Artificial Intelligence in Finance

at

Hong Kong University of Science and Technology

Course Infomation

- Course web:
  - http://aifin-hkust.github.io/2020/
- Time:
  - Every Tuesday, 3:00-5:50pm
- Venue:
  - Zoom Meetings: https://hkust.zoom.us/j/933177593
  - Until new notification
- Instructors:
  - Anthony Woo (Alpha Intelligence Capital, HK)
  - Yuan Yao (HKUST)

MAFS 6010U: Artificial Intelligence in Finance

Module Description	Course Focus
This course explores the basic concepts and underlying principles of artificial intelligence (AI), delving into the **fundamentals of machine learning** with insights from **case studies** of relevant technologies. Allowing for the **experimentation of applications** of machine learning, this course is designed to encourage students to devise creative ways to **put readily-available AI technologies to use to tackle problems in real life**.	The module aims to provide students with an understanding of artificial intelligence through: • Examining the history as well as **key concepts and theories** of AI and the enablers of the technology • Reviewing various types of neural networks, and analyzing the relevant use cases of AI **across industry verticals**, including robotic process automation, finance, cybersecurity, computer vision, and autonomous driving

Artificial Intelligence, Machine Learning, and Deep Learning

- AI is born in 1950s, when a handful of pioneers from the nascent field of computer science started asking **whether computers could be made to "think"**—a question whose ramifications we're still exploring today.

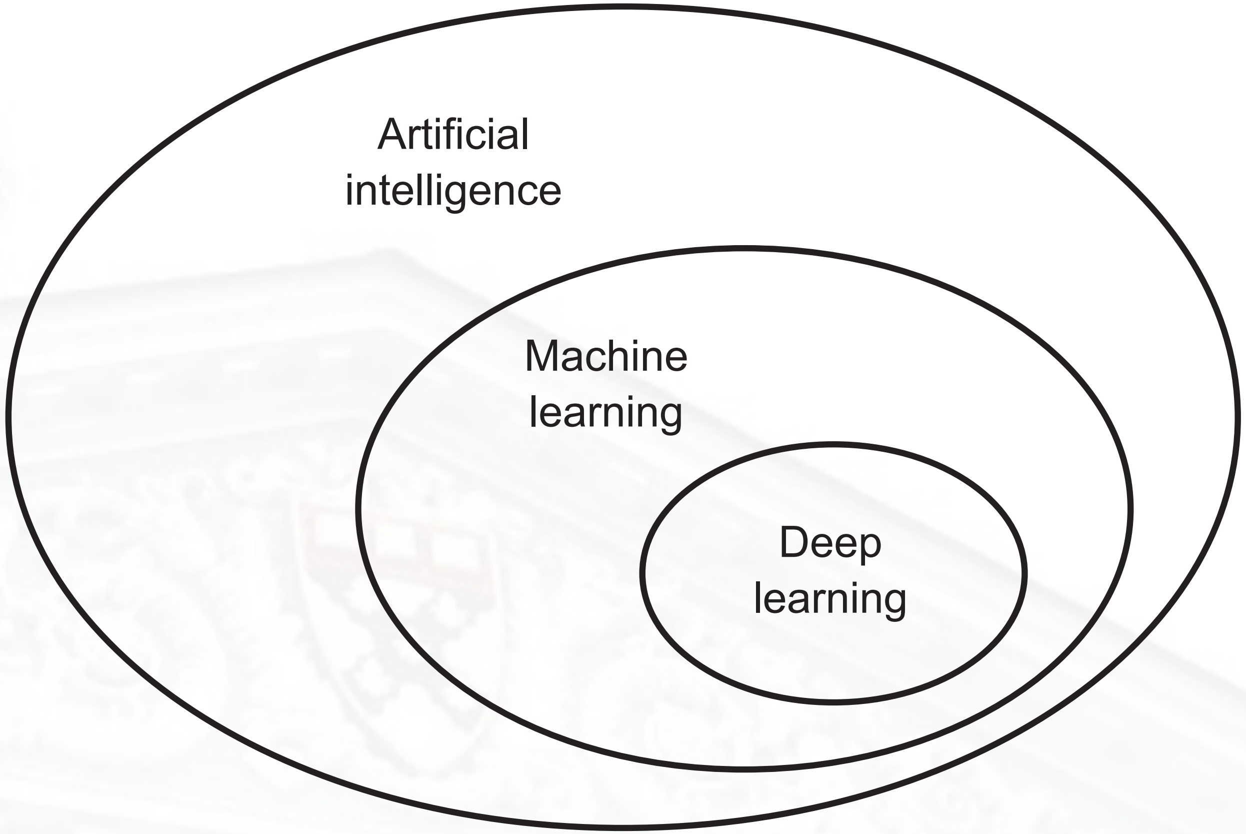

Machine Learning is a new paradigm of computer programming

- During 1950s-1980s, two competitive ideas of realizing AI exist
  - Rule based inference, or called Expert System
  - Statistics based inference, or called Machine Learning
- 1990s- Machine Learning becomes dominant

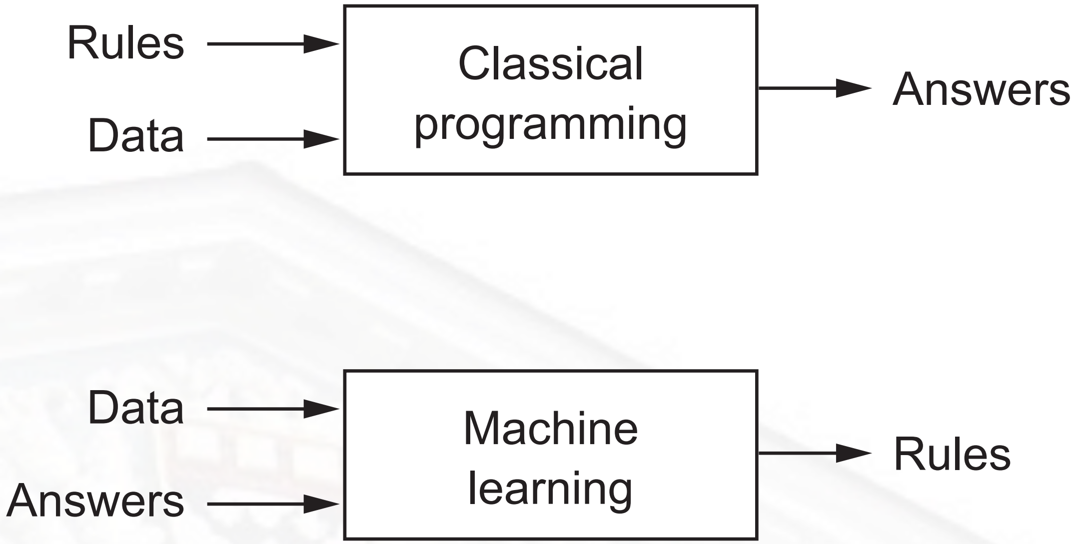

The 1st machine learning method: Least Squares

- Invention:
  - **Carl Friederich Gauss** (~1795/1809/1810),
  - Adrien-Marie Legendre (1805)
  - Robert Adrain (1808)
- Application:
  - Prediction of the location of asteroid Ceres after it emerged from behind the sun (Franz Xaver von Zach 1801)
  - Orbits of planets, Newton Laws
  - Statistics,
  - …

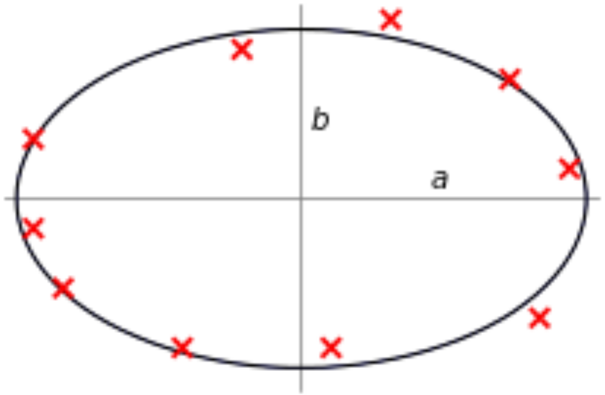

A Brief History of Neural Networks

(Deep Learning)

Perceptron: single-layer

- Invented by Frank Rosenblatt (1957)

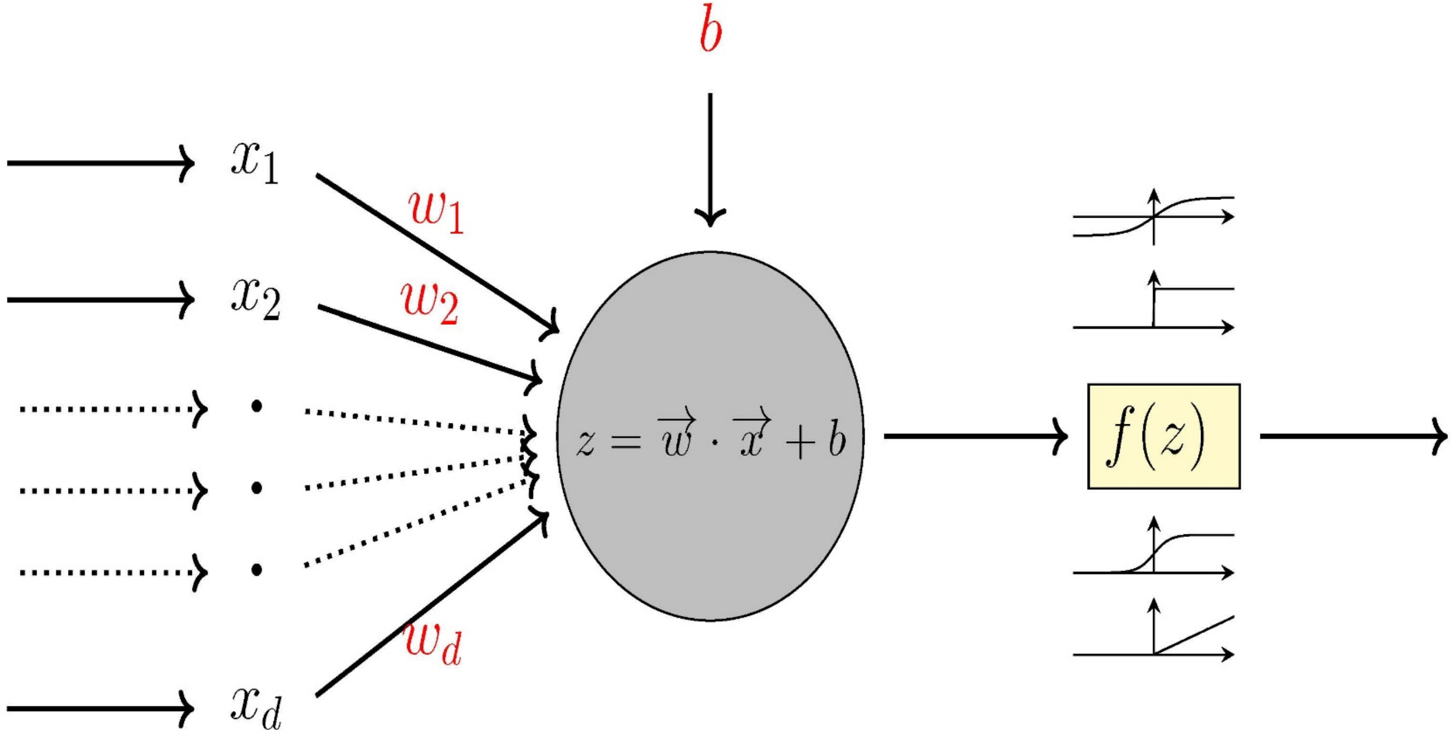

The Perceptron Algorithm for classification

$$\ell(w) = -\sum_{i \in \mathcal{M}_w} y_i \langle w, \mathbf{x}_i \rangle , \quad \mathcal{M}_w = \{i : y_i \langle \mathbf{x}_i, w \rangle < 0, y_i \in \{-1, 1\}\}.$$

The Perceptron Algorithm is a *Stochastic Gradient Descent* method (**Robbins-Monro 1951**):

$$\begin{aligned} w_{t+1} &= w_t - \eta_t \nabla_i \ell(w) \\ &= \begin{cases} w_t - \eta_t y_i \mathbf{x}_i, & \text{if } y_i w_t^T \mathbf{x}_i < 0, \\ w_t, & \text{otherwise.} \end{cases} \end{aligned}$$

Finiteness of Stopping Time and Margin

The perceptron convergence theorem was proved by Block (1962) and Novikoff (1962). The following version is based on that in Cristianini and Shawe-Taylor (2000).

Theorem 1 (Block, Novikoff). *Let the training set* $S = \{(\mathbf{x}_1, t_1), \ldots, (\mathbf{x}_n, t_n)\}$ *be contained in a sphere of radius* R *about the origin. Assume the dataset to be linearly separable, and let* $\mathbf{w}_{\text{opt}}$, $\|\mathbf{w}_{\text{opt}}\| = 1$, *define the hyperplane separating the samples, having functional margin* $\gamma > 0$. *We initialise the normal vector as* $\mathbf{w}_0 = \mathbf{0}$. *The number of updates, k, of the perceptron algorithms is then bounded by*

$$k \leq \left(\frac{2R}{\gamma}\right)^2 . \quad (10)$$

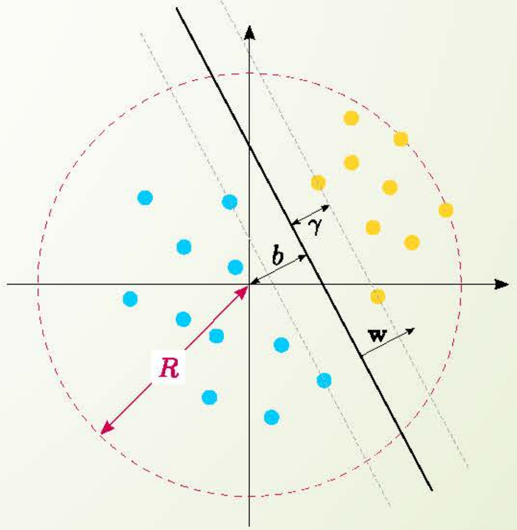

Input ball: $R = \max_i \|\mathbf{x}_i\|.$

Margin: $\gamma := \min_i y_i f(x_i)$

Hilbert's 13th Problem

Algebraic equations (under a suitable transformation) of degree up to 6 can be solved by functions of two variables. What about

$$x^7 + ax^3 + bx^2 + cx + 1 = 0?$$

Hilbert's conjecture: $x(a, b, c)$ cannot be expressed by a superposition (sums and compositions) of bivariate functions.

Question: can every continuous (analytic, C^∞, etc) function of n variables be represented as a superposition of continuous (analytic, C^∞, etc) functions of $n - 1$ variables?

Theorem (D. Hilbert)

There is an analytic function of three variables that cannot be expressed as a superposition of bivariate ones.

Kolmogorov's Superposition Theorem

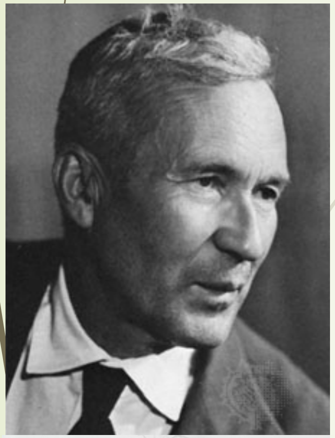

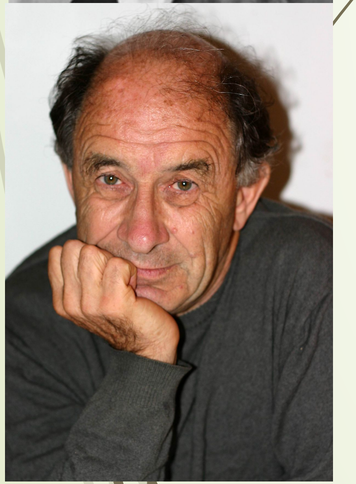

Theorem (A. Kolmogorov, 1956; V. Arnold, 1957)

Given $n \in \mathbb{Z}^+$, every $f_0 \in C([0,1]^n)$ can be reprensented as

$$f_0(x_1, x_2, \cdots, x_n) = \sum_{q=1}^{2n+1} g_q \left(\sum_{p=1}^{n} \phi_{pq}(x_p) \right),$$

where $\phi_{pq} \in C[0,1]$ are increasing functions independent of f_0 and $g_q \in C[0,1]$ depend on f_0.

- Can choose g_q to be all the same $g_q \equiv g$ (Lorentz, 1966).
- Can choose ϕ_{pq} to be Hölder or Lipschitz continuous, but not C^1 (Fridman, 1967).
- Can choose $\phi_{pq} = \lambda_p \phi_q$ where $\lambda_1, \cdots, \lambda_n > 0$ and $\sum_p \lambda_p = 1$ (Sprecher, 1972).

If *f* is a multivariate continuous function, then *f* can be written as a superposition of composite functions of mixtures of continuous functions of single variables:
finite **composition** of continuous functions of a **single variable** and the **addition**.

Kolmogorov's Exact Representation is not stable or smooth

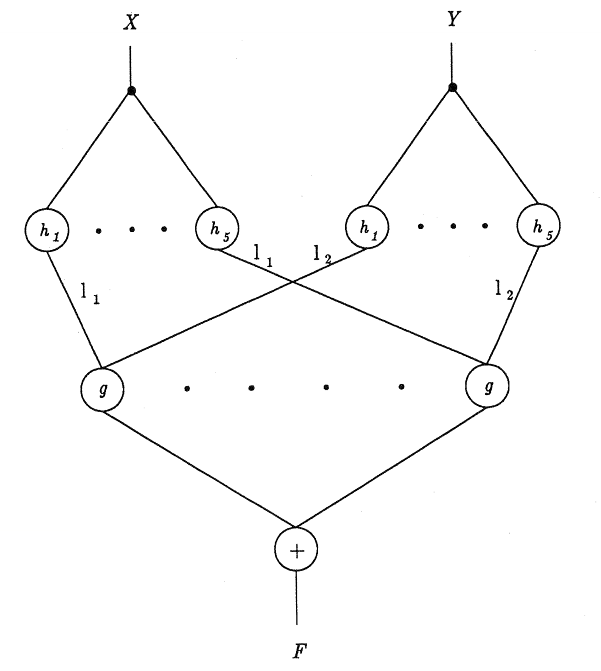

Figure 1: The network representation of an improved version of Kolmogorov's theorem, due to Kahane (1975). The figure shows the case of a bivariate function. The Kahane's representation formula is $f(x_1, \ldots, x_n) = \sum_{q=1}^{2n+1} g[\sum_{p=1}^{n} l_p h_q(x_p)]$ where h_q are strictly monotonic functions and l_p are strictly positive constants smaller than 1.

- [Girosi-Poggio'1989] Representation Properties of Networks: Kolmogorov's Theorem Is Irrelevant, https://www.mitpressjournals.org/doi/pdf/10.1162/neco.1989.1.4.465
- Lacking smoothness in *h* and *g* [Vitushkin'1964] fails to guarantee the **generalization ability (stability)** against noise and perturbations
- The representation is **not universal** in the sense that *g* and *h* both depend on the function F to be represented.

Universal Approximate Representation

[Cybenko'1989, Hornik et al. 1989, Poggio-Girosi'1989, ...]

For continuous $f : [0,1]^N \to \mathbb{R}$ and $\varepsilon > 0$ there exists

$$F(x) = \alpha^\top \sigma(Wx + \beta)$$
$$= \sum_i \alpha_i \sigma\left(\sum_j W_{i,j}\, x_j + \beta_i\right)$$

such that for all x in $[0,1]^N$ we have $|F(x) - f(x)| < \varepsilon$.

Complexity (regularity, smoothness) thereafter becomes the central pursuit in Approximation Theory.

Locality or Sparsity of Computation

Minsky and Papert, 1969

Perceptron can't do **XOR** classification

Perceptron needs infinite global information to compute **connectivity**

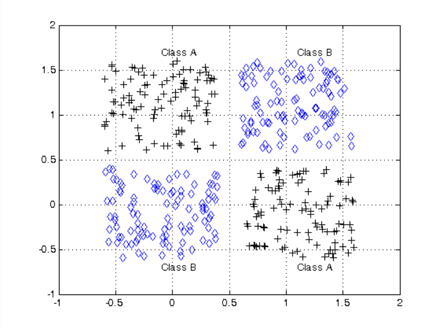

Locality or **Sparsity** is important:

Locality in time?

Locality in space?

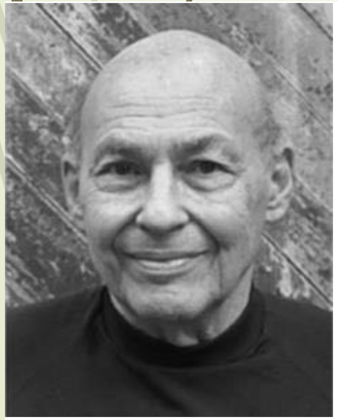

Marvin Minsky

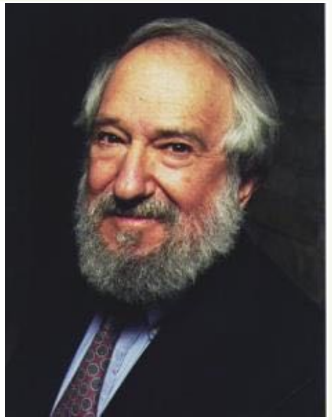

Seymour Papert

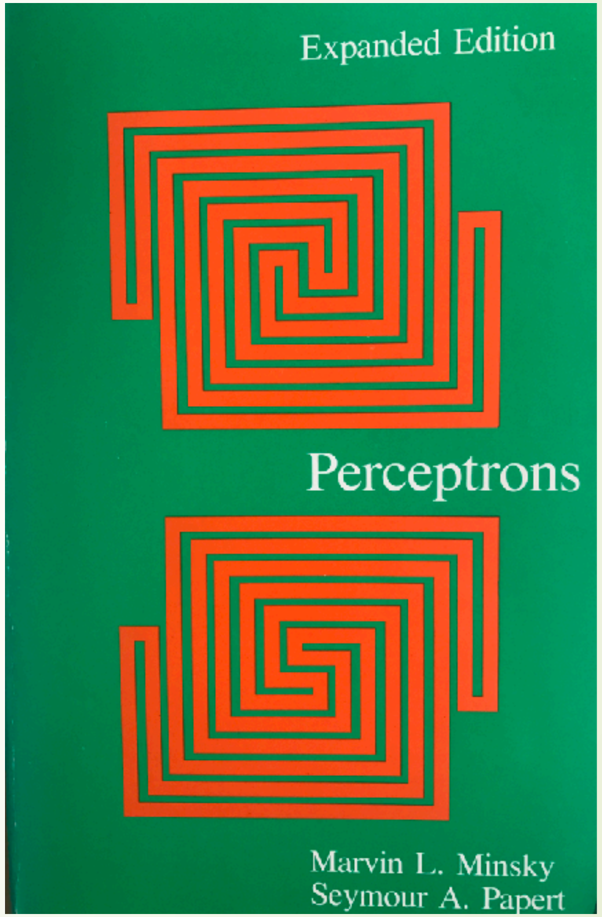

Multilayer Perceptrons (MLP) and Back-Propagation (BP) Algorithms

D.E. Rumelhart, G. Hinton, R.J. Williams (1986)
Learning representations by back-propagating errors, Nature, 323(9): 533-536

BP algorithms as **stochastic gradient descent** algorithms (**Robbins–Monro 1950; Kiefer-Wolfowitz 1951**) with Chain rules of Gradient maps

Deep network may classify **XOR**. Yet **topology**?

We address complexity and geometric invariant properties first.

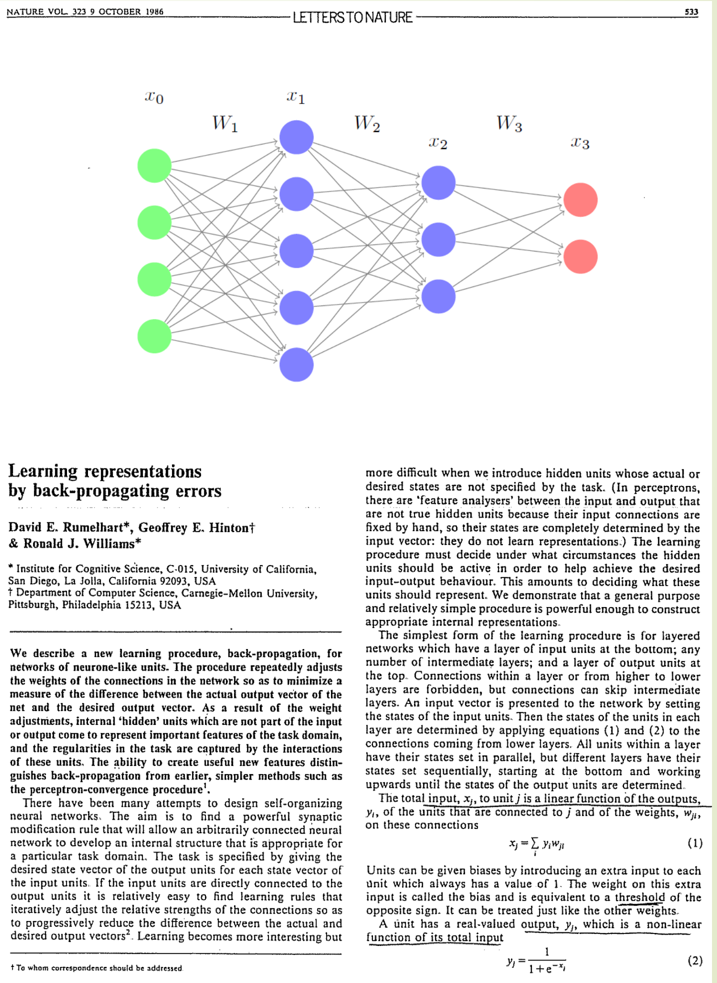

Learning representations by back-propagating errors

David E. Rumelhart*, Geoffrey E. Hinton† & Ronald J. Williams*

* Institute for Cognitive Science, C-015, University of California, San Diego, La Jolla, California 92093, USA
† Department of Computer Science, Carnegie-Mellon University, Pittsburgh, Philadelphia 15213, USA

We describe a new learning procedure, back-propagation, for networks of neurone-like units. The procedure repeatedly adjusts the weights of the connections in the network so as to minimize a measure of the difference between the actual output vector of the net and the desired output vector. As a result of the weight adjustments, internal 'hidden' units which are not part of the input or output come to represent important features of the task domain, and the regularities in the task are captured by the interactions of these units. The ability to create useful new features distinguishes back-propagation from earlier, simpler methods such as the perceptron-convergence procedure[1].

There have been many attempts to design self-organizing neural networks. The aim is to find a powerful synaptic modification rule that will allow an arbitrarily connected neural network to develop an internal structure that is appropriate for a particular task domain. The task is specified by giving the desired state vector of the output units for each state vector of the input units. If the input units are directly connected to the output units it is relatively easy to find learning rules that iteratively adjust the relative strengths of the connections so as to progressively reduce the difference between the actual and desired output vectors[2]. Learning becomes more interesting but more difficult when we introduce hidden units whose actual or desired states are not specified by the task. (In perceptrons, there are 'feature analysers' between the input and output that are not true hidden units because their input connections are fixed by hand, so their states are completely determined by the input vector: they do not learn representations.) The learning procedure must decide under what circumstances the hidden units should be active in order to help achieve the desired input–output behaviour. This amounts to deciding what these units should represent. We demonstrate that a general purpose and relatively simple procedure is powerful enough to construct appropriate internal representations.

The simplest form of the learning procedure is for layered networks which have a layer of input units at the bottom; any number of intermediate layers; and a layer of output units at the top. Connections within a layer or from higher to lower layers are forbidden, but connections can skip intermediate layers. An input vector is presented to the network by setting the states of the input units. Then the states of the units in each layer are determined by applying equations (1) and (2) to the connections coming from lower layers. All units within a layer have their states set in parallel, but different layers have their states set sequentially, starting at the bottom and working upwards until the states of the output units are determined.

The total input, x_j, to unit j is a linear function of the outputs, y_i, of the units that are connected to j and of the weights, w_{ji}, on these connections

$$x_j = \sum_i y_i w_{ji} \qquad (1)$$

Units can be given biases by introducing an extra input to each unit which always has a value of 1. The weight on this extra input is called the bias and is equivalent to a threshold of the opposite sign. It can be treated just like the other weights.

A unit has a real-valued output, y_j, which is a non-linear function of its total input

$$y_j = \frac{1}{1+e^{-x_j}} \qquad (2)$$

† To whom correspondence should be addressed.

Parallel Distributed Processing

by Rumelhart and McClelland, 1986

Minsky and Papert set out to show which functions can and cannot be computed by this class of machines. They demonstrated, in particular, that such perceptrons are unable to calculate such mathematical functions as parity (whether an odd or even number of points are on in the retina) or the topological function of connectedness (whether all points that are on are connected to all other points that are on either directly or via other points that are also on) without making use of absurdly large numbers of predicates. The analysis is extremely elegant and demonstrates the importance of a mathematical approach to analyz-

of multilayer networks that compute parity). Similarly, it is not difficult to develop networks capable of solving the connectedness or inside/outside problem. Hinton and Sejnowski have analyzed a version of such a network (see Chapter 7).

Essentially, then, although Minsky and Papert were exactly correct in their analysis of the *one-layer perceptron*, the theorems don't apply to systems which are even a little more complex. In particular, it doesn't apply to multilayer systems nor to systems that allow feedback loops.

BP algorithm = Gradient Descent Method

- Training examples $\{x_0^i\}_{i=1}^n$ and labels $\{y^i\}_{i=1}^n$
- Output of the network $\{x_L^i\}_{i=1}^m$
- Objective   Square loss, cross-entropy loss, etc.

$$J(\{W_l\},\{b_l\}) = \frac{1}{n}\sum_{i=1}^{n}\frac{1}{2}\|y^i - x_L^i\|_2^2 \tag{1}$$

- Gradient descent

$$W_l = W_l - \eta\frac{\partial J}{\partial W_l}$$

$$b_l = b_l - \eta\frac{\partial J}{\partial b_l}$$

In practice: use Stochastic Gradient Descent (SGD)

Derivation of BP: Lagrangian Multiplier

LeCun et al. 1988

Given n training examples $(I_i, y_i) \equiv$ (input,target) and L layers

- Constrained optimization

$$\min_{W,x} \quad \sum_{i=1}^{n} \|x_i(L) - y_i\|_2$$

$$\text{subject to} \quad x_i(\ell) = f_\ell\Big[W_\ell x_i\left(\ell-1\right)\Big],$$

$$i = 1, \ldots, n, \quad \ell = 1, \ldots, L, \; x_i(0) = I_i$$

- Lagrangian formulation (Unconstrained)

$$\min_{W,x,B} \mathcal{L}(W, x, B)$$

$$\mathcal{L}(W, x, B) = \sum_{i=1}^{n} \Bigg\{ \|x_i(L) - y_i\|_2^2 + \sum_{\ell=1}^{L} B_i(\ell)^T \Big(x_i(\ell) - f_\ell\Big[W_\ell x_i\left(\ell-1\right)\Big]\Big)\Bigg\}$$

http://yann.lecun.com/exdb/publis/pdf/lecun-88.pdf

BP Algorithm: Forward Pass

- Cascade of repeated [linear operation followed by coordinatewise nonlinearity]'s
- Nonlinearities: sigmoid, hyperbolic tangent, (recently) ReLU.

Algorithm 1 Forward pass

Input: x_0
Output: x_L

1: **for** $\ell = 1$ to L **do**
2: $\quad x_\ell = f_\ell(W_\ell x_{\ell-1} + b_\ell)$
3: **end for**

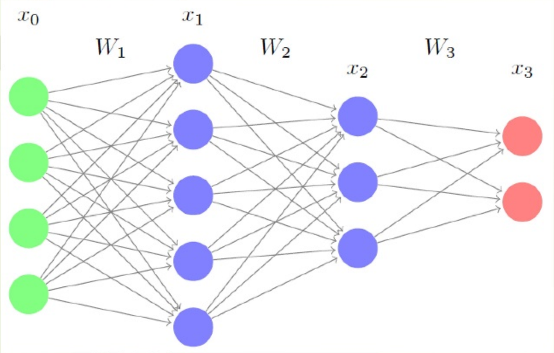

back-propagation – derivation

- $\frac{\partial \mathcal{L}}{\partial B}$

Forward pass

$$x_i(\ell) = f_\ell\Big[\underbrace{W_\ell x_i(\ell-1)}_{A_i(\ell)}\Big] \quad \ell = 1,\dots,L, \quad i = 1,\dots,n$$

- $\frac{\partial \mathcal{L}}{\partial x}, z_\ell = [\nabla f_\ell] B(\ell)$

Backward (adjoint) pass

$$z(L) = 2\nabla f_L\Big[A_i(L)\Big](y_i - x_i(L))$$

$$z_i(\ell) = \nabla f_\ell\Big[A_i(\ell)\Big] W_{\ell+1}^T z_i(\ell+1) \quad \ell = 0,\dots,L-1$$

- $W \leftarrow W + \lambda \frac{\partial \mathcal{L}}{\partial W}$

Weight update

$$W_\ell \leftarrow W_\ell + \lambda \textstyle\sum_{i=1}^n z_i(\ell) x_i^T(\ell-1)$$

Convolutional Neural Networks: shift invariances and locality

- Can be traced to *Neocognitron* of Kunihiko Fukushima (1979)
- Yann LeCun combined convolutional neural networks with back propagation (1989)
- Imposes **shift invariance** and **locality** on the weights
- Forward pass remains similar
- Backpropagation slightly changes – need to sum over the gradients from all spatial positions

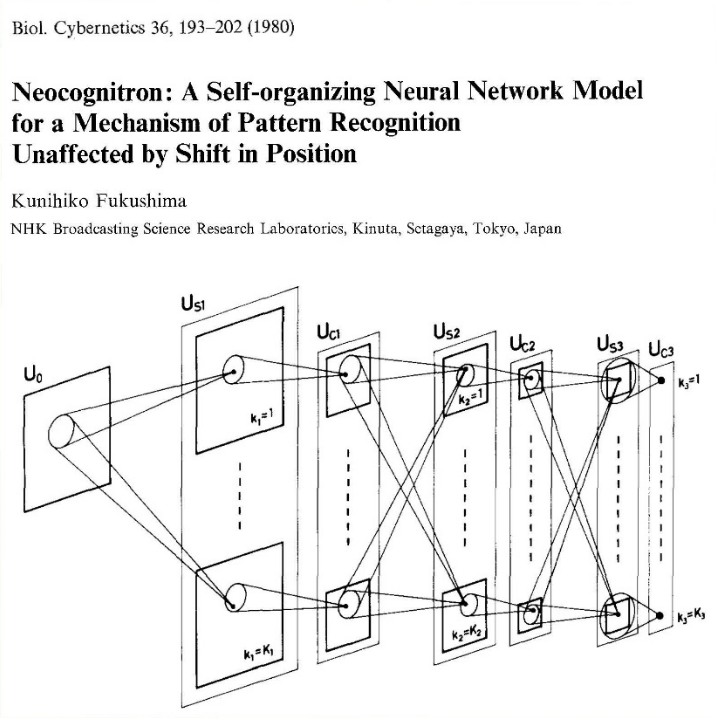
Biol. Cybernetics 36, 193–202 (1980)

Neocognitron: A Self-organizing Neural Network Model for a Mechanism of Pattern Recognition Unaffected by Shift in Position

Kunihiko Fukushima

NHK Broadcasting Science Research Laboratories, Kinuta, Setagaya, Tokyo, Japan

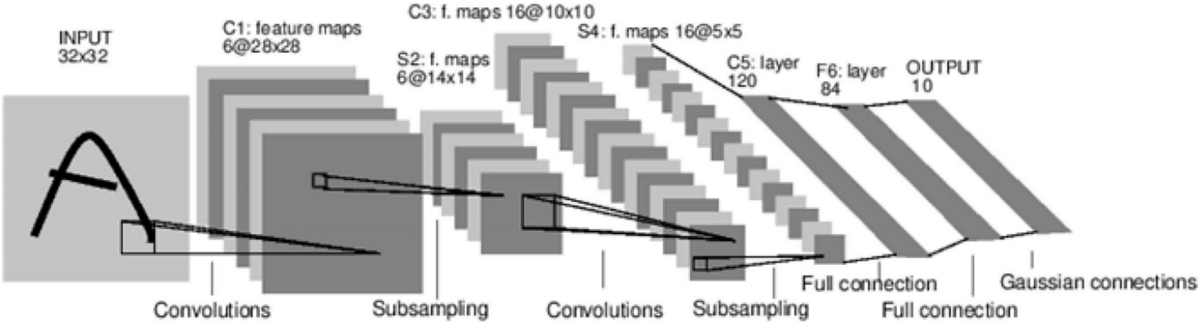

Max-Margin Classifier (SVM)

$$\text{minimize}_{\beta_0,\beta_1,\ldots,\beta_p} \|\beta\|^2 := \sum_j \beta_j^2$$

subject to $y_i(\beta_0 + \beta_1 x_{i1} + \ldots + \beta_p x_{ip}) \geq 1$ for all i

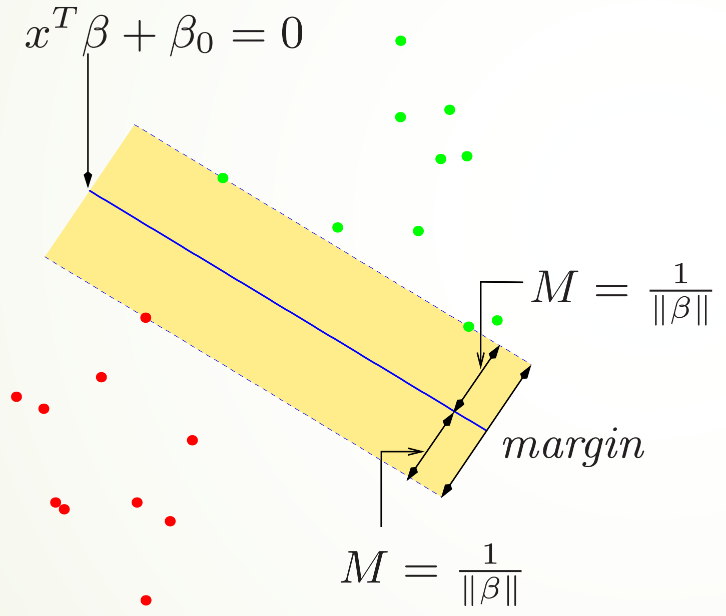

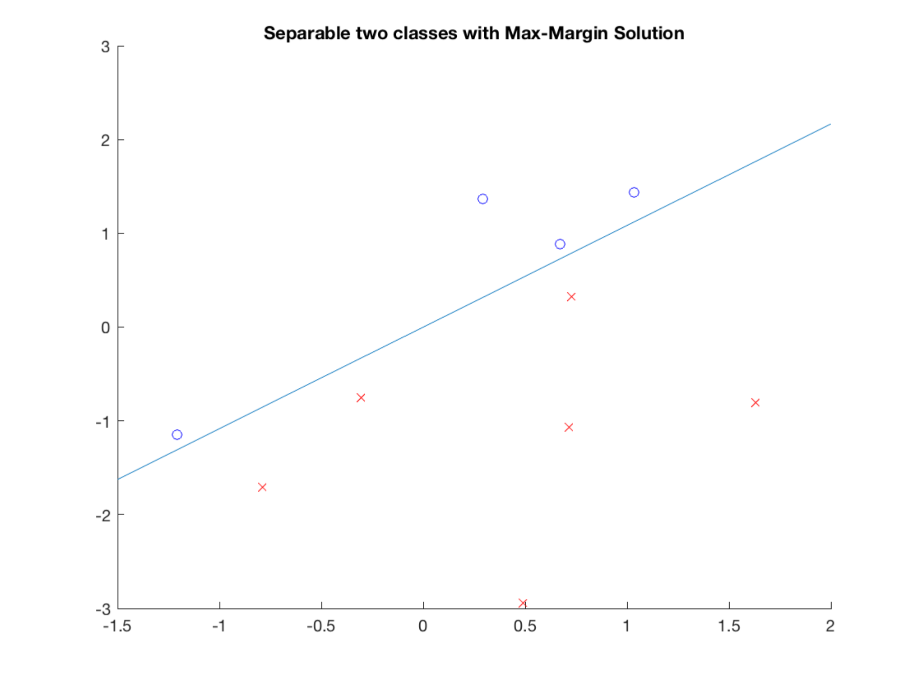

Vladmir Vapnik, 1994

MNIST Dataset Test Error
LeCun et al. 1998

Simple SVM performs as well as Multilayer Convolutional Neural Networks which need careful tuning (LeNets)

Dark era for NN: 1998-2012

2000-2010: The Era of SVM, Boosting, … as nights of Neural Networks

Decision Trees and Boosting

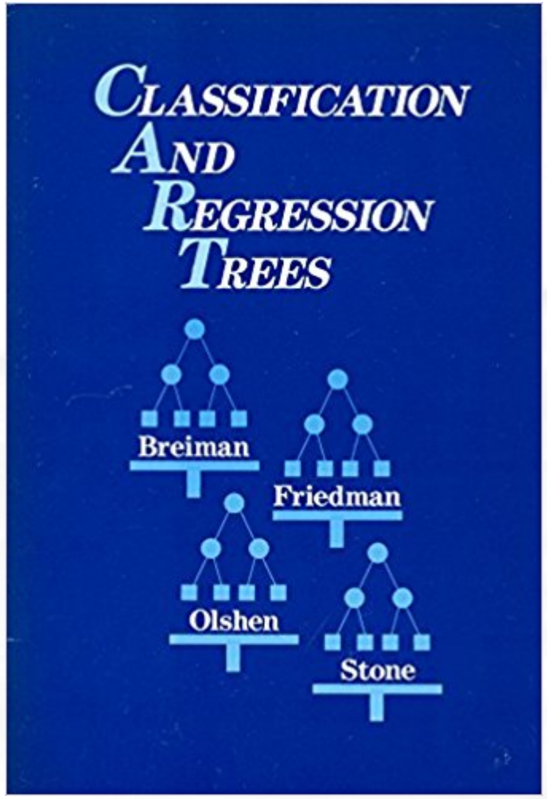

- Breiman, Friedman, Olshen, Stone, (1983): CART
- ``The Boosting problem’’ (M. Kearns & L. Valiant): **Can a set of weak learners create a single strong learner**? (三个臭皮匠顶个诸葛亮？)
- Breiman (1996): Bagging
- Freund, Schapire (1997): **AdaBoost** (“the best off-the-shelf algorithm” by Breiman)
- Breiman (2001): **Random Forests**

Around the year of 2012: return of NN as `deep learning'

Speech Recognition: TIMIT

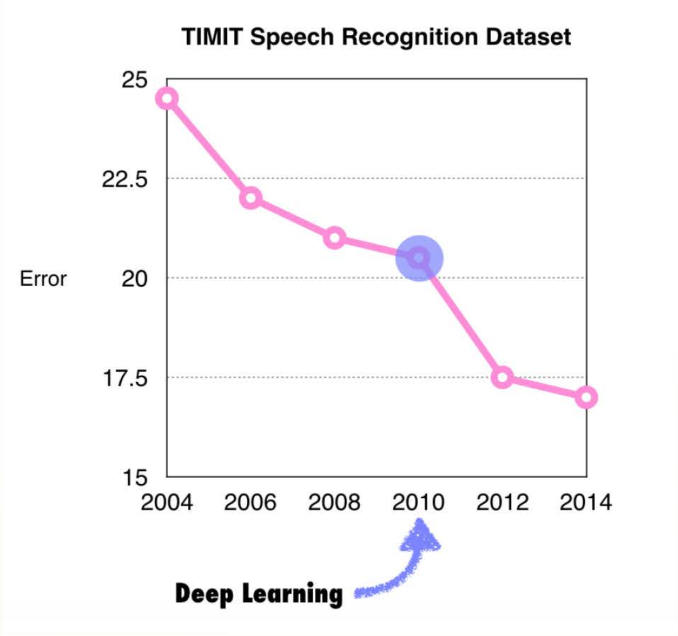

Computer Vision: ImageNet

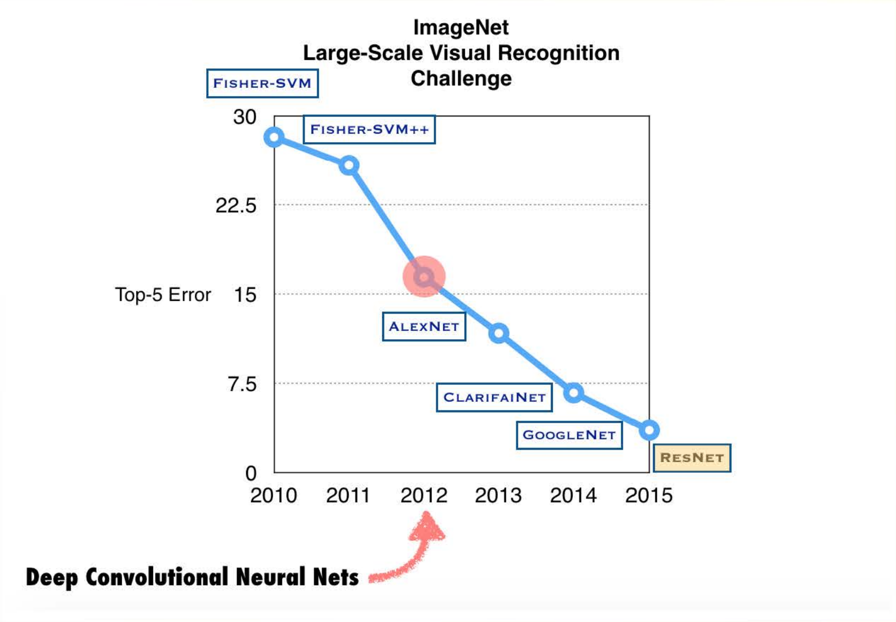

Depth as function of year

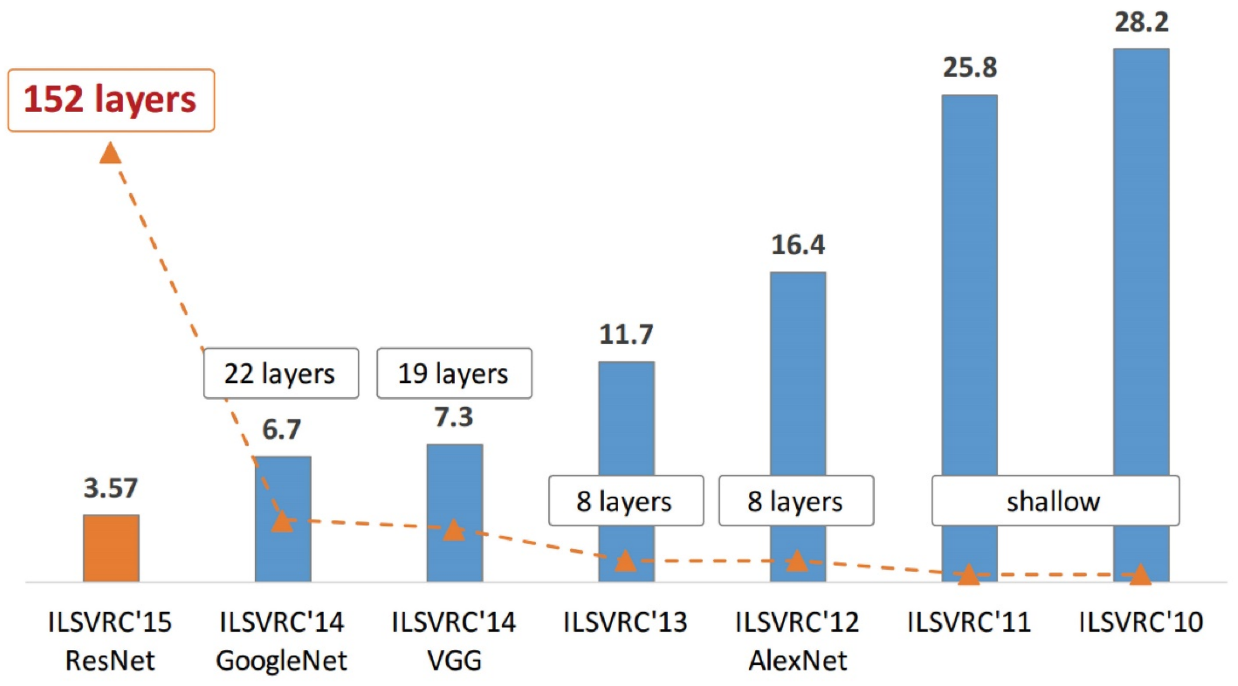

[He et al., 2016]

ILSVRC ImageNet Top 5 errors

- ImageNet (subset):
  - 1.2 million training images
  - 100,000 test images
  - 1000 classes
- ImageNet large-scale visual recognition Challenge

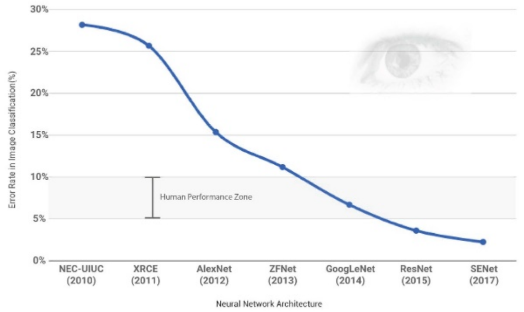

source: https://www.linkedin.com/pulse/must-read-path-breaking-papers-image-classification-muktabh-mayank

Reaching Human Performance Level in Games

Deep Blue in 1997

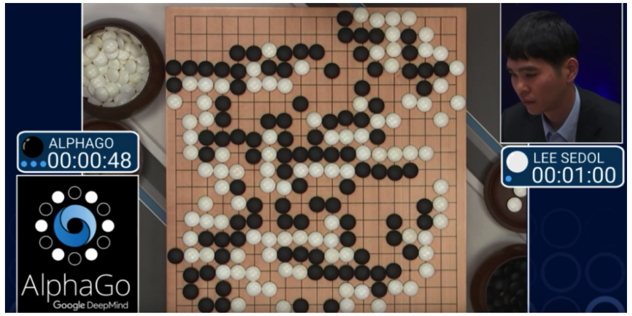

AlphaGo "LEE" 2016

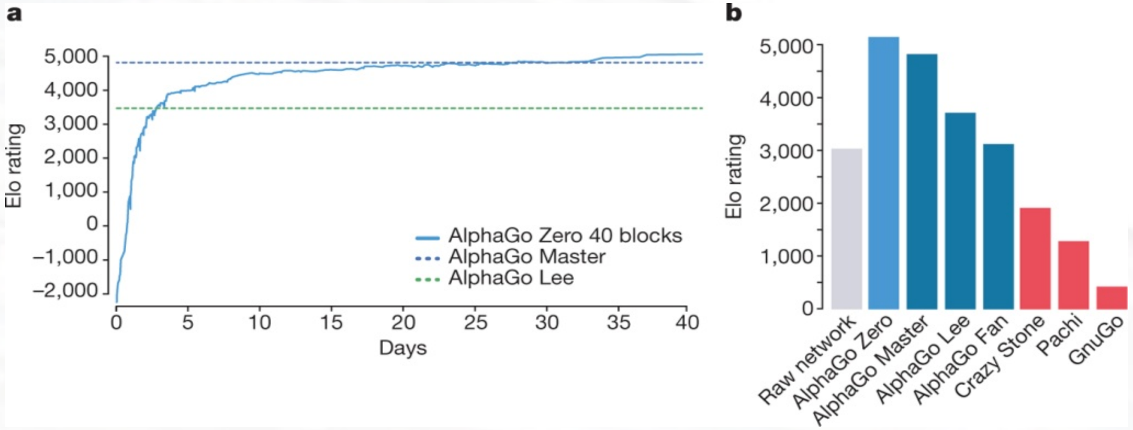

AlphaGo "ZERO" D Silver *et al.* *Nature* **550,** 354–359 (2017) doi:10.1038/nature24270

Number of AI papers on arXiv, 2010-2019

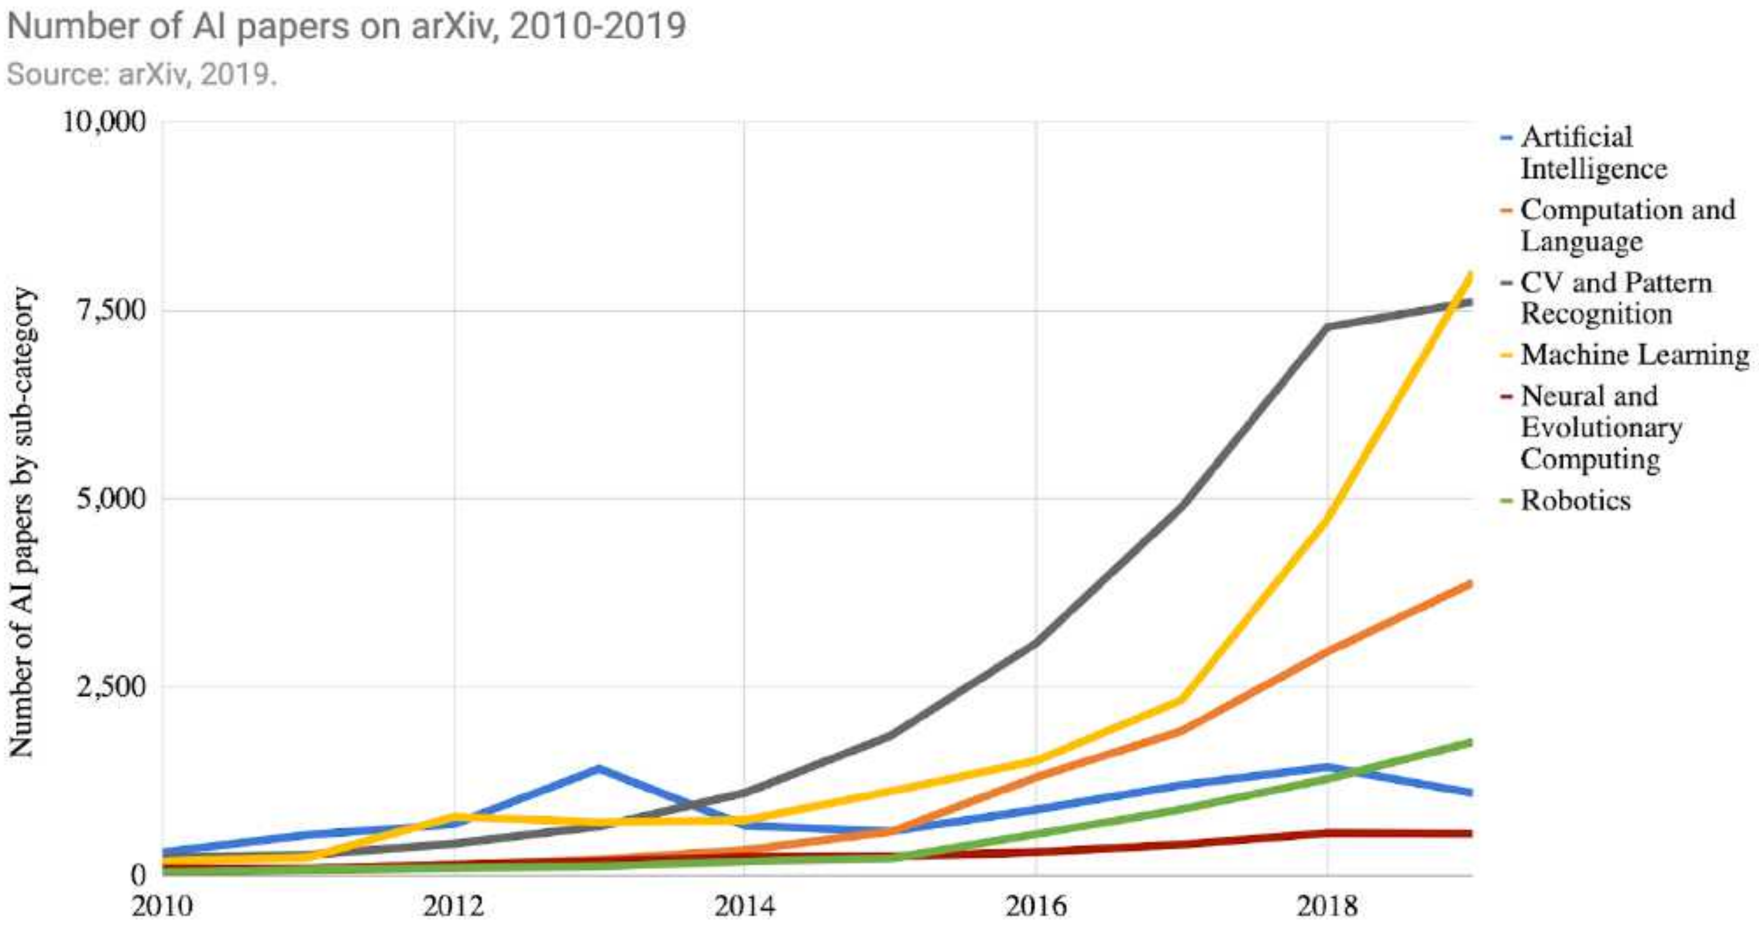

Fig. 1.6.

Growth of Deep Learning

'Deep Learning' is coined by Hinton et al. in their Restricted Boltzman Machine paper, *Science* 2006, not yet popular until championing ImageNet competitions.

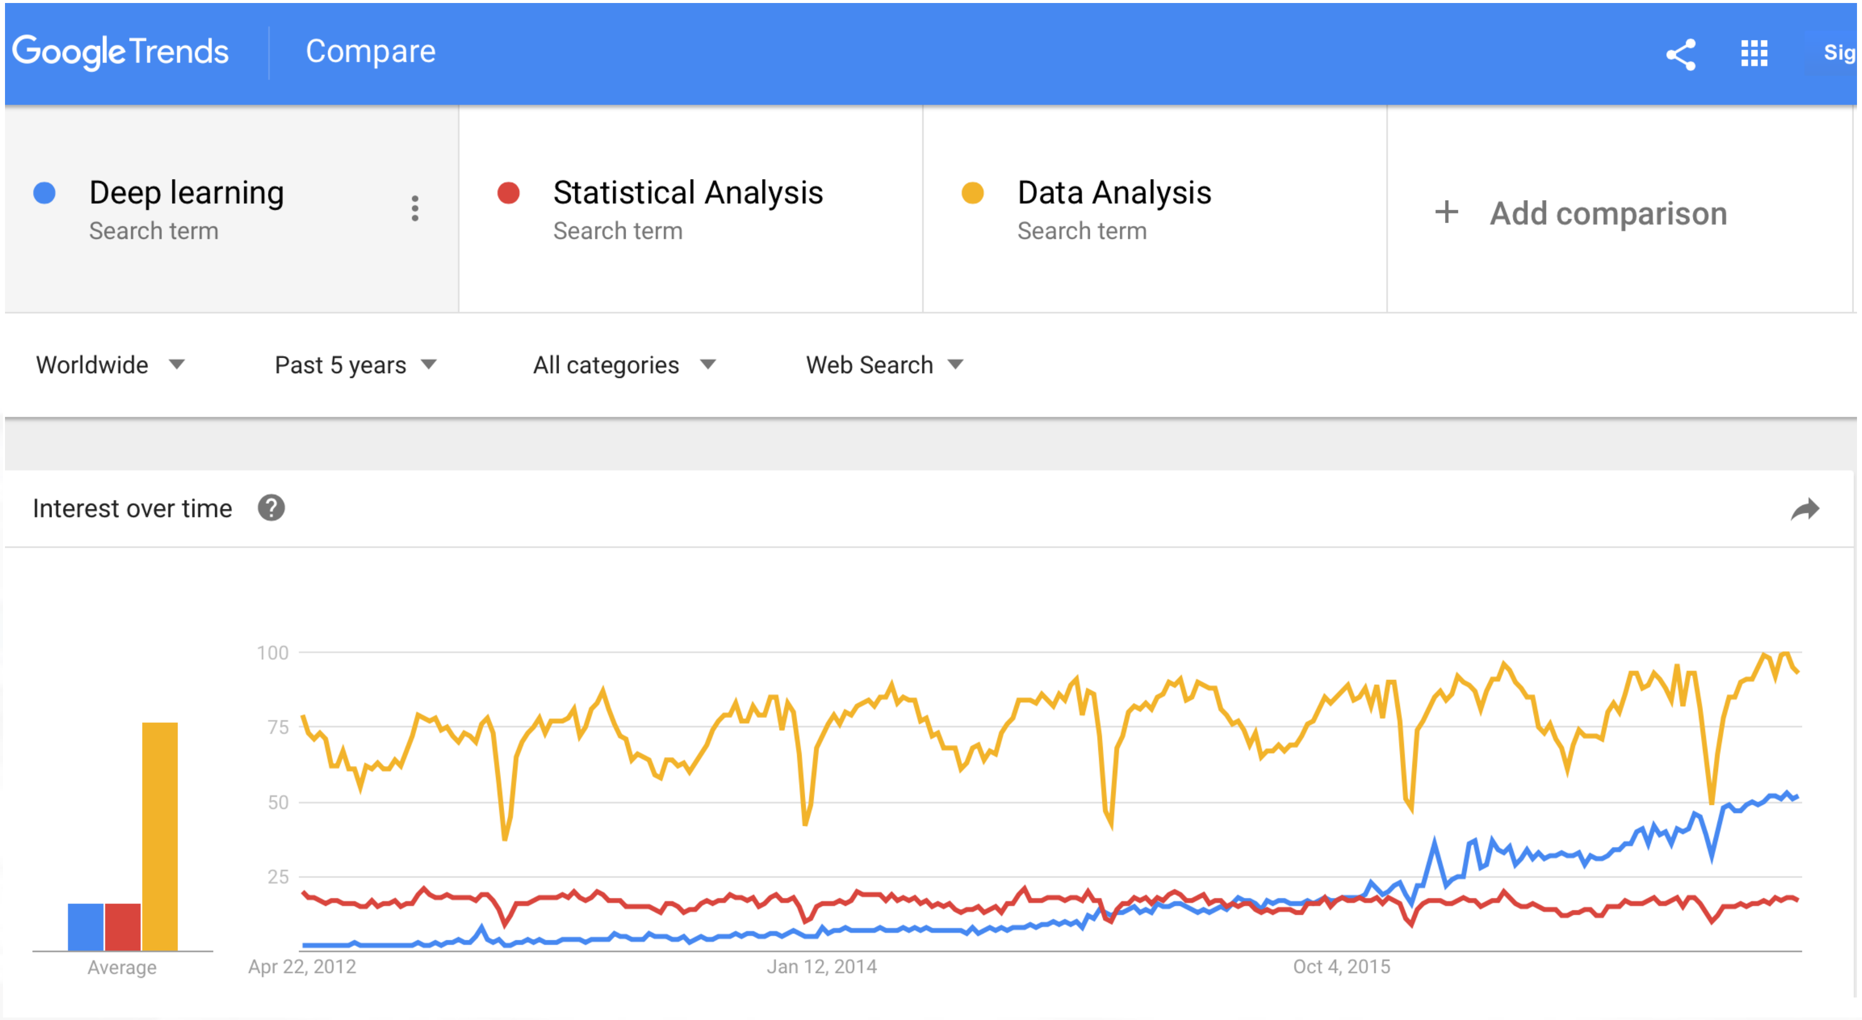

Some Cold Water: Tesla Autopilot Misclassifies Truck as Billboard

Problem: Why? How can you trust a blackbox?

Deep Learning may be fragile in generalization against noise!

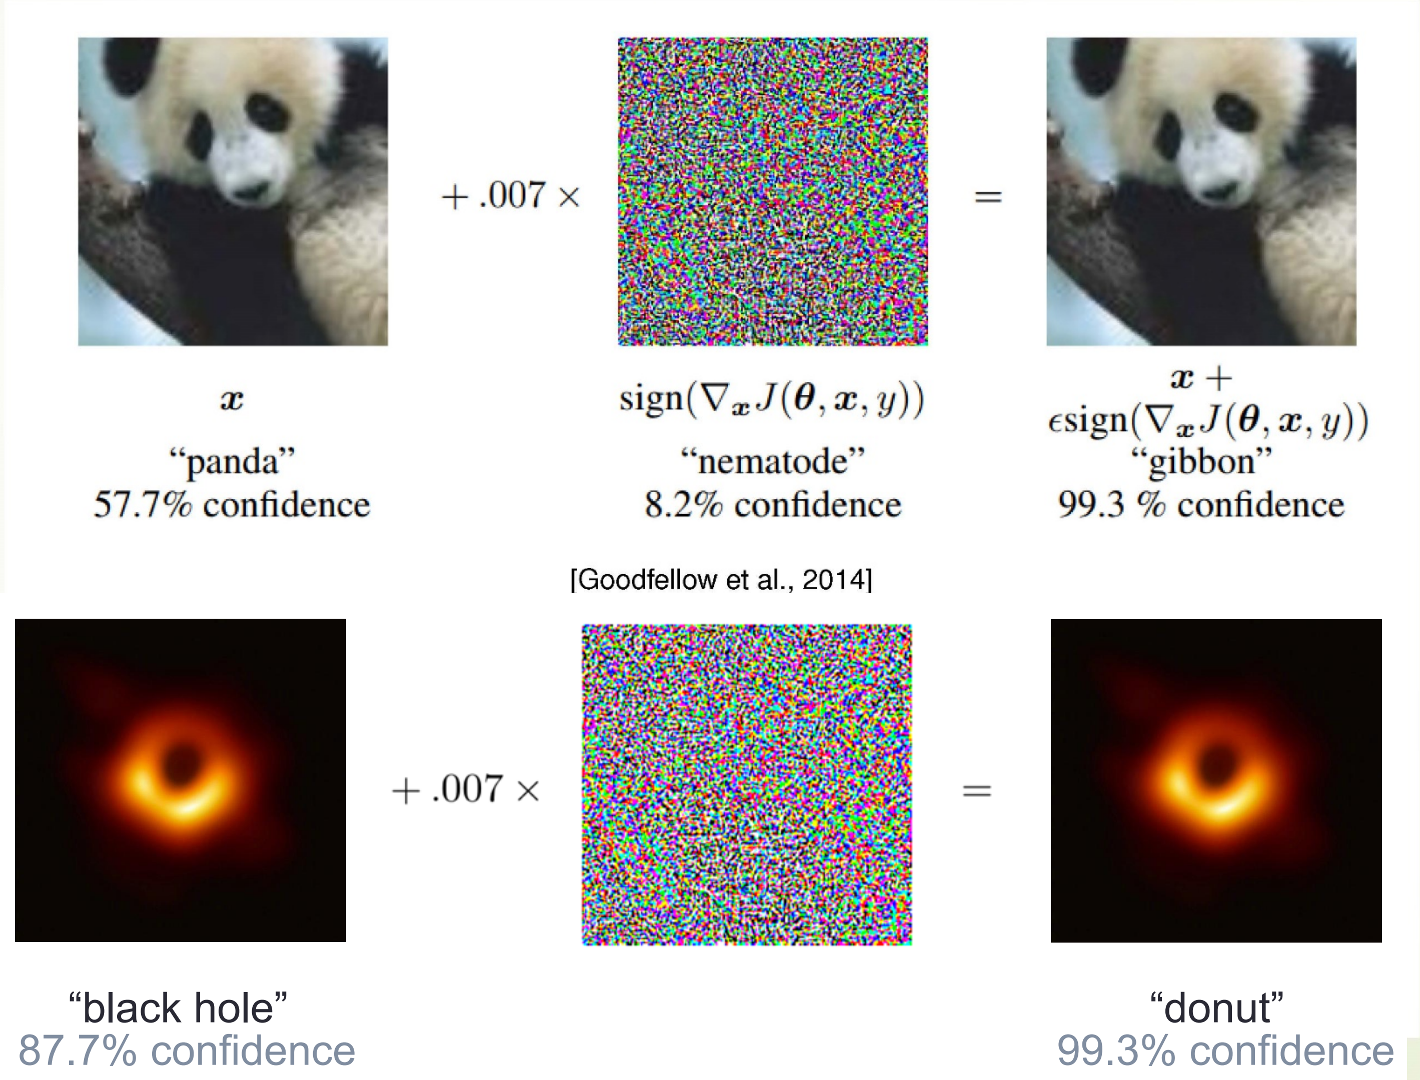

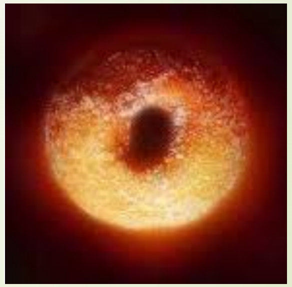

CNN learns texture features, not shapes

(a) Texture image
81.4% **Indian elephant**
10.3% indri
8.2% black swan

(b) Content image
71.1% **tabby cat**
17.3% grey fox
3.3% Siamese cat

(c) Texture-shape cue conflict
63.9% **Indian elephant**
26.4% indri
9.6% black swan

Geirhos et al. ICLR 2019

https://videoken.com/embed/W2HvLBMhCJQ?tocitem=46

1:16:47

Overfitting causes privacy leakage

- Model inversion attack leaks privacy

Figure: Recovered (Left), Original (Right)

Fredrikson et al. Proc. CCS, 2016

What's wrong with deep learning?

Ali Rahimi NIPS'17: Machine (deep) Learning has become **alchemy**.
https://www.youtube.com/watch?v=ORHFOnaEzPc

Yann LeCun CVPR'15, invited talk: **What's wrong with deep learning?**
One important piece: **missing some theory (clarity in understanding)**!

http://techtalks.tv/talks/whats-wrong-with-deep-learning/61639/

Being alchemy is certainly not a shame, not wanting to work on advancing to chemistry is a shame! -- **by Eric Xing**

"Shall we see soon an emergence from Alchemy to Science in deep leaning?"

How can we teach our students in the next generation science rather than alchemy?

Course Schedule		
Session	**Topic**	**Application & Case Study**
1	– History and Overview of Artificial Intelligence	– Case study on HireVue
2	– Introduction to Supervised Learning, Linear Regression and Classification	– Google Experiments: Draw! – Chatbots
3	– Model Assessment and Selection with Regularization: Ridge Regression and LASSO	– Google Experiments: Vision Sensing – Case study: WorkFusion (Robotic Process Automation)
4	– Decision Trees, Random Forests and Boosting	– Credit analysis
5	– Support Vector Machines	– Tutorial on Machine Learning with Python – Tutorial on GPU server
6	– An Introduction to Convolutional Neural Networks	– FinTech & Blockchain
7	– Transfer Learning and Neurostyle	– Natural language processing – Case study: Deep Instinct (Cybersecurity)
8	– An Introduction to Recurrent Neural Networks (RNN) and Long Short Term Memory (LSTM)	– Google Image Recognition – Case study: SenseTime (Computer Vision)
9	– An Introduction to Reinforcement Learning	– Competition of Cryptocurrency Trading with Deep Learning – Introduction to Deep Reinforcement Learning Trading
10	– Introduction to Unsupervised Learning: PCA, AutoEncoder, VAE and GANs	
11	– Investment Trends and FinTech Outlook	– Sales and Trading Business in Global Investment Banks – Ripe for Disruption by AI?
12	– Tutorial on Deep Learning in Python	– Exercise on Python Notebook
13	– Class Wrap	

Note: Details may change depending on class progress, development of relevant technologies, as well as information and feedback from students' surveys.

Featured Code Competition

Two Sigma Financial Modeling Challenge

Can you uncover predictive value in an uncertain world?

$100,000
Prize Money

Two Sigma · 2,066 teams · 3 years ago

Overview

How can we use the world's tools and intelligence to forecast economic outcomes that can never be entirely predictable? This question is at the core of countless economic activities around the world – including at Two Sigma Investments, who has been applying technology and systematic strategies to financial trading since 2001.

For over 15 years, Two Sigma has been at the forefront of applying technology and data science to financial forecasts. While their pioneering advances in big data, AI, and machine learning in the financial world have been pushing the industry forward, as with all other scientific progress, they are driven to make continual progress. Through this exclusive partnership, Two Sigma is excited to explore what untapped value Kaggle's diverse data science community can discover in the financial markets.

Economic opportunity depends on the ability to deliver singularly accurate forecasts in a world of uncertainty. By accurately predicting financial movements, Kagglers will learn about scientifically-driven approaches to unlocking significant predictive capability. Two Sigma is excited to find predictive value and gain a better understanding of the skills offered by the global data science crowd.

What is a Code Competition?

Welcome to Kaggle's very first Code Competition! In contrast to our traditional competitions, where competitors submit only prediction outputs, participants in Code Competitions will submit their code via Kaggle Kernels. All kernels are private by default in Code Competitions. You can build your models in Kernels by running them on a training set and, once you're ready to submit your code, your model's performance will be evaluated against the test set and your score and public leaderboard position revealed. As with our traditional competitions, we still maintain a private leaderboard test set, which your code is also evaluated against for final scoring, but is not revealed until the competition closes.

Since Code Competitions are brand new, we ask for your patience if you encounter bugs or frustrating platform quirks. Please report any issues you find in the forums and we'll do our best to respond.

Emerging Technologies

Emerging Technologies 前沿科技	Foundational 底层科技	Cumulative and exponential 叠加和倍增	Capital Intensive 资本密集	We know “how” to do it 实施方法明确	We know the “end game” 最终目标明确	Adoption 接纳程度
Artificial Intelligence 人工智能	Yes	Yes	No	No	No	Relatively easy
Blockchain 区块链	Yes	No	No	Yes	Yes	Easy & Difficult[1]
AR/VR 虚拟现实和增强现实	No	No	Yes	Yes	Yes	Moderate
Robotics 机器人	No	No	Yes	Yes	Yes	Moderate
Internet of Things (IoT) 物联网	Yes	Yes	Yes	Depends[2]	Yes	Relatively easy
Space Tech 太空科技	No	No	Yes	Yes & No[3]	Yes	Difficult

Notes:
1. Adoption at the front end is relatively straightforward, while the back end can be challenging due to legacy issues.
2. IoT is dependent on a lot of other technologies, systems, and infrastructure (e.g. 5G, smart contracts).
3. Current technology (i.e. combustion-based) makes Mars reachable, but not beyond.

History of A.I.

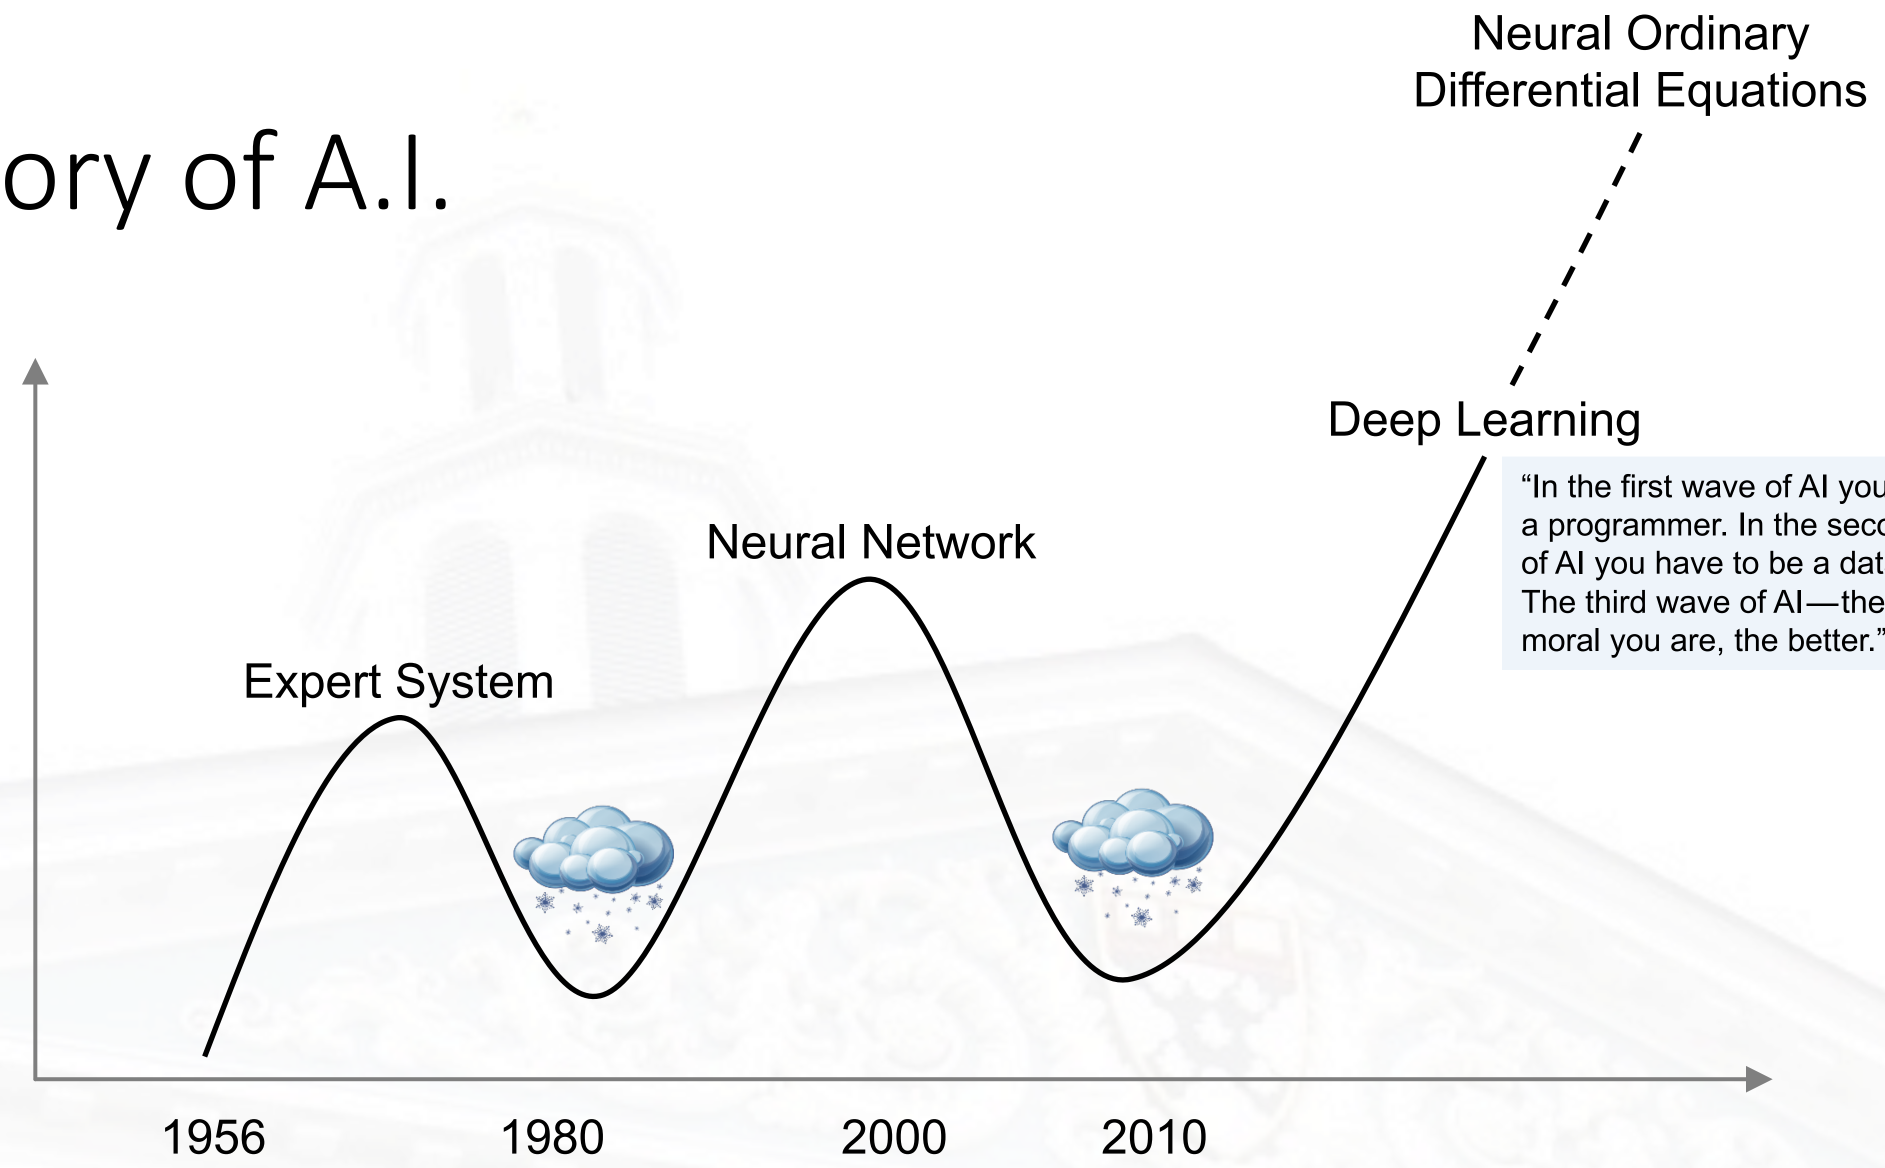

"In the first wave of AI you had to be a programmer. In the second wave of AI you have to be a data scientist. The third wave of AI—the more moral you are, the better."

Source: Dr. Lee Kai-Fu • Sinovation Ventures and 2018 Conference on Neural Information Processing Systems

Credit: https://www.pinterest.es/pin/570620215256233288/

Credit: https://www.slideshare.net/bretmc/machine-learning-with-google-machine-learning-apis-puppy-or-muffin

Credit: https://www.slideshare.net/bretmc/machine-learning-with-google-machine-learning-apis-puppy-or-muffin

Looking for:

- Dog (or breed)
- Muffin

Credit: https://www.slideshare.net/bretmc/machine-learning-with-google-machine-learning-apis-puppy-or-muffin

source: boredpanda.com

A Tale of Two AI Camps

Property	Symbolic AI	Connectionist AI
Knowledge Acquisition	✓ Theoretical knowledge insertion can be made in a simple and direct manner. It is sufficient to clarify, convert, and formalize the knowledge.	× Theoretical knowledge may not useful in constructing neural networks, while examples are always required for knowledge acquisition.
Processing Mode	× Processing is sequential. Answer and consultation times are long.	✓ Neural networks consist of a set of units with information processing completed in a parallel fashion.
New Knowledge Insertion	✓ Insertion of knowledge (rules) can be made very quickly once experts have already processed them.	× Training process is often time-consuming as weights and biases are trained gradually.
Training	× Training is not a basic process. Knowledge acquisition is done by explicitation, with potential bottleneck issues.	✓ Training and generalization from examples are fundamental and integrate processes.
Results Explanation	✓ Reasoning process allows for explanability. Knowledge is coded in a language close to natural language, and therefore easily interpretable.	× Neural networks are “black” boxes, where knowledge is coded in weights and interconnections, with a lack of access to a form that is interpretable by humans.

Source: Souici-Meslati, Labiba & Sellami, Mokhtar. (2004). A hybrid approach for Arabic literal amounts recognition. The Arabian Journal for Science and Engineering. 29.

A Tale of Two AI Camps

Property	Symbolic AI	Connectionist AI
Information Processing	× Theoretical knowledge must be complete beforehand, and the approach is not conducive to approximate or incomplete information processing.	✓ Neural networks are conducive to approximate and incomplete information processing (i.e. fuzzy logic).
Knowledge Coding	✓ Knowledge is represented by rules and data structures.	✓ Knowledge is coded in networks representing the relationships among input variables.
Development	× Long development cycles with domain experts are typical.	× Architecture and (hyper)parameters derivation and tuning can be time-consuming and difficult.
Maintenance	× Managing and maintaining large databases of rules are challenging. Adding new rules and updating existing ones may be difficult.	✓ Maintenance and management are often easy, and networks can be retrained based on changes in situational factors.

Source: Souici-Meslati, Labiba & Sellami, Mokhtar. (2004). A hybrid approach for Arabic literal amounts recognition. The Arabian Journal for Science and Engineering. 29.

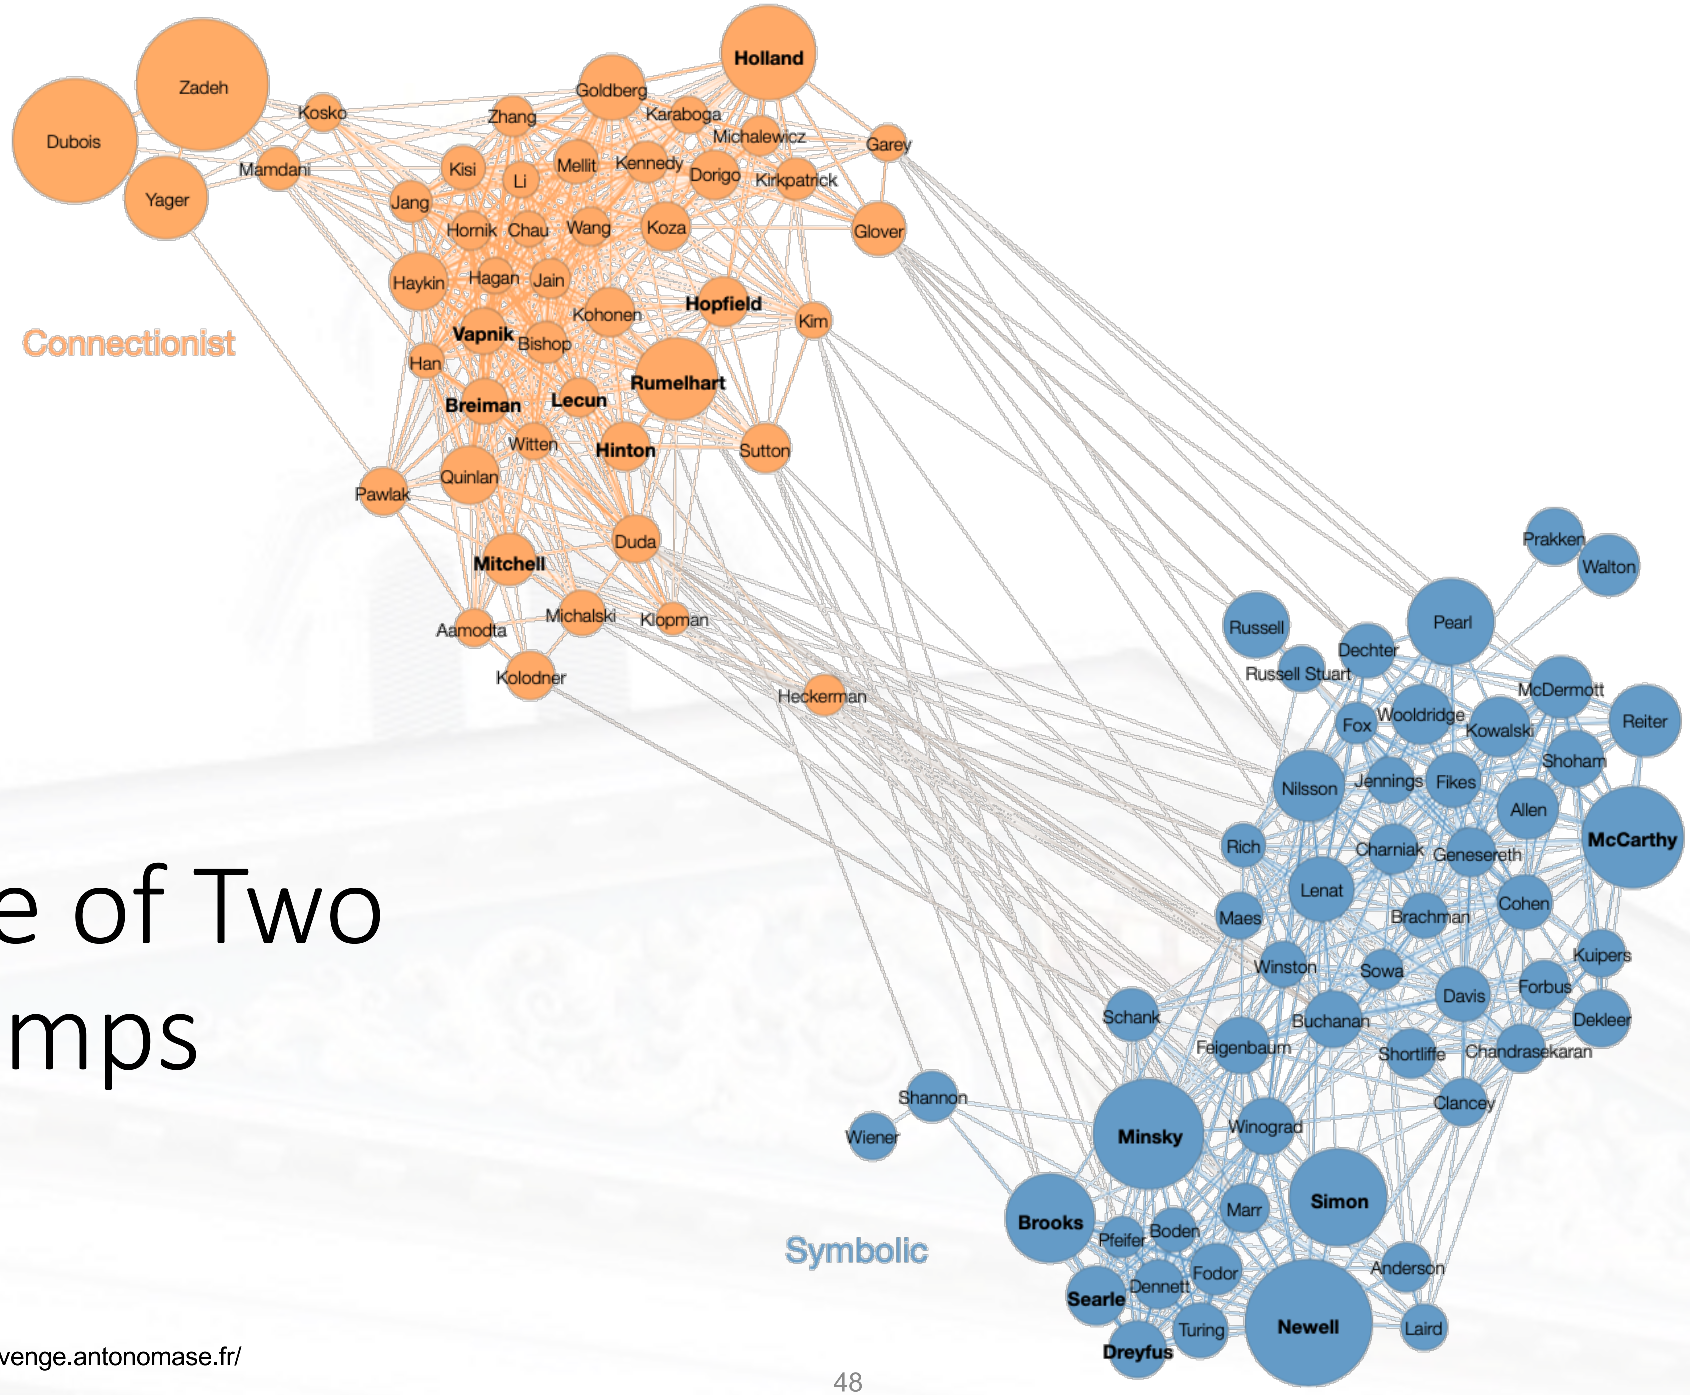

A Tale of Two AI Camps

Source: https://neurovenge.antonomase.fr/

A Tug of War

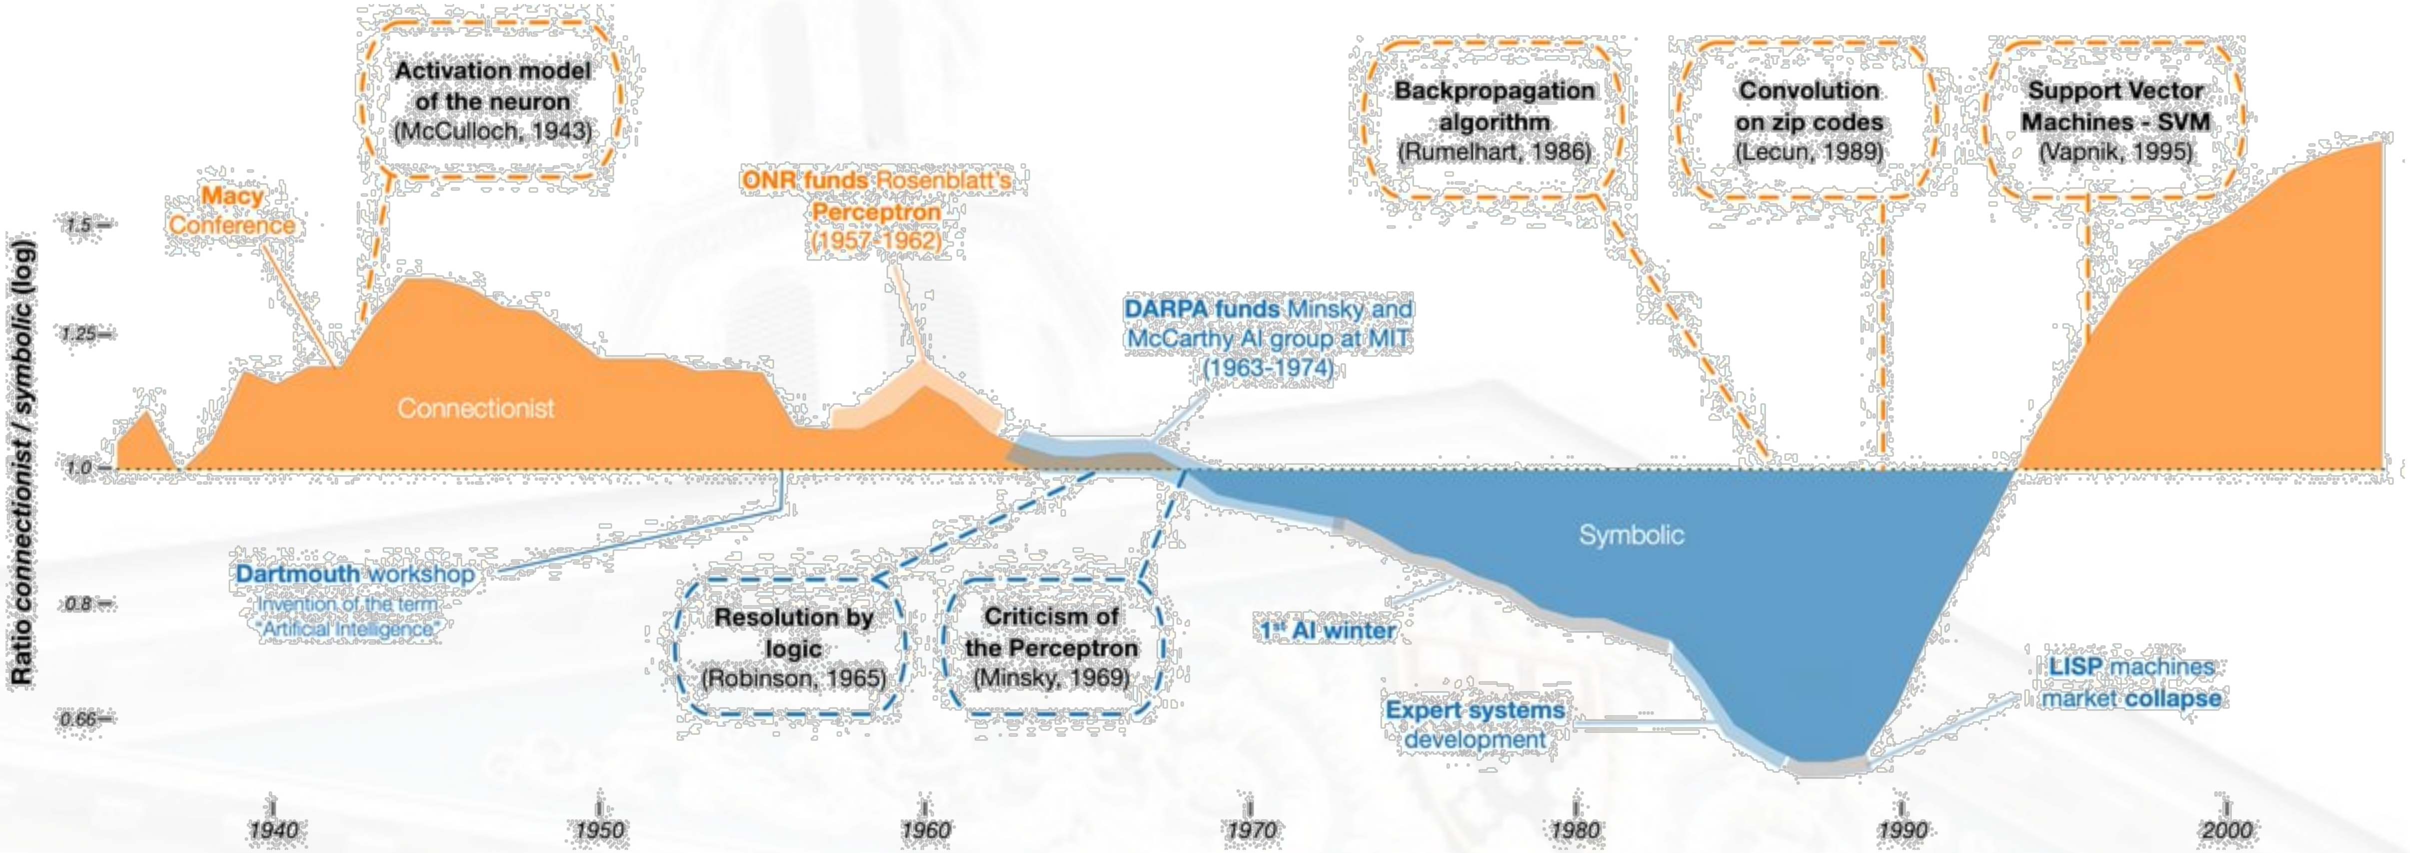

Source: https://neurovenge.antonomase.fr/

The Five Elements

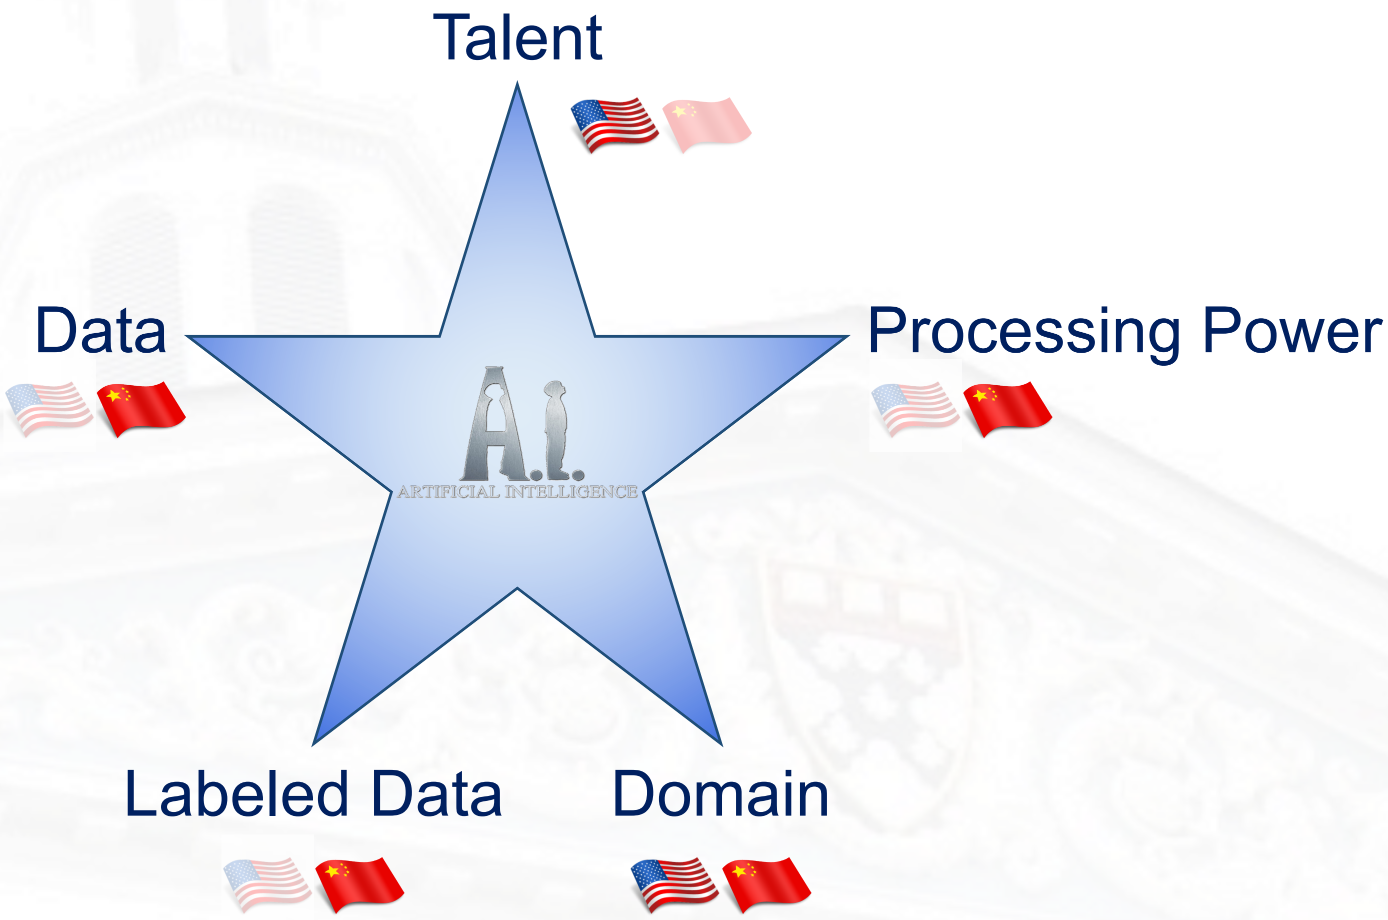

China seeks semiconductor security in wake of ZTE ban

Stubborn technology gap has frustrated Beijing's bid to build a world-class chip sector

Save

Taiwan Semiconductor Manufacturing, the world's biggest contract chipmaker, says Chinese companies will struggle to catch up to foreign rivals © FT montage / Reuters

Edward White in Hsinchu 11 HOURS AGO

15

When the US administration shut down Chinese telecoms equipment maker ZTE in April — putting the future of a $17bn company and 75,000 jobs at risk after sanctions cut the supply of key microchips — it highlighted a vulnerability in the Chinese economy: it depends on foreign-made chips.

The country remains dependent on imports to build the phones, telecoms gear,

商汤 sensetime

AI for a Better Tomorrow

SenseTime at a Glance

- Largest AI pure-play globally in terms of sales revenue and enterprise value, offering AI-embedded recognition software and integrated solutions with diverse revenues from 10+ verticals and 700+ top-notch customers
- Fastest growing technology platform company in Asia, with 3000+ employees operating in 10 countries, headquartered in Hong Kong
- Asia's largest AI research team with 200 PhDs (including 40 professors) leading 1100 researchers
- ~3x yearly revenue growth in the past 3 consecutive years; profitable since 2017
- Attracted $2.3bn funding in 2018 from Softbank, Silverlake, Temasek, Fidelity, Alibaba, Qualcomm, etc.

Key Talents

Founding Team — Research scientists from the Multimedia Lab at The Chinese University of Hong Kong; published numerous papers in computer vision and was the earliest research team in Asia to study deep learning

Research — Largest AI research team in Asia with 200 PhDs (incl. 40 professors) leading 1100 researchers; ranked #3 in the world and #1 in Asia in terms of computer vision

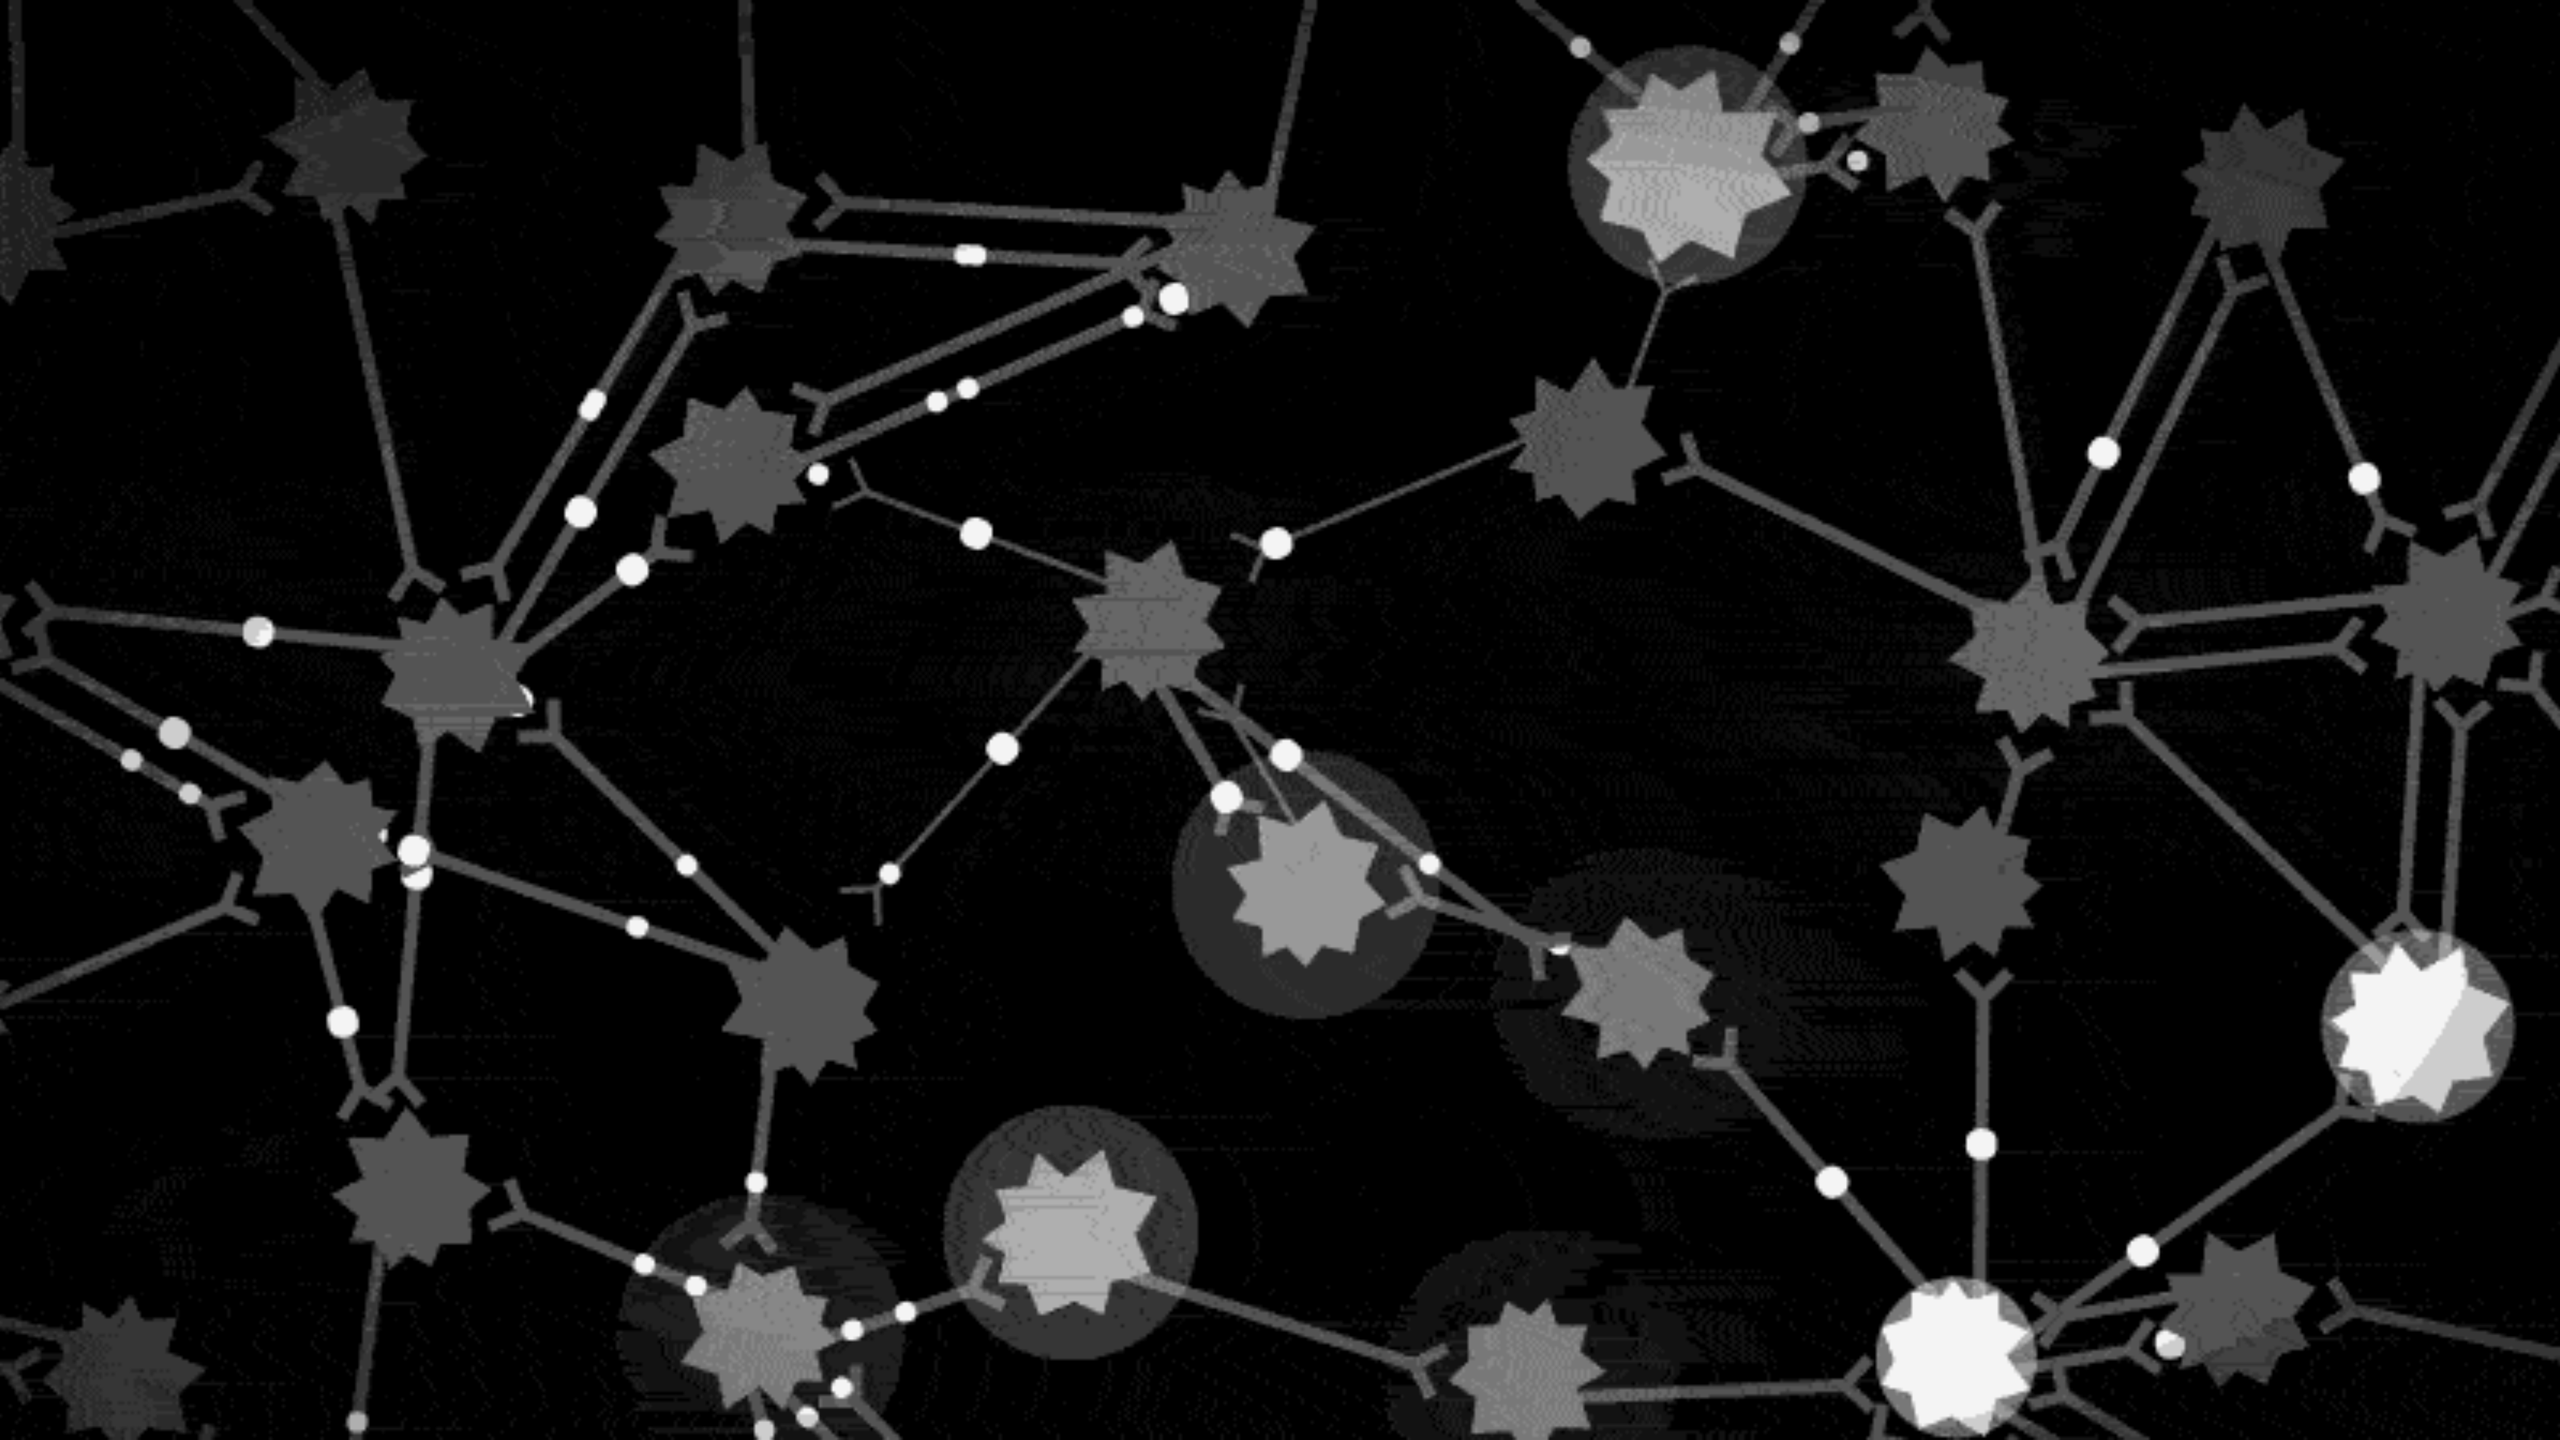

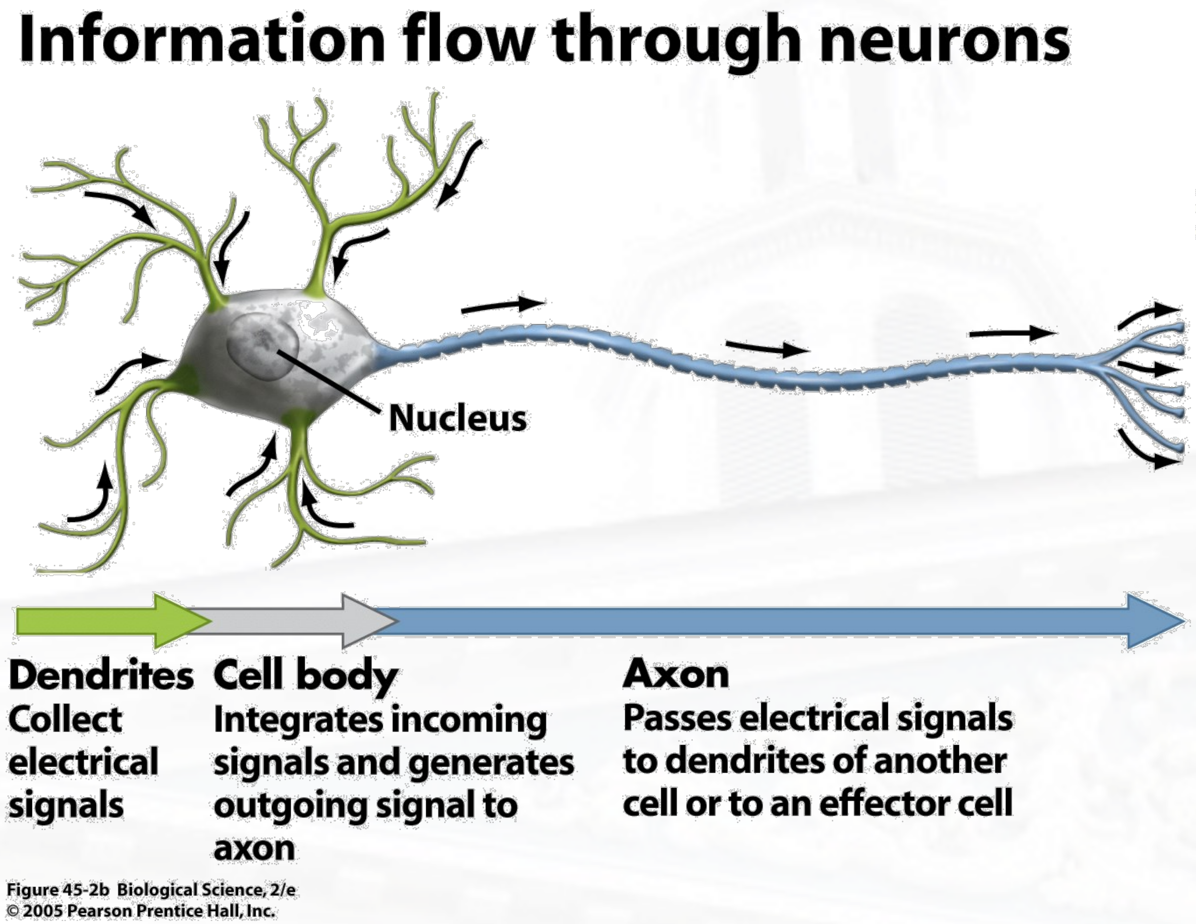

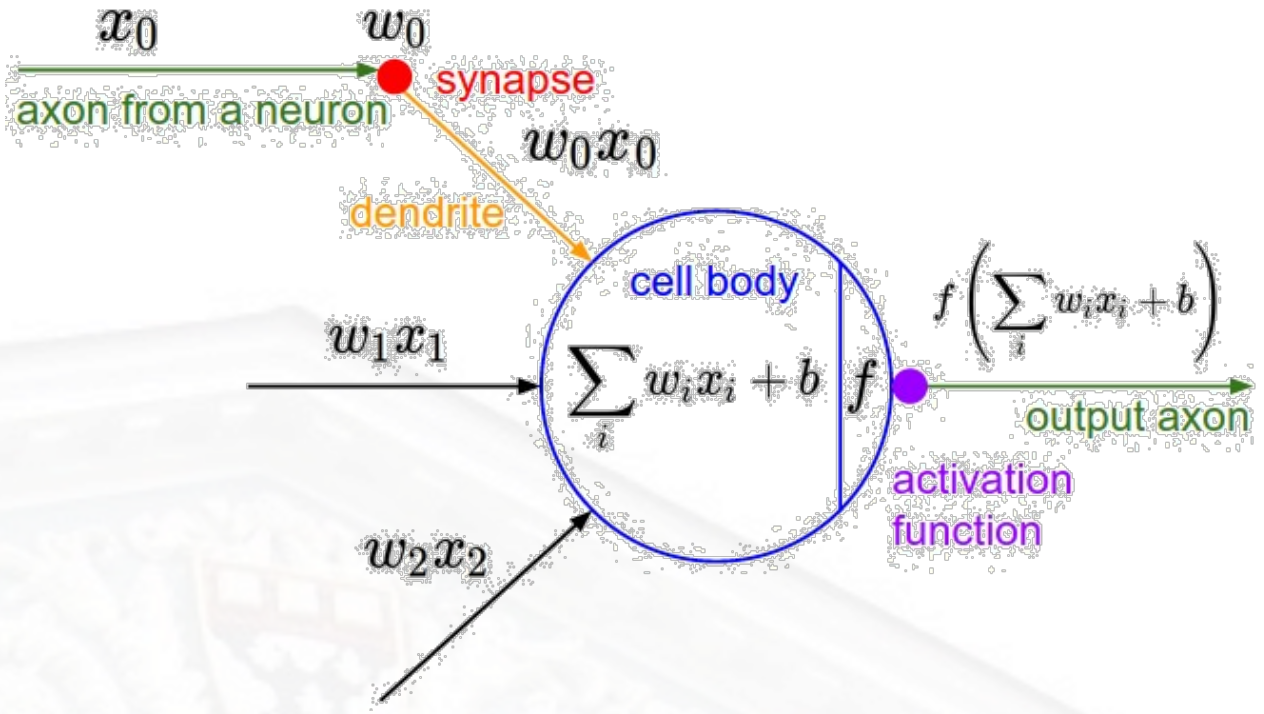

Credit (Left): https://www.psychologyinaction.org/psychology-in-action-1/2011/04/01/conventional-wisdom-upset-persistent-action-potential-firing-in-distal-axons
Credit (Right): http://cs231n.github.io/convolutional-networks/

AI Development Timeline

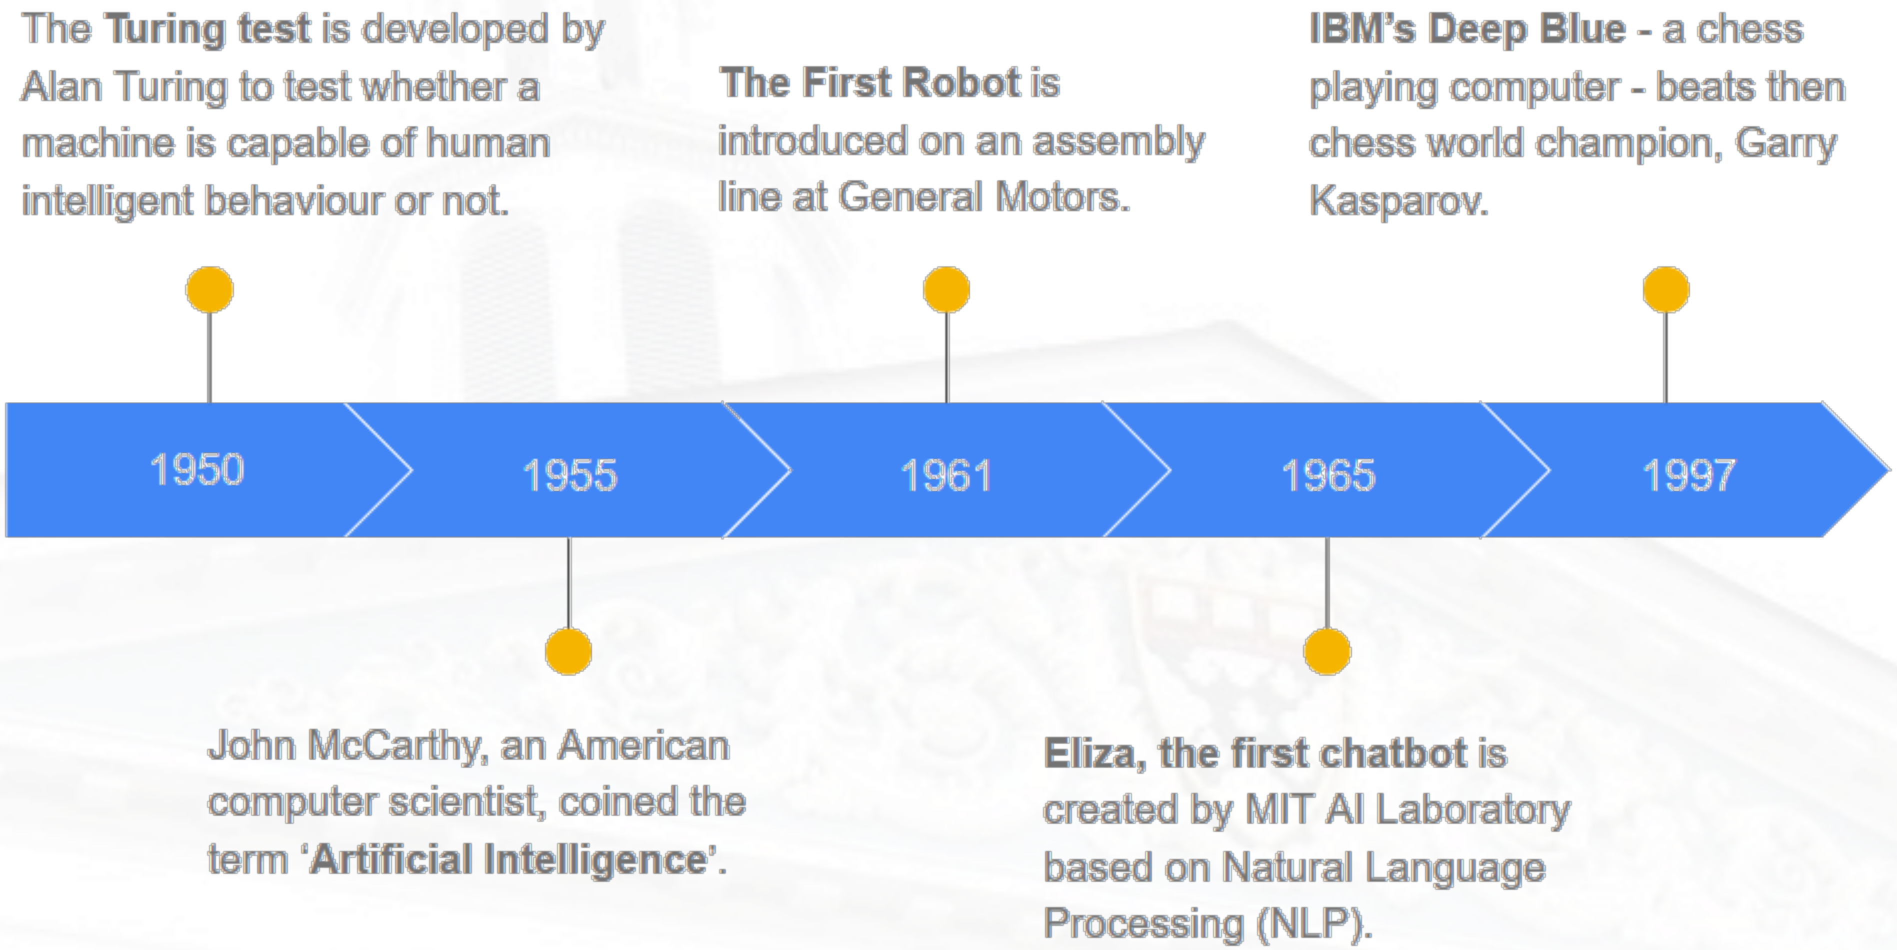

Credit: https://www.analyticsinsight.net/top-artificial-intelligence-technologies-past-present-future/

AI Development Timeline (Cont'd)

1999 – **First emotional AI demo - Kismet** - by MIT AI Laboratory

2009 – Google starts building **self-driving cars.**

2011 – **IBM Watson** beats *Jeopardy* champions. Same year **Siri, Google Now and Cortana** become mainstream.

2016 – **Google DeepMind's AlphaGo beats world champion Lee Sedol in Go**, an ancient Chinese board game.

2017 – **Medical breakthroughs with AI** identify skin cancer, diagnose irregular heart rhythms.

Credit: https://www.analyticsinsight.net/top-artificial-intelligence-technologies-past-present-future/

Deep Learning Timeline

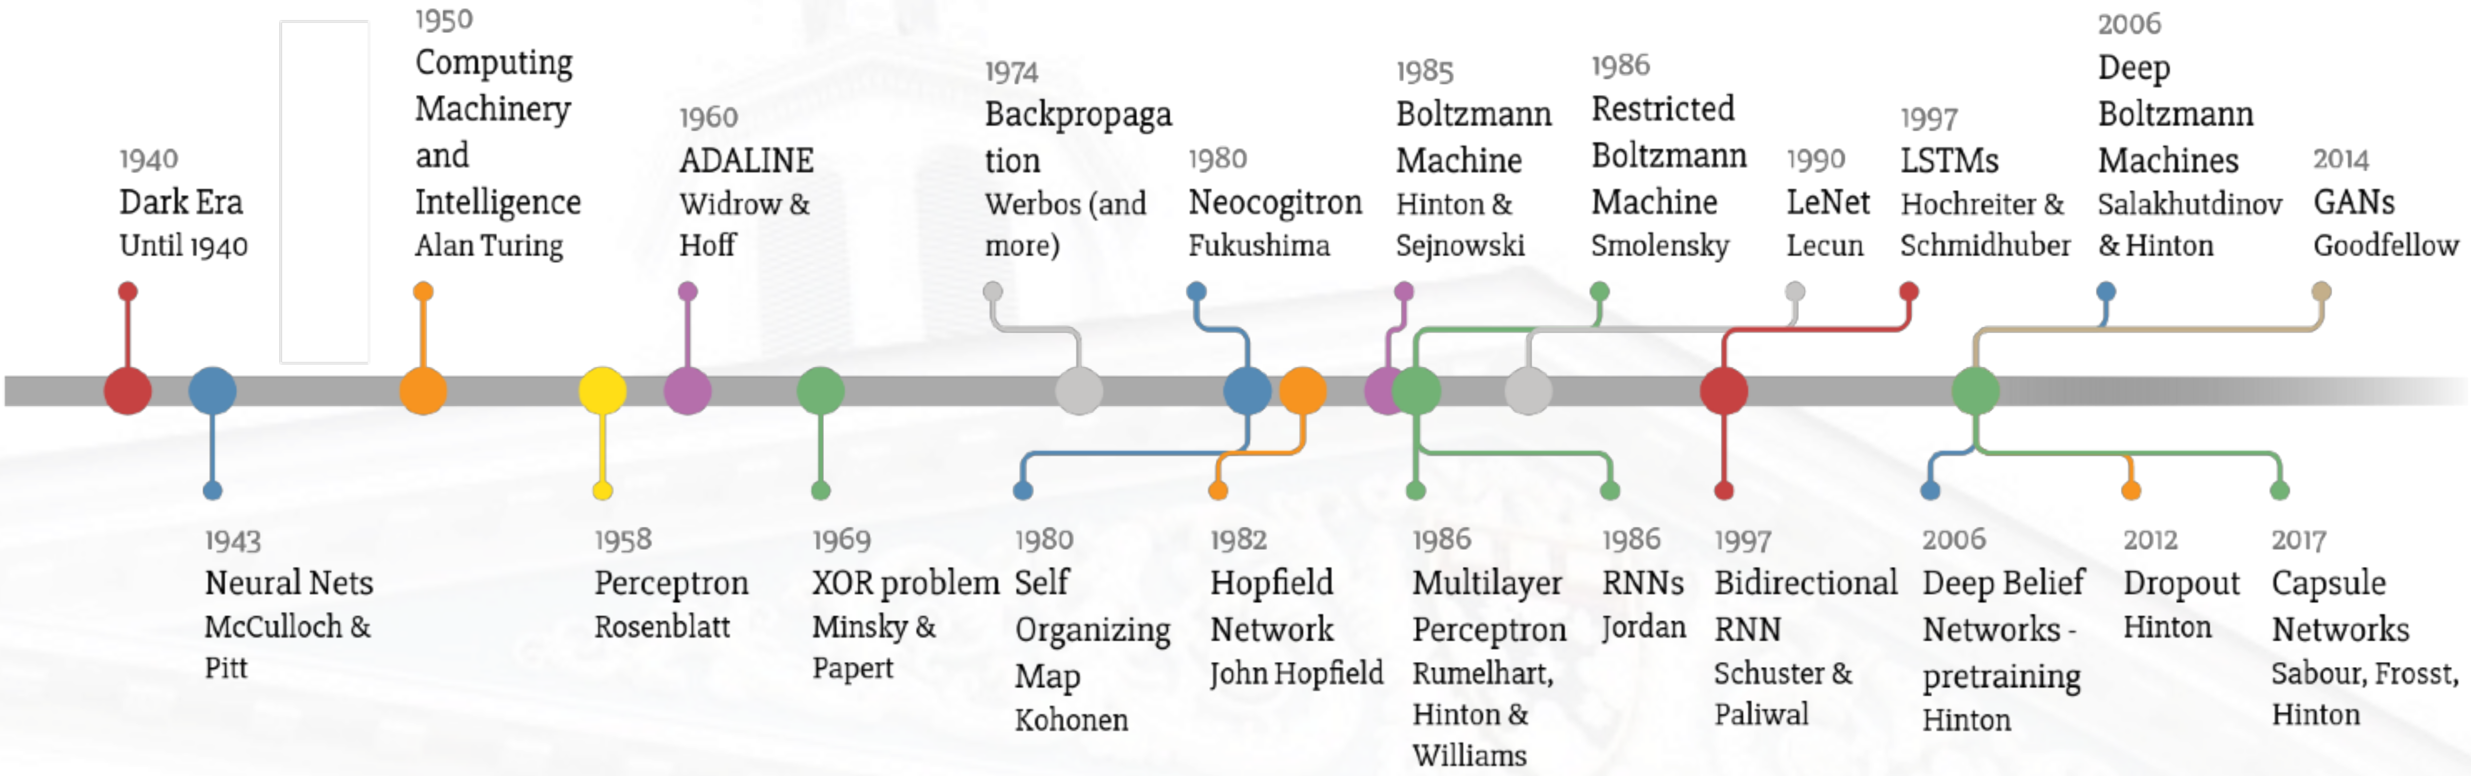

Made by Favio Vázquez

Credit: https://towardsdatascience.com/a-weird-introduction-to-deep-learning-7828803693b0

https://quickdraw.withgoogle.com/

Can a neural network learn to recognize doodling?

Help teach it by adding your drawings to the world's largest doodling data set, shared publicly to help with machine learning research.

Let's Draw!

www.slido.com
#UST
sli.do

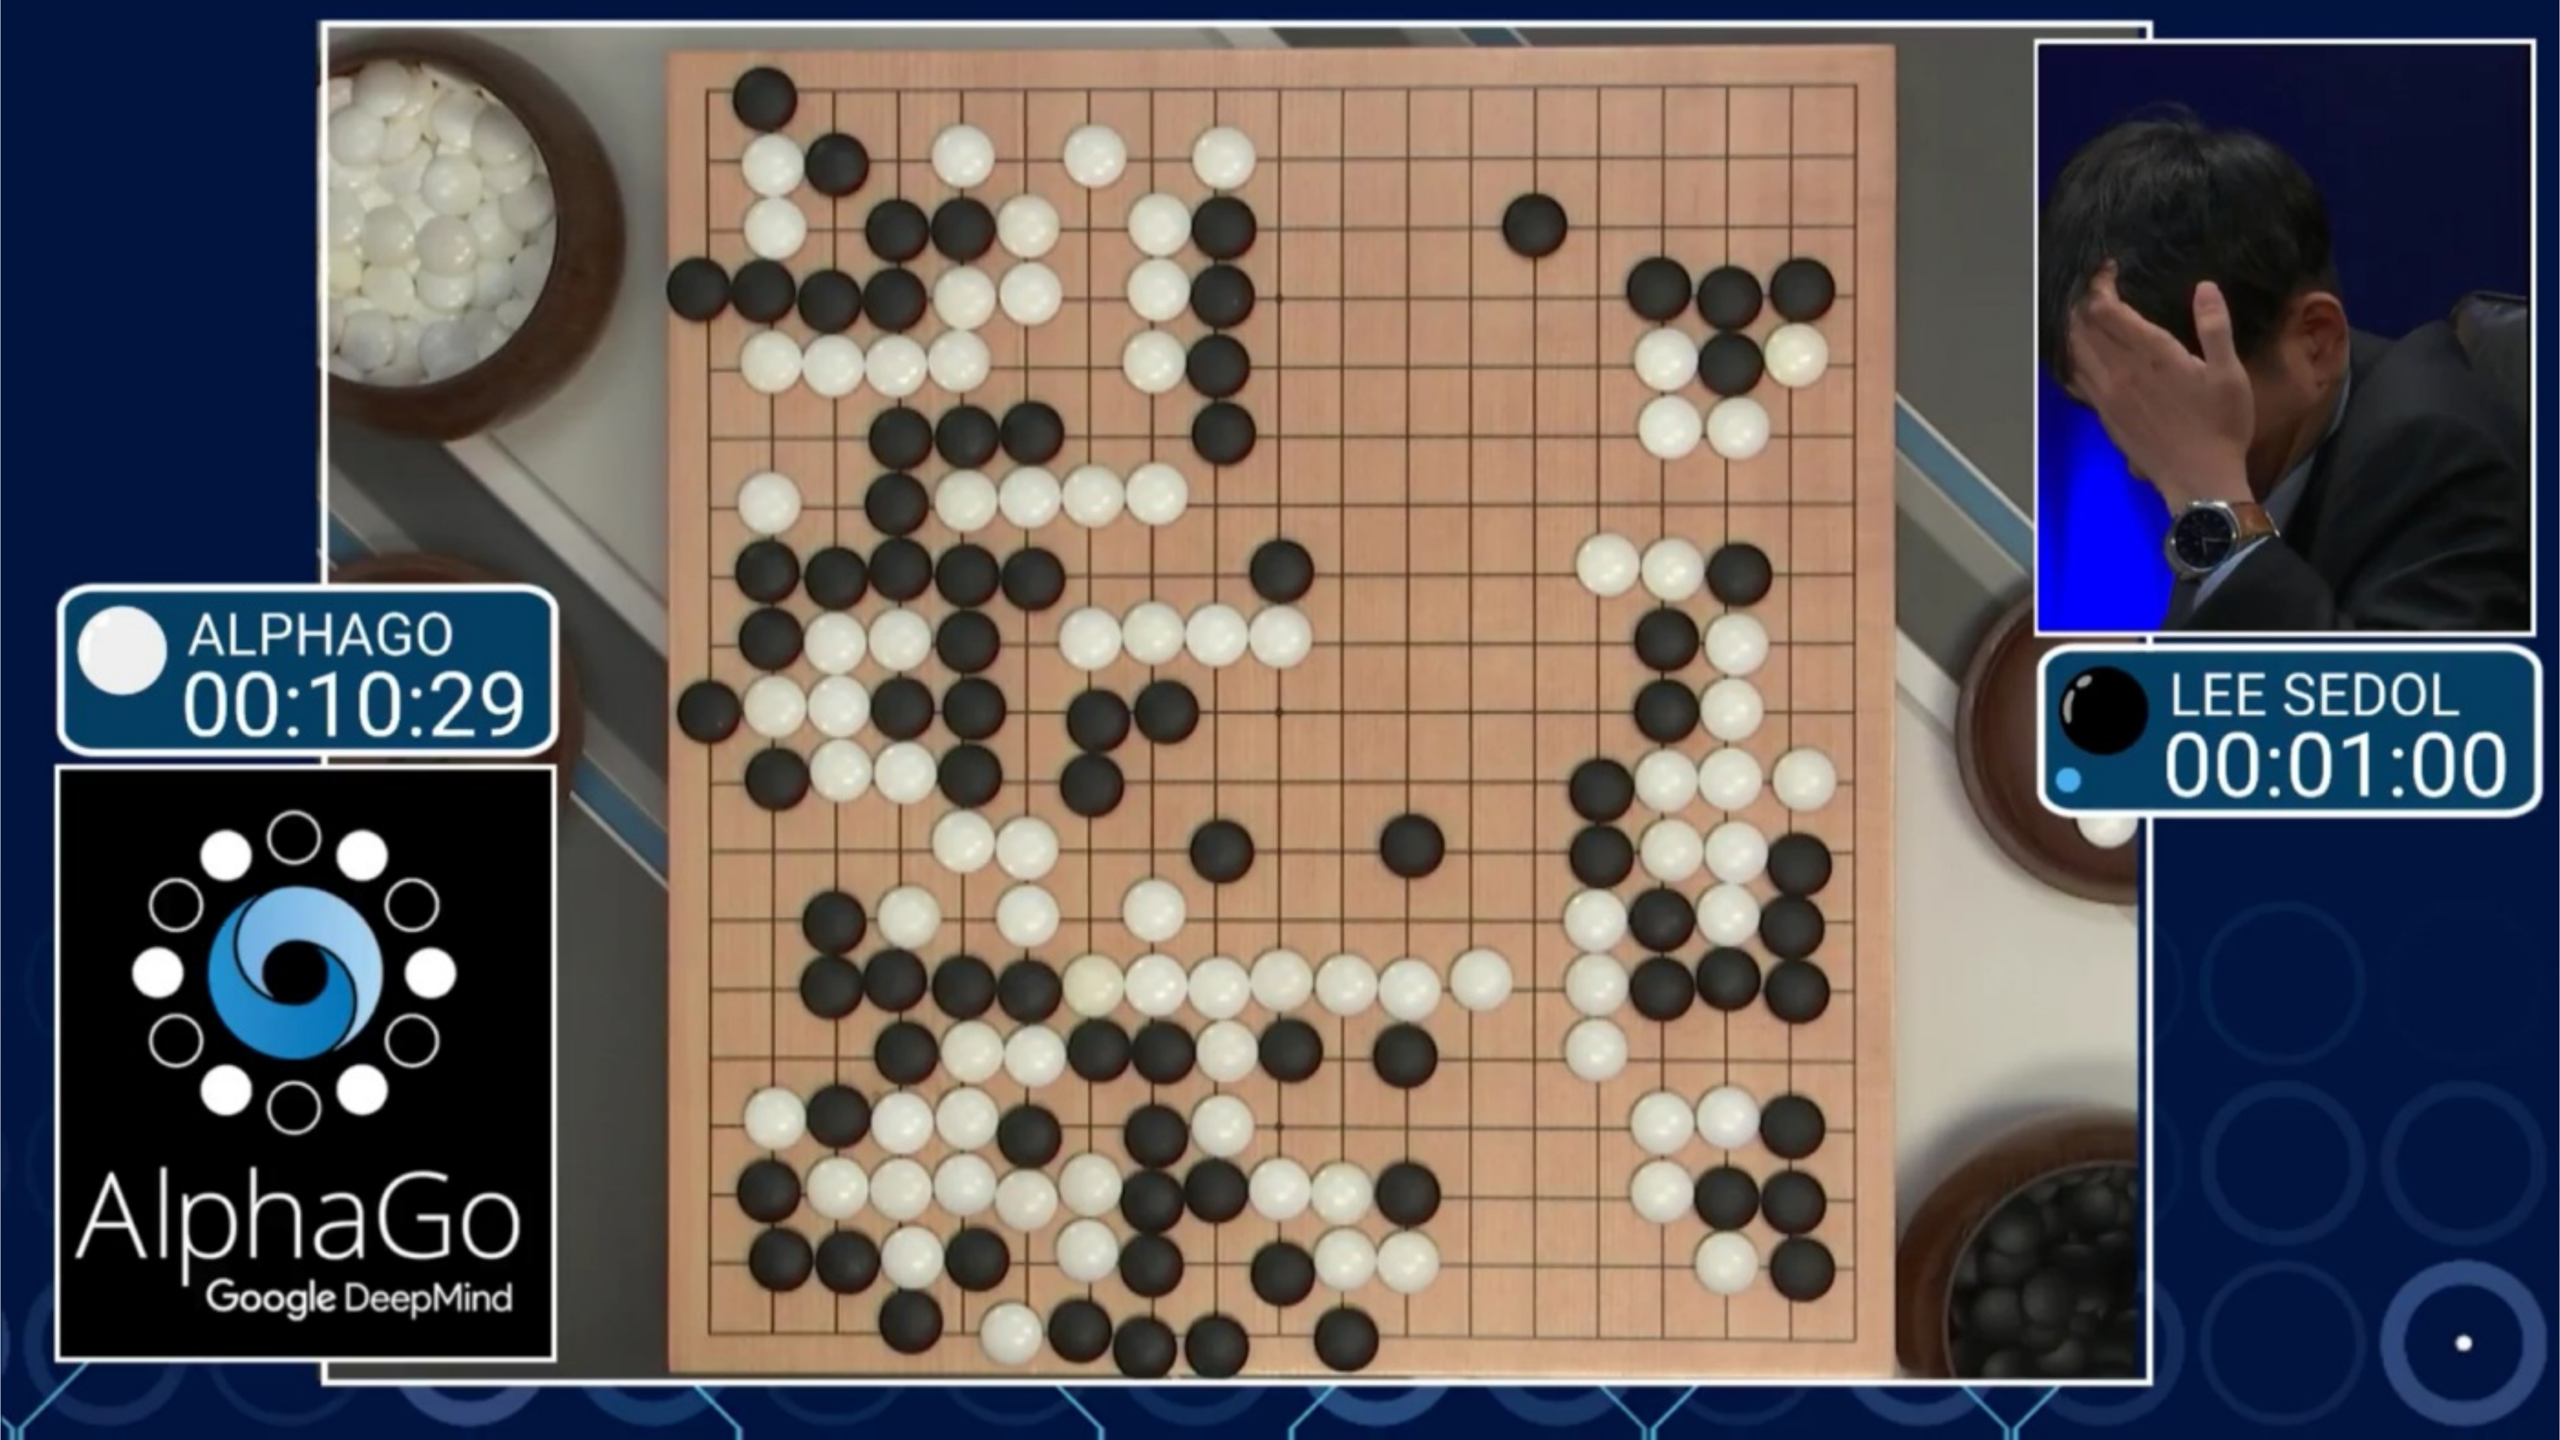
ALPHAGO
00:10:29
LEE SEDOL
00:01:00
AlphaGo
Google DeepMind

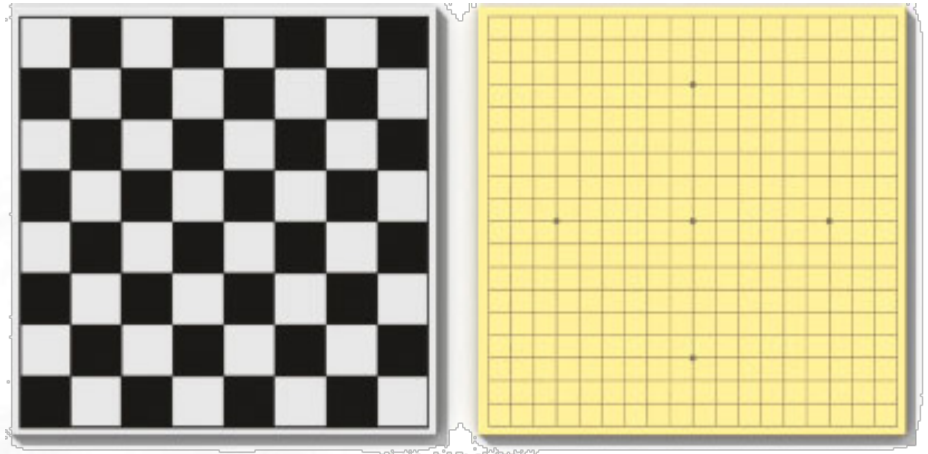

GRID SIZE

8 x 8	19 x 19

AVERAGE NUMBER OF MOVE CHOICES PER TURN

35	200–300

LENGTH OF TYPICAL GAME

60 moves	200 moves

NUMBER OF POSSIBLE GAME POSITIONS

10^{44}	10^{170}

EXPLOSION OF CHOICES
(starting from average game position)

35	Move 1	200
1225	Move 2	40 000
42 875	Move 3	8 000 000
1 500 625	Move 4	1 600 000 000

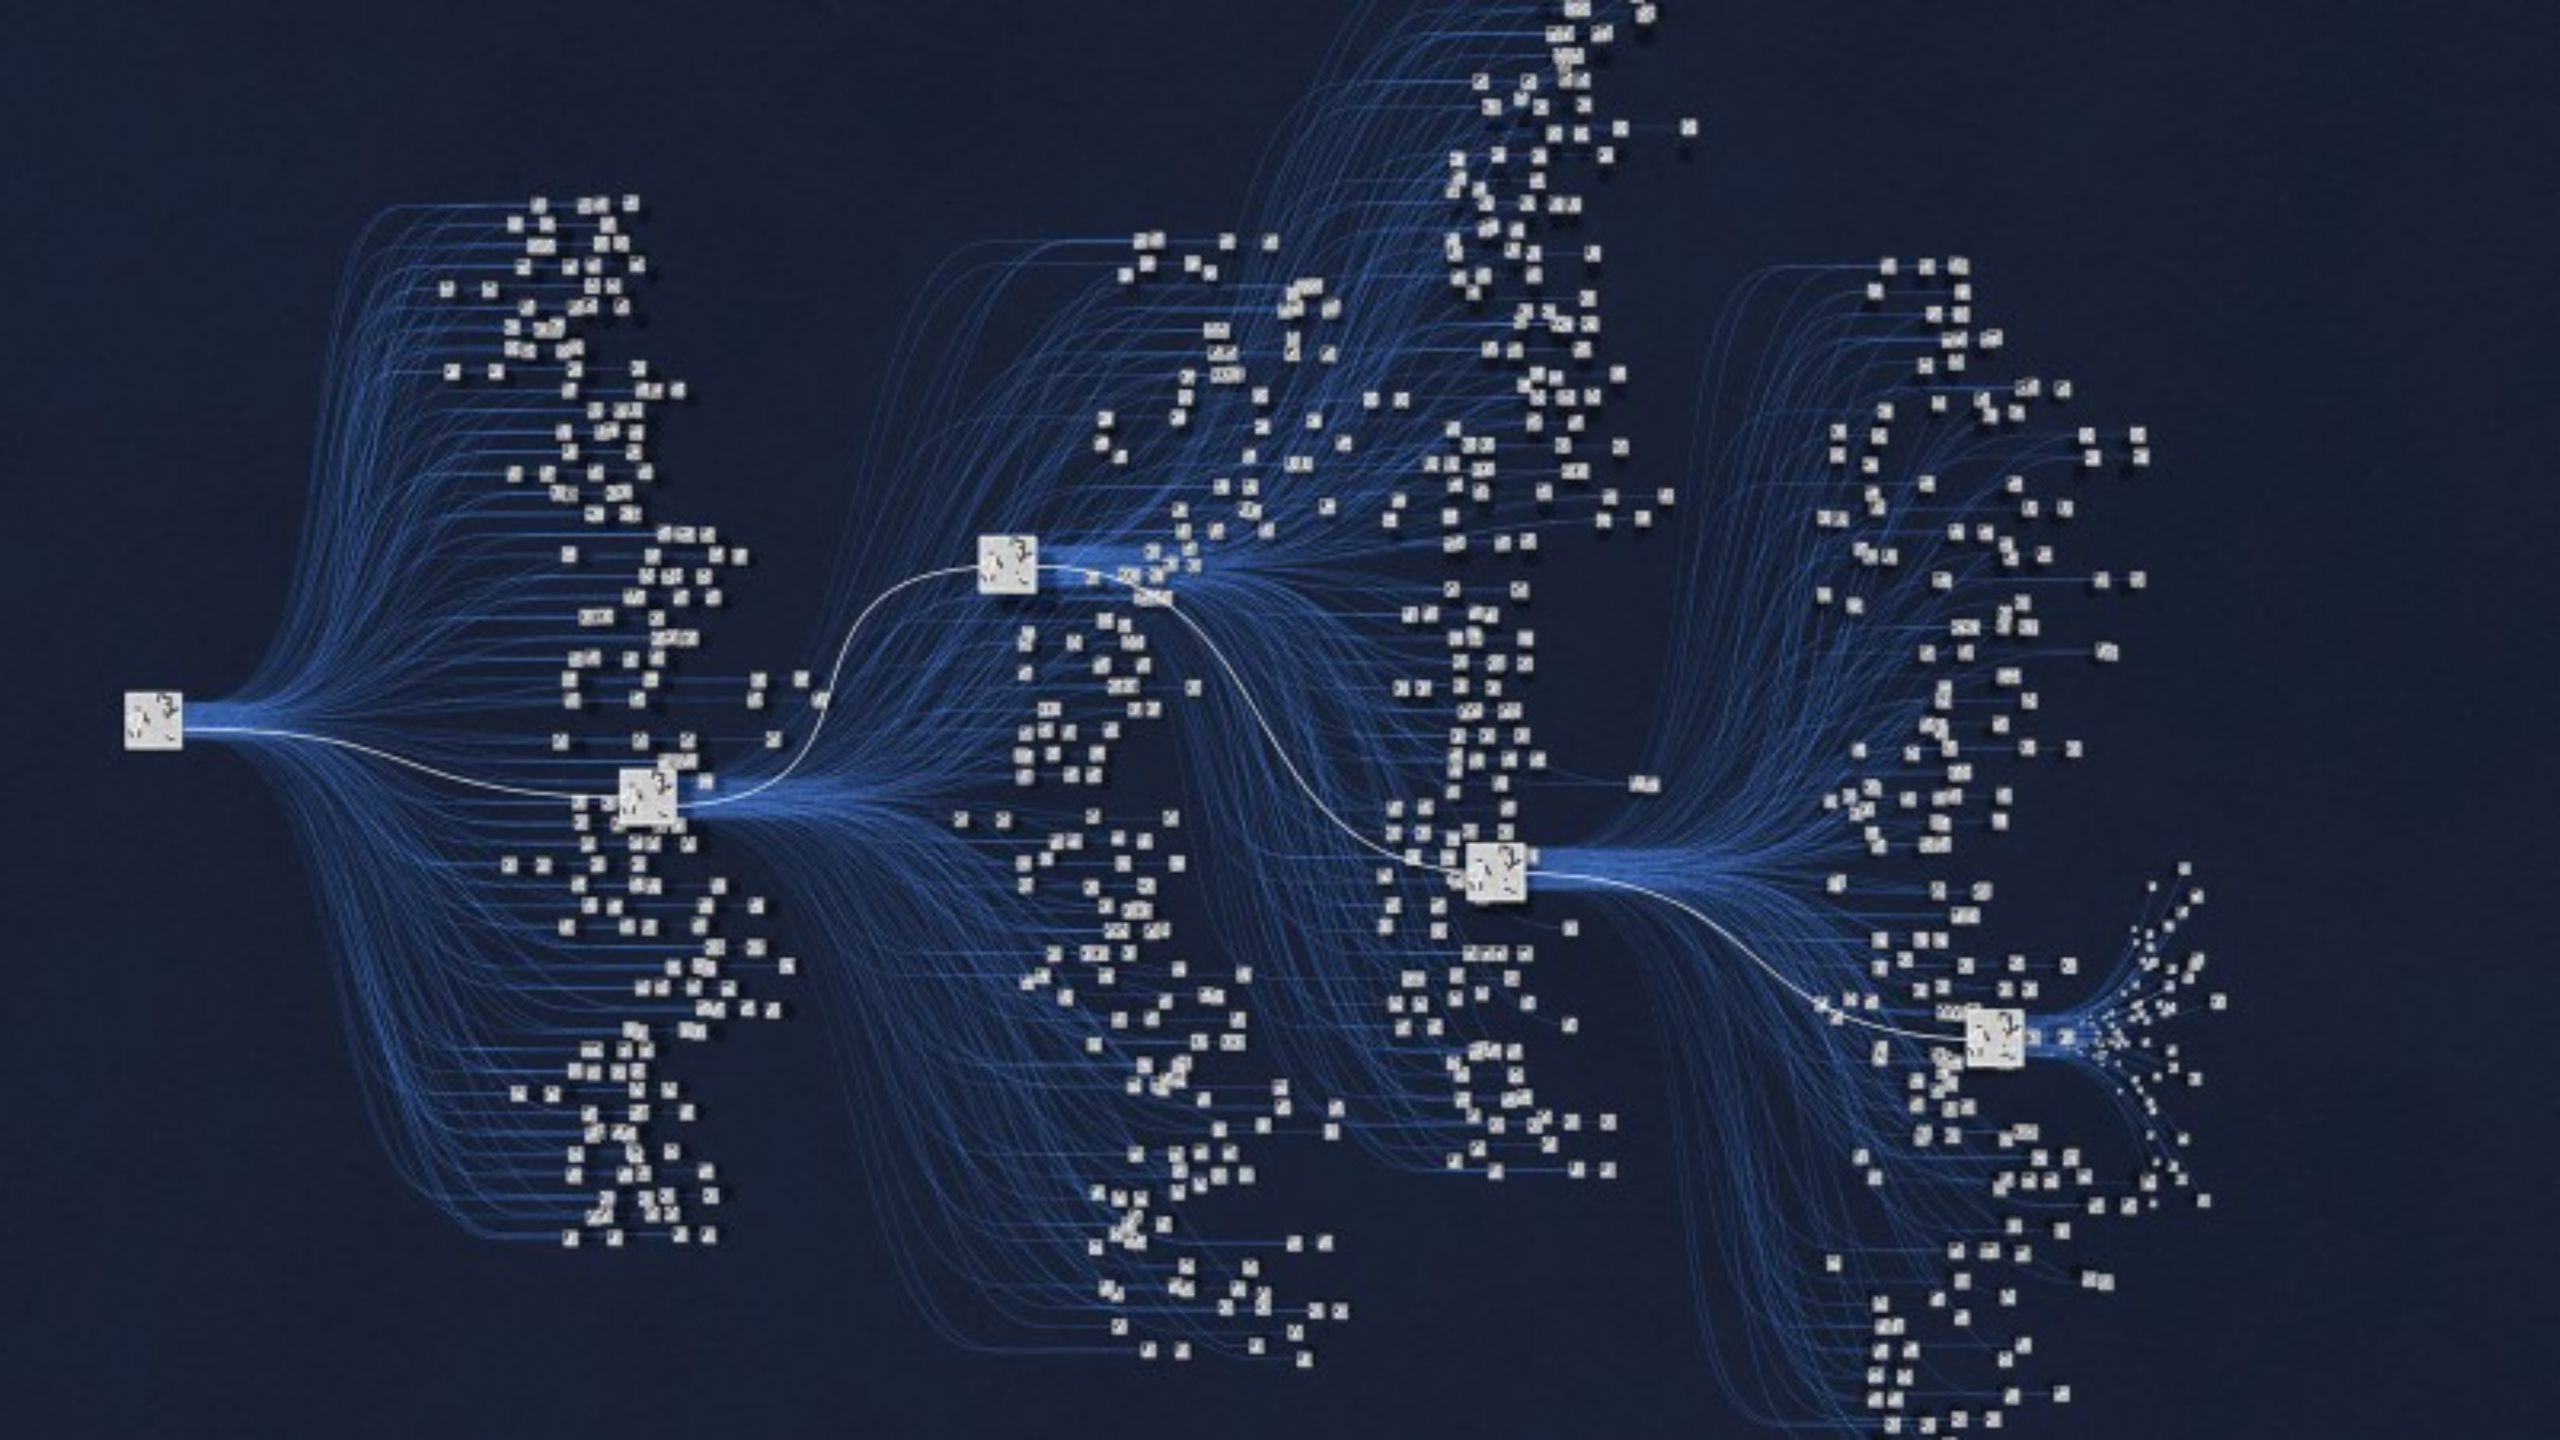

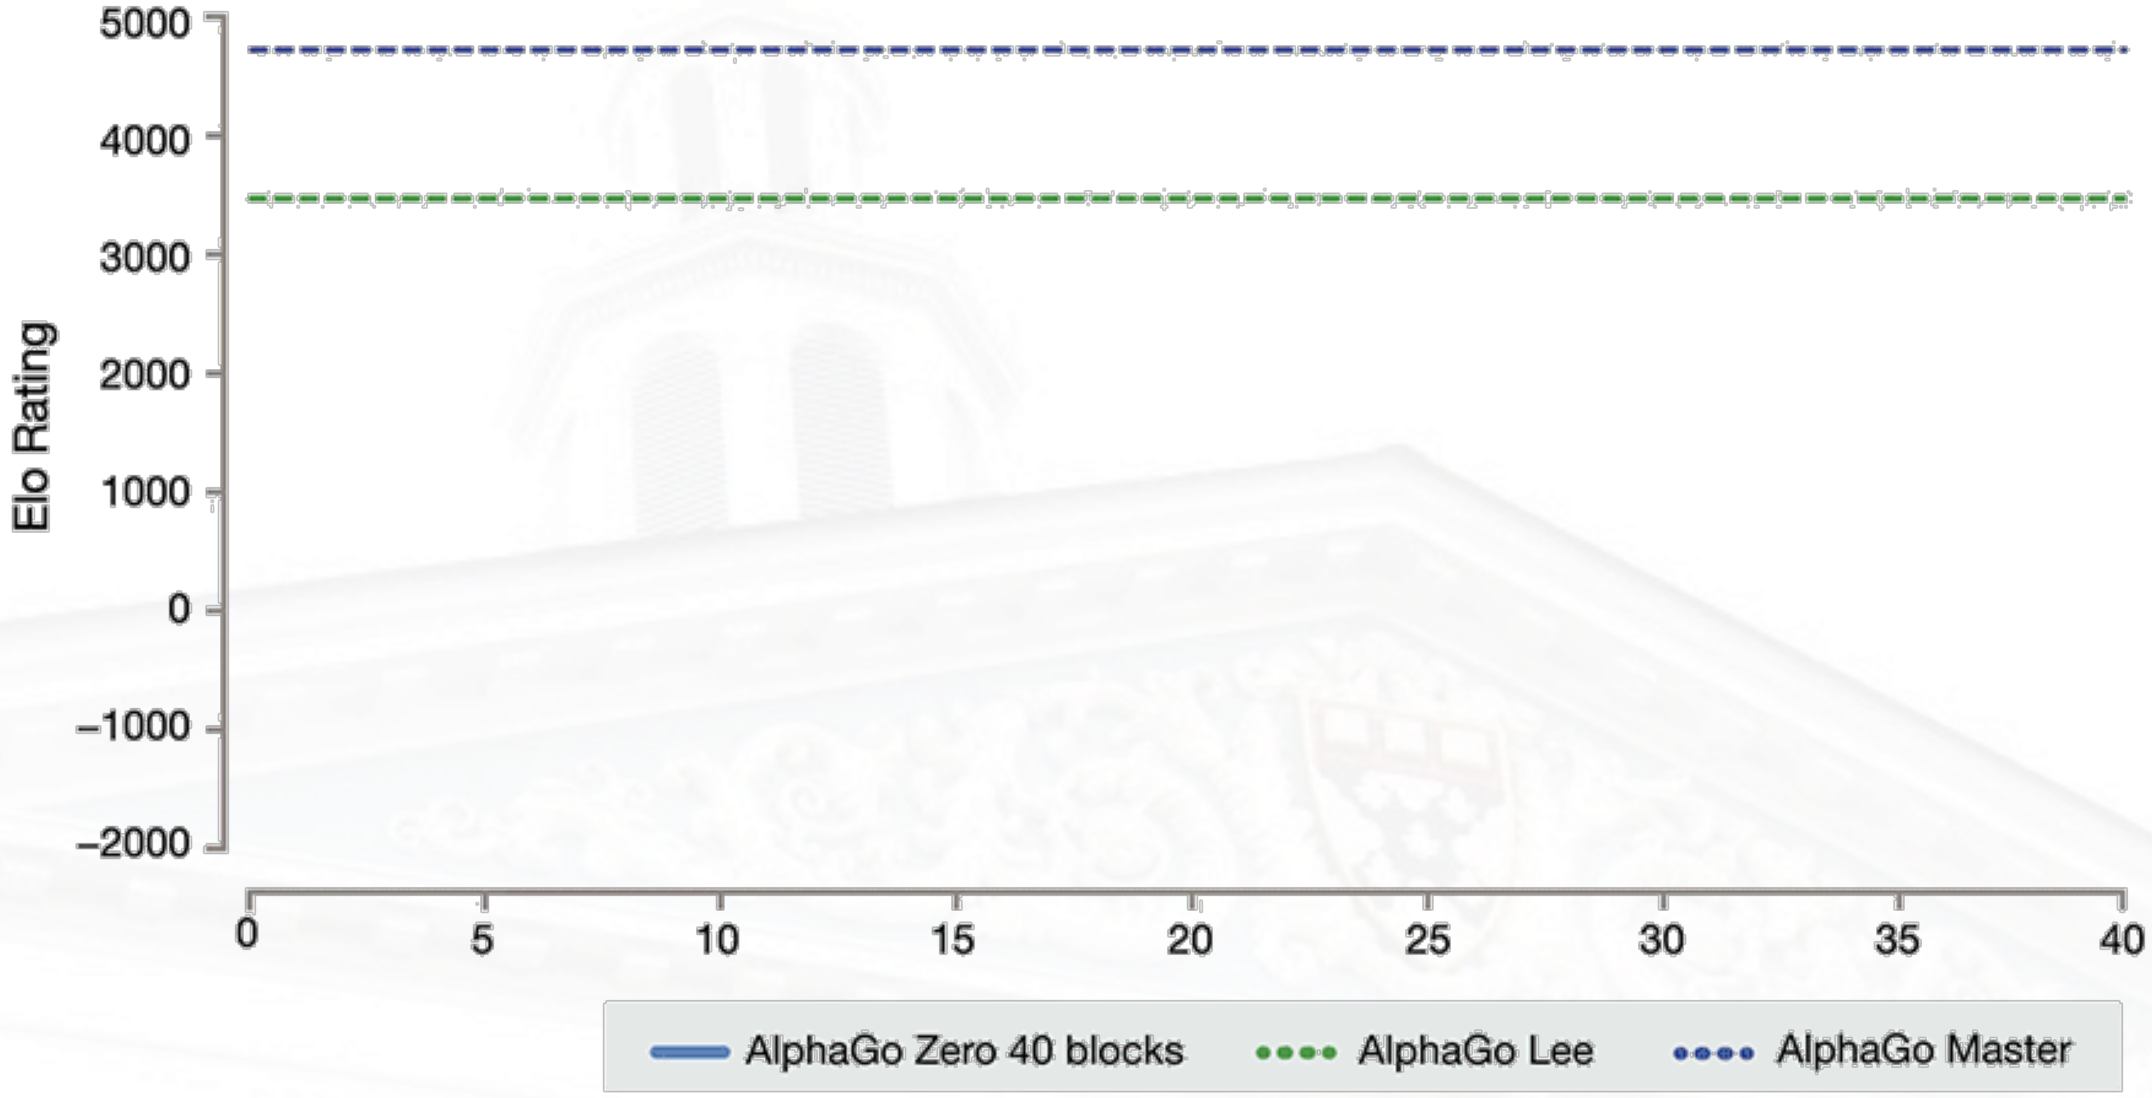
5000
4000
3000
2000
1000
0
–1000
–2000
Elo Rating
0
5
10
15
20
25
30
35
40
AlphaGo Zero 40 blocks
AlphaGo Lee
AlphaGo Master

A.I. Landscape: China vs. U.S.

	China	United States
Institutional Norm	Large volume of data via proprietary systems, yet to focus on building innovation capacity. China tends to be averse to adopting the standard metrics structures used by most multinationals. Chinese A.I. initiatives are good at developing facial recognition as well as tools for surveillance and tracking.	A.I. ecosystem with unified standards and cross-platform sharing. More conducive to developing international standards for what is acceptable for law enforcement use of big data and A.I., and how they will be held accountable for abuse. Developing A.I. tools for surveillance and tracking remains a sensitive topic in U.S.
Regulatory Environment	Tends to set regulations *after* product commercialization. Pursues a strategy of "military-civil fusion" in A.I., as China devises a range of policy mechanisms to incentivize industry cooperation. A looser approach to digital regulations means that companies have more freedom to experiment.	Tends to set regulations *before* product goes to market. The White House has so far been characterized as "missing-in-action" in terms of policymaking for A.I.
Industry Structure	592 A.I. companies (23% of global). Came second in the total number of A.I. enterprises in the world in 2017, and contributed 48% of the world's total A.I. startup funding. A.I. Potential Index = 17.	Ranked first with 1,078 A.I. firms (42% of global). Provided 38% of the funding for A.I. startups globally in 2017. A.I. Potential Index = 33.

A.I. Landscape: China vs. U.S. (Cont'd)

	China	United States
Institutional Norm		
	• China has a **large volume of data** via proprietary systems. • China has yet to focus on building its innovation capacity. But the nation has been supporting different research and workforce development, and reportedly aiming for international collaboration and expansion. • In China, "outside companies do not plug in," but become part of the business as one of hundreds of players in an ecosystem. • China tends to be **averse to adopting the standard metrics structures used by most multinationals**; local suppliers, distributors or customers become partners to help them achieve success in an uncertain business environment. • With regarding to developing software and hardware in A.I., China looks at it from a marathon perspective. • Chinese A.I. initiatives are good at developing **facial recognition as well as tools for surveillance and tracking**.	• The U.S. is in process of creating a **data-friendly ecosystem with unified standards and cross-platform sharing**. • The U.S is producing **more influential A.I. research**, with a **more robust ecosystem** nurturing more competitive A.I. startups. • Companies in U.S. tend create platforms which external parties either plug into or put to use directly. • The U.S. believes it is essential to develop **international standards** for what is acceptable for law enforcement use of big data and A.I., and how they will be held accountable for abuse. • The U.S. is driving A.I. innovation across the spectrum in both software and hardware, with more early adopters and innovators. • In U.S., companies in A.I. tend to be **averse to developing tools for surveillance and tracking**.

A.I. Landscape: China vs. U.S. (Cont'd)

	China	United States
Regulatory Environment		
	• China can be the leader in introducing new regulations for the A.I. industry in the world, suggested by Jeffrey Ding, Macrostrategy Researcher at Future of Humanity Institute in Oxford University. • The Chinese government **sets regulations after product commercialization in the market**. • China pursues a strategy of "**military-civil fusion**" in A.I., as it wields a range of policy mechanisms to incentivize industry cooperation. • The Chinese government is willing to give private entrepreneurs the opportunity to test ideas, e.g. creating policy frameworks, providing subsidies and setting preferential policies to help them. • A **looser approach** to digital regulations means that **companies can experiment more freely**.	• The White House has so far been characterized as "**missing-in-action**". • **Regulations are often devised before the product goes to the market**. • U.S. companies with the best A.I. technology are often considerably **less willing to invest in national security applications**. • In 2017, the U.S. government drafted the first policy to move the U.S. public sector beyond acknowledging the significance of A.I., and toward fully embracing A.I. technologies. • More emphasis placed on **digital regulations**, e.g. tighter cryptocurrency regulations.

A.I. Landscape: China vs. U.S. (Cont'd)

	China	United States
Industry Structure		
	• With **592 A.I. companies** (23%), China came second in the total number of A.I. enterprises among the world in 2017. • Chinese A.I. companies received RMB 63.5 billion (USD 10.1 billion) in funding as of June 2017, and collectively ranked No. 2 in the world in terms of capital raised. Most funds were raised from domestic sources. • China overtook the U.S. in terms of A.I. startup funding, with the former contributing **48% of the world's total A.I. startup funding in 2017**. But in terms of individual deals, China only accounted for 9% of the total. • **Chinese production of semiconductors is only 4% of the total global market share**. • Most Chinese companies tend to only hire Chinese people, focus on the China market, and may lack an international vision. • The "AI Potential Index" of China is 17, almost half of that of U.S., according to an analysis at the University of Oxford.	• The U.S. ranked first with **1,078 A.I. firms**, representing 42% of the total worldwide. • About 50% of global A.I. investments went to U.S. startups, reaching RMB 97.8 billion (USD 15.5 billion) as of June 2017, and leading the world in terms of funding. • In 2017, the U.S. provided 38% of the funding for A.I. startups globally, and **led in both the total number of A.I. startups and total overall funding**. • **50% of semiconductors in the world is produced by the U.S.** • Silicon Valley companies are dominated by a diversified culture. • The "AI Potential Index" of U.S. is 33.

A.I. Implementation Matrix

	Process A	Process B	Process C	Process D	Process E
Function & Nature	Strategic direction exploration	Sales lead management	Quality control	Risk management	Customer service (non-chatbot)
Data Availability					
AI Function & Role	Classification and Sentiment analysis	Facial recognition	Fault identification	Fraud detection and forecasting	Text to image processing
AI Data Type	Unstructured	Structured and clean	Structured and clean	Structured but sparse	Unstructured
AI Model	• Support Vector Machines (SVM) • Clustering • Recurrent Neural Networks (RNN)	• Convolutional Neural Networks (CNN)	• Support Vector Machines (SVM) • Clustering	• K-Nearest Neighbors (kNN) • Recurrent Neural Networks (RNN)	• Generative Adversarial Network (GAN)
AI Readiness					
Schedule	TBD	Q3 2019	Q1 2019	TBD	Q1 2020

Management

A.I. Experts

A.I. Dashboard

ConvnetJS demo: toy 2d classification with 2-layer neural network

The simulation below shows a toy binary problem with a few data points of class 0 (red) and 1 (green). The network is set up as:

```
layer_defs = [];
layer_defs.push({type:'input', out_sx:1, out_sy:1, out_depth:2});
layer_defs.push({type:'fc', num_neurons:6, activation: 'tanh'});
layer_defs.push({type:'fc', num_neurons:2, activation: 'tanh'});
layer_defs.push({type:'softmax', num_classes:2});

net = new convnetjs.Net();
net.makeLayers(layer_defs);

trainer = new convnetjs.SGDTrainer(net, {learning_rate:0.01, momentum:0.1, batch_size:10, l2_decay:0.001});
```

change network

Feel free to change this, the text area above gets eval()'d when you hit the button and the network gets reloaded. Every 10th of a second, all points are fed to the network multiple times through the trainer class to train the network. The resulting predictions of the network are then "painted" under the data points to show you the generalization.

On the right we visualize the transformed representation of all grid points in the original space and the data, for a given layer and only for 2 neurons at a time. The number in the bracket shows the total number of neurons at that level of representation. If the number is more than 2, you will only see the two visualized but you can cycle through all of them with the cycle button.

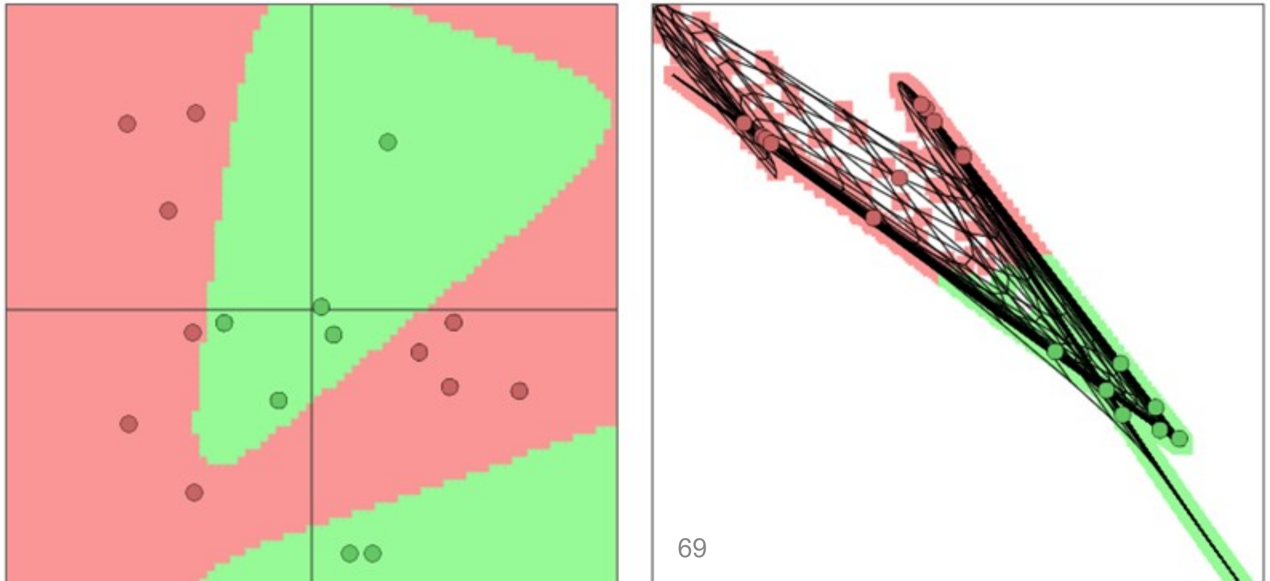

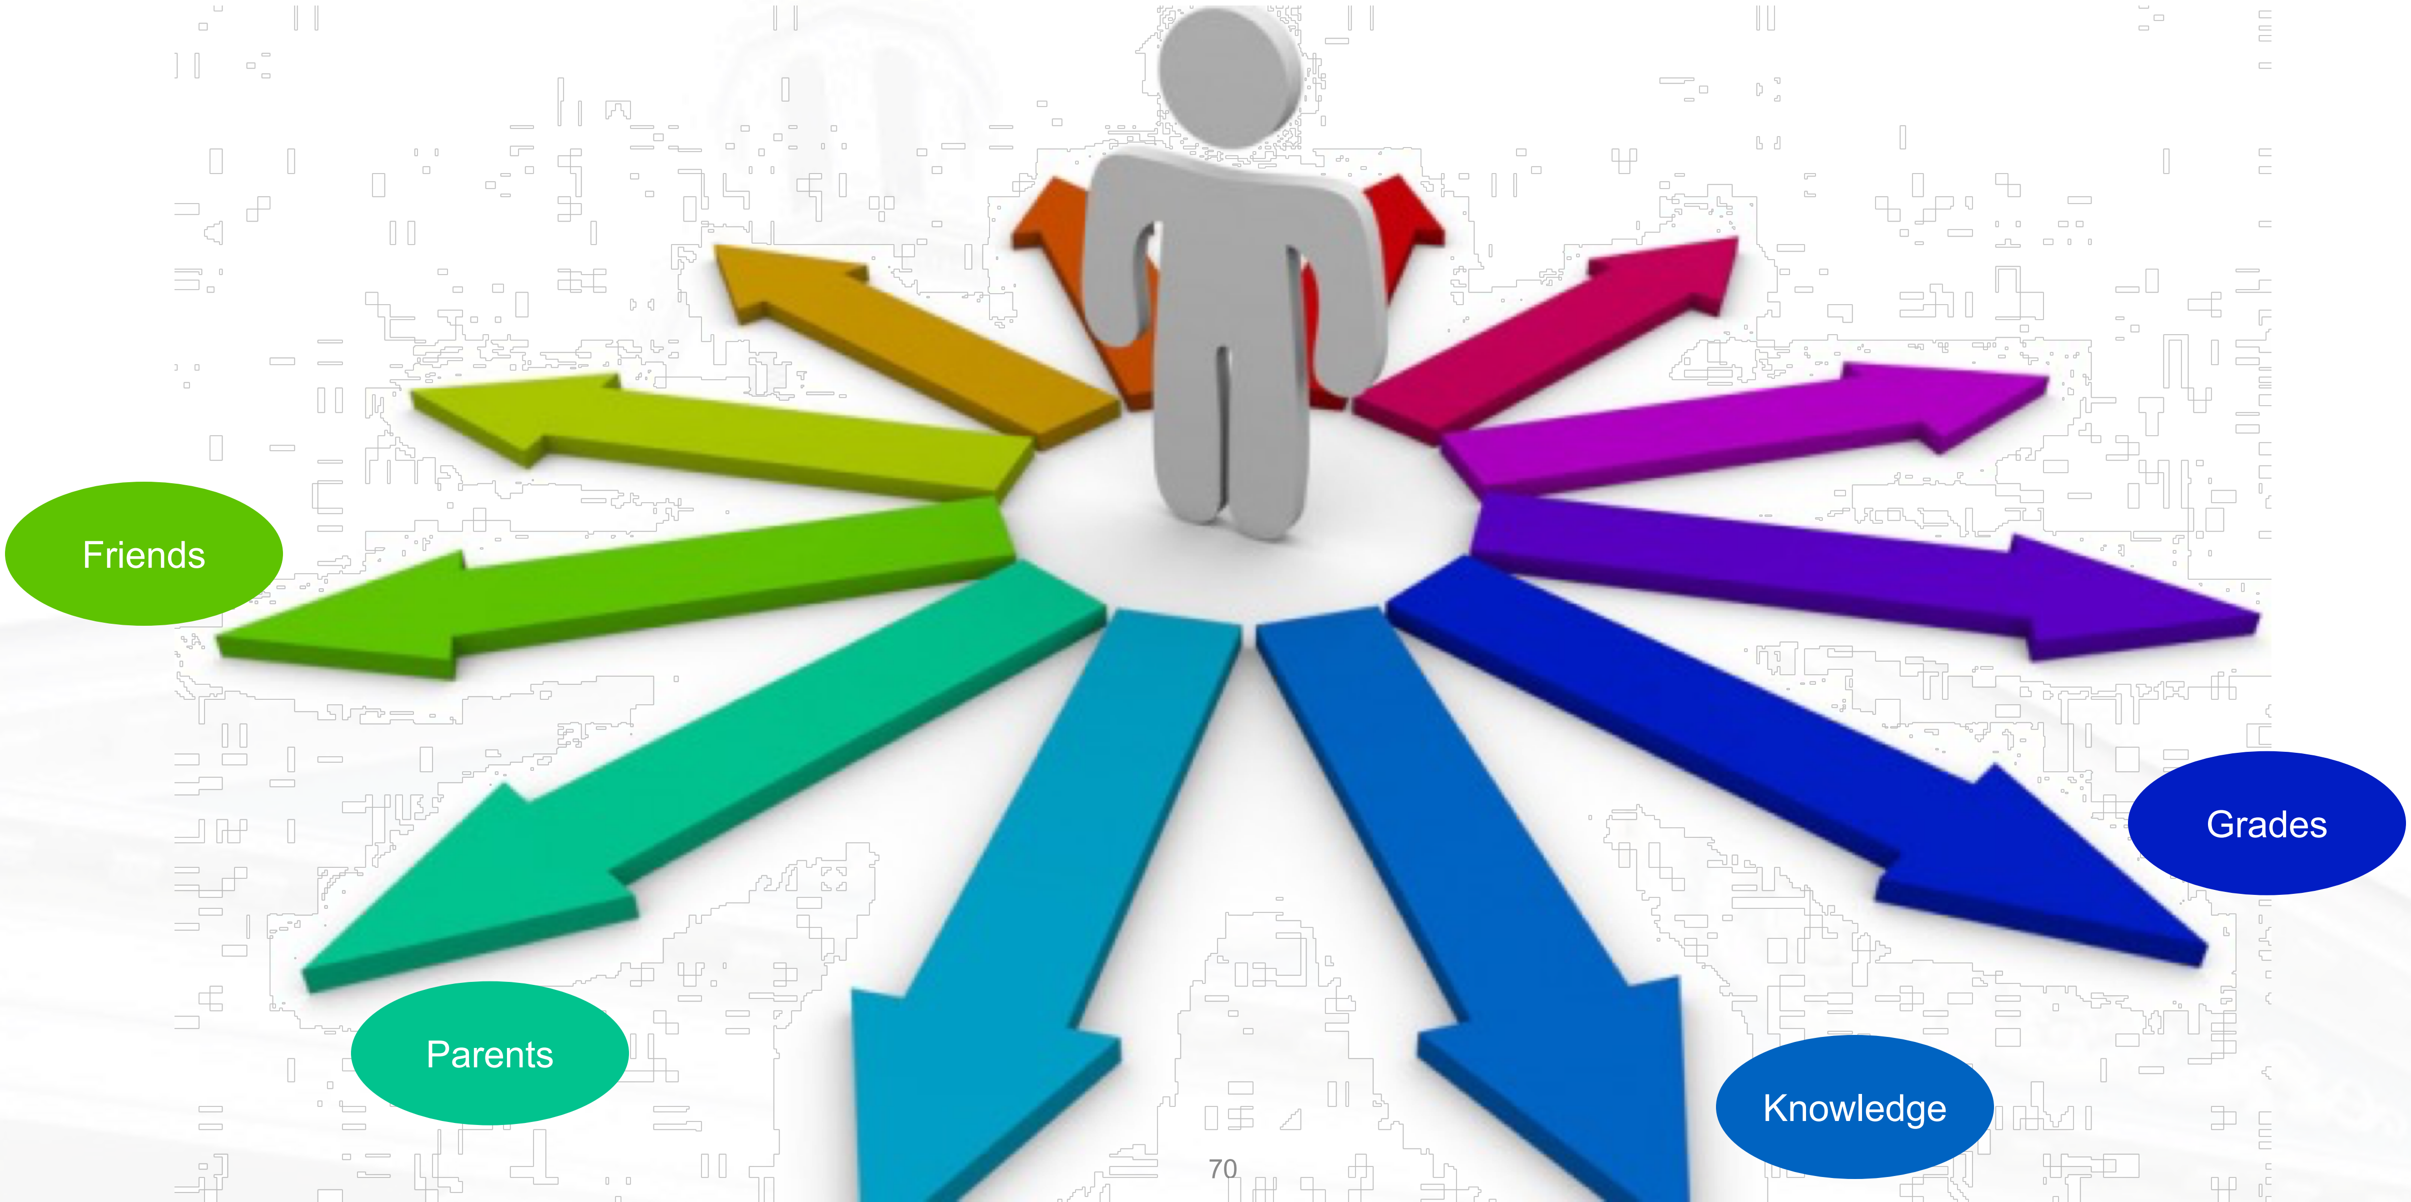
Friends
Parents
Knowledge
Grades

"We're looking for someone with your exact qualifications, but a mechanical version."

Connecting Academia with Reality

Goldman Sachs

YF CAPITAL
云 锋 基 金

sensetime

CogX

8-10
JUN 20

Get Involved · Highlights 2019 · CogX 2020 tickets

The World's Most Exciting Celebration of Innovation and Transformational Opportunities

| Festival of AI, Blockchain and Breakthrough Technologies
Mon 8 - Wed 10 June 2020, King's Cross, London

Last Chance Super Early Bird CogX 2020 Tickets

Robotic Process Automation
机器人与工序自动化
Full-stack AI-enabled Robotic Process Automation (RPA) solution suite for business-critical tasks spanning multiple applications and data silos

Quantum Computing
量子计算
Solution provider and platform developer for quantum and classical computing for predictive analytics, forecasting, and optimization

Computer Vision
计算机视觉
Developer of deep learning technology-based computer vision solutions aimed at a broad range of consumer and enterprise applications

Knowledge Graph
知识图表
Dynamically evolving knowledge graphs that provides inference strength across concepts, events and themes derived from a wide variety of information services

Real-Time Robotics Automation
实时机器人自动化
Deep reinforcement learning-based AI software platform that enables enhanced perception, reaction and control in real-time robotics environments

AI Chips
人工智能芯片
Deep reinforcement learning-based AI software platform that enables enhanced perception, reaction and control in real-time robotics environments

Cybersecurity
网络安全
Advanced deep learning technology-based cybersecurity products and solutions for threat detection and prevention

Music Augmentation
音乐强化
Developer of a music augmentation technology that transforms linear music to dynamically personalized music for consumers, ad-agencies, music labels, and producers

Course Assessments & Objectives

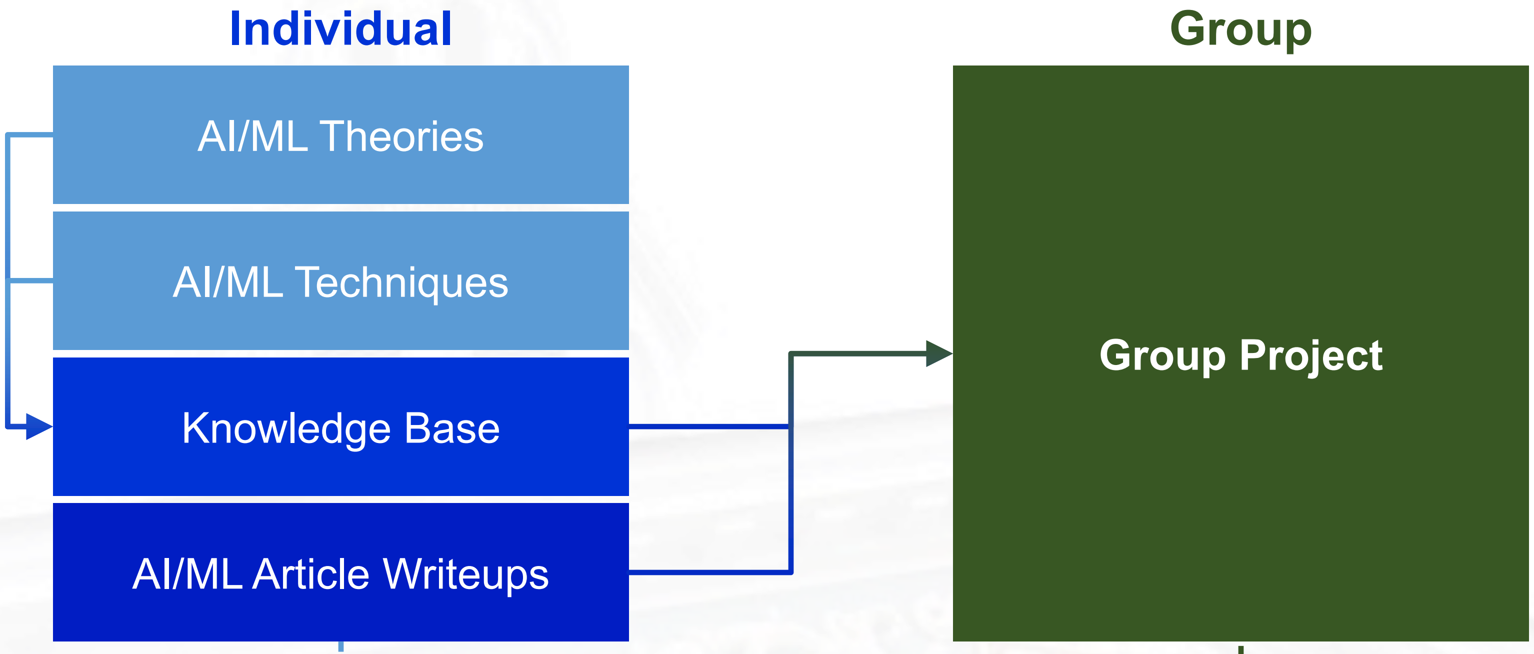

Jan 2019 – Present

MOSAIC FINANCIAL TECHNOLOGY LTD. *(AI/ML startup)* **HONG KONG, CHINA**

Co-Founder and Chief Technology Officer

- Compiled pitch deck and presented Mosaic's vision and underlying technology to potential investors and AI specialists at HKUST. Created a proof-of-concept (PoC) and demonstrated a prototype
- Devised strategy to drive adoption and compiled reports on the applications of AI/ML techniques. Elaborated on the plans for future product launches. Established Advisory Board with industry experts

Assessments

Deliverables	Details	Weight	Due Date	Mode	Rationale
Presentation 10 core slides	– Project presentation with 10 core slides and supplementary materials in appendix – Selected teams to present to a panel of judges	40%	May 31	Individual or group	CV building and interview preparation
Writeup of recent AI journal articles No more than 2 pages in total	– Review of 1-2 AI journal articles, including a critical analysis and potential applications	20%	Last day of each month (Mar, Apr and May)	Individual	Industry knowledge
Extra credit activities and assignments	– 5-min presentation of journal-article analysis in-person or via Zoom – Other assignments and activities TBA	20%	Various	Individual	Rewarding those who are proactive
Participation	– Class attendance with 10% base and 10% based on contribution	20%	Various	Individual	Encourage active contributing in class

Typical Class Flow

15 mins	1 hour	15 mins	1 hour	15-20 mins
Review	**Theory & Technical Details**	**Class Break**	**Applications & Case Study**	**Q&A** **CV Clinic**

Speakers & Judges (Tentative)

Mr. Antoine Blondeau

Managing Partner
Alpha Intelligence Capital

Founder
Sentient Technologies

Investor
SpaceX, Dianrong, PeerIQ

Advisory Board Member
Zeroth.ai

Dr. Gregg Li

Founder & Principal
Sinoalpha Ventures

Adjunct Professor
University of Hong Kong

Former Principal Consultant
PricewaterhouseCoopers

Independent Non-Exec. Director
Cyberport

Honorary President
Invotech

Mr. Jacob Wai

Chief Data & Risk Officer
MoneySQ

Visiting Lecturer
Hong Kong Polytechnic University

Corporate Trainer & Lecturer
Hong Kong Institute of Bankers

Chairman of the Financial Technology SIG
Hong Kong Computer Society

Mr. Leo Tong

FinTech Adviser & Compliance Officer
Private Equity Funds & Family Offices

Certified M&A Dealmaker

Certified Information Systems Security Professional

Judge & Chief Assessor
Hong Kong ICT Awards

Co-Chairperson of FinTech SIG
Hong Kong Computer Society

Speakers & Judges (Tentative)

Mr. Ian Huang

Chairman
Co-operatives of Innovative Intellectuals

Former Visiting Chief Architect
Singapore National Science & Technology Board (NSTB, now called A*STAR)

Fellow
Hong Kong Computer Society

Fellow
Hong Kong Institute of Directors

Ir. Dr. Daniel Ng

Director

Engineer

Controller

Visiting Lecturer and Examiner

Forensics and Business Intelligence Machine Learning Researcher

Mr. Dominic Wu

Managing Director & Senior Risk Manager
BNY Mellon Asia

President
Asia Financial Risk Think Tank

Mr. Philip Leung

Founder
Small Talks Circles

Founder
Big Data & AI Startup

Honorary Career Advisor
Centre of Development and Resources for Students (CEDARS)

Speakers & Judges (Tentative)

Mr. Allan Lee

Director
Training and Master Family Legacy Planner of the Legacy Academy

Accredited Administrator
MBTI

Registered Corporate Coach

Mr. Winston Lam

Chairman
Invotech

Chairman
Berkeley Club of Hong Kong

Member
Entrepreneurship Committee Advisory Group
Cyberport

Mr. Jeffrey Hui

Chairperson
Hong Kong Institute of Marketing

Managing Director
InnoSights

Adjunct Professor
Chinese University of Hong Kong

Mr. Justin Lao

Founder
K.U.G. Education Technology

Member
Chinese People's Political Consultative Conference
Chongming District
Shanghai

CO-CREATION

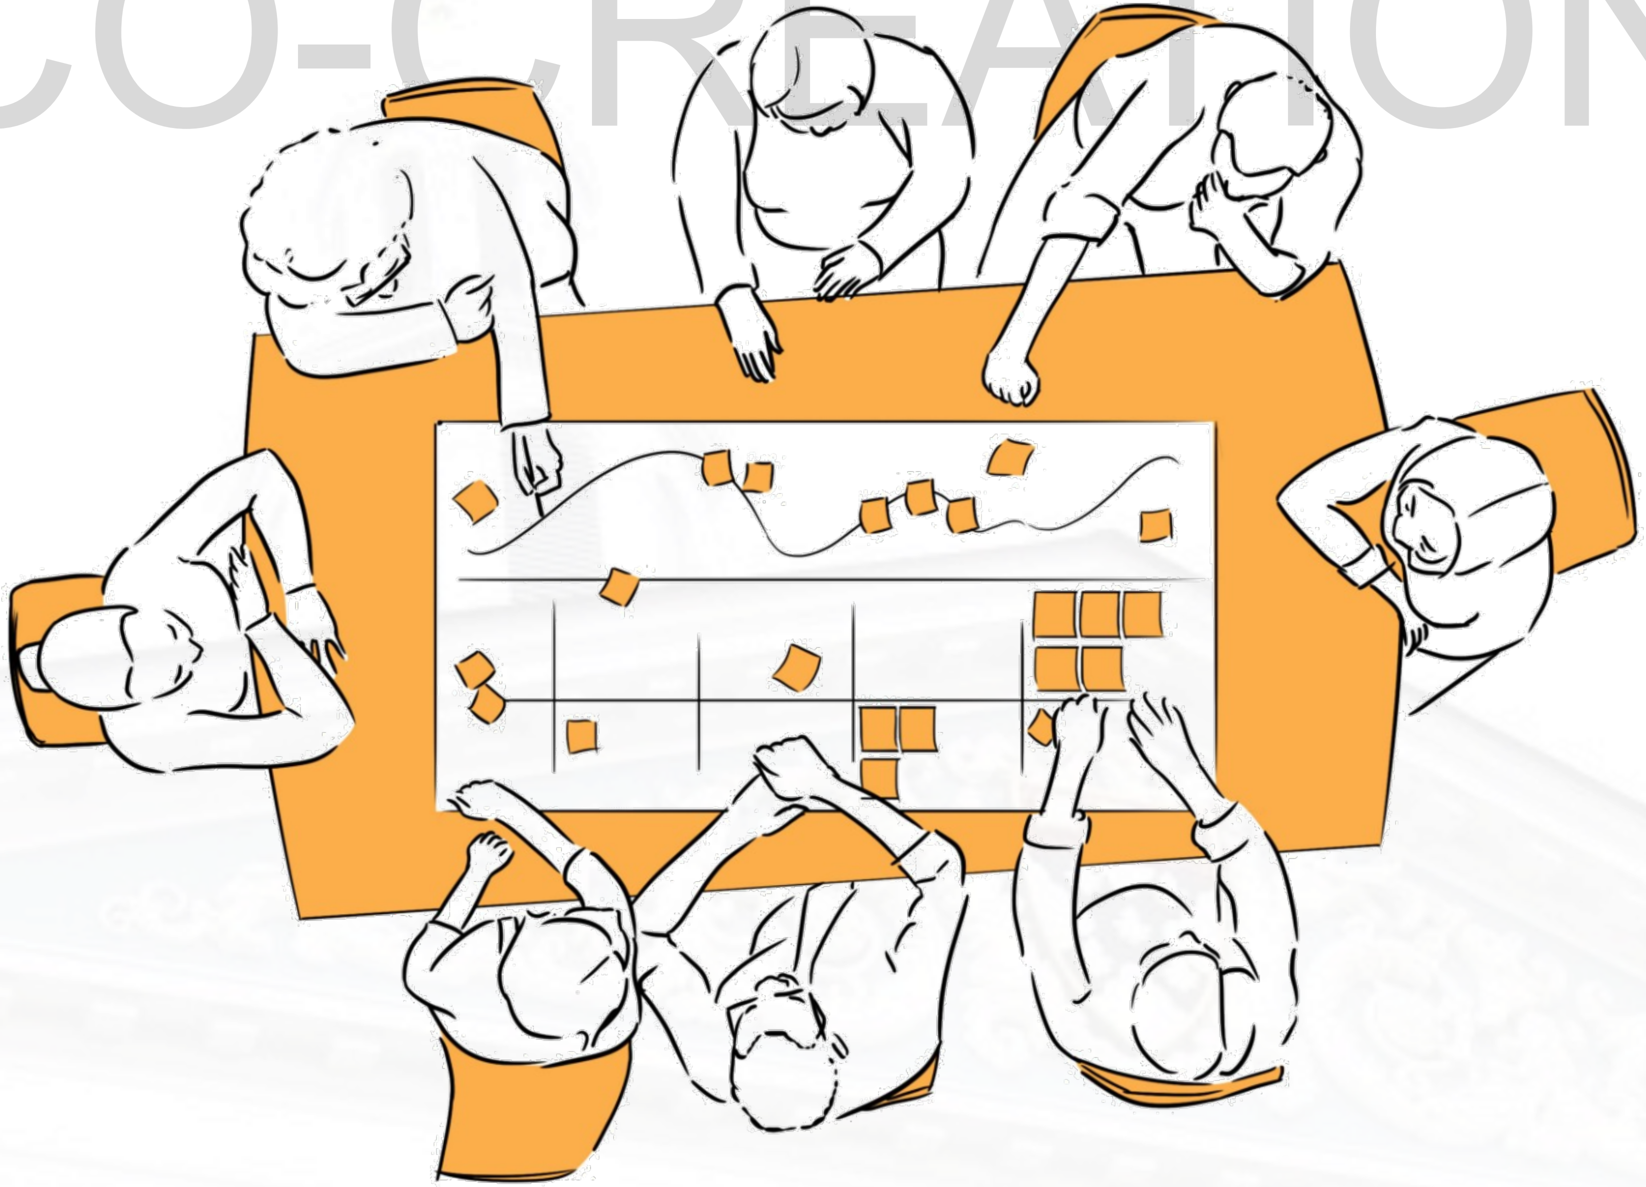

HireVue Overview

This will be a new experience for many applicants because top banks like Goldman Sachs and JP Morgan have recently begun to use HireVue. It apparently adds 13% more top performers above the client’s industry.

The first interview is provided by HireVue; however, it will not completely replace the more traditional, intensive recruiting process. If the first interview is successful, then **a representative from the bank will contact the candidate for a second interview**. From then on, any upcoming interviews will be part of the regular interview rounds, involving live interactions with analysts, associates and even VPs and MDs.

You will be given about **20-30 seconds for each question to think of a response**. After that, you’ll have about **3 minutes to record your answer**. The amount of time given really depends on the questions. For instance, a question about why you would be the best candidate for the role will require a much longer and thoughtful response than answering a question about what your overall GPA is.

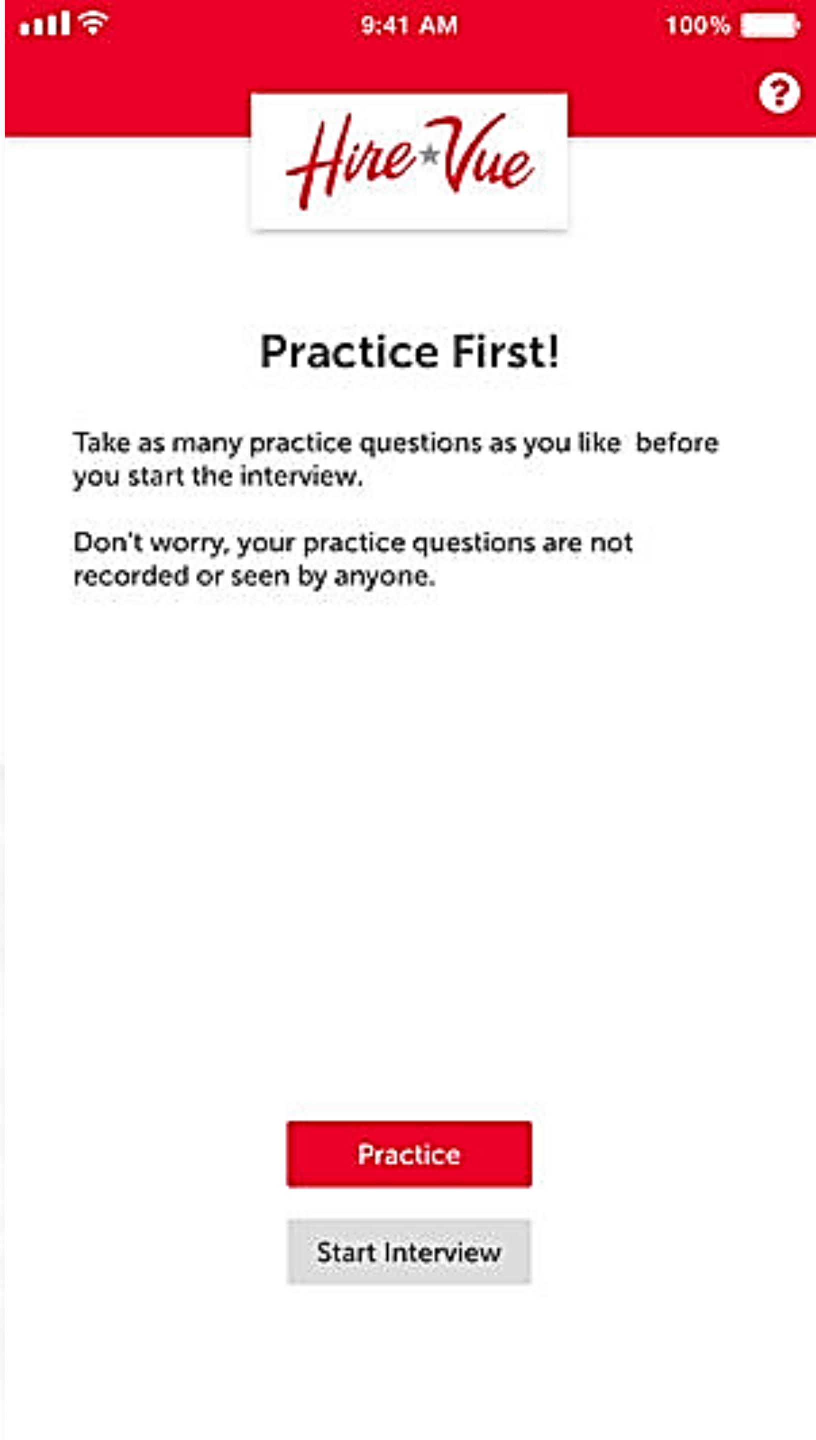

IBM Watson | HILTON WORLDWIDE | Unilever | rackspace | Keurig Dr Pepper

URBAN OUTFITTERS | ATLANTA PUBLIC SCHOOLS Making A Difference | Carnival | BOSTON RED SOX | UNDER ARMOUR

vodafone | SHIPT | Sonic Automotive | maxis | blackbaud

dunkin' brands | CATHAY PACIFIC | Children's Healthcare of Atlanta | ORACLE | HBO

HireVue Iris™ Deep Learning Analytics Engine

15,000
PREDICTIVE ATTRIBUTES

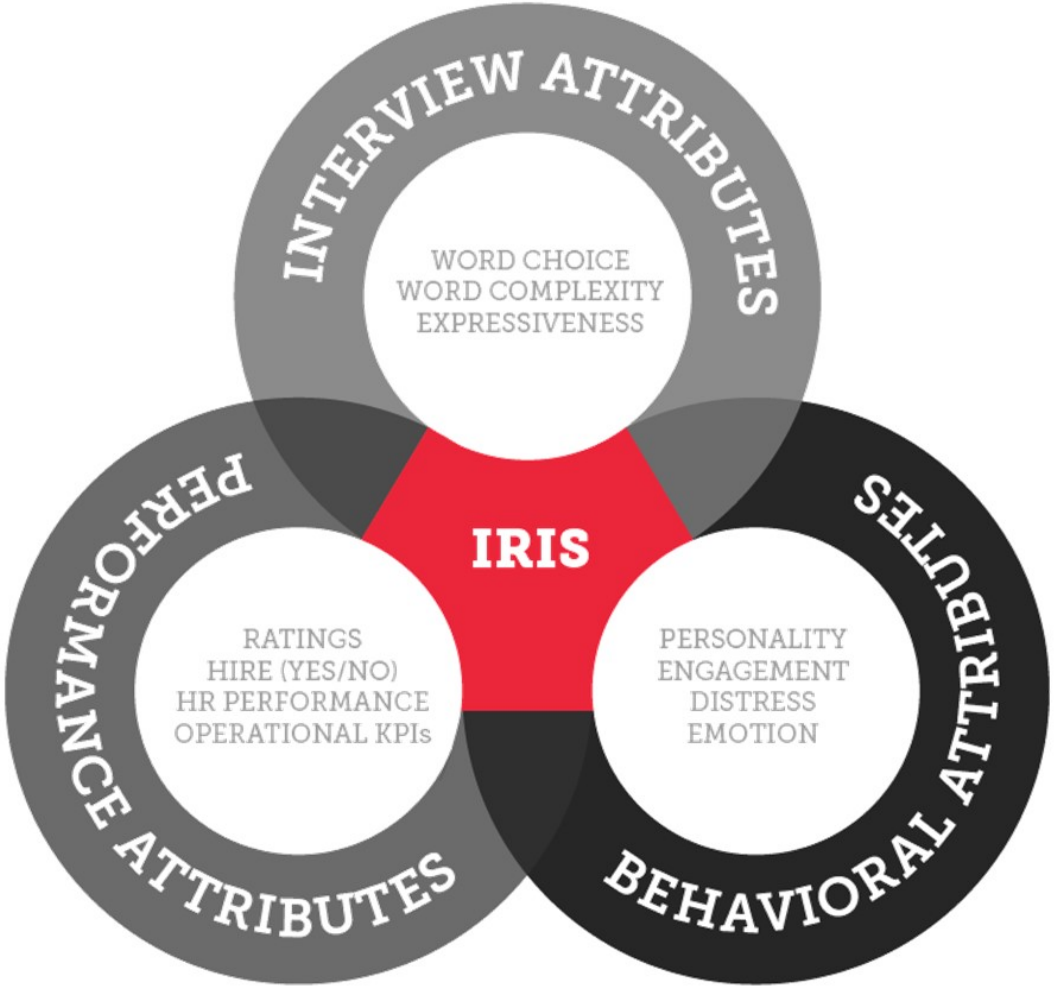

100,000X
MORE DATA THAN A RESUME

Predictive Power of a Traditional Assessment

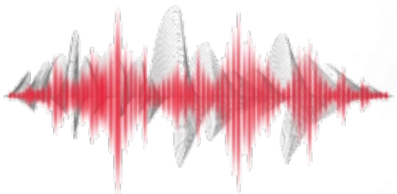

150 words/minute * 20 minutes

+

50 features/sample * 10 samples/sec * 60 seconds * 20 minutes

+

30 features/frame * 24 frames/sec * 60 seconds * 20 minutes

= 1.4 M Features

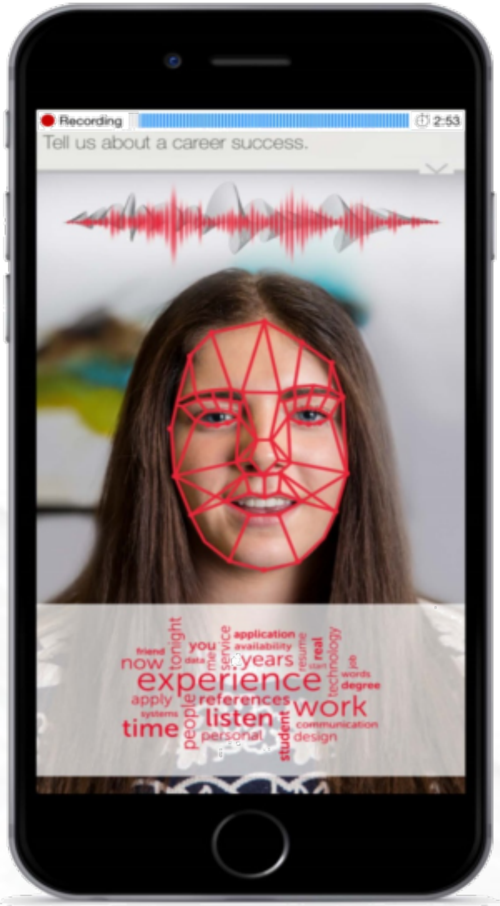

>4000x More Features Available than a 300 Question Assessment

Underlying Mechanics

HireVue Iris, a patented deep learning analytics engine that powers HireVue Insights, analyzes a unique data set of interactions, feedback and outcomes that never before existed. Developed by HireVue's data science team, Iris was built based on over **3 million interview responses**. Each candidate interview contains 100,000 times more bytes of data than the resume or profile traditionally used for identifying job candidates. The platform examines attributes in three major categories: **interview attributes**, **behavioral attributes**, and **performance attributes**. Iris's proprietary algorithms discover patterns and learn which attributes predict performance, then **scores each candidate on how they compare to existing top performers**.

Video interview, the recruiter can see the personality, drive, and work ethic of a candidate. Furthermore, this **attracts and captures more candidates from more schools**, because the talent isn't just found from target schools such as those belonging to Ivy League group.

HireVue Assessments evaluates tens of thousands of data points, studying both **verbal and nonverbal content**, including:

- **Word choice and vocabulary**
- **Intonation**
- **Inflection**
- **Facial expressions**

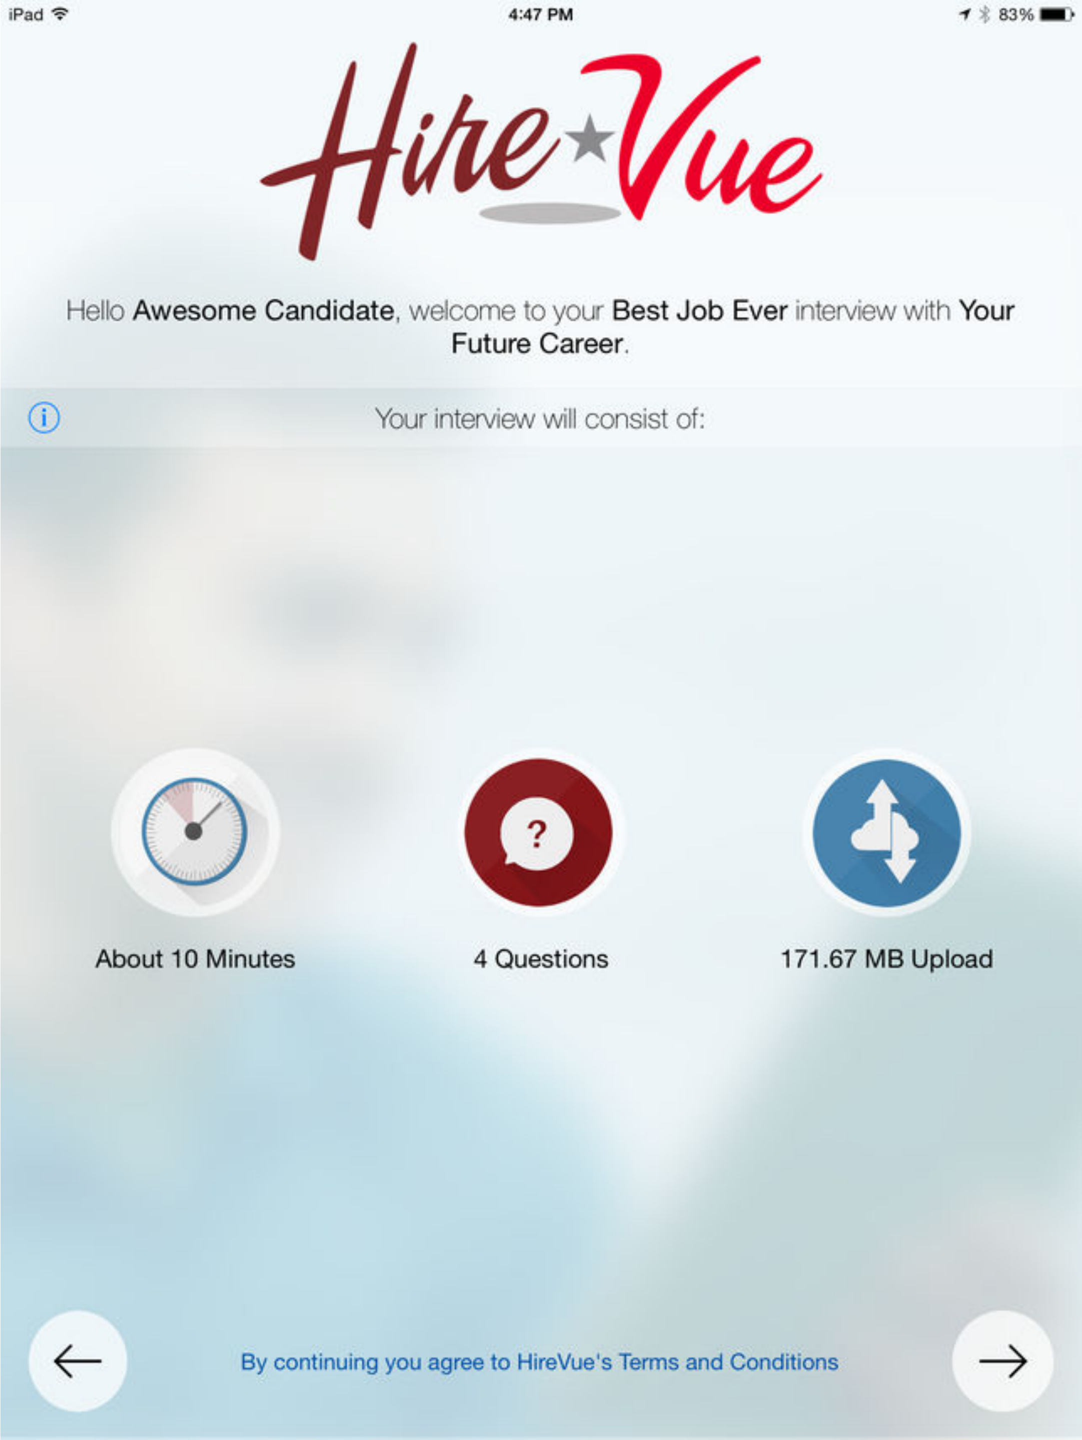
iPad
4:47 PM
83%
HireVue
Hello Awesome Candidate, welcome to your Best Job Ever interview with Your Future Career.
Your interview will consist of:
About 10 Minutes
4 Questions
171.67 MB Upload
By continuing you agree to HireVue's Terms and Conditions

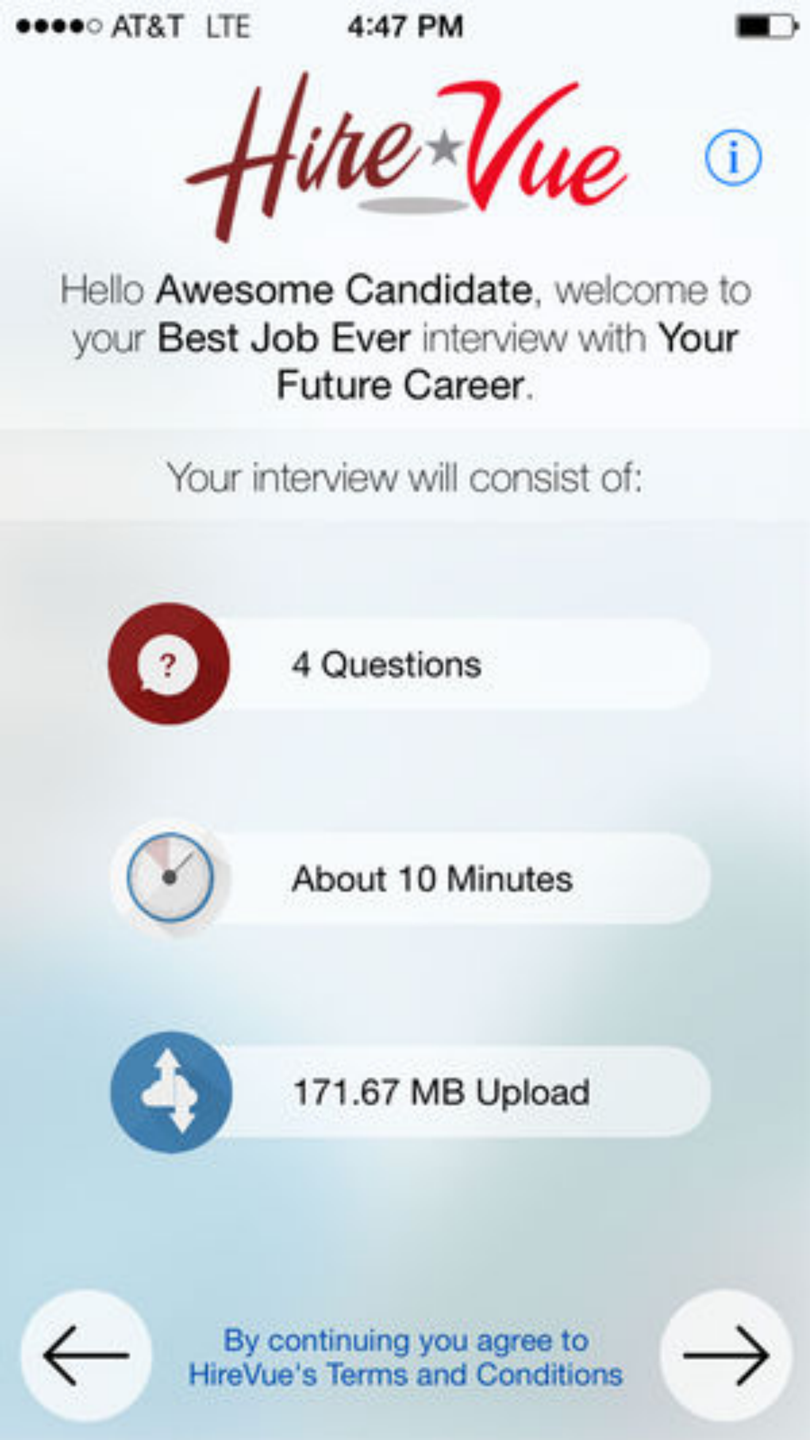
AT&T LTE
4:47 PM
HireVue
Hello Awesome Candidate, welcome to your Best Job Ever interview with Your Future Career.
Your interview will consist of:
4 Questions
About 10 Minutes
171.67 MB Upload
By continuing you agree to HireVue's Terms and Conditions

Deep Learning on Bespoke Video Assessments

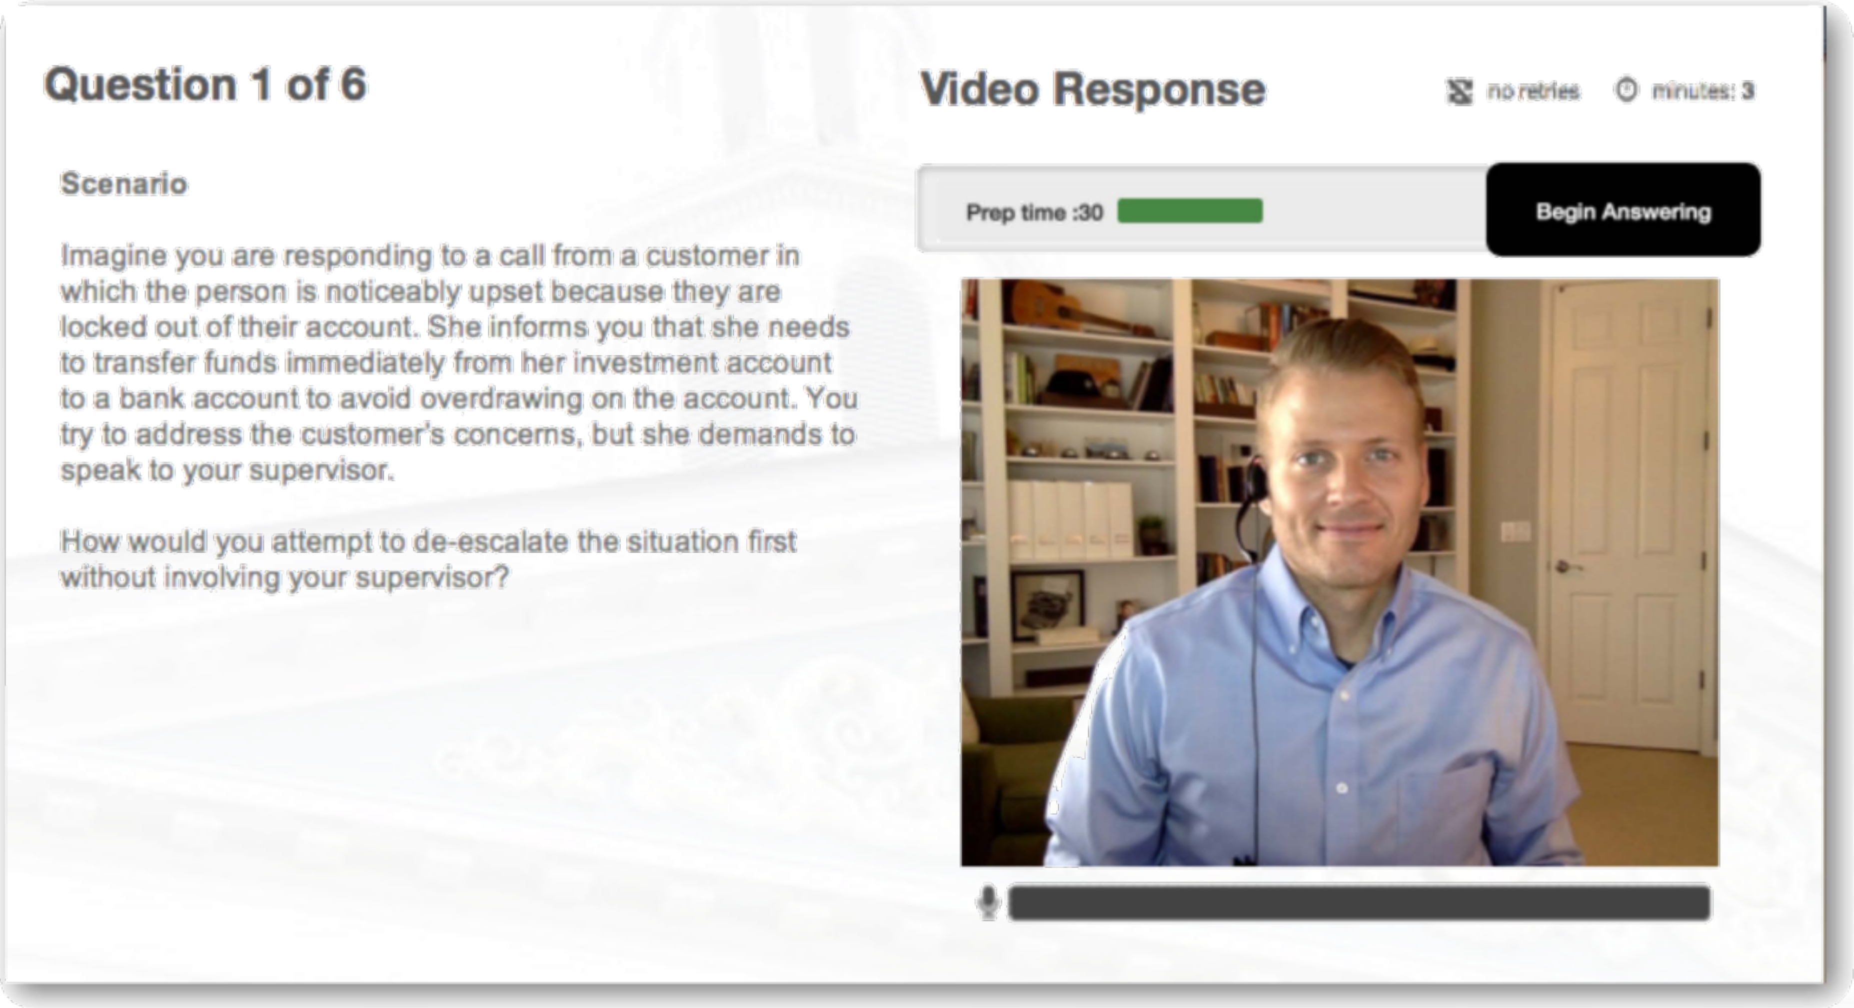

TEXT

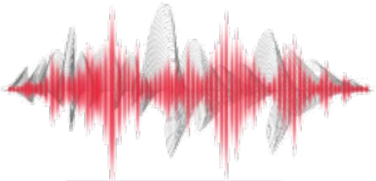

AUDIO

VIDEO

Sample Questions

- Why are you applying for this position?
- How did you deal with a difficult co-worker in your previous work experiences?
- How did you handle a drastic change in role to achieve a goal?
- Why do you think you are the right candidate for this position?
- What current events are you following at the moment? Why are they interesting?
- Do you have an expertise or unique experience that can benefit our team?
- What is the most important leadership experience you have? And why?
- What efforts do you make to keep abreast of financial markets and business news?
- Why do you want to work for our company?
- What relevant skills have you gained from your past work or internship that are easily transferable and directly beneficial to the new role you're applying for?
- Why are you looking for a new role in our company?
- What were your top responsibilities at your current/previous position?
- What are your three main weaknesses?
- How will you use your background and skills to succeed in his role

ICAR Catalogue

Version 1.0
06 I 17

This catalogue will list all ICAR measures that have been made available to qualified users on the ICAR website. Information about the item types, including both brief introduction and sample items, will be regularly updated.

1. Three-Dimensional Rotation

The 24 Three-Dimensional Rotation items present participants with cube renderings and ask participants to identify which of the response choices is a possible rotation of the target stimuli.

Sample item:

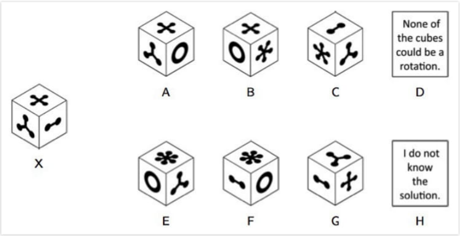

2. Letter and Number Series

The 9 Letter and Number Series items prompt participants with short digit or letter sequences and ask them to identify the next position in the sequence from among six choices.

Sample item:
In the following alphanumeric series, what letter comes next? I J L O S
(1) T (2) U (3) V (4) X (5) Y (6) Z

3. Matrix Reasoning

The 11 Matrix Reasoning items contain stimuli that are similar to those used in Raven's Progressive Matrices. The stimuli are 3x3 arrays of geometric shapes with one of the nine shapes missing. Participants are instructed to identify which of six geometric shapes presented as response choices will best complete the stimuli.

Sample item:

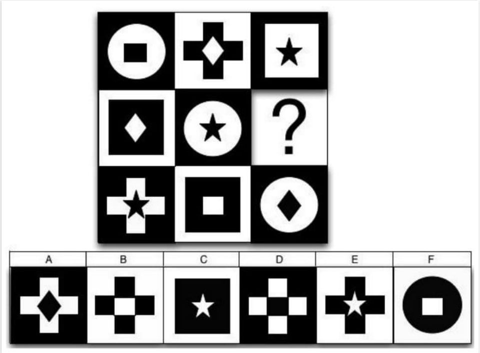

4. Verbal Reasoning

The 16 Verbal Reasoning items include a variety of logic, vocabulary and general knowledge questions.

Sample item:
IF the day after tomorrow is two days before Thursday, then what day is it today?

(1) Friday (2) Monday (3) Wednesday (4) Saturday (5) Tuesday (6) Sunday

Types 1-4 taken from:

Condon, D. M. and Revelle, W. (2014). The International Cognitive Ability Resource: Development and initial validation of a public domain measure Intelligence, 43, 52–64.

5. Progressive Matrices

27 items have been designed based on similar rules follow
Raven's Progressive Matrices. Each item has 8 distractors.

Sample item:

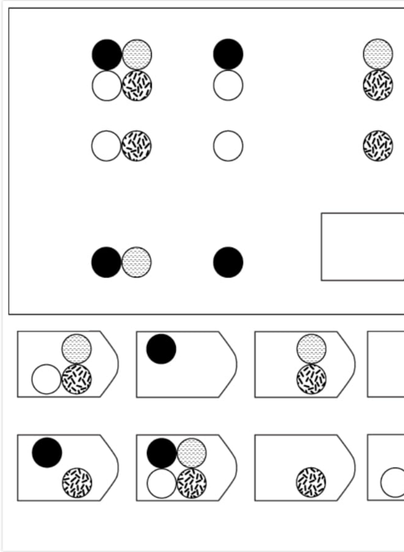

6. Face Detection

The newly developed Mooney-Verhallen Test (Verhaller
2015) comprises 140 items, divided into four sections of e
(35 trials each), whereby each section is preceded by one p
(all four practice trials are unique). The presentation of th
is randomised for each participant, and the performance
interest is the number of trials on which participants corre
either of the eyes of the face (not just on the correct image

Take a sample Wechsler test and see where you fall on the Wechsler Adult Intelligence Scale!

Quick Wechsler Test

25 questions in 6 minutes

START

Full Wechsler Test

50 questions in 12 minutes

START

Wechsler Intelligence Scale For Children Vs. Adults

David Wechsler (1896 - 1981) was an American psychologist whose work frequently specialized in intelligence testing. He developed two intelligence scales - the Wechsler Adult Intelligence Scale (WAIS) and the Wechsler Intelligence Scale for Children (WISC) - that still remained use, although in updated versions. He was also an influential theorist whose research regarding the human intellect remains important and relevant, especially in the later development of theory of multiple intelligences by Howard Gardner in the 1980's.

Wechsler Intelligence Scale for Children (WISC)

Best Practices

Logistics
- 30 secs to prepare; 3 mins to answer; do not have to use all 3 mins; a minimum of 10 secs
- Speak slowly with clear pronunciation and polished diction
- Clean camera lens and keep camera at eye level
- Notepad to minimize rambling
- Resume and job description in front of you
- Try not to prop up your phone
- Phone fully charged and strong Wifi (350 Kbps per second)
- Good background, clean desk and computer, bookshelf
- Quiet locale (e.g. carpet and "things" to minimize echo)

Interview Preparation
- Prepare questions in advance
- Fully utilize sample interview before the real one
- On-demand video interviews are always structured interviews, i.e. same questions, in the same order as everyone else for the position

Light
- Near window for natural light
- No light source from behind or below

Dress Code
- Dress up (i.e. professional business attire)

Demeanor
- Practice to avoid awkwardness and self-consciousness
- Show enthusiasm and maintain eye contact, enthusiasm (i.e. no downcast)
- Keep video feed near camera
- Stand up and "present" if you can
- Not evaluated based on stage presence

Examples
- Toastmasters International, Enactus, TED

Understanding Features of Candidate Scores

Low Scoring Tier Candidates

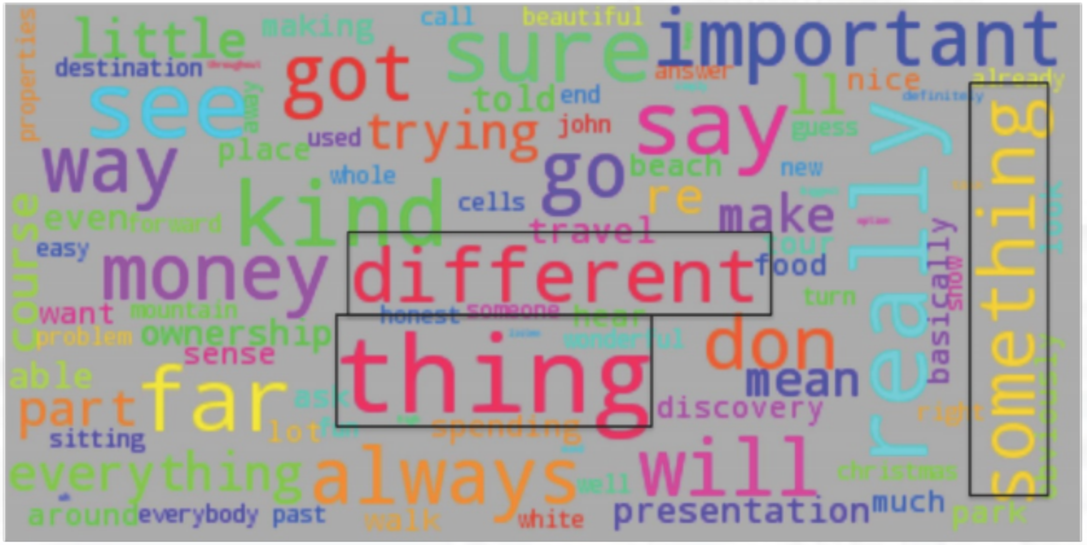

Top Scoring Tier Candidates

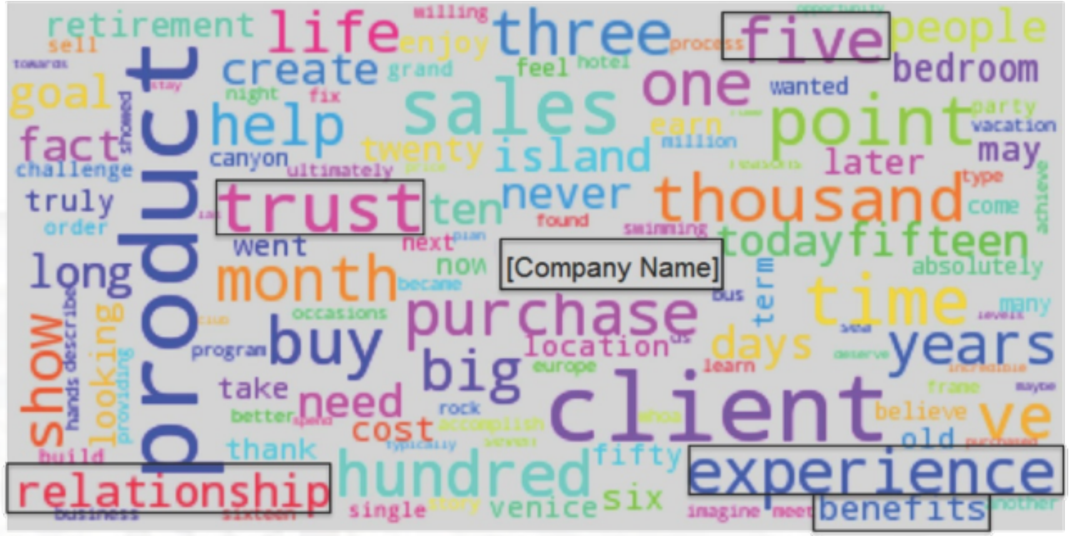

Don’t Let Artificial Intelligence Pick Your Employees

A Stanford GSB scholar shares why algorithms aren’t sophisticated enough to make these strategic decisions … yet.

February 8, 2019 | by Dylan Walsh

Editor's Picks

A Different Kind of Funding Bias

Reshaping Markets to Solve Poverty and Inequality

How Politicians Hurt Your Public Pension Plan

Related

WE'RE HIRING!

Alpha Intelligence Capital

Key Responsibilities

Students to form groups of five (5) to complete the following:

- Develop an understanding of the AI/ML industry landscape and relevant technologies
- Creation of a conference chatbot application to be deployed in June 2020

Deliverables

- Knowledge assessment of NLP techniques and related applications via in-person interviews
- Monthly progress reports and regular updates
- Final group presentation on-campus or Central

Position Requirements

- Undergraduate or postgraduate students, preferably majoring in Artificial Intelligence, Mathematics, Statistics, Computer Science, Business Analytics, Finance, and/or Economics

Duration: Spring and summer 2020
Format: Groups of five (5) with multidisciplinary (and complementary) skillsets
Opportunity to convert to full-time based on performance

Anthony Woo
CFA CAIA FRM
Associate Director
Alpha Intelligence Capital
aw@aicapital.ai

WE'RE HIRING!

Alpha Intelligence Capital

Key Responsibilities

- Keen interest and experience in venture capital transactions and financial analysis
- Familiar with transaction legal documents and investment structure
- Understanding of the commercialization of technologies from an investor perspective
- Conduct due diligence for potential investments in the AI space
- Assist in the preparation of deal-related materials (e.g. data collection, industry research, comparable company and valuation analysis)

Position Requirements

- Undergraduate or postgraduate students, preferably majoring in Finance, Investments, Legal Studies, and/or Business Analytics
- Experience in using Bloomberg, S&P Capital IQ, Pitchbook, and Wind preferred
- Exposure and understanding of the high-tech industry (especially AI/ML) a plus
- Strong sense of responsibility and attention to details

Availability: Spring and Summer 2020
Number of Openings: 1-2
Renumeration: Market competitive
Office: Central • Hong Kong

Anthony Woo
CFA CAIA FRM
Associate Director
Alpha Intelligence Capital
aw@aicapital.ai

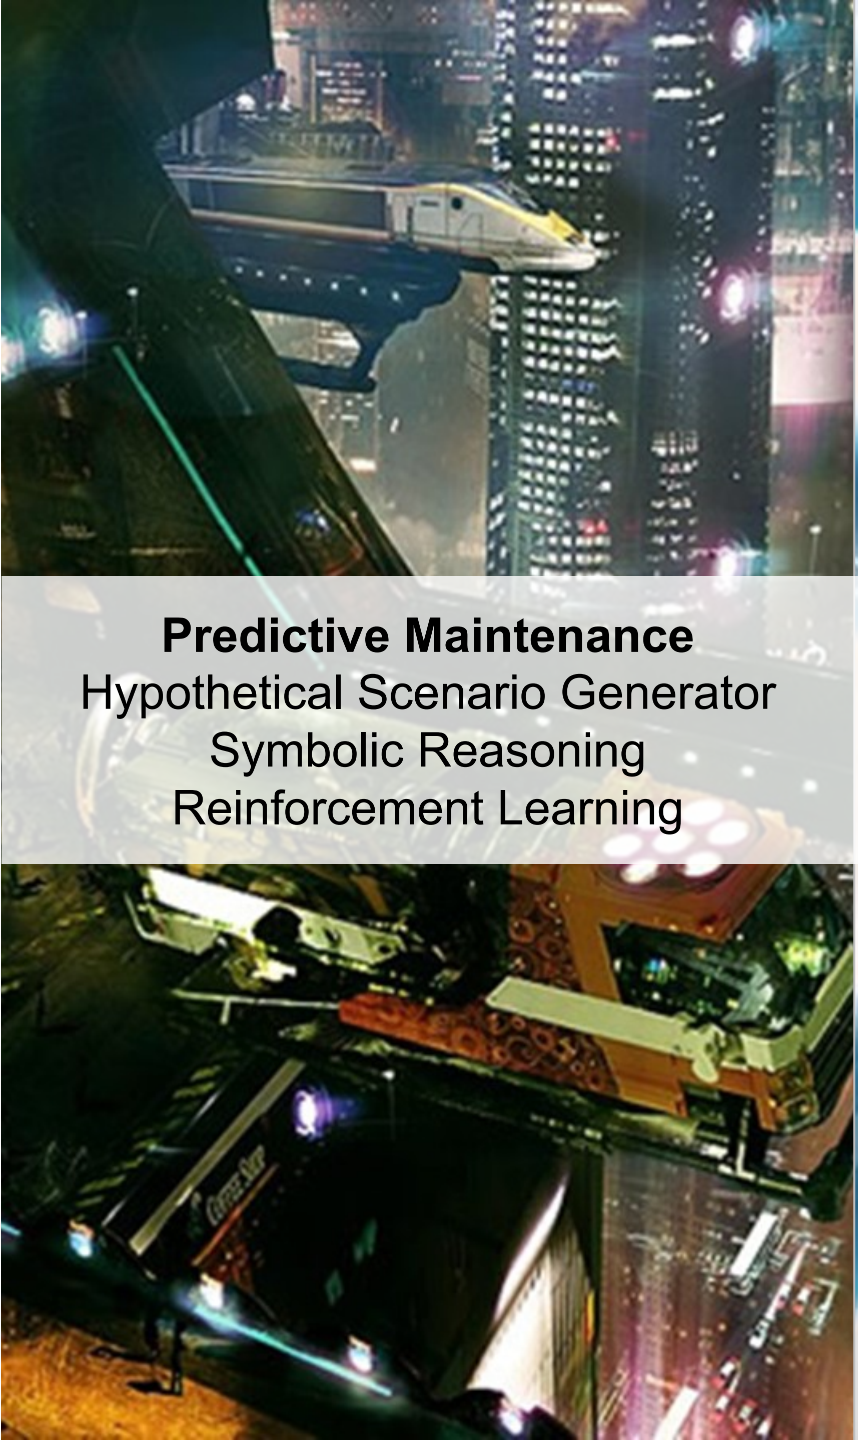

Predictive Maintenance
Hypothetical Scenario Generator
Symbolic Reasoning
Reinforcement Learning

Drug Discovery
Generative Adversarial Networks

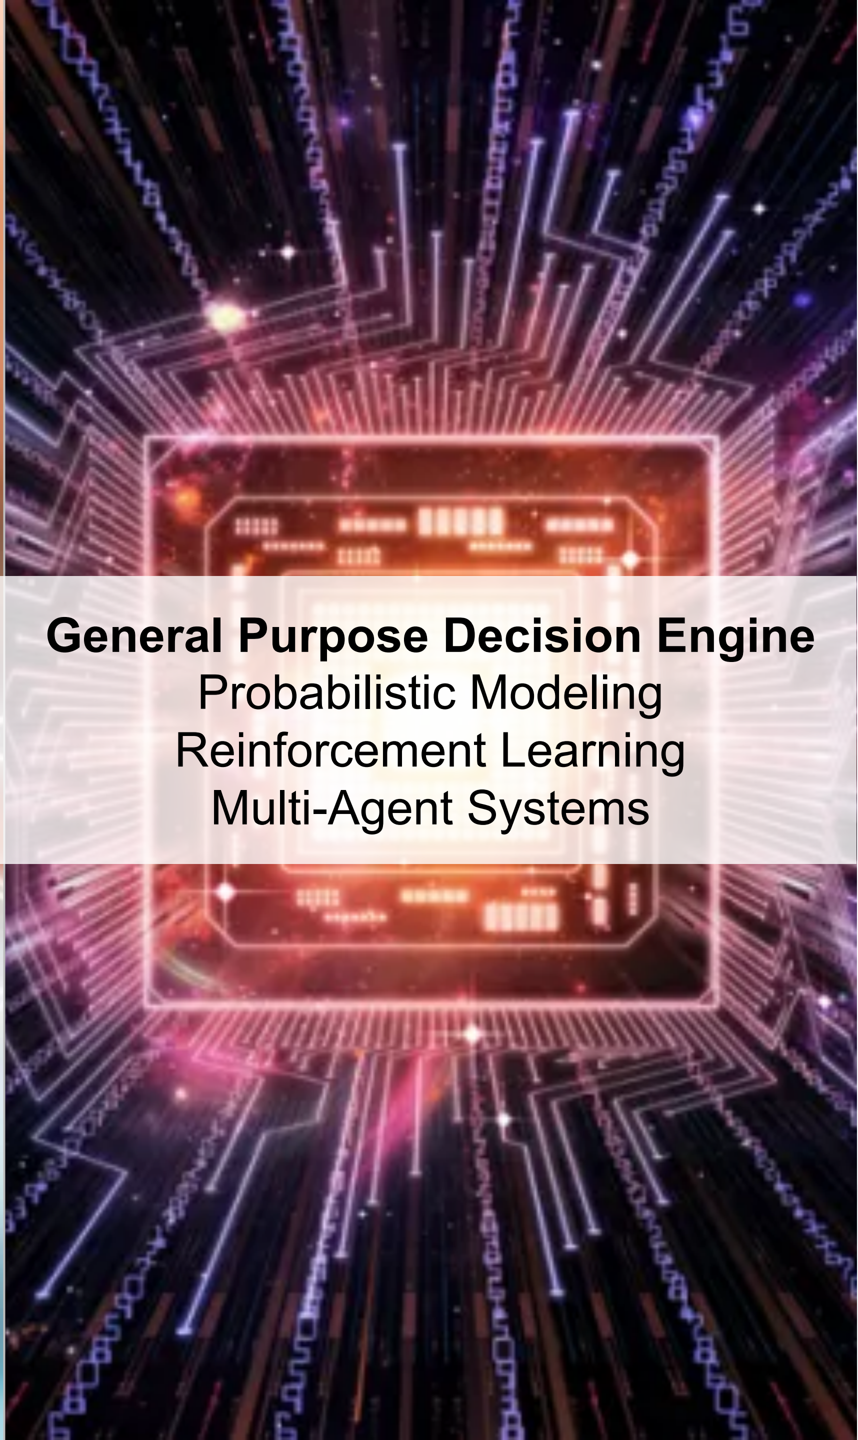

General Purpose Decision Engine
Probabilistic Modeling
Reinforcement Learning
Multi-Agent Systems

CFA Exam Adds Fintech to 2019 Curriculum

Come 2019, the wealth managers and financial analysts aspiring to add the Chartered Financial Analyst designation to their credentials have one more subject to deal with. The CFA Institute has decided to **add fintech to its 2019 exam curriculum**. The new curriculum contains a section called fintech and adds study material on hot industry topics such as **robo-advisors, big data, artificial intelligence and data analysis**. The new questions will appear in the CFA exam that will be administered in 2019.

1. **Financial Analysis Technology**: This includes how the financial analysis landscape is changing with things such as big data analysis, artificial intelligence, machine learning, and algorithmic trading.
2. **Portfolio Management Technology**: This includes robo-advisors, technology in enterprises such as asset management companies
3. **Capital Formation**: This includes peer-to-peer lending, shadow banking, and crowd funding.
4. **Market Infrastructure**: This includes innovations such as cryptocurrencies, blockchain technology, high-frequency trading, and regulatory-related technology

Source: https://financetrain.com/cfa-exam-adds-fintech-big-data-data-analysis-2019-curriculum/

Search

Home My Network Jobs Messaging Notifications Me Work Free Upgrade to Premium

Anthony Woo, CFA CAIA FRM

Associate Director at Alpha Intelligence Capital

888
Who's viewed your profile

7,567
Connections
Manage your network

Access exclusive tools & insights
Free Upgrade to Premium

Share an article, photo, video or idea

Write an article | Images | Video | Post

Sort by: Top

Germain Chastel
Developing the First Open Knowledge Access Platform
16h

The market for solar homes has skyrocketed - and these green tech disruptors are winning the race to reduce our carbon footprint:

The Market For Solar Homes Has Skyrocketed, and These Green Tech Disruptors Are Winning the Race to Reduce our Carbon Footprint
Germain Chastel on LinkedIn
By Germain Chastel and Sascha Eder NewtonX recently conducted an extensive panel on the pos...

18 Likes

Like Comment Share

1010 Corporate Solutions
3,087 followers
Promoted
+ Follow

Add to your feed

Investment Banking + Follow

Hussain Ali
HR Manager at Canadian Medical... + Follow

Ovilia. Mendonca.
Vice President-Flight Services⭐ Corporate... + Follow

View all recommendations

Messaging

Resources for a Career in Finance

1. FinanceAsia (http://www.financeasia.com/)
2. Dealogic (http://www.dealogic.com/)
3. McKinsey Insights (http://www.mckinsey.com/insights)
4. 清科集团：投资界 (http://www.pedaily.cn/)
5. Wind Financial Terminal

HireVue Iris™ Deep Learning Analytics Engine

15,000
PREDICTIVE ATTRIBUTES

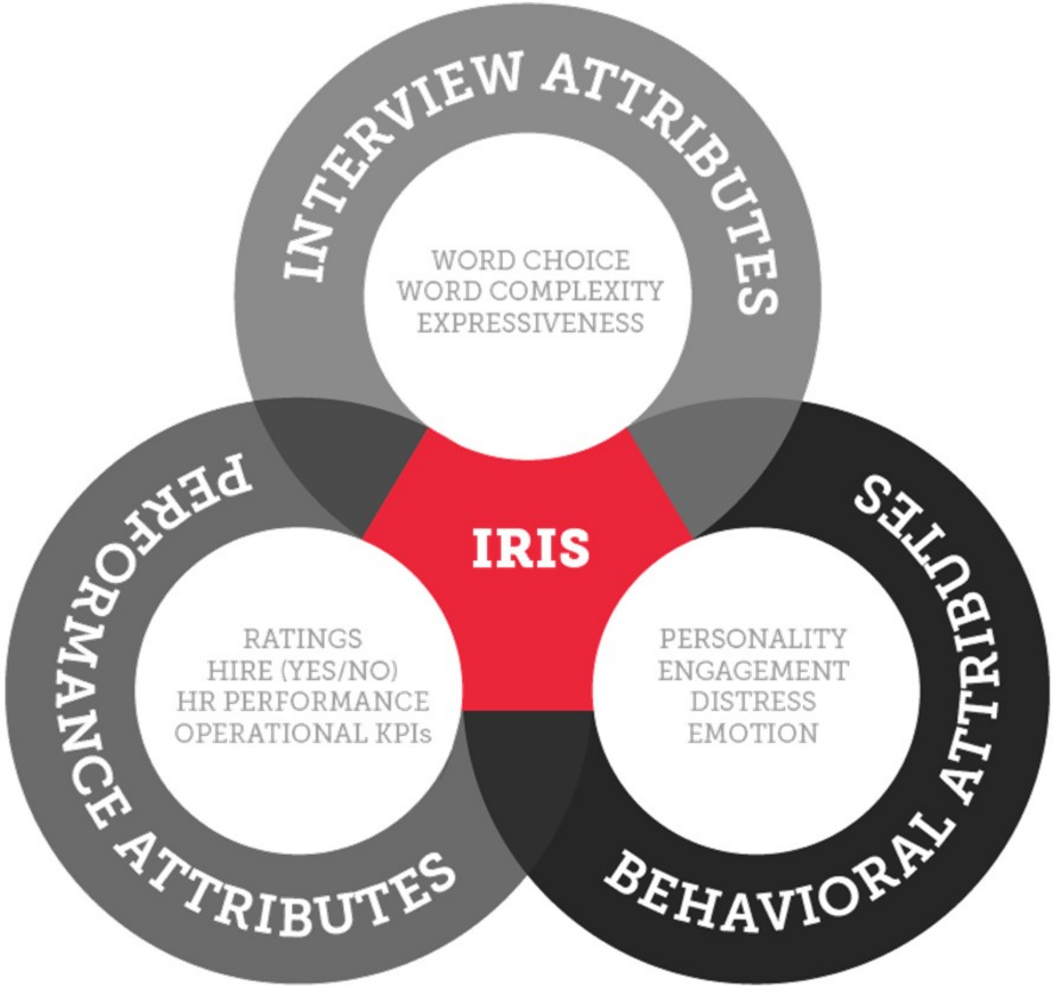

100,000X
MORE DATA THAN A RESUME

W

Aug 2011 – May 2013 | **MORGAN STANLEY ASIA** | **HONG KONG, CHINA**

Associate, Investment Banking Division – Hong Kong Corporate Finance Coverage Team

- Shangri-La Asia inaugural issuance of US$600MM under US$3Bn Medium Term Note Program
- Multiple senior unsecured bond offerings for Hong Kong corporates, such as Kerry Properties (US$600MM), Hang Lung Properties (US$500MM), Nan Fung (US$300MM), and PCCW (US$500MM)

Summer 2010 — **Summer Associate, Investment Banking Division – Technology, Media & Telecommunications Group**

- US$272MM IPO of Dangdang Inc., China's largest B2C e-commerce company (equivalent of Amazon)

2006 – 20 — **J.P. MORGAN CHASE & CO.**

2008 – 2009 — **Investment Strategy Analyst, J.P. Morgan Private Wealth Management** — **NEW YORK, NY**

- Sole analyst directly supporting the global Chief Investment Officer (CIO) and Chief Economist of PWM

2006 – 2008 — **Financial Analyst, J.P. Morgan Private Bank** — **SAN FRANCISCO, CA**

Summer 2007 — **Financial Analyst, J.P. Morgan Private Bank, EMEA Equity Derivatives Group** — **LONDON, UK**

E

Education

2015 – 20 — **UNIVERSITY OF HONG KONG** — **HONG KONG, CHINA**

Master of Science in Information Technology in Education (Specialist Strand: e-Leadership), *Distinction*.

2009 – 20 — **HARVARD BUSINESS SCHOOL** — **BOSTON, MA**

MBA. Co-producer, Asian Cultural Show. Advisor, Harvard Innovation Lab (iLab)

2002 – 20 — **UNIVERSITY OF CALIFORNIA, BERKELEY – HAAS SCHOOL OF BUSINESS** — **BERKELEY, CA**

Bachelor of Science in Business Administration, *summa cum laude* (cu[…]lative GPA: 3.9, top 3% of class). Dean's Honor List (02-06). President, California Investment Associat[…]as-sponsored investment fund)

T

Technology-related Certifications

Certificate on Machine Learning for [...]ata & Text Processing at MIT Computer Science and Artificial Intelligence Laboratory (CSAIL). Ce[...]e on Deep Learning and Machine Learning with TensorFlow. Certified Bitcoin Professional (CBP). Conducted research into Probabilistic Topic Modeling using R

Career Roadmap & CV Clinic

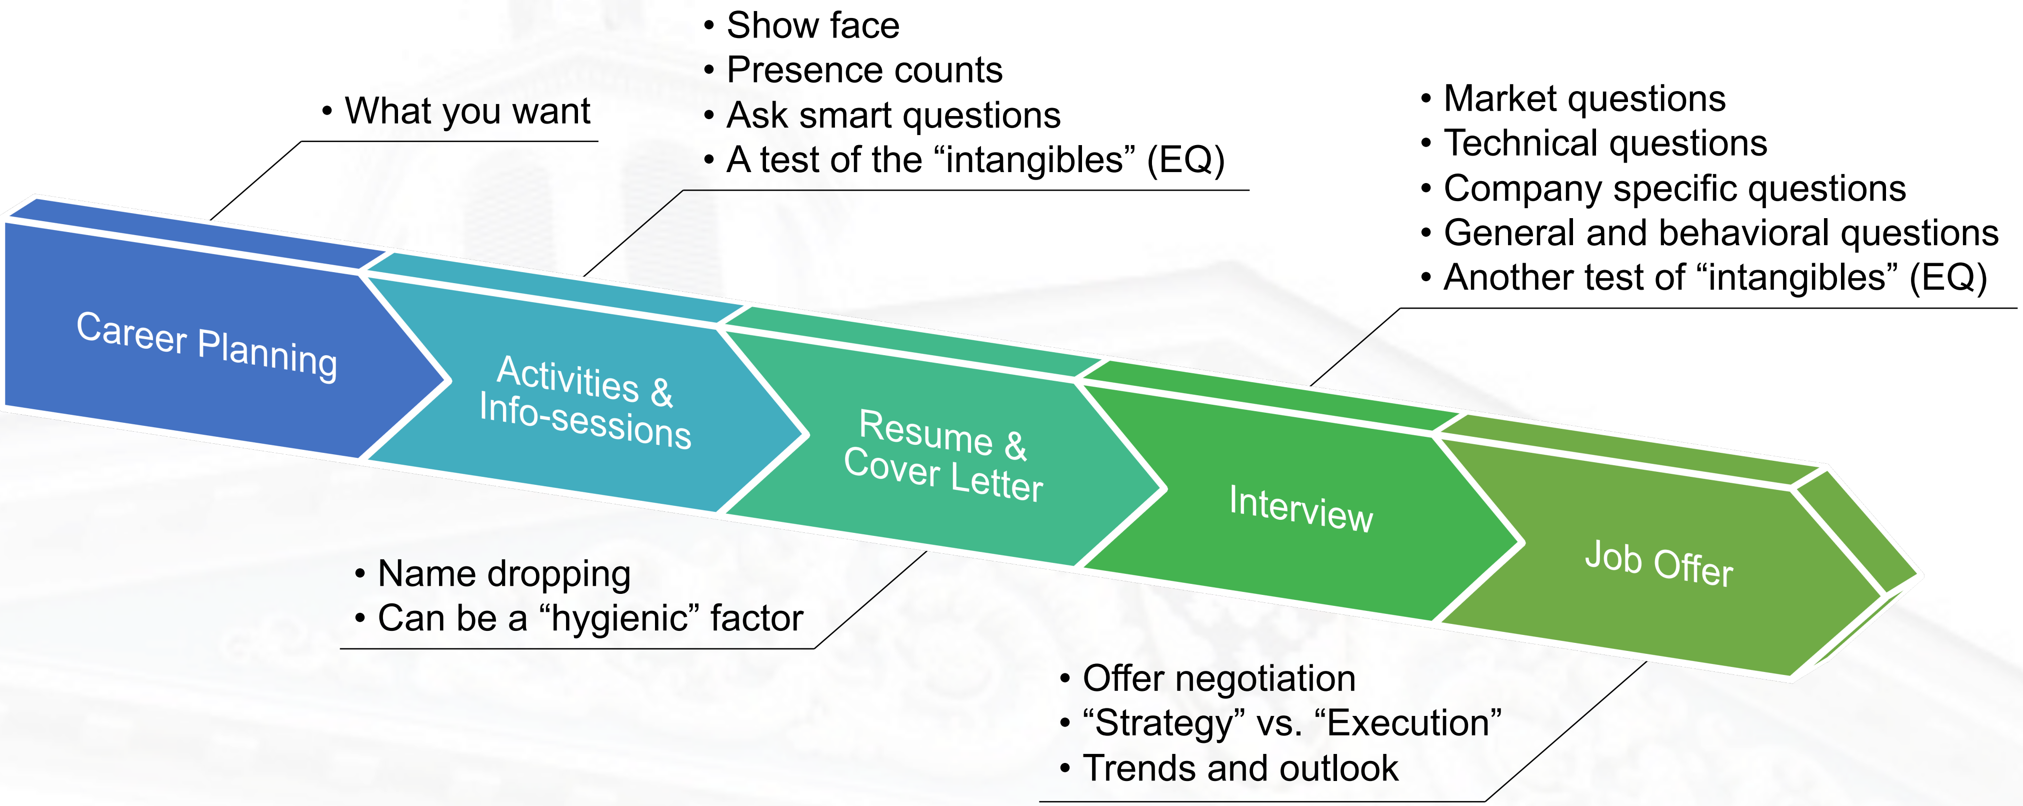

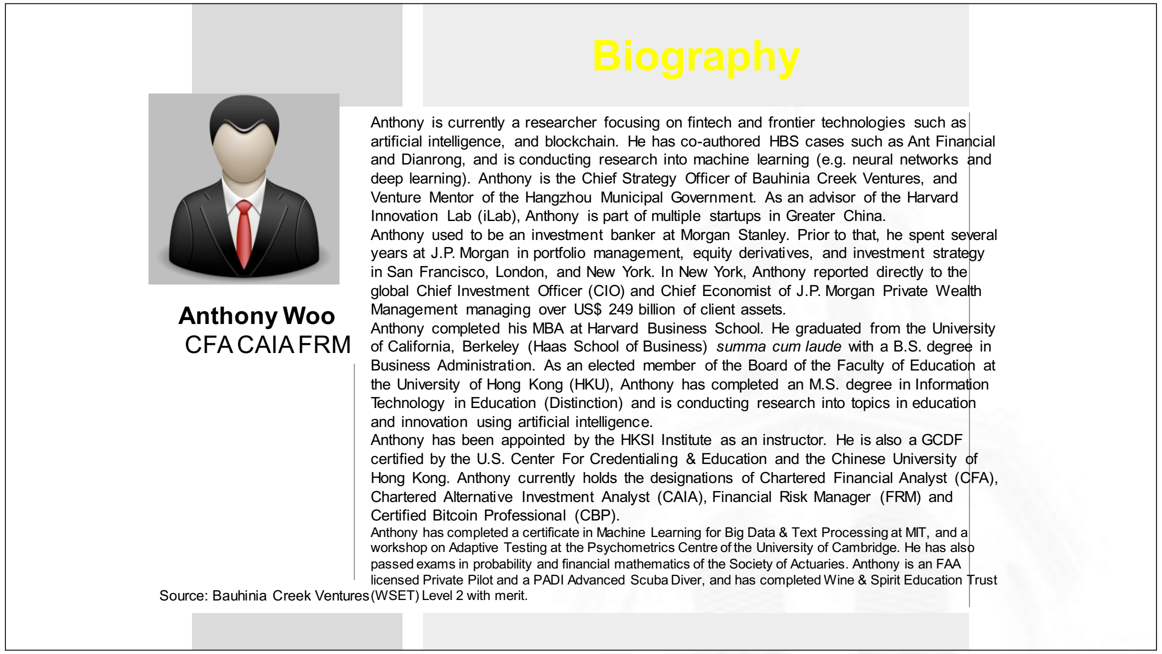
Biography
Anthony Woo
CFA CAIA FRM
Anthony is currently a researcher focusing on fintech and frontier technologies such as artificial intelligence, and blockchain. He has co-authored HBS cases such as Ant Financial and Dianrong, and is conducting research into machine learning (e.g. neural networks and deep learning). Anthony is the Chief Strategy Officer of Bauhinia Creek Ventures, and Venture Mentor of the Hangzhou Municipal Government. As an advisor of the Harvard Innovation Lab (iLab), Anthony is part of multiple startups in Greater China.
Anthony used to be an investment banker at Morgan Stanley. Prior to that, he spent several years at J.P. Morgan in portfolio management, equity derivatives, and investment strategy in San Francisco, London, and New York. In New York, Anthony reported directly to the global Chief Investment Officer (CIO) and Chief Economist of J.P. Morgan Private Wealth Management managing over US$ 249 billion of client assets.
Anthony completed his MBA at Harvard Business School. He graduated from the University of California, Berkeley (Haas School of Business) summa cum laude with a B.S. degree in Business Administration. As an elected member of the Board of the Faculty of Education at the University of Hong Kong (HKU), Anthony has completed an M.S. degree in Information Technology in Education (Distinction) and is conducting research into topics in education and innovation using artificial intelligence.
Anthony has been appointed by the HKSI Institute as an instructor. He is also a GCDF certified by the U.S. Center For Credentialing & Education and the Chinese University of Hong Kong. Anthony currently holds the designations of Chartered Financial Analyst (CFA), Chartered Alternative Investment Analyst (CAIA), Financial Risk Manager (FRM) and Certified Bitcoin Professional (CBP).
Anthony has completed a certificate in Machine Learning for Big Data & Text Processing at MIT, and a workshop on Adaptive Testing at the Psychometrics Centre of the University of Cambridge. He has also passed exams in probability and financial mathematics of the Society of Actuaries. Anthony is an FAA licensed Private Pilot and a PADI Advanced Scuba Diver, and has completed Wine & Spirit Education Trust (WSET) Level 2 with merit.
Source: Bauhinia Creek Ventures


Poor contrast and illegible
Text overflow
Too many words and full sentences not helpful
Poor image selection
Lack of emphasis
Anthony Woo
CFA CAIA FRM
Misalignment
Unexplained jargon
Inconsistent font size
Inadequate referencing
Missing page number
Source: Bauhinia Creek Ventures

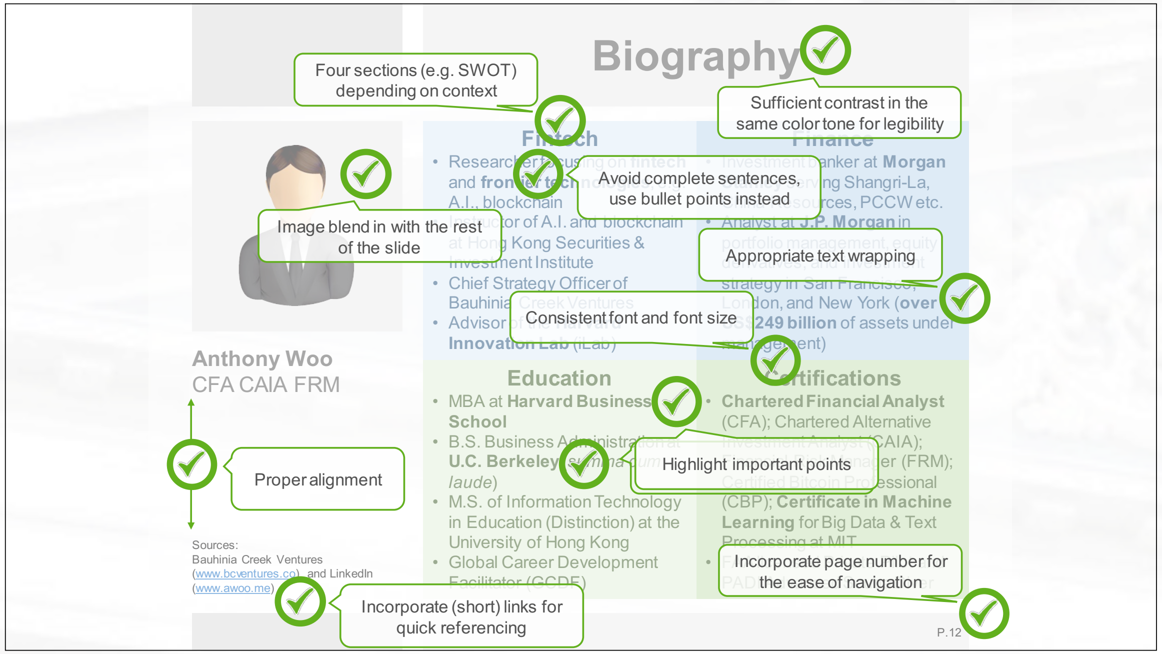
Biography
Four sections (e.g. SWOT) depending on context
Sufficient contrast in the same color tone for legibility
Fintech
Finance
Avoid complete sentences, use bullet points instead
Image blend in with the rest of the slide
Appropriate text wrapping
Consistent font and font size
Anthony Woo
CFA CAIA FRM
Education
Certifications
Highlight important points
Proper alignment
Sources:
Bauhinia Creek Ventures (www.bcventures.co), and LinkedIn (www.awoo.me)
Incorporate page number for the ease of navigation
Incorporate (short) links for quick referencing
P.12

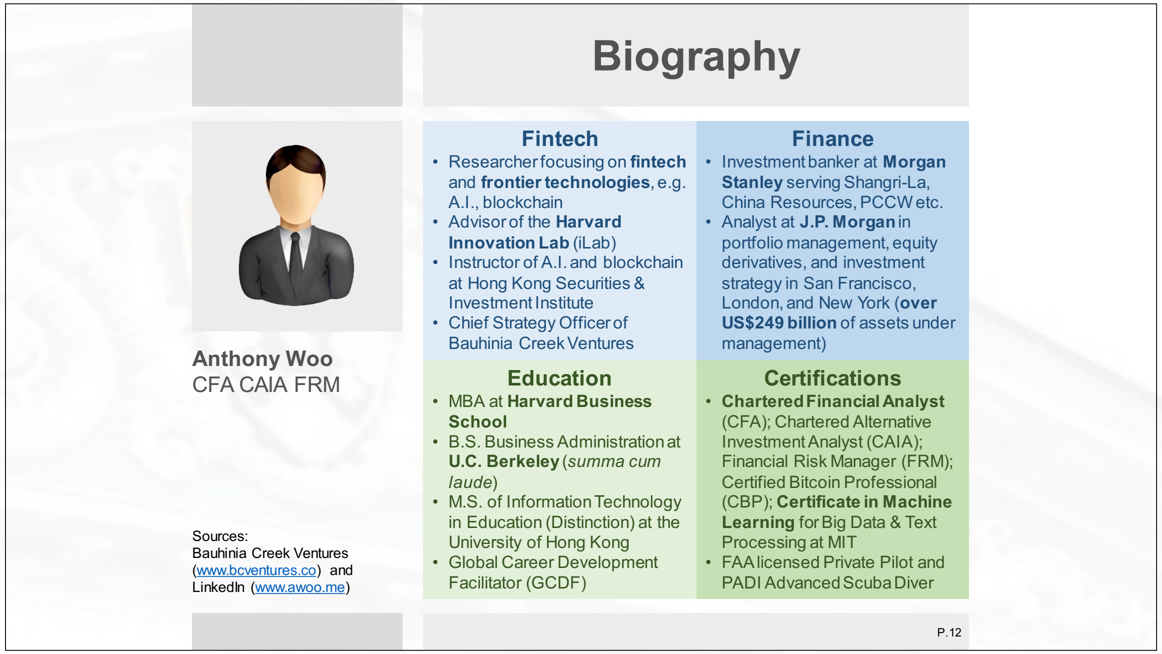
Biography
Anthony Woo
CFA CAIA FRM
Fintech
• Researcher focusing on fintech and frontier technologies, e.g. A.I., blockchain
• Advisor of the Harvard Innovation Lab (iLab)
• Instructor of A.I. and blockchain at Hong Kong Securities & Investment Institute
• Chief Strategy Officer of Bauhinia Creek Ventures
Finance
• Investment banker at Morgan Stanley serving Shangri-La, China Resources, PCCW etc.
• Analyst at J.P. Morgan in portfolio management, equity derivatives, and investment strategy in San Francisco, London, and New York (over US$249 billion of assets under management)
Education
• MBA at Harvard Business School
• B.S. Business Administration at U.C. Berkeley (summa cum laude)
• M.S. of Information Technology in Education (Distinction) at the University of Hong Kong
• Global Career Development Facilitator (GCDF)
Certifications
• Chartered Financial Analyst (CFA); Chartered Alternative Investment Analyst (CAIA); Financial Risk Manager (FRM); Certified Bitcoin Professional (CBP); Certificate in Machine Learning for Big Data & Text Processing at MIT
• FAA licensed Private Pilot and PADI Advanced Scuba Diver
Sources:
Bauhinia Creek Ventures (www.bcventures.co) and LinkedIn (www.awoo.me)
P.12