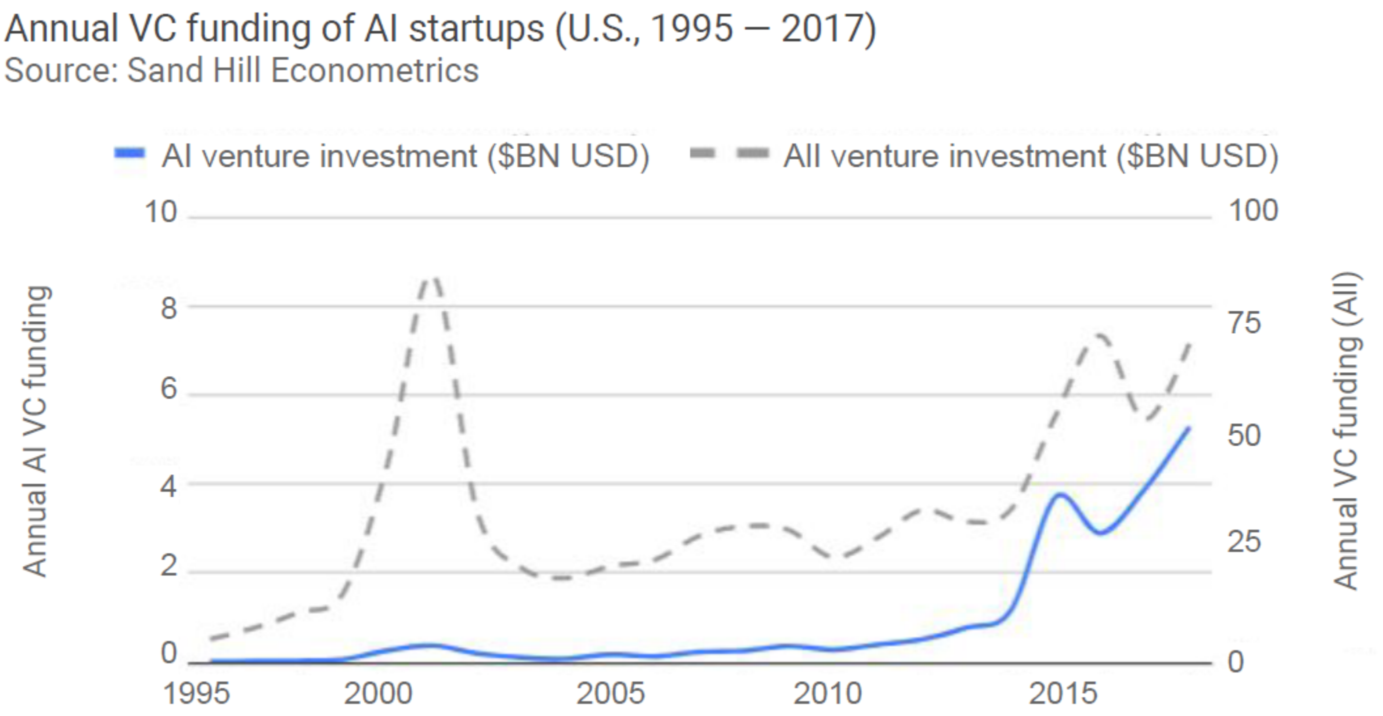

Impact of AI on IED/Jobs

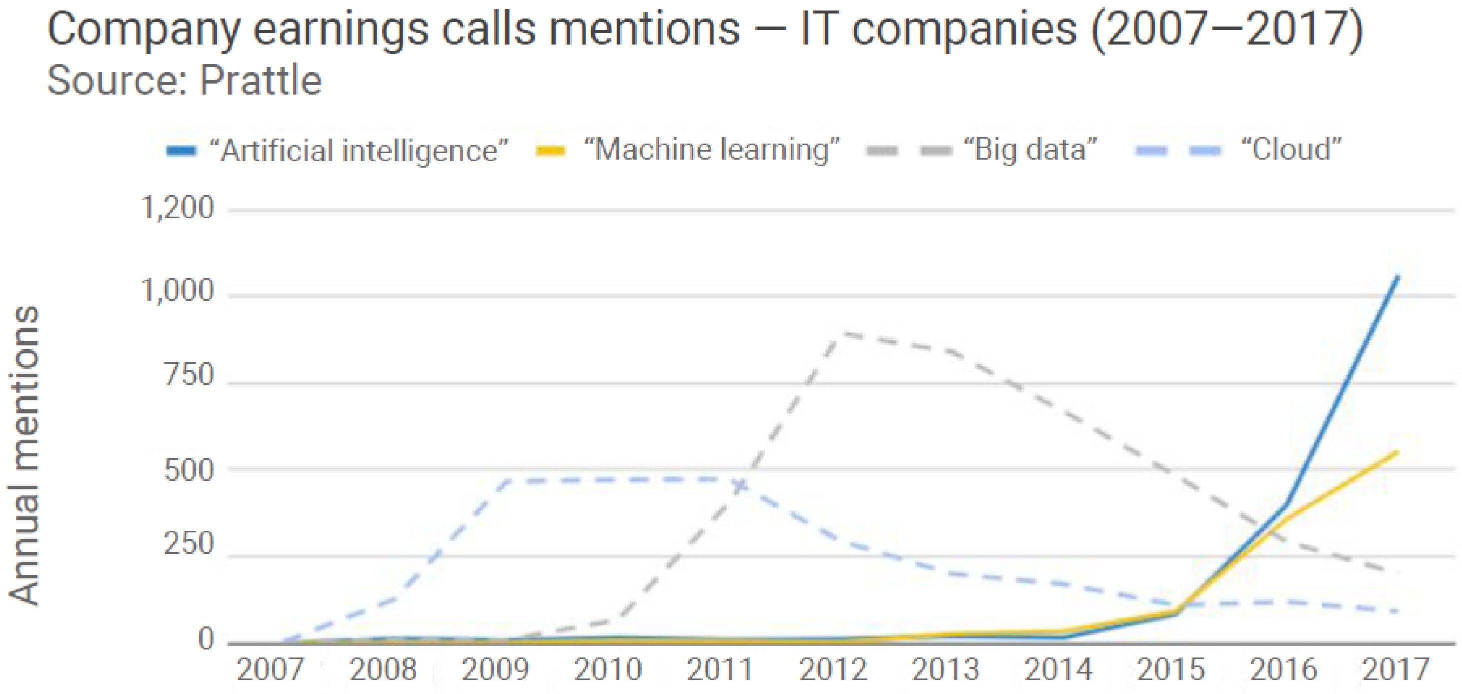

جامعة الملك عبدالله
للعلوم والتقنية
King Abdullah University of
Science and Technology

Impact of AI on IED/Jobs

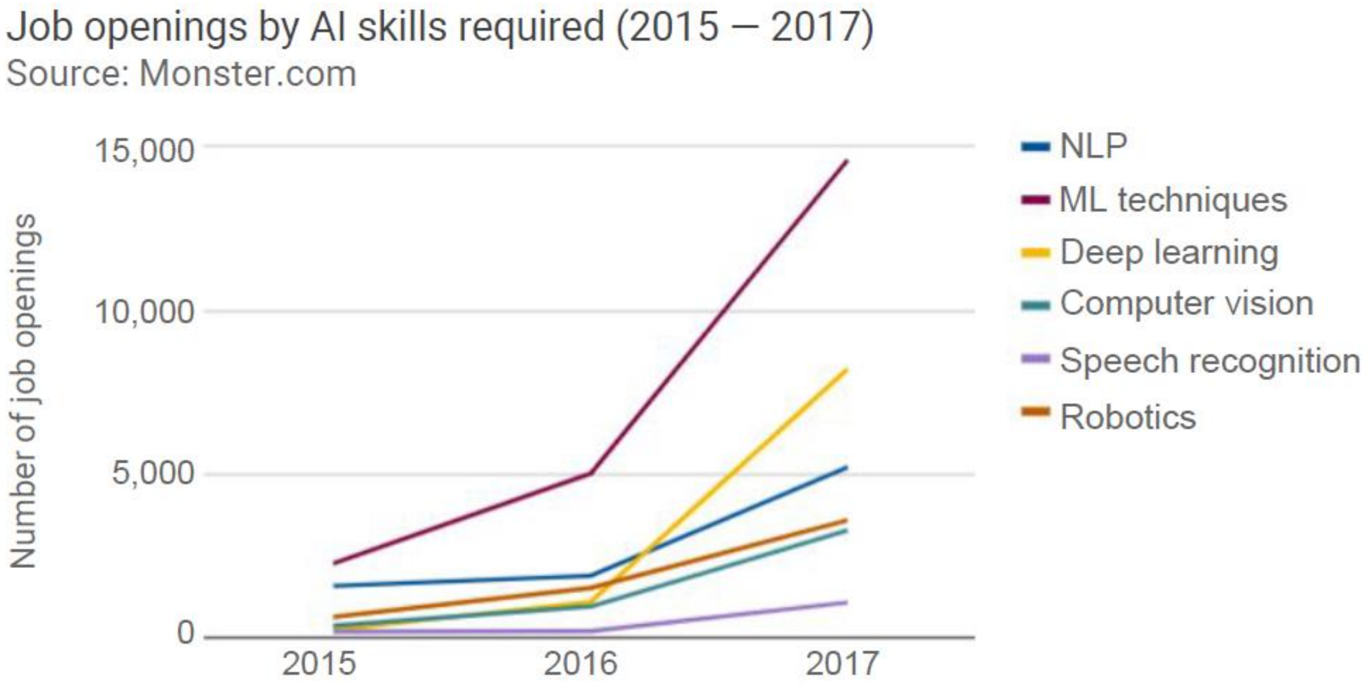

Impact of AI on IED/Jobs

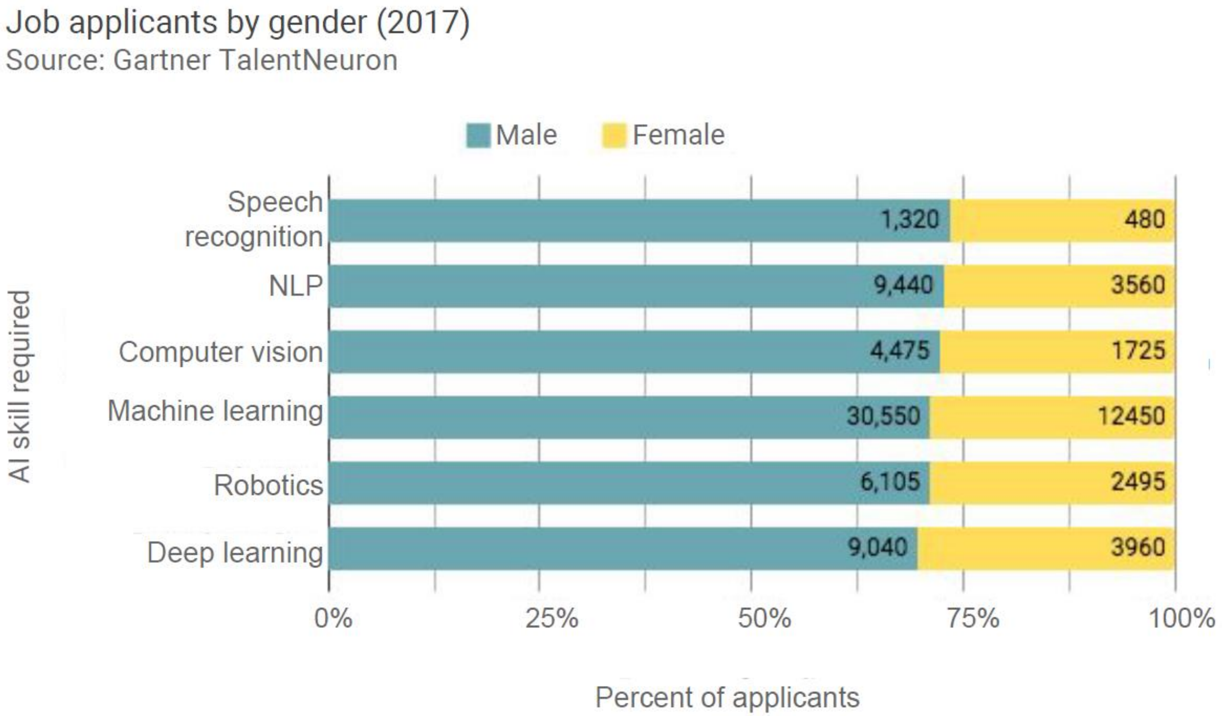

Why has AI/ML been so impactful?

Perfect “storm”

- Large-scale data
- Tools (e.g. deep networks)
- Mathematical foundations
- Hardware/Software systems (e.g. GPUs, TensorFlow)
- Applications and industry adoption

Where is KAUST in all this?

AMiner Conference & Journal Ranking in: Computer Science / AI & Pattern Recognition

Rank	Conference (Full Name)		Short Name	H5-Index
1	IEEE Conference on Computer Vision and Pattern Recognition	★	CVPR	112.00
2	IEEE Transactions on Pattern Analysis and Machine Intelligence	★	TPAMI	101.00
3	Expert Systems with Applications		Expert Syst. Appl.	59.00
4	International Journal of Computer Vision	★	IJCV	58.00
5	IEEE International Conference on Robotics and Automation	★	ICRA	58.00
6	International Conference on Computer Vision	★	ICCV	58.00
7	International Conference on Machine Learning	★	ICML	56.00
8	Annual Conference on Neural Information Processing Systems	★	NIPS	51.00
9	Journal of Machine Learning Research	★	JMLR	49.00
10	Annual Meeting of the Association for Computational Linguistics		ACL	48.00
11	IEEE Transactions on Audio, Speech, and Language Processing		TASLP	46.00
12	Conference on Empirical Methods in Natural Language Processing		EMNLP	45.00
13	AAAI Conference on Artificial Intelligence	★	AAAI	44.00
14	IEEE Transactions on Fuzzy Systems		TFS	44.00
15	Decision Support Systems		DSS	43.00
16	Neurocomputing			39.00
17	IEEE Transactions on Evolutionary Computation		TEVC	38.00
18	Pattern Recognition Letters		PRL	37.00
19	Autonomous Agents and Multi-Agent Systems		AAMAS	36.00
20	International Joint Conference on Artificial Intelligence	★	IJCAI	35.00

https://aminer.org/ranks/conf

Predicting Binding Specificity of RNA Constituents

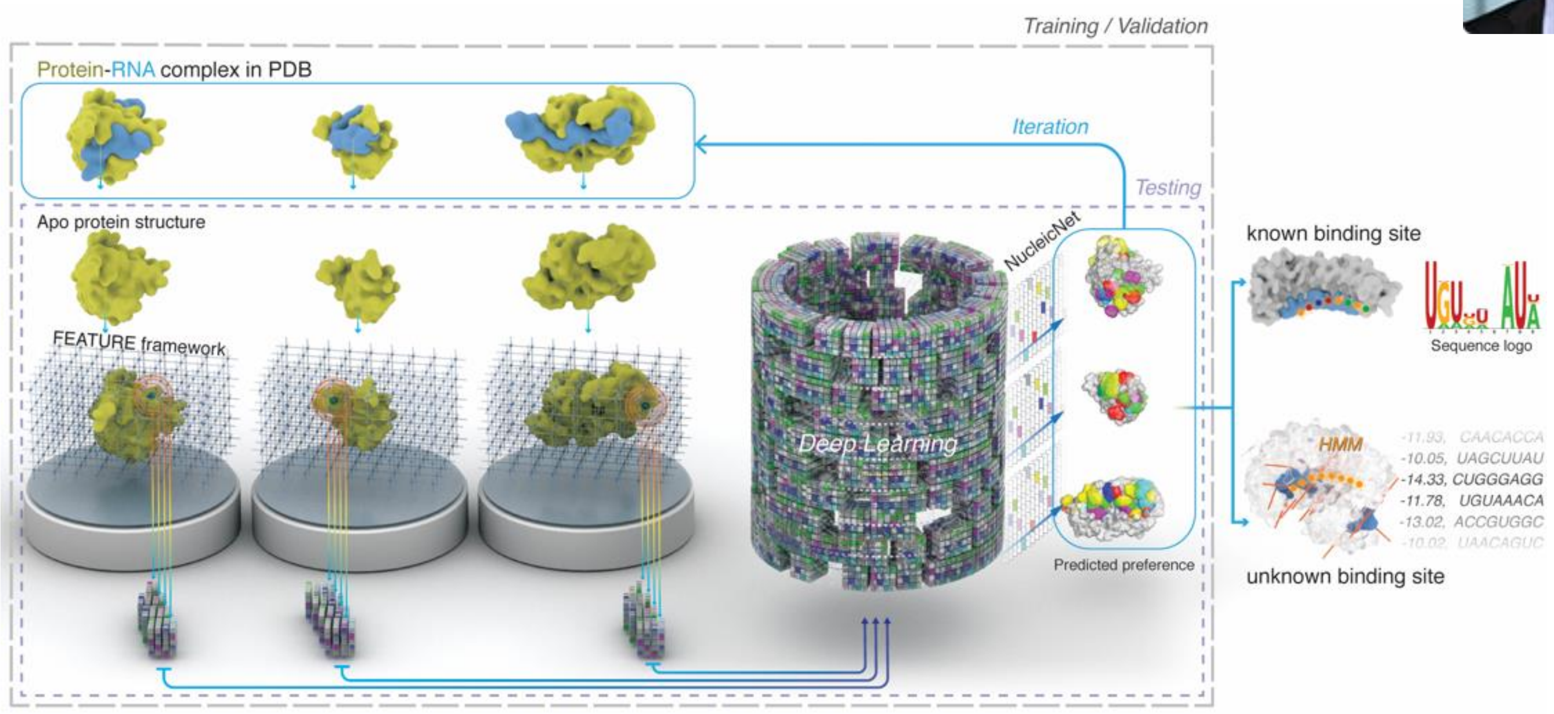

Learn to Profile Users in Social Media

Input: A stream of tweets generated across the time

Twitter Users

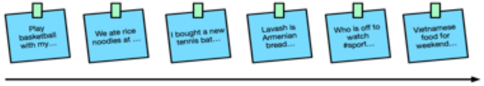

Tweets over time

Output: A set of **keywords** to profile the user at different point in time

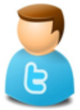

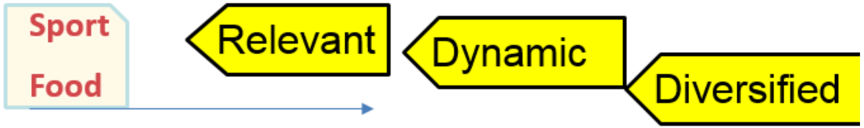

Tags over time

Input: Attributed Network G^t

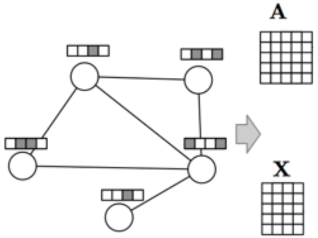

Output:

- **user** representation
- **attribute** (e.g., keyword) representation

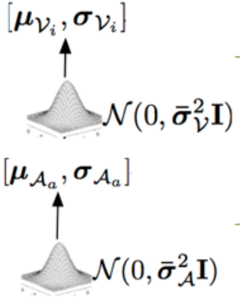

compare in the same space

Two papers published at
[SIGKDD 2018, WSDM 2019]

Prediction of preignition events in spark ignited engines using deep learning

- Pre-ignition is defined as ignition of the fuel/air mixture prior to spark plug firing
- Highly stochastic event that can lead to engine destruction
- We developed a deep learning model to predict preignition early in the engine cycle to avoid damage.

Multivariate time-series data

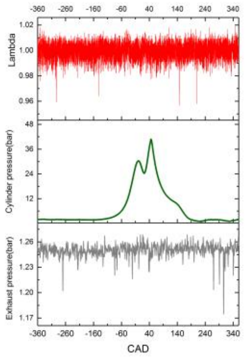

Supervised Binary Classification Problem

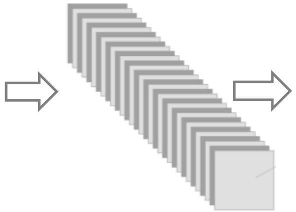

CNN layers

Hard Negative Mining

Output

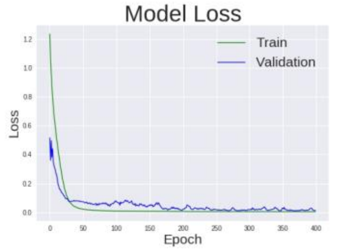

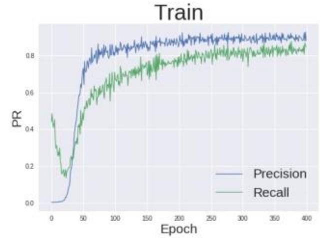

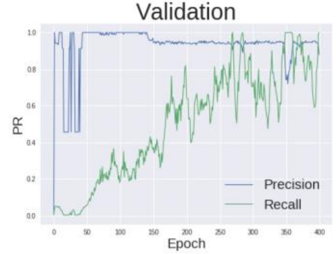

Federated Learning

TASK: Train machine learning models on millions of mobile devices without sharing private data

Collaboration with Google

Downloaded by 1bn+ Android users

Horizon 2020 call dedicated to FL

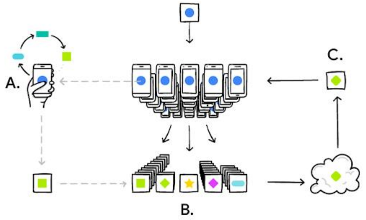

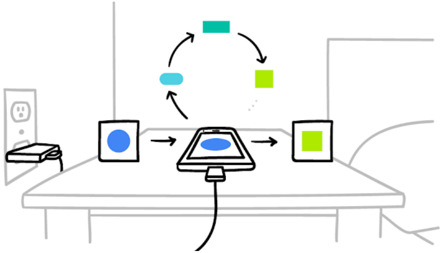

- Forbes
- The Verge
- Quartz
- TechRepublic
- ZDNet
- Computer Business Review
- Motherboard Vice
- Infoworld
- Silicon.co.uk
- Venturebeat
- Engadget
- Tech Narratives
- GadgetHacks
- BGR
- AndroidAuthority
- AndroidHeadlines
- Tom's Guide
- Digital Trends
- The Exponential View
- vvcat
- 9to5google

The CEO of Google, Sundar Pichai, said (May 2017):

"... **We continue to set the pace in machine learning and AI research. We introduced a new technique for training deep neural networks on mobile devices called Federated Learning.** This technique enables people to run a shared machine learning model, while keeping the underlying data stored locally on mobile phones..."

NIPS 2016

Federated Learning: Strategies for Improving Communication Efficiency

Jakub Konečný, H. Brendan McMahan, Felix X. Yu, Ananda Theertha Suresh & Dave Bacon
Google

Peter Richtárik
King Abdullah University of Science and Technology (KAUST), Thuwal, Saudi Arabia
University of Edinburgh, Edinburgh, Scotland

Abstract

1 Introduction

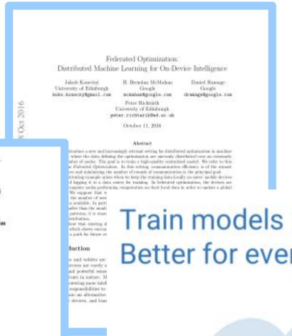

Federated Optimization:
Distributed Machine Learning for On-Device Intelligence

Jakub Konečný, University of Edinburgh
H. Brendan McMahan, Google
Daniel Ramage, Google
Peter Richtárik, University of Edinburgh

October 11, 2016

Abstract

https://ai.googleblog.com/2017/04/federated-learning-collaborative.html

Scale ML to 1000's of GPUs

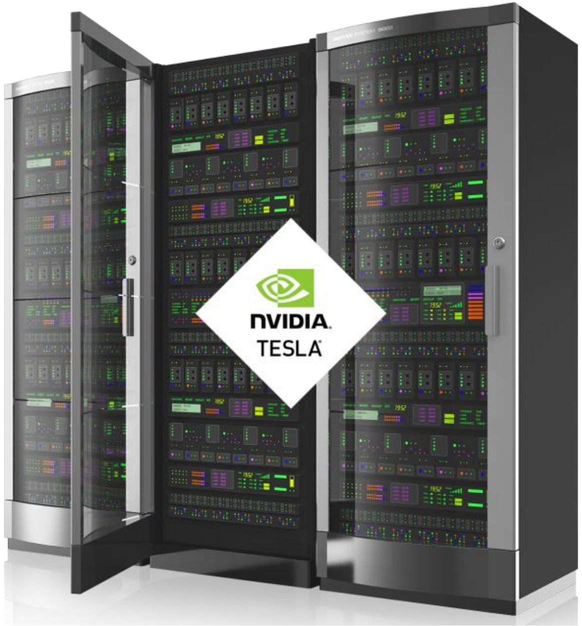

- Accelerate training
  - Minimize communication
  - Use very large mini-batch sizes
  - Split a single model in multiple GPUs
- Accelerate Serving (inference)
  - Many instances of the same model
  - Exploit common computation paths
- ML for scientific data

Scaling Distributed ML with In-Network Aggregation

Distributed ML training scales poorly due to communication costs

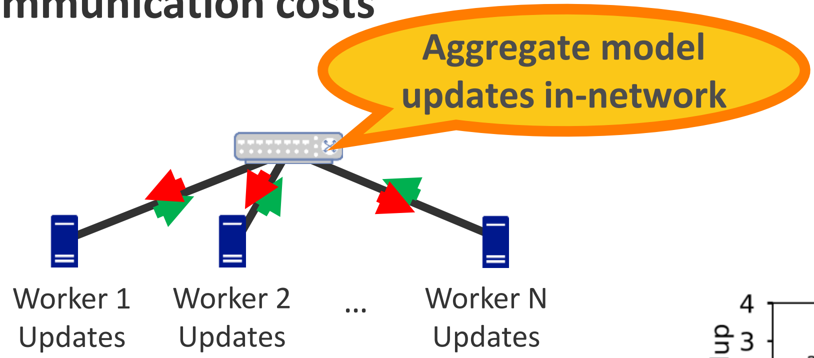

Key idea: co-design ML and networking to maximize communication efficiency

https://sands.kaust.edu.sa/daiet/

Designed and built **SwitchML**

Workers stream model updates to **a 6.4 Tbps programmable switch** that aggregates and redistributes combined parameters

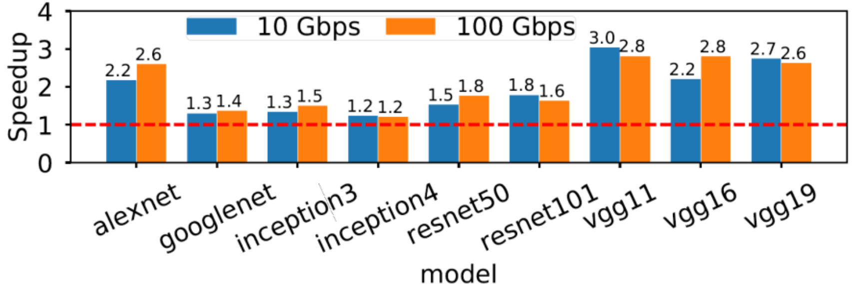

SwitchML speeds up training by up to 300%

A KAUST collaboration with

Microsoft

BAREFOOT NETWORKS

UNIVERSITY of WASHINGTON

Fish Species Classification Underwater

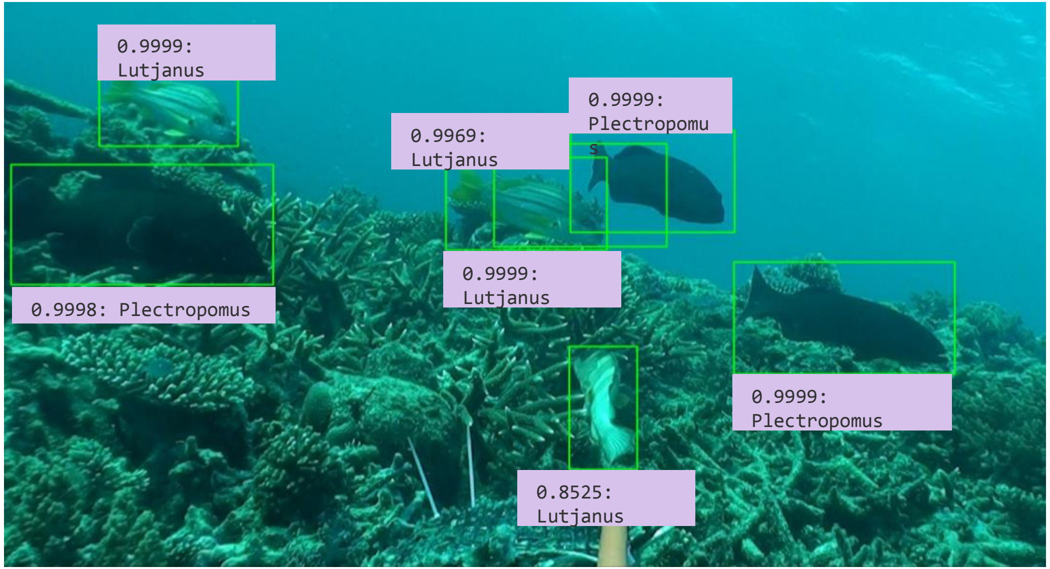

Tools and Foundations for Understanding Images and Videos

Video Understanding

- Activity Detection
- Efficient Search
- Object Tracking

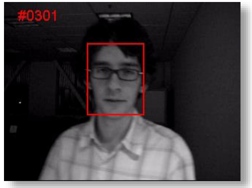

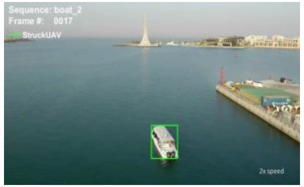

ACTIVITYNET

Vision for Automation

- Sim4CV
- Transfer Learning
- Applications

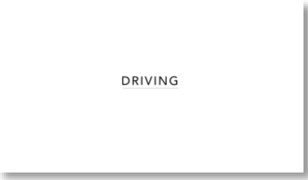

Fundamentals

- Optimization for CV&ML (sparse, low-rank, integer)
- Deep DNN Understanding

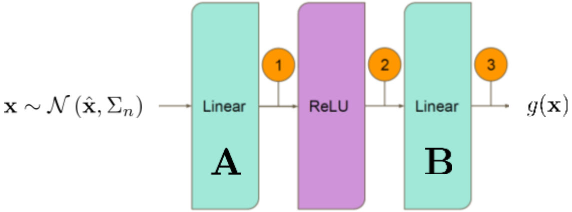

$$\min_{\mathbf{x}} \; f(\mathbf{x}) \quad \text{s.t.} \;\; \mathbf{x} \in \{1, -1\}^n ; \mathbf{x} \in \Omega$$

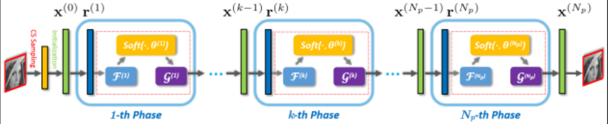

Large-Scale Activity Detection

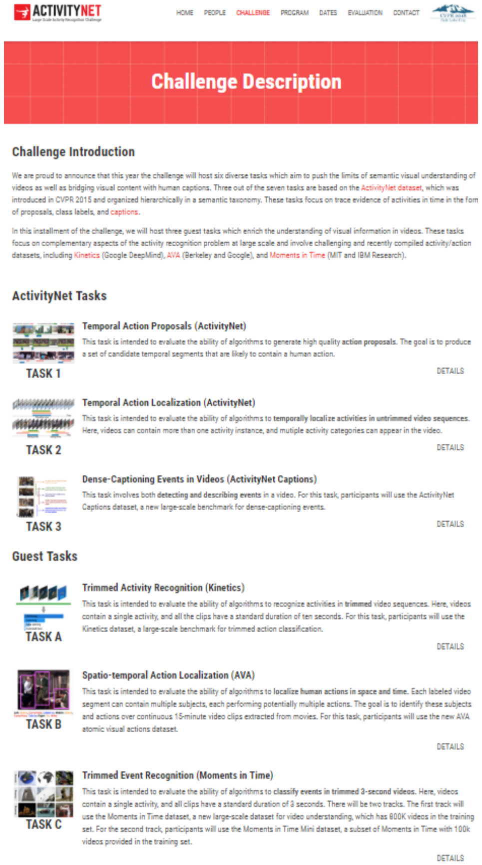

ACTIVITYNET

HOME PEOPLE CHALLENGE PROGRAM DATES EVALUATION CONTACT

Challenge Description

Challenge Introduction

We are proud to announce that this year the challenge will host six diverse tasks which aim to push the limits of semantic visual understanding of videos as well as bridging visual content with human captions. Three out of the seven tasks are based on the ActivityNet dataset, which was introduced in CVPR 2015 and organized hierarchically in a semantic taxonomy. These tasks focus on trace evidence of activities in time in the fom of proposals, class labels, and captions.

In this installment of the challenge, we will host three guest tasks which enrich the understanding of visual information in videos. These tasks focus on complementary aspects of the activity recognition problem at large scale and involve challenging and recently compiled activity/action datasets, including Kinetics (Google DeepMind), AVA (Berkeley and Google), and Moments in Time (MIT and IBM Research).

ActivityNet Tasks

TASK 1

Temporal Action Proposals (ActivityNet)

This task is intended to evaluate the ability of algorithms to generate high quality **action proposals**. The goal is to produce a set of candidate temporal segments that are likely to contain a human action.

DETAILS

TASK 2

Temporal Action Localization (ActivityNet)

This task is intended to evaluate the ability of algorithms to **temporally localize activities in untrimmed video sequences**. Here, videos can contain more than one activity instance, and mutiple activity categories can appear in the video.

DETAILS

TASK 3

Dense-Captioning Events in Videos (ActivityNet Captions)

This task involves both **detecting and describing events** in a video. For this task, participants will use the ActivityNet Captions dataset, a new large-scale benchmark for dense-captioning events.

DETAILS

Guest Tasks

TASK A

Trimmed Activity Recognition (Kinetics)

This task is intended to evaluate the ability of algorithms to recognize activities in **trimmed** video sequences. Here, videos contain a single activity, and all the clips have a standard duration of ten seconds. For this task, participants will use the Kinetics dataset, a large-scale benchmark for trimmed action classification.

DETAILS

TASK B

Spatio-temporal Action Localization (AVA)

This task is intended to evaluate the ability of algorithms to **localize human actions in space and time**. Each labeled video segment can contain multiple subjects, each performing potentially multiple actions. The goal is to identify these subjects and actions over continuous 15-minute video clips extracted from movies. For this task, participants will use the new AVA atomic visual actions dataset.

DETAILS

TASK C

Trimmed Event Recognition (Moments in Time)

This task is intended to evaluate the ability of algorithms to **classify events in trimmed 3-second videos**. Here, videos contain a single activity, and all clips have a standard duration of 3 seconds. There will be two tracks. The first track will use the Moments in Time dataset, a new large-scale dataset for video understanding, which has 800K videos in the training set. For the second track, participants will use the Moments in Time Mini dataset, a subset of Moments in Time with 100k videos provided in the training set.

DETAILS

Key
Detection
Ground-truth
Time

(Background is sped up)

Sponsors: DeepMind Google AI facebook

ActivityNet: A Large-Scale Video Benchmark for Human Activity Understanding [CVPR'15] [Google Faculty Research Award (only one from SA)]

Natural Genomics Circuits - are complicated - we explore foundational research on Machine Intelligence and Causal Reasoning to facilitate <u>Scientific Discoveries</u> from our Single-Cell Genomics Experiments

Home > News > Developing an AI system that thinks like a scientist

Developing an AI system that thinks like a scientist

Feb 03, 2019 Research

EMAIL FACEBOOK LINKEDIN TWITTER

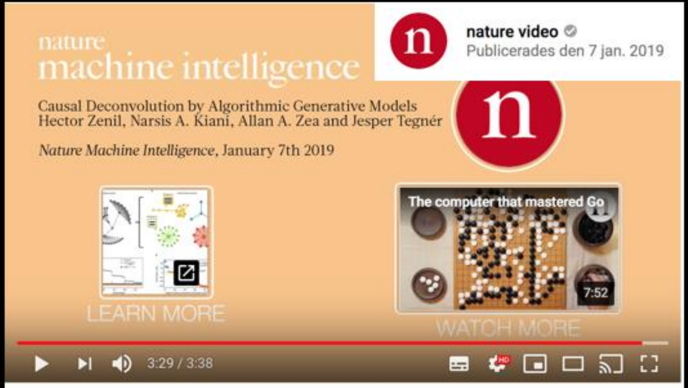

Remodelling machine learning: An AI that thinks like a scientist

- ***<u>Technically</u>: Approximating uncomputable Mathematics, in a Church-Turing sense, to create an <u>ORACLE</u>***

- https://github.com/allgebrist/Causal-Deconvolution-of-Networks/
- https://www.youtube.com/watch?v=rkmz7DAA-t8&feature=youtu.be

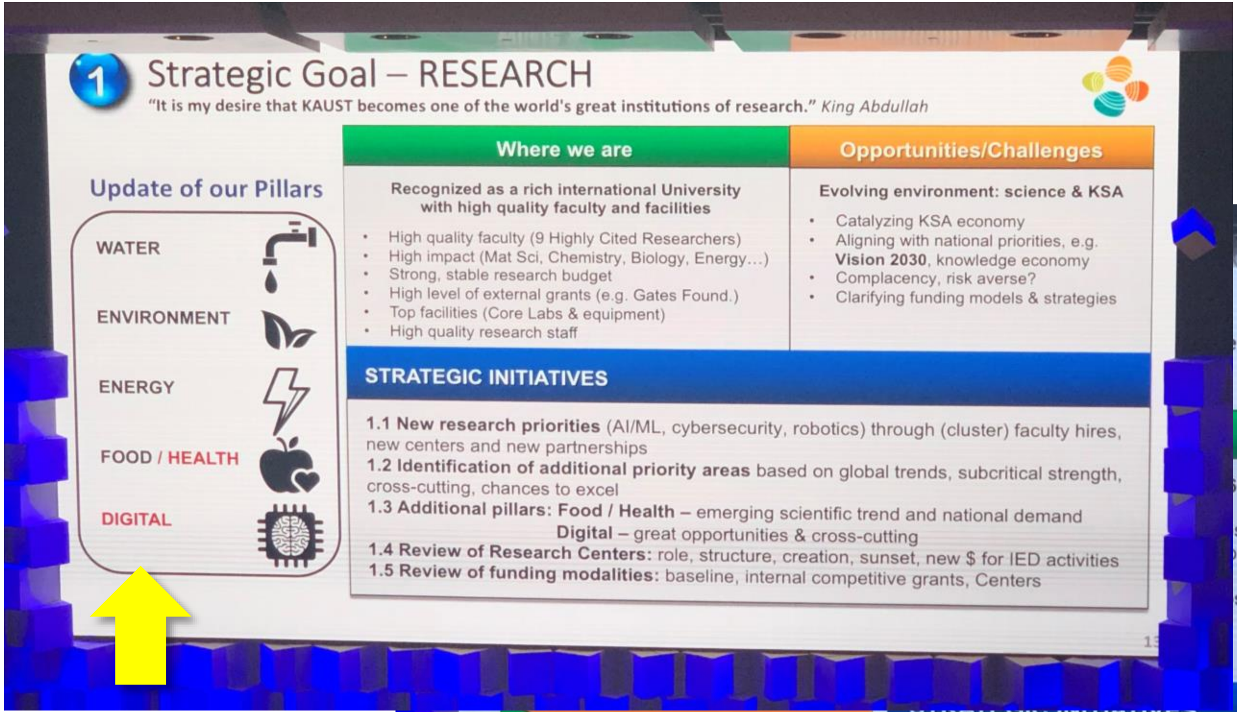
1
Strategic Goal – RESEARCH
"It is my desire that KAUST becomes one of the world's great institutions of research." King Abdullah
Update of our Pillars
WATER
ENVIRONMENT
ENERGY
FOOD / HEALTH
DIGITAL
Where we are
Recognized as a rich international University with high quality faculty and facilities
High quality faculty (9 Highly Cited Researchers)
High impact (Mat Sci, Chemistry, Biology, Energy…)
Strong, stable research budget
High level of external grants (e.g. Gates Found.)
Top facilities (Core Labs & equipment)
High quality research staff
Opportunities/Challenges
Evolving environment: science & KSA
Catalyzing KSA economy
Aligning with national priorities, e.g. Vision 2030, knowledge economy
Complacency, risk averse?
Clarifying funding models & strategies
STRATEGIC INITIATIVES
1.1 New research priorities (AI/ML, cybersecurity, robotics) through (cluster) faculty hires, new centers and new partnerships
1.2 Identification of additional priority areas based on global trends, subcritical strength, cross-cutting, chances to excel
1.3 Additional pillars: Food / Health – emerging scientific trend and national demand
Digital – great opportunities & cross-cutting
1.4 Review of Research Centers: role, structure, creation, sunset, new $ for IED activities
1.5 Review of funding modalities: baseline, internal competitive grants, Centers

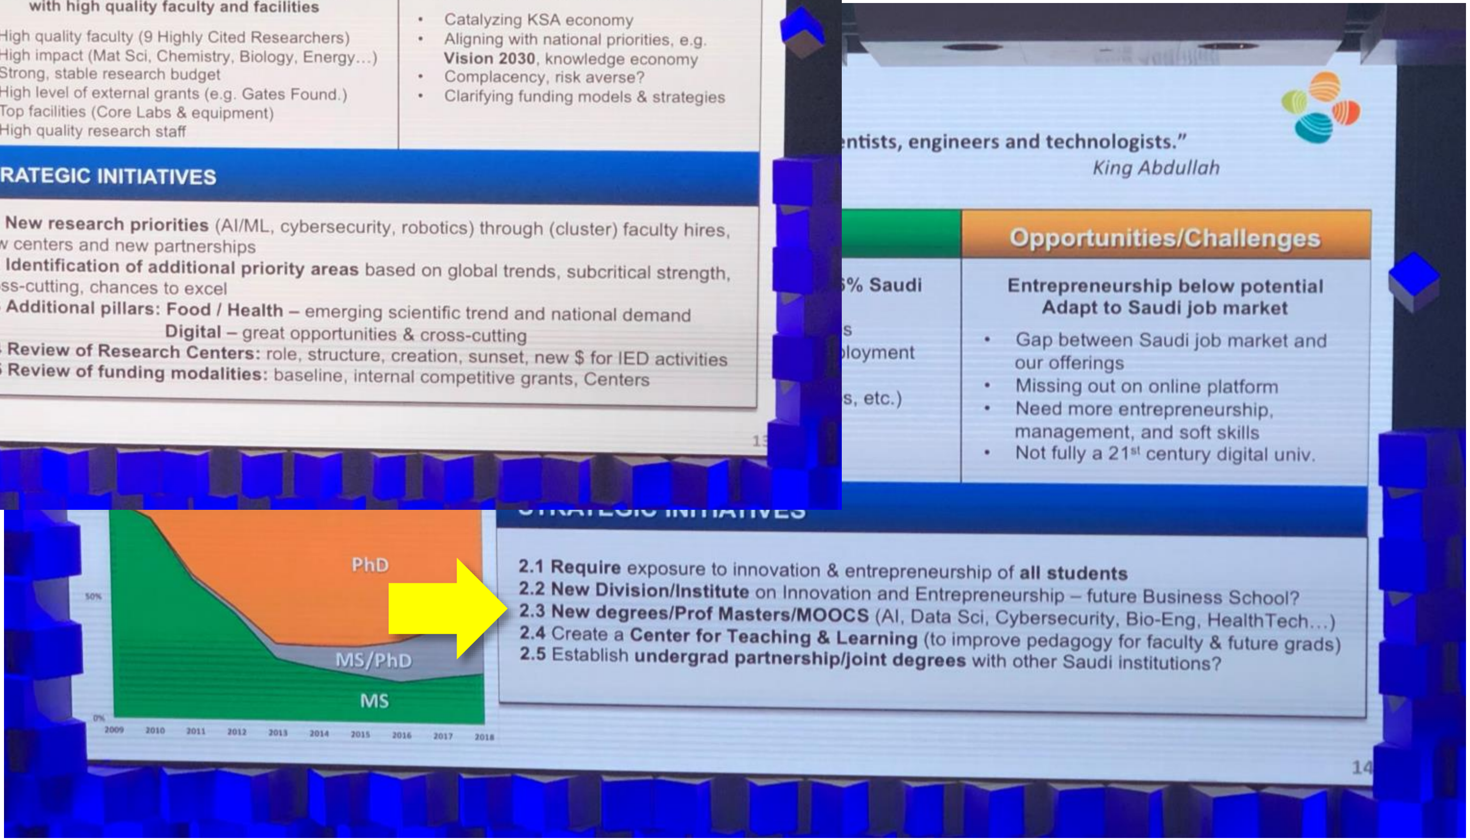
entists, engineers and technologists."
King Abdullah
% Saudi
s
ployment
s, etc.)
Opportunities/Challenges
Entrepreneurship below potential
Adapt to Saudi job market
Gap between Saudi job market and our offerings
Missing out on online platform
Need more entrepreneurship, management, and soft skills
Not fully a 21st century digital univ.
PhD
MS/PhD
MS
50%
0%
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2.1 Require exposure to innovation & entrepreneurship of all students
2.2 New Division/Institute on Innovation and Entrepreneurship – future Business School?
2.3 New degrees/Prof Masters/MOOCS (AI, Data Sci, Cybersecurity, Bio-Eng, HealthTech…)
2.4 Create a Center for Teaching & Learning (to improve pedagogy for faculty & future grads)
2.5 Establish undergrad partnership/joint degrees with other Saudi institutions?
14

The ML Hub Offerings

One-stop-shop for AI/ML on campus

- regular seminars and activities
  - invited speakers (academic & industry)
  - new research and trends in the areas
  - participation events (e.g. hackathon)

The ML Hub Offerings

One-stop-shop for AI/ML on campus

- network of AI/ML enthusiasts and experts
  - mentorship and consulting
  - brainstorming and collaboration
  - vehicle for disseminating information (e.g. news, research, job opportunities, etc.)

The ML Hub Offerings

One-stop-shop for AI/ML on campus

- hands-on tutorials and workshops
  - practical starters
  - various levels of AI/ML competence and use
  - hardware and software awareness

The ML Hub Offerings

One-stop-shop for AI/ML on campus

- the AI/ML KAUST window to the world
  - aggregation of our work in the field
  - showcase our growing position in these areas

What's coming up next?

- **Sci-Café: AI**
  - February 20th
- **Hub Seminar (Prof. Zheng Qu, HKU)**
  - March 20th

Webpage: ml.kaust.edu.sa **Twitter:** @kaustMLhub **Slack:** the-ml-hub.slack.com