

The Sorcerer's Apprentice

or

Why Semantics still matters and why we should not only trust in ML

Prof. Dr. Harald Sack
Sustainable Knowledge Graphs and AI
Gothenburg, 26. Apr. 2019

**“Any sufficiently advanced technology
is indistinguishable from magic.”**

Arthur C. Clarke, Profiles of the Future (1973)





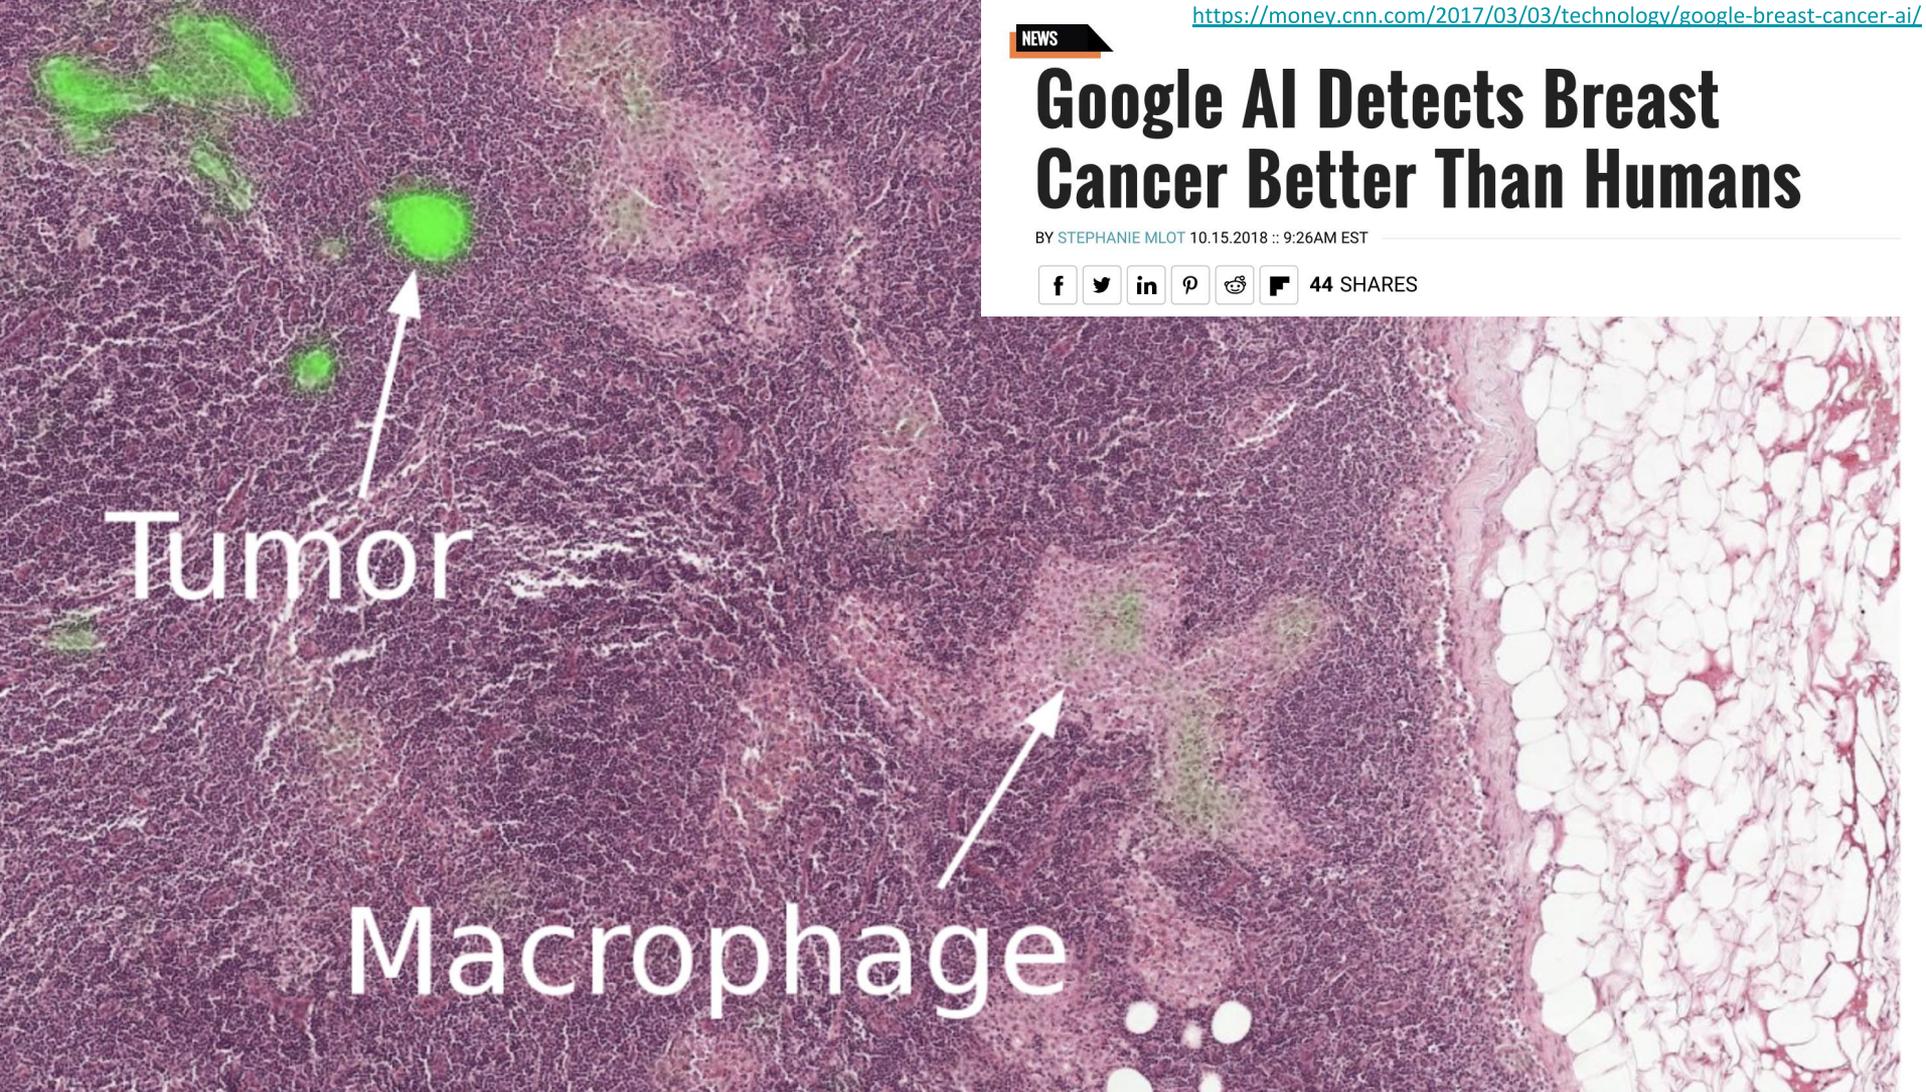
Can you find the cancer?

NEWS

Google AI Detects Breast Cancer Better Than Humans

BY STEPHANIE MLOT 10.15.2018 :: 9:26AM EST

      44 SHARES



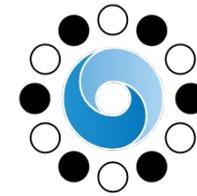
Tumor

Macrophage

AlphaGo Zero: Google DeepMind supercomputer learns 3,000 years of human knowledge in 40 days



17



AlphaGo



<http://www.telegraph.co.uk/science/2017/10/18/alphago-zero-google-deepmind-supercomputer-learns-3000-years/>



<https://www.christies.com/features/A-collaboration-between-two-artists-one-human-one-a-machine-9332-1.aspx>

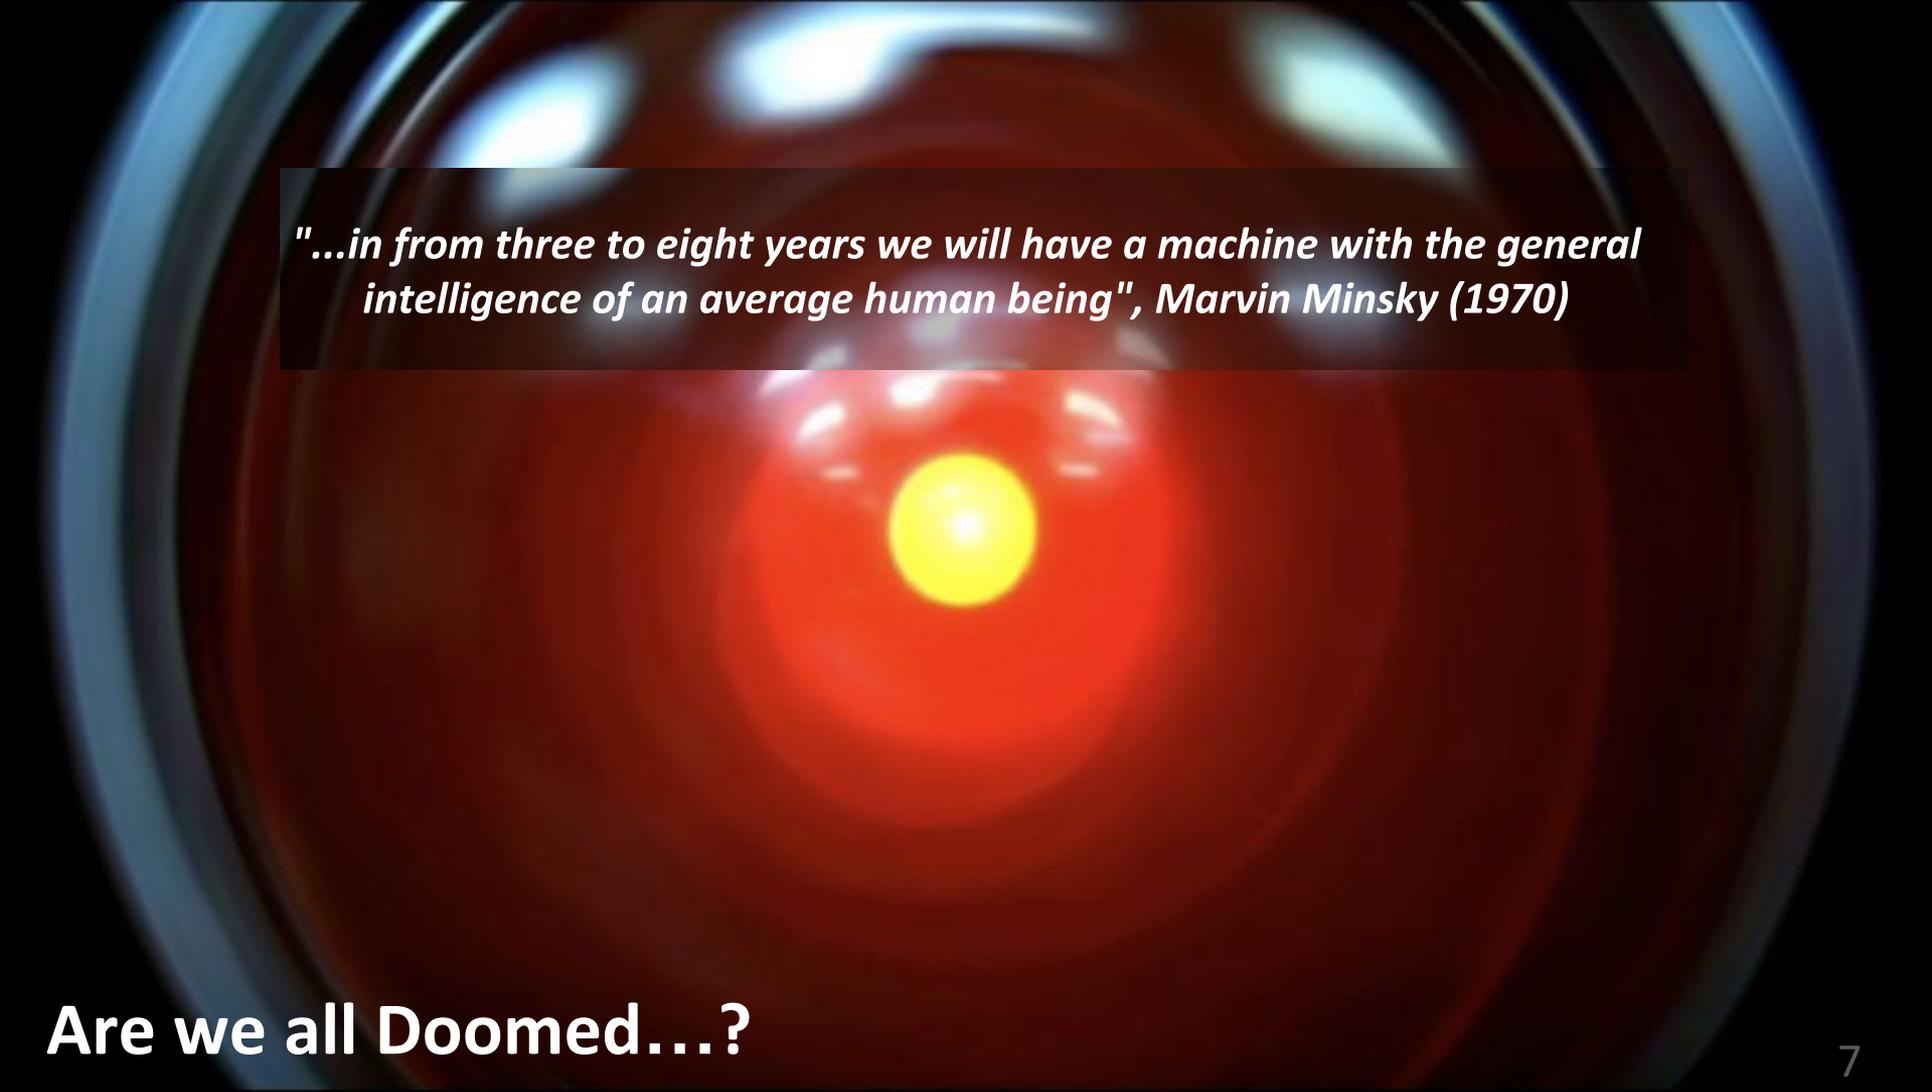


Is artificial intelligence set to become art's next medium?

16 October 2018

PHOTOGRAPHS & PRINTS |

AI artwork sells for \$432,500 — nearly 45 times its high estimate — as Christie's becomes the first auction house to offer

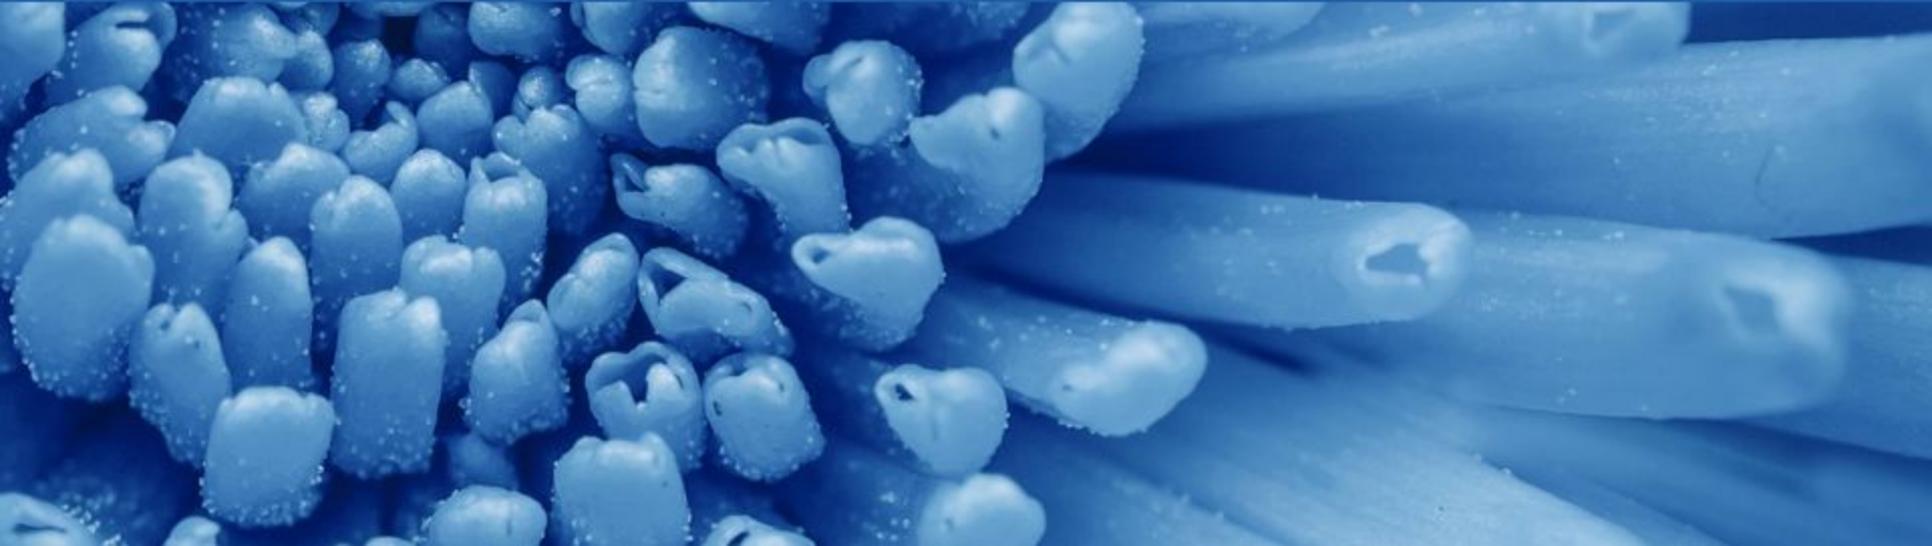


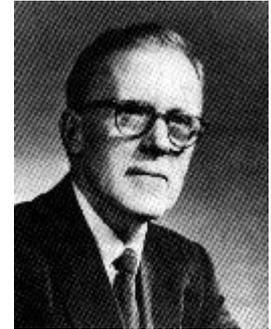
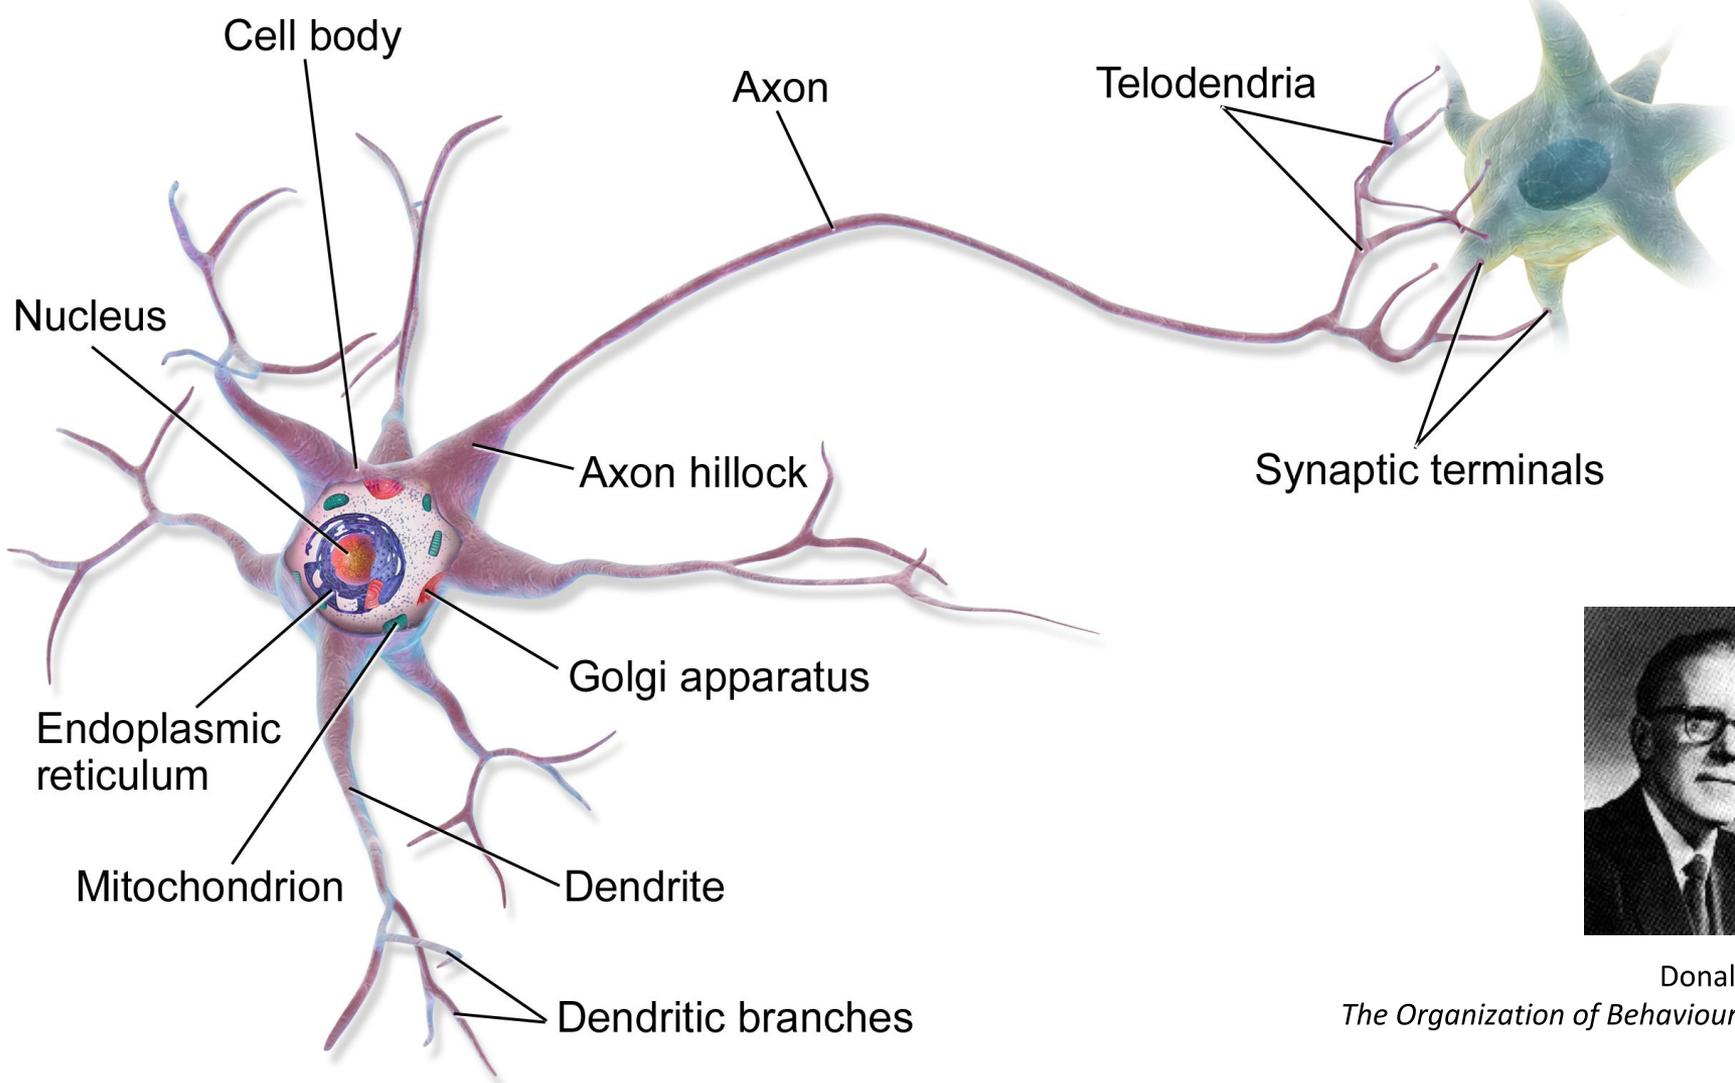
"...in from three to eight years we will have a machine with the general intelligence of an average human being", Marvin Minsky (1970)

Are we all Doomed...?

Where Did the Magic Come From?

A Short History of Machine Learning





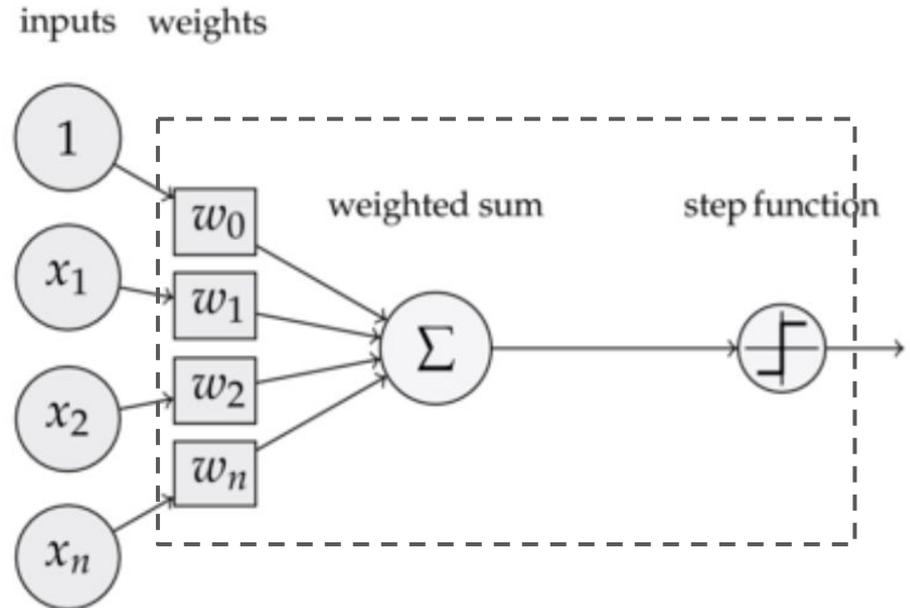
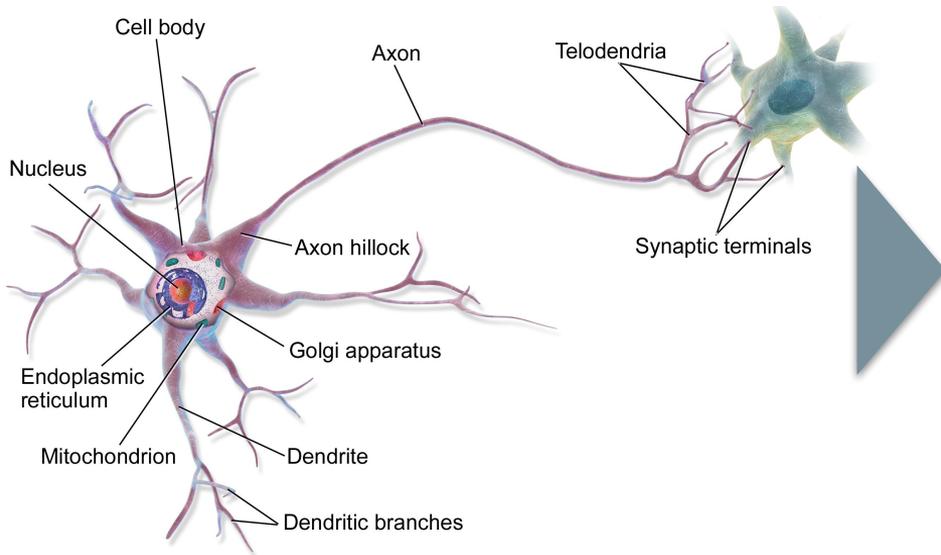
Donald Hebb

The Organization of Behaviour (1949)

From Biology to Mathematical Models



Warren Sturgis McCulloch & Walter Pitts
A Logical Calculus of the Ideas Immanent in Nervous Activity (1943)

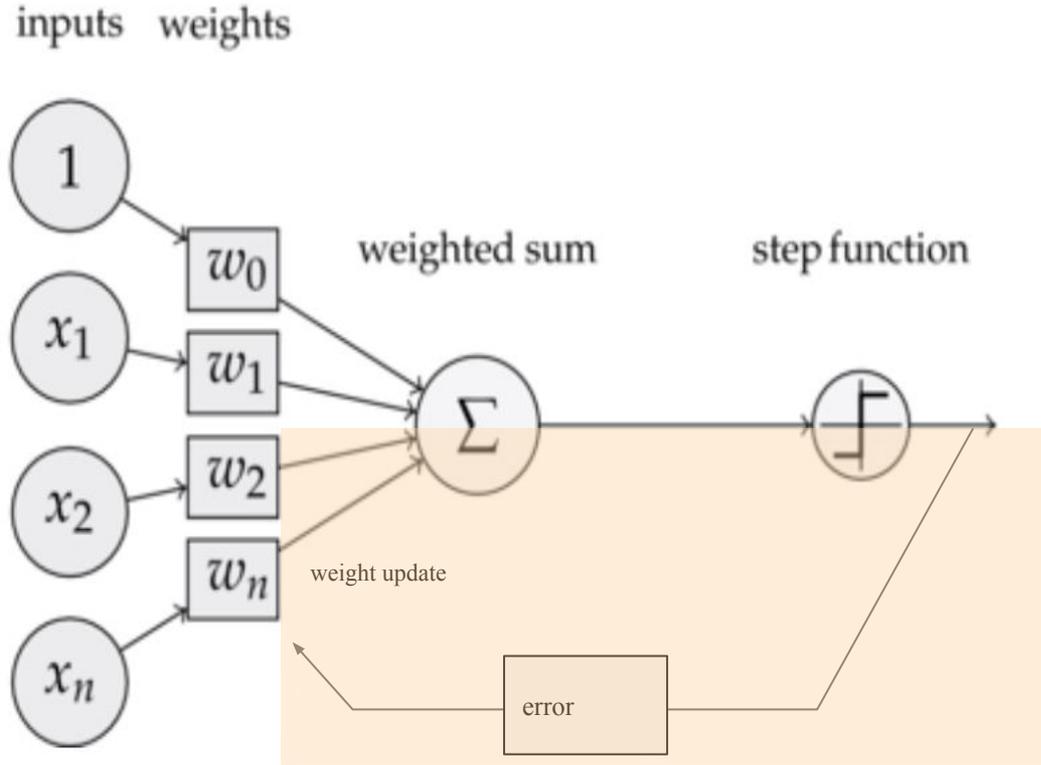


Perceptron Algorithm



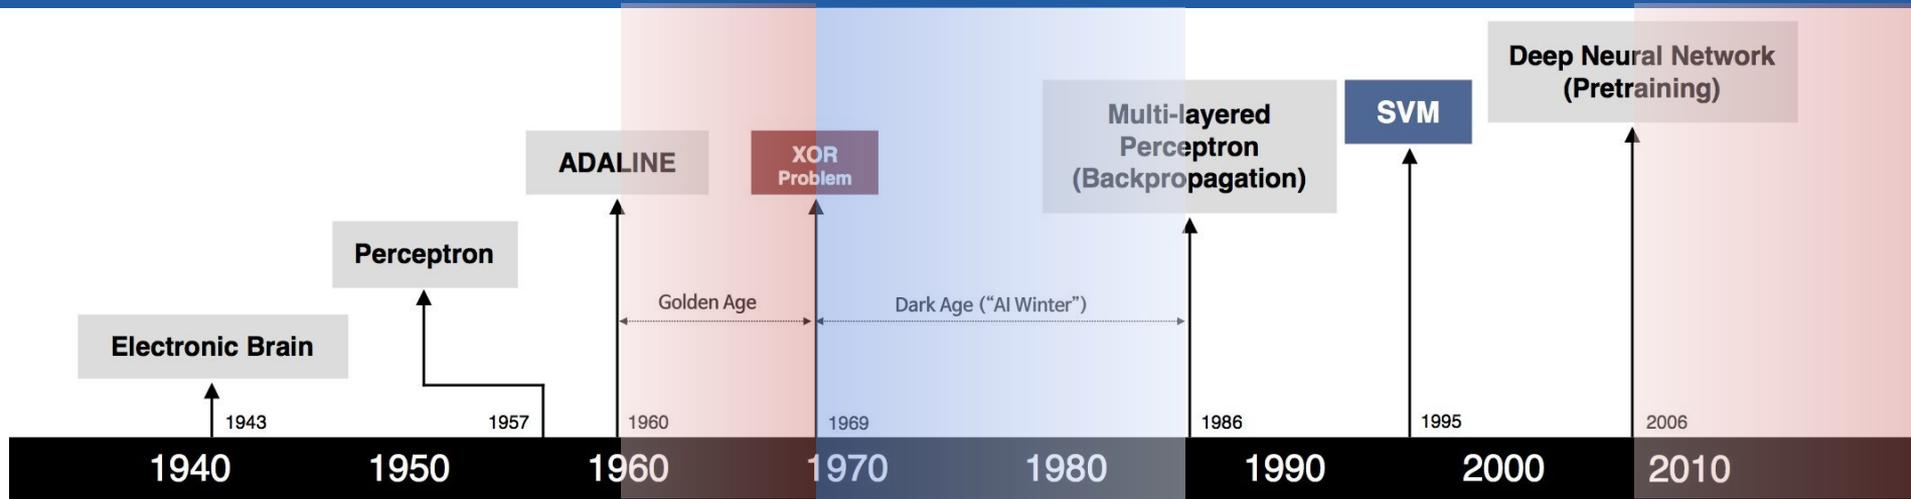
Frank Rosenblatt

The perceptron: a probabilistic model for information storage and organization in the brain. (1958)

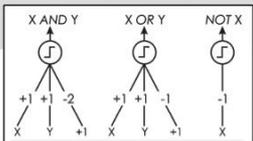


$$w_{ij}^{\text{neu}} = w_{ij}^{\text{alt}} + \Delta w_{ij},$$
$$\Delta w_{ij} = \alpha \cdot (t_j - o_j) \cdot x_i.$$

The Rise of the Connectivist Paradigm



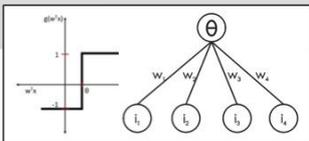
S. McCulloch – W. Pitts



- Adjustable Weights
- Weights are not Learned



F. Rosenblatt



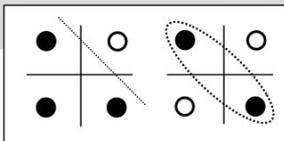
- Learnable Weights and Threshold



B. Widrow – M. Hoff



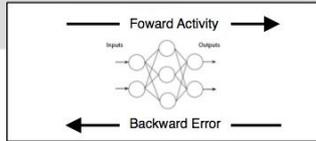
M. Minsky – S. Papert



- XOR Problem



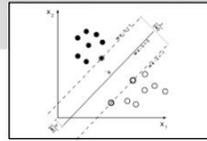
D. Rumelhart – G. Hinton – R. Williams



- Solution to nonlinearly separable problems
- Big computation, local optima and overfitting



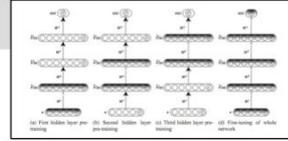
V. Vapnik – C. Cortes



- Limitations of learning prior knowledge
- Kernel function: Human Intervention

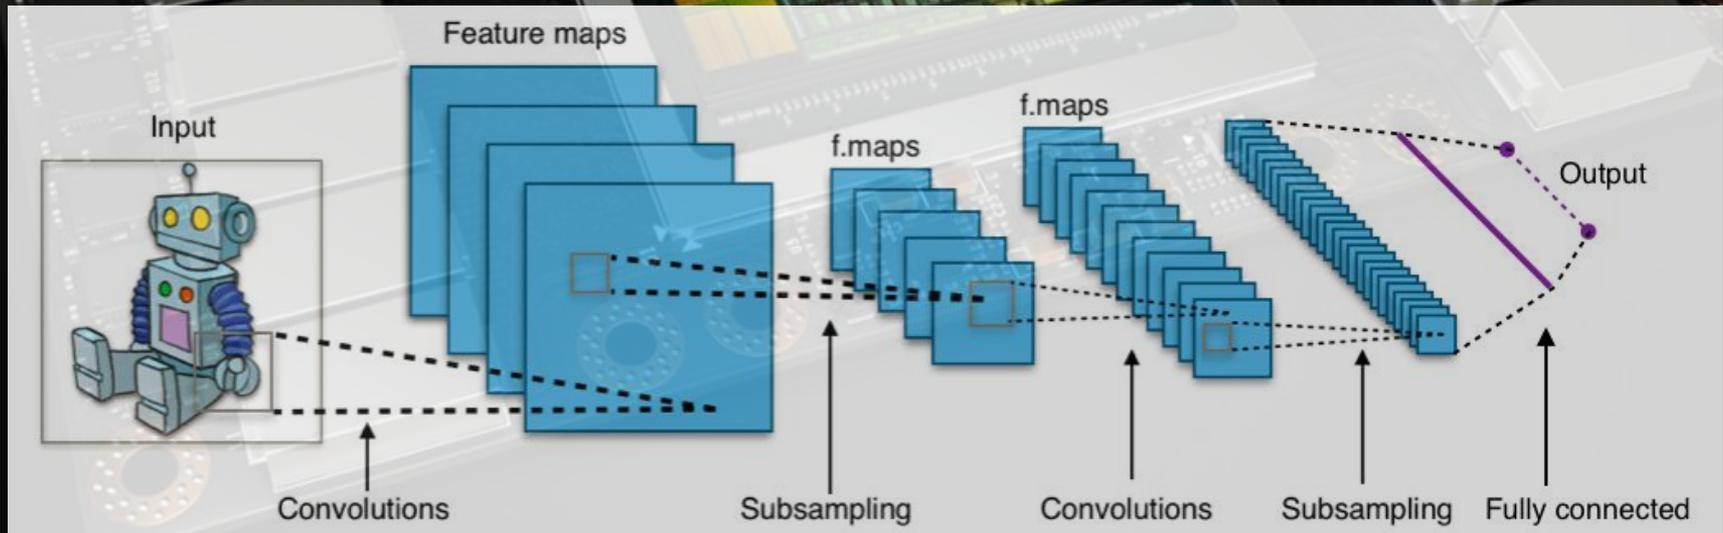


G. Hinton – S. Ruslan



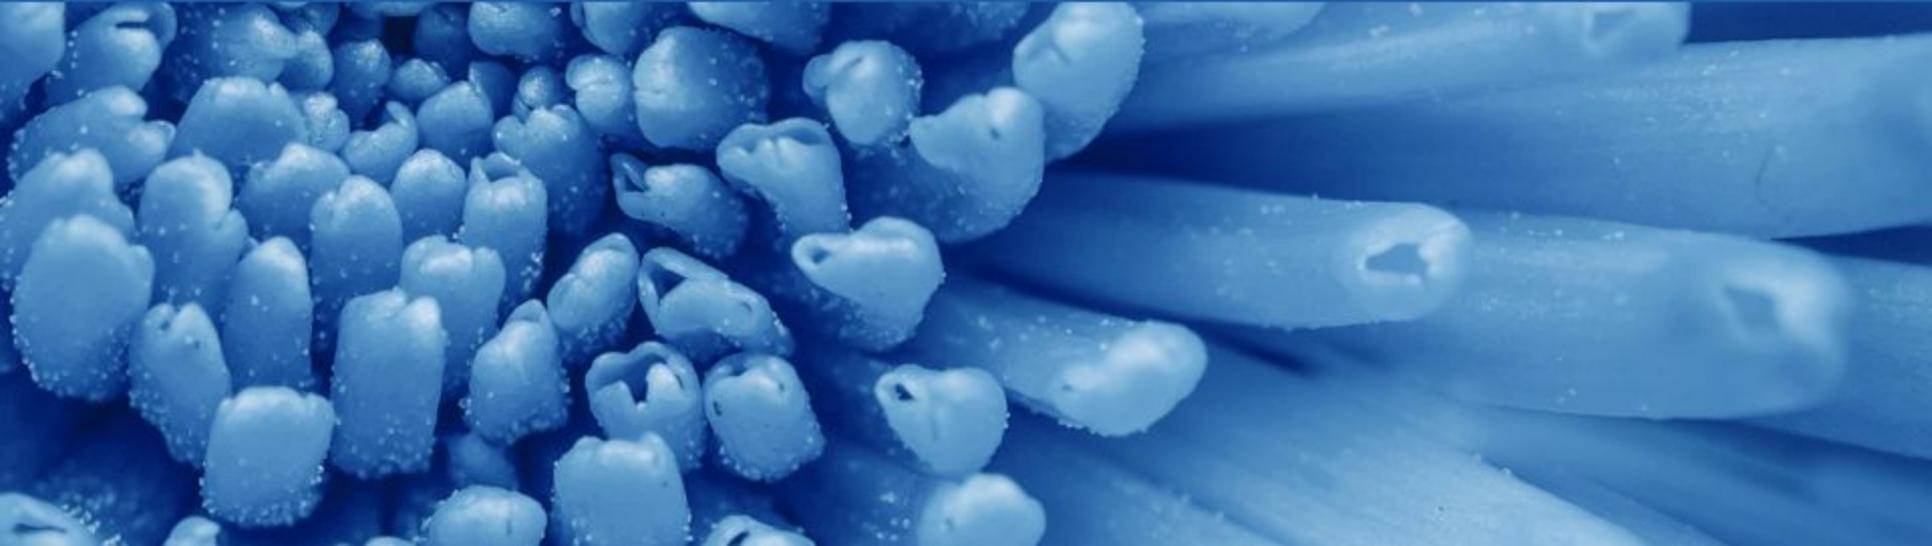
- Hierarchical feature Learning

Convolutional Deep Neural Networks and GPU Supercomputers

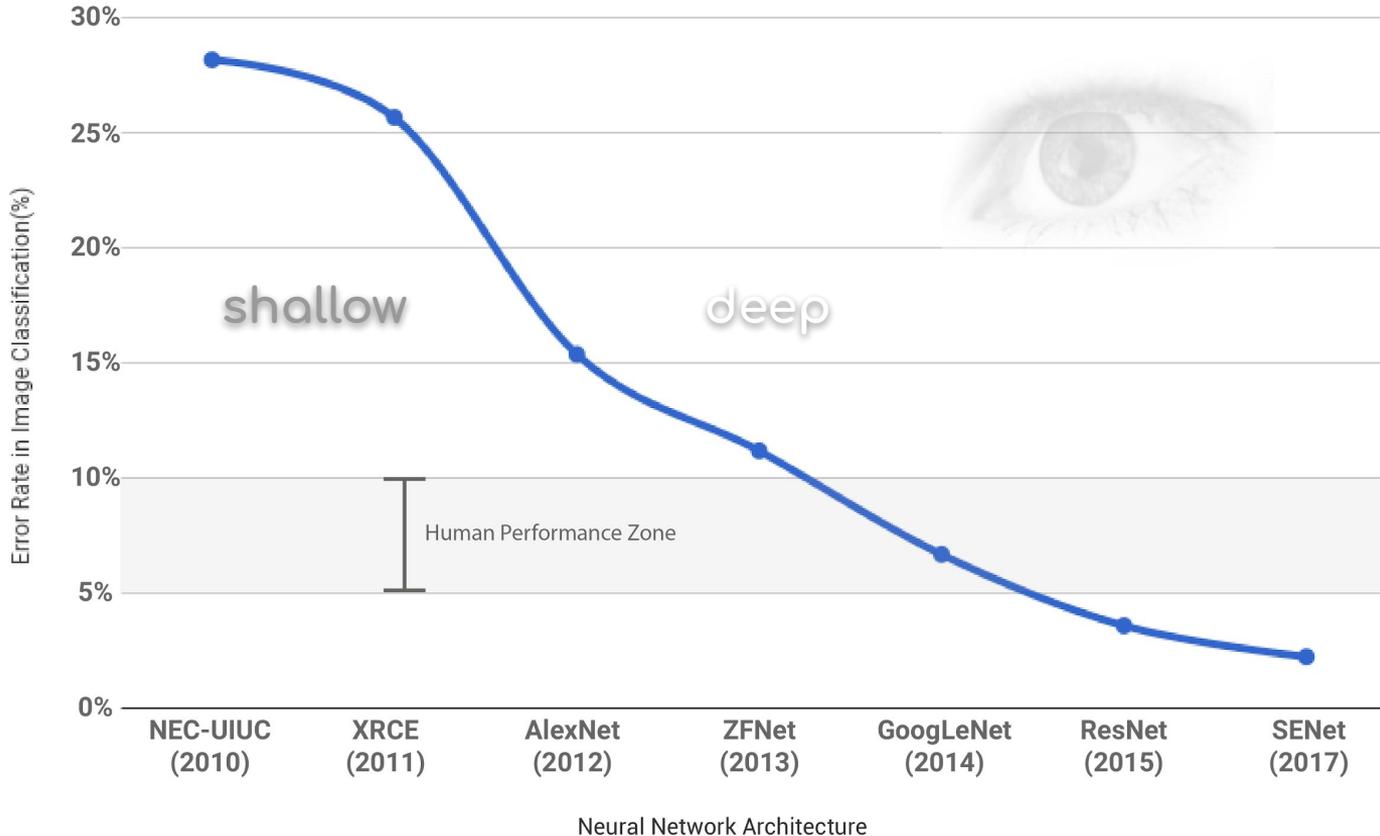


**“Many secrets of art and nature are thought
by the unlearned to be magical.”**

Roger Bacon (c. 1219/20 – c. 1292)



Deep Learning for Visual Analysis



IMAGENET

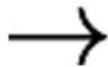
Large Scale Visual Recognition
Challenge (ILSVRC)

<http://image-net.org/challenges/LSVRC/>

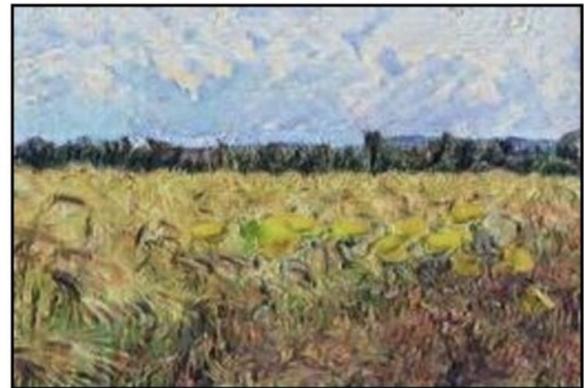
From Classification to Generation



Photograph



Monet



Van Gogh



Cezanne



Ukiyo-e

From Classification to Generation

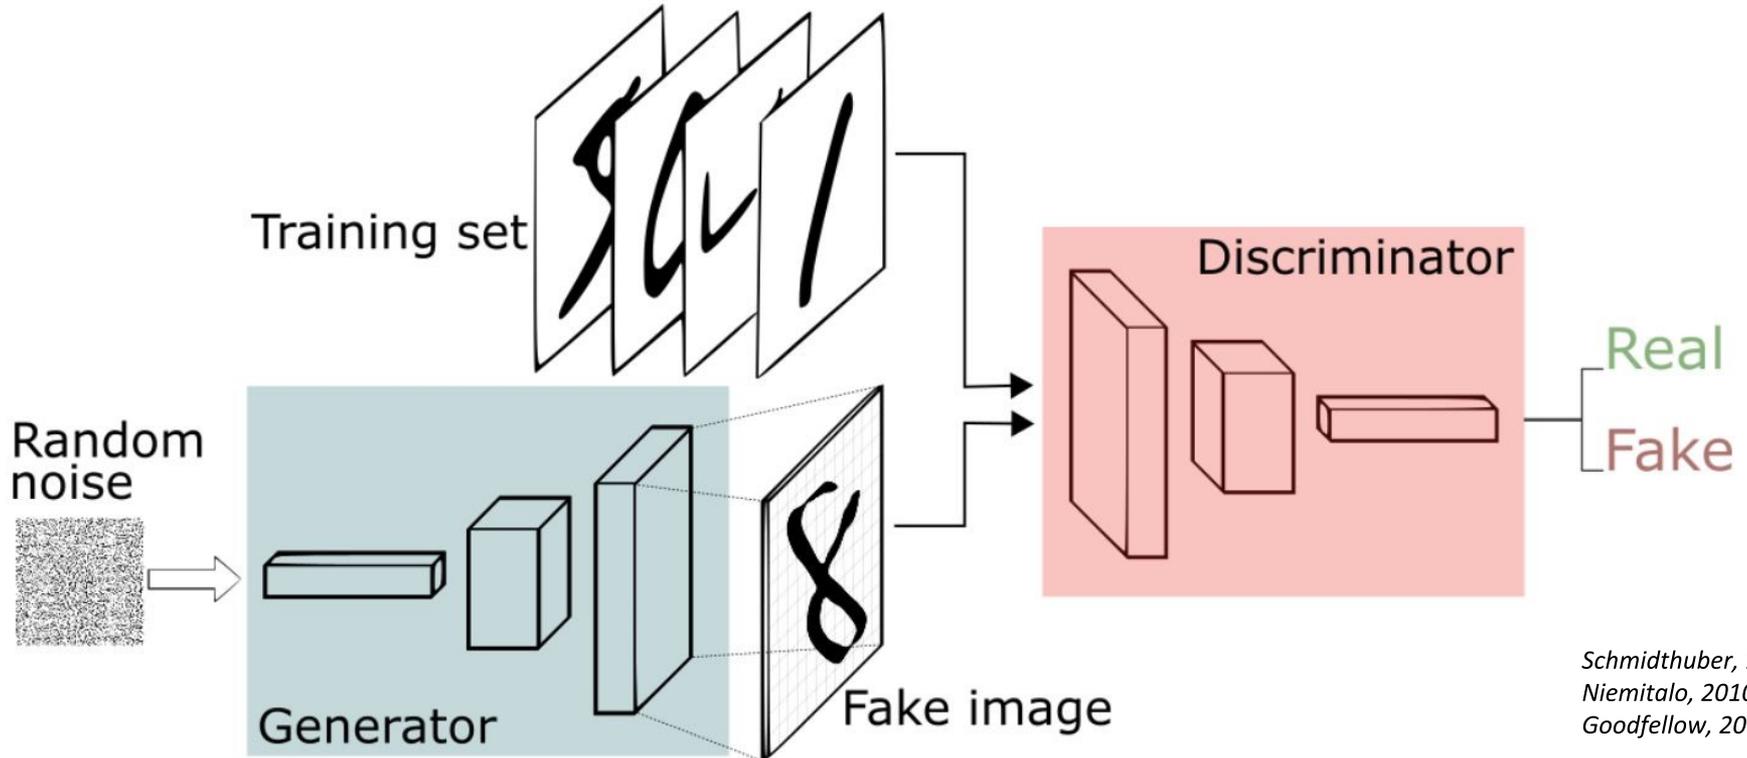
Monet



Photograph



Comparative Learning - Generative Adversarial Networks



Schmidhuber, 1992
Niimitalo, 2010
Goodfellow, 2014



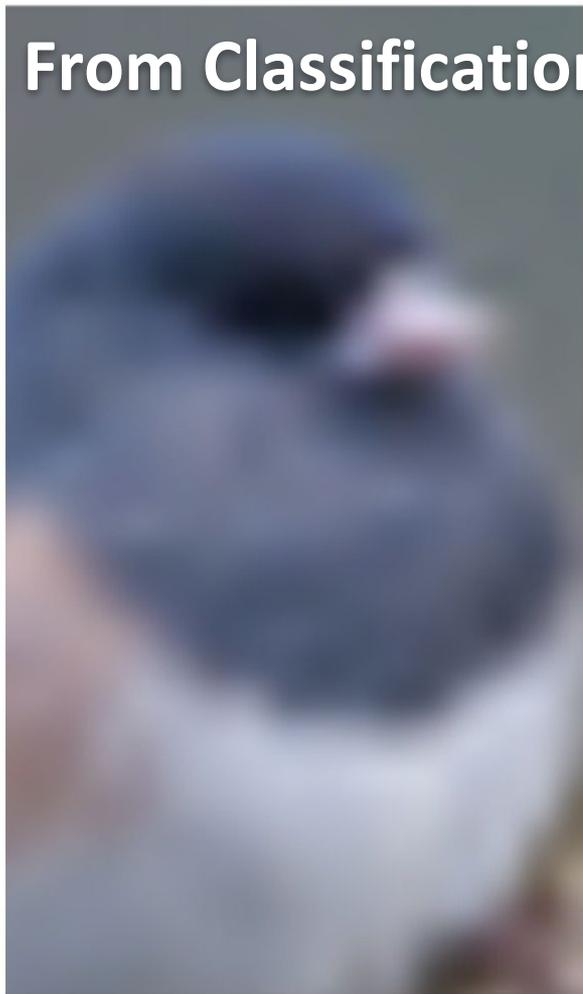
From Classification to Generation

From Classification to Generation

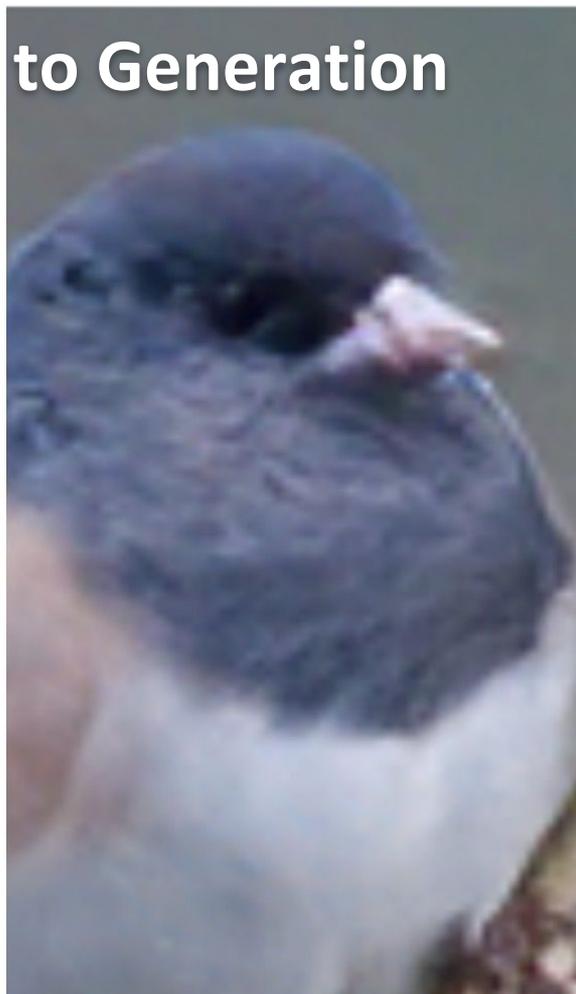


<https://iunvanz.github.io/CycleGAN/>

From Classification to Generation



Low-resolution input



Our result



Original image

Input



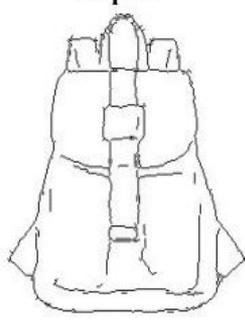
Ground truth



Output



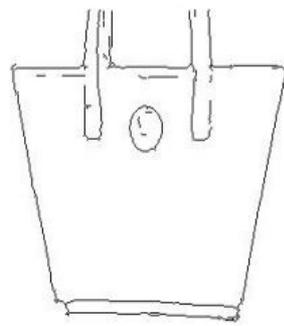
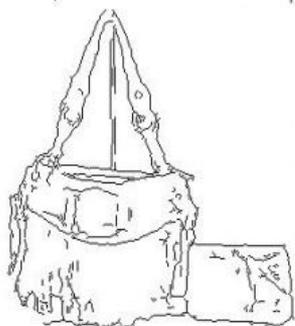
Input



Ground truth



Output



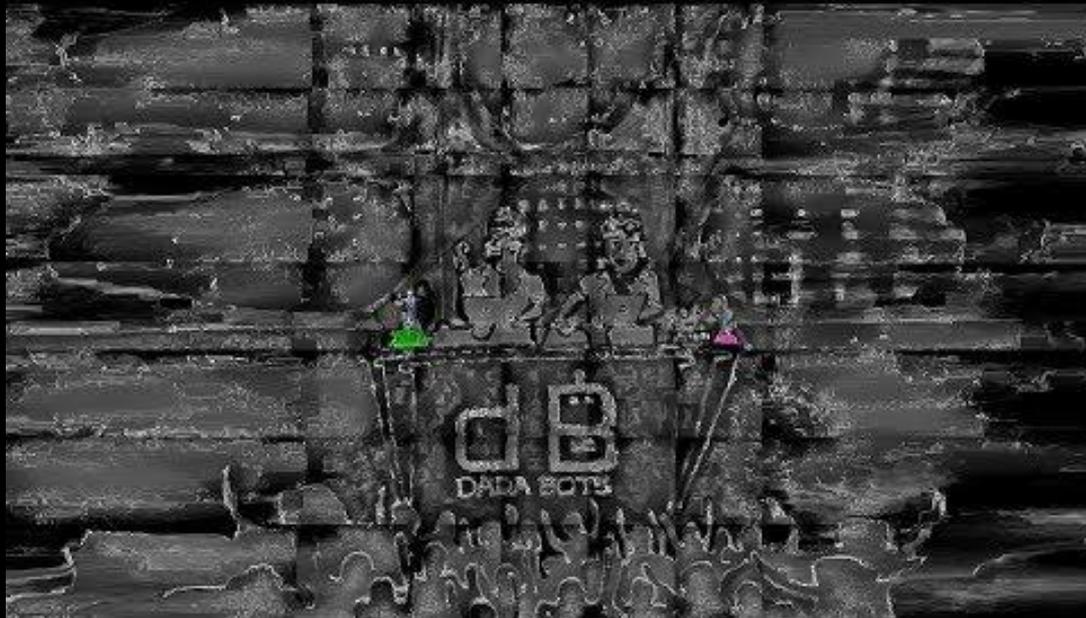
Can you tell the difference between Bach and RoboBach?

Artificial intelligence has become very good at imitating human composers

<https://www.theverge.com/2016/12/23/14069382/ai-music-creativity-bach-deepbach-csl>



Bach OR computer?



AI generates non-stop stream of death metal

by [ENGADGET RSS FEED](#) on APR 21, 2019

GAMING & CULTURE —

Movie written by algorithm turns out to be hilarious and intense



WATCH

Sunspring | A Sci-Fi Short Film Starring Thomas Middleditch

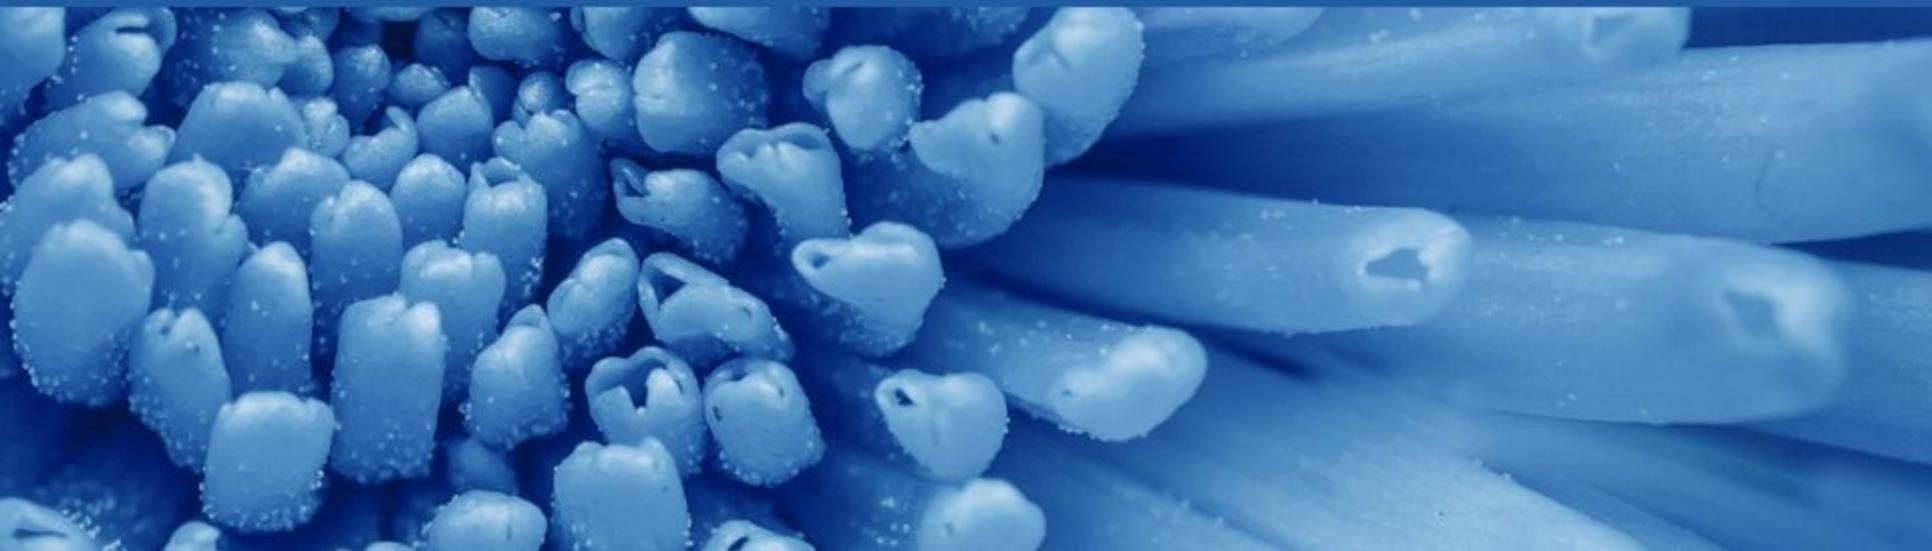
<https://arstechnica.com/gaming/2016/06/an-ai-wrote-this-movie-and-its-strangely-moving/>

What Deep Learning has achieved so far...

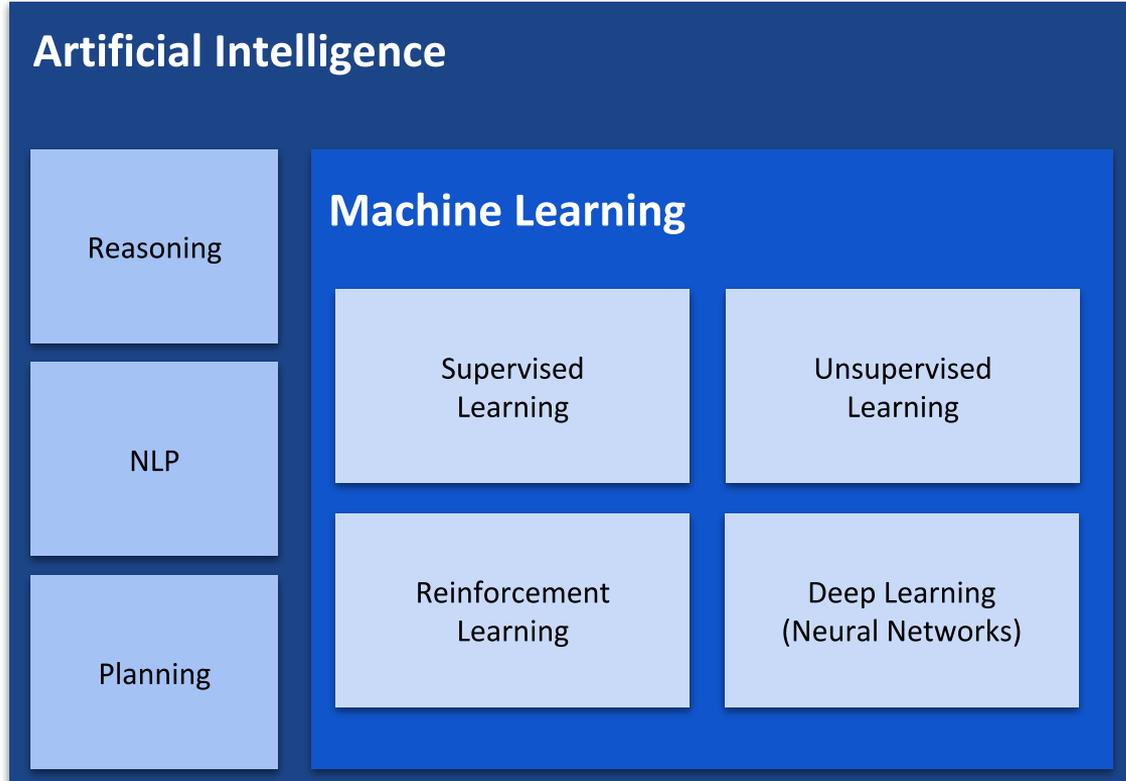
- Near-human to superhuman level **image classification**
- Near-human level **speech recognition**
- Near-human level **handwriting transcription**
- Improved **machine translation**
- Improved **text-to-speech conversion**
- **Digital assistants** such as Google Now or Amazon Alexa
- Near-human level **autonomous driving**
- Superhuman Go playing

**“First rule of magic:
Don't let anyone know your real name.”**

Neil Gaiman, The invisible Labyrinth (1990)



Artificial Intelligence and Machine Learning



“The Goal of AI is to develop machines that behave as though they were intelligent.”

- John McCarthy (1955)

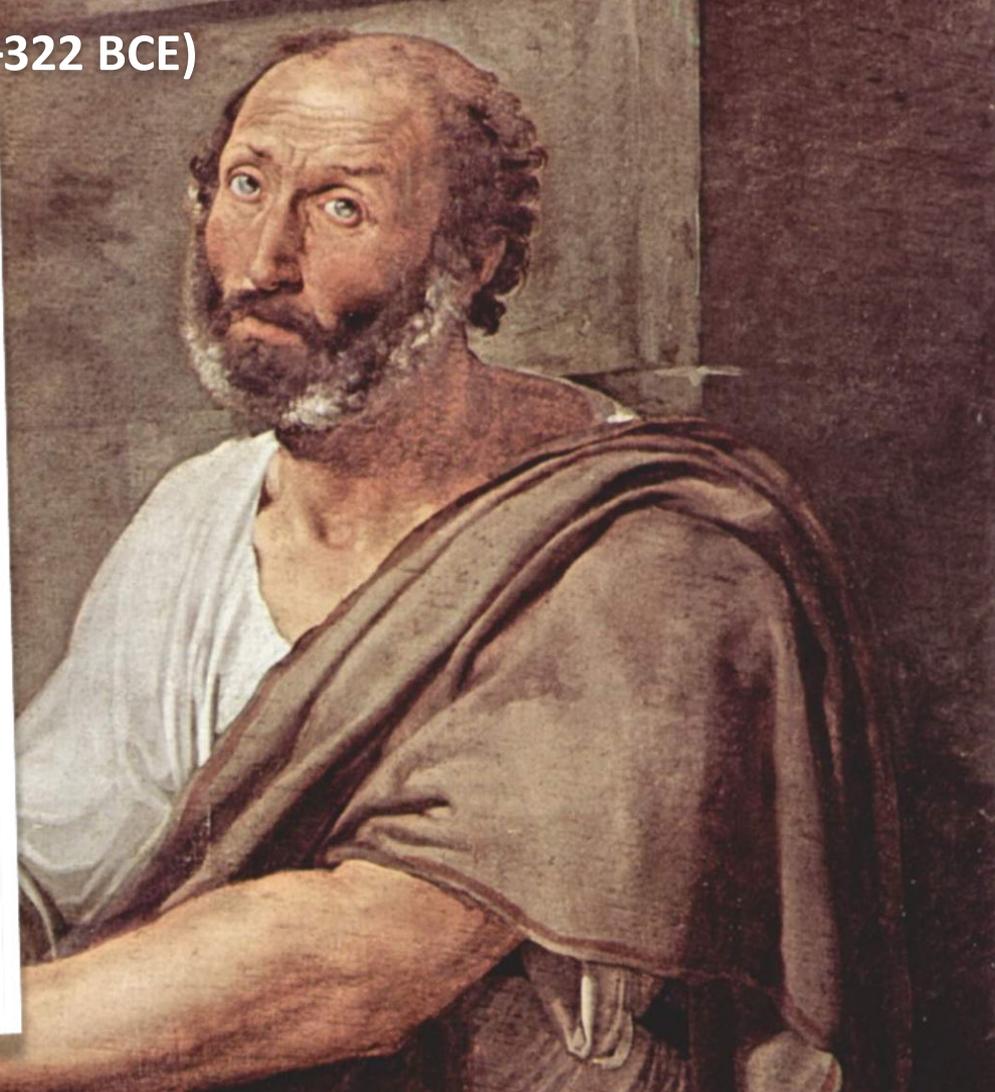
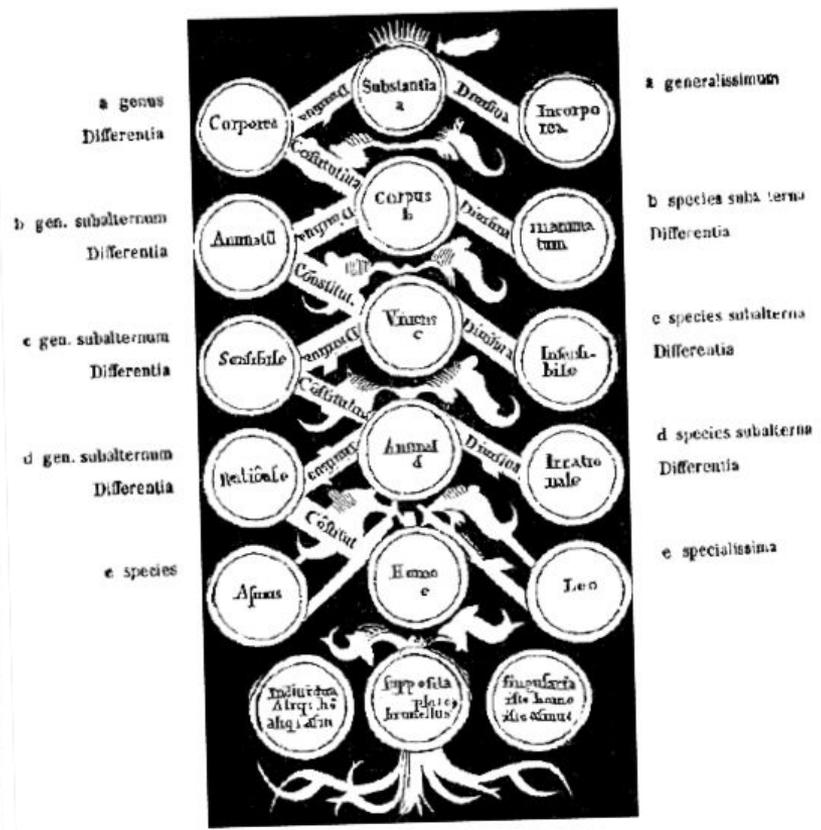
P A R E N T A L

ADVISORY

EXPLICIT SEMANTICS

The Universal Categories - Aristotle (384–322 BCE)

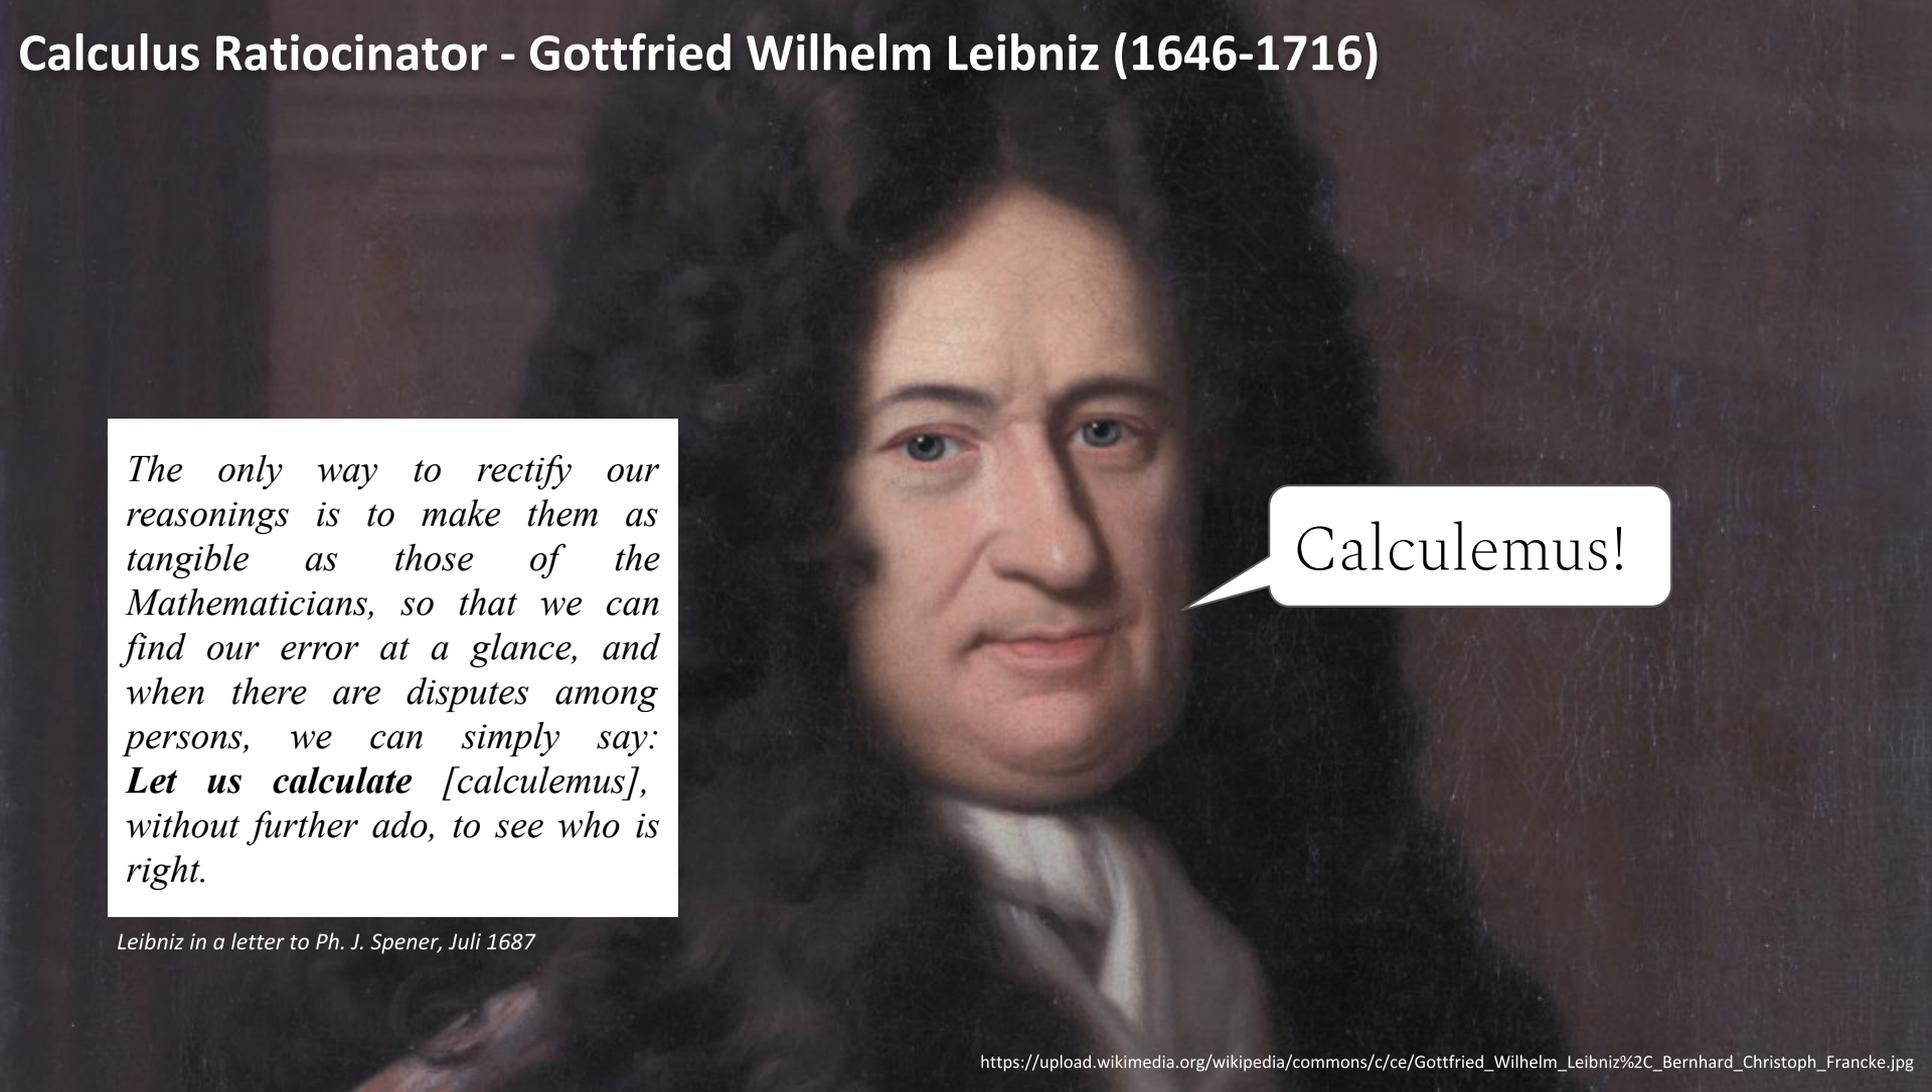
IN PORPHYRIUM DIALOGUS I.



Calculus Ratiocinator - Gottfried Wilhelm Leibniz (1646-1716)

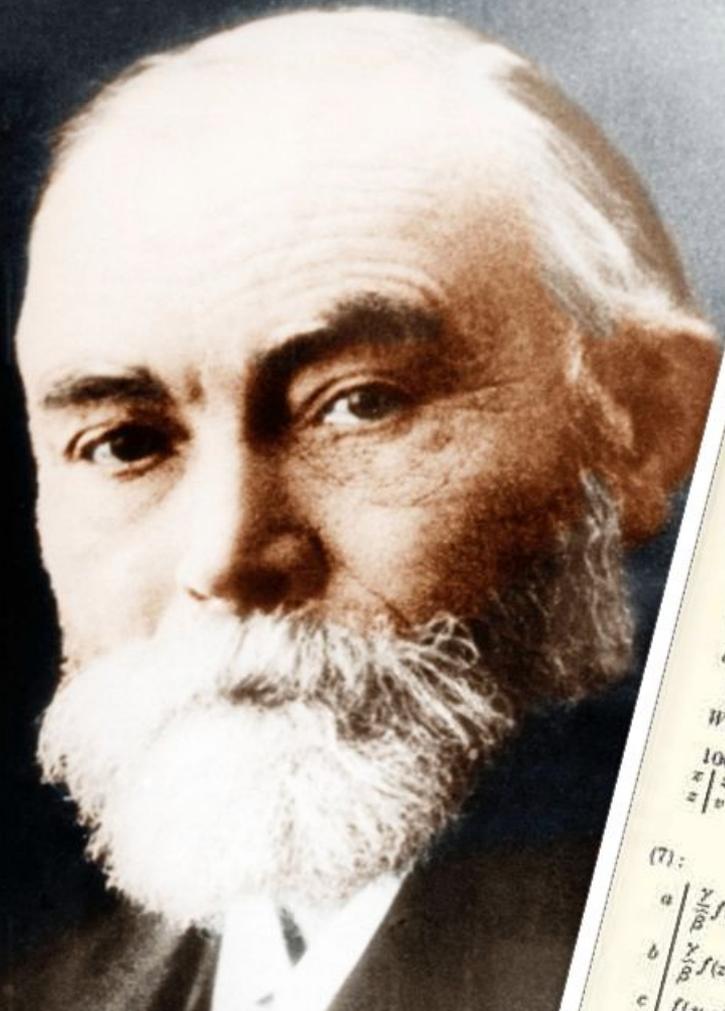
*The only way to rectify our reasonings is to make them as tangible as those of the Mathematicians, so that we can find our error at a glance, and when there are disputes among persons, we can simply say: **Let us calculate** [calculemus], without further ado, to see who is right.*

Leibniz in a letter to Ph. J. Spener, Juli 1687

A portrait of Gottfried Wilhelm Leibniz, a German philosopher, mathematician, and scientist. He is shown from the chest up, wearing a dark coat and a white cravat. He has long, dark, wavy hair and is looking slightly to the right of the viewer with a neutral expression.

Calculemus!

Begriffsschrift - Gottlob Frege (1848-1925)



71

BEGRIFFSSCHRIFT

(55) ::

$$\begin{array}{l} d \mid x \\ c \mid z \end{array}$$

$$\begin{array}{l} \vdash \\ \quad \vdash (x \equiv z) \\ \quad \quad \vdash \frac{\gamma}{\beta} f(x, z) \\ \quad \quad \quad \vdash \frac{\gamma}{\beta} f(x, z) \end{array}$$

§ 30. 99

(52):

$$\begin{array}{l} f(I) \mid \Gamma \\ c \end{array}$$

$$\begin{array}{l} \vdash \\ \quad \vdash (z \equiv x) \\ \quad \quad \vdash \frac{\gamma}{\beta} f(x, z) \\ \quad \quad \quad \vdash \frac{\gamma}{\beta} f(x, z) \end{array}$$

$$\left[\begin{array}{l} \vdash \\ \quad \vdash (z \equiv x) \\ \quad \quad \vdash \frac{\gamma}{\beta} f(x, z) \end{array} \right] \equiv \frac{\gamma}{\beta} f(x, z)$$

(37):

$$\begin{array}{l} a \mid \frac{\gamma}{\beta} f(x, z) \\ b \mid (z \equiv x) \\ c \mid \frac{\gamma}{\beta} f(x, z) \end{array}$$

$$\begin{array}{l} \vdash \\ \quad \vdash \frac{\gamma}{\beta} f(x, z) \\ \quad \quad \vdash (z \equiv x) \\ \quad \quad \quad \vdash \frac{\gamma}{\beta} f(x, z) \end{array}$$

(104).

$$\begin{array}{l} \vdash \\ \quad \vdash \frac{\gamma}{\beta} f(x, z) \\ \quad \quad \vdash \frac{\gamma}{\beta} f(x, z) \end{array}$$

(105).

Whatever follows x in the f -sequence belongs to the f -sequence beginning with x .

$$\begin{array}{l} 106 \\ x \mid z \\ z \mid v \end{array}$$

$$\begin{array}{l} \vdash \\ \quad \vdash \frac{\gamma}{\beta} f(z, v) \\ \quad \quad \vdash \frac{\gamma}{\beta} f(z, v) \end{array}$$

(7):

$$\begin{array}{l} a \mid \frac{\gamma}{\beta} f(z, v) \\ b \mid \frac{\gamma}{\beta} f(z, v) \\ c \mid f \dots \end{array}$$

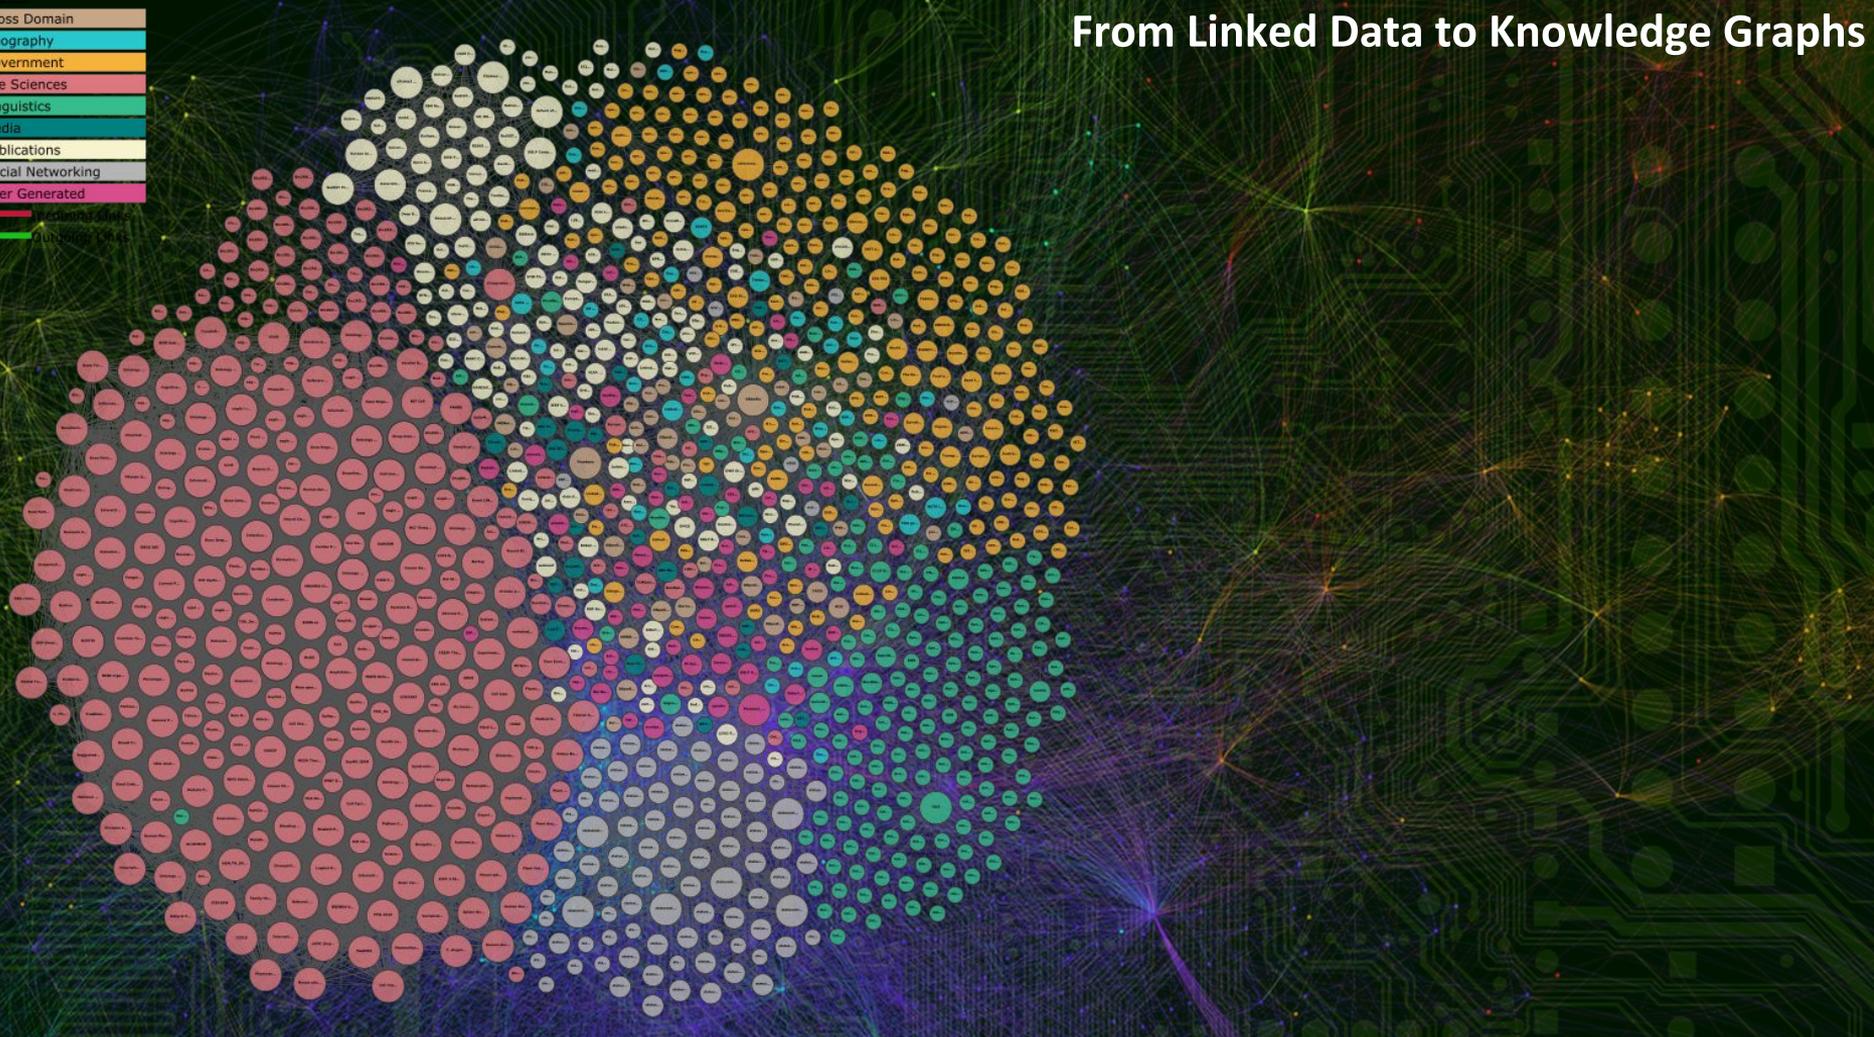
$$\begin{array}{l} \vdash \\ \quad \vdash \frac{\gamma}{\beta} f(z, v) \end{array}$$

(106).

Frame Logic for Knowledge Representation - Marvin Minsky (1974)



From Linked Data to Knowledge Graphs



Knowledge Graph Applications



T. Tietz, J. Jäger, J. Waitelonis, H. Sack, *Semantic Annotation and Information Visualization for Blogposts with refer*, (VOILA 2016)

Relation Browser Timeline

Data Integration

Visualizations

Semantic & Exploratory Search

Recommender Systems

Question Answering

- ### 15 Recommended Articles
- #1 The Case of J. Robert Oppenheimer
 - #2 Wilhelm Pfeffer and Plant Physiology
 - #3 Karl Pearson and Mathematical Statistics
 - #4 Raphael and the School of Athens
 - #5 Jerzy Neyman – Architect of Modern Theoretical Statistics
 - #6 Christian Gottfried Ehrenberg – Father of Microzoology

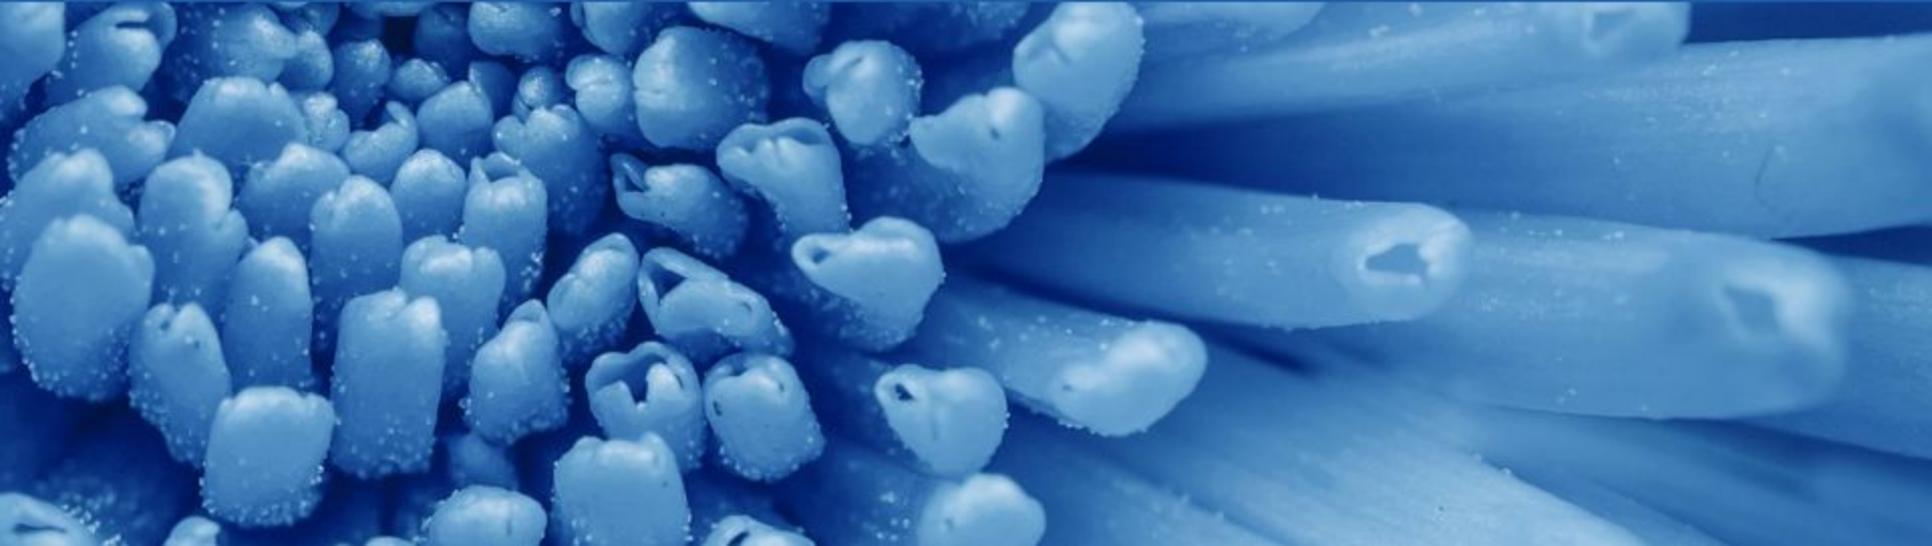
Gottfried Wilhelm Leibniz

Gottfried Wilhelm von Leibniz (German: [ˈɡɔʦfʁiːt ˈvɪlhɛlm fɔn ˈlaɪbnɪts] or [ˈlaɪpnɪts]) (July 1, 1646 – November 14, 1716) was a German mathematician and philosopher. He occupies a prominent place in the history of mathematics and the history of philosophy. Leibniz developed the infinitesimal calculus independently of Isaac Newton, and Leibniz's mathematical notation has been widely used ever since it was published. It was only in the 20th century that his Law of Continuity and Transcendental Law of Homogeneity found mathematical implementation (by means of non-standard analysis). He became one of the most prolific inventors in the field of mechanical calculators. While working on adding automatic multiplication and division to Pascal's calculator, he was the first to describe a pinwheel calculator in 1685 and invented the Leibniz wheel, used

DBpedia: Gottfried Wilhelm Leibniz

“There ain’t no such thing as a free Lunch.”

Robert A. Heinlein, The Moon is a Harsh Mistress (1966)

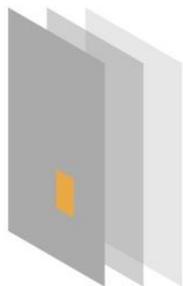


The Clever Hans Effect

or Why we shouldn't always trust ML

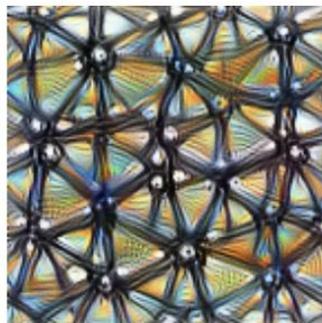


What do Neural Networks Really See?



Neuron

`layern[x,y,z]`



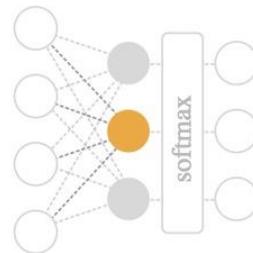
Channel

`layern[:, :, z]`



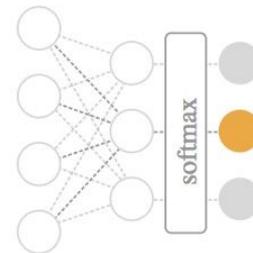
Layer/DeepDream

`layern[:, :, :]2`



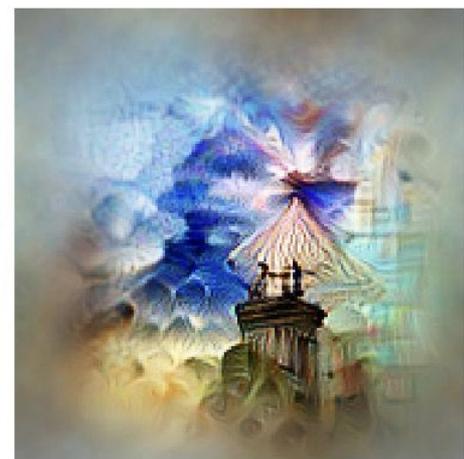
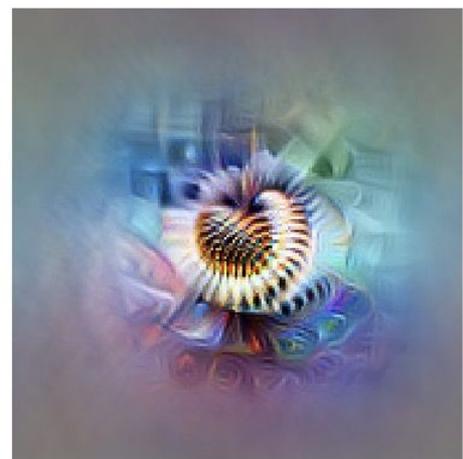
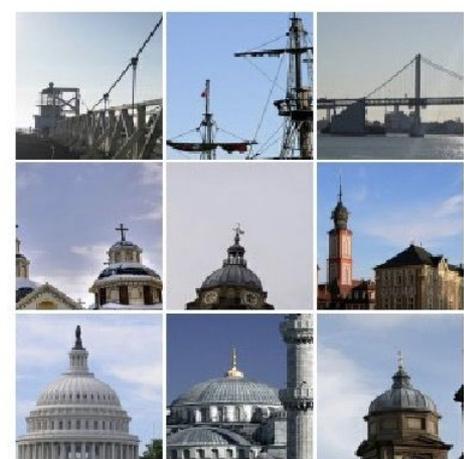
Class Logits

`pre_softmax[k]`



Class Probability

`softmax[k]`



Baseball—or stripes?

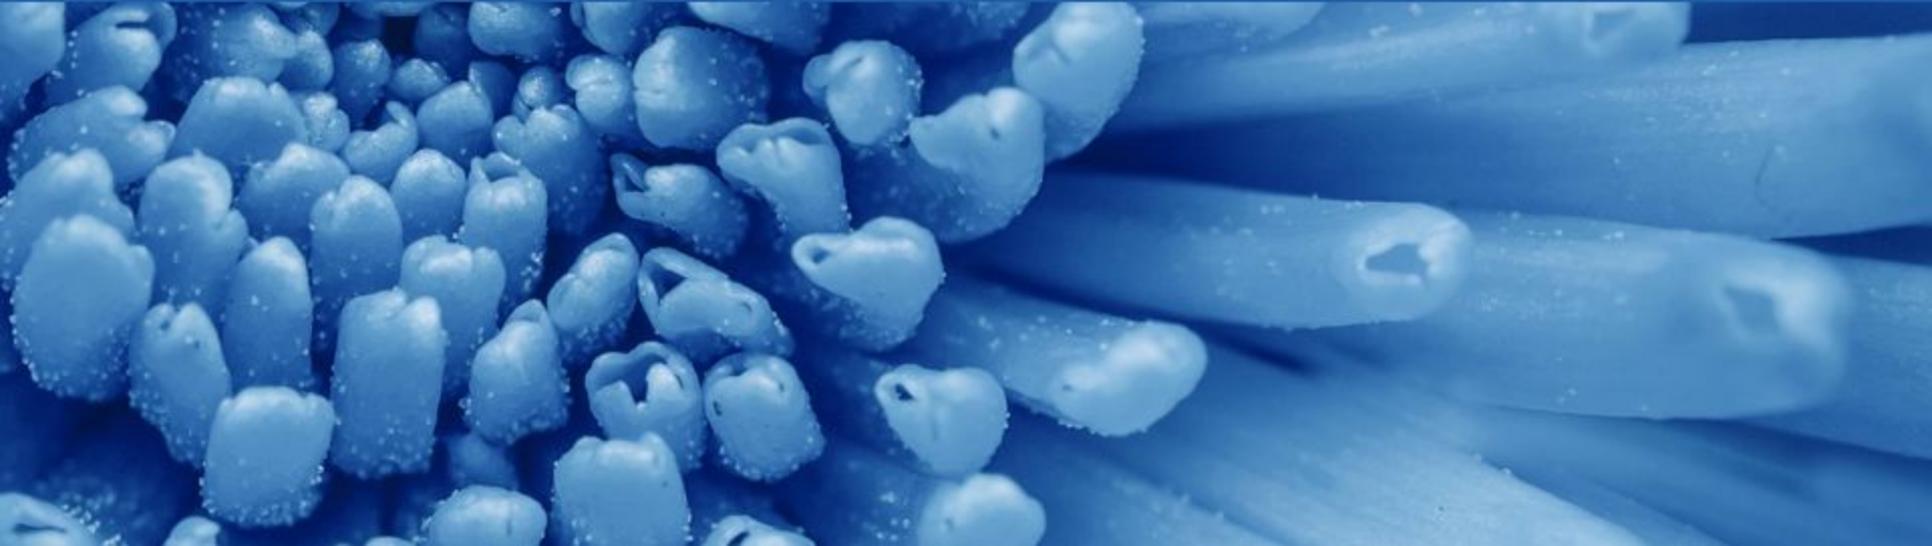
Animal faces—or snouts?

Clouds—or fluffiness?

Buildings—or sky?

“90% of most Magic merely consists of knowing one extra fact.”

Terry Pratchett, Night Watch (2002)



Generative Adversarial Networks - Example

Text description

This bird is blue with white and has a very short beak

This bird has wings that are brown and has a yellow belly

A white bird with a black crown and yellow beak

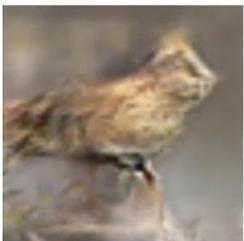
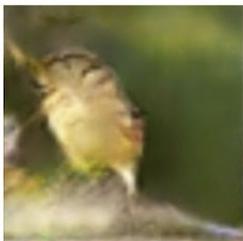
This bird is white, black, and brown in color, with a brown beak

The bird has small beak, with reddish brown crown and gray belly

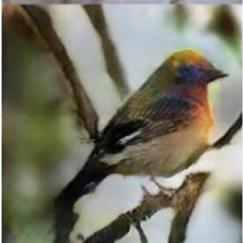
This is a small, black bird with a white breast and white on the wingbars.

This bird is white black and yellow in color, with a short black beak

Stage-I images



Stage-II images



Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, Dimitris N. Metaxas:

[StackGAN++: Realistic Image Synthesis with Stacked Generative Adversarial Networks.](#)

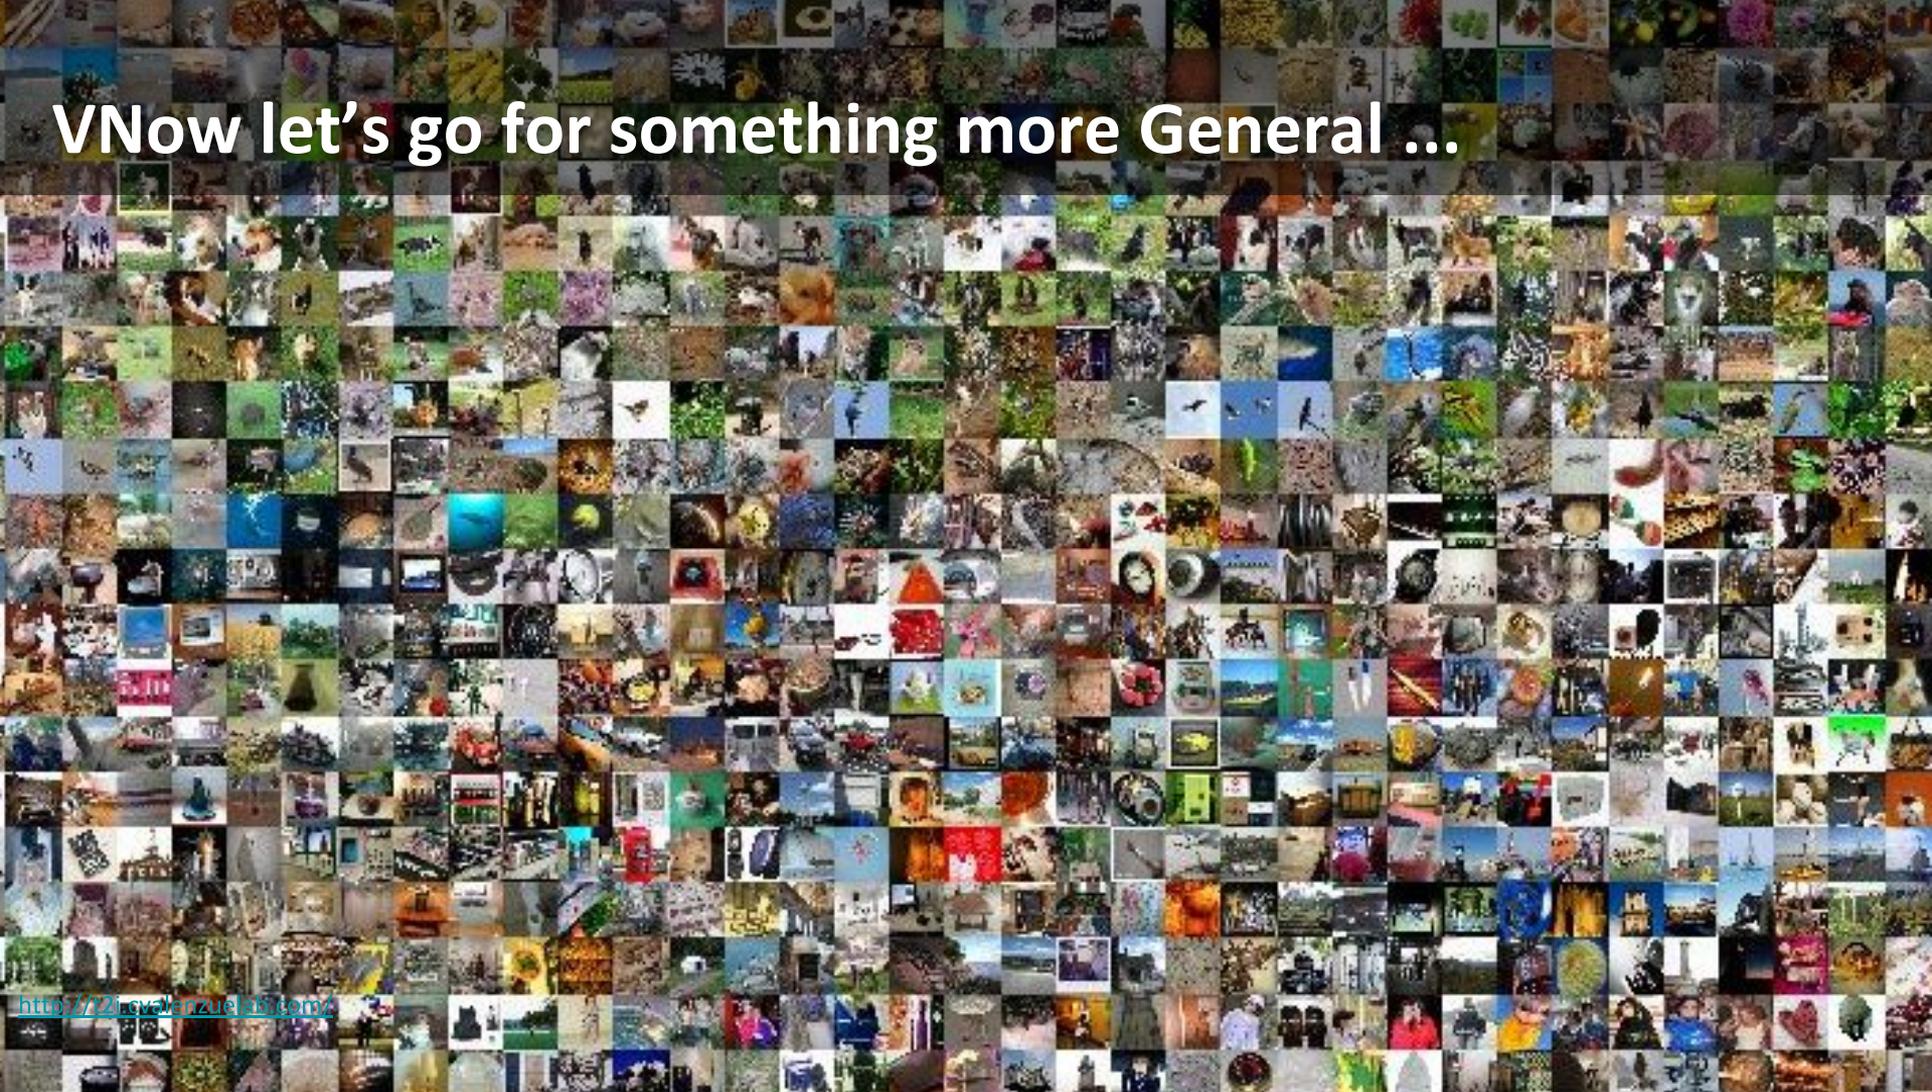
CoRR abs/1710.10916 (2017)

Generative Adversarial Networks - Example

This is a small light gray bird with a small head and green crown nape and some green coloring on its wings



VNow let's go for something more General ...

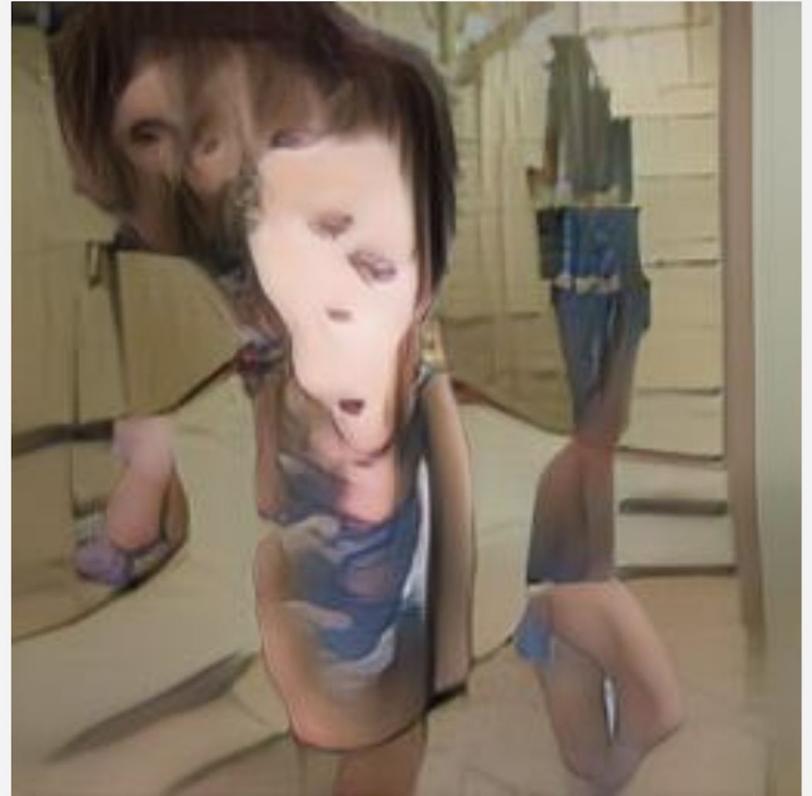


A flock of sheep on green meadows



<http://t2i.cvalenzuelab.com/>

a girl
watching tv



<http://t2i.cvalenzuelab.com/>

a girl with
two blue eyes,
one upturned
nose, and a
red mouth

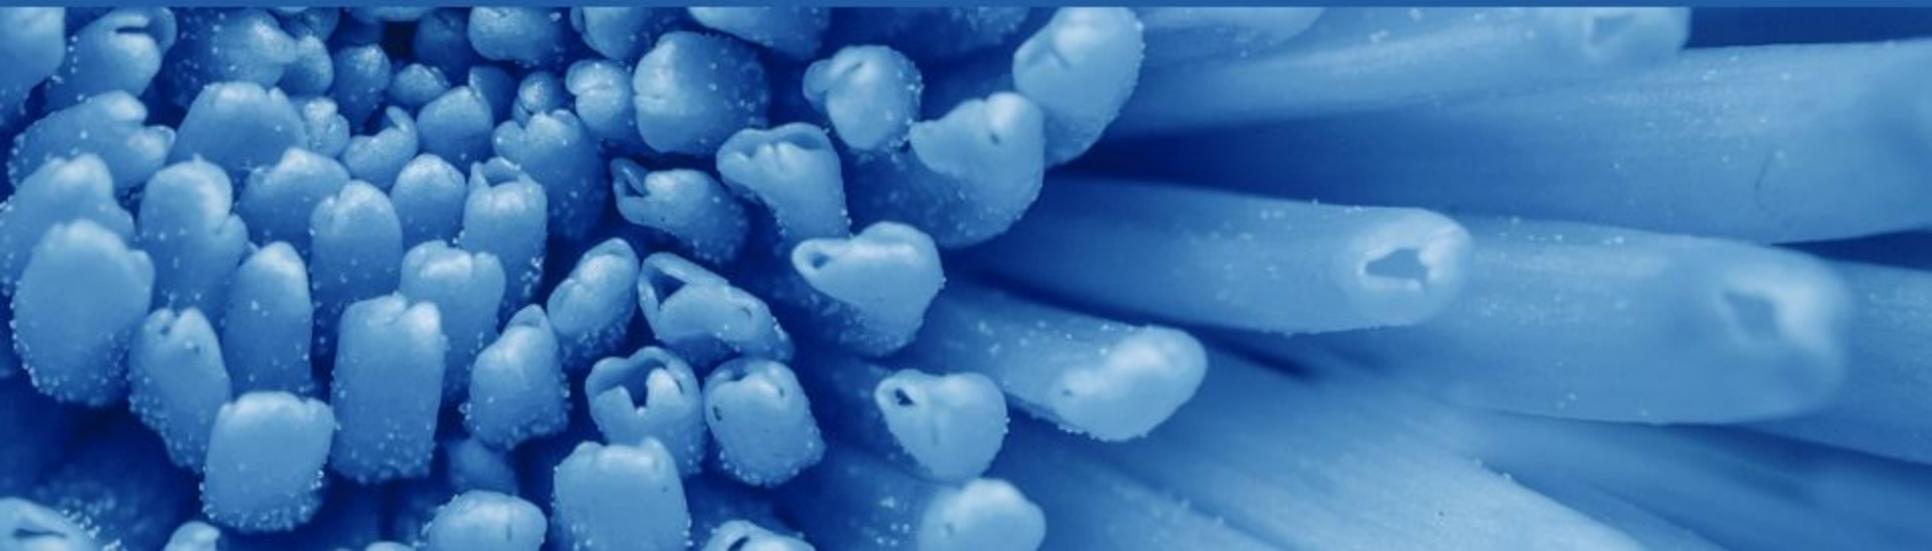


<http://t2i.cvalenzuelab.com/>

melting clocks over the dessert

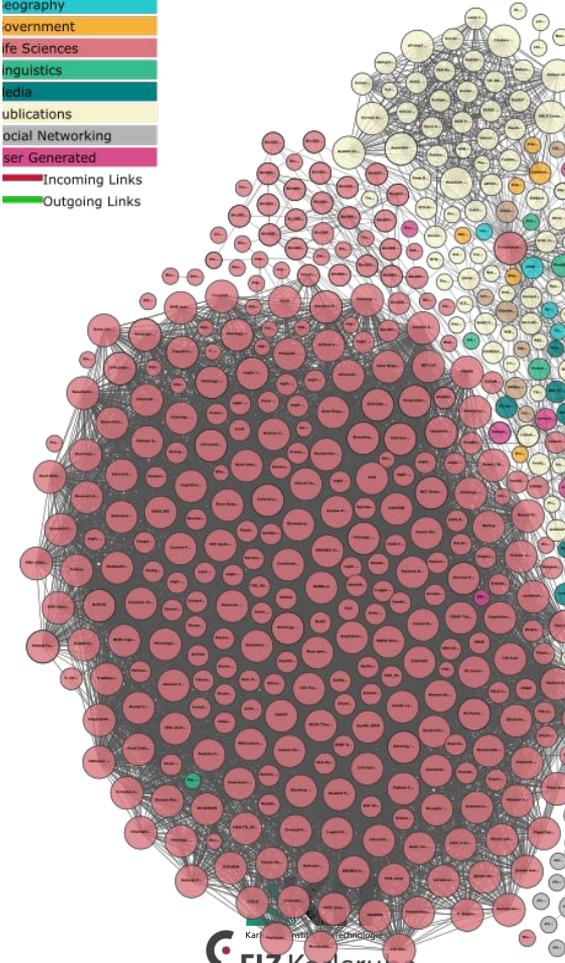


Quo vadis AI?



Deep Learning for Knowledge Graphs

- NLP and Knowledge Extraction via Deep Learning to **populate and extend Knowledge Graphs**
- NLP and Knowledge Extraction via Deep Learning for **Ontology Learning** to **extend and refine Knowledge Graphs**
- NLP and Graph Analysis supported by Deep Learning for **Ontology Alignment** and **Link Discovery** to **combine and integrate Knowledge Graphs**



Neil Armstrong

From Wikipedia, the free encyclopedia

For other people named Neil Armstrong, see *Neil Armstrong (disambiguation)*.

Neil Alden Armstrong (August 5, 1930 – August 25, 2012) was an American astronaut and the first person to walk on the Moon.

his bachelor's degree at [Purdue University](#) and served as a test pilot at the [National Advisory Committee for Aeronautics \(NACA\) High-Speed Flight Station](#), where he logged over 900 flights. He later completed graduate studies at the [University of Southern California](#).

A participant in the [U.S. Air Force](#), he made his first space flight as command pilot [David Scott](#).^[1] This mission was a space emergency.

Armstrong's second and last space flight was [Apollo 11](#), where he descended to the lunar surface with [Collins](#) and [Aldrin](#), Armstrong was awarded the [Presidential Medal of Freedom](#) by [President Richard Nixon](#). [President Jimmy Carter](#) presented Armstrong the [Congressional Space Medal of Honor](#) in 1978. Armstrong and his former crewmates received the [Congressional Gold Medal](#) in 2009.

Armstrong died in [Cincinnati, Ohio](#) on August 25, 2012, at the age of 82, after complications from [coronary artery bypass surgery](#).^{[2][3]}

```
dbr:Neil_Armstrong rdfs:label "Neil Alden Armstrong"@en ;
dbo:birthDate "1930-08-05"^^xsd:date ;
dbo:deathDate "2012-08-25"^^xsd:date ;
dbo:occupation dbo:Astronaut .
```

ps in 1962. He made two spacecraft, with er, in the first in-flight role pilot [Buzz Aldrin](#) vice Module. Along



Armstrong in July 1969



Contents [\[hide\]](#)

- 1 Early years
- 2 Navy service
- 3 College years
- 4 Test pilot
- 5 Astronaut career
 - 5.1 Gemini program
 - 5.1.1 *Gemini 8*
 - 5.1.2 Gemini 11
 - 5.2 Apollo program
 - 5.2.1 Apollo 11
 - 5.2.1.1 *Voyage to the Moon*
 - 5.2.1.2 First Moon walk
 - 5.2.1.3 Return to Earth
- 6 Life after Apollo
 - 6.1 Teaching

```
dbr:Neil_Armstrong rdfs:label "Neil Alden Armstrong"@en ;
dbo:birthDate "1930-08-05"^^xsd:date ;
dbo:birthPlace dbr:Wapakoneta,Ohio ;
dbo:deathDate "2012-08-25"^^xsd:date ;
dbo:deathPlace dbr:Cincinnati,Ohio ;
dbo:occupation dbo:Astronaut ;
dbo:almaMater dbr:Purdue_University ,
dbr:University_of_Southern_California .
```

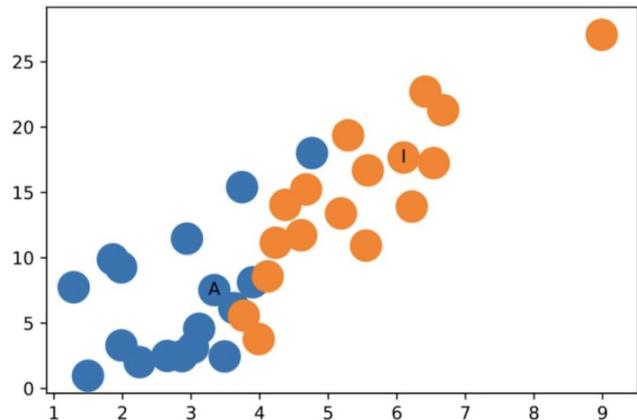
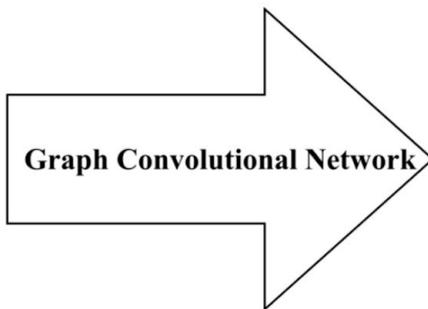
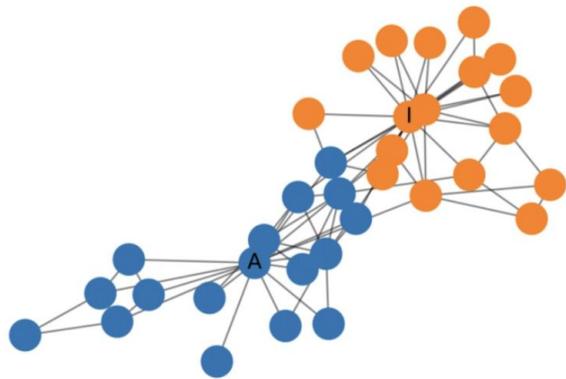
...

USAF / NASA astronaut

Born	Neil Alden Armstrong August 5, 1930 Near Wapakoneta, Ohio , U.S.
Died	August 25, 2012 (aged 82) Cincinnati, Ohio , U.S.
Previous occupation	Naval aviator, test pilot
Alma mater	Purdue University , B.S. 1955 University of Southern California , M.S. 1970
Rank	Lieutenant (junior grade), United States Navy

Knowledge Graphs for Deep Learning?

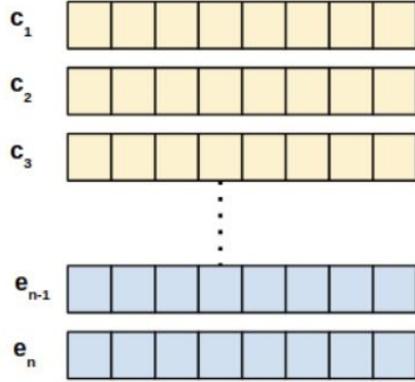
- Use **Graph Embeddings** for a latent semantic representation of **Knowledge Graphs**
- Combining latent semantic representations of **different (symbolic) representations (Hybrid Embeddings)**
- **Graph Convolutional Neural Networks**



input: "IBM adds midrange server to eServer lineup"



Entity and Category Embeddings



IBM adds midrange server to eServer lineup

Mention Detection Based on Anchor-Text Dictionary

IBM adds midrange server to eServer lineup
 m_1 m_2 m_3

Candidate Generation

IBM

midrange server

eServer

e_1 : [IBM](#)
 e_2 : [IBM_Notes](#)
 e_3 : [IBM_AIX](#)

e_1' : [Midrange_computer](#)
 e_2' : [Mid-range_speaker](#)
 e_3' : [Mid-range](#)

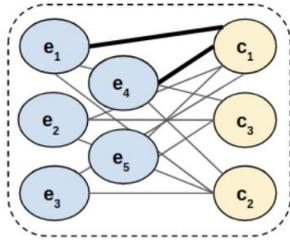
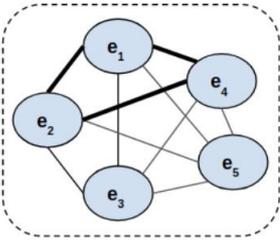
e''_1 : [IBM_eServer](#)
 e''_2 : [Server](#)
 e''_3 : [Web_server](#)

Category List
{Sports, Technology, Culture, World}



Entity-Entity Network

Entity-Category Network



Heterogeneous Entity and Category Network

Probabilistic Model

Technology

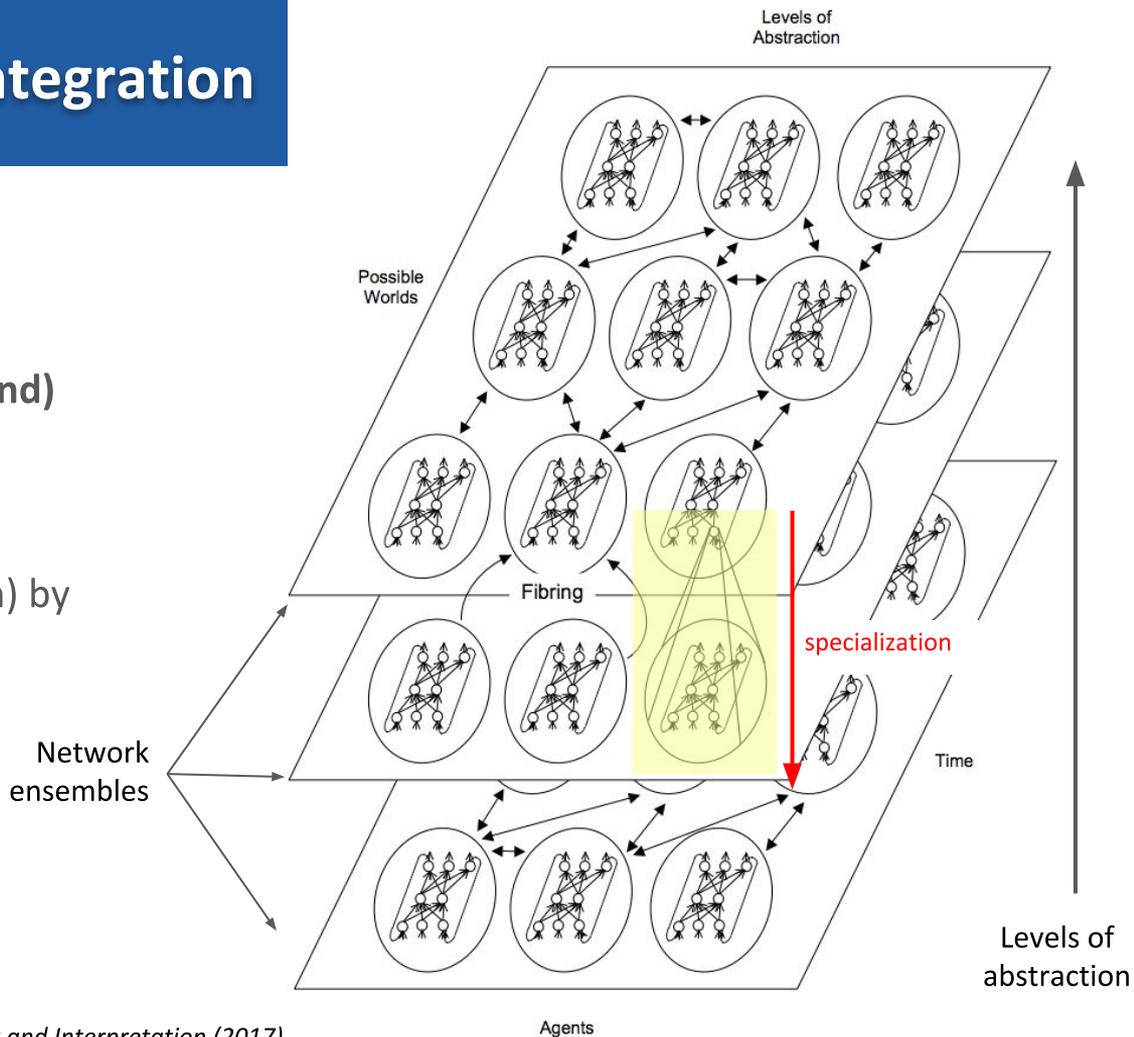
output:

R. Türker, L. Zhang, M. Koutraki, H. Sack:
Knowledge-Based Short Text Categorization Using Entity
and Category Embedding (ESWC 2019)

Towards Neuro-Symbolic Integration

Neuro-Symbolic Systems

1. Translation of **symbolic (background) knowledge into the network**
2. **Learning of additional knowledge** from examples (and generalisation) by the network
3. **Executing the network** (i.e. reasoning), and
4. **Symbolic knowledge extraction** from the network.



Model	AG (title)	AG (title+description)	Snippets
Dataless ESA [14]	53.5	64.1	48.5
Dataless Word2Vec [14]	49.5	52.7	52.4
NB+TF-IDF	86.6	90.2	64.4
SVM+TF-IDF	87.6	91.9	69.1
LR+TF-IDF	87.1	91.7	63.6
KBSTC+Our Embedding	67.9	80.5	72.0

Short Term Goals to keep in Mind...

- **Knowledge Graphs:**
Don't try to make Everything Explicit
- **Deep Learning:**
Try to make the Implicit Explicit
- otherwise.....

otherwise...

**“All those moments will be lost in time,
like tears in the rain.”** *(Bladerunner, 1982)*

Prof. Dr. Harald Sack

*The Sorcerer's Apprentice or
Why Semantics still matters and
why we should not only trust in ML*

harald.sack@fiz-karlsruhe.de

twitter: [lysander07](https://twitter.com/lysander07)

Sustainable Knowledge Graphs and AI
Gothenburg, 26. Apr. 2019





SEMANTiCS 2019

The Power of AI and Knowledge Graphs

Karlsruhe, Germany, 9-12 Sep 2019