FIZ Karlsruhe
Leibniz-Institut für Informationsinfrastruktur

Deep Learning and Semantics
The two Faces of Artificial Intelligence

Prof. Dr. Harald Sack
STN Basel Info Day 2020
October 26, 2020

“Any sufficiently advanced technology is indistinguishable from magic.”

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

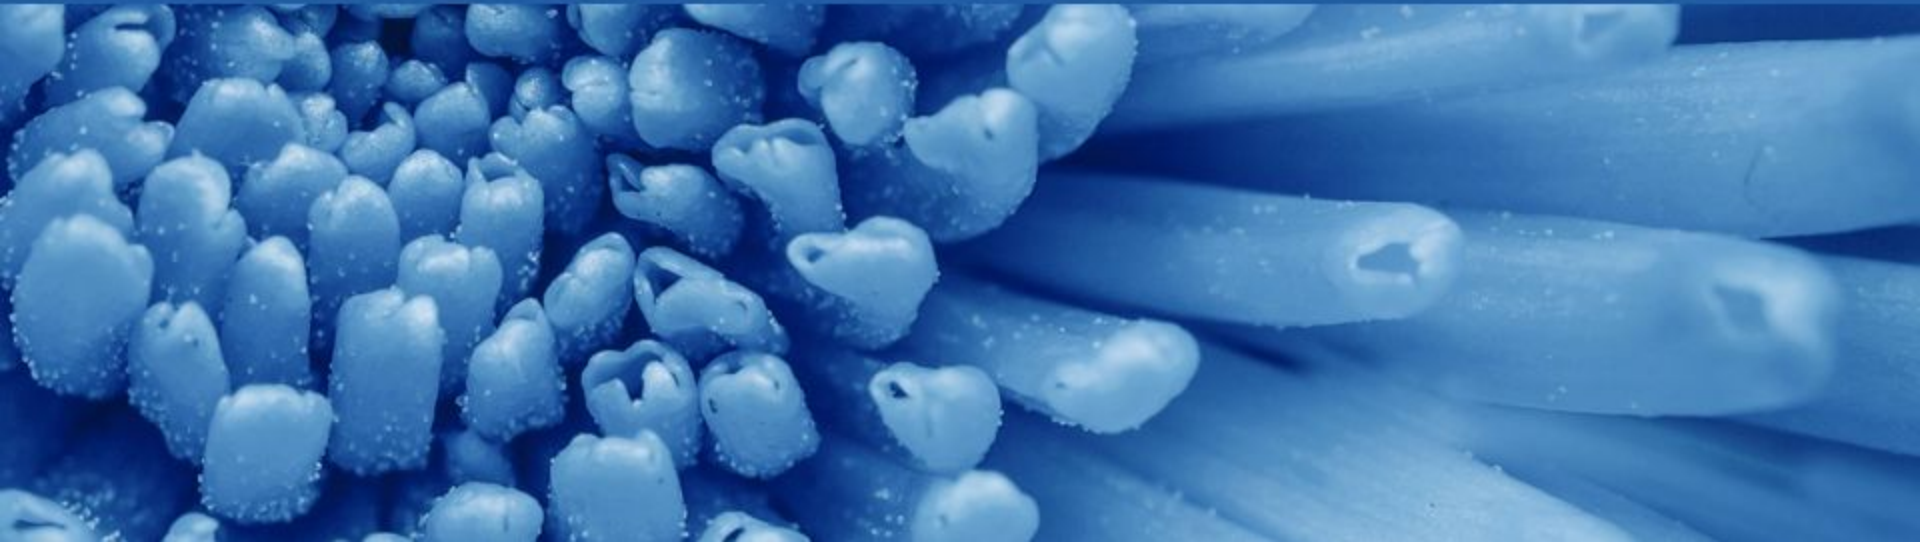

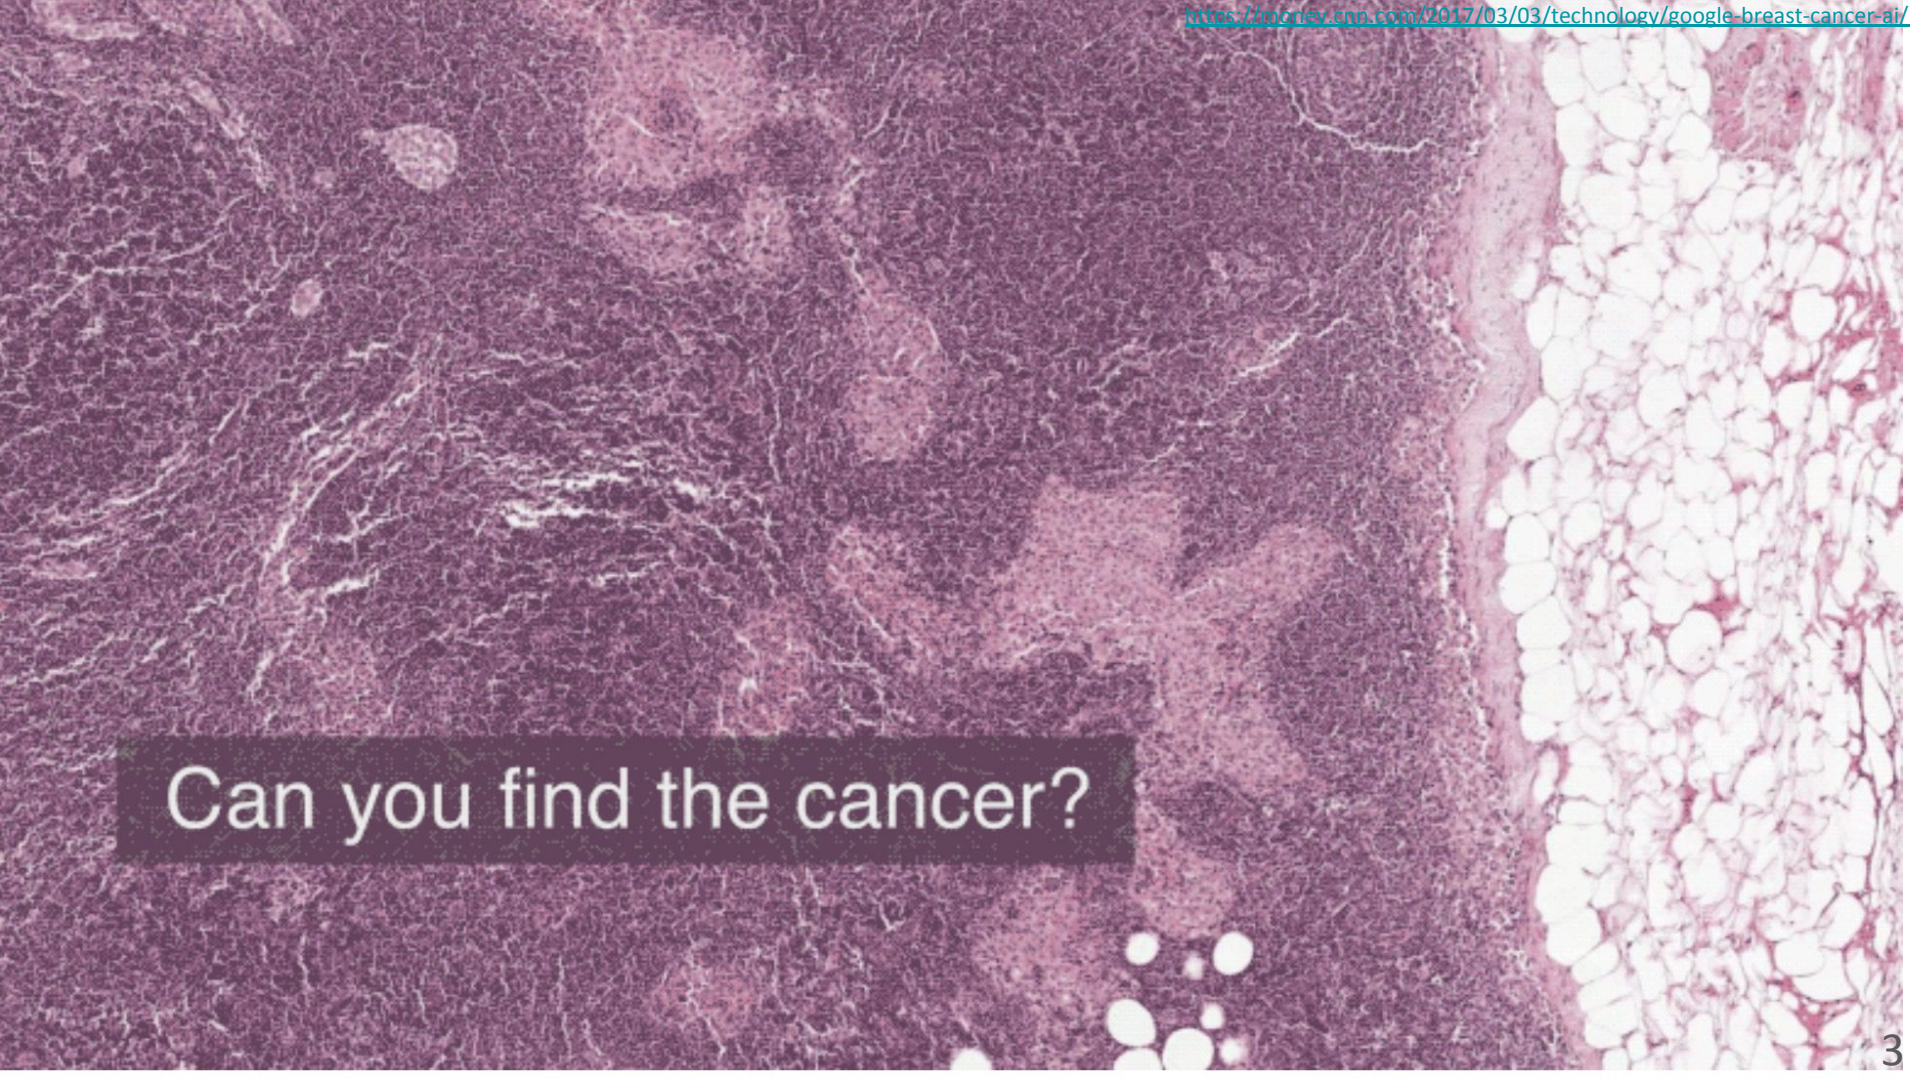
https://money.cnn.com/2017/03/03/technology/google-breast-cancer-ai/
Can you find the cancer?

https://money.cnn.com/2017/03/03/technology/google-breast-cancer-ai/

NEWS

Google AI Detects Breast Cancer Better Than Humans

BY STEPHANIE MLOT 10.15.2018 :: 9:26AM EST

44 SHARES

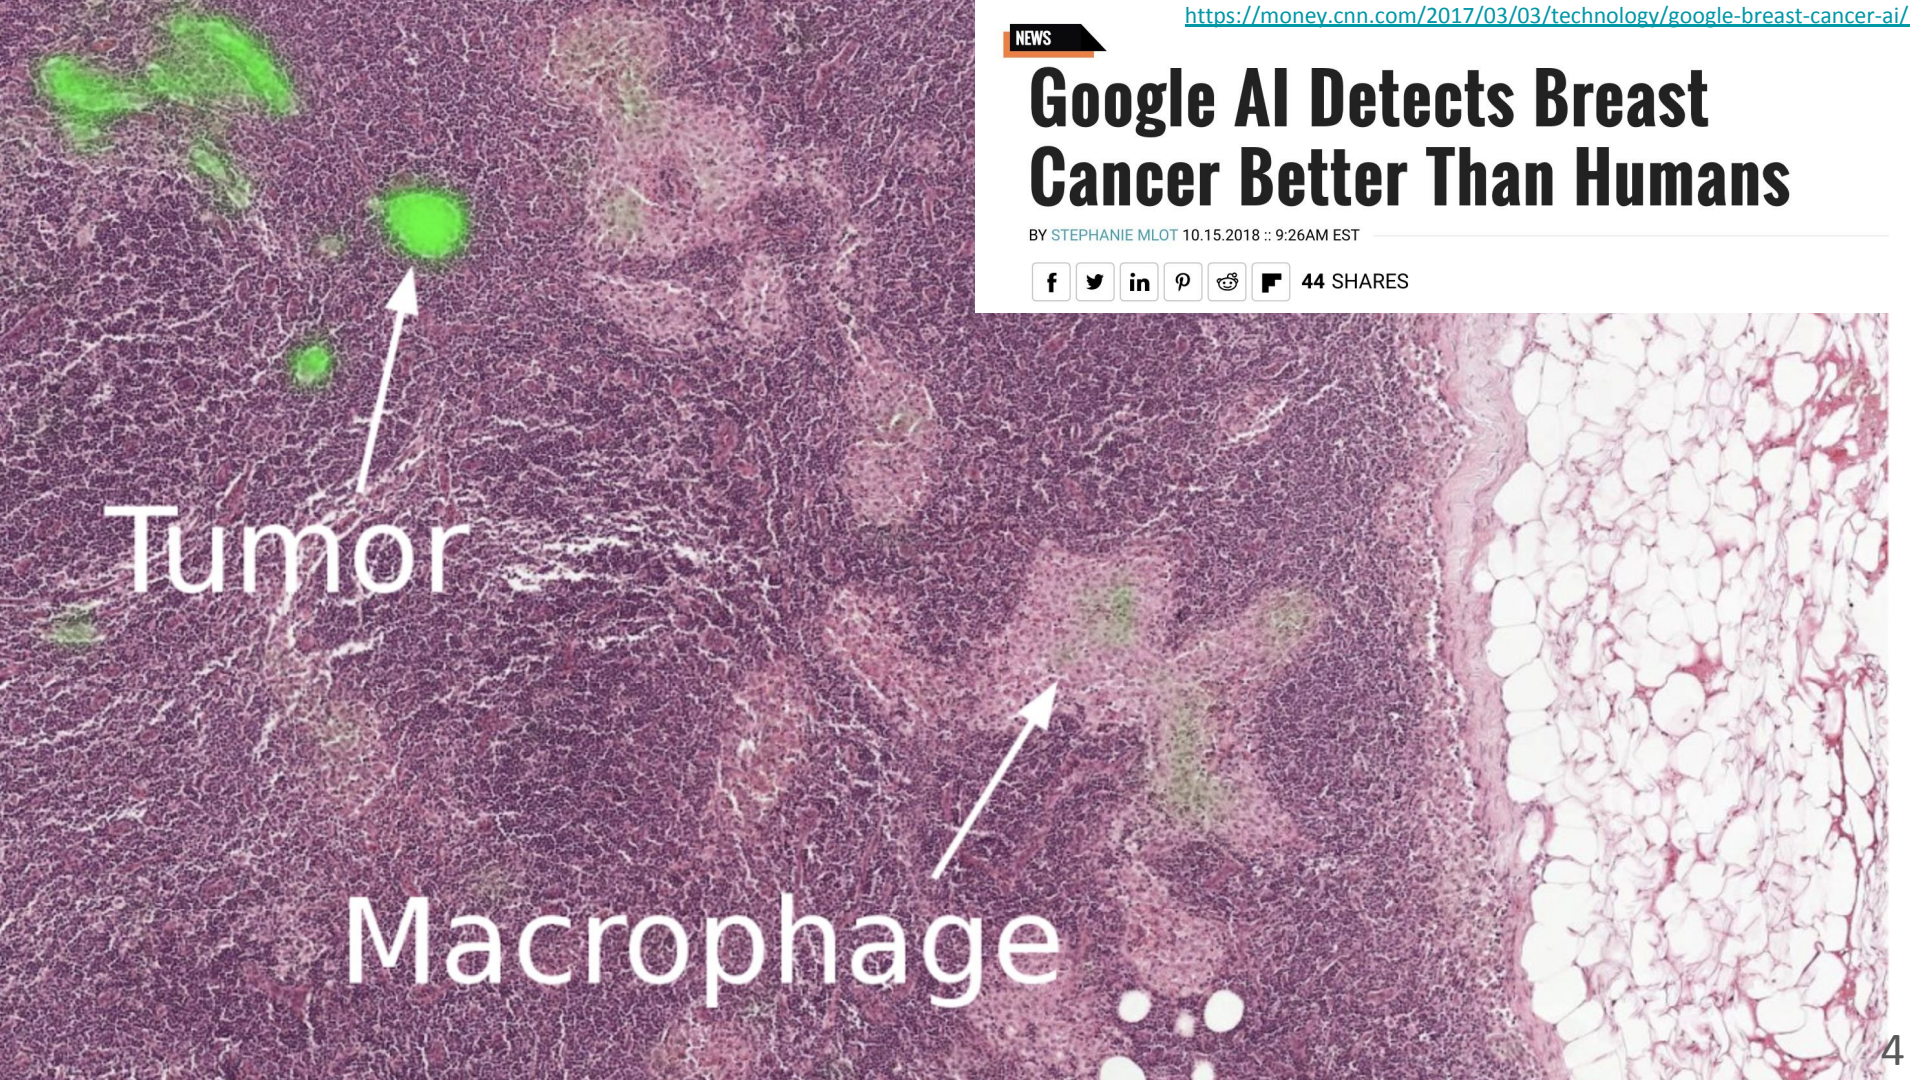

Science

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

share

17

AlphaGo

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

FIZ Karlsruhe
Leibniz-Institut für Informationsinfrastruktur

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

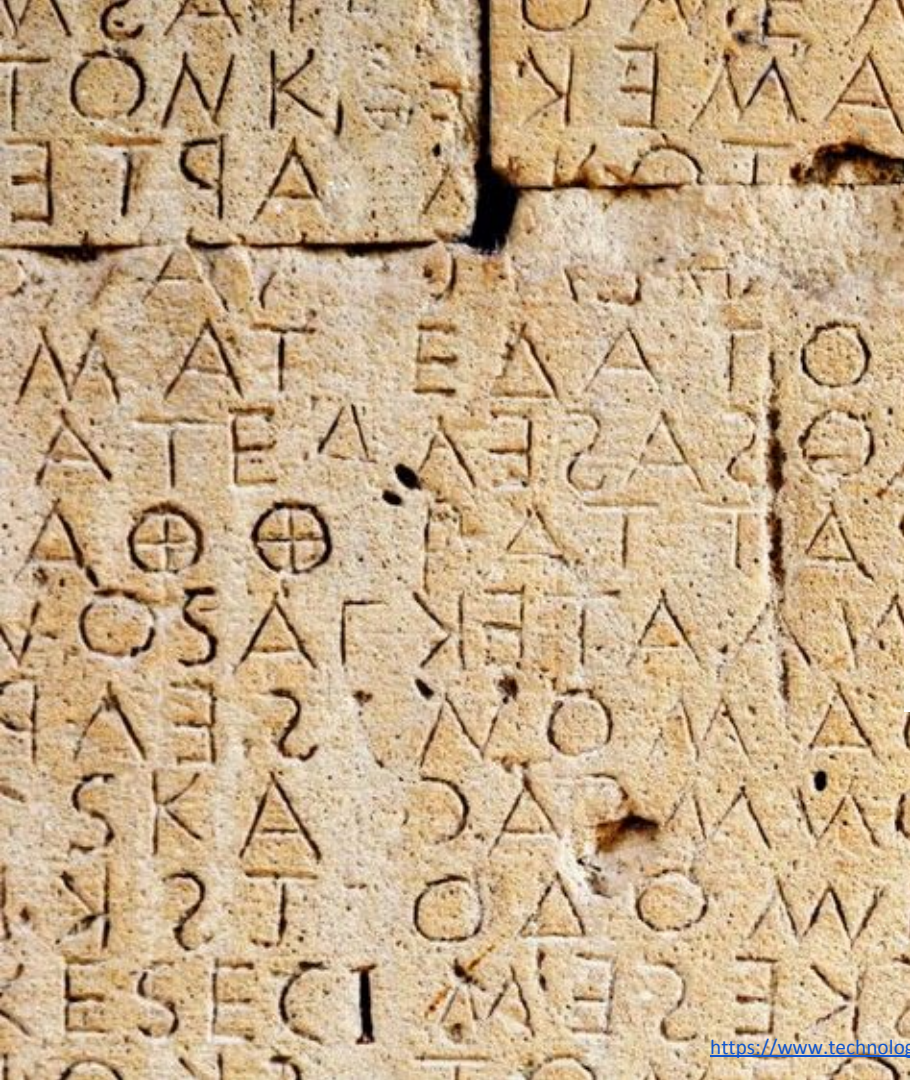

Humans and Technology

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.

by **Emerging Technology from the arXiv** Jul 1, 2019

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

OpenAI's GPT-3 may be the biggest thing since bitcoin

JUL 18, 2020

Summary: I share my early experiments with OpenAI's new language prediction model (GPT-3) beta. I explain why I think GPT-3 has disruptive potential comparable to that of blockchain technology.

OpenAI, a non-profit artificial intelligence research company backed by Peter Thiel, Elon Musk, Reid Hoffman, Marc Benioff, Sam Altman and others, released its third generation of language prediction model (GPT-3) into the open-source wild. Language models allow computers to produce random-ish sentences of approximately the same length and grammatical structure as those in a given body of text.

In my early experiments with GPT-3 I found that GPT-3's predicted sentences, when published on the bitcointalk.org forum, attracted lots of positive attention from posters there, including suggestions that the system must have been intelligent (and/or sarcastic) and that it had found subtle patterns in their posts. I imagine that similar results can be obtained by republishing GPT-3's outputs to other message boards, blogs, and social media.

I predict that, unlike its two predecessors (PTB and OpenAI GPT-2), OpenAI GPT-3 will eventually be widely used to pretend the author of a text is a person of interest, with unpredictable and amusing effects on various communities. I further predict that this will spark a creative gold rush among talented amateurs to train similar models and adapt them to a variety of purposes, including: mock news, "researched journalism", advertising, politics, and propaganda.

Are you being served?

I chose bitcointalk.org as the target environment for my experiments for a variety of reasons: It is a popular forum with many types of posts and posters.

https://maraoz.com/2020/07/18/openai-gpt3/

"...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.I (1854)

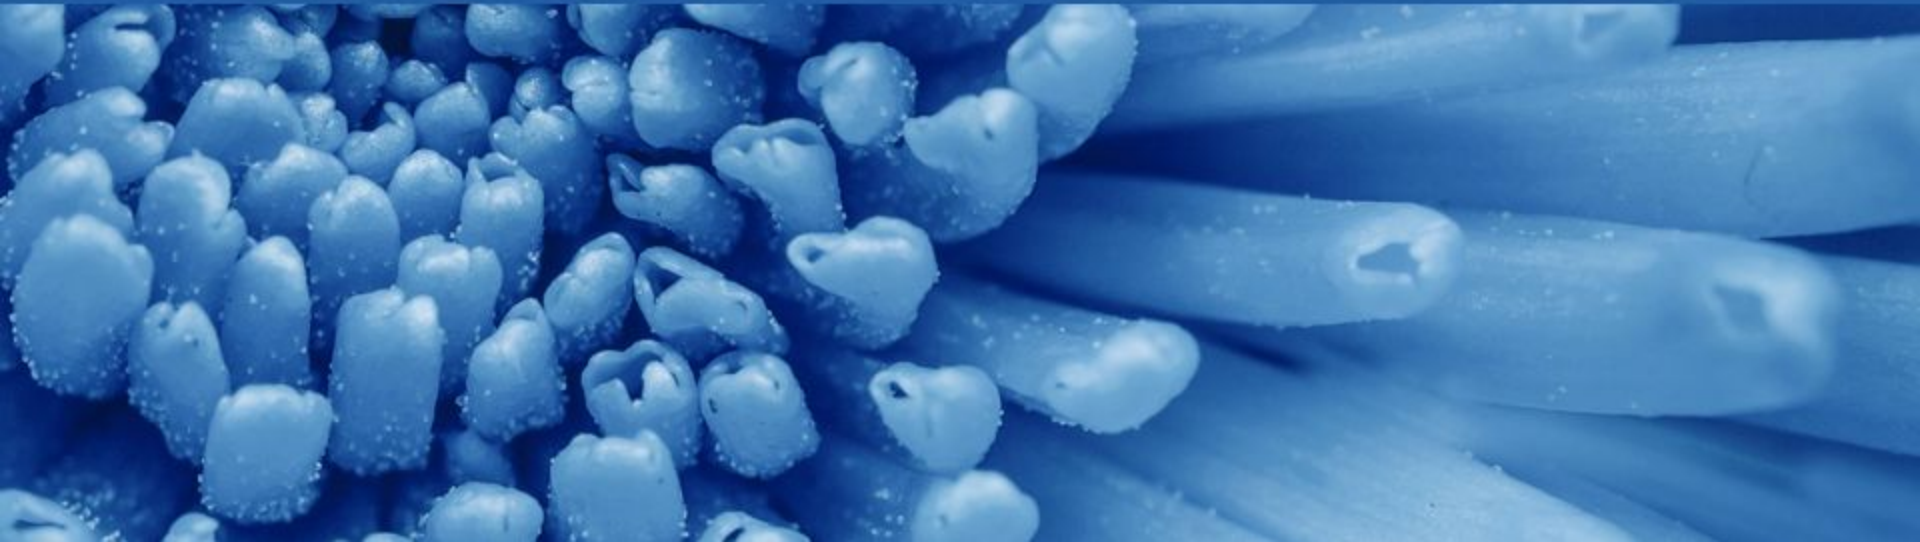

Transferring Biology into a Mathematical Model

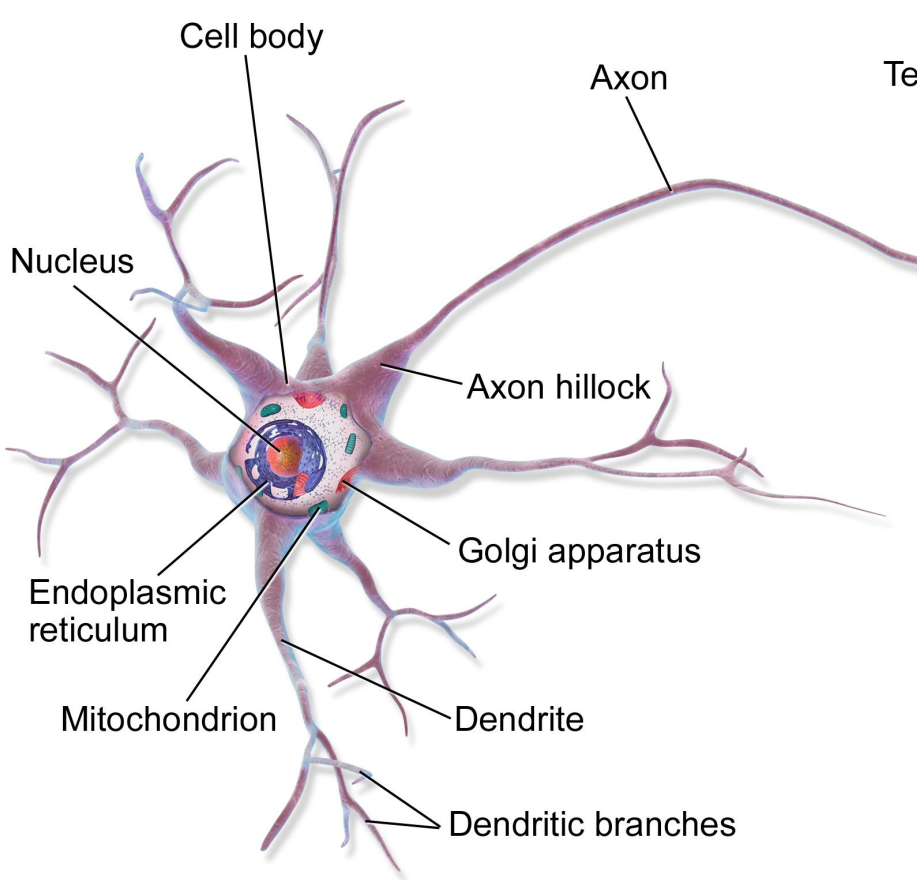

Donald Hebb (1940)
McCulloch & Pitts (1943)

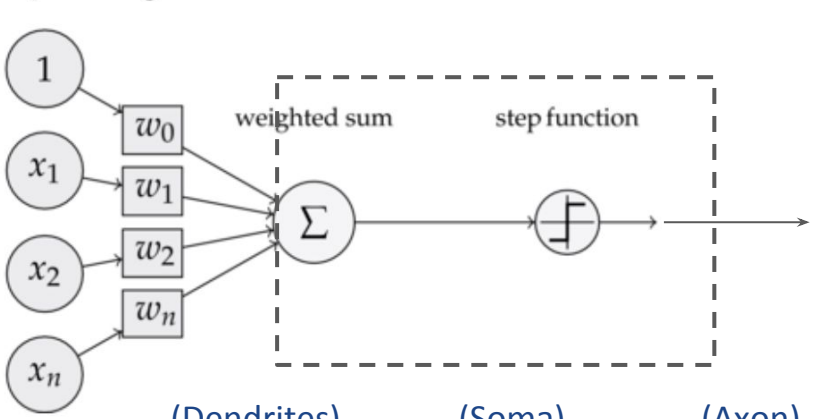

Cognitive Computing - The MARK 1 Perceptron

(Rosenblatt, 1957)

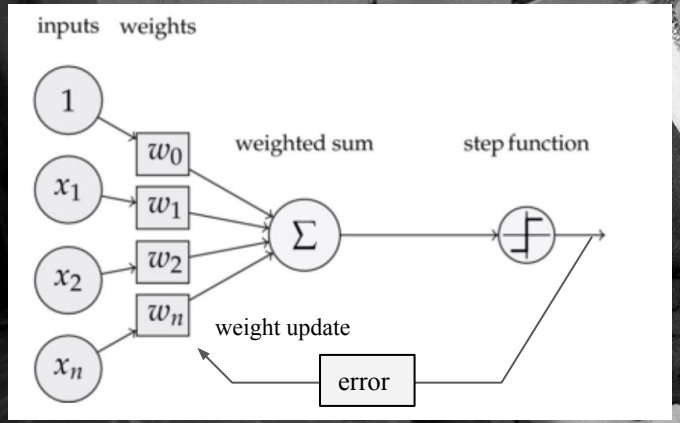

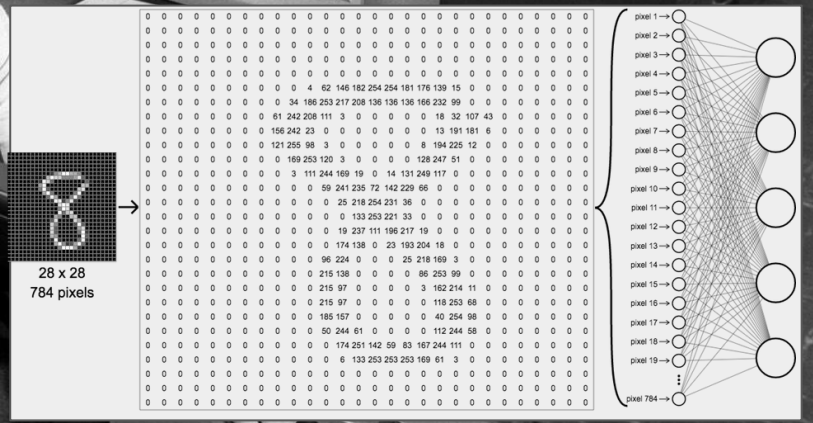

The Triumph of the Connectionist Paradigm

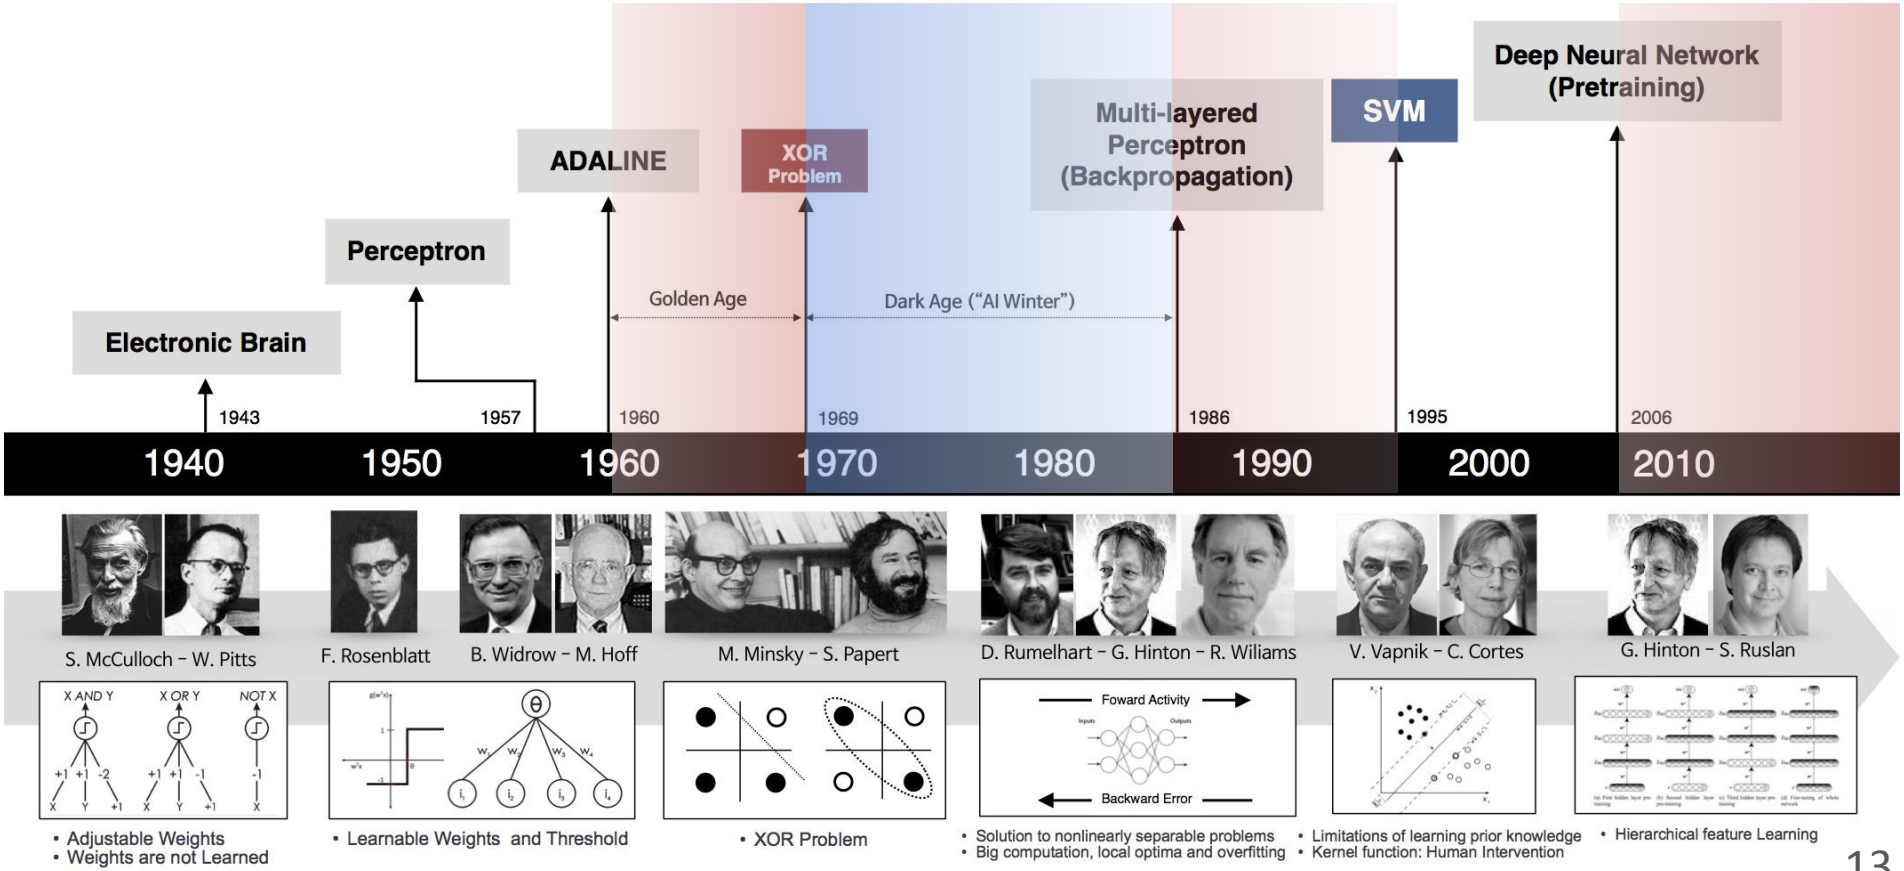

“90% of most Magic merely consists of knowing one extra fact.”

Terry Pratchett, Night Watch (2002)

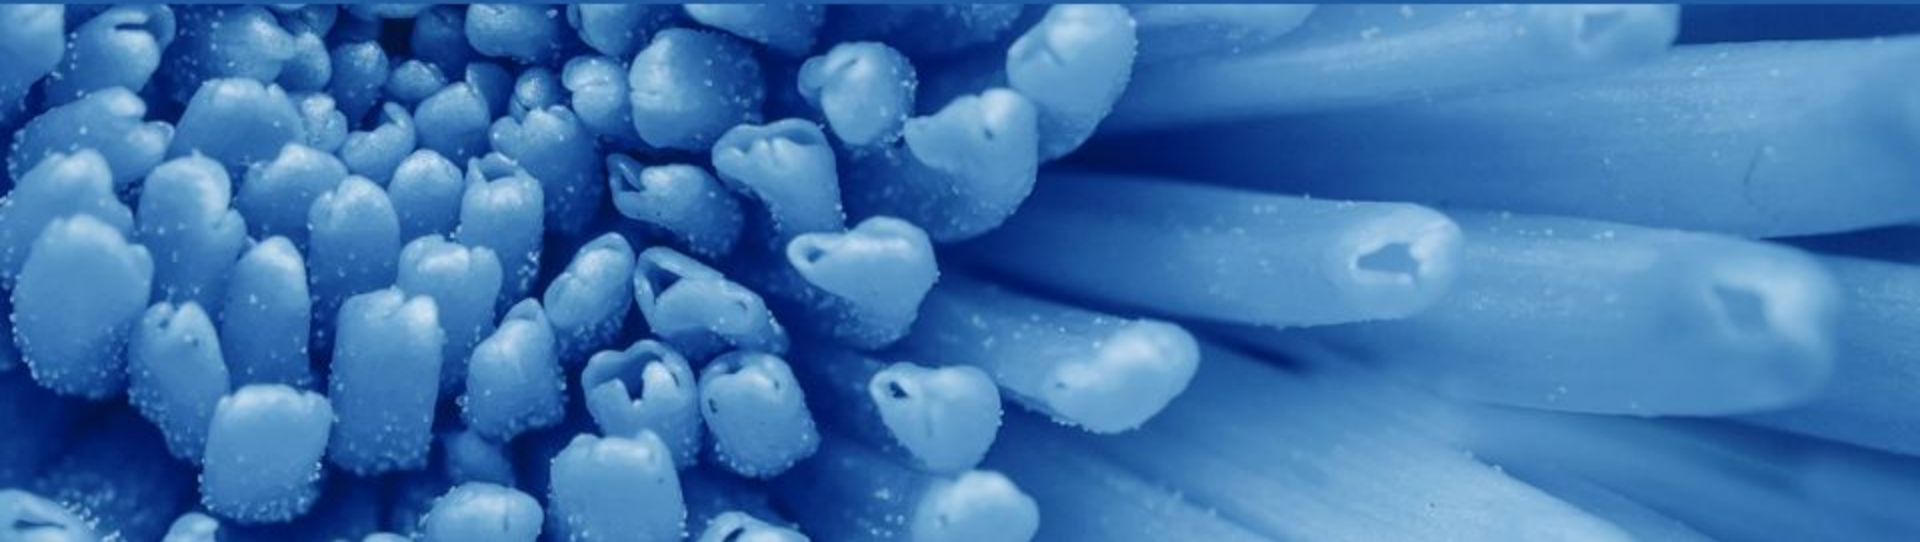

Why has Deep Learning become so successful?

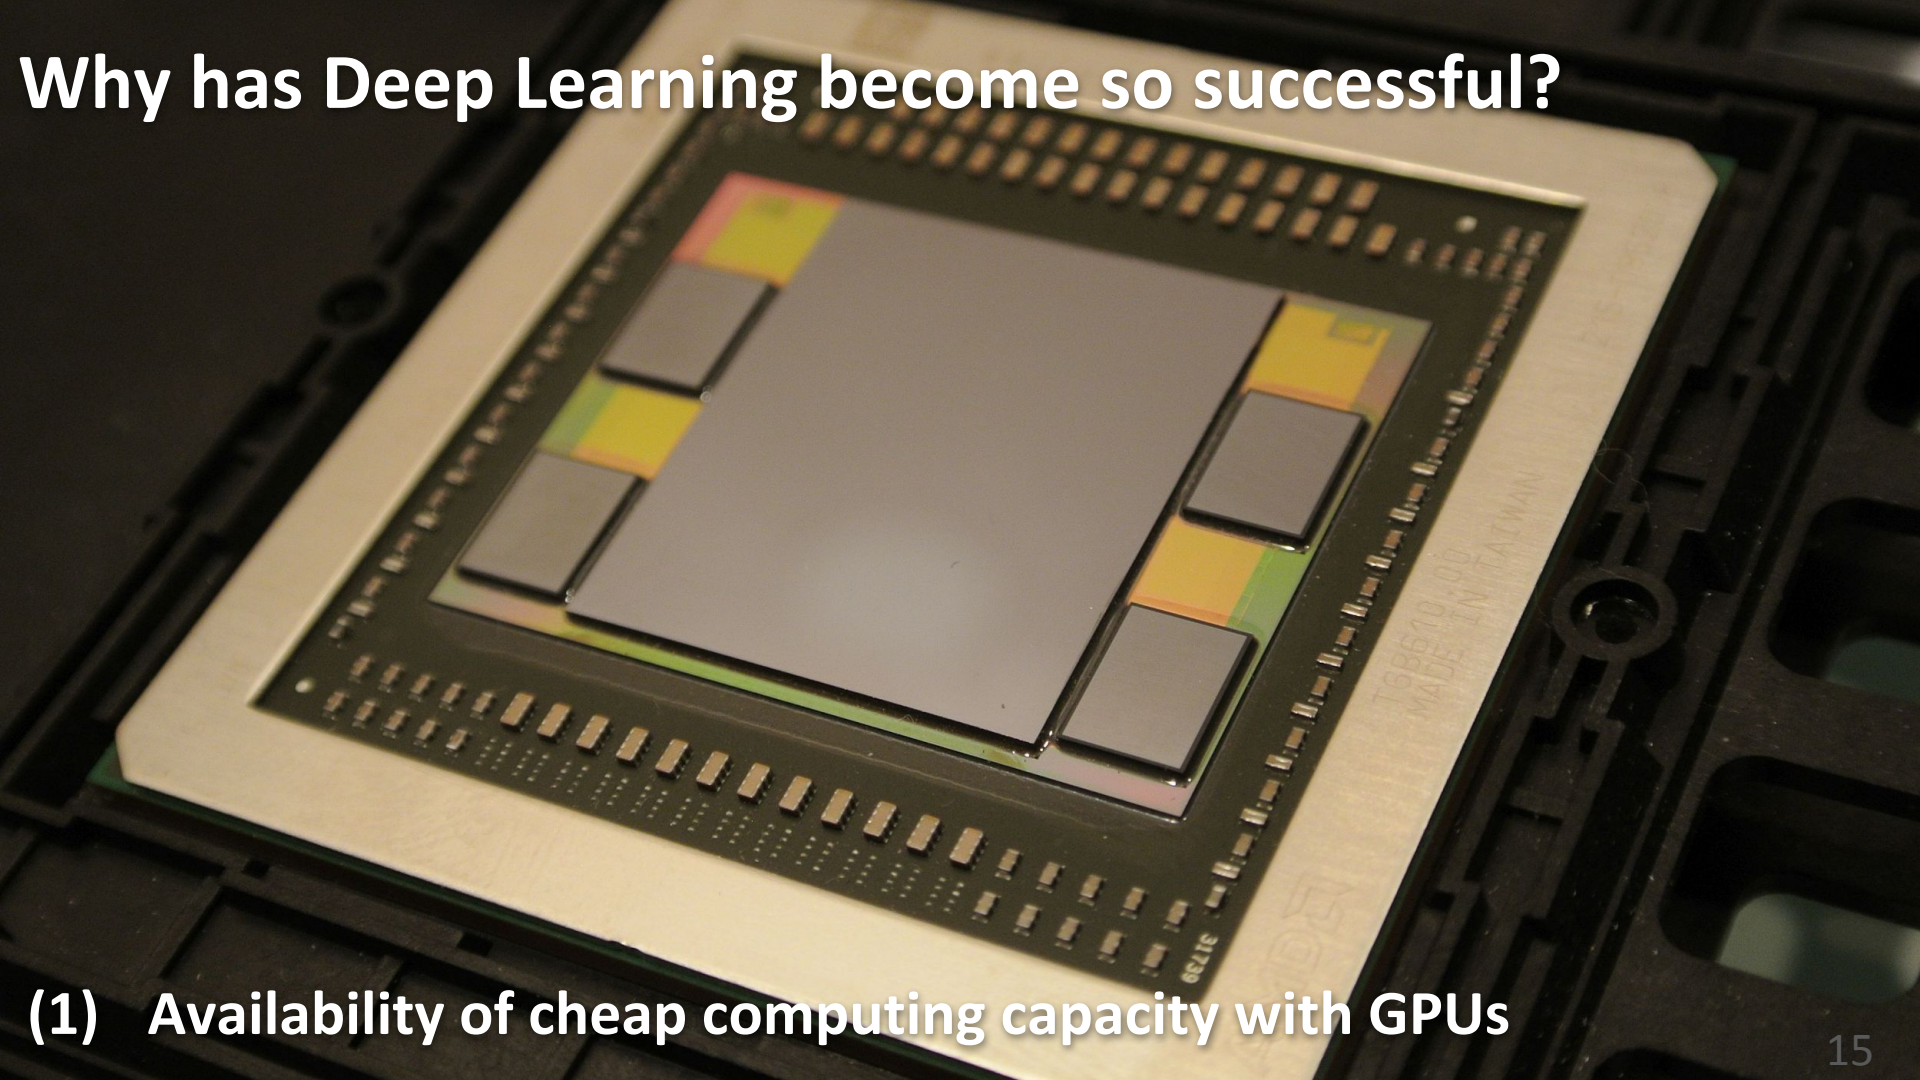

(1) Availability of cheap computing capacity with GPUs

(2) Reusable Highly Trained Complex Models

“AlexNet”

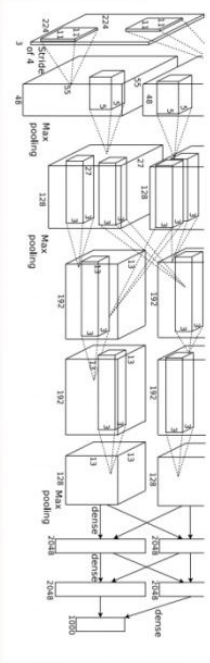

[Krizhevsky et al. NIPS 2012]

“GoogLeNet”

[Szegedy et al. CVPR 2015]

“VGG Net”

[Simonyan & Zisserman, ICLR 2015]

“ResNet”

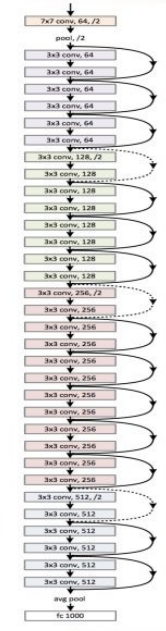

[He et al. CVPR 2016]

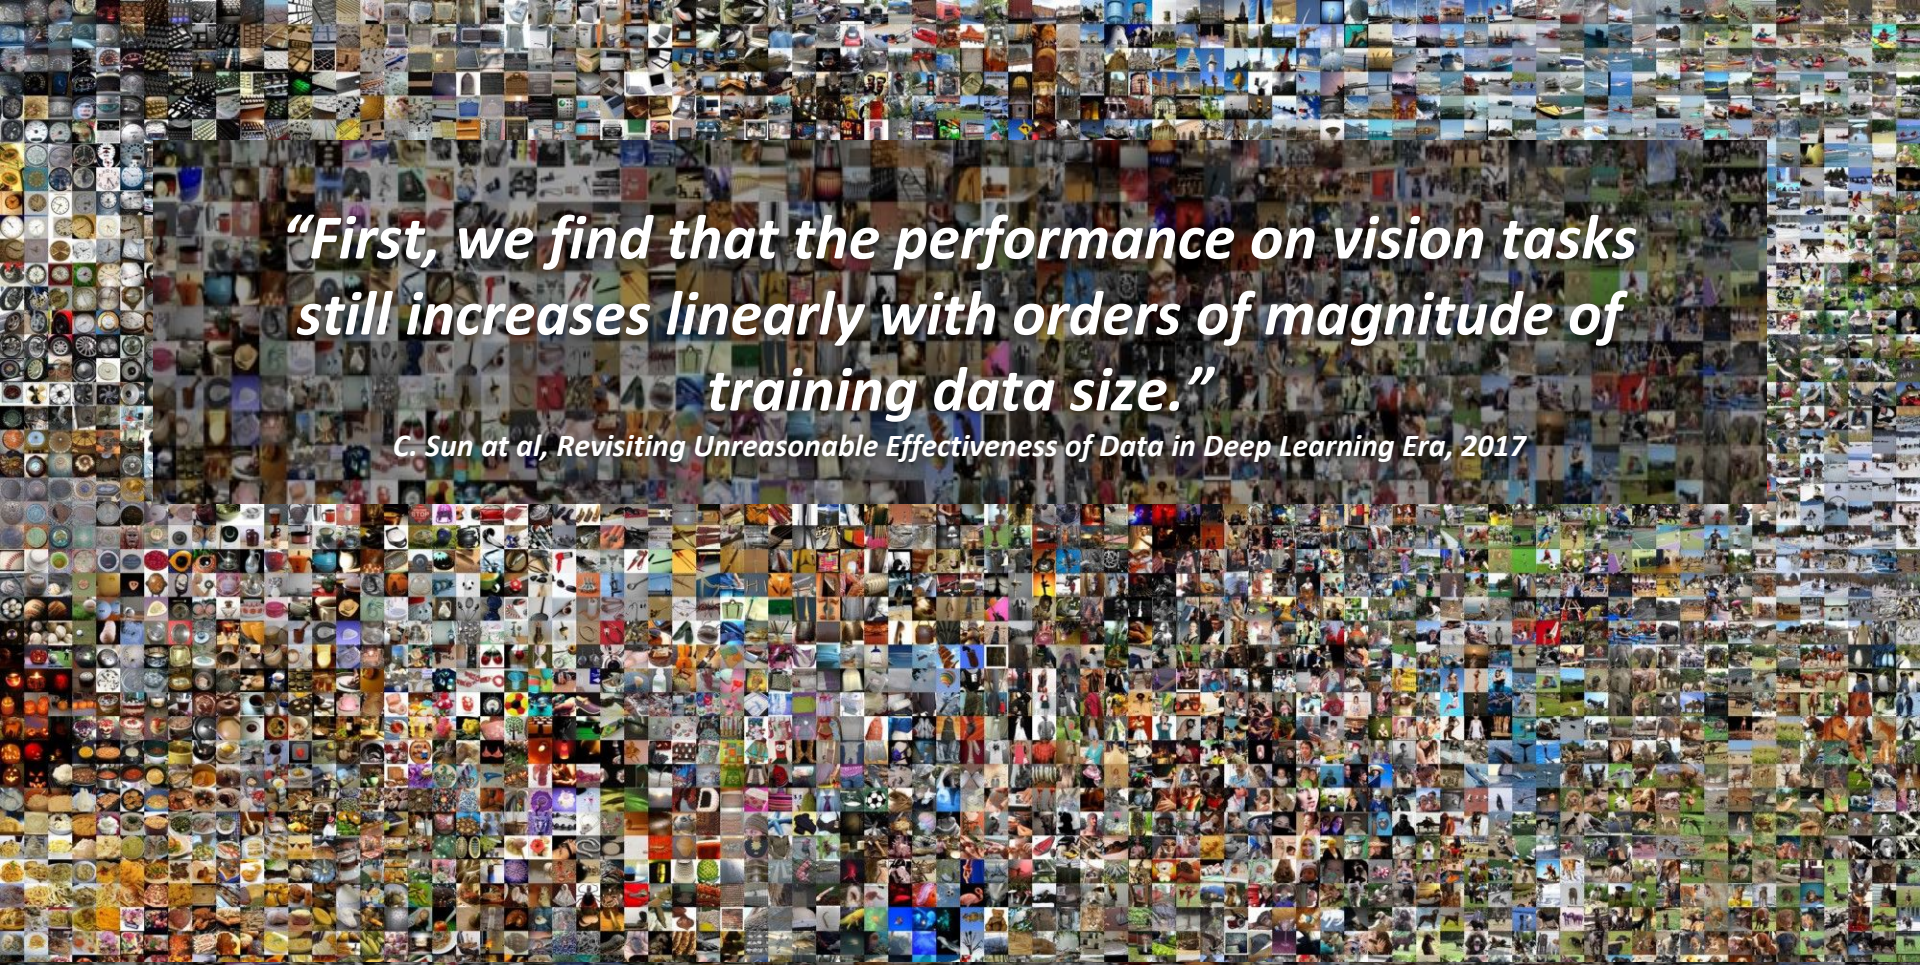

(3) Availability of Large Scale Labeled Data

Deep Learning for Visual Analysis

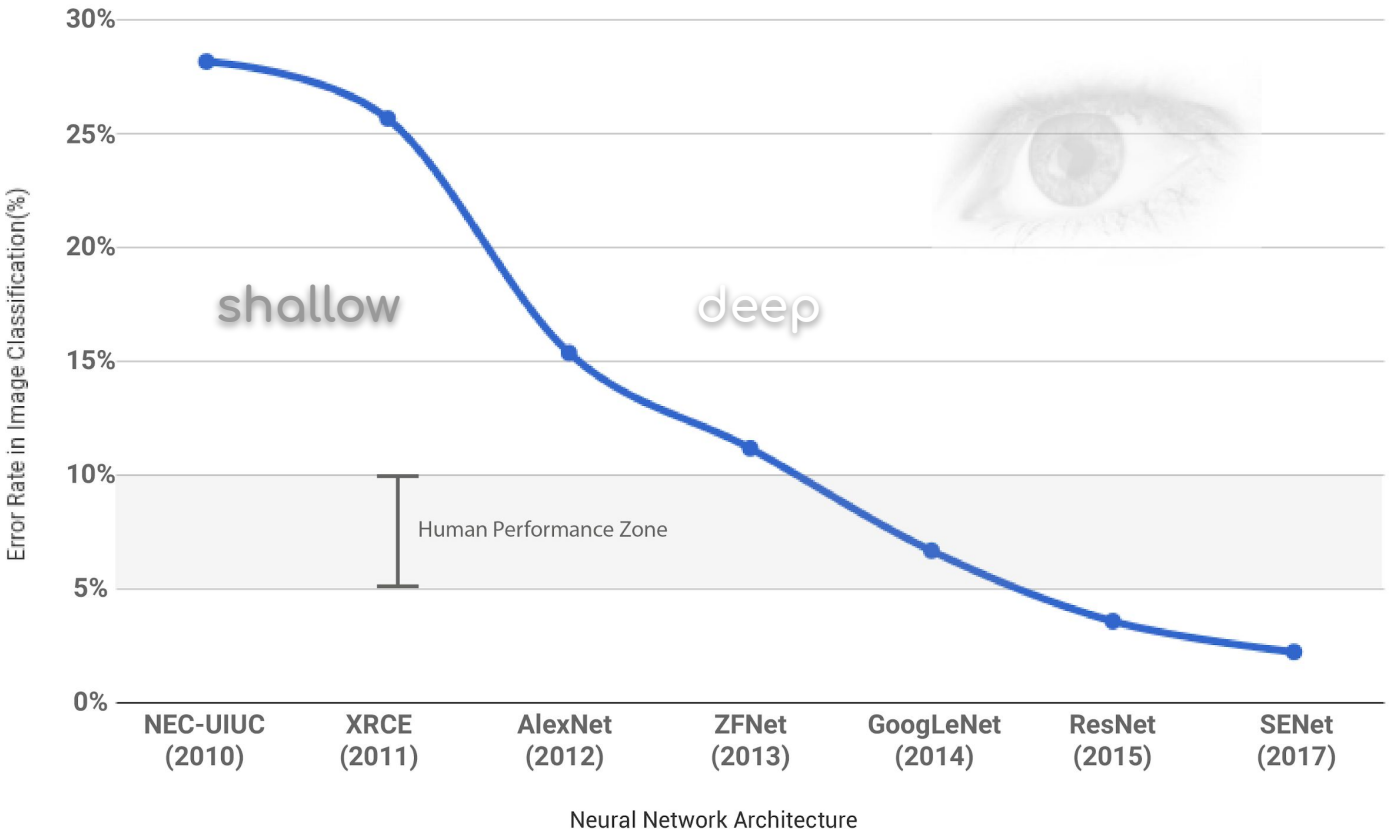

IMAGENET

Large Scale Visual Recognition Challenge (ILSVRC)

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

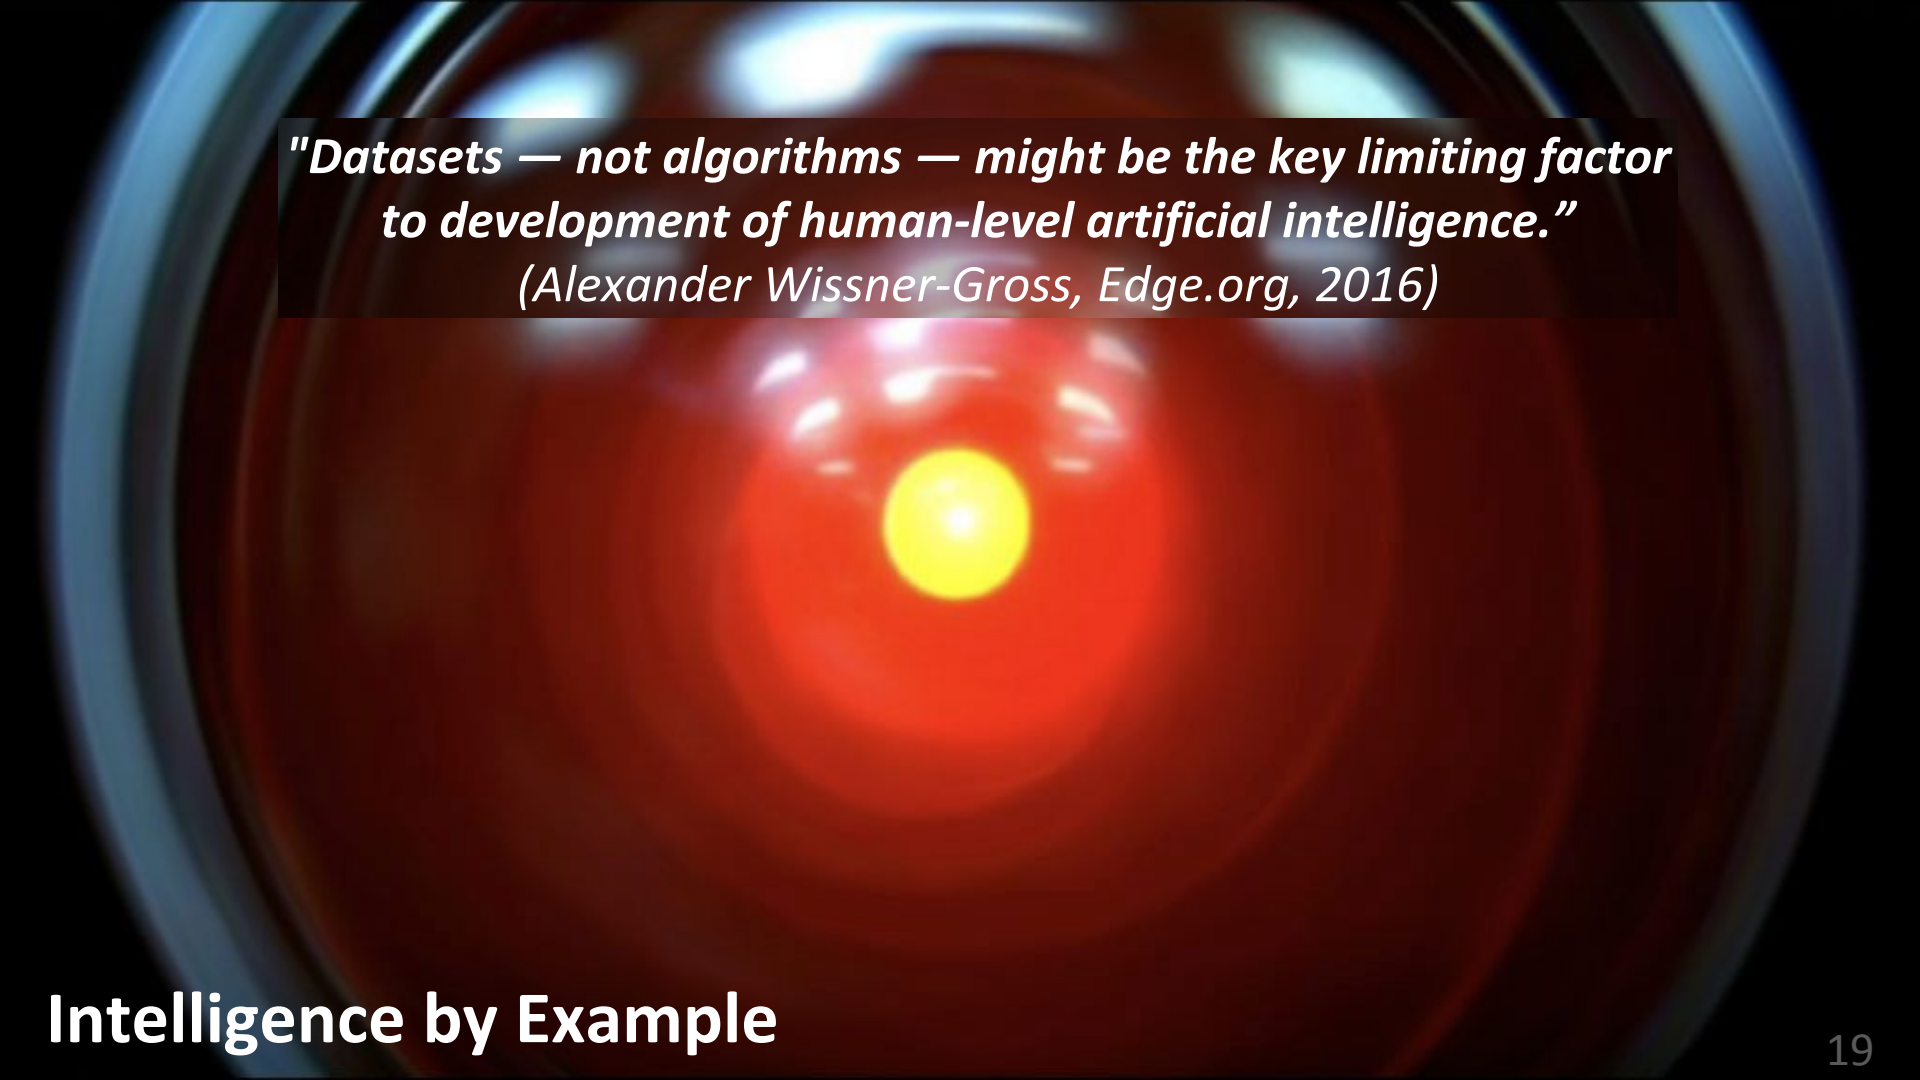

Intelligence by Example

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

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

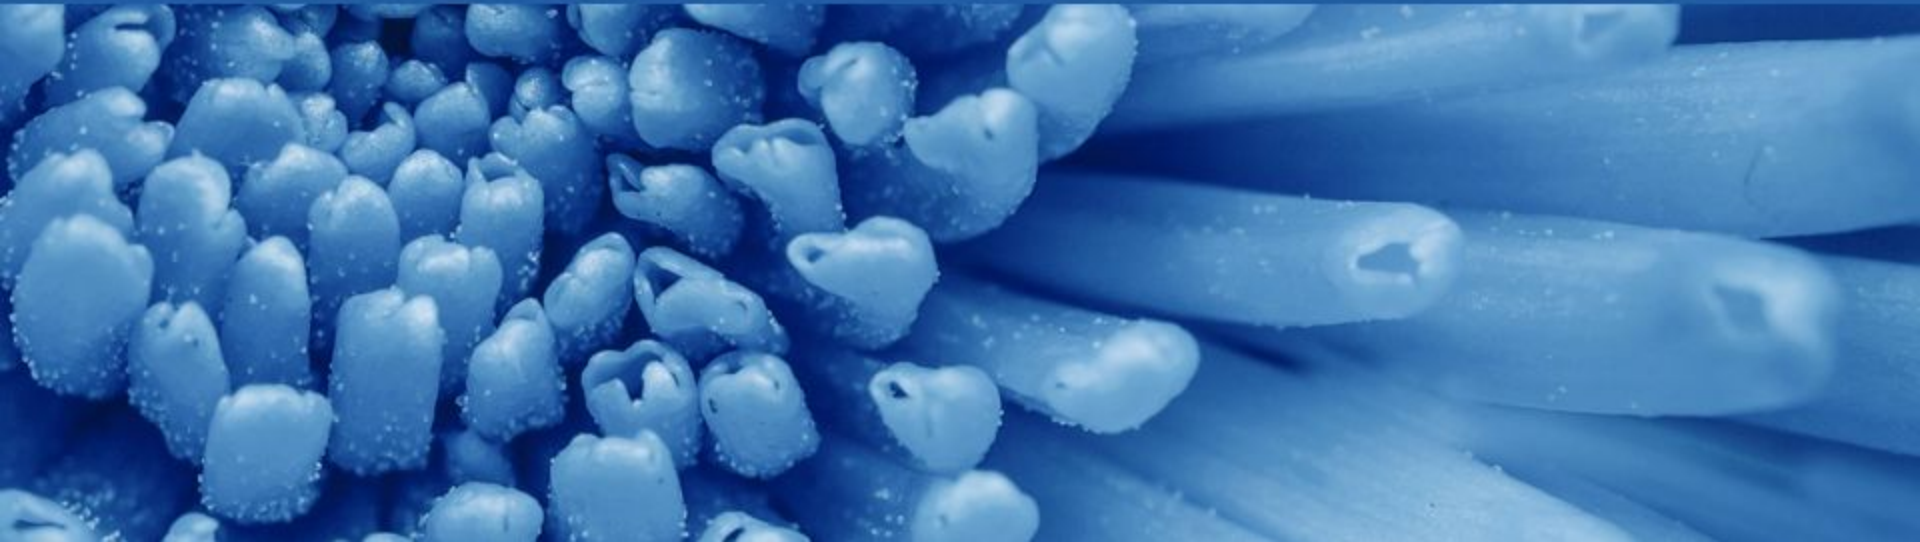

From Classification to Generation

Leon A. Gatys, Alexander S. Ecker, Matthias Bethge:
A Neural Algorithm of Artistic Style. CoRR abs/1508.06576 (2015)

Creation of New Content - Cross Domain Transfer

Monet → Photograph

https://junyanz.github.io/CycleGAN/images/teaser_high_res.jpg

Comparative Learning - Generative Adversarial Networks

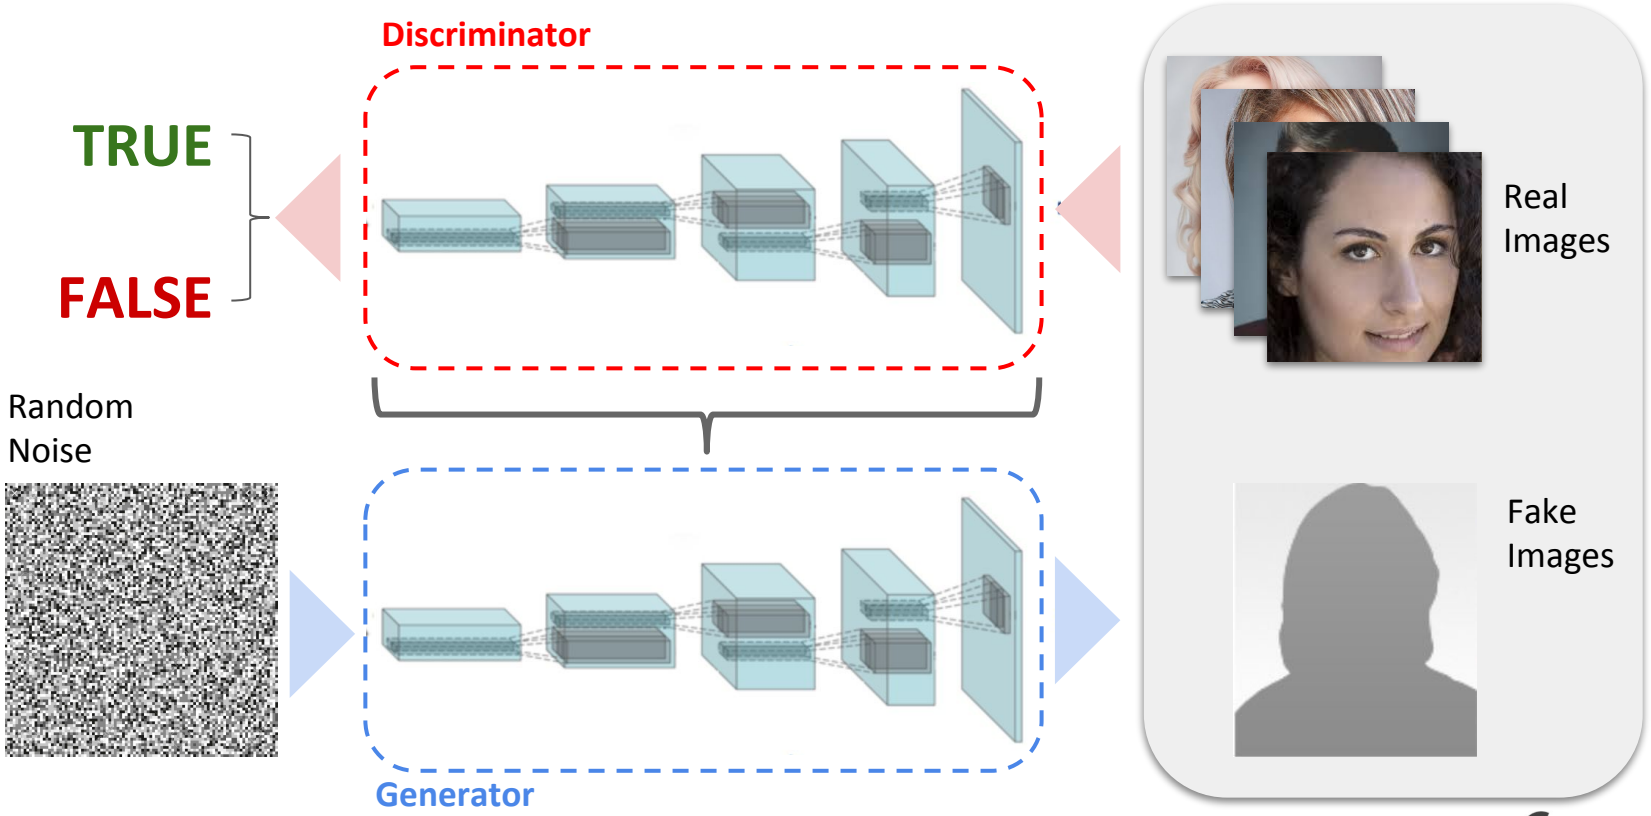

Schmidthuber, 1992
Niemitalo, 2010
Goodfellow, 2014

https://towardsdatascience.com/generative-adversarial-networks-explained-34472718707a

Creation of New Content - Cross Domain Transfer

https://junyanz.github.io/CycleGAN/

https://github.com/jantic/DeOldify

Generation of New Content - Colourization

Generation of New Content - Super Resolution

bicubic
(21.59dB/0.6423)

SRResNet
(23.53dB/0.7832)

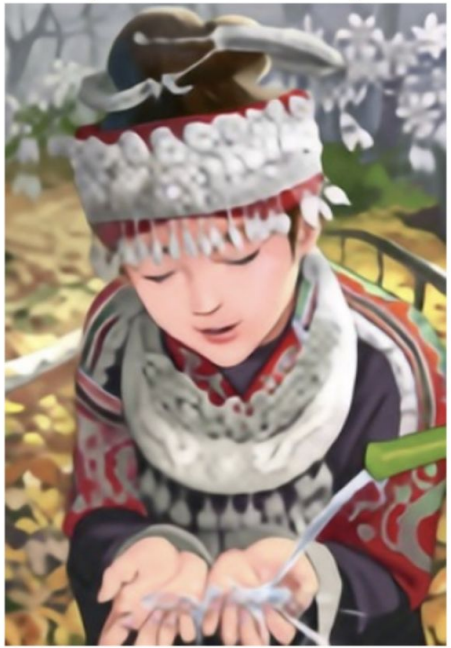

SRGAN
(21.15dB/0.6868)

original

https://arxiv.org/pdf/1609.04802.pdf

Generation of New Content - Image Completion

https://syncedreview.com/2020/06/18/from-texts-to-kitties-openais-gpt-language-model-tackles-image-generation/

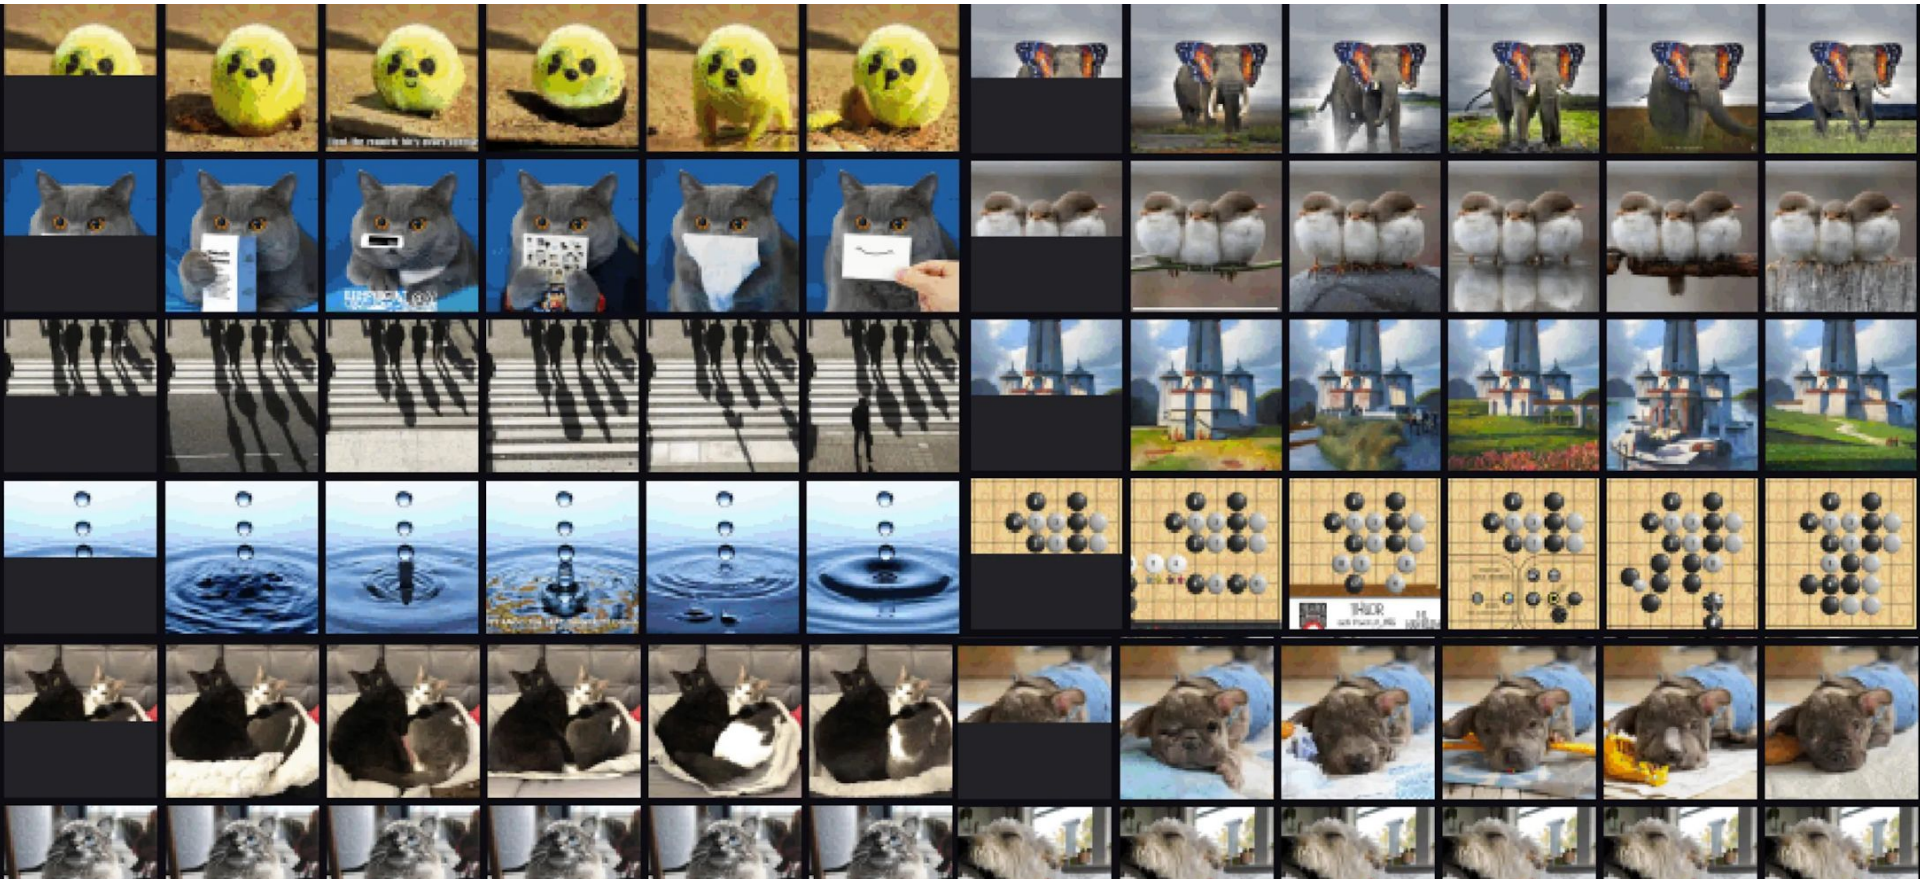

Generation of New Content - Text to Images

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)

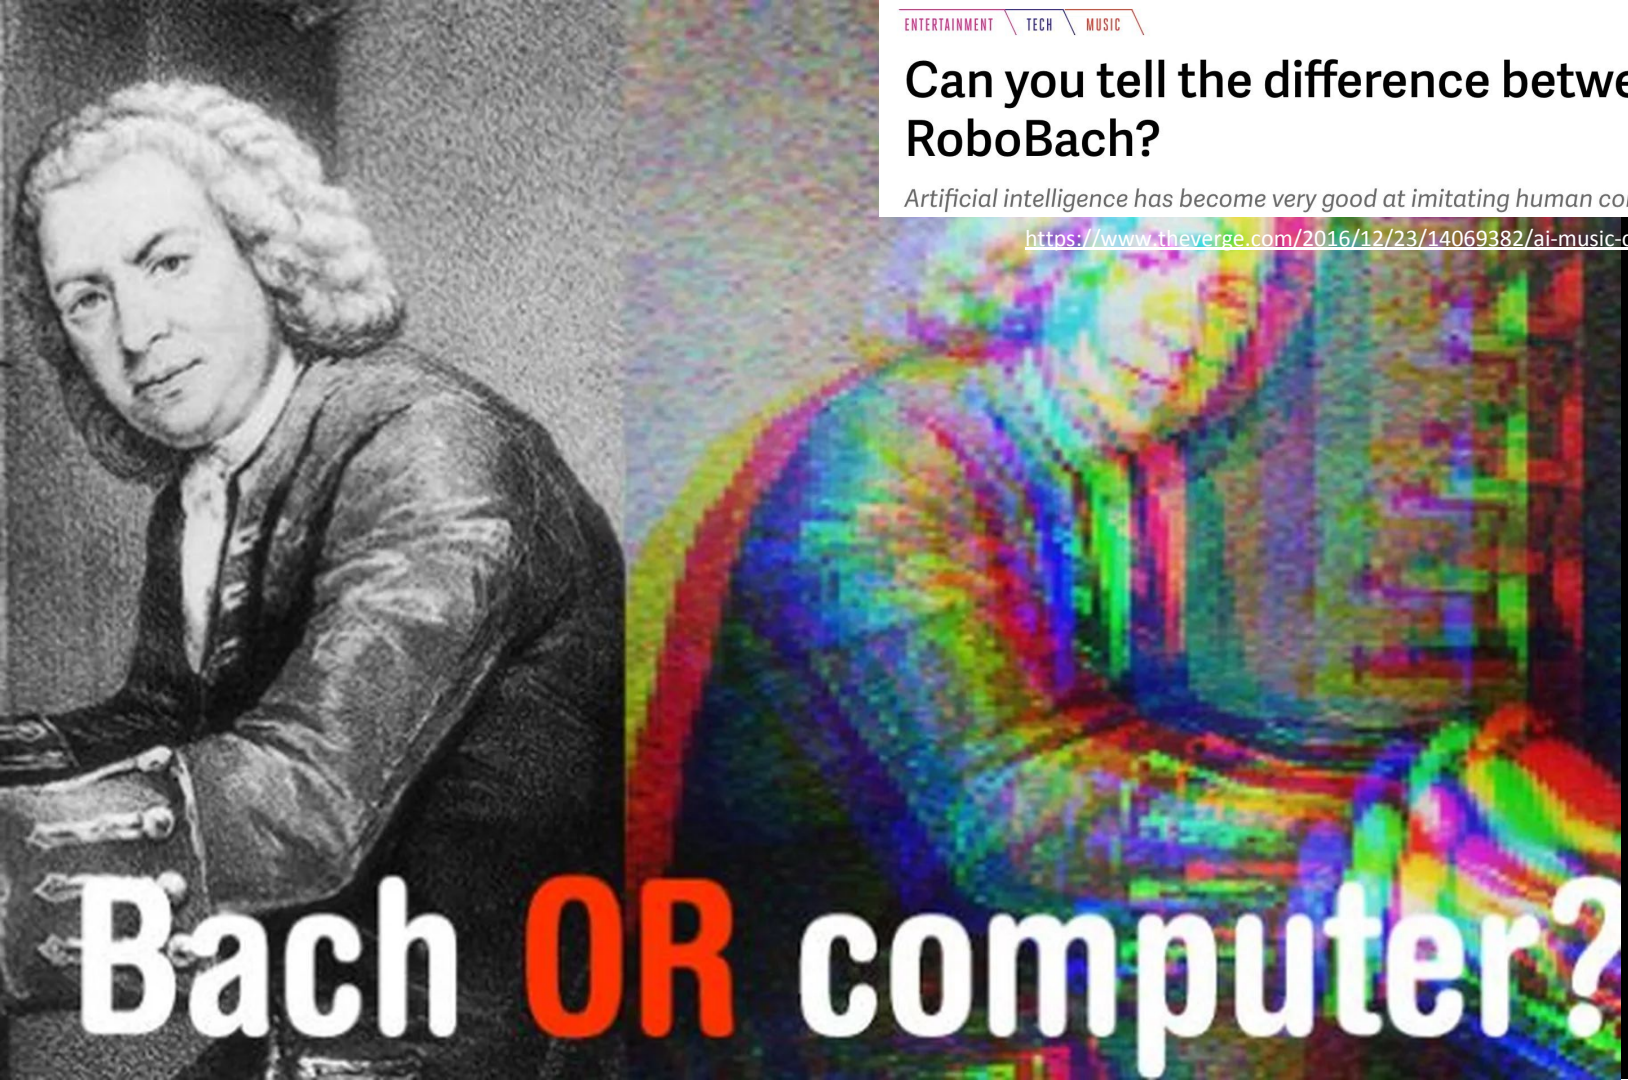

ENTERTAINMENT | TECH | MUSIC

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

https://www.youtube.com/watch?v=MwtVkPKx3RA

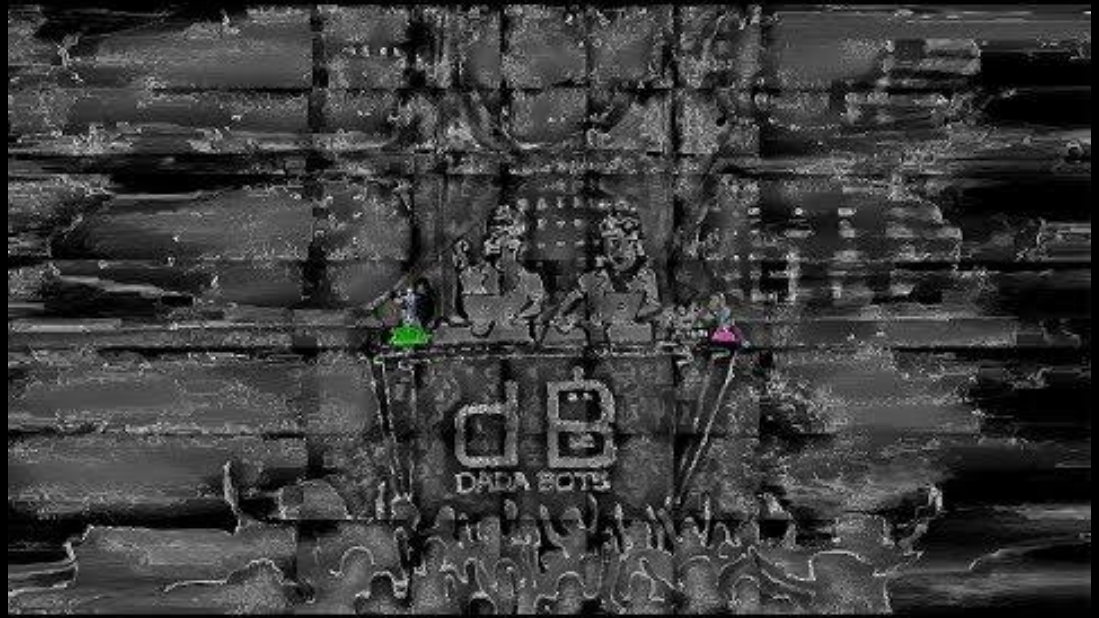

AI generates non-stop stream of death metal

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

The Clever Hans Effect

or Why we shouldn't always trust ML

https://commons.wikimedia.org/wiki/File:Osten_und_Hans.jpg

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

Neil Gaiman, The Invisible Labyrinth (1990)

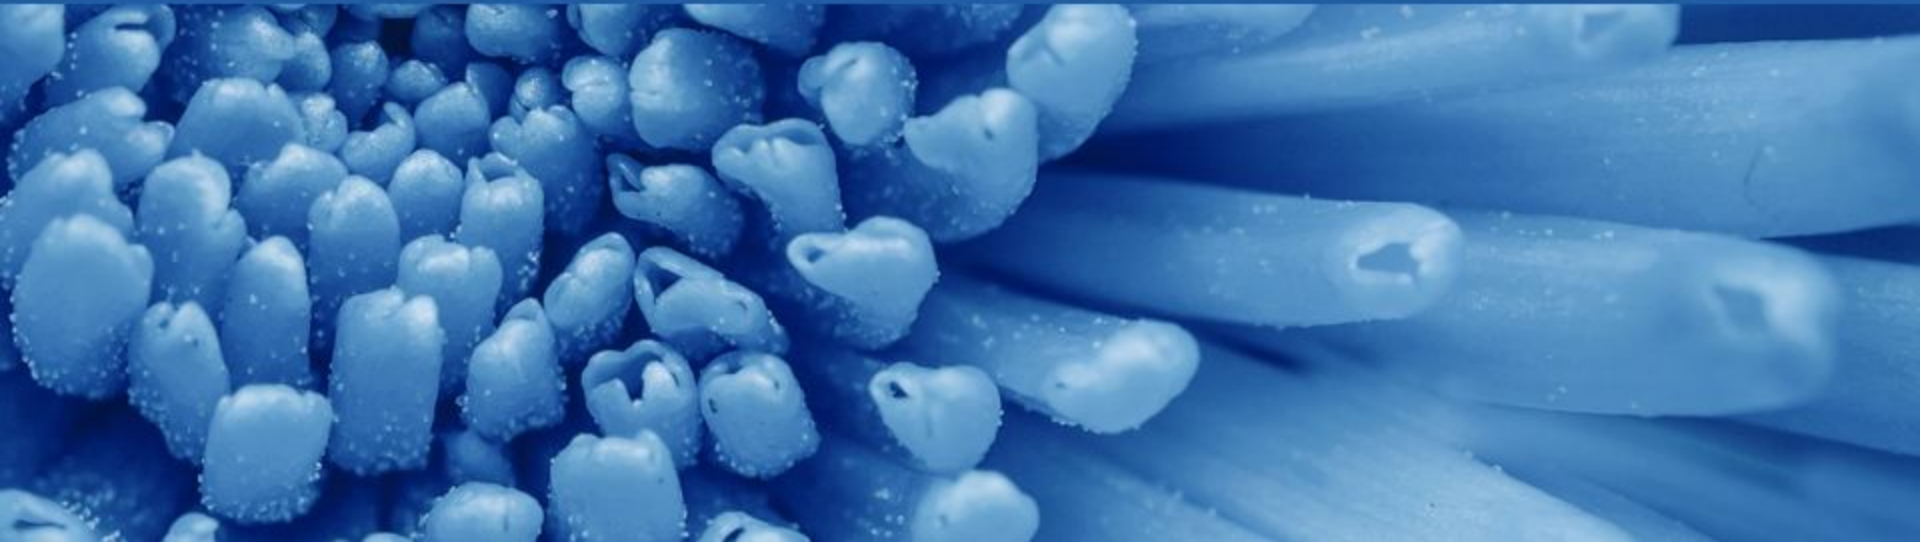

Artificial Intelligence vs Machine Learning

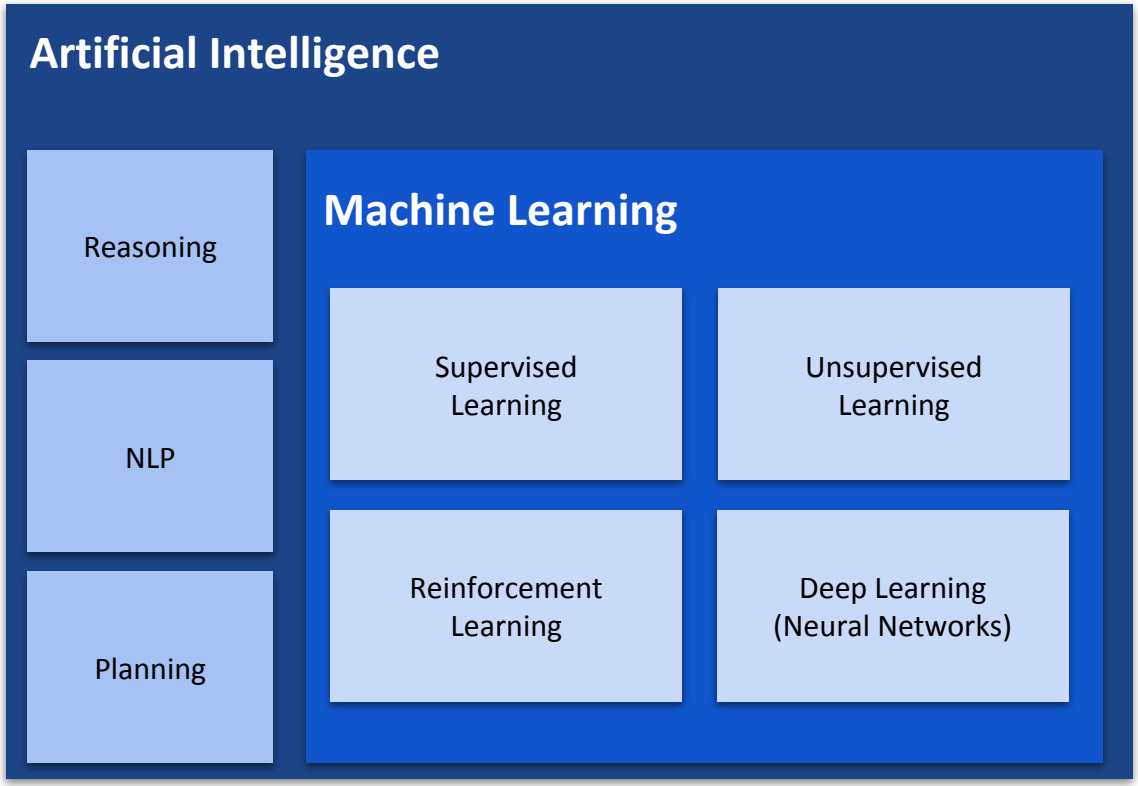

"The Goal of AI is to develop machines that behave as though they were intelligent."

- John McCarthy (1955)

34

Indexing with Pretrained State-of-the-Art Models

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

(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),
```

FIZ Karlsruhe
Leibniz-Institut für Informationsinfrastruktur

PARENTAL
ADVISORY
EXPLICIT SEMANTICS

The Universal Categories - Aristotle (384–322 BCE)

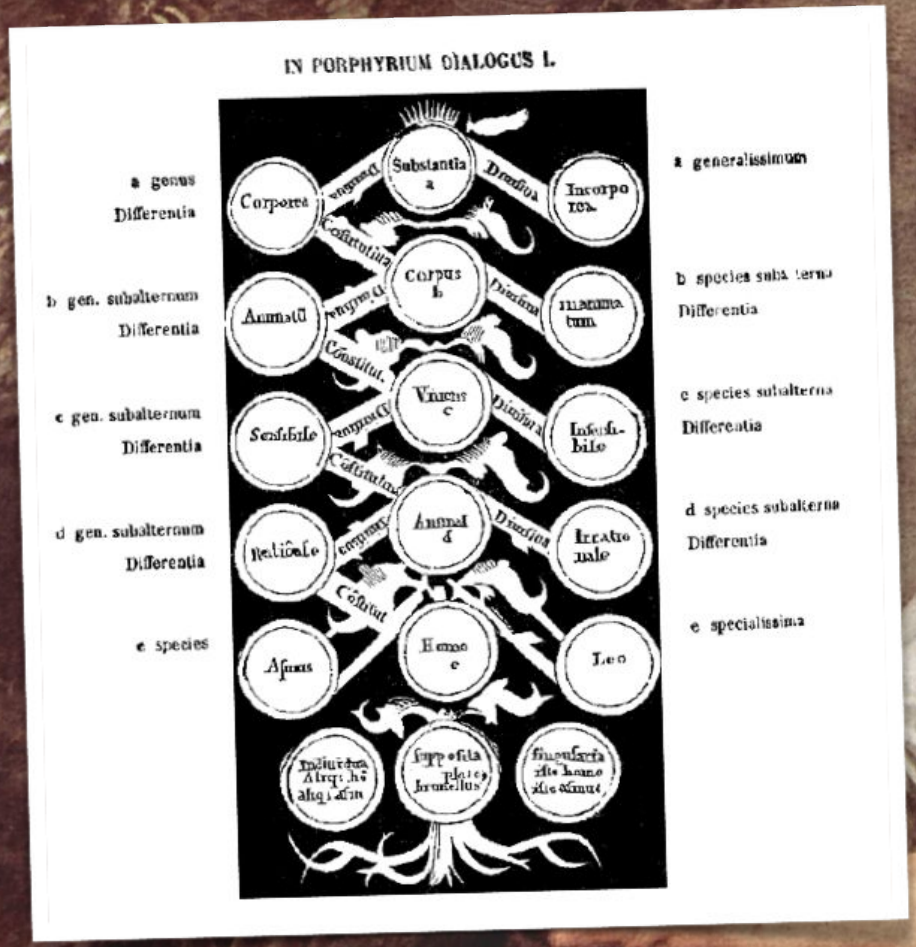

Calculus Ratiocinator - Gottfried Wilhelm Leibniz (1646-1716)

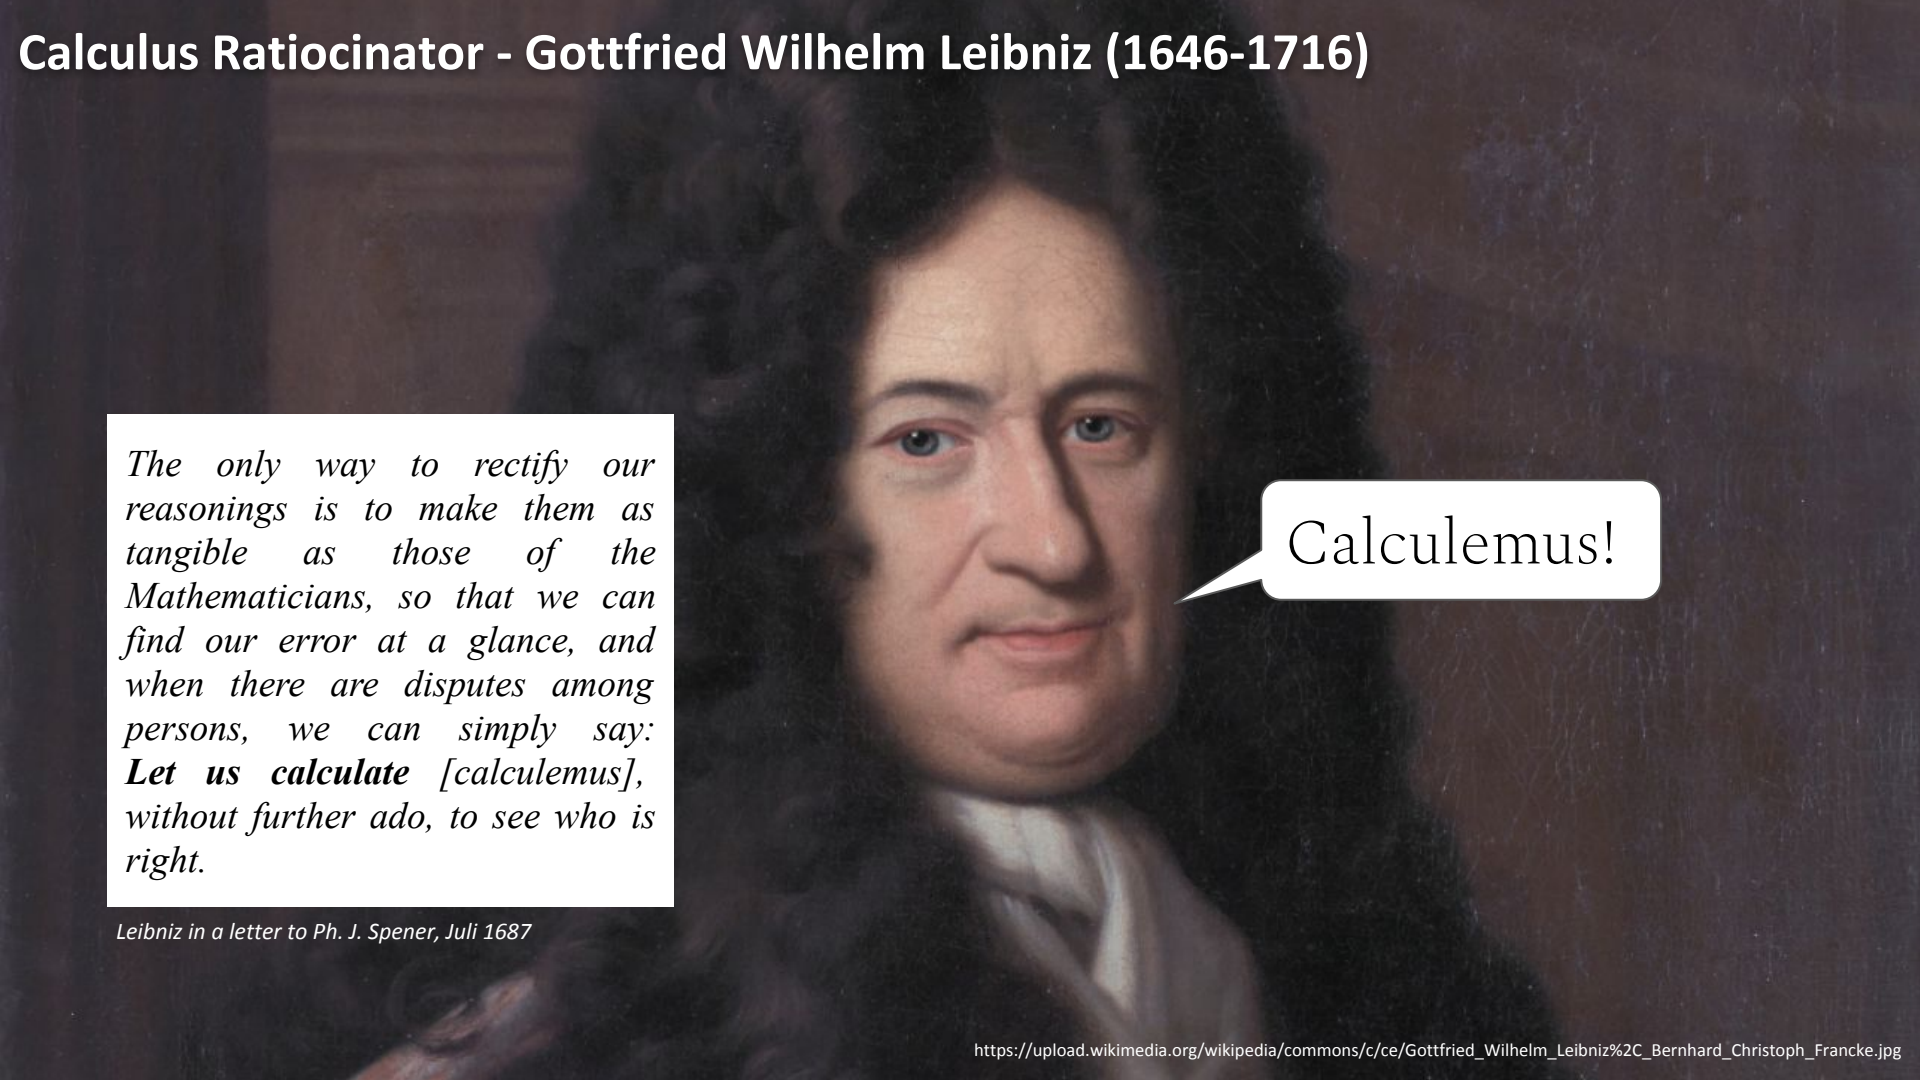

The only way to rectify our reasonings is to make them as tangible as those of the Mathematicians, so that we can find our error at a glance, and when there are disputes among persons, we can simply say: ***Let us calculate*** *[calculemus], without further ado, to see who is right.*

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

https://upload.wikimedia.org/wikipedia/commons/c/ce/Gottfried_Wilhelm_Leibniz%2C_Bernhard_Christoph_Francke.jpg

Begriffsschrift - Gottlob Frege (1848-1925)

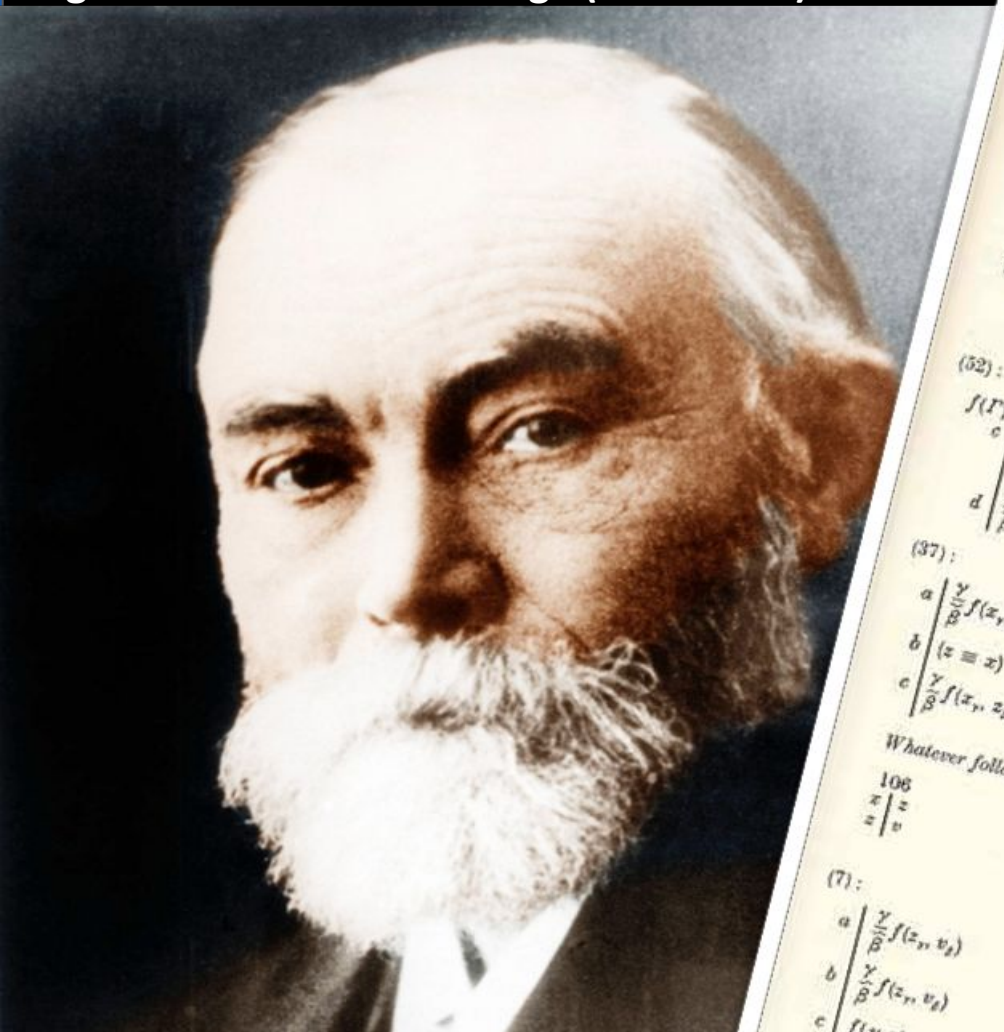

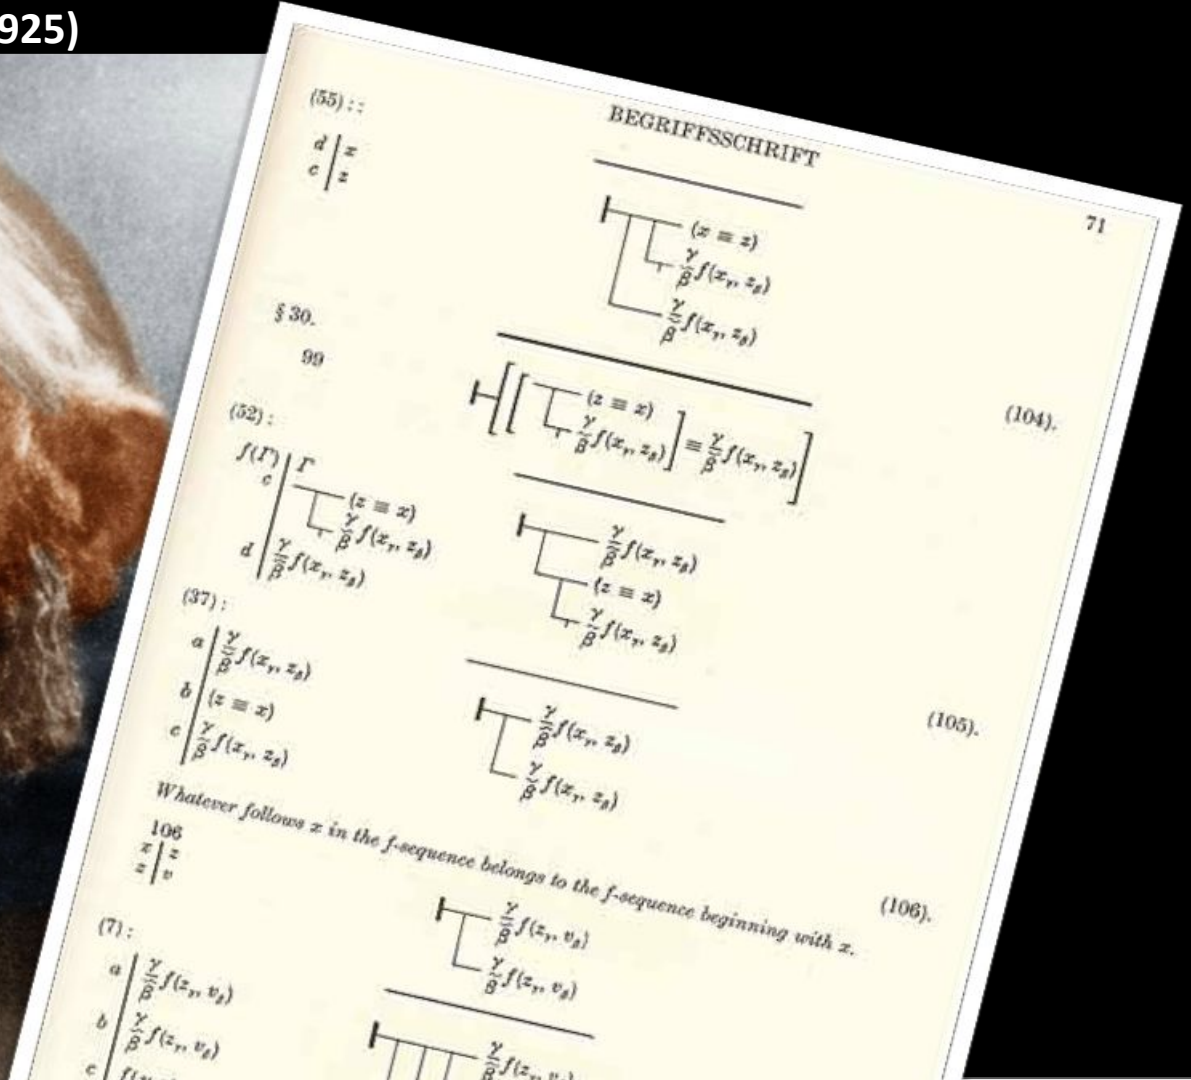

BEGRIFFSSCHRIFT

71

(55) : :

$d \mid x$
$c \mid z$

$(x \equiv z)$
$\frac{\gamma}{\beta} f(x_\gamma, z_\beta)$
$\frac{\gamma}{\beta} f(x_\gamma, z_\beta)$

§ 30.

99

$\left[\left[(z \equiv x) \quad \frac{\gamma}{\beta} f(x_\gamma, z_\beta) \right] \equiv \frac{\gamma}{\beta} f(x_\gamma, z_\beta) \right]$ (104).

(52) :

$f(\Gamma) \mid \Gamma$
$c \mid (z \equiv x) \quad \frac{\gamma}{\beta} f(x_\gamma, z_\beta)$
$d \mid \frac{\gamma}{\beta} f(x_\gamma, z_\beta)$

$\frac{\gamma}{\beta} f(x_\gamma, z_\beta)$
$(z \equiv x)$
$\frac{\gamma}{\beta} f(x_\gamma, z_\beta)$

(37) :

$a \mid \frac{\gamma}{\beta} f(x_\gamma, z_\beta)$
$b \mid (z \equiv x)$
$c \mid \frac{\gamma}{\beta} f(x_\gamma, z_\beta)$

$\frac{\gamma}{\beta} f(x_\gamma, z_\beta)$
$\frac{\gamma}{\beta} f(x_\gamma, z_\beta)$ (105).

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

106

$x \mid z$
$z \mid v$

$\frac{\gamma}{\beta} f(z_\gamma, v_\beta)$
$\frac{\gamma}{\beta} f(z_\gamma, v_\beta)$ (106).

(7) :

$a \mid \frac{\gamma}{\beta} f(z_\gamma, v_\beta)$
$b \mid \frac{\gamma}{\beta} f(z_\gamma, v_\beta)$
$c \mid f(\ldots$

$\frac{\gamma}{\beta} f(z_\gamma, v_\beta)$

Frame Logic for Knowledge Representation - Marvin Minsky (1974)

https://www.flickr.com/photos/x-ray_delta_one/

Ontologies

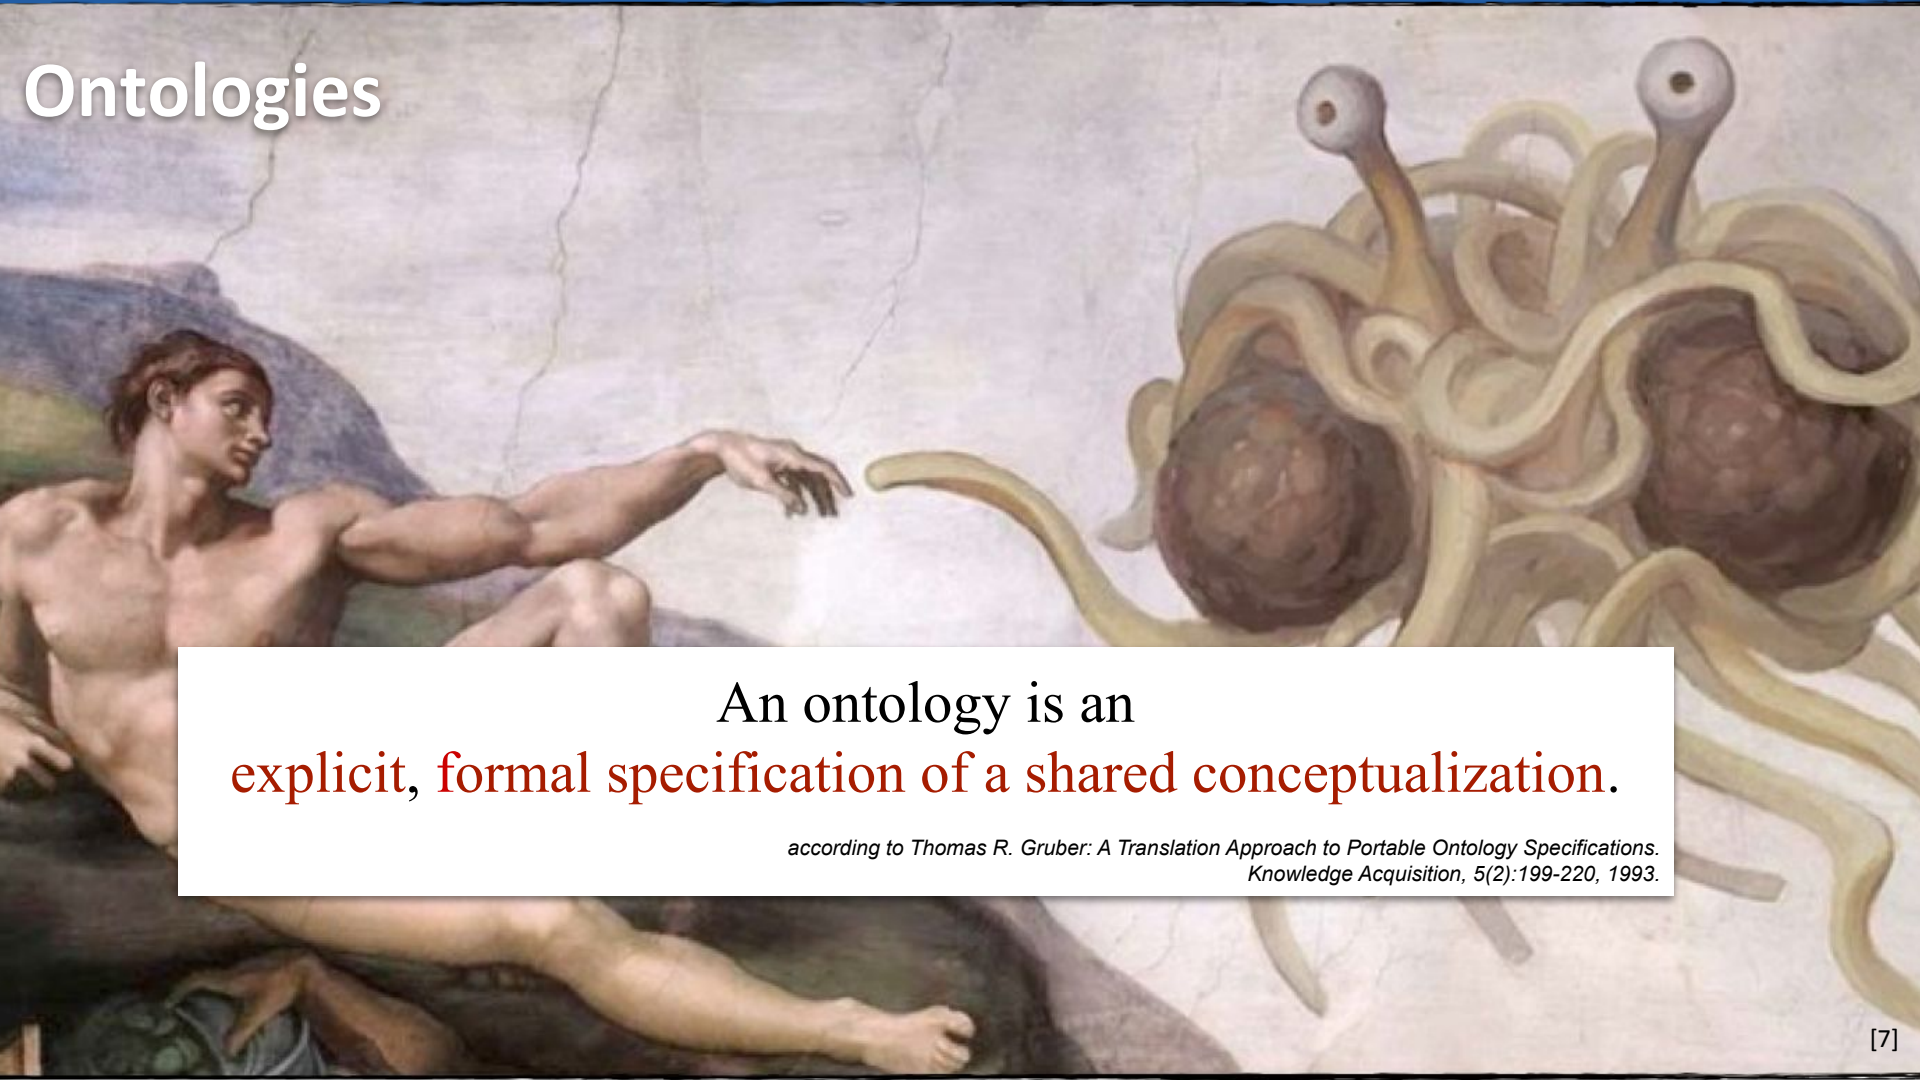

An ontology is an
explicit, formal specification of a shared conceptualization.

according to Thomas R. Gruber: A Translation Approach to Portable Ontology Specifications. Knowledge Acquisition, 5(2):199-220, 1993.

What is an Ontology?

An ontology is an
explicit, formal specification of a shared conceptualization.

according to Thomas R. Gruber: A Translation Approach to Portable Ontology Specifications. Knowledge Acquisition, 5(2):199-220, 1993.

conceptualization:	abstract model (domain, identified relevant concepts, relations)
explicit:	meaning of all concepts must be defined
formal:	machine understandable
shared:	consensus about ontology

From Linked Data to Knowledge Graphs

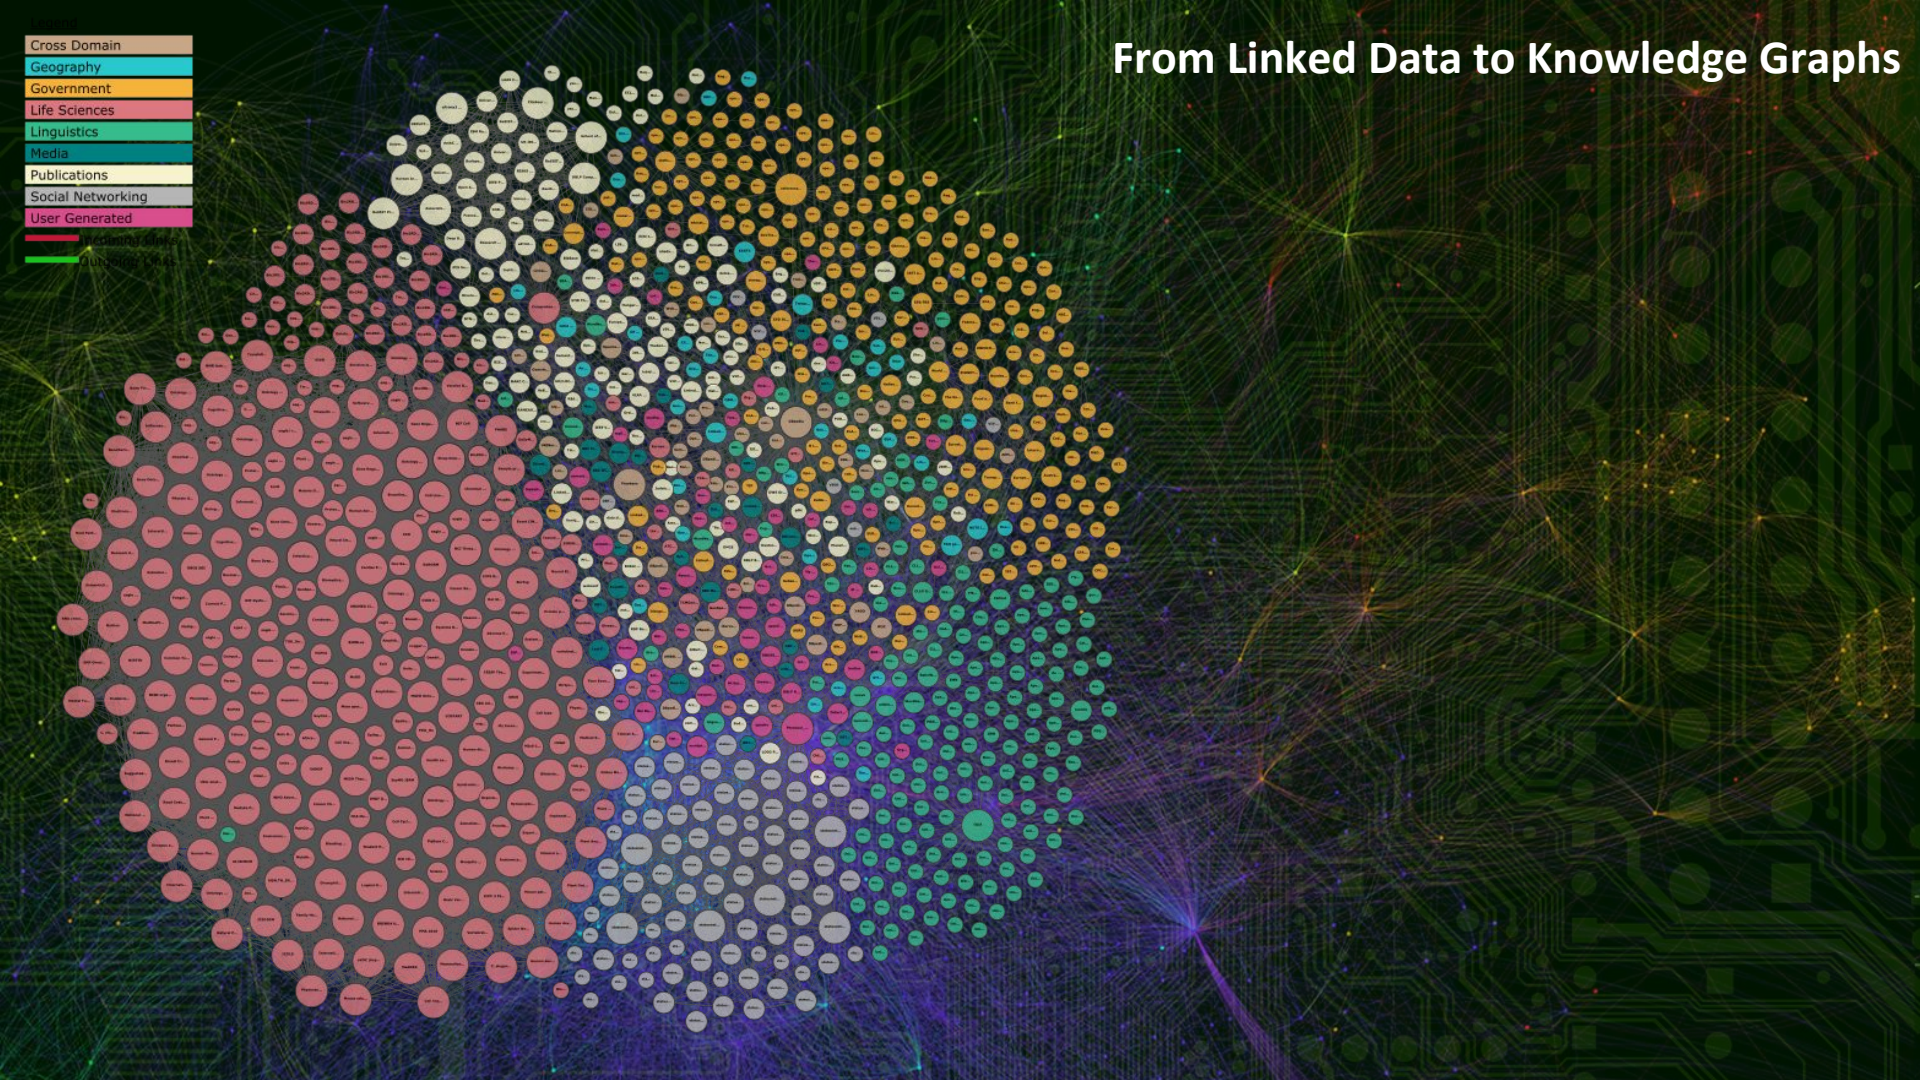

Knowledge Graph Applications

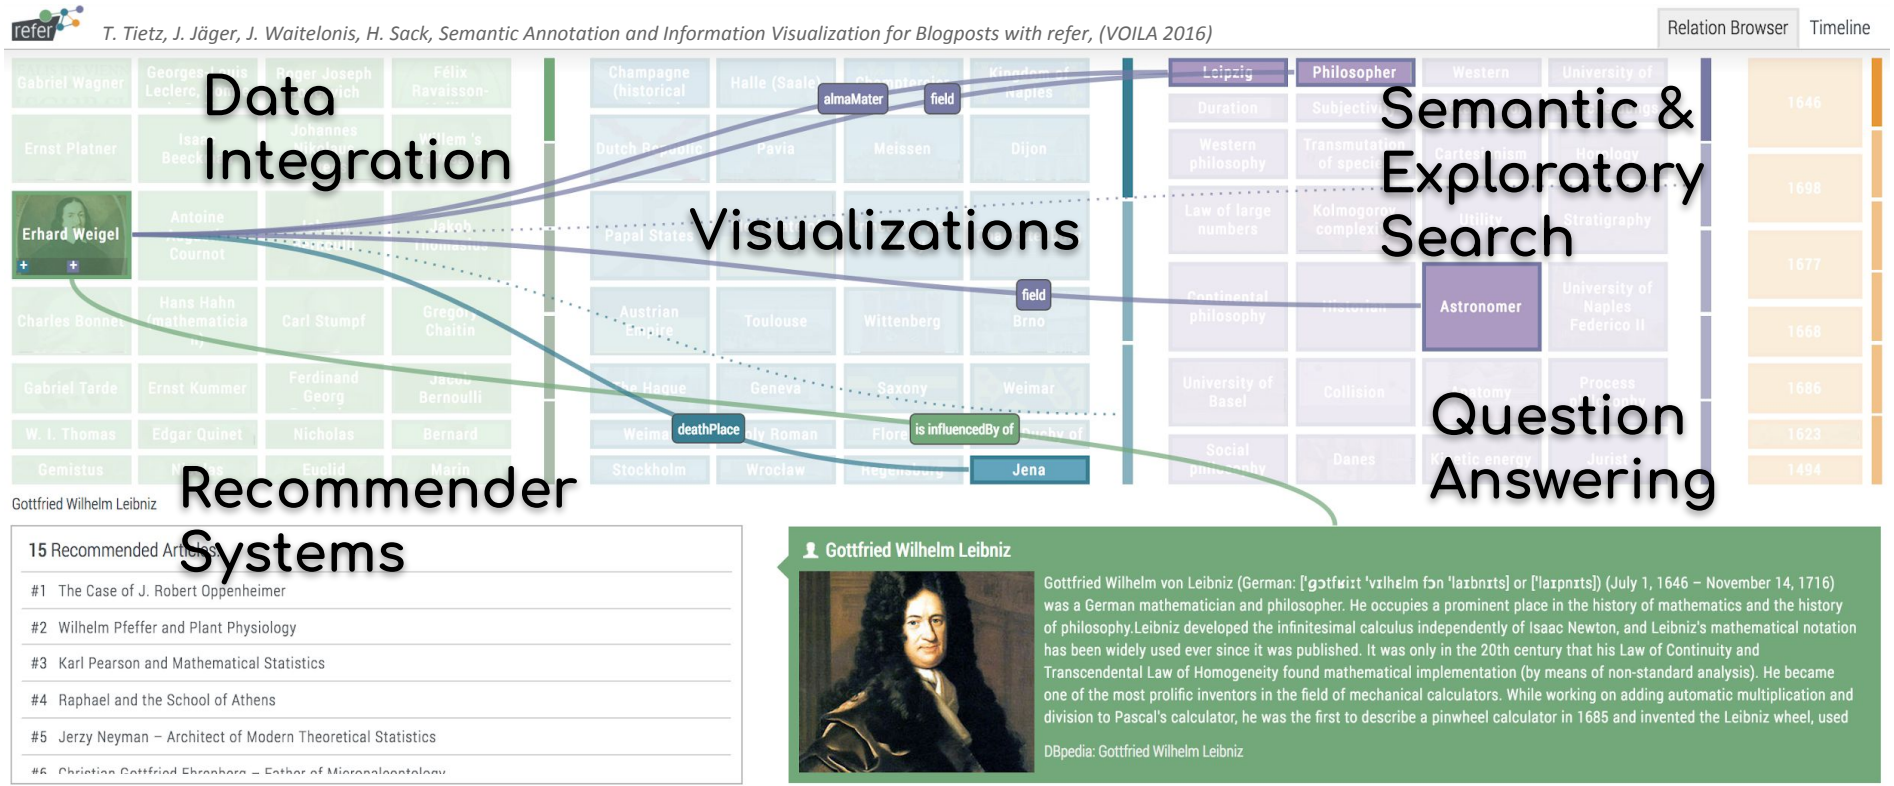

“There ain’t no such thing as a free Lunch.”

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

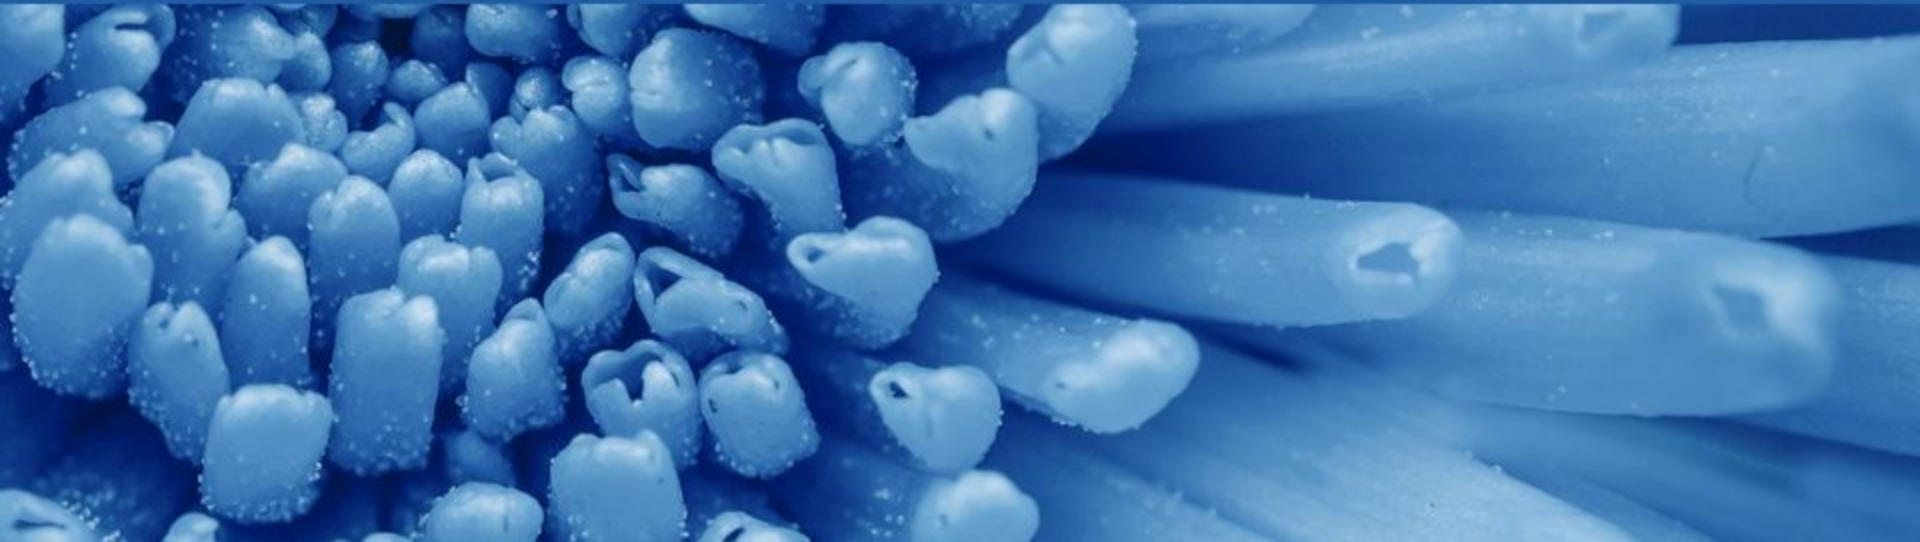

Generation of New Content - Text to Images

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

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)

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

http://2i.ovalenzuelan.com/

A flock of sheep on green meadows

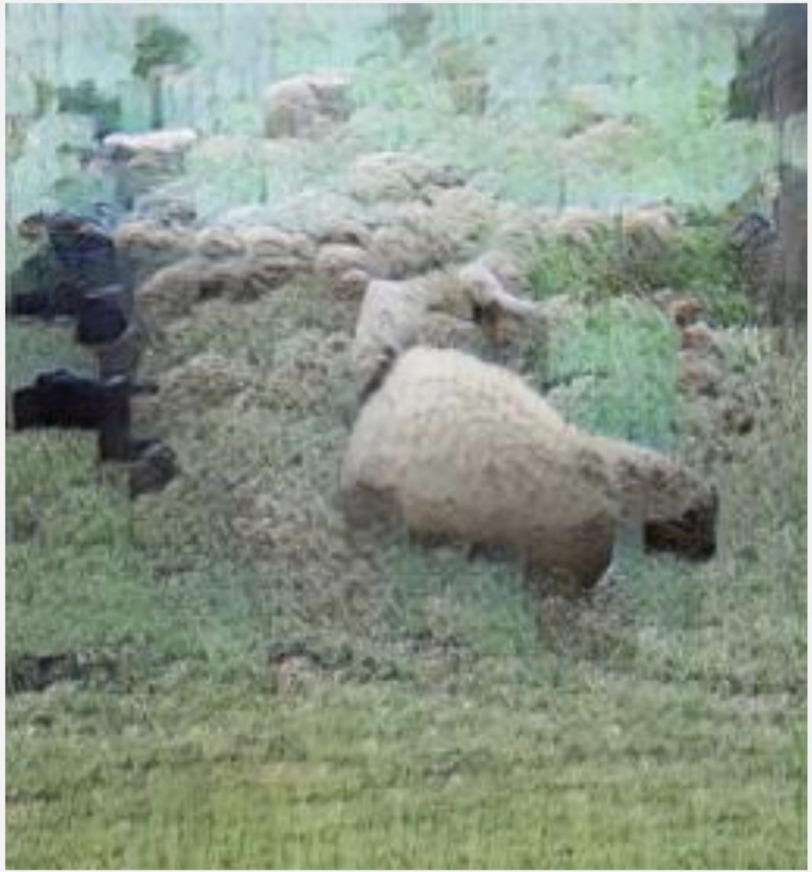

http://t2i.cvalenzuelab.com/

a girl
watching tv

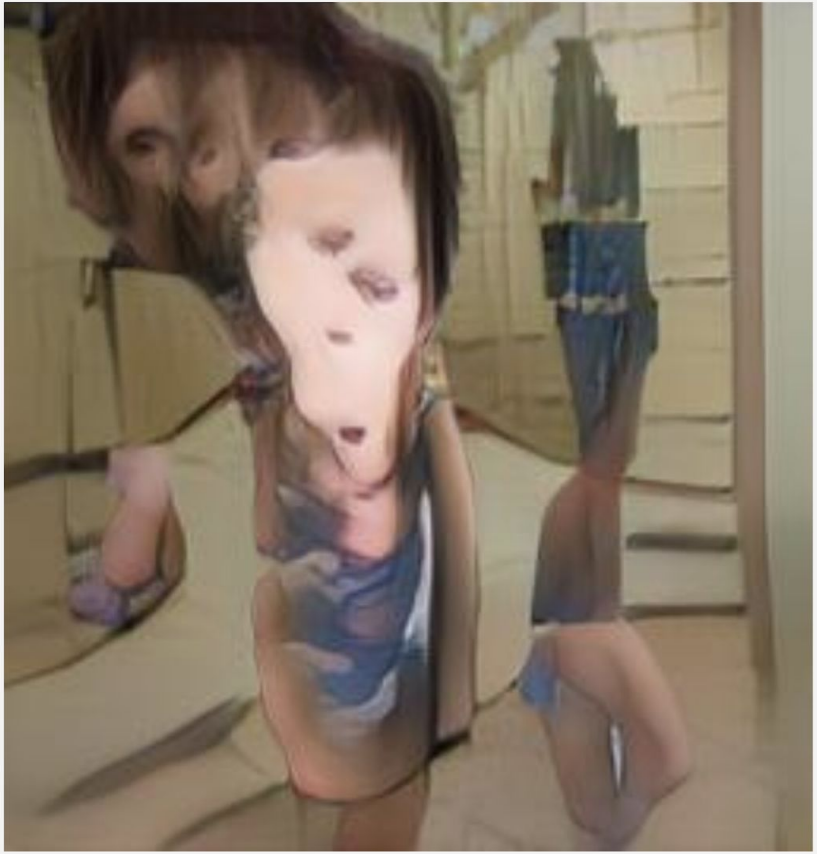

http://t2i.cvalenzuelab.com/

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

http://t2i.cvalenzuelab.com/

melting
clocks over
the dessert

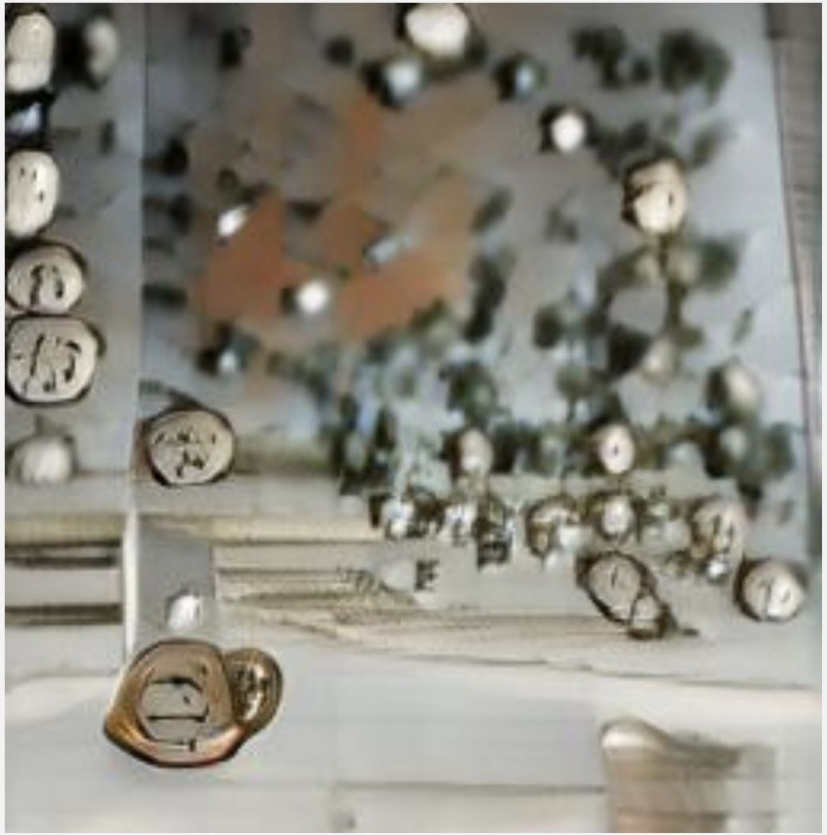

http://t2i.cvalenzuelab.com/

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

Alexandre Dumas, The Three Musketeers, (1844)

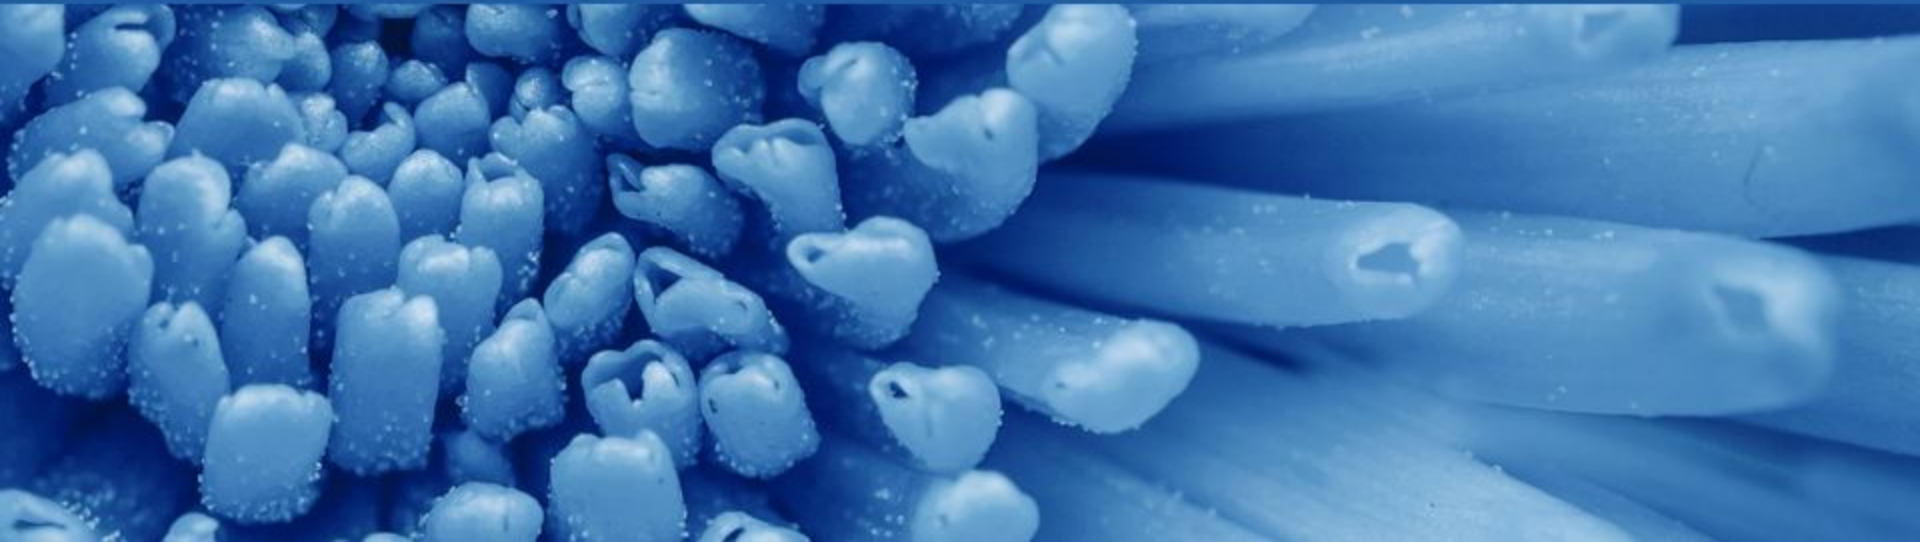

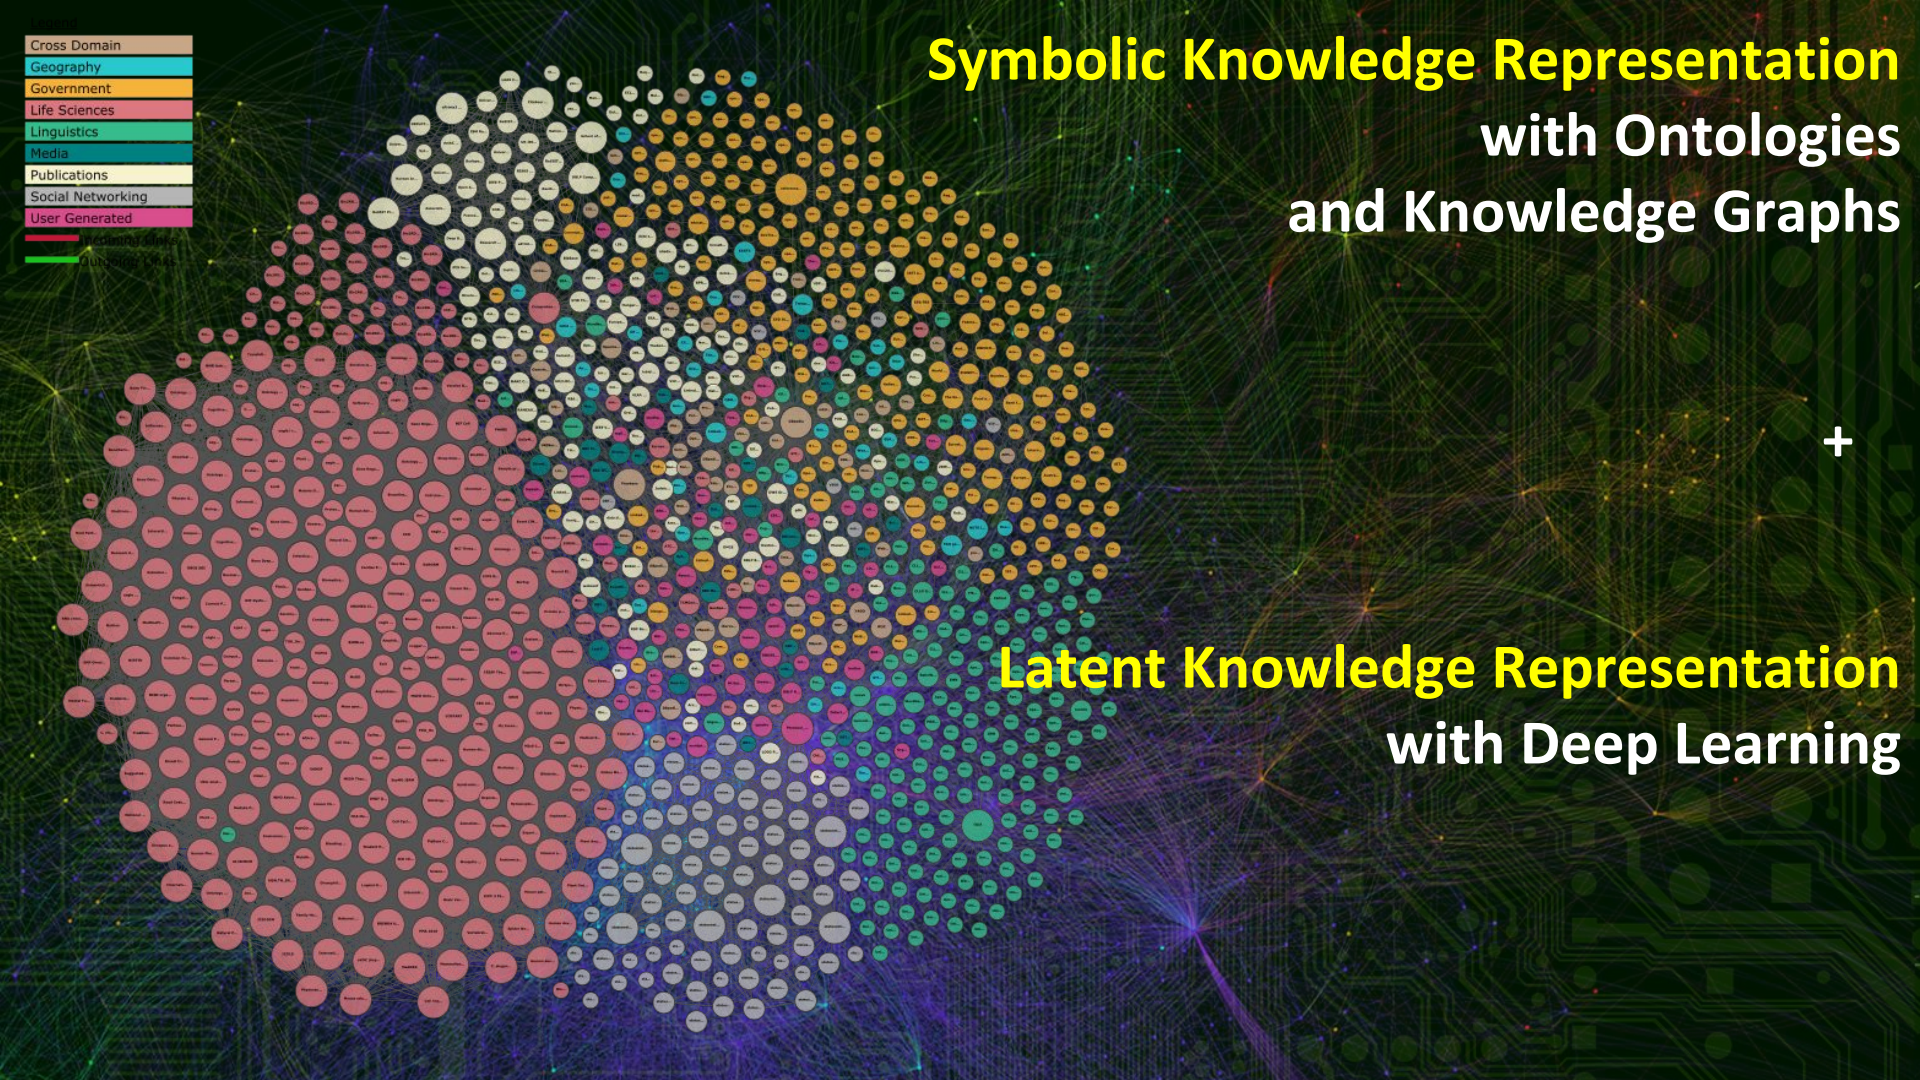
Legend
Cross Domain
Geography
Government
Life Sciences
Linguistics
Media
Publications
Social Networking
User Generated
Symbolic Knowledge Representation
with Ontologies
and Knowledge Graphs
+
Latent Knowledge Representation
with Deep Learning

Combining Deep Learning and Semantics

Automated Image captioning:

a **person** holding a **skateboard**

Entity Linking

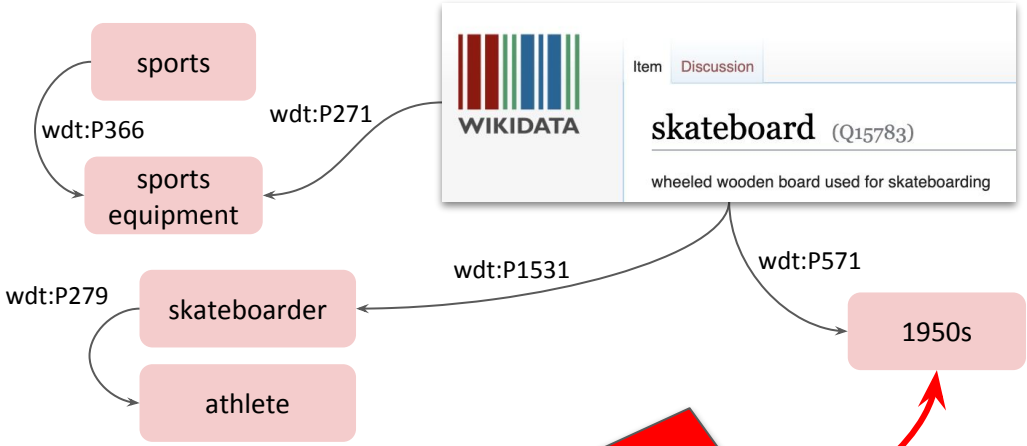

Metadata

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

Knowledge Graph Based Exploration

FIZ Karlsruhe
Leibniz-Institut für Informationsinfrastruktur

Karlsruher Institut für Technologie

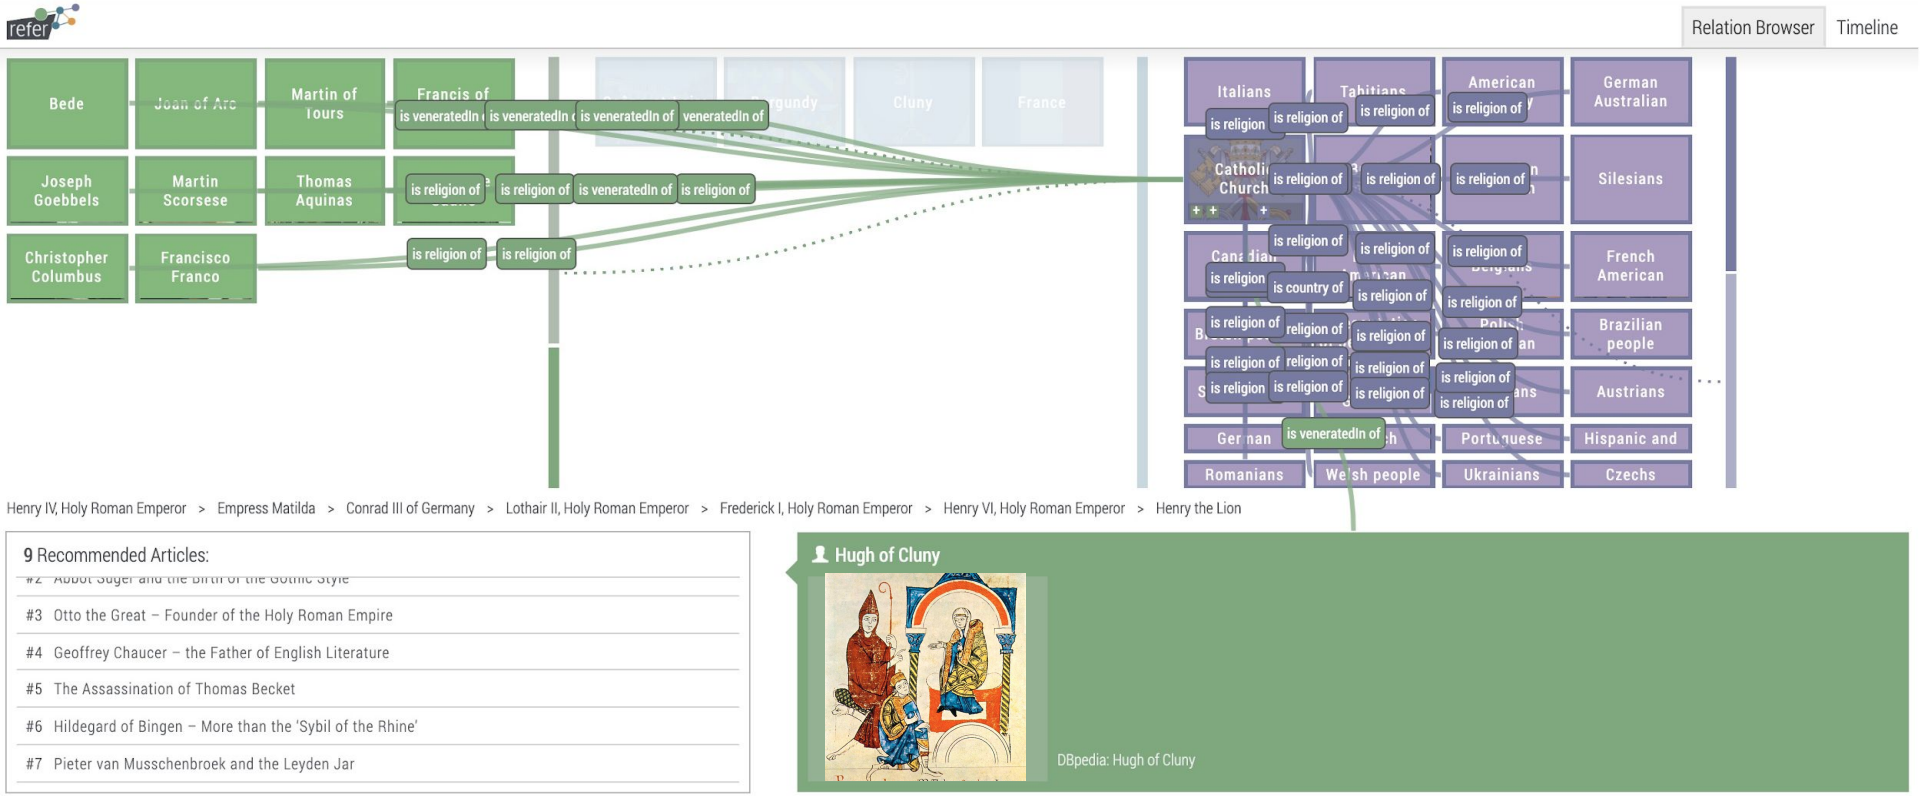

http://scihi.org/

Knowledge Graph Based Exploration

FIZ Karlsruhe
Leibniz-Institut für Informationsinfrastruktur

Karlsruher Institut für Technologie

Q4060918
Q11398608
Yaroslav II of Kiev
Friedrich IV.
Kasimir I.
Richardis von Babenberg
Agnes von Österreich
Béla, Heir of Hungary
Q11801056
Heinrich
Heinrich II.
Judith de Lorraine
Q9140803
Boleslaw I.
Konrad III.
Mieszko I.
Stephan III.
Q933892
Beatrice of Auxonne
Ludmilla von Böhmen
Agnes of Bohemia
Adalbert III. von Böhmen
Jaroslaw
Konrad II. von Babenberg
Johann I.
Boleslaus, son of Dépolt
Balduin V.
Berta von Boll
Agnes von Babenberg
Dulcia II.
Gertrude von Babenberg
Friedrich
Sophie von Böhmen
Heinrich II.
Hedwig von Lothringen
Beatrice, Countess of Chalon
Konrad
Anna von Burgund, Gräfin von Savoyen
Ottone del Carretto, Marquis of Savona
Enrico II Del Carretto
Adalbert II.
Agnes von Waiblingen
Heinrich IV.
Leonor de Aragón
Ferdinando d'Aragona
Wilhelm von Montferrat
Vladimir Mstislavich
Richenza von Everstein
Heinrich V.
Beatrix
Azalaïs of Montferrat
Leopold
Sancha d'Aragona
Beatrice del Monferrato
Roman
Sancha von Kastilien
Guigues VI of Viennois
Jutta von Babenberg
Mstislaw II.
Beatrice Schwäbische
Otto von Freising
Ferdinand of Castille
Peter II.
Konstanze von Aragón
Ladislaus III.
Ambrogio Del Carretto
Konrad
Wsewolod Mstyslawytsch
Alfons II.
Bonifacio I
Rainer von Montferrat
Cunigunda of Lorraine
Heinrich IV.
Friedrich V. von Schwaben
Vsevolod Vsevolodovich
Maria von Montferrat
Bonifatius I.
Heinrich VI.
Friedrich II.
Vasylko Yaropolkpvych
https://w.wiki/CTJ

Visual Embeddings

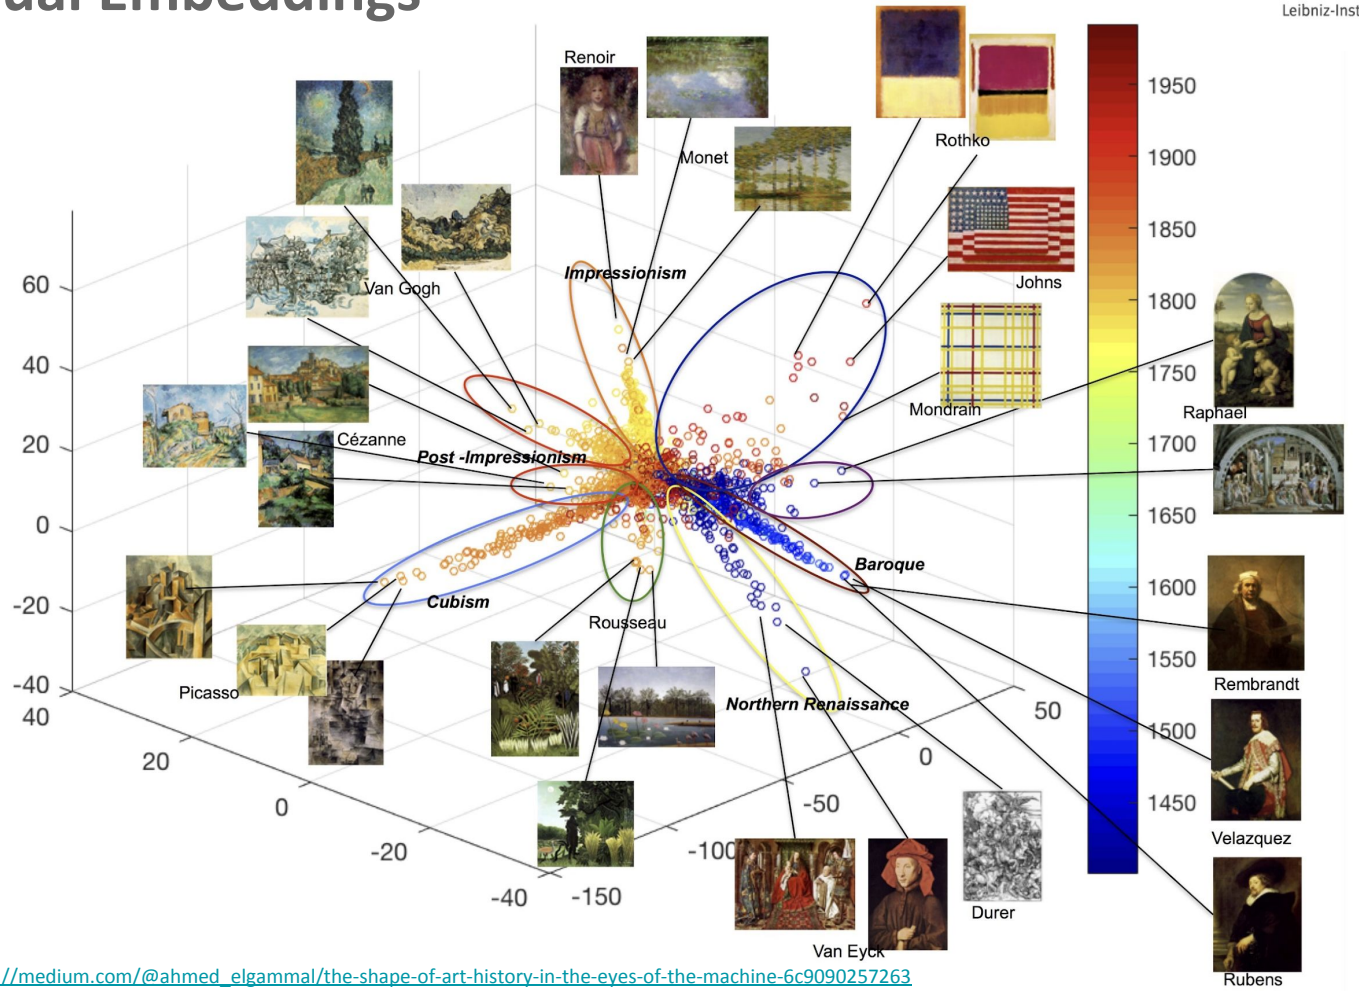

https://medium.com/@ahmed_elgammal/the-shape-of-art-history-in-the-eyes-of-the-machine-6c9090257263

Similarity Based Search and Exploration

https://wikiview.net/

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**
- **Deep Learning benefits from exploiting explicit Semantics**
- **Semantics benefits from leveraging Deep Learning**

“Technology presumes there's just one right way to do things and there never is.”

Robert M. Pirsig, Zen and the Art of Motorcycle Maintenance (1974)

Prof. Dr. Harald Sack

Deep Learning and Semantics

The two Faces of AI

harald.sack@fiz-karlsruhe.de
twitter: lysander07

STN Basel Info Day
26. Oct. 2020

FIZ Karlsruhe
Leibniz-Institut für Informationsinfrastruktur

Mitglied der Leibniz Gemeinschaft

Image References:

[3,4] Stephanie Mlot, Google AI Detects BReast Cancer better than Humans, 10/15/2018, https://money.cnn.com/2017/03/03/technology/google-breast-cancer-ai/

[5] AlphaGo Zero: Google DeepMind supercomputer learns 3,000 years of human knowledge in 40 days, UK Daily Telegraph, 10/18/2017,, http://www.telegraph.co.uk/science/2017/10/18/alphago-zero-google-deepmind-supercomputer-learns-3000-years/

[6] Is artificial intelligence set to become art's next medium?, Christie's, 12/12/2018, https://www.christies.com/features/A-collaboration-between-two-artists-one-human-one-a-machine-9332-1.aspx

[7] Machine learning has been used to automatically translate long-lost languages, MIT Technology Review, 7/1/2019, https://www.technologyreview.com/s/613899/machine-learning-has-been-used-to-automatically-translate-long-lost-languages/

[8] OpenAI's GPT-3 may be the biggest thing since bitcoin, maroz.com, 7/18/2020, https://maraoz.com/2020/07/18/openai-gpt3/

[9,19] James Vaughan, 1968 - "2001 - Hal's Eye", Flickr, https://www.flickr.com/photos/40143737@N02/4128131032

[10] BruceBlaus, Multipolar Neuron, CC-BY-3.0, https://commons.wikimedia.org/wiki/File:Blausen_0657_MultipolarNeuron.png

[12] (March 21, 1957) Man and woman shown working with IBM type 704 electronic data processing machine used for making computations for aeronautical research., Public Domain, https://commons.wikimedia.org/wiki/File:IBM_Electronic_Data_Processing_Machine_(9467782802).jpg

[15] Photo of the AMD Fiji GPU package which includes the Fiji GPU, HBM memory, interposer, substrate and the rest of the package. C. Spille/pcgameshardware.de, CC-BY-SA, https://commons.wikimedia.org/wiki/File:AMD_Fiji_GPU_package_with_GPU,_HBM_memory_and_interposer.jpg

[21] Leon A. Gatys, Alexander S. Ecker, Matthias Bethge: A Neural Algorithm of Artistic Style. CoRR abs/1508.06576 (2015)

[22,24] Zhu, Jun-Yan; Park, Taesung; Isola, Phillip; Efros, Alexei A., Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks, arXiv.org 2017, https://arxiv.org/abs/1703.10593

[25] DeOldyfy, FitHub 2018, https://github.com/jantic/DeOldify

[26] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, Wenzhe Shi, Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network, arxiv.org, 2016, https://arxiv.org/abs/1609.04802

[27] From Texts to Kitties: OpenAI's GPT Language Model Tackles Image Generation, Synced, 6/18/2020, https://syncedreview.com/2020/06/18/from-texts-to-kitties-openais-gpt-language-model-tackles-image-generation/

[28,44,45] 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)

Image References (cont.):

[29] James Vincent, Can you tell the difference between Bach and RoboBach?, The Verve, 12/23/2016, https://www.theverge.com/2016/12/23/14069382/ai-music-creativity-bach-deepbach-csl

[30] RELENTLESS DOPPELGANGER \m/ \m/ \m/ \m/ \m/ \m/ \m/ \m/ \m/ \m/ \m/ \m/ \m/ \m/, https://www.youtube.com/watch?v=MwtVkPKx3RA

[34] Parental Advisory, Wikimedia Commons, Open Source, https://commons.wikimedia.org/wiki/File:Parental_Advisory_-_explicit_content_-_label.jpg

[35] Aristotle by Francesco Hayez. Oil on canvas, 1811. Public Domain, https://commons.wikimedia.org/wiki/File:Francesco_Hayez_001.jpg

[35] Arbor porphyrii (probably from one of Boethius' translations), Public Domain, https://commons.wikimedia.org/wiki/File:Arbor_porphyrii_(probably_from_one_of_Boethius%27_translations).png

[39,40] Niklas Jansson, Touched by His Noodly Appendage, https://commons.wikimedia.org/wiki/File:Touched_by_His_Noodly_Appendage_HD.jpg

[41,53] The Linked Data Cloud, 2019, https://lod-cloud.net]

[42, 56] SciHi Blog, http://scihi.org

[51] Wilhelm von Osten and Clever Hans, 1908, Public Domain, https://commons.wikimedia.org/wiki/File:Osten_und_Hans.jpg

[54,55, 56, 57, 59] Hugh of Cluny, Holy Roman Emperor Henry IV, and Matilda of Tuscany, Cod. Vat. lat. 4922, fol. 49r (completed in 1115 AD), public domain, https://commons.wikimedia.org/wiki/File:Hugo-v-cluny_heinrich-iv_mathilde-v-tuszien_cod-vat-lat-4922_1115ad.jpg

[58] Ahmad ElGammal, The Shape of Art History in the Eyes of the Machine, 2/12/2018, https://medium.com/@ahmed_elgammal/the-shape-of-art-history-in-the-eyes-of-the-machine-6c9090257263

[59] WikiView, https://wikiview.net/