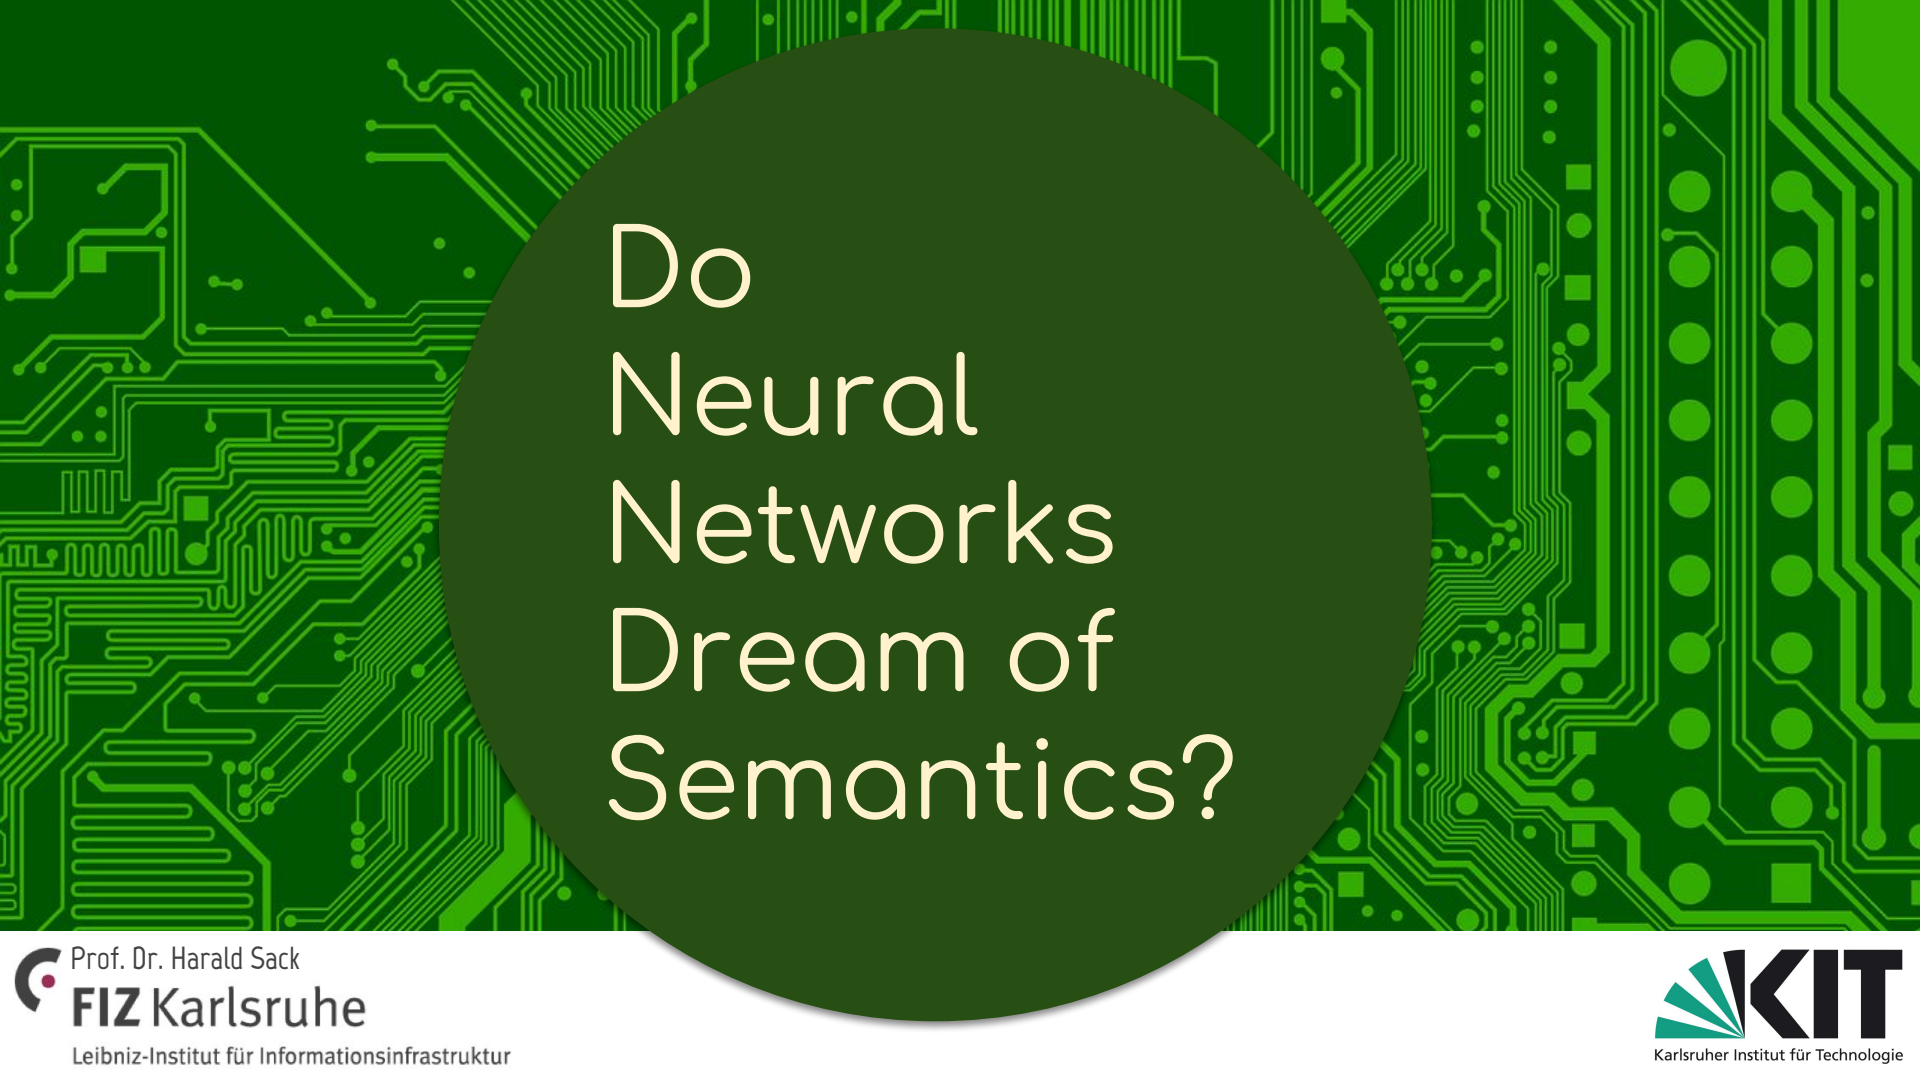

Prof. Dr. Harald Sack
FIZ Karlsruhe
Leibniz-Institut für Informationsinfrastruktur

Karlsruher Institut für Technologie

Science

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

AlphaGo

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

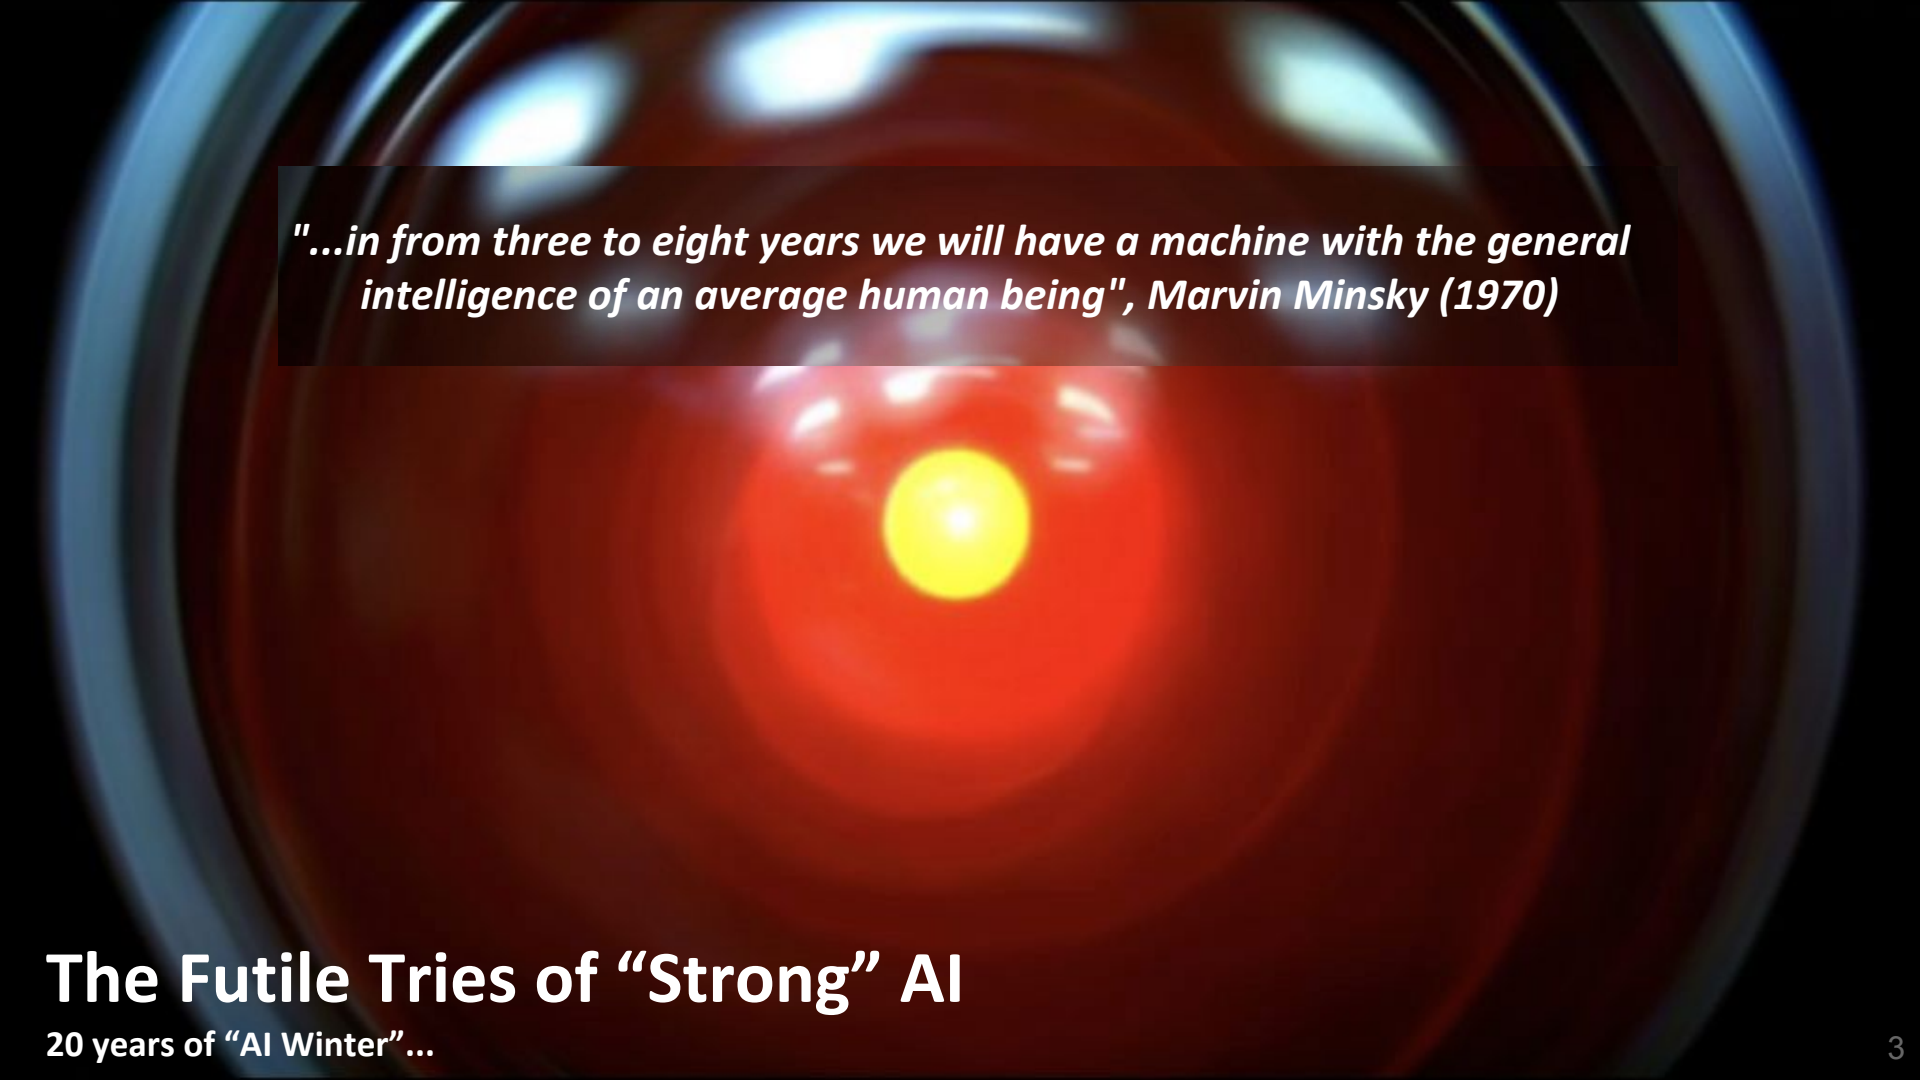

"...in from three to eight years we will have a machine with the general intelligence of an average human being", Marvin Minsky (1970)

The Futile Tries of “Strong” AI

20 years of “AI Winter”...


Where
did
the Magic
come
from...?

Transfering Biology into a Mathematical Model

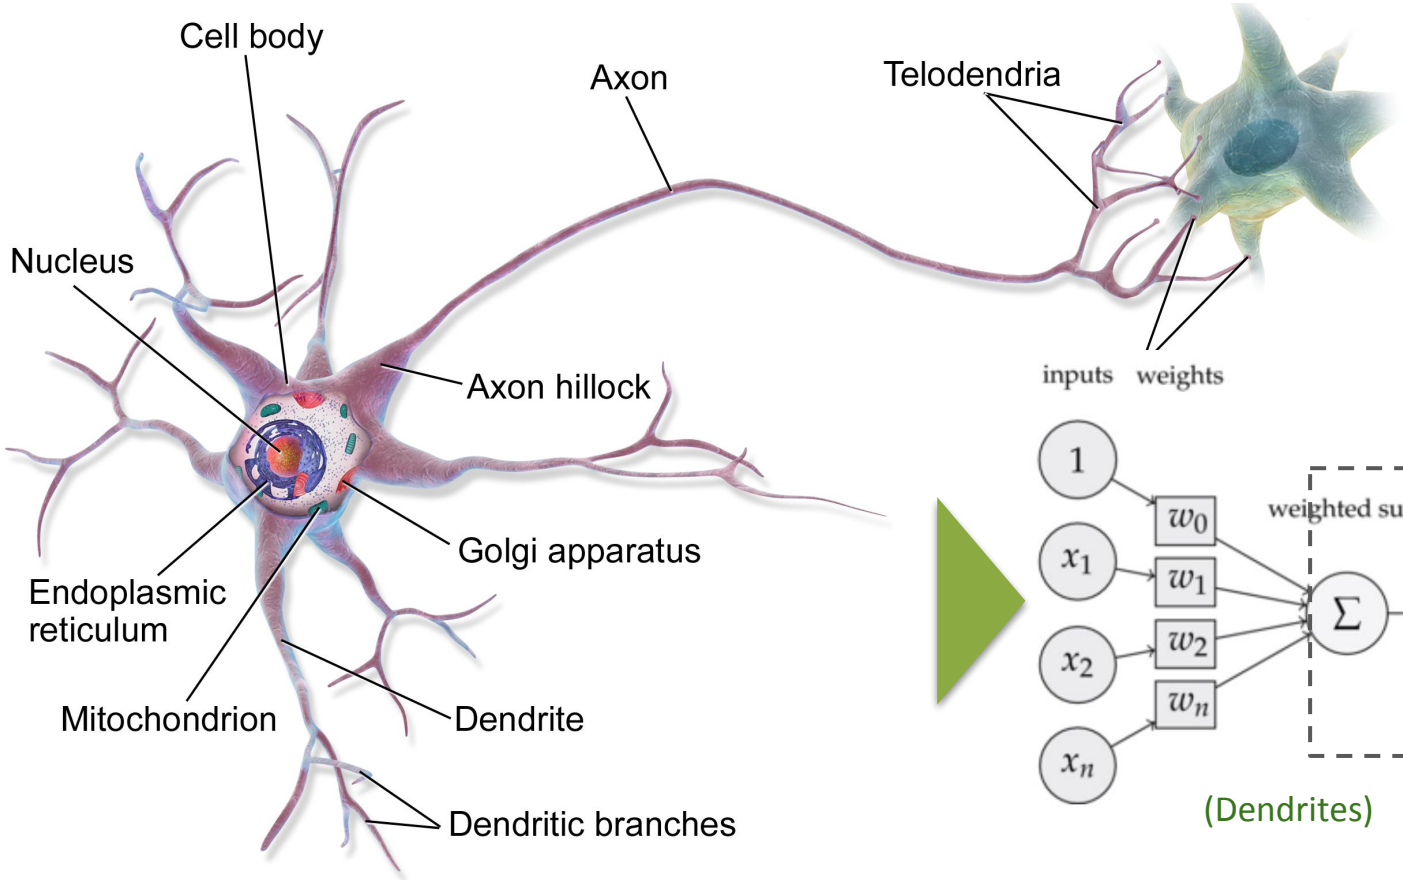

Donald Hebb (1940)
McCulloch & Pitts (1943)

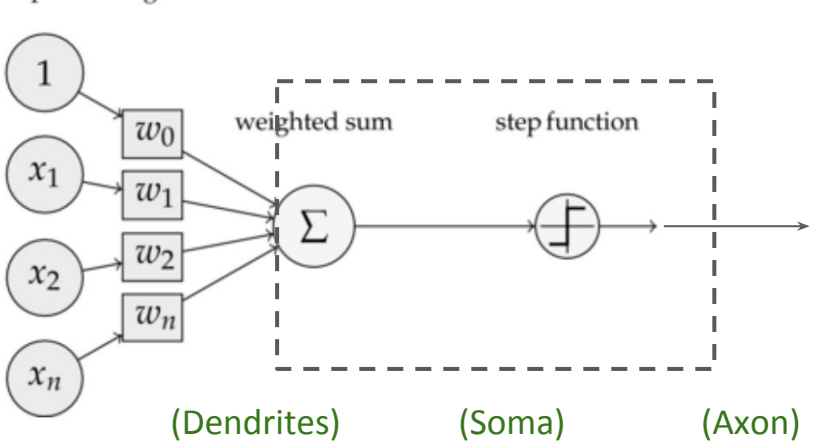

Cognitive Computing - The MARK 1 Perceptron (Rosenblatt, 1957)

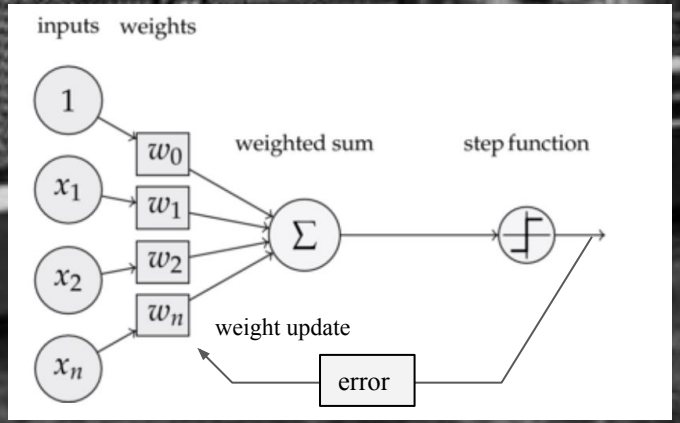

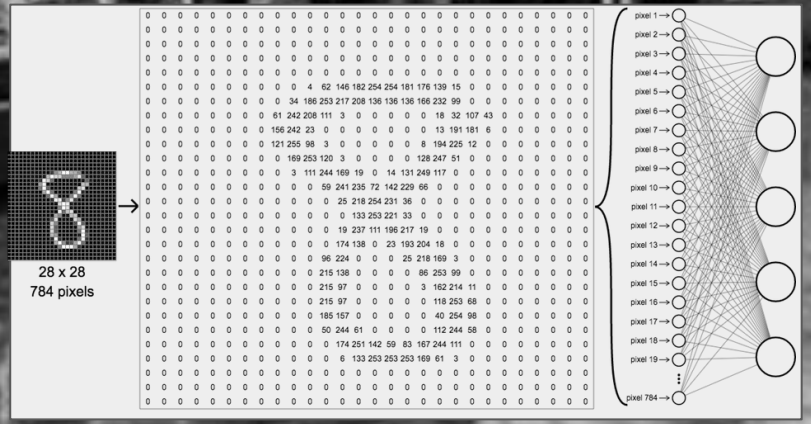

The Triumphant Progress of Neural Network

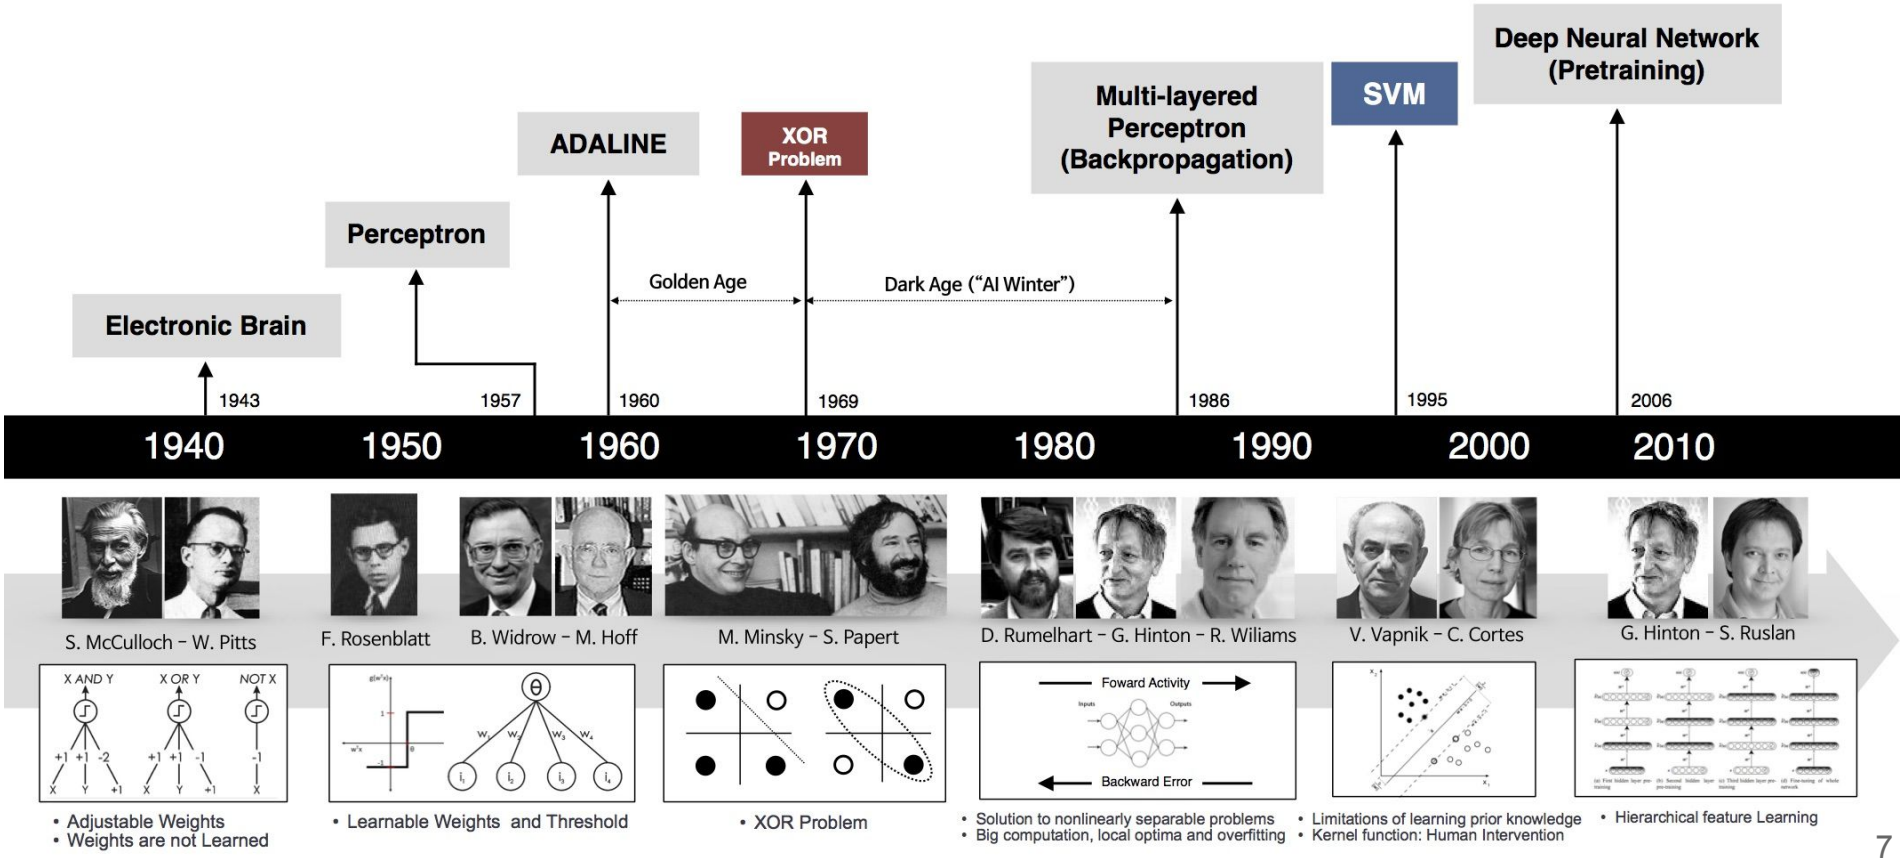

- Adjustable Weights
- Weights are not Learned

- Learnable Weights and Threshold

- XOR Problem

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

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

- Hierarchical feature Learning

Deep Convolutional Neural Networks on GPU Supercomputers

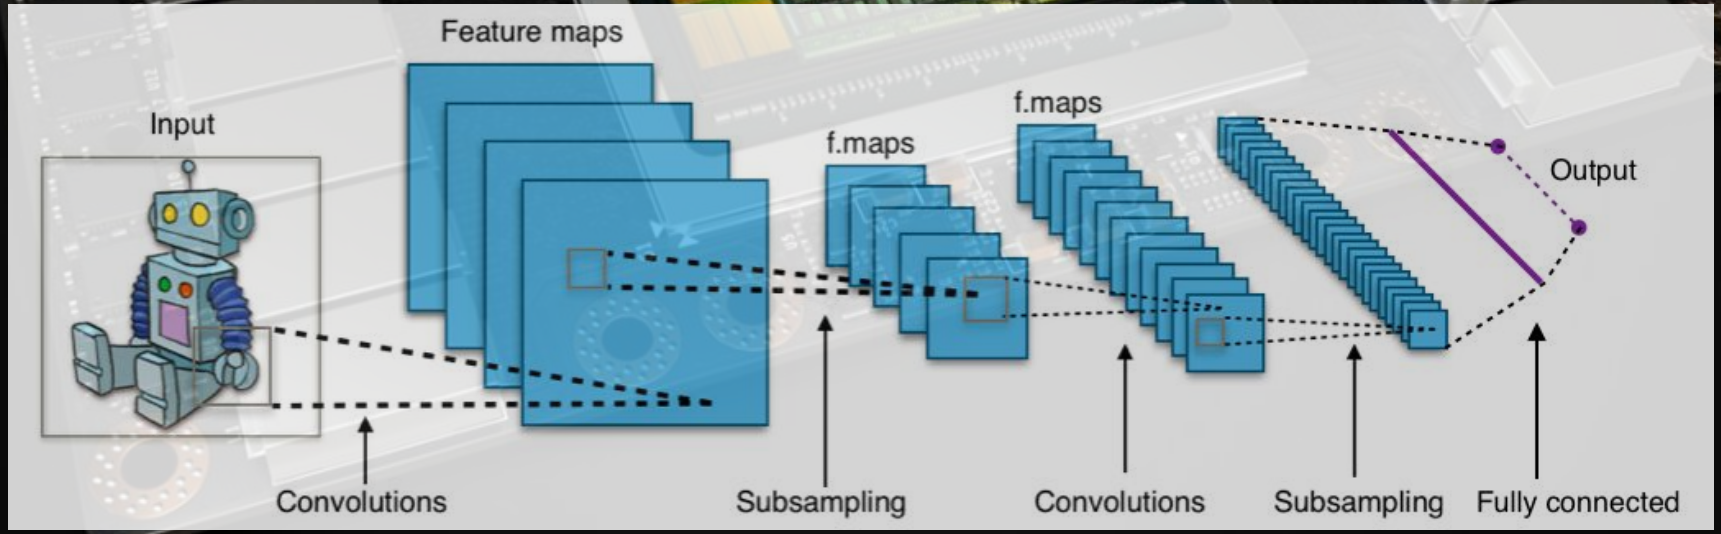

Any sufficiently advanced technology is indistinguishable from magic.

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

Neural Network in Visual Analysis

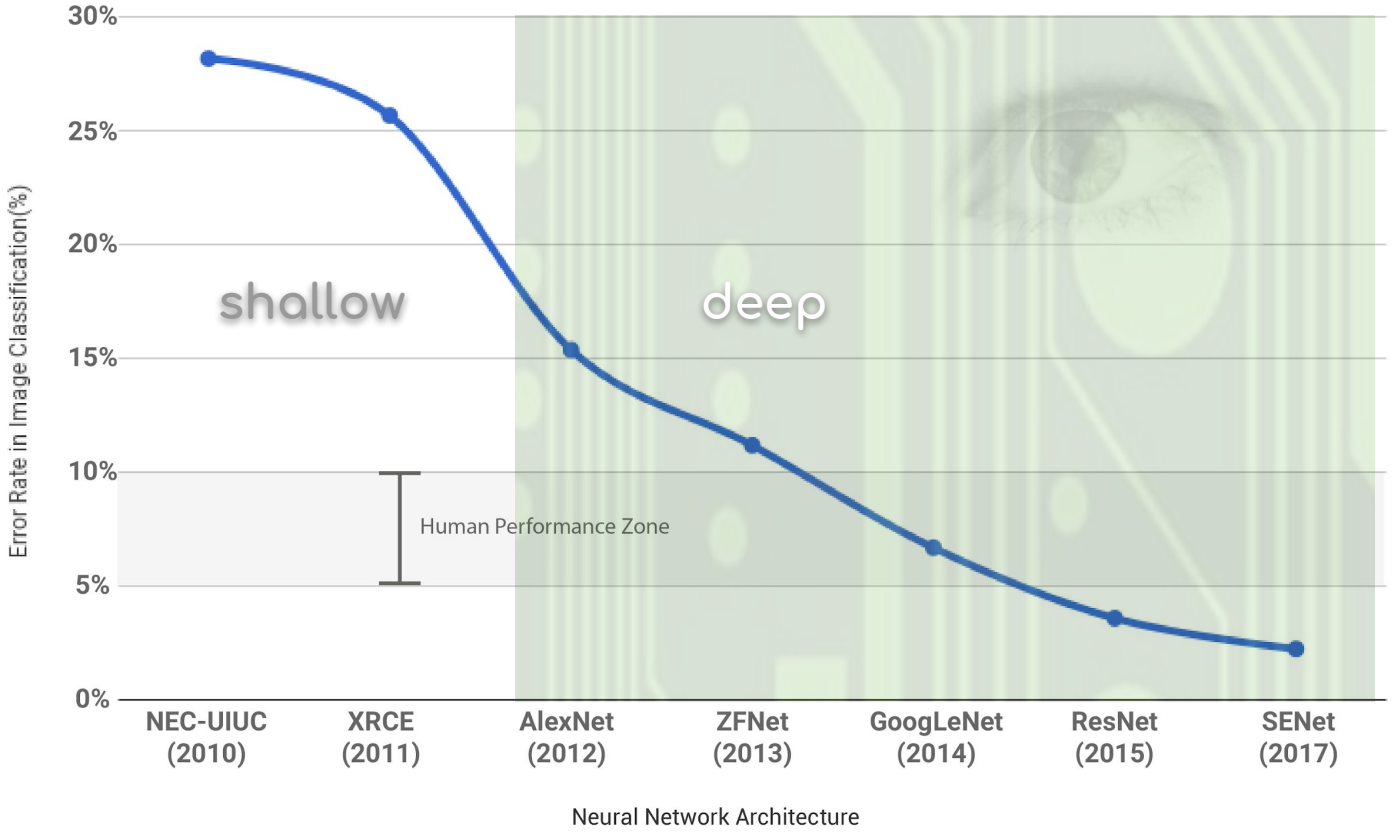

Karlsruher Institut für Technologie

FIZ Karlsruhe
Leibniz-Institut für Informationsinfrastruktur

Combining Visual Analysis and Linguistics

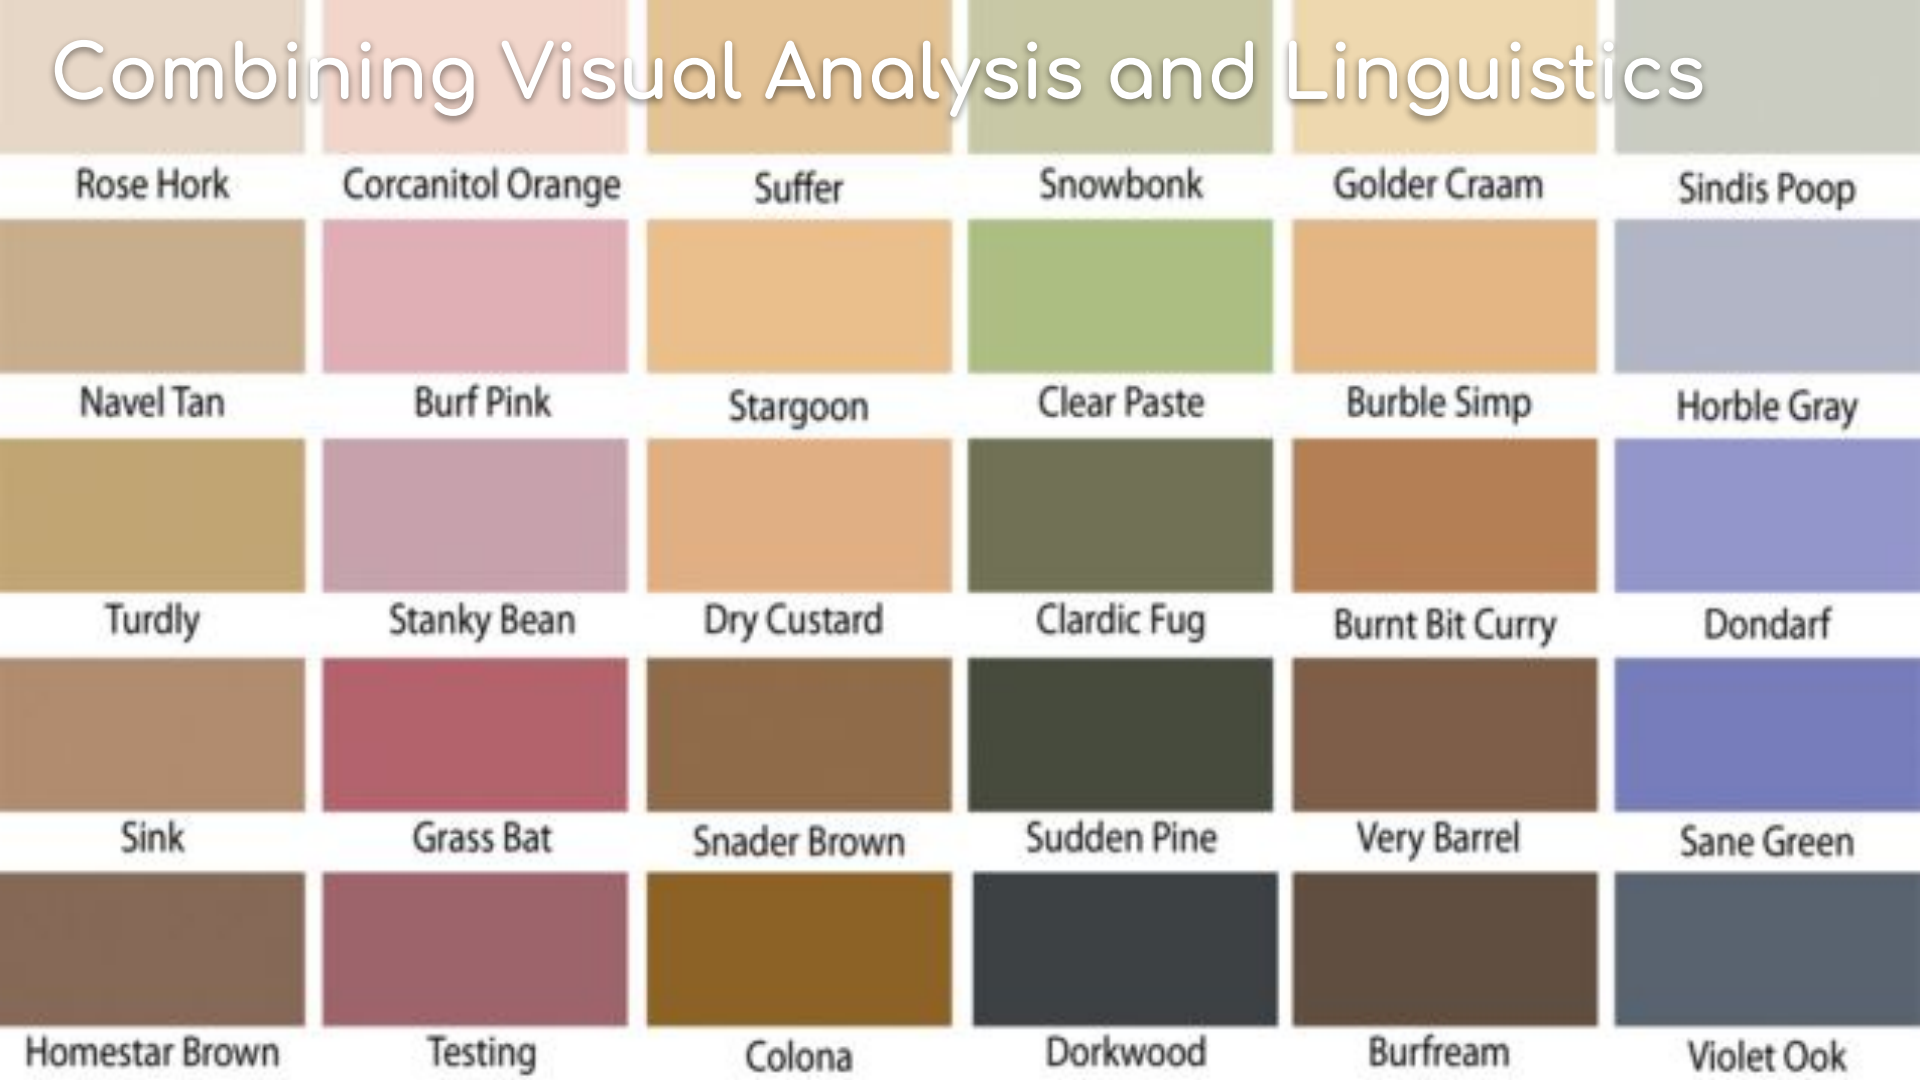

What Deep Learning has achieved so far

- 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

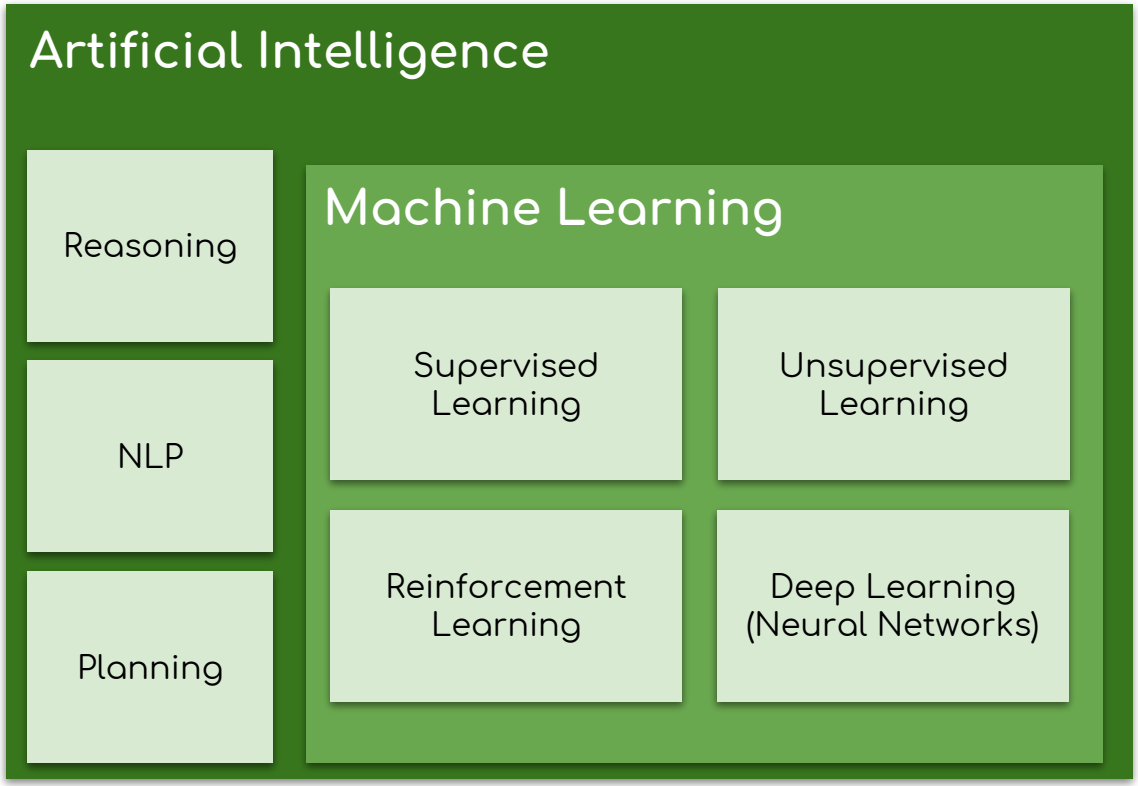

"The Goal of AI is to develop machines that behave as though they were intelligent."

- John McCarthy (1955)


PARENTAL
ADVISORY
EXPLICIT SEMANTICS

The Universal Categories - Aristotle (384–322 BC)

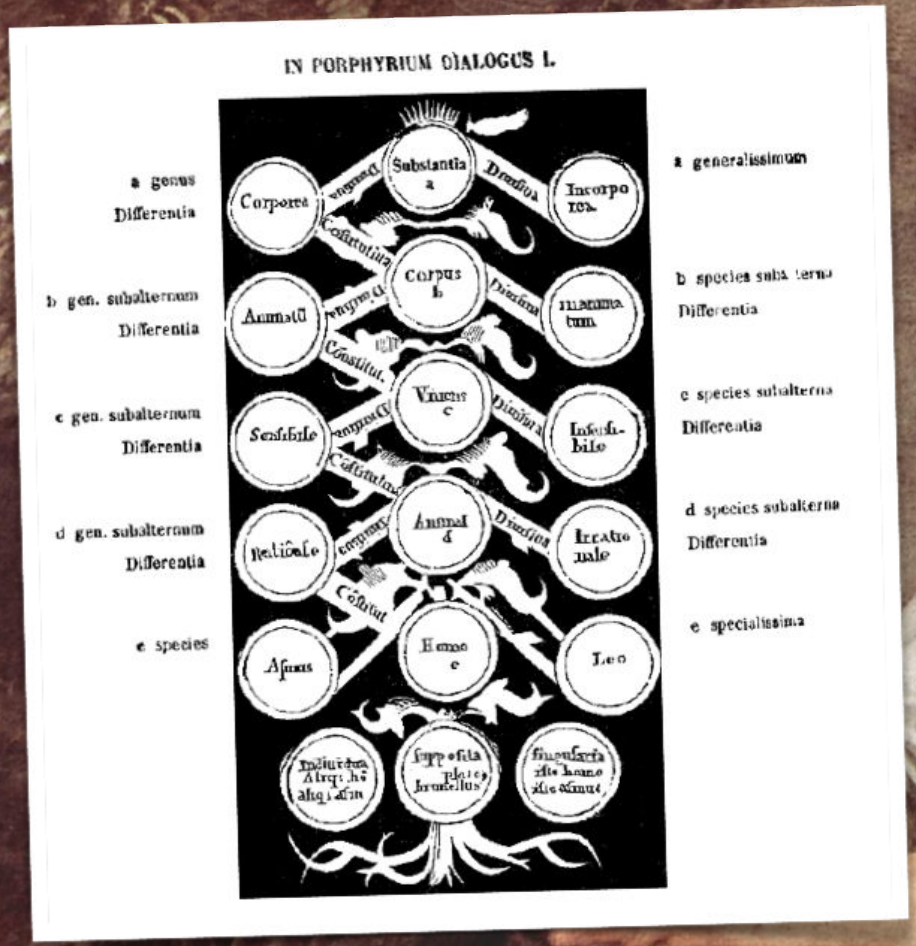

Calculus Ratiocinator - Gottfried Wilhelm Leibniz (1646-1716)

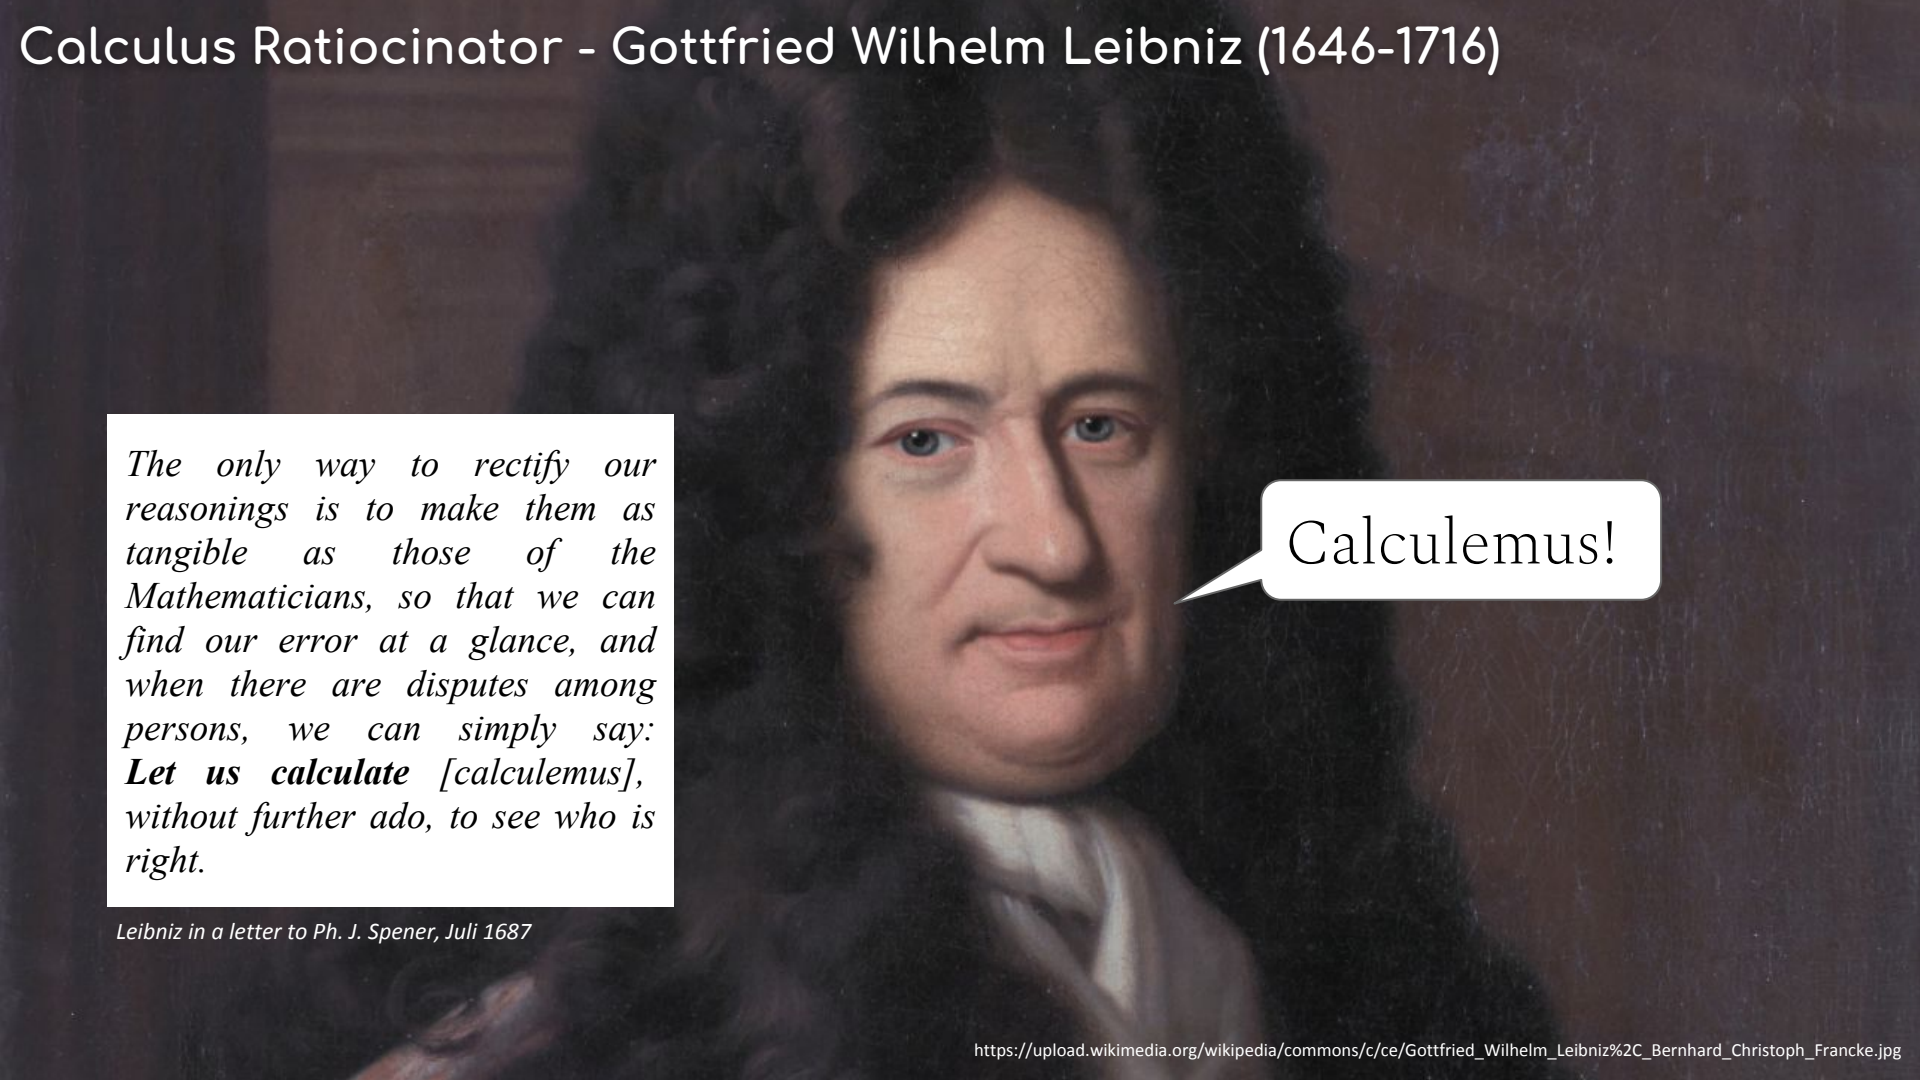

The only way to rectify our reasonings is to make them as tangible as those of the Mathematicians, so that we can find our error at a glance, and when there are disputes among persons, we can simply say: ***Let us calculate*** *[calculemus], without further ado, to see who is right.*

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)

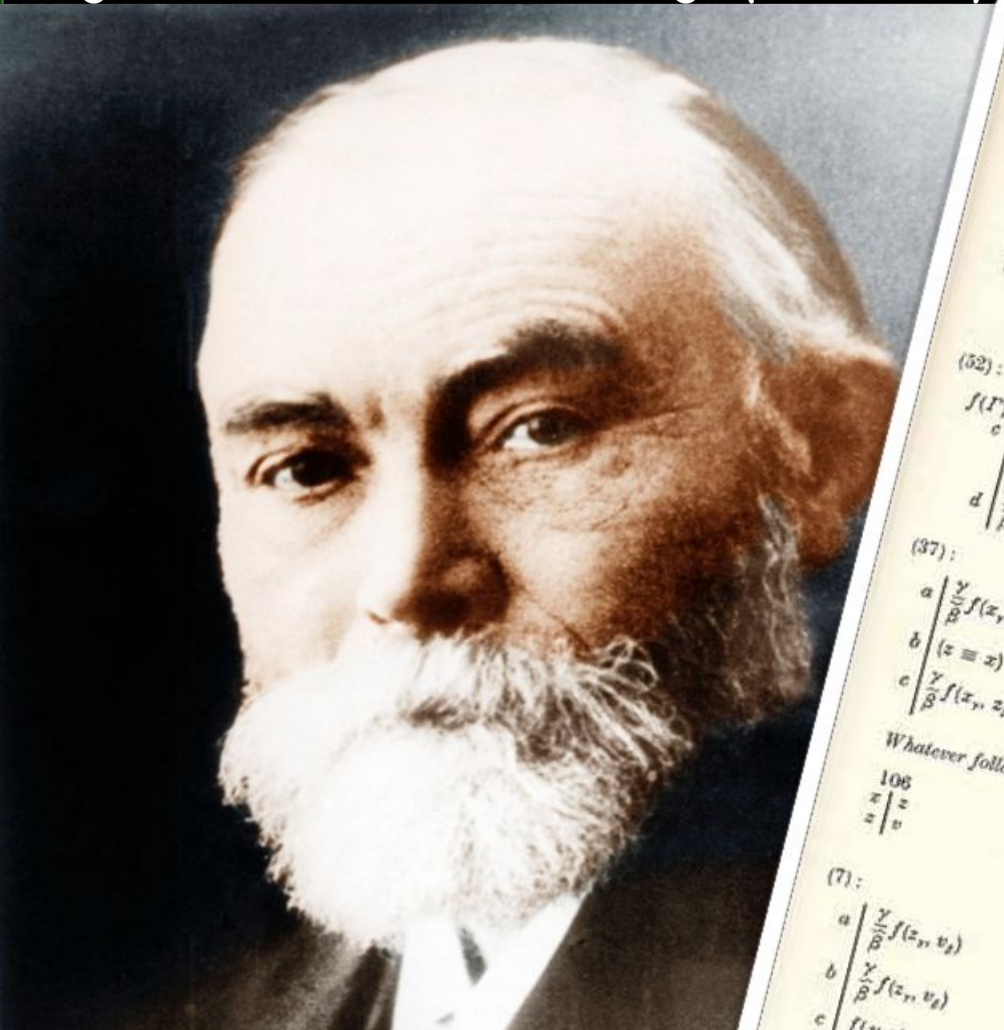

BEGRIFFSSCHRIFT

71

§ 30.

99

(55) : :

$$\begin{array}{l|l} d & z \\ c & z \end{array}$$

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

(52) :

$$\begin{array}{l|l} f(\Gamma) & \Gamma \\ c & \quad (z \equiv x) \\ d & \frac{\gamma}{\beta} f(x_\gamma, z_\beta) \end{array}$$

$$\begin{array}{l} \\ \hline \\ \left[\left[\quad (z \equiv x) \right. \right. \\ \quad \left. \left. \frac{\gamma}{\beta} f(x_\gamma, z_\beta) \right] \equiv \frac{\gamma}{\beta} f(x_\gamma, z_\beta) \right] \end{array} \qquad (104).$$

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

(37) :

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

$$\begin{array}{l} \\ \hline \\ \quad \frac{\gamma}{\beta} f(x_\gamma, z_\beta) \\ \frac{\gamma}{\beta} f(x_\gamma, z_\beta) \end{array} \qquad (105).$$

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

106

$$\begin{array}{l|l} x & z \\ z & v \end{array}$$

$$\begin{array}{l} \\ \hline \\ \quad \frac{\gamma}{\beta} f(z_\gamma, v_\beta) \\ \frac{\gamma}{\beta} f(z_\gamma, v_\beta) \end{array} \qquad (106).$$

(7) :

$$\begin{array}{l|l} a & \frac{\gamma}{\beta} f(z_\gamma, v_\beta) \\ b & \frac{\gamma}{\beta} f(z_\gamma, v_\beta) \\ c & f(z_\gamma, v_\beta) \end{array}$$

$$\begin{array}{l} \\ \hline \\ \quad \frac{\gamma}{\beta} f(z_\gamma, v_\beta) \\ \quad \frac{\gamma}{\beta} f(z_\gamma, v_\beta) \\ \quad f(z_\gamma, v_\beta) \end{array}$$

Frames for Represent Knowledge - Marvin Minsky (1974)

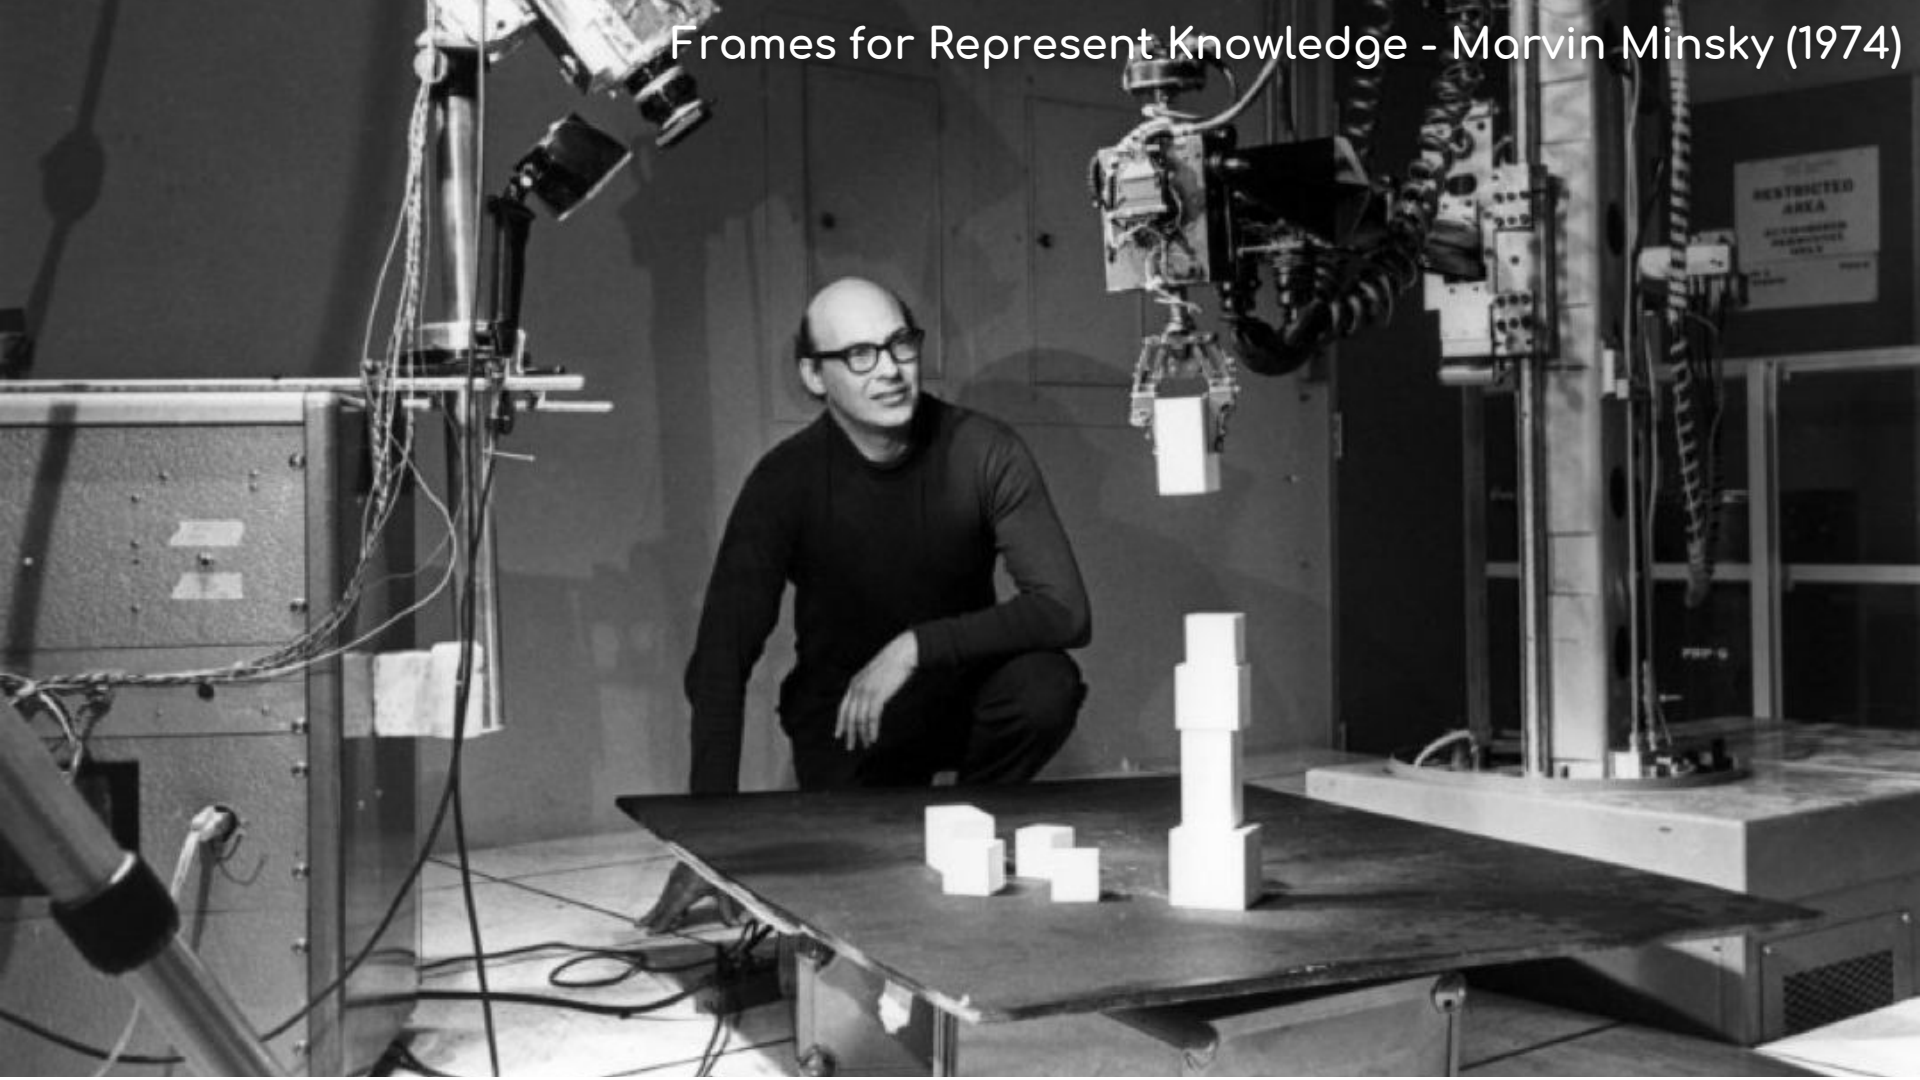

From Linked Data to Knowledge Graphs

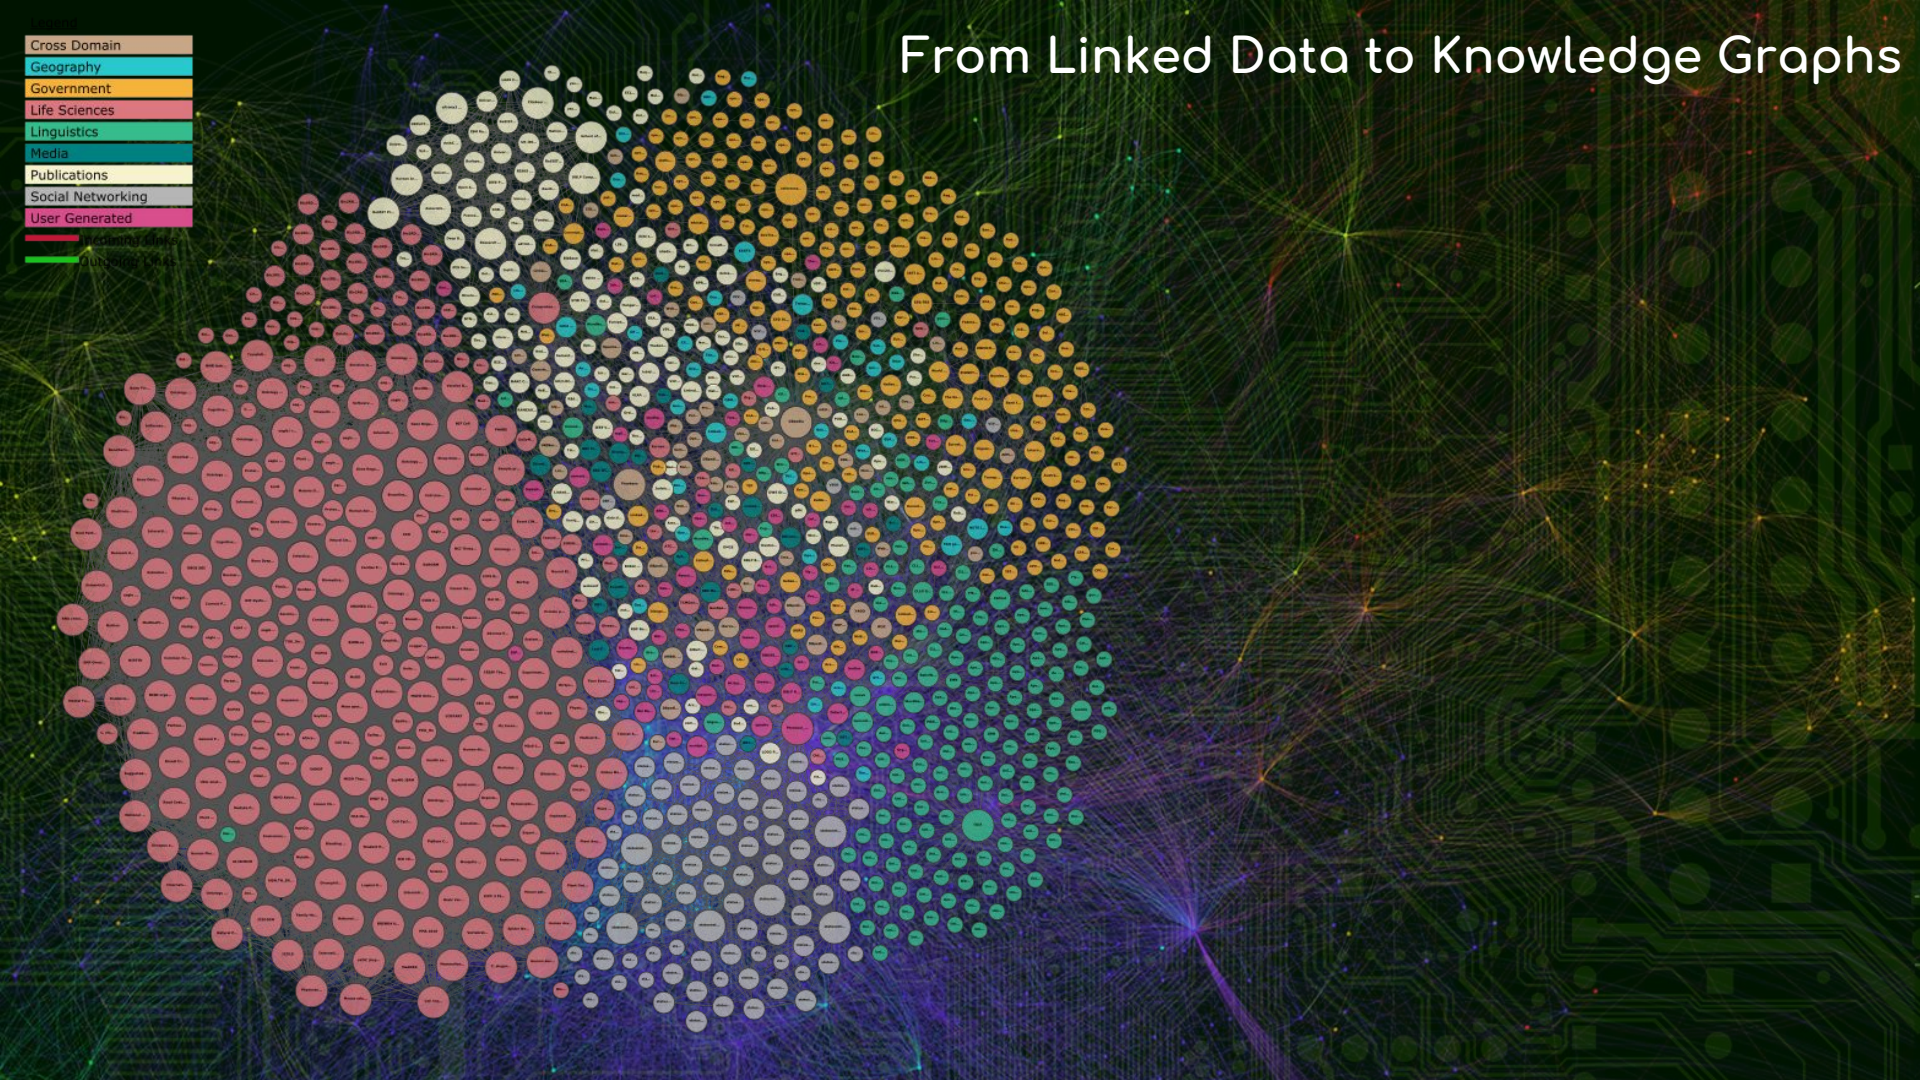

Knowledge Graph Applications

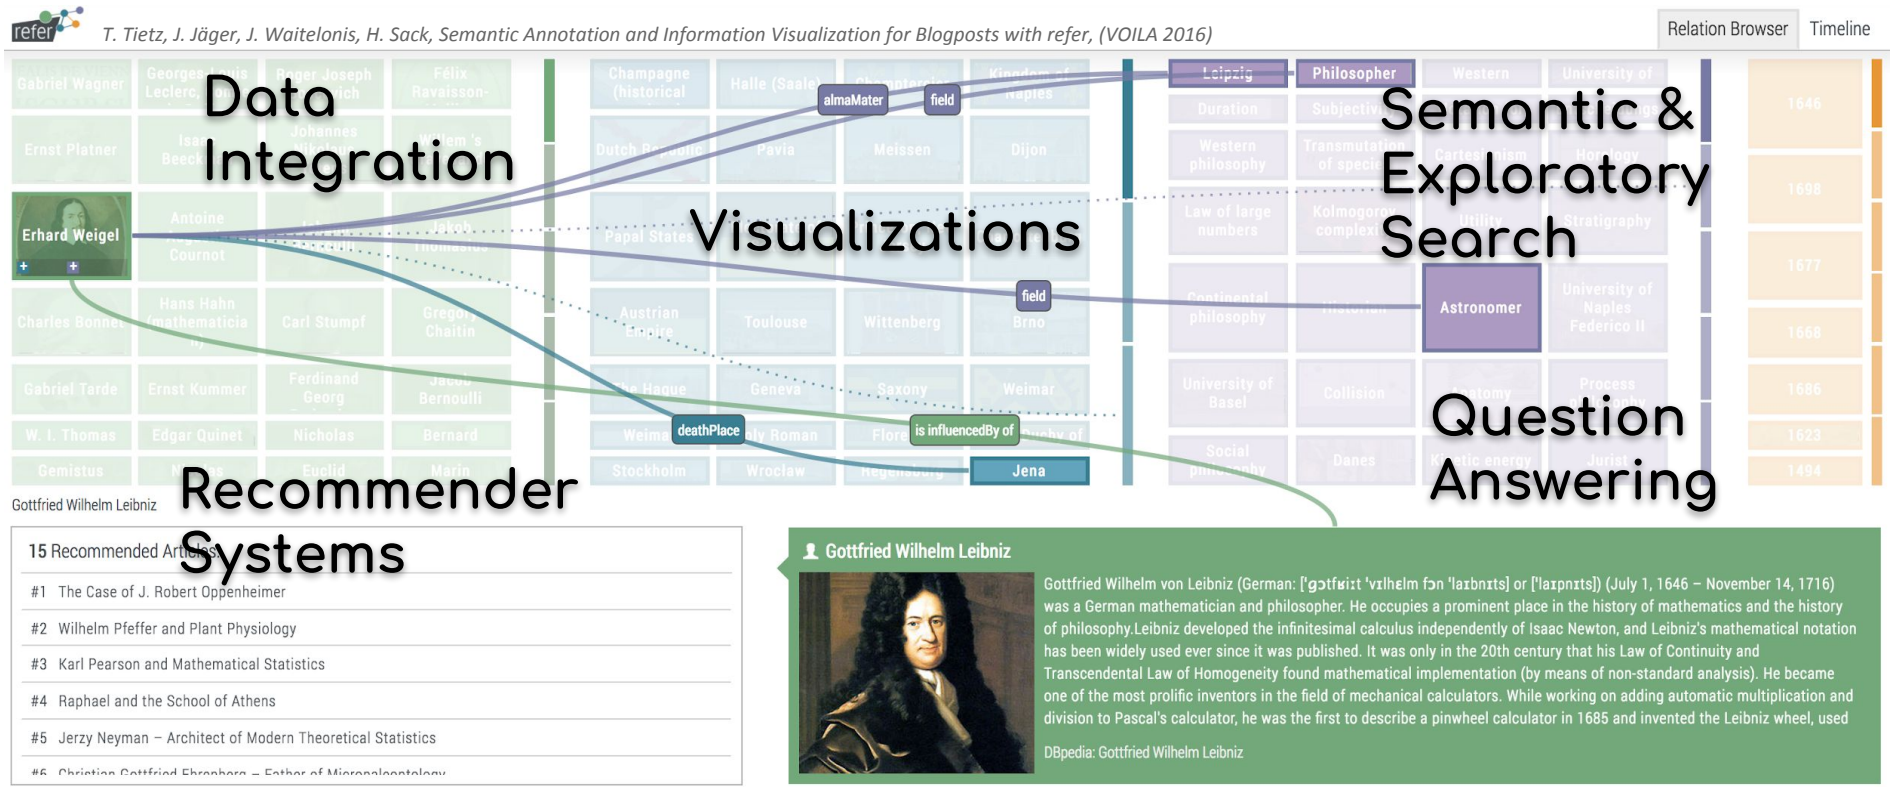


“There ain’t no such thing as a free Lunch.”
Robert A. Heinlein,
The Moon is a Harsh Mistress (1966)

What do Neural Networks (really) see?

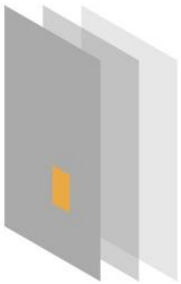

Neuron

`layer`$_n$`[x,y,z]`

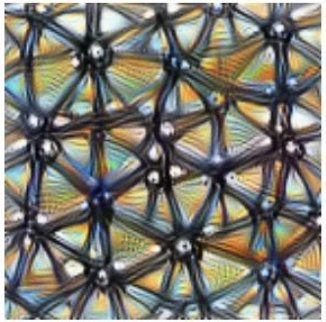

Channel

`layer`$_n$`[:,:,z]`

Layer/DeepDream

`layer`$_n$`[:,:,:]`^{2}

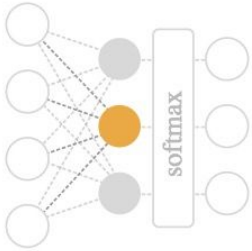

Class Logits

`pre_softmax[k]`

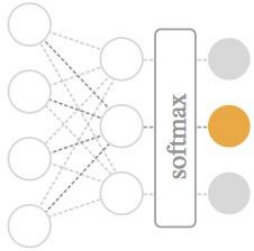

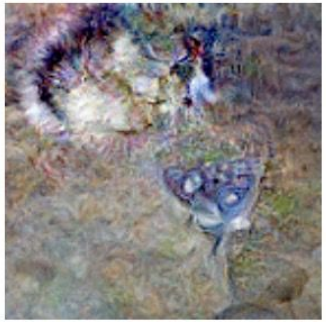

Class Probability

`softmax[k]`

https://distill.pub/2017/feature-visualization/

What do Neural Networks (really) see?

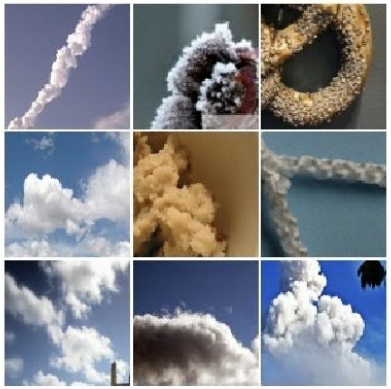

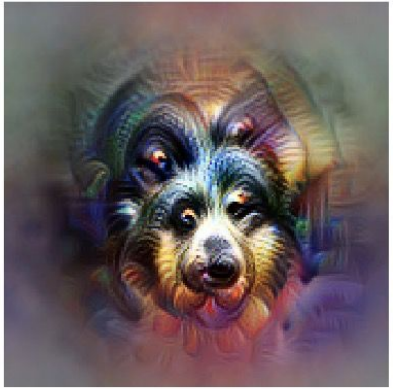

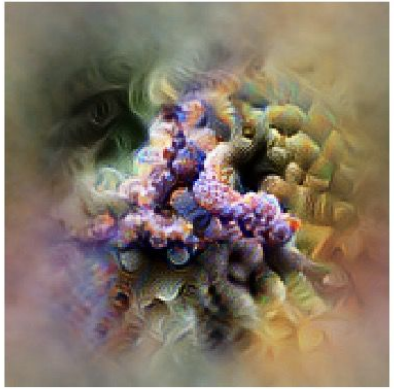

Baseball—or stripes?

Animal faces—or snouts?

Clouds—or fluffiness?

Buildings—or sky?

https://distill.pub/2017/feature-visualization/

Knowledge Graphs: Local Details vs. Global Coverage

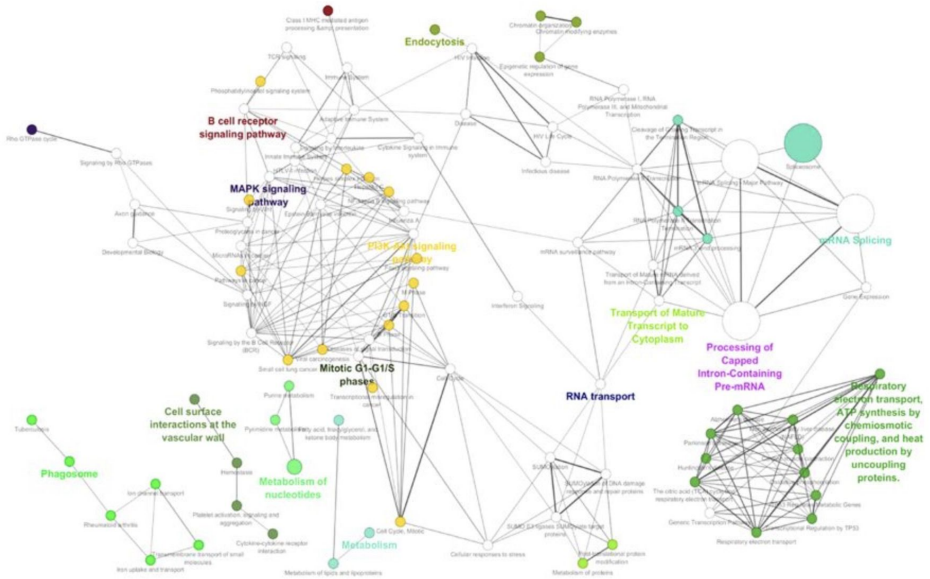

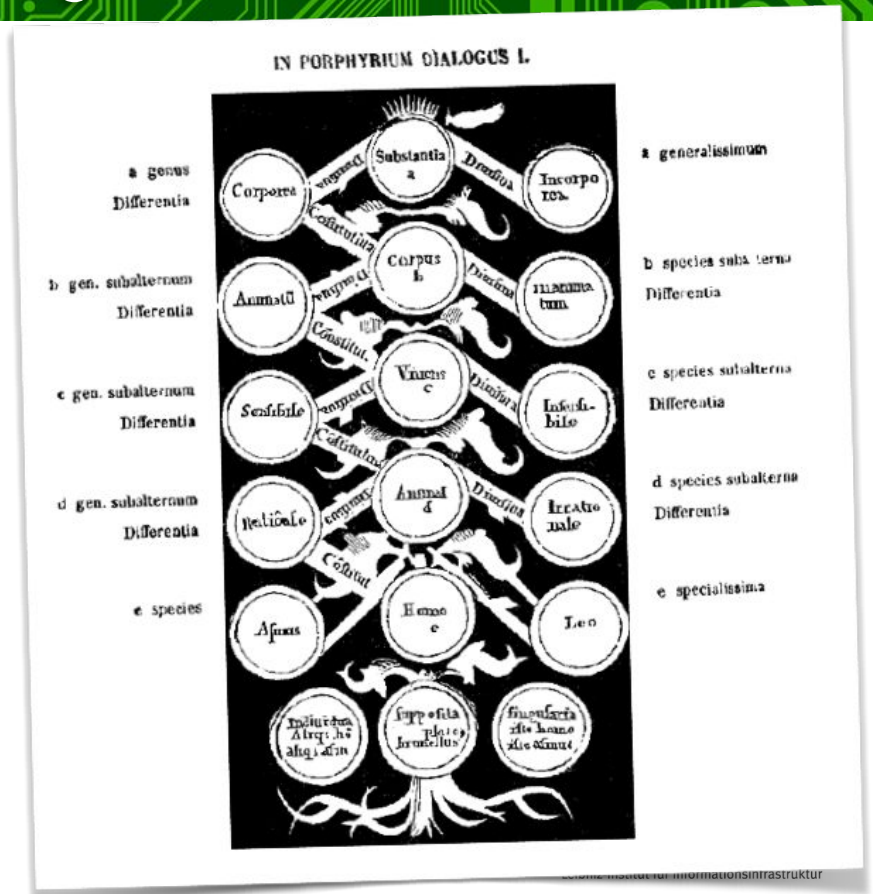

90%
of most
Magic merely
consists of
knowing one
extra fact.

Terry Pratchett, Night Watch (2002)

Generative Adversarial Networks Example

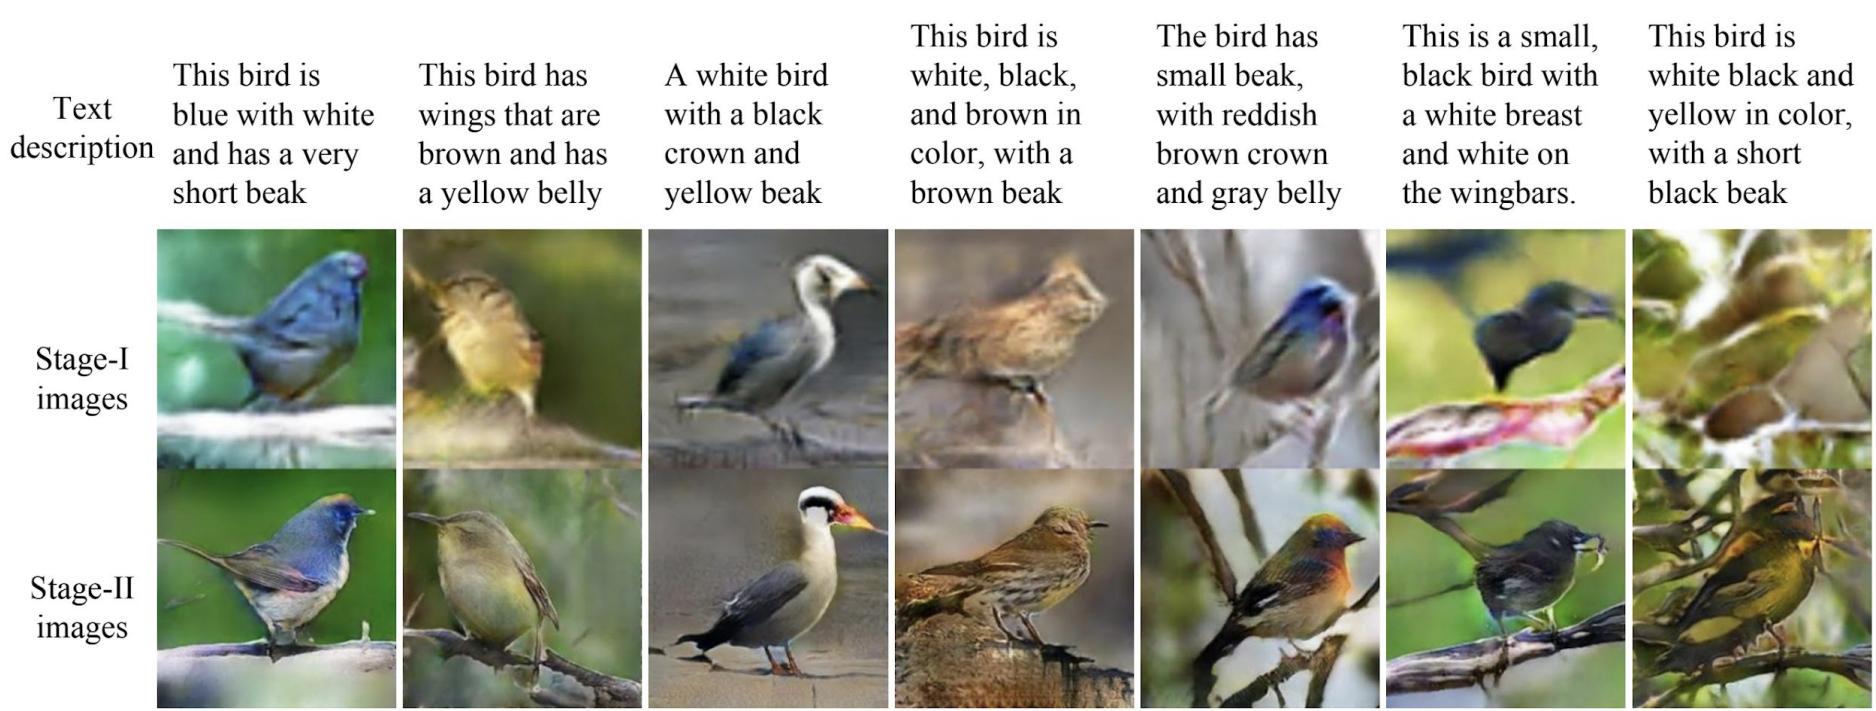

Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, Dimitris N. Metaxas:
StackGAN++: Realistic Image Synthesis with Stacked Generative Adversarial Networks.
CoRR abs/1710.10916 (2017)

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

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://i2cv.vanzuelab.com/

A flock of sheep on green meadows

http://t2i.cvalenzuelab.com/

a girl watching tv

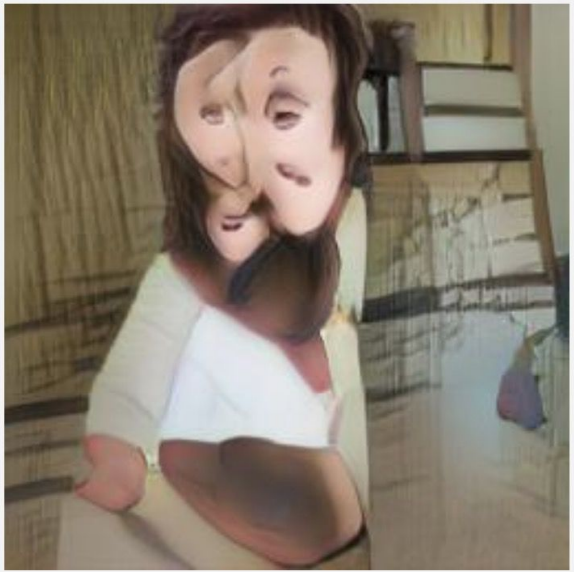

http://t2i.cvalenzuelab.com/

A girl wearing a red skirt

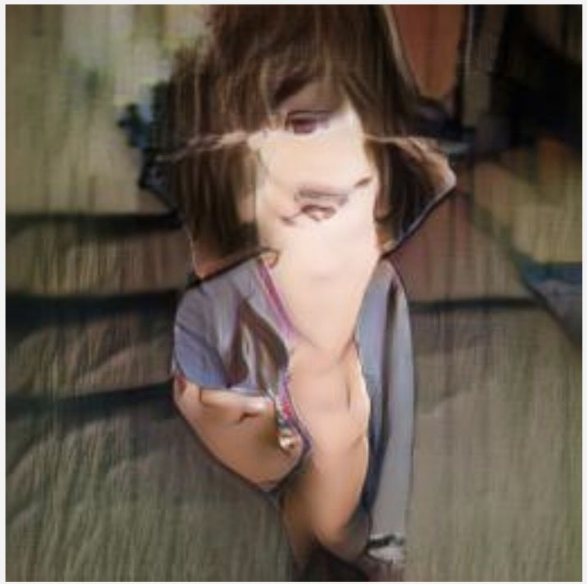

http://t2i.cvalenzuelab.com/

A girl with two eyes one mouth and one nose

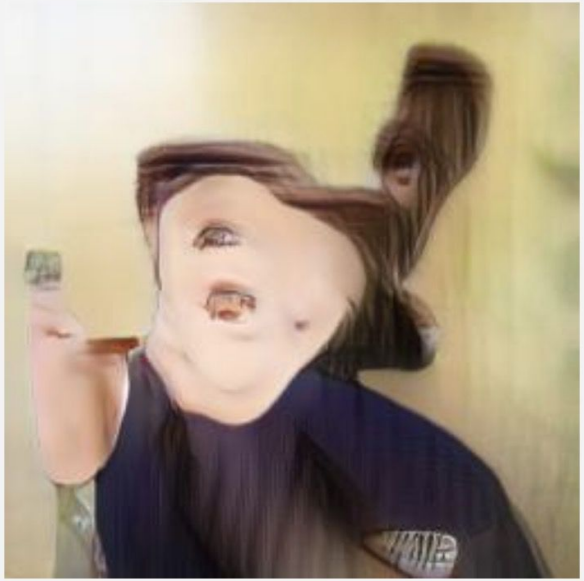

http://t2i.cvalenzuelab.com/

A man with white hair, a white beard and a white shirt, wearing glasses

http://t2i.cvalenzuelab.com/

melting clocks flying over a large dessert

http://t2i.cvalenzuelab.com/

(Long Term) Goals

- Combining **Semantics** (a.k.a. Symbolic Reasoning) and **Machine Learning** (**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." (Bladerunner, 1982)