

Do Neural Networks Dream of Semantics?

“Any
sufficiently advanced
technology is
indistinguishable
from magic.”

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



Can you find the cancer?

Google AI Beats Doctors at Breast Cancer Detection—Sometimes

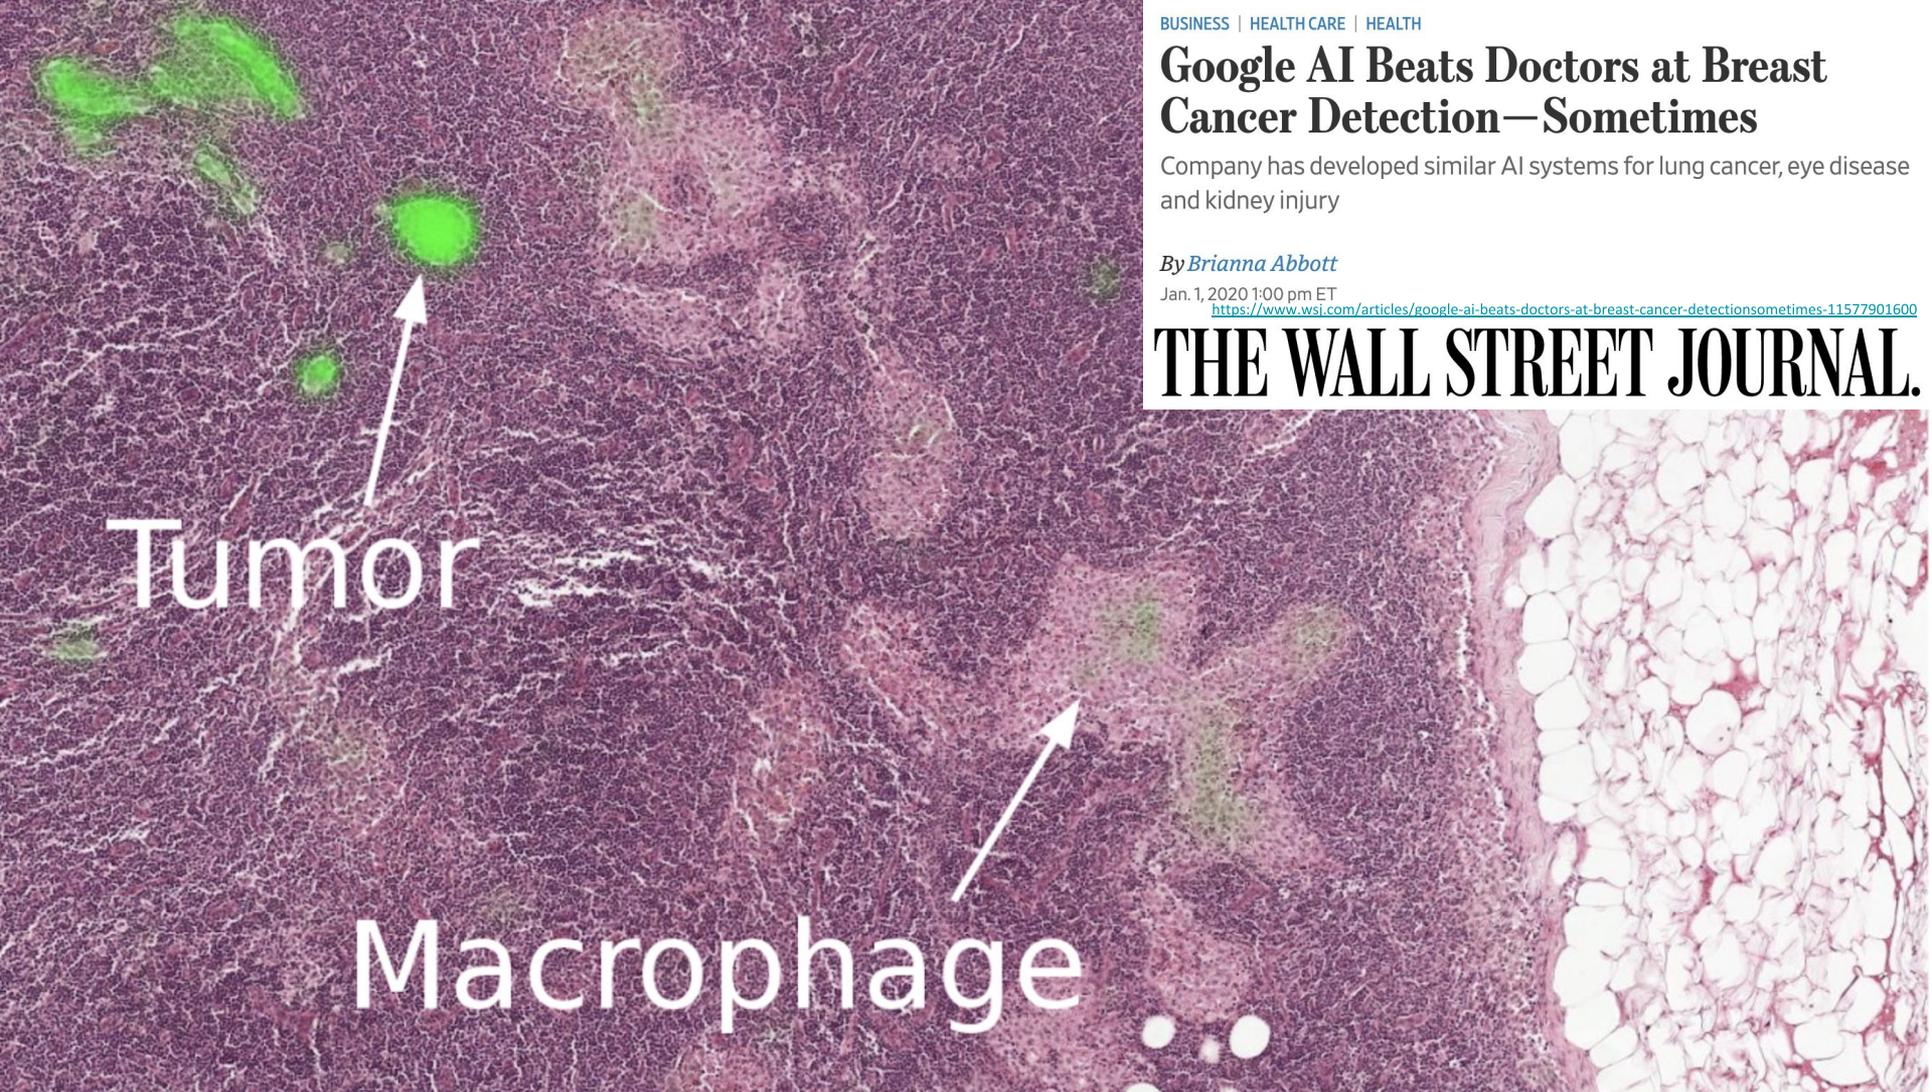
Company has developed similar AI systems for lung cancer, eye disease and kidney injury

By *Brianna Abbott*

Jan. 1, 2020 1:00 pm ET

<https://www.wsj.com/articles/google-ai-beats-doctors-at-breast-cancer-detection-sometimes-11577901600>

THE WALL STREET JOURNAL.



Tumor

Macrophage

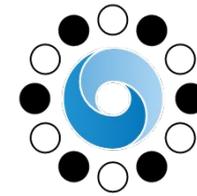
AlphaGo Zero: Google DeepMind supercomputer learns 3,000 years of human knowledge in 40 days



The Telegraph



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

Machine learning has been used to automatically translate long-lost languages

Some languages that have never been deciphered could be the next ones to get the machine translation treatment.

<https://www.technologyreview.com/s/613899/machine-learning-has-been-used-to-automatically-translate-long-lost-languages/>

by **Emerging Technology from the arXiv**

Jul 1, 2019

"...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...?

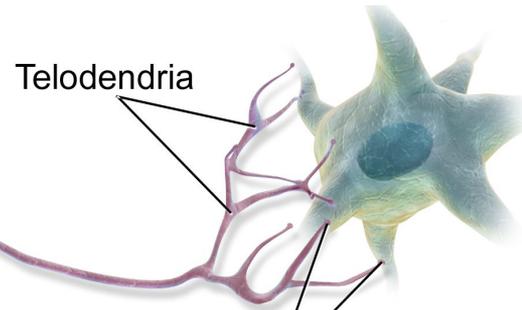
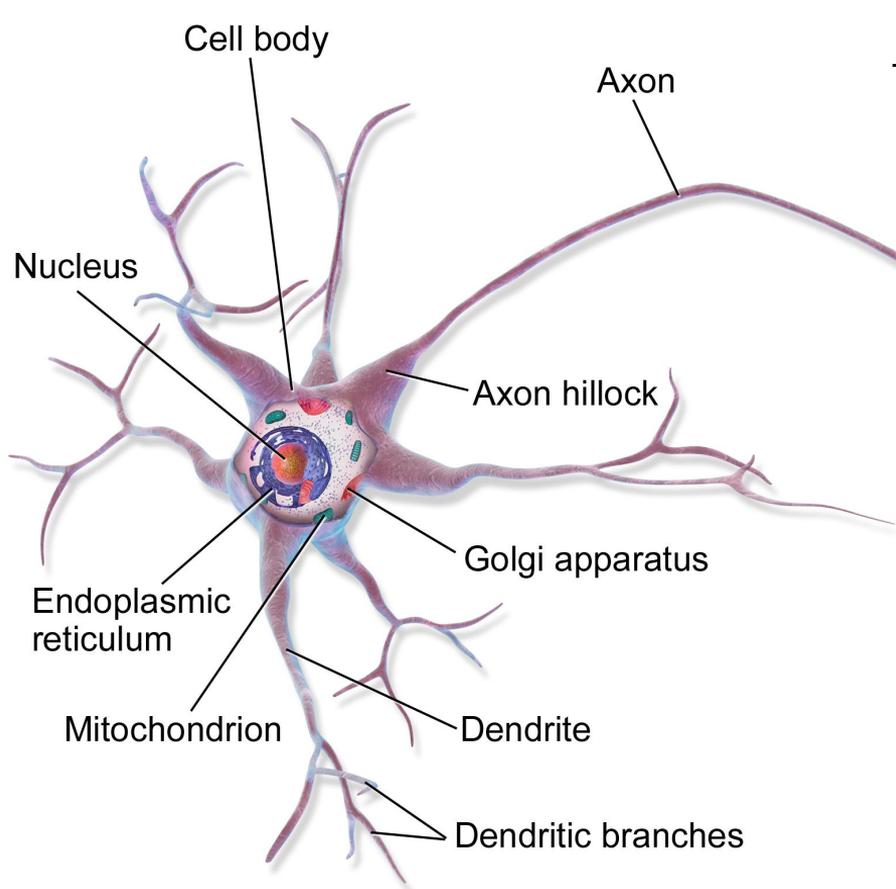
...or do we simply have a tendency to overestimate technology?

“Nothing succeeds
like success.”

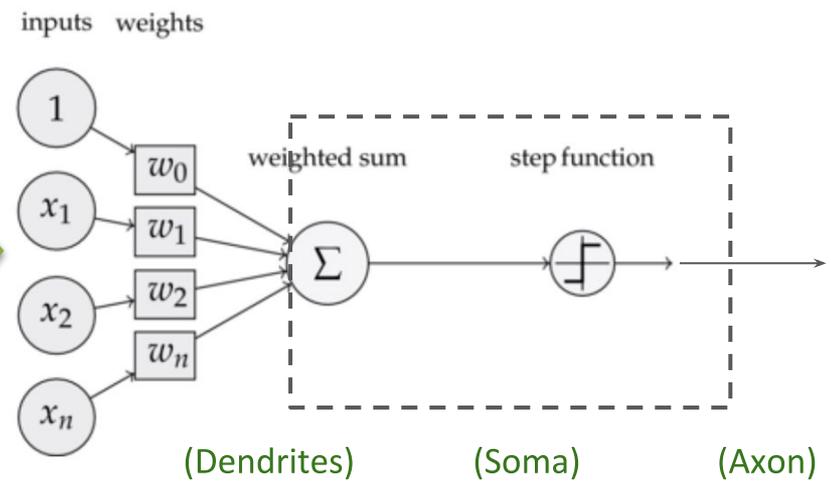
Alexandre Dumas,
Ange Pitou, Vol. 1,
(1854)



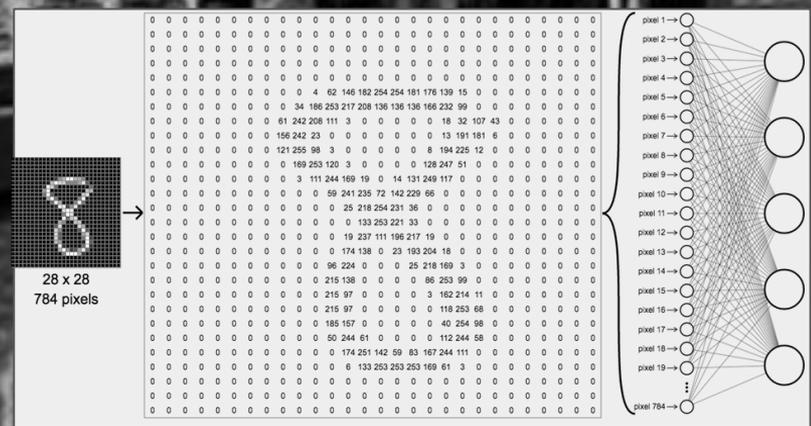
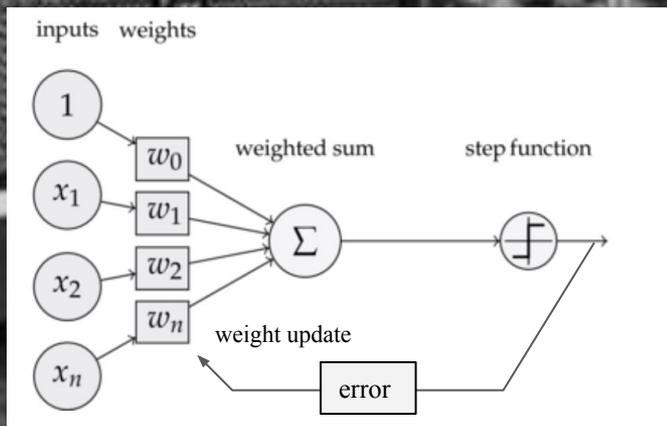
Transferring Biology into a Mathematical Model



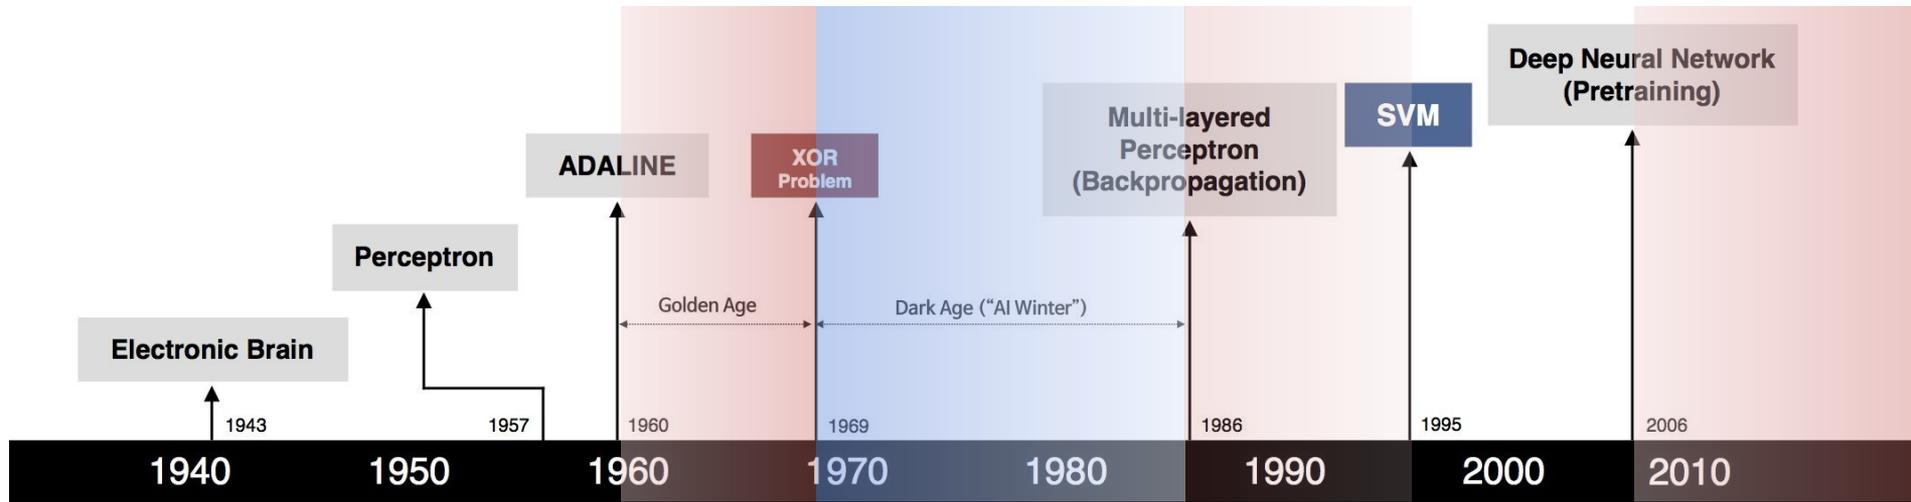
Donald Hebb (1940)
McCulloch & Pitts (1943)



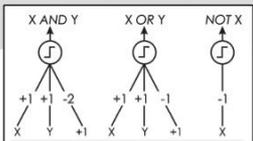
Cognitive Computing - The MARK 1 Perceptron (Rosenblatt, 1957)



The Triumph of the Connectionist Paradigm



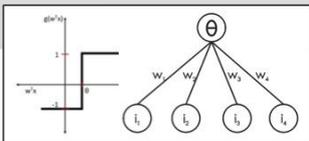
S. McCulloch – W. Pitts



- Adjustable Weights
- Weights are not Learned



F. Rosenblatt



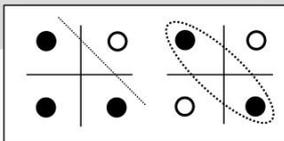
- Learnable Weights and Threshold



B. Widrow – M. Hoff



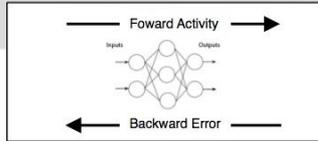
M. Minsky – S. Papert



- XOR Problem



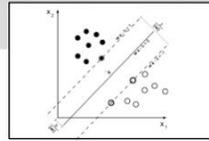
D. Rumelhart – G. Hinton – R. Williams



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



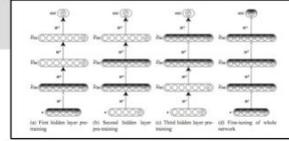
V. Vapnik – C. Cortes



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



G. Hinton – S. Ruslan



- Hierarchical feature Learning

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

Terry Pratchett,
Night Watch (2002)

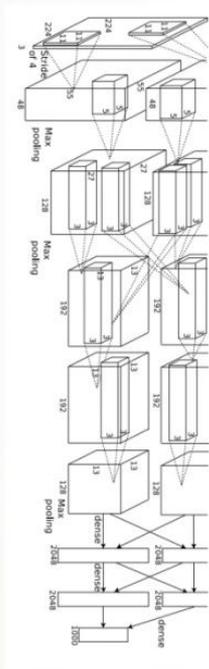
Why has Deep Learning become so successful?



(1) Availability of cheap computing capacity with GPUs

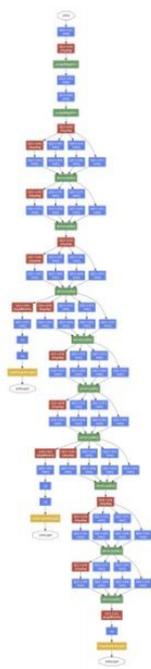
(2) Reusable Highly Trained Complex Models

“AlexNet”



[Krizhevsky et al. NIPS 2012]

“GoogLeNet”



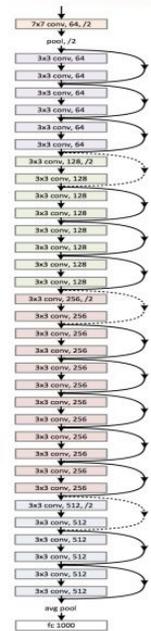
[Szegedy et al. CVPR 2015]

“VGG Net”



[Simonyan & Zisserman, ICLR 2015]

“ResNet”



[He et al. CVPR 2016]

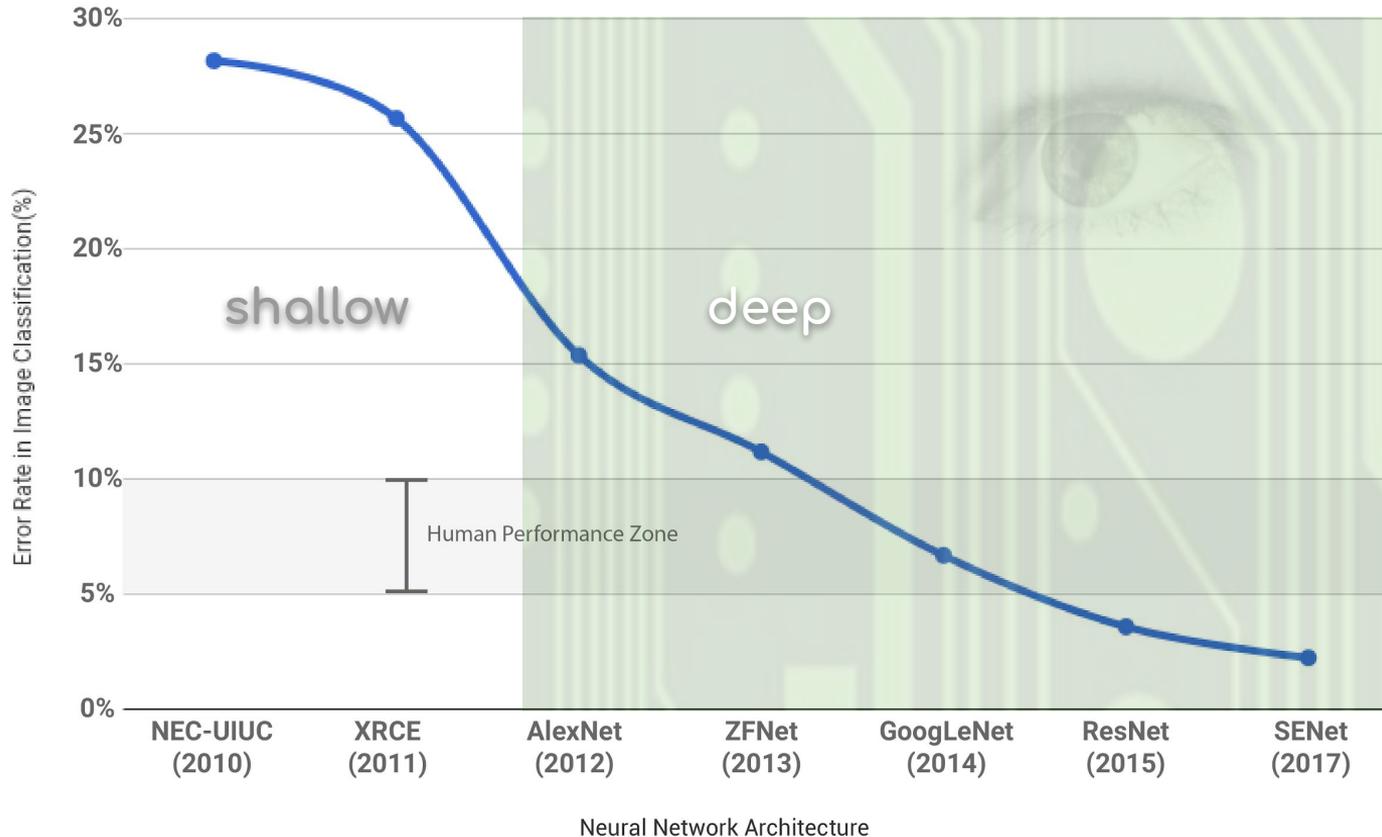


***“First, we find that the performance on vision tasks
still increases linearly with orders of magnitude of
training data size.”***

C. Sun et al, Revisiting Unreasonable Effectiveness of Data in Deep Learning Era, 2017

(3) Availability of Large Scale Labeled Data

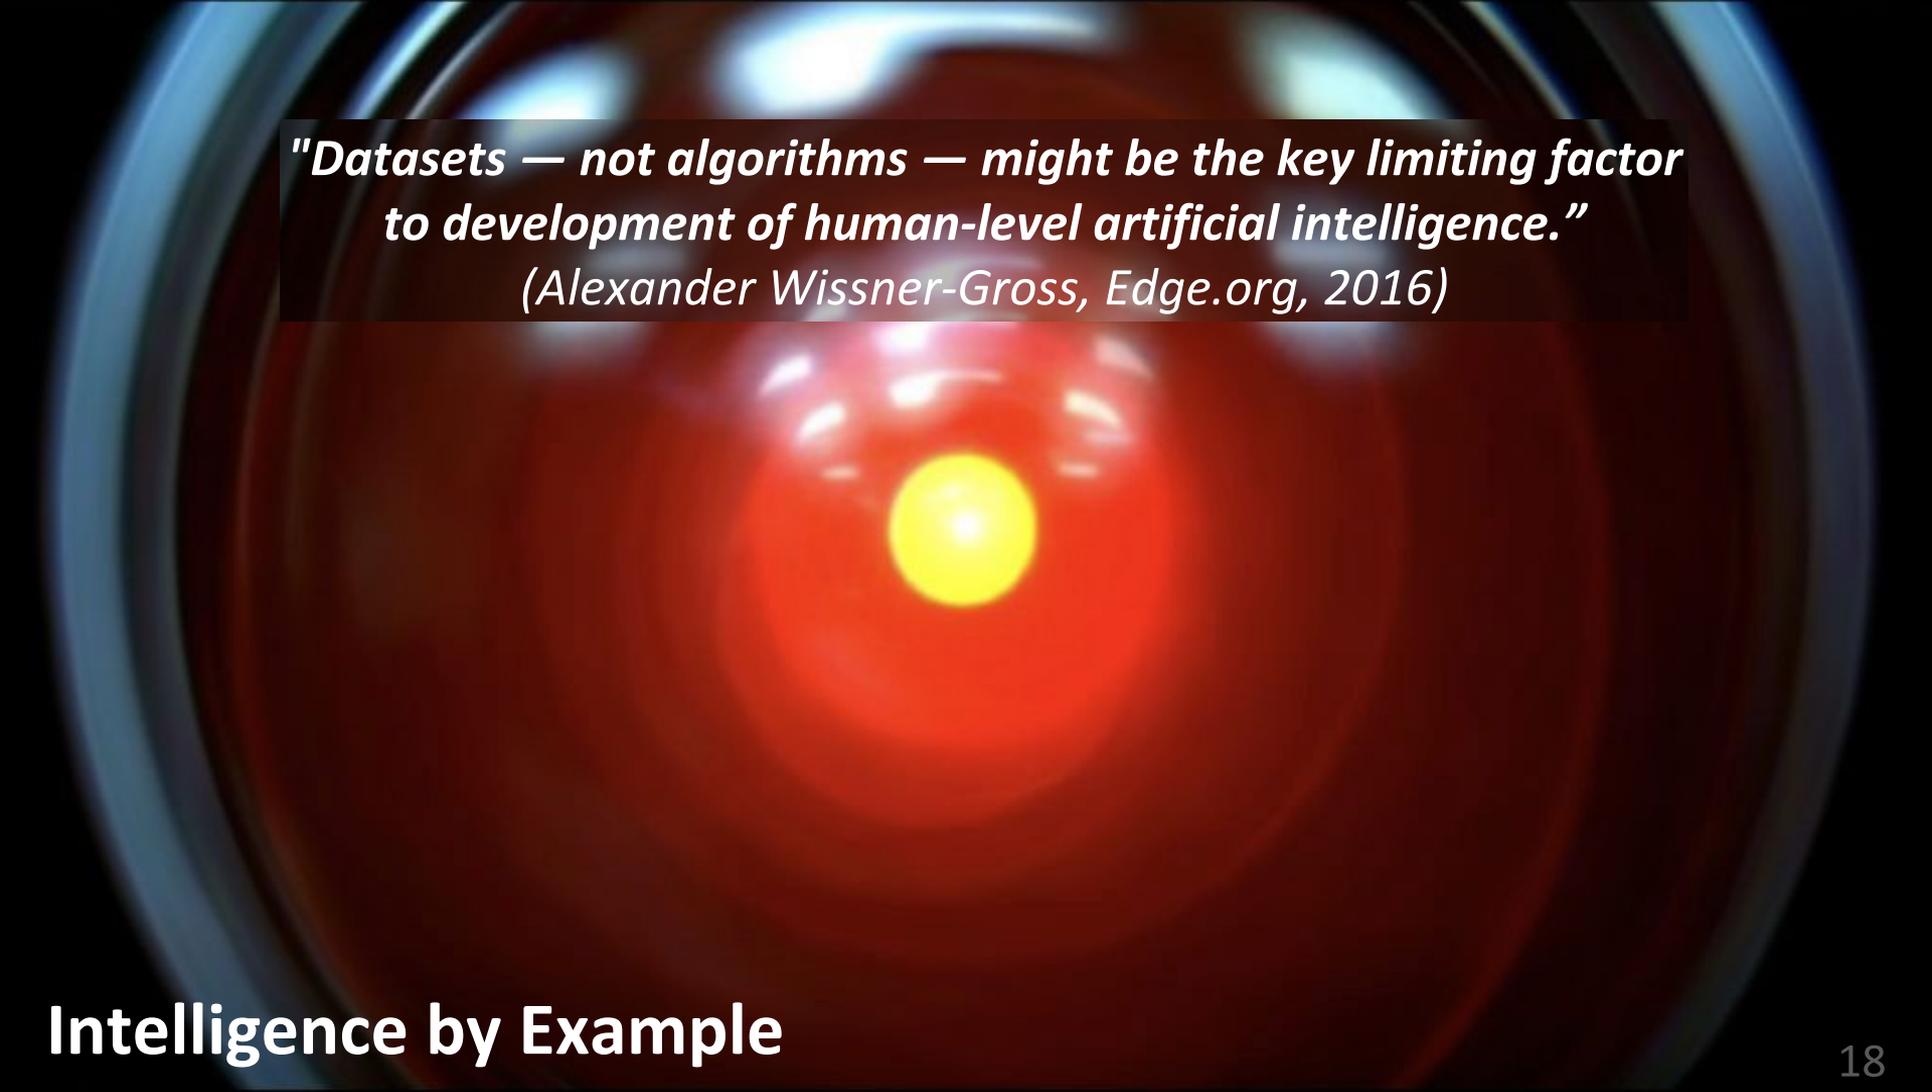
Deep Learning for Visual Analysis



IMAGENET

Large Scale Visual Recognition
Challenge (ILSVRC)

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



***"Datasets — not algorithms — might be the key limiting factor to development of human-level artificial intelligence."
(Alexander Wissner-Gross, Edge.org, 2016)***

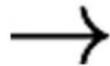
“Creation
is never complete.
It started once,
but it will never stop.”

Immanuel Kant,
General Natural History
and Theory of Heaven (1755)

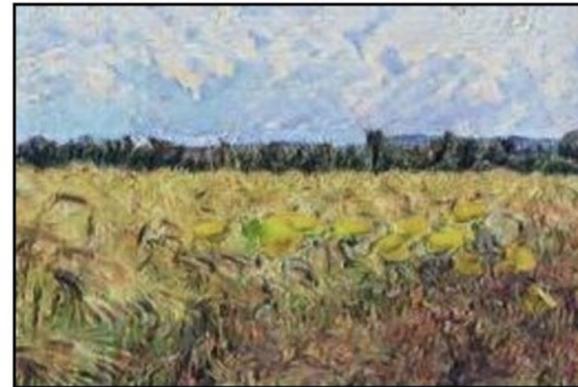
Creation of New Content - Cross Domain Transfer



Photograph



Monet



Van Gogh



Cezanne



Ukiyo-e

Creation of New Content - Cross Domain Transfer

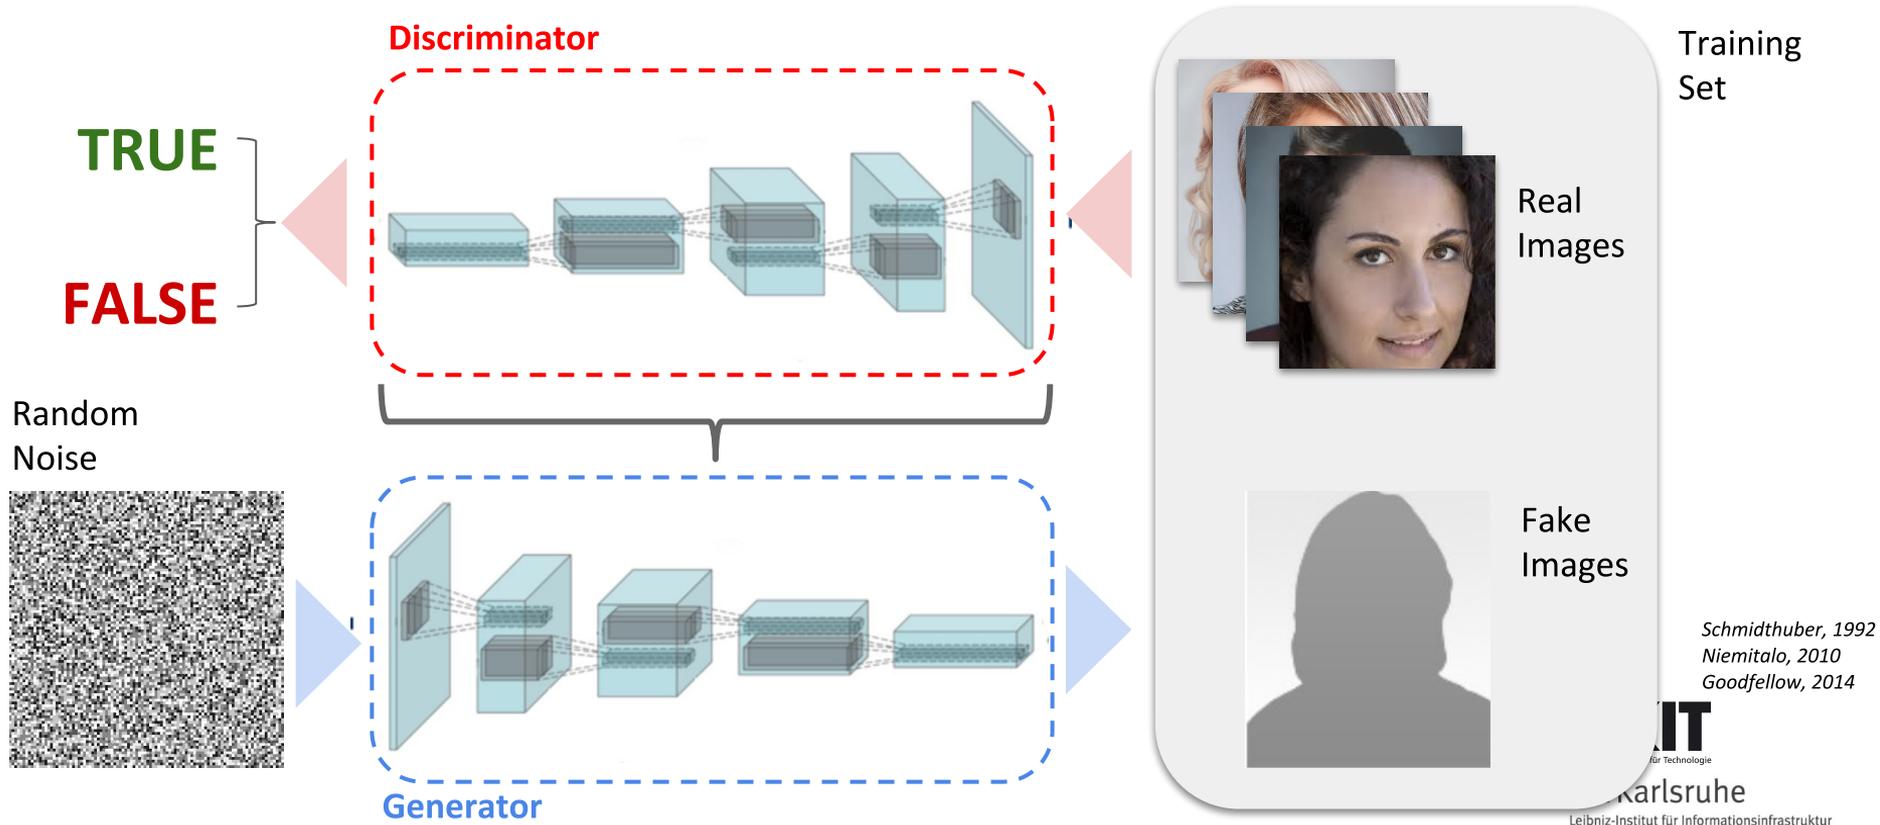
Monet



Photograph



Comparative Learning - Generative Adversarial Networks



Creation of New Content - Cross Domain Transfer



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



Generation of New Content - Colourization

Generation of New Content - Super Resolution

bicubic
(21.59dB/0.6423)



SRResNet
(23.53dB/0.7832)



SRGAN
(21.15dB/0.6868)



original



Generation of New Content - Text to Images

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:

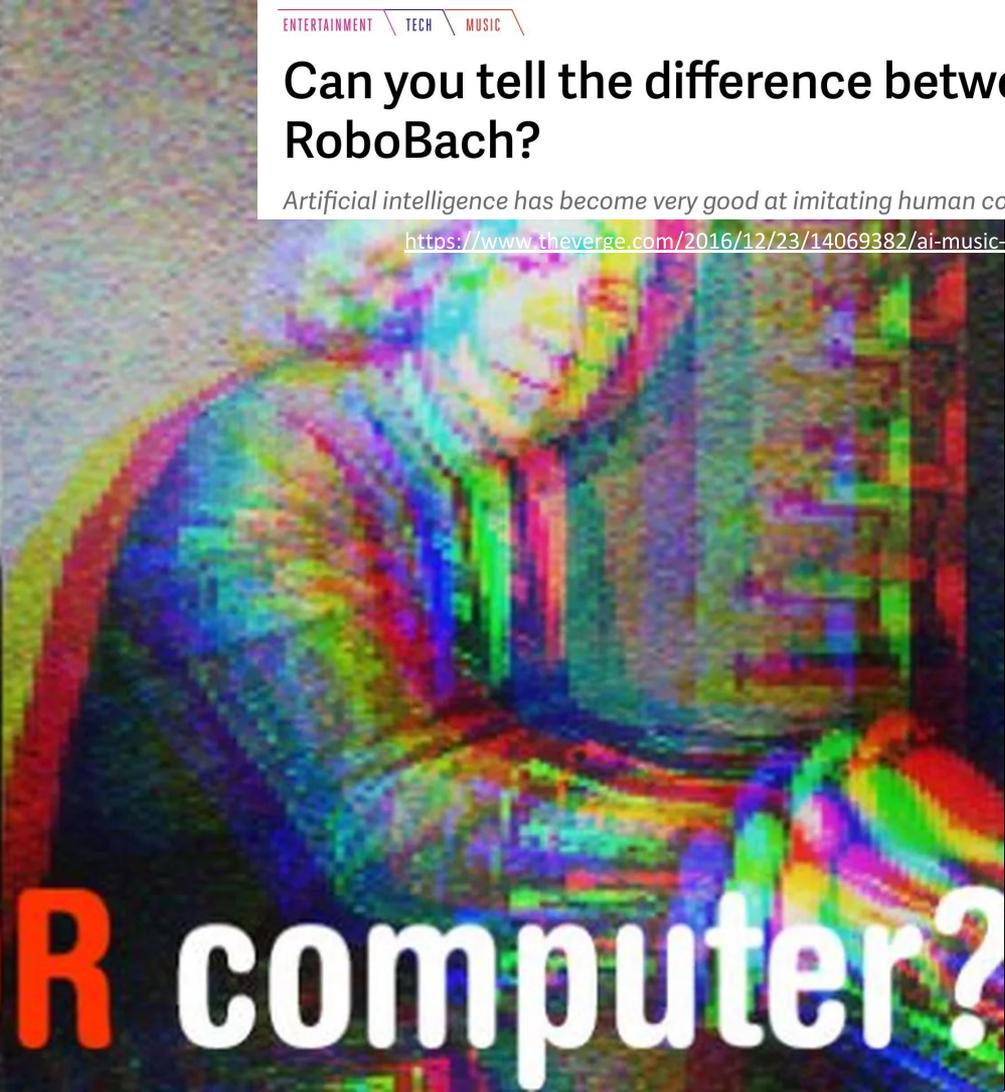
[StackGAN++: Realistic Image Synthesis with Stacked Generative Adversarial Networks.](#)

CoRR abs/1710.10916 (2017)

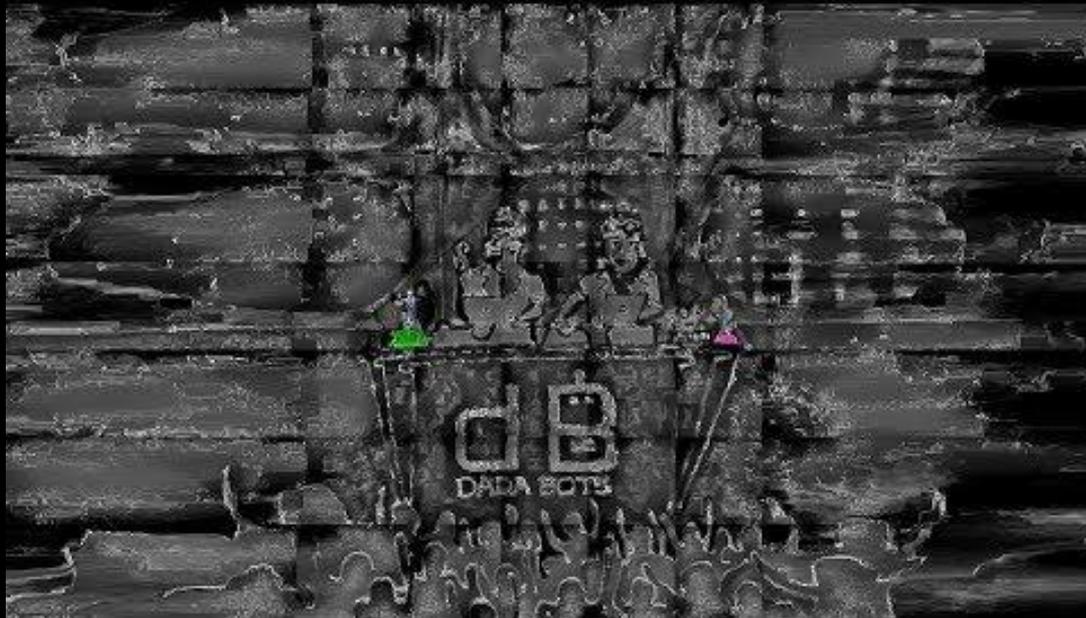
Can you tell the difference between Bach and RoboBach?

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

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** with speech interface
- 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

Artificial Intelligence

Reasoning

NLP

Planning

Machine Learning

Supervised Learning

Unsupervised Learning

Reinforcement Learning

Deep Learning
(Neural Networks)

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

- John McCarthy (1955)

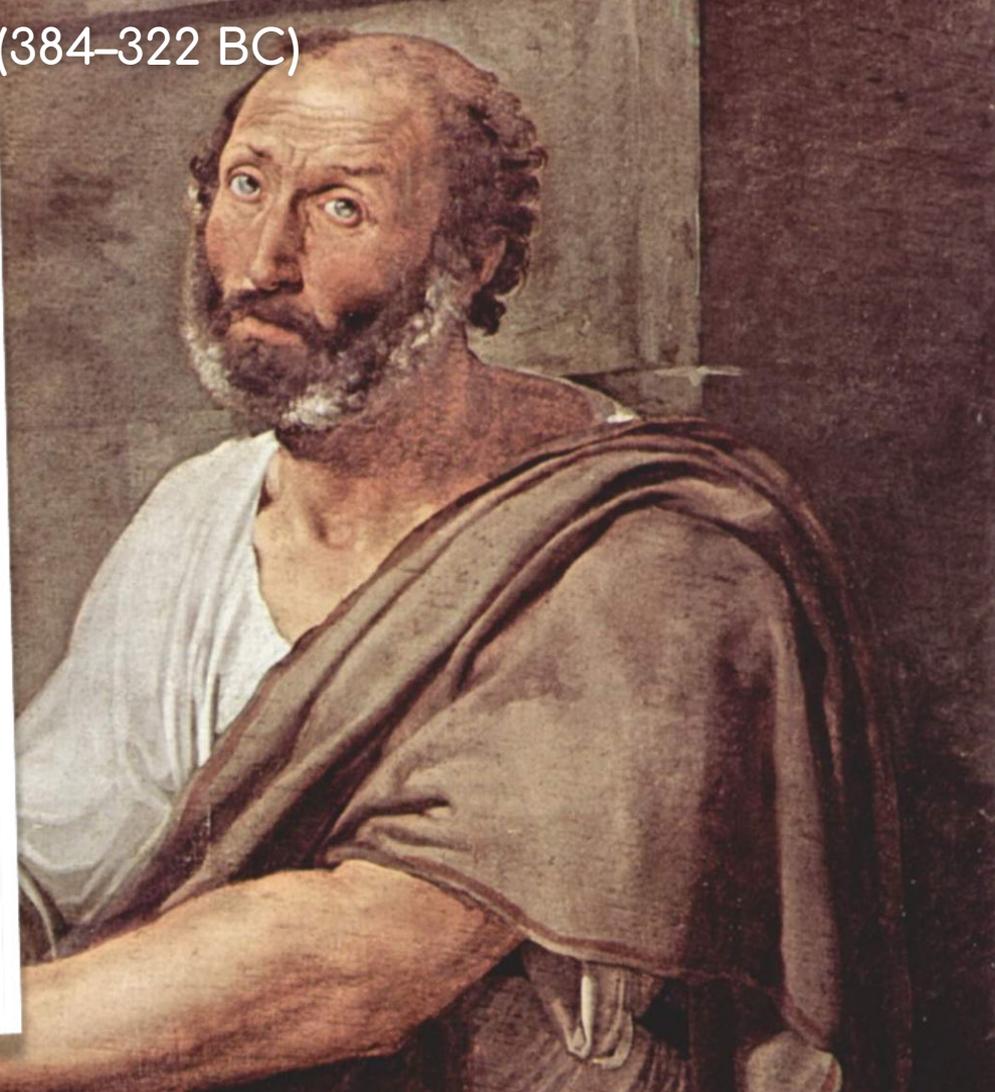
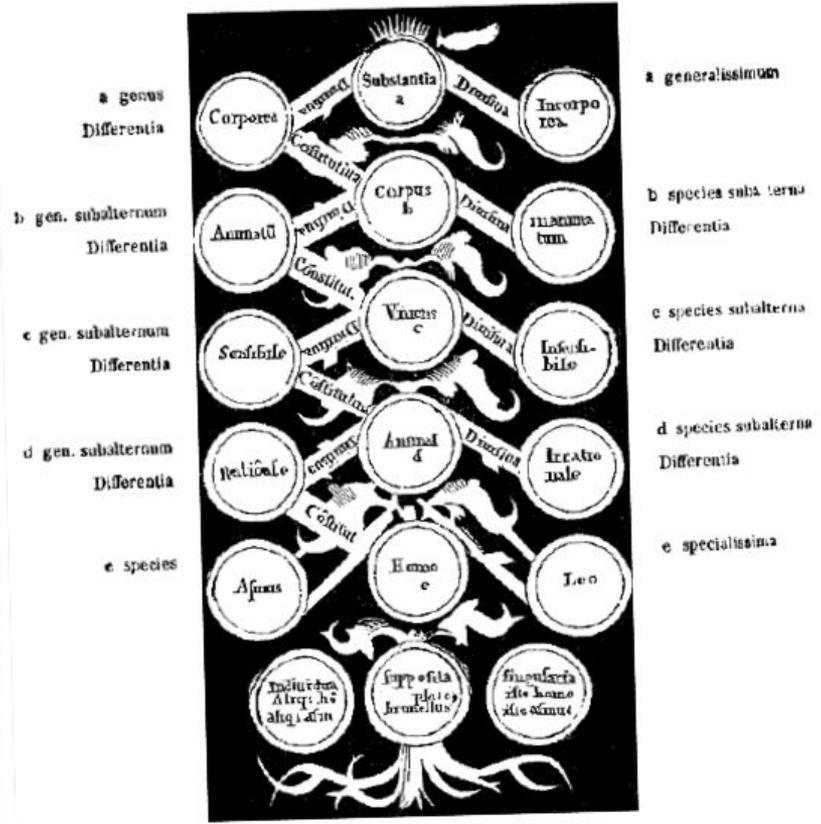
P A R E N T A L

ADVISORY

EXPLICIT SEMANTICS

The Universal Categories - Aristotle (384–322 BC)

IN PORPHYRIUM DIALOGUS I.



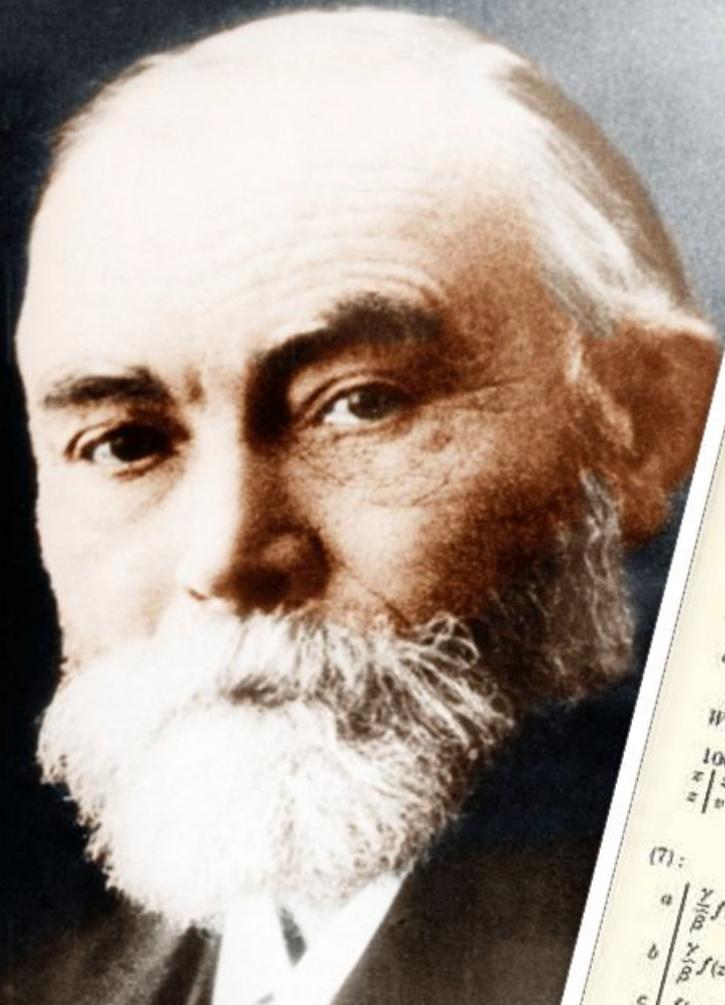
Calculus Ratiocinator - Gottfried Wilhelm Leibniz (1646-1716)

*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

Calculemus!

Begriffsschrift - Gottlob Frege (1848-1925)



71

BEGRIFFSSCHRIFT

(55) ::

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

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

§ 30. 99

(52):

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

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

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

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

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

(104).

(105).

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

106

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

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

(7):

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

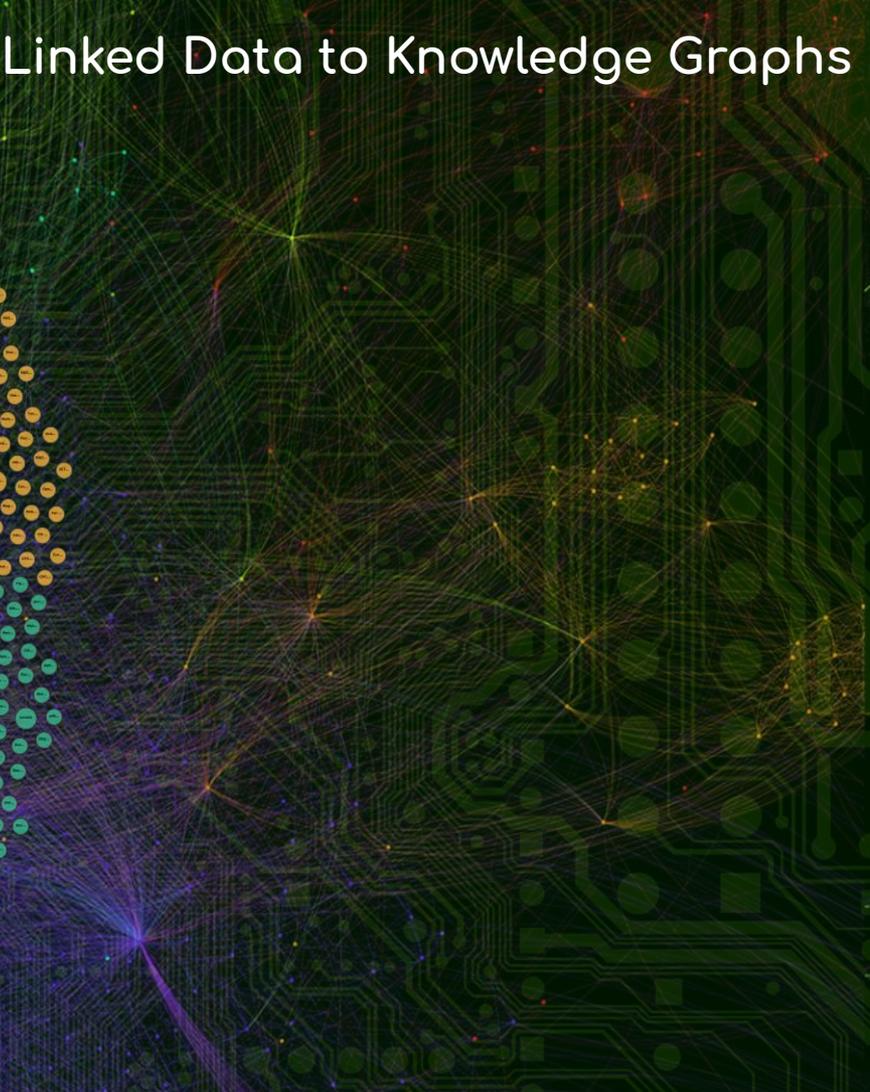
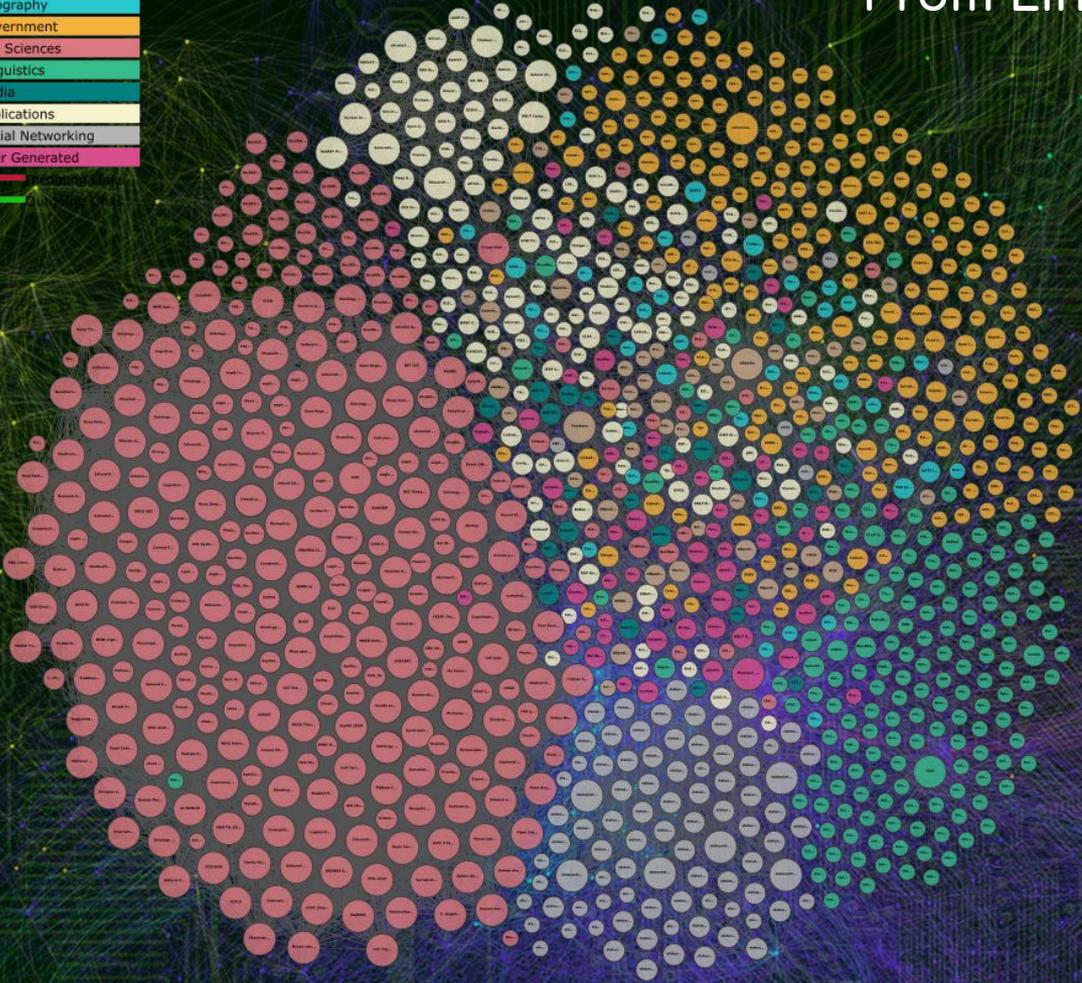
$$\begin{array}{l} \vdash \frac{\gamma}{\beta} f(z, v) \\ \vdash \frac{\gamma}{\beta} f(z, v) \end{array}$$

(106).

Frame Logic for Knowledge Representation - Marvin Minsky (1974)



From Linked Data to Knowledge Graphs

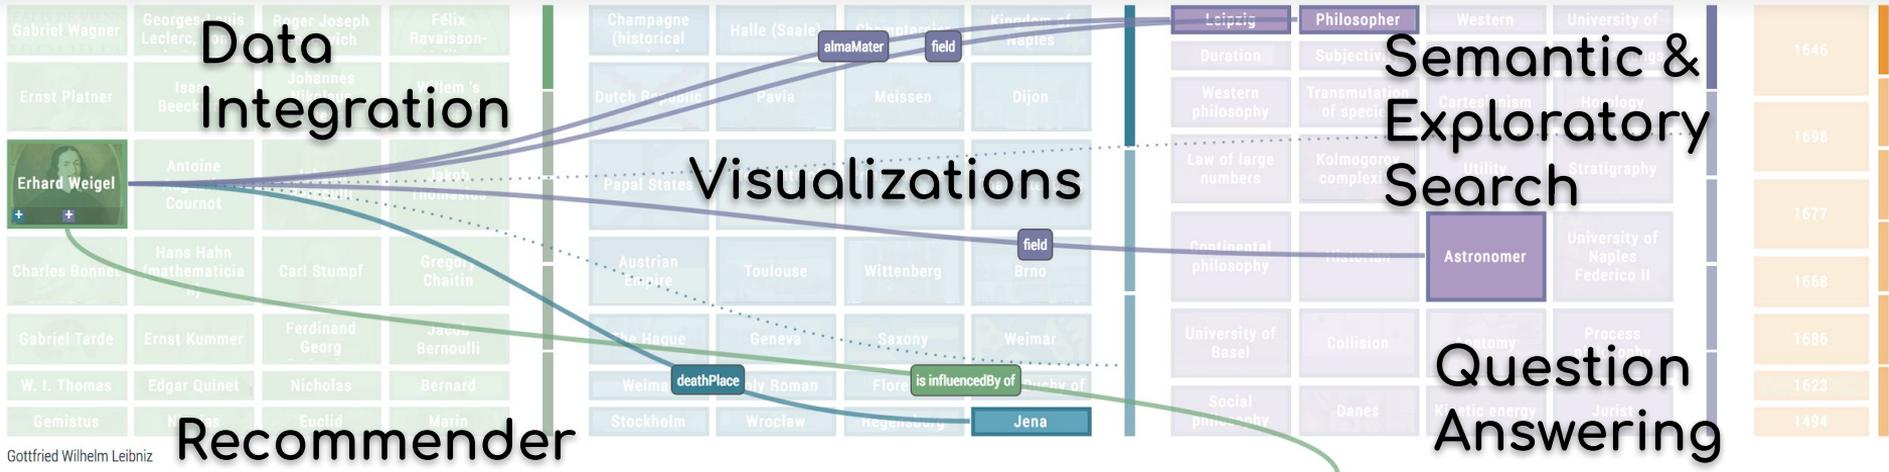


Knowledge Graph Applications



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

Relation Browser Timeline



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 ˈlaɪbnɪtʃ] or [ˈlaɪpnɪtʃ]) (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



1.1. or 12. or 13. or 14. or 15. or 16. or 17. f
2.1. w 22. d 23. v 24. i 25. m 26. / 27. g
31. f 32. i 33. i 34. j 35. k 36. l 37. m
41. w 42. o 43. o 44. p 45. q 46. r 47. f
54. p 55. p 56. p 57. l
67. r 68. m 69. q 70. g

$\frac{2}{3} + \frac{3}{4} =$
26743:8=
12986 x 3 =

Generation of New Content - Text to Images

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)

Generation of New Content - Text to Images

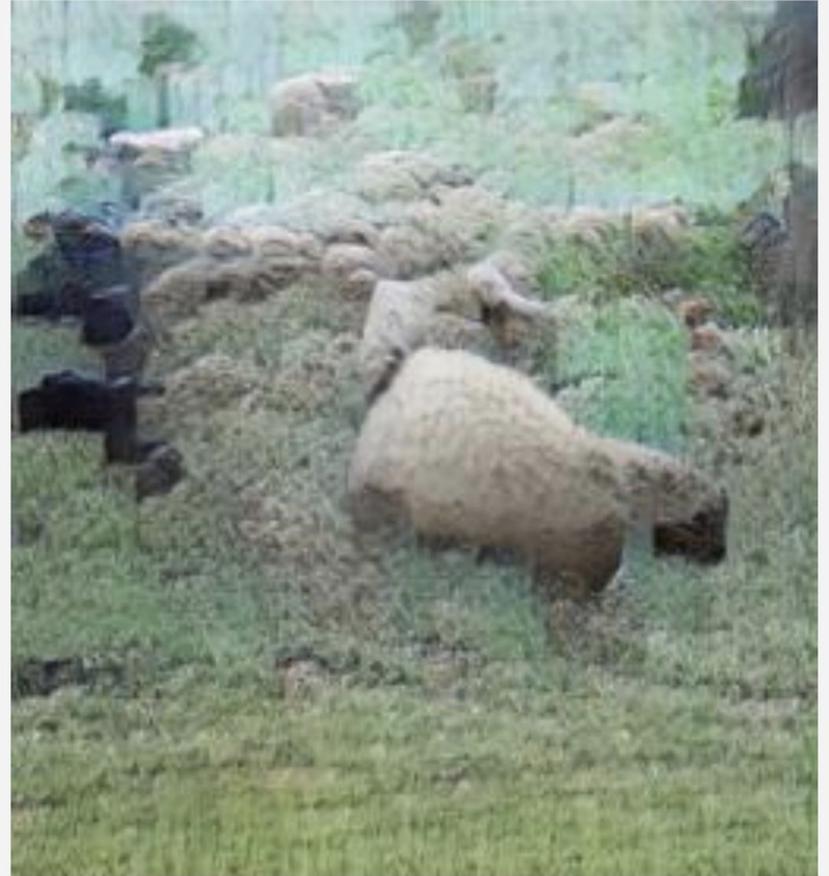
This is a small light gray bird with a small head and green crown nape and some green coloring on its wings



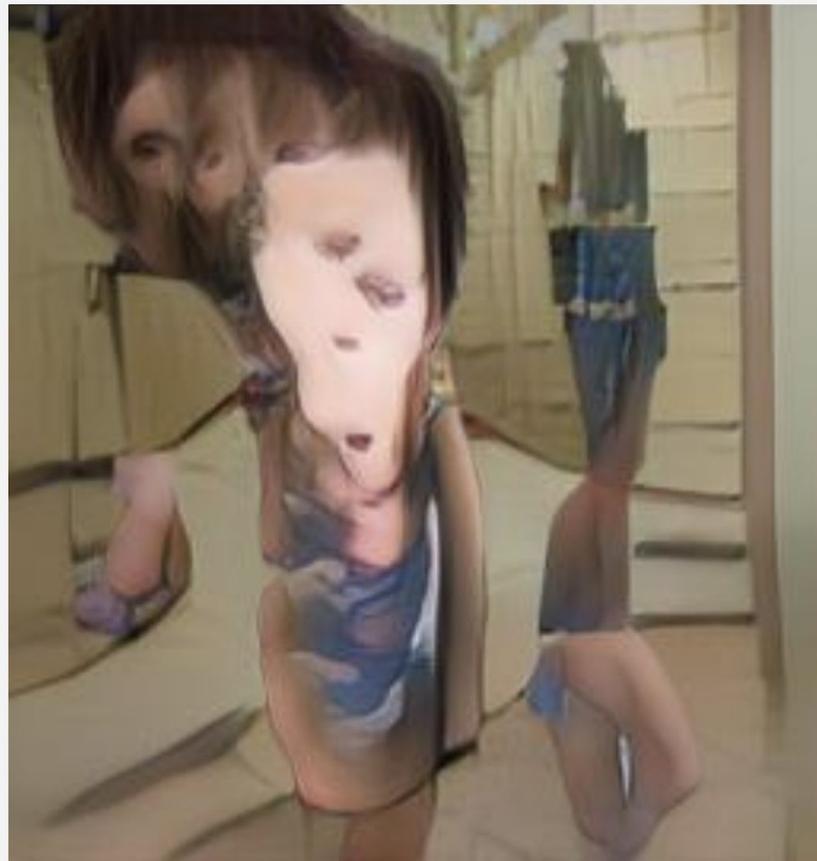
Now let's go for something more General ...



A flock of sheep on green meadows



a girl
watching tv



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



melting clocks over the dessert



“All for one,
one for all,
that is our motto.”

Alexandre Dumas,
The Three Musketeers,
(1844)



Prof. Dr. Harald Sack

FIZ Karlsruhe

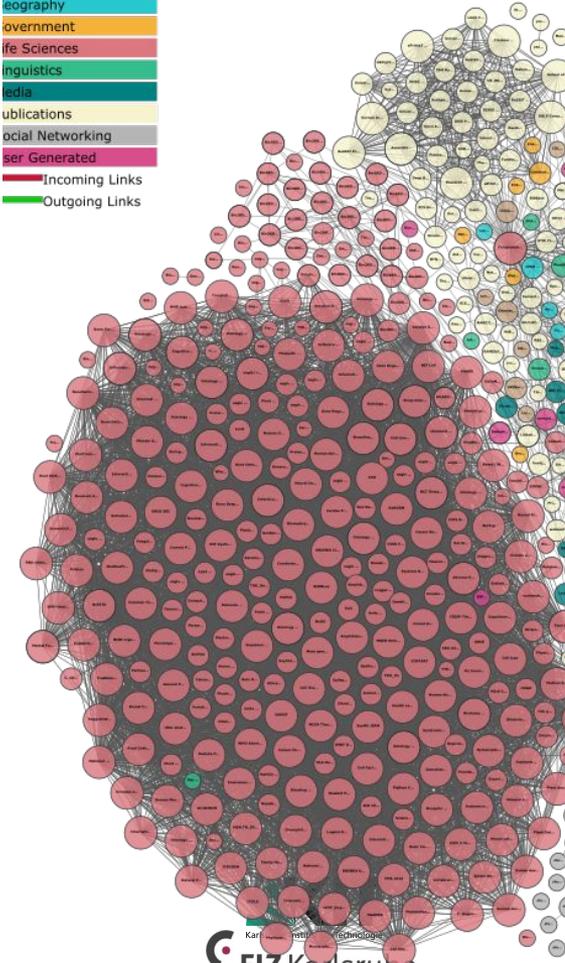
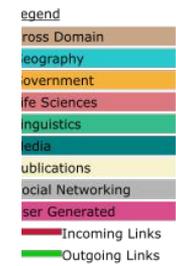
Leibniz-Institut für Informationsinfrastruktur



Karlsruher Institut für Technologie

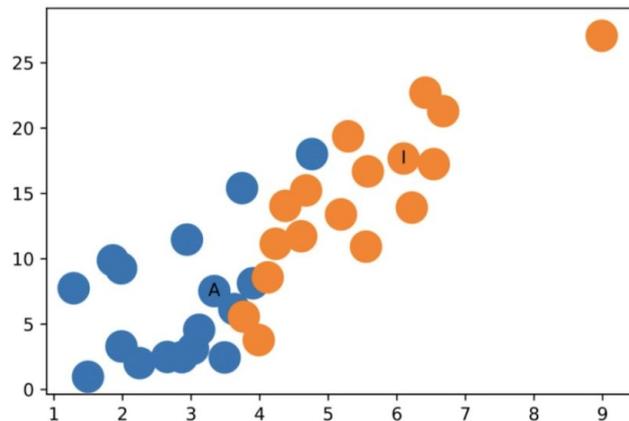
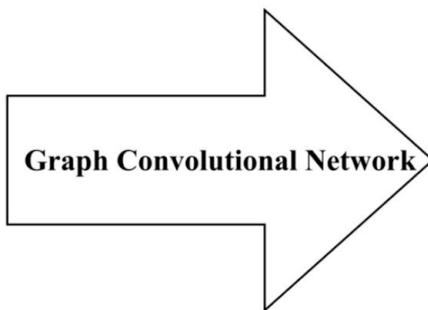
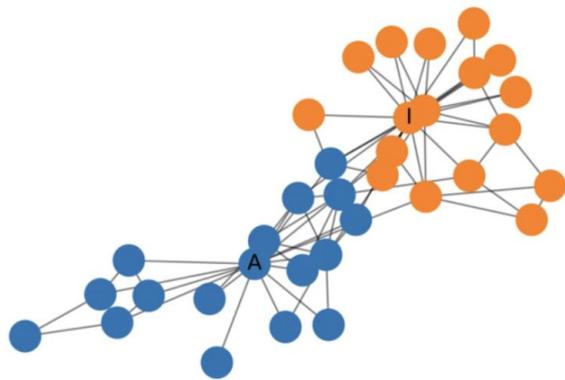
Deep Learning for Knowledge Graphs

- 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**
- Graph Analysis and NLP with Deep Learning for **Ontology Alignment** and **Link Discovery** to **combine and integrate Knowledge Graphs**



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 & Knowledge Graph Embeddings**



Indexing with Pretrained State-of-the-Art Models



(1) Image Captioning (resnet50, 5m iterations)

- 0) a collage of photos of a person holding a skateboard (p=0.000002)
- 1) a collage of photos with a bunch of different pictures (p=0.000002)
- 2) a collage of photos of a person holding a skateboard . (p=0.000001)

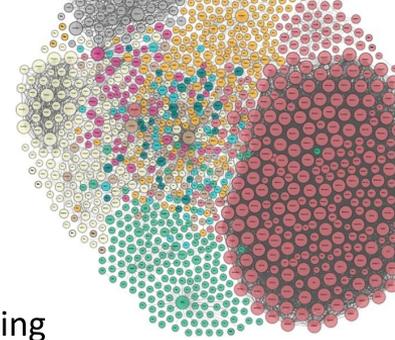
a collage of photos of a person holding a skateboard

(2) Visual Concept Detection (based on ImageNet 1K)

```
inception_v3: [[('n06596364', 'comic_book', 0.29700932), ('n07248320', 'book_jacket', 0.21479161),
vgg16       : [[('n03291819', 'envelope', 0.80103236), ('n07248320', 'book_jacket', 0.12616517),
vgg19       : [[('n03291819', 'envelope', 0.71845376), ('n06596364', 'comic_book', 0.21161233), ('
Resnet50    : [[('n03291819', 'envelope', 0.5337895), ('n06596364', 'comic_book', 0.20693506), ('
Mobilenet_v2: [[('n06596364', 'comic_book', 0.3172333), ('n03598930', 'jigsaw_puzzle', 0.16213572),
Densenet    : [[('n03291819', 'envelope', 0.2729636), ('n06596364', 'comic_book', 0.19083193), ('
NASAnet     : [[('n03291819', 'envelope', 0.4995414), ('n03485794', 'handkerchief', 0.25641188),
```

Abbot Hugo de Cluny, Margravine Mathilda of Tuscany and Henry IV, miniature from the manuscript Vita Mathildis (c 1115)

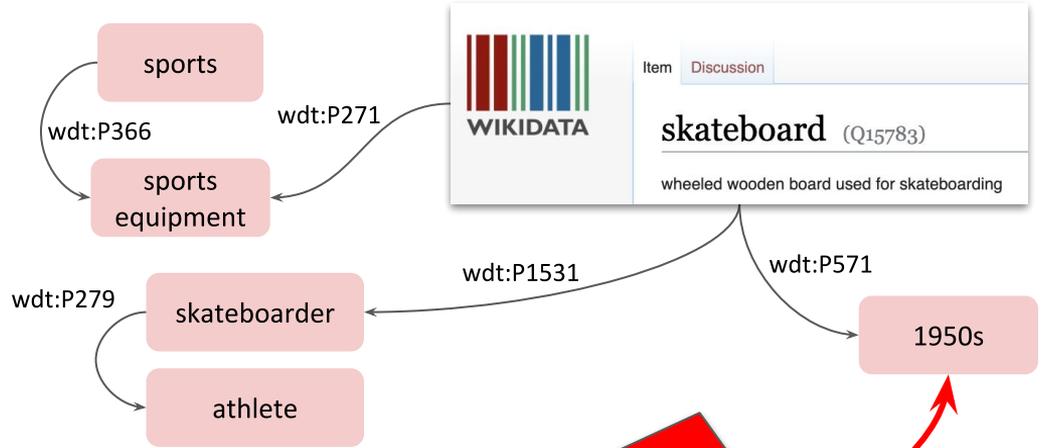
Combining Deep Learning and Semantics



Automated Image captioning:

a person holding a **skateboard**

Entity Linking



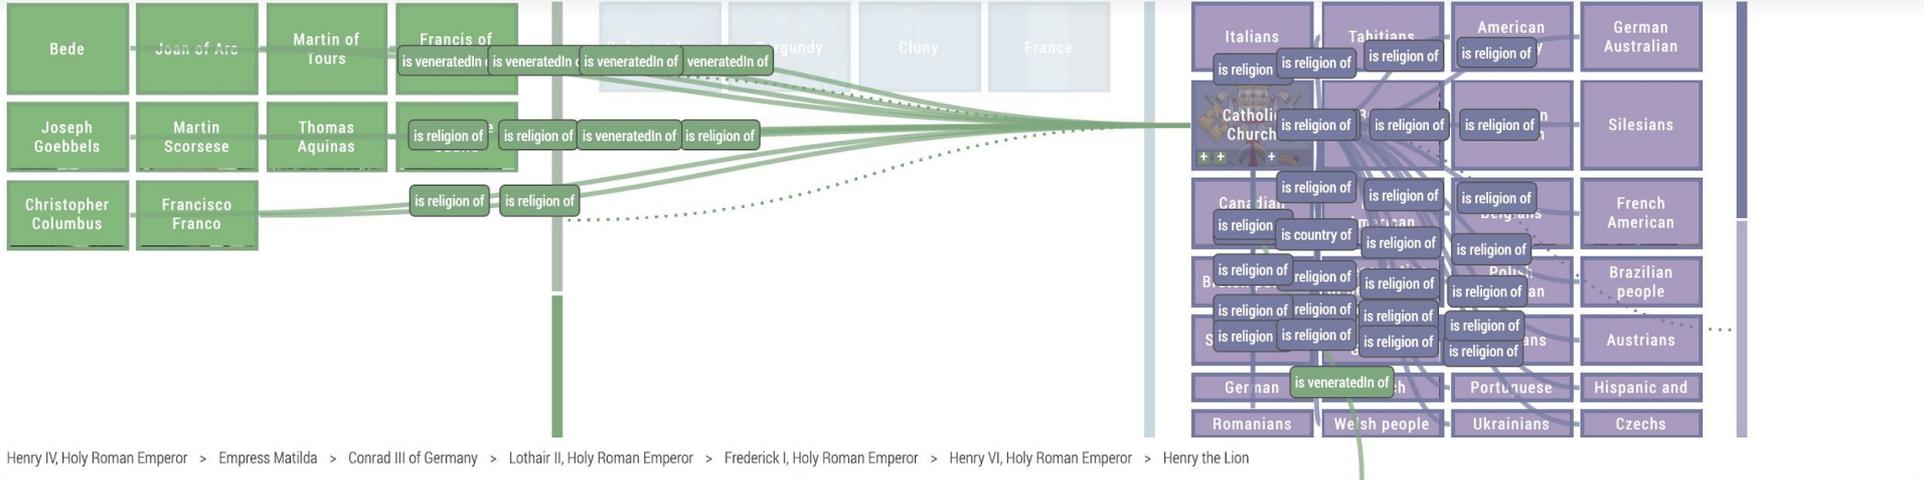
Metadata

title: Vita Mathildis
author: Donizio
date: 1115
location: Bibliotheca Apostolica Vaticana...

Knowledge Graph Based Exploration



Relation Browser Timeline



9 Recommended Articles:

- #2 Abbot Suger and the Birth of the Gothic Style
- #3 Otto the Great – Founder of the Holy Roman Empire
- #4 Geoffrey Chaucer – the Father of English Literature
- #5 The Assassination of Thomas Becket
- #6 Hildegard of Bingen – More than the 'Sybil of the Rhine'
- #7 Pieter van Musschenbroek and the Leyden Jar

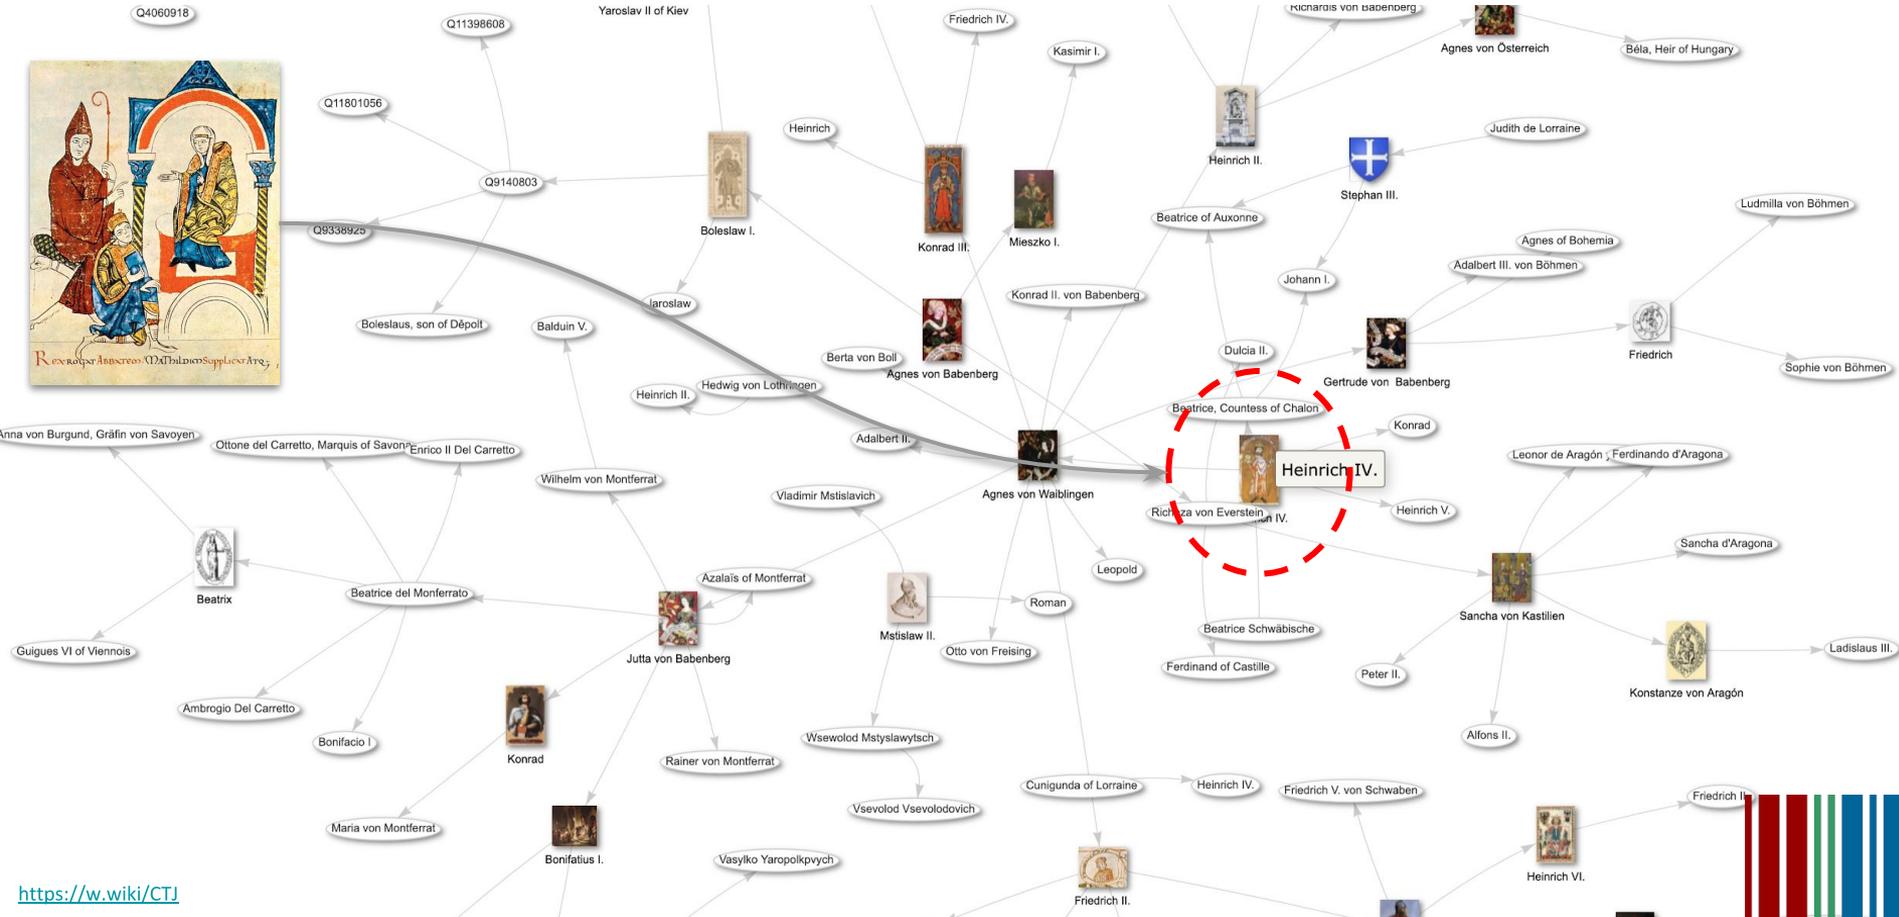
Hugh of Cluny



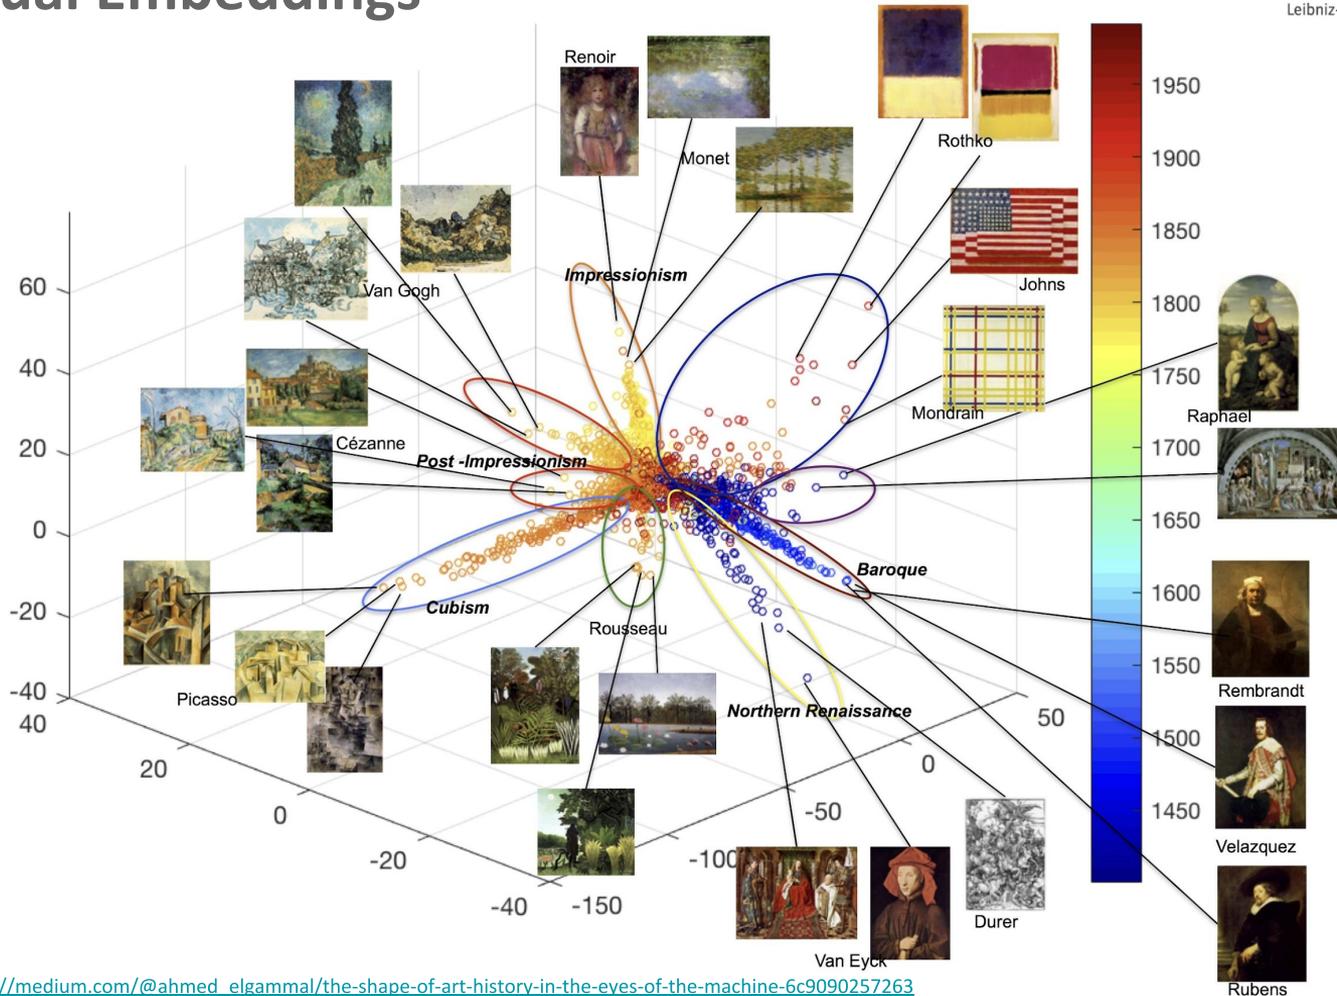
DBpedia: Hugh of Cluny



Knowledge Graph Based Exploration



Visual Embeddings



Similarity Based Search and Exploration

wikiview.net

hugo of cluny

Add search term...

art religion culture ancient faith church
image religious background god

BACK TO RESULTS



Rex rogat abbate[m] Matildem supplicatque Atq[ue]

SHOW SIMILAR WIKIMEDIA

Hugo-v-cluny heinrich-iv
mathilde-v-tuszien cod-vat-lat-
4922 1115ad



Navigation and control interface including a blue pyramid icon, a search icon, and a grid icon.

Take Home Messages:

- Deep Learning is a **game changer**
- The quality of your achieved results always depends on the **quality (and quantity) of your training data**
- (Out-of-the-box Deep Learning) models are **easy to use and work quite well** (however..not always)
- **Deep Learning benefits from exploiting explicit Semantics**
- **Semantics benefits from leveraging Deep Learning**

[Long Term] Goal

- Combine **Semantics** (a.k.a. Symbolic Reasoning) and **Machine Learning** into hybrid systems (**Neuro-Symbolic Integration**)

[Short Term] Goals

- **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.” [Ridley Scott, Bladerunner, 1982]

Prof. Dr. Harald Sack

Do Neural Networks Dream of Semantics?

harald.sack@fiz-karlsruhe.de

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

PLDN Knowledge Graphs & Linked Data Event 2020

Hilversum, 22. Jan. 2020

