Semantics, Deep Learning, and the Transformation of Business

Steve Omohundro, Ph.D.
PossibilityResearch.com
SteveOmohundro.com
SelfAwareSystems.com

http://discovermagazine.com/~/media/Images/Issues/2013/Jan-Feb/connectome.jpg

Economic Impact

Deep Learning

Neats vs. Scruffies

Semantics

The Future

https://www.flickr.com/photos/danielfoster/14758510078/

McKinsey: $50 Trillion to 2025

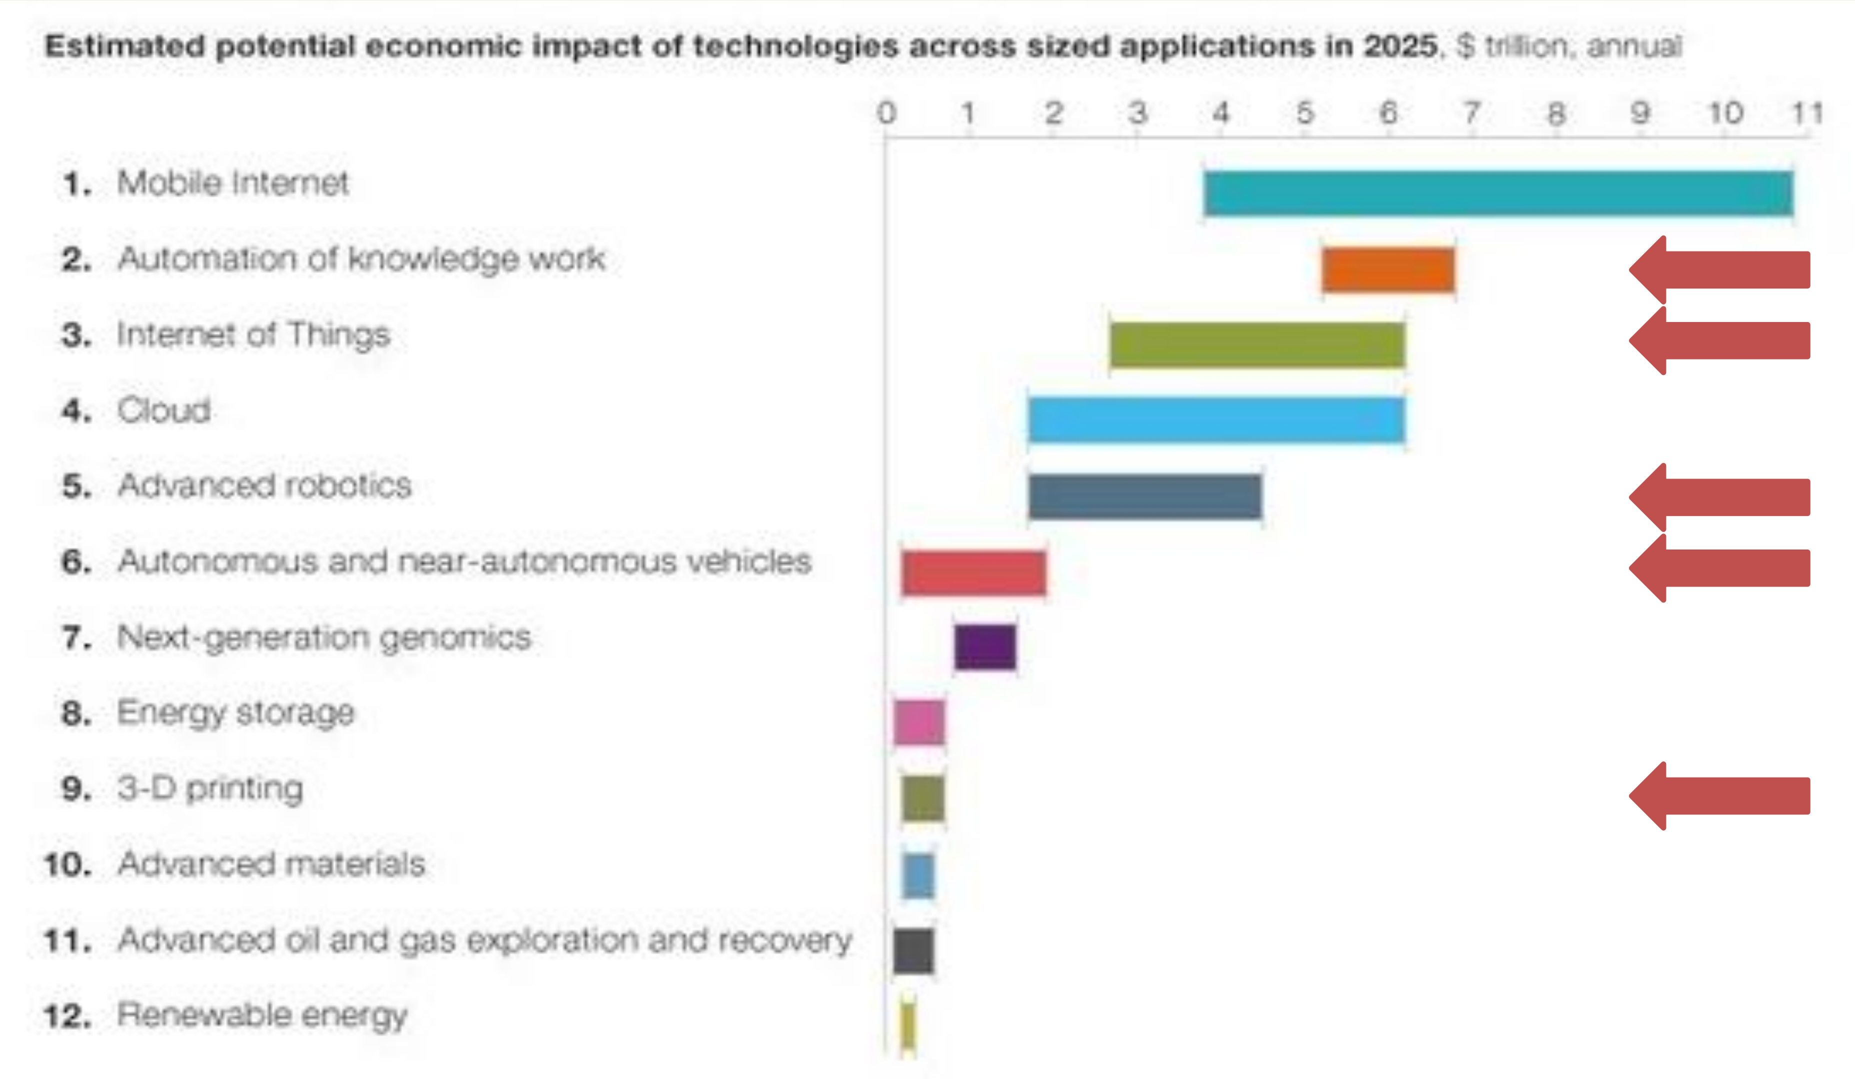

http://www.mckinsey.com/insights/business_technology/disruptive_technologies


AI Knowledge Work: $25 Trillion to 2025
Marketing, ERP, Big Data, Smart Assistants
http://thismasquerade.me/wp-content/uploads/blogger/-84ftTg_azwY/Um8up3j4zTI/AAAAAAAABJ8/9_-dxg23UaY/s1600/9979402_ml.jpg

Internet of Things: $15 Trillion to 2025

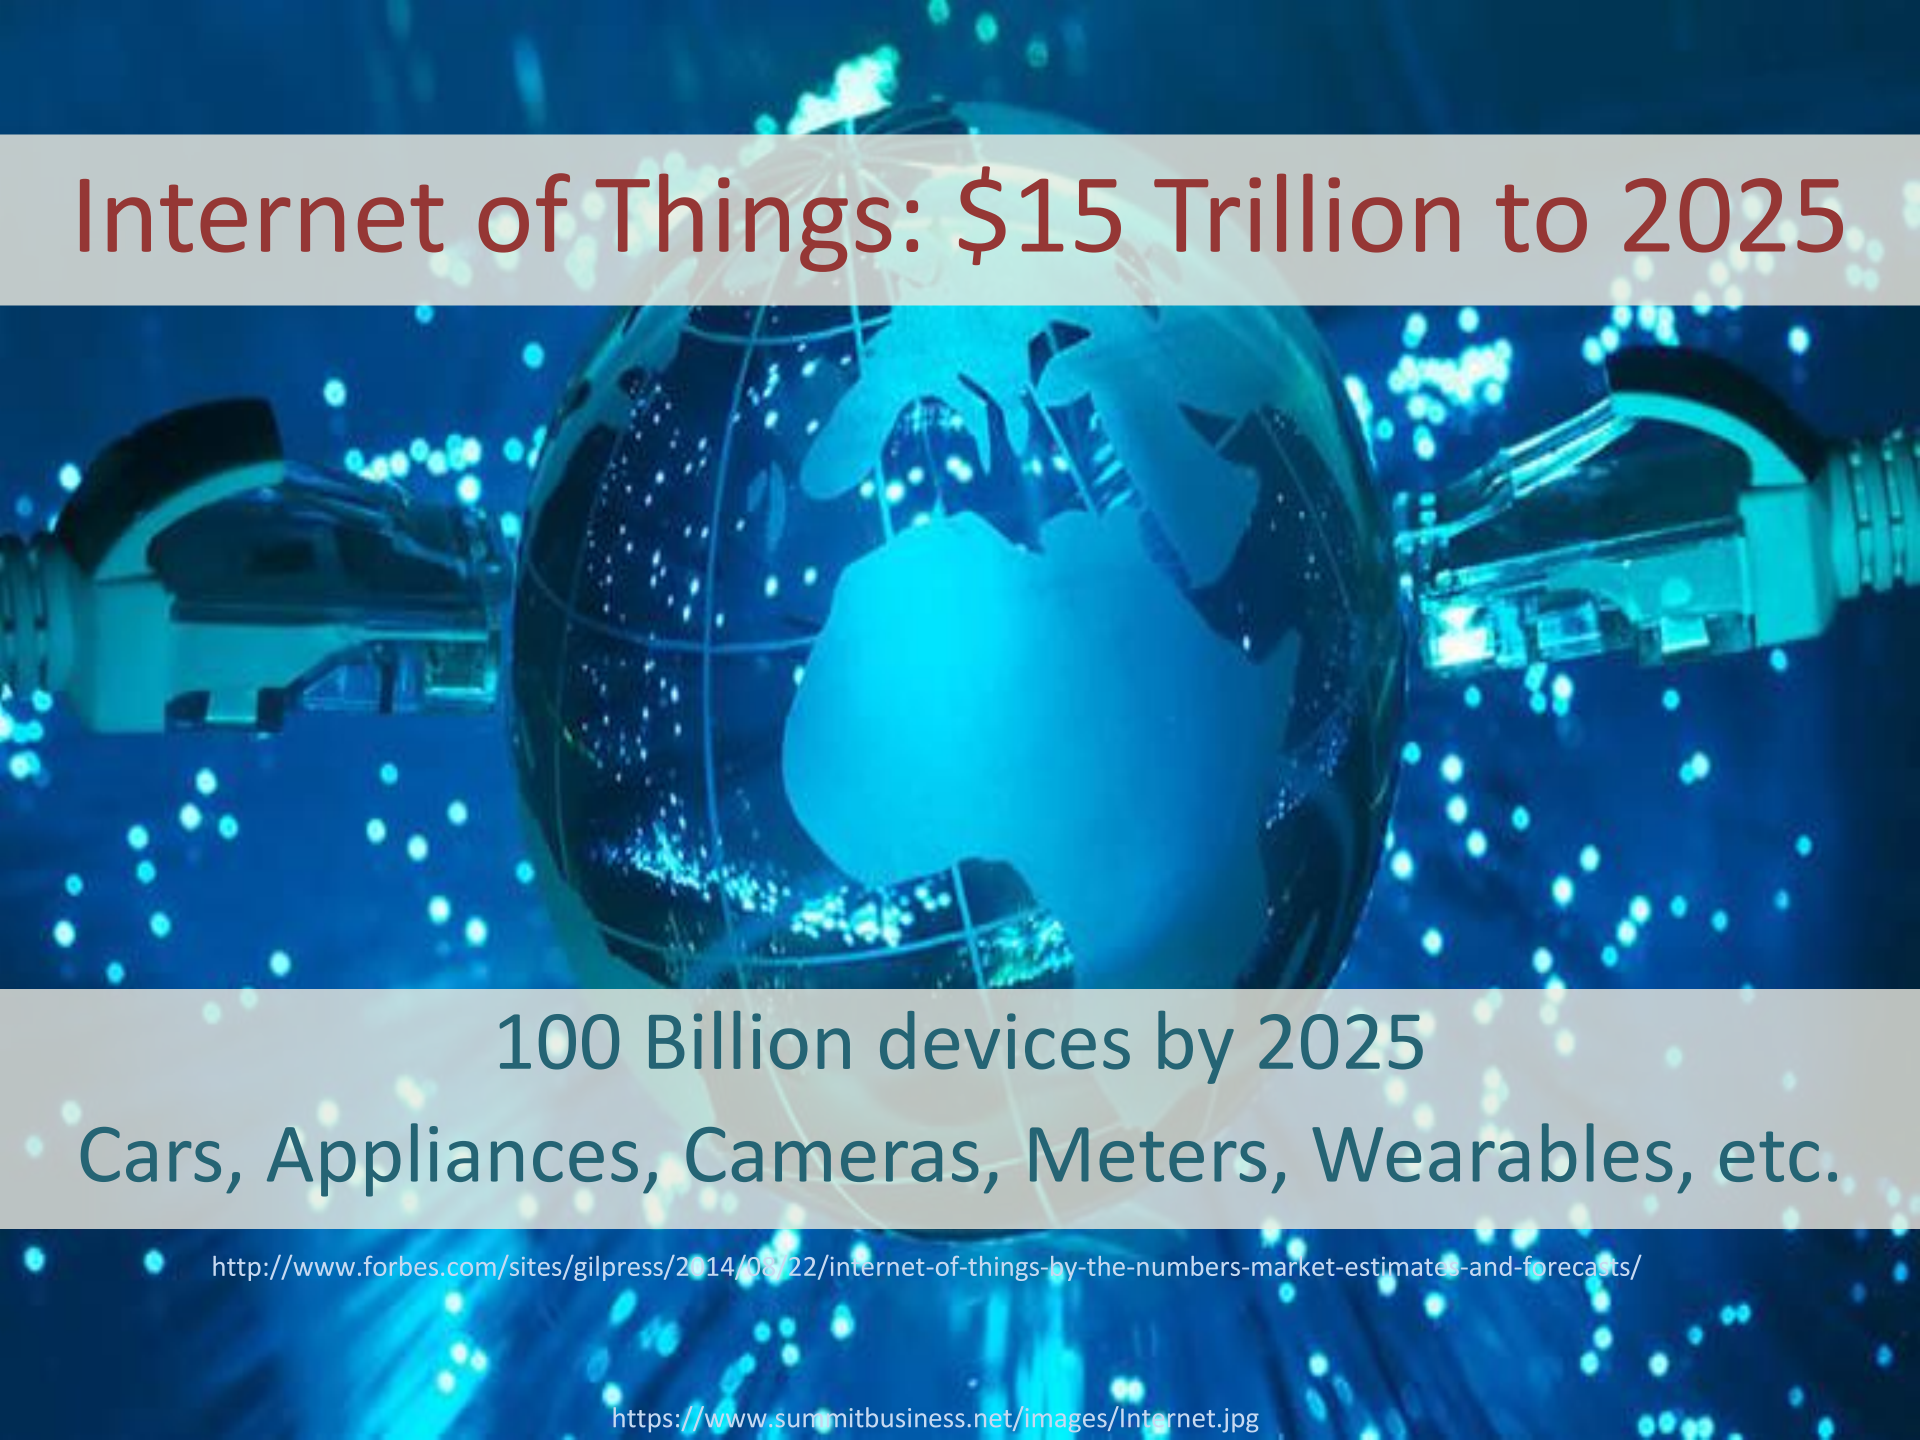

100 Billion devices by 2025

Cars, Appliances, Cameras, Meters, Wearables, etc.

http://www.forbes.com/sites/gilpress/2014/08/22/internet-of-things-by-the-numbers-market-estimates-and-forecasts/

https://www.summitbusiness.net/images/Internet.jpg

Robot Manufacturing: $10 Trillion to 2025

420 Chinese robot companies

Foxconn building 30K robots per year

1500 Dongguan “Robot Replace Human” factories

http://thisisrealmedia.com/2014/06/19/robotics-and-ethics-the-smart-car-by-ron-parlato/

Health Care: $10 Trillion to 2025

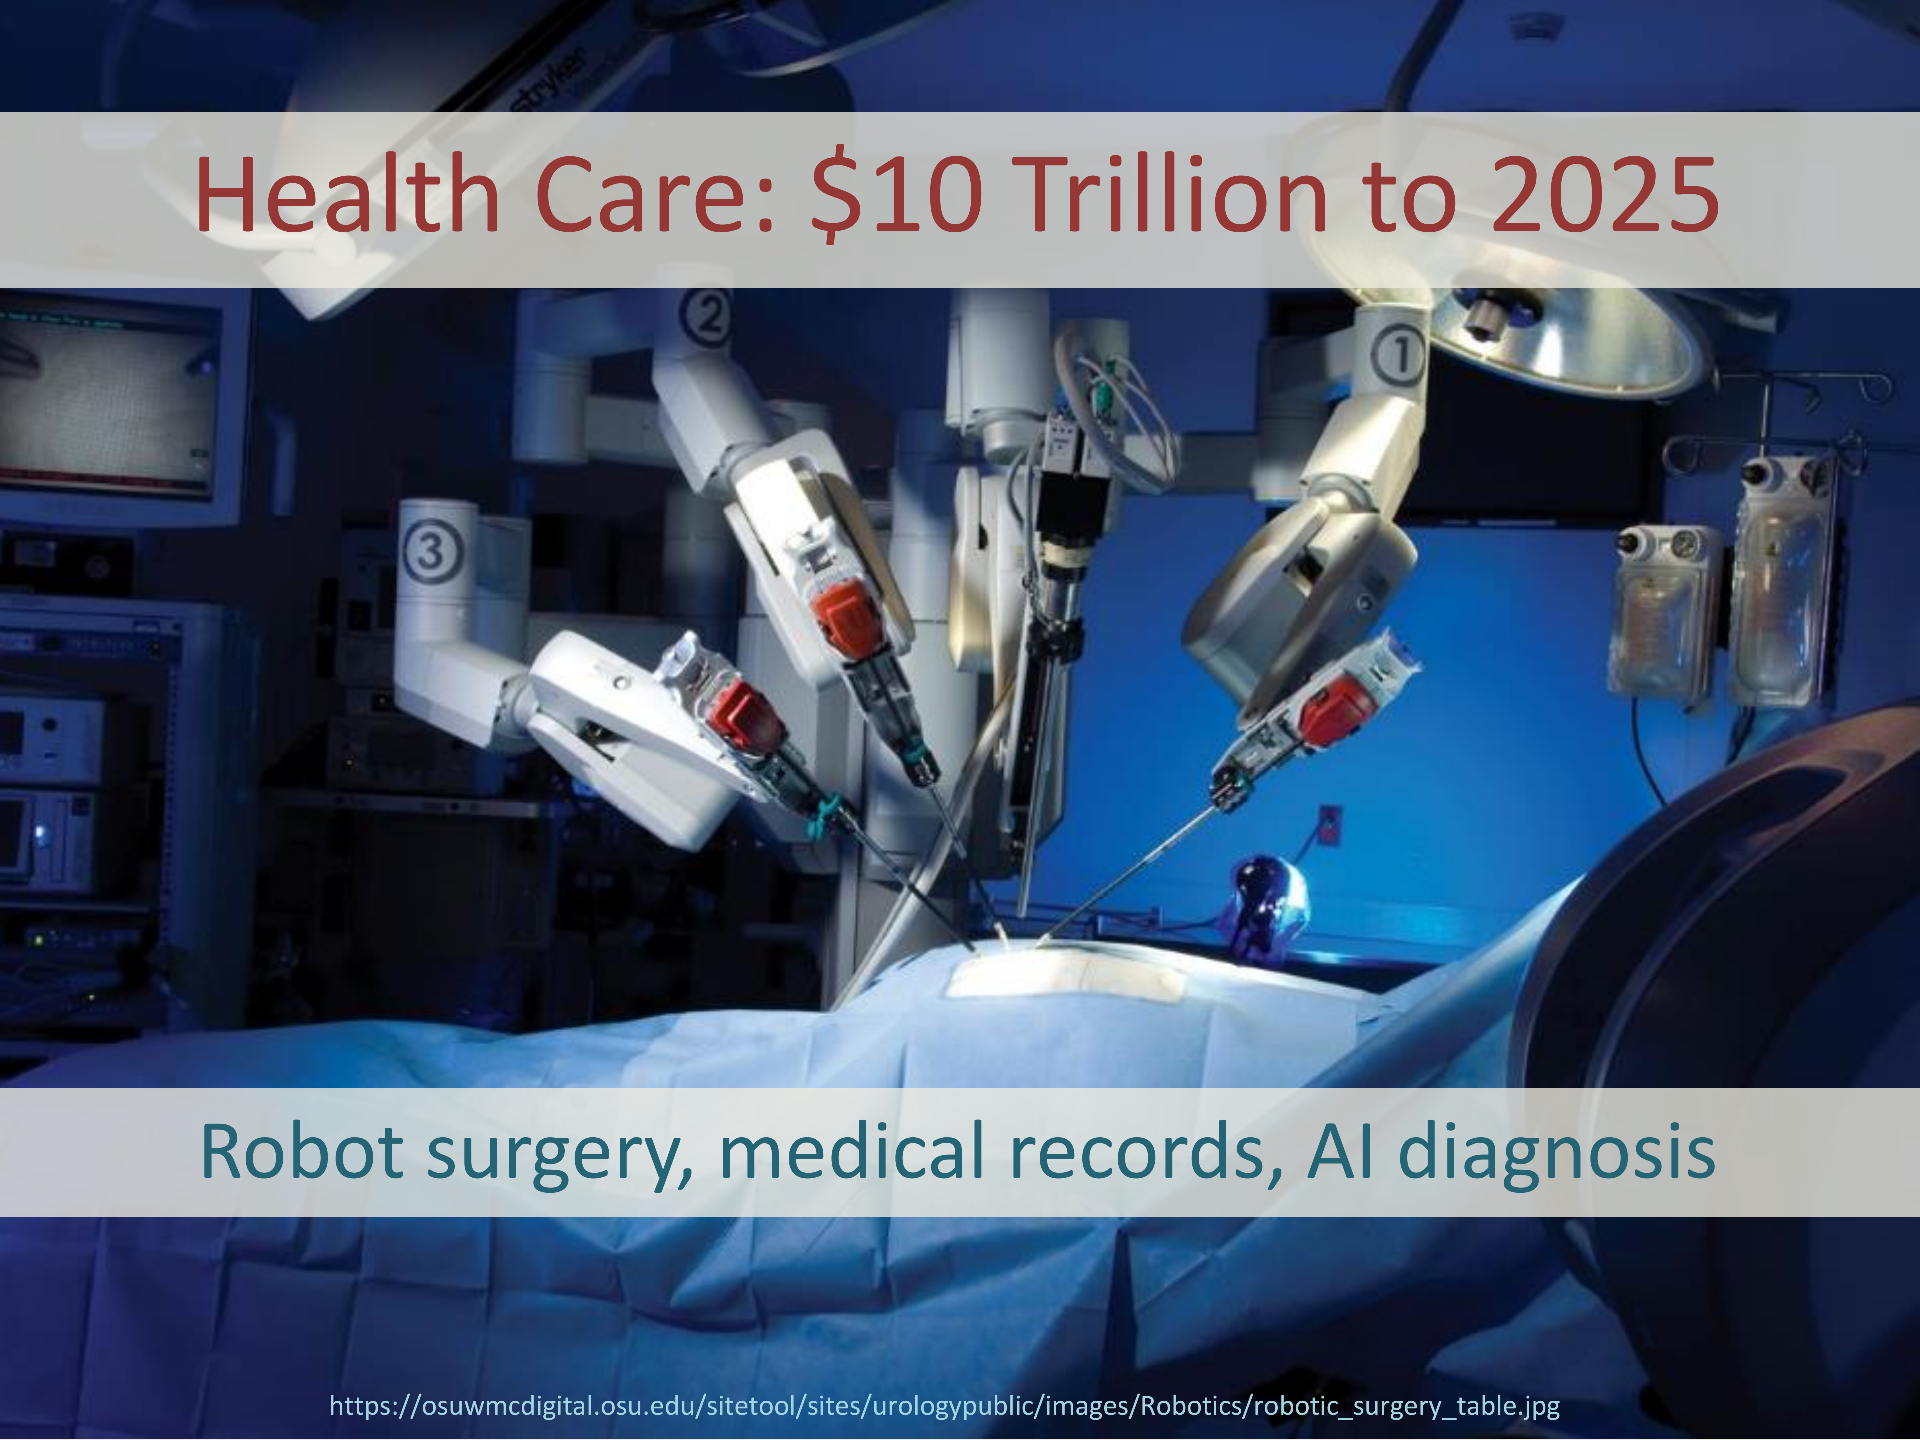

Robot surgery, medical records, AI diagnosis

https://osuwmcdigital.osu.edu/sitetool/sites/urologypublic/images/Robotics/robotic_surgery_table.jpg

Self-Driving Vehicles: $10 Trillion by 2025

Disrupt Dealers, Insurance, Parking, Finance, Trucking, Taxis

10 million jobs

http://www.theverge.com/2014/5/28/5756852/googles-self-driving-car-isnt-a-car-its-the-future

http://zackkanter.com/2015/01/23/how-ubers-autonomous-cars-will-destroy-10-million-jobs-by-2025/

3D Printing: $2 Trillion by 2025

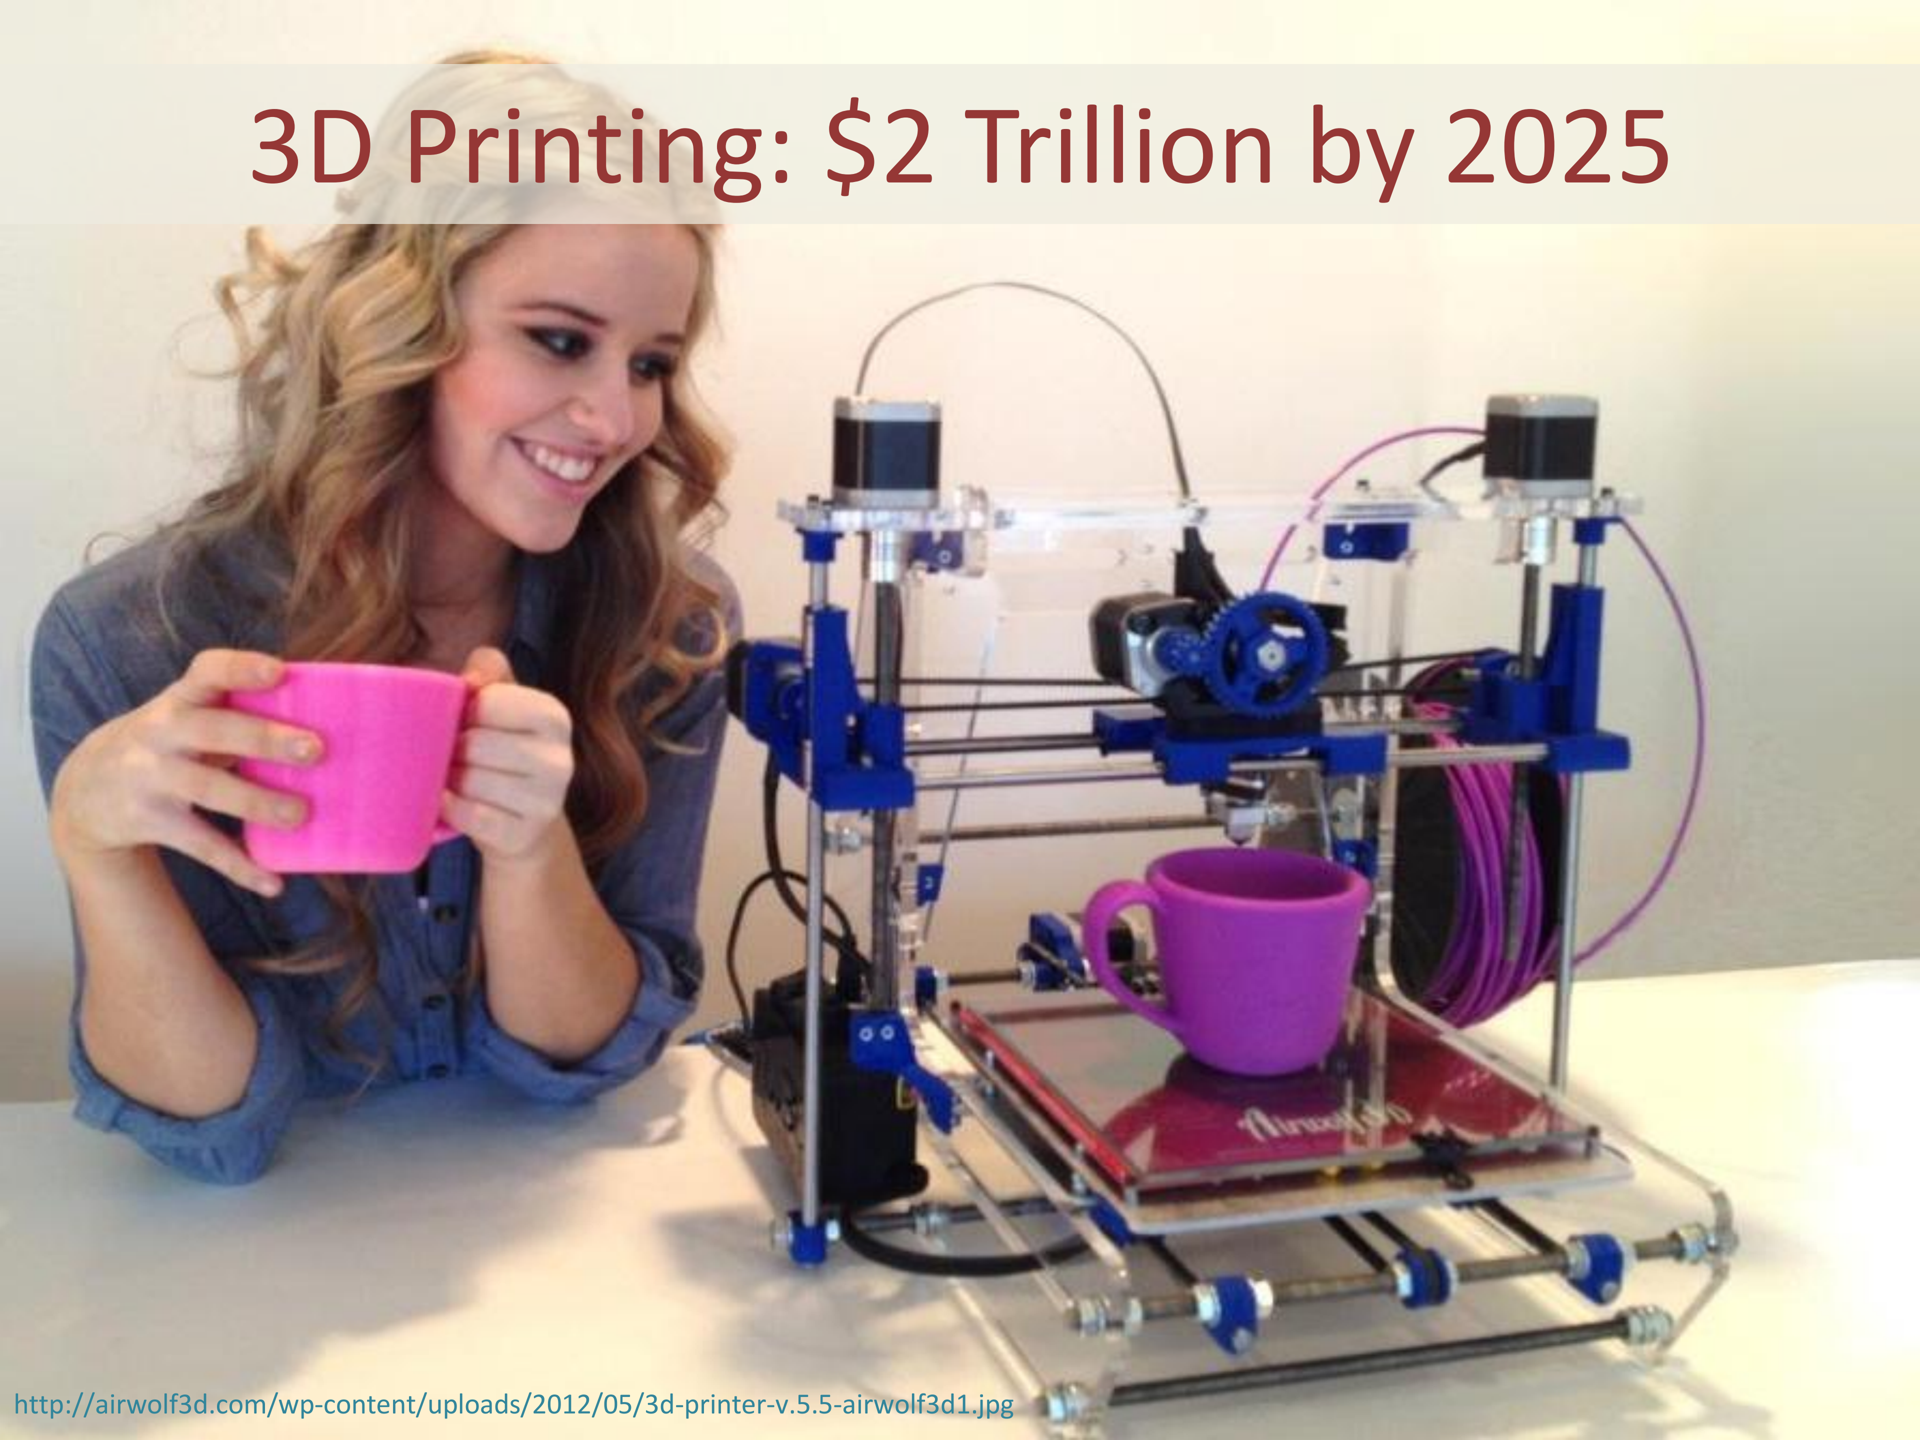

http://airwolf3d.com/wp-content/uploads/2012/05/3d-printer-v.5.5-airwolf3d1.jpg

WinSun 3D printed 12,000 sq ft villa

US Building construction: $1 Trillion/yr
5.8 million employees

http://3dprint.com/38144/3d-printed-apartment-building/

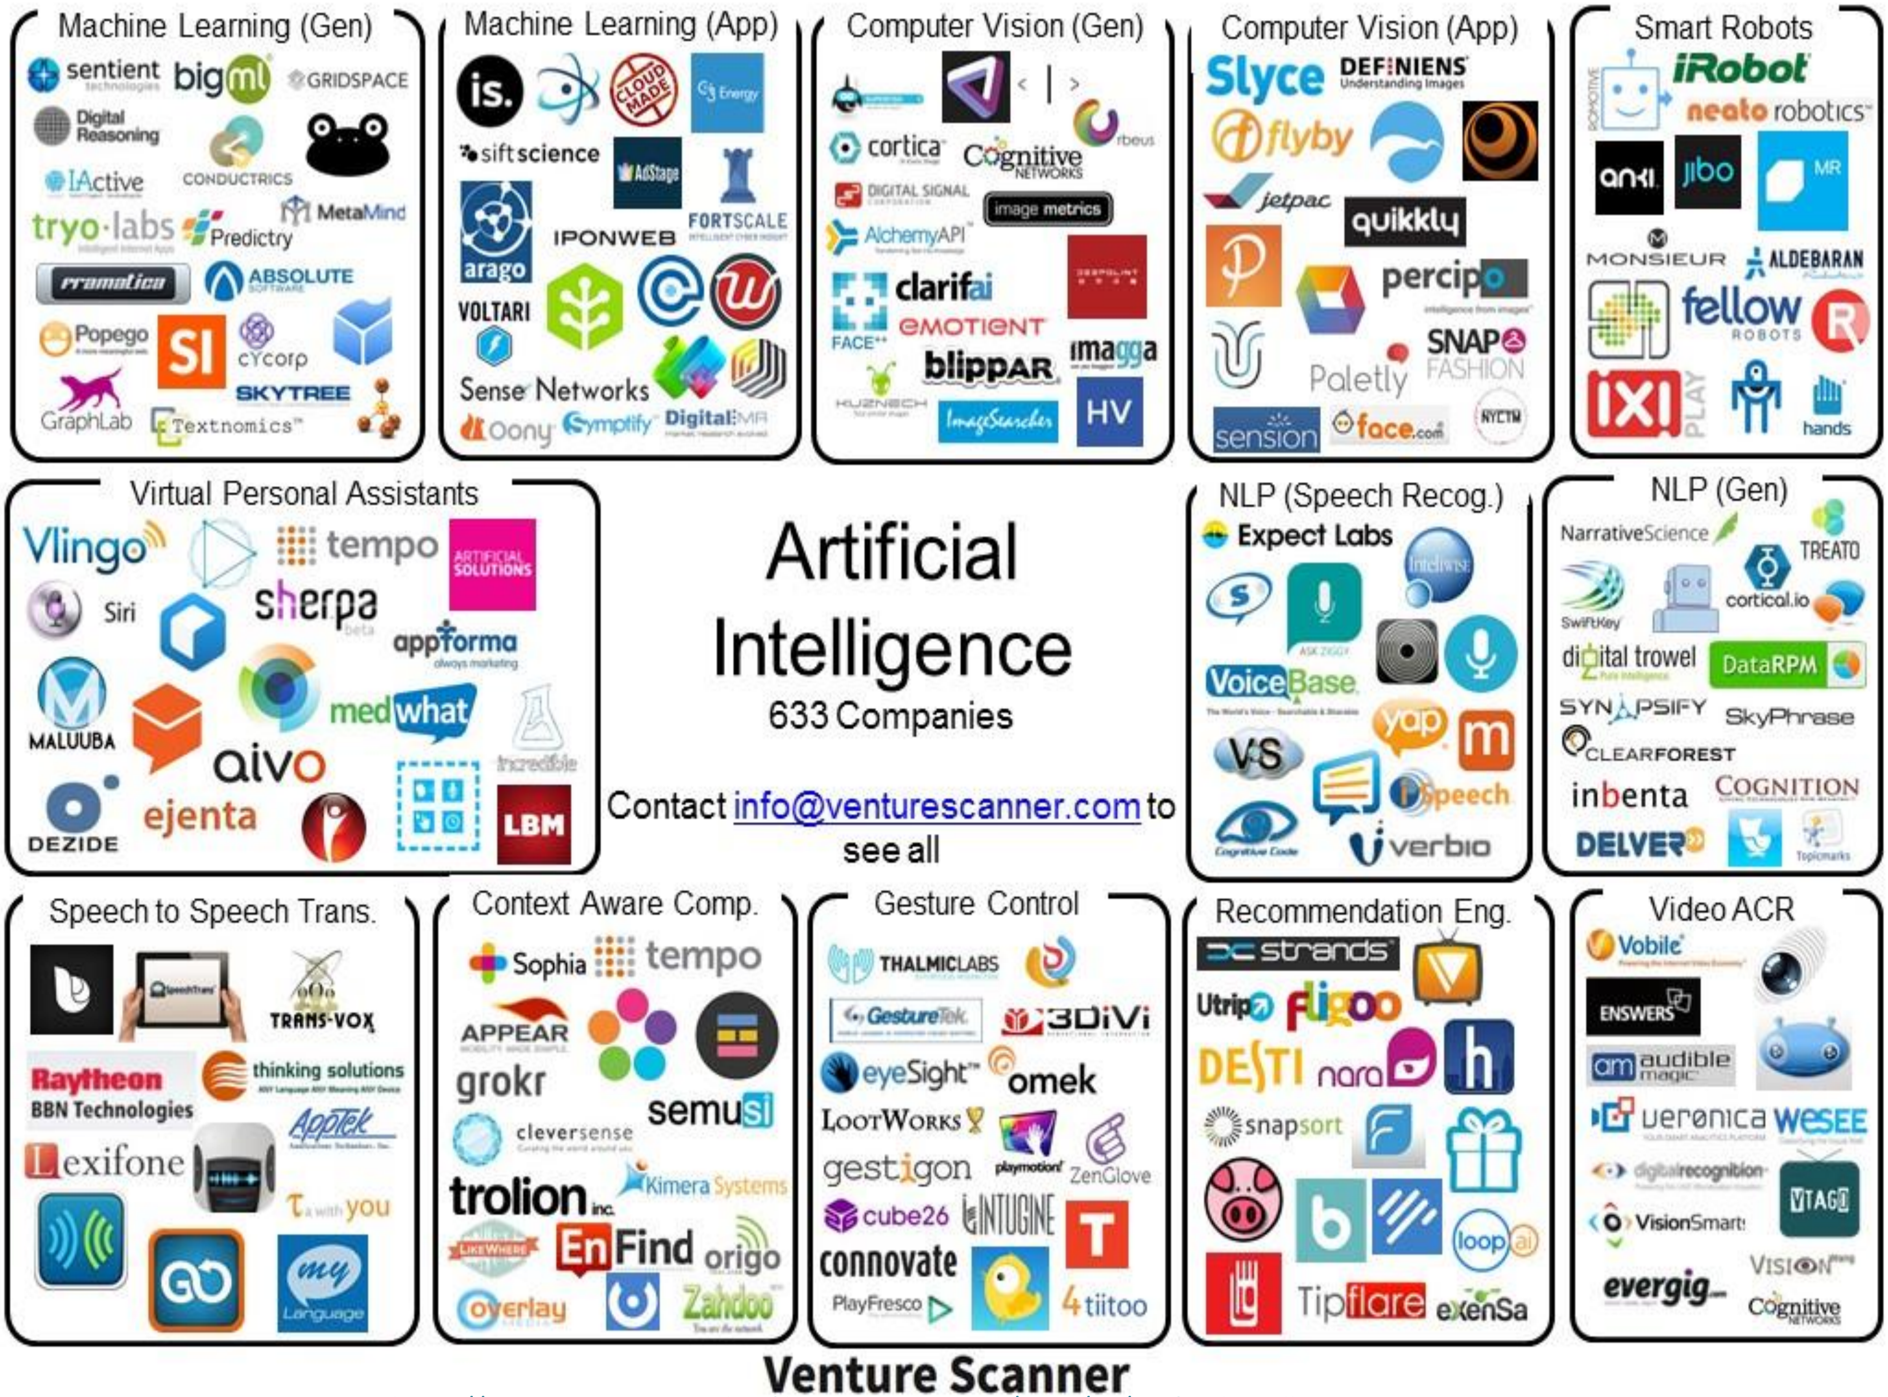

https://venturescannerinsights.files.wordpress.com/2015/01/artificial-intelligence-map.jpg https://venturescannerinsights.files.wordpress.com/2015/09/ai4.jpeg

Venture Investing in Artificial Intelligence

Venture Scanner

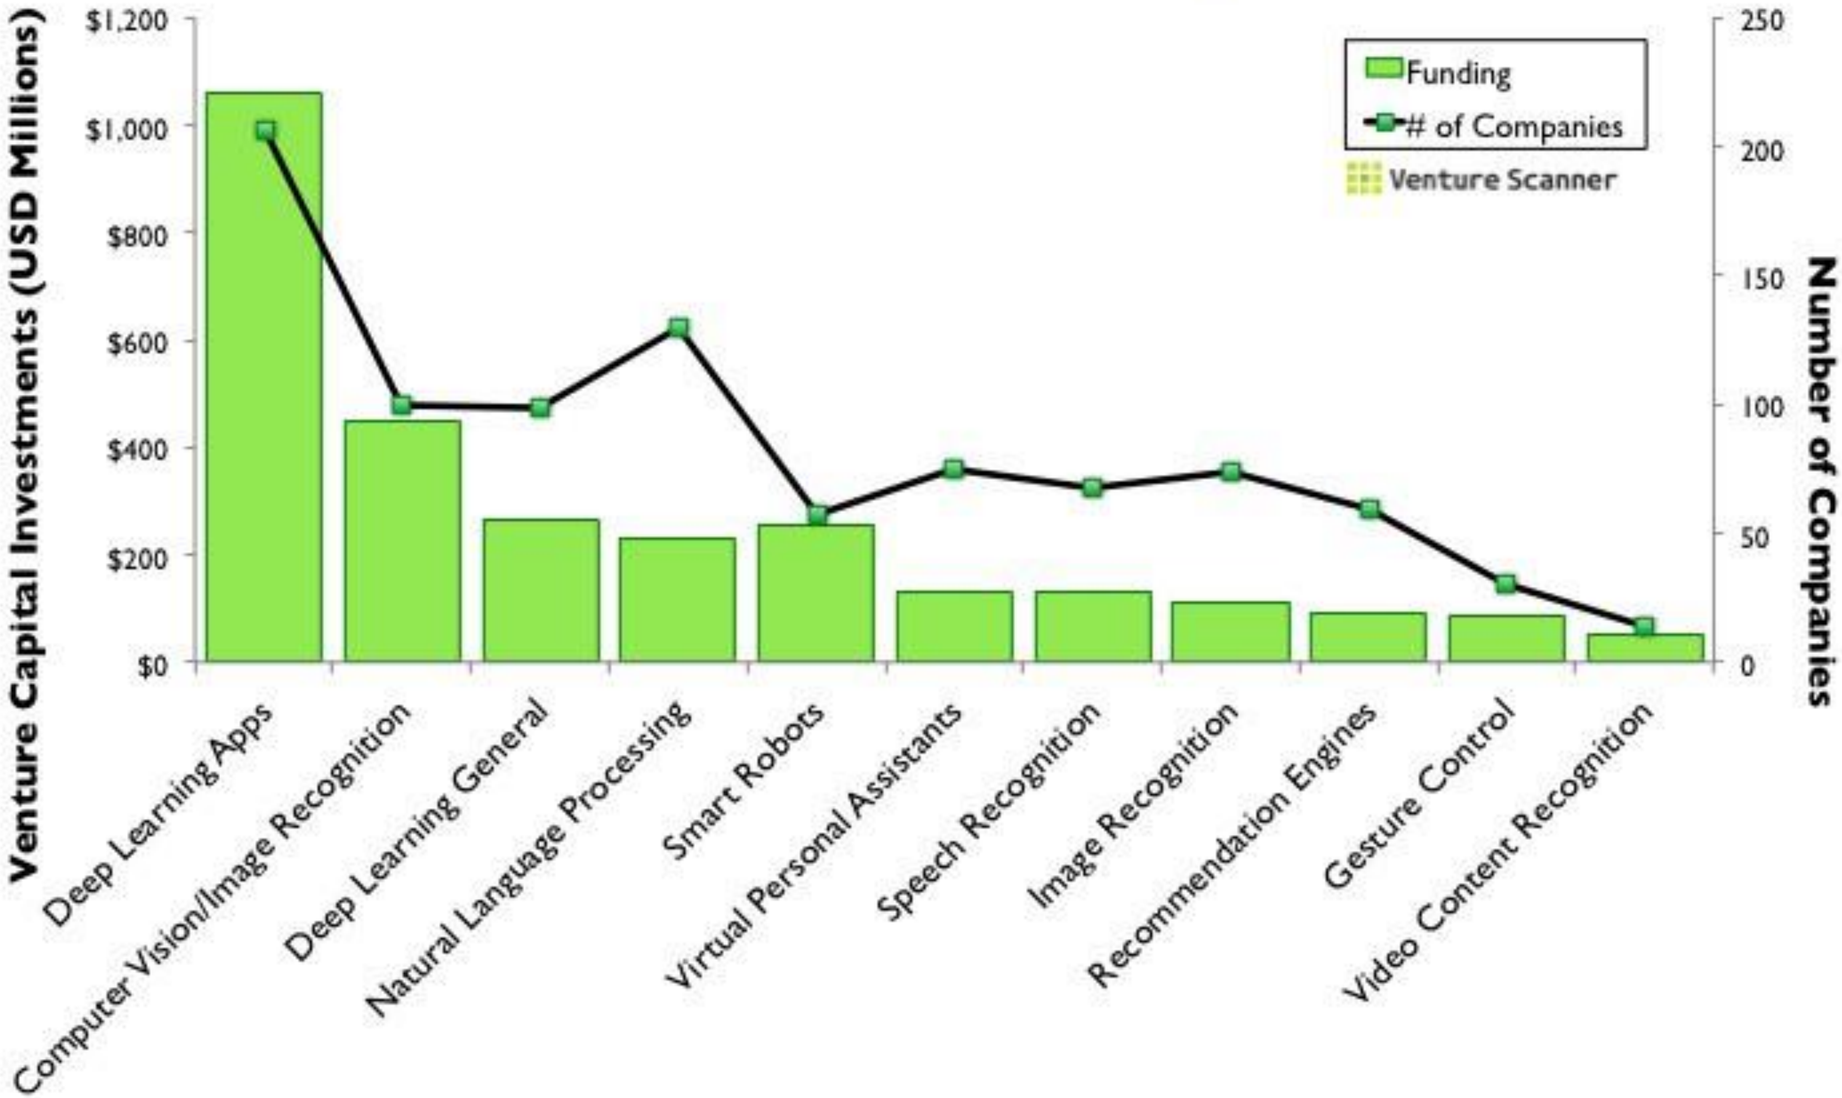

Contact us at info@venturescanner.com to see all 855 AI Startups

Deep Learning Neural Nets

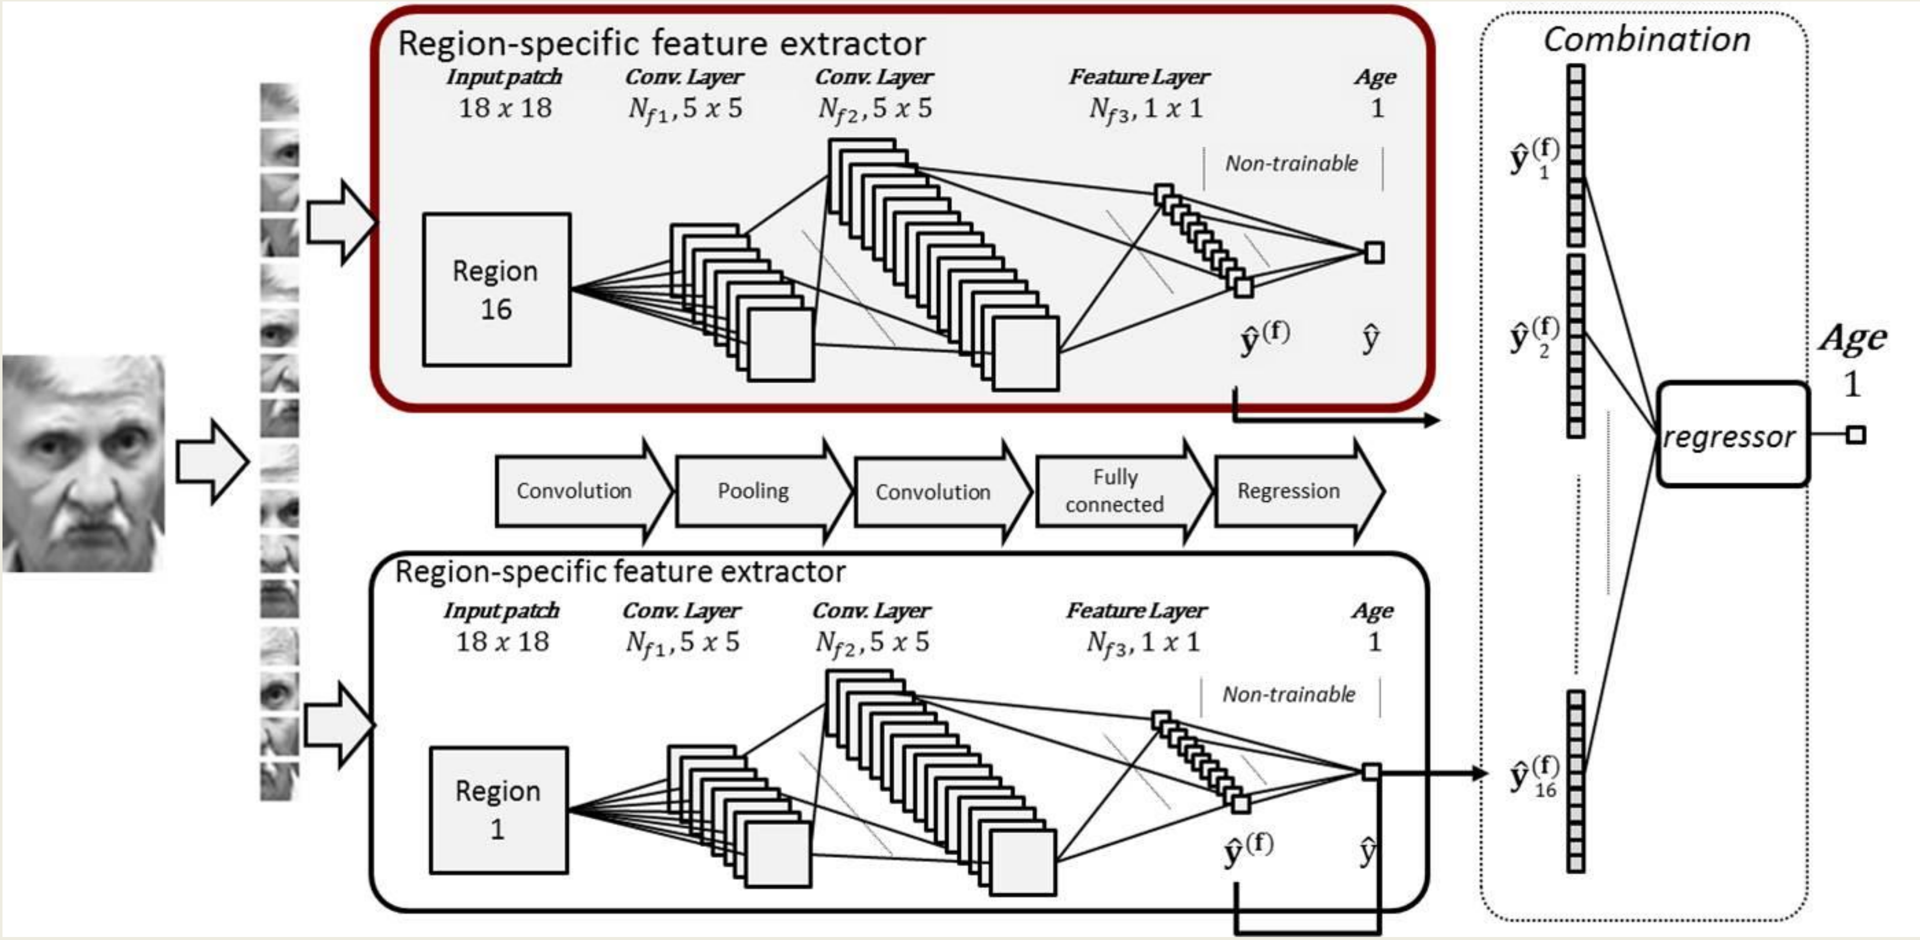

http://www.nicta.com.au/content/uploads/2015/02/deep.jpg

Deep Learning Successes

- Speech Recognition TIMIT 2009: Cortana, Skype, Google Now, Siri, Baidu, Nuance, etc.
- Image Recognition ImageNet 2012
- Image Captioning 2014
- Natural Language: Sentiment 2013, Translation 2014, Semantics 2014
- Drug Discovery: Merck Challenge 2012
- DeepMind 49 Atari Video Games 2015

http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html

A PROPOSAL FOR THE DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I.B.M. Corporation
C.E. Shannon, Bell Telephone Laboratories

August 31, 1955

2 months, 10 people, $13,500

Automate: Natural Language, Neural Nets, Self-Improvement, Abstraction, Creativity

http://www.dartmouth.edu/~vox/0607/0724/images/ai50.jpg

“Neats” vs. “Scruffies”

http://news.stanford.edu/news/2003/june18/mccarthy-618.html

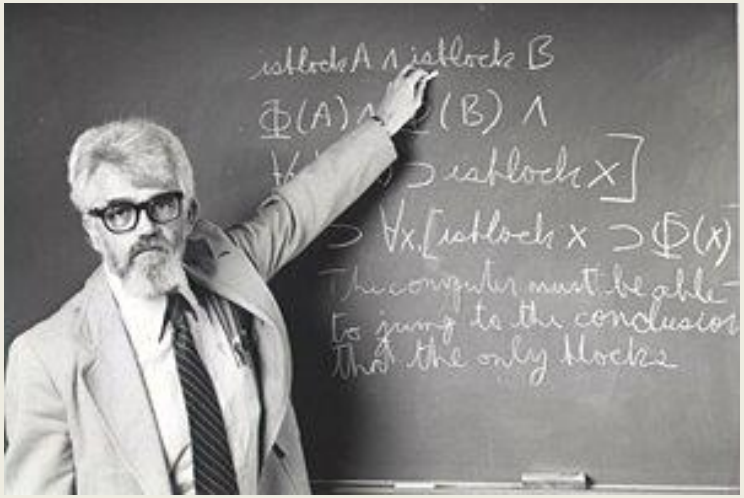

John McCarthy
1963 Stanford AI Lab

Mathematically Precise
Thinking = Logical Inference
Semantic Representations

http://www.bbc.co.uk/timelines/zq376fr

Marvin Minsky
1963 MIT MAC AI Group

Self-Organized
Adaptive Elements
Machine Learning
Emergent Semantics

“Neats” Rise/Fall/Rise/Fall

https://upload.wikimedia.org/wikipedia/commons/a/ae/Aristotle_Altemps_Inv8575.jpg

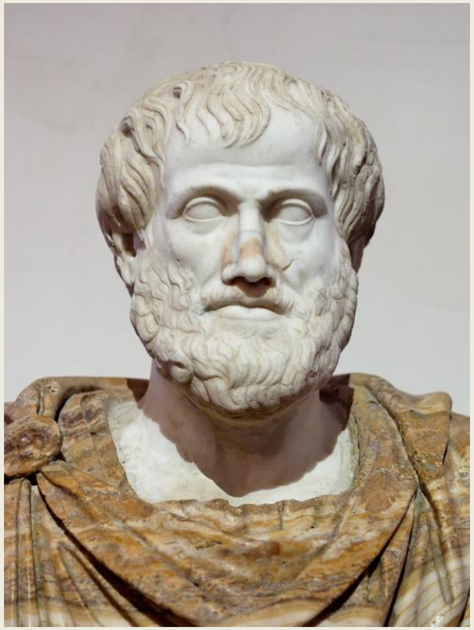

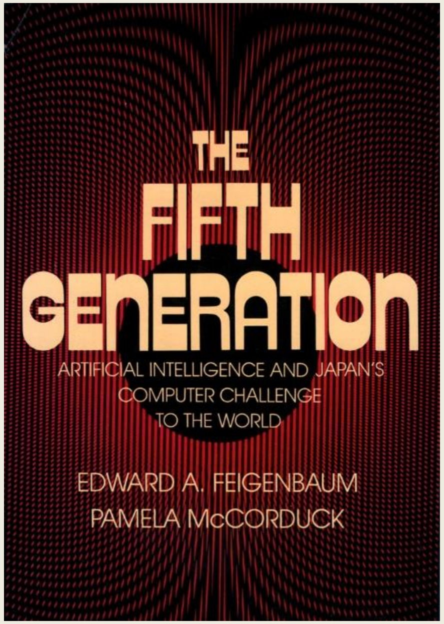

384BC Aristotle
1677 Leibniz
1879 Frege
1879 Cantor
1908 Zermelo
1936 Turing
1957 Chomsky
1959 McCarthy

1974 First AI Winter
1973 Lighthill Report
US, British funding cuts

1980 Expert Systems
1982 Fifth Generation
Prolog
1985 Bayes Nets

1989 Second AI Winter
1993 Expert Systems
Collapse
1990 Fifth Gen Fades
US funding cuts

https://upload.wikimedia.org/wikipedia/en/b/b3/Lighthill_3.jpeg

http://www.computerhistory.org/timeline/media/img/timeline_ai.robotics_1992.fifthgeneration.jpg

“Scruffies” Rise/Fall/Rise/Fall/Rise

http://www.nature.com/polopoly_fs/7.14689.1389093731!/image/deep-learning-graphic.jpg_gen/derivatives/landscape_400/deep-learning-graphic.jpg

http://www.rutherfordjournal.org/article040101.html

http://opticalengineering.spiedigitallibrary.org/article.aspx?articleid=1714547

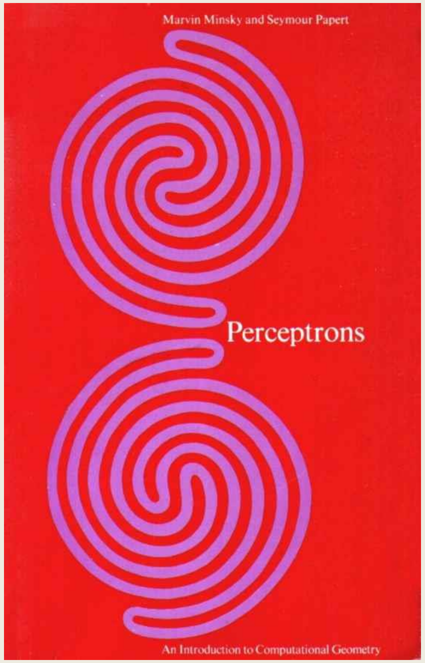

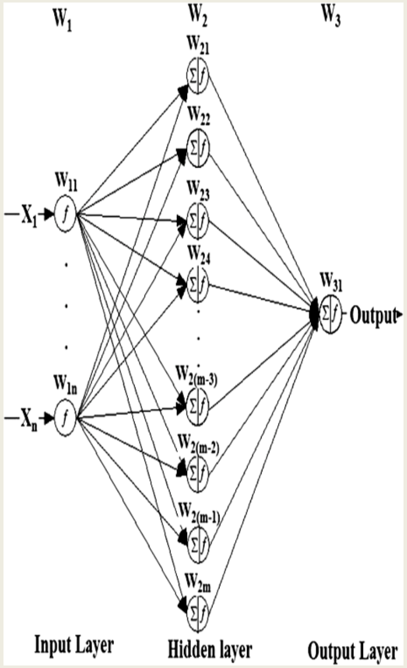

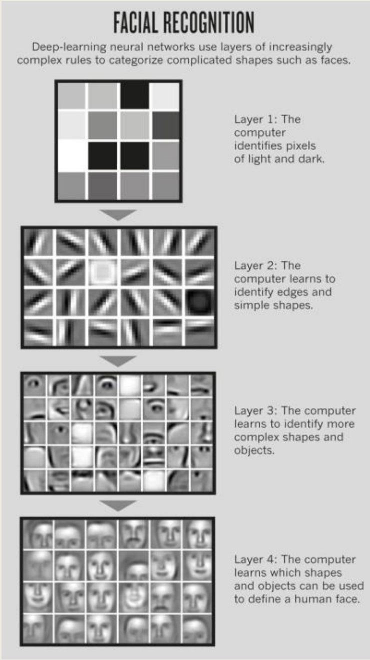

1957 Rosenblatt Perceptron

“The embryo of an electronic computer that [the Navy] expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence."

https://en.wikipedia.org/wiki/Perceptron

1969 Minsky & Papert
Can’t do XOR!

https://constructingkids.files.wordpress.com/2013/05/minsky-papert-71-csolomon-x640.jpg
http://www.i-programmer.info/images/stories/BabBag/AI/book.jpg

Backpropagation
1986 Rumelhart
(1963 Bryson
1974 Werbos)

Deep Learning
2007 Hinton
(1989 LeCun
1992 Schmidhuber)

“Neat” Software Compiler

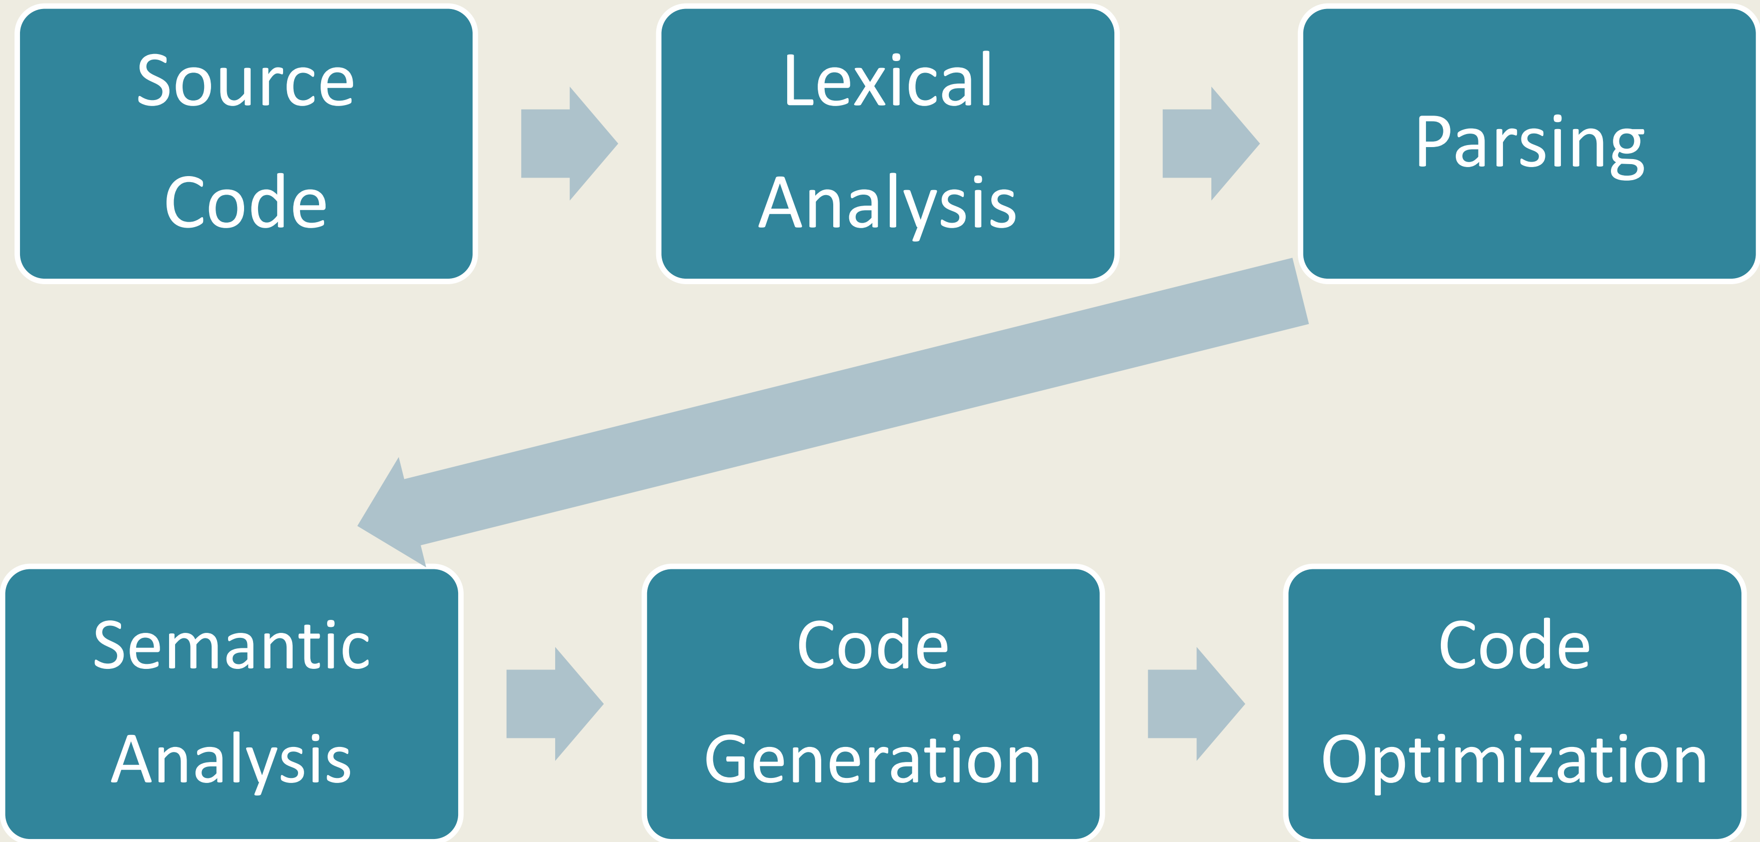

“Neat” Language Translator?

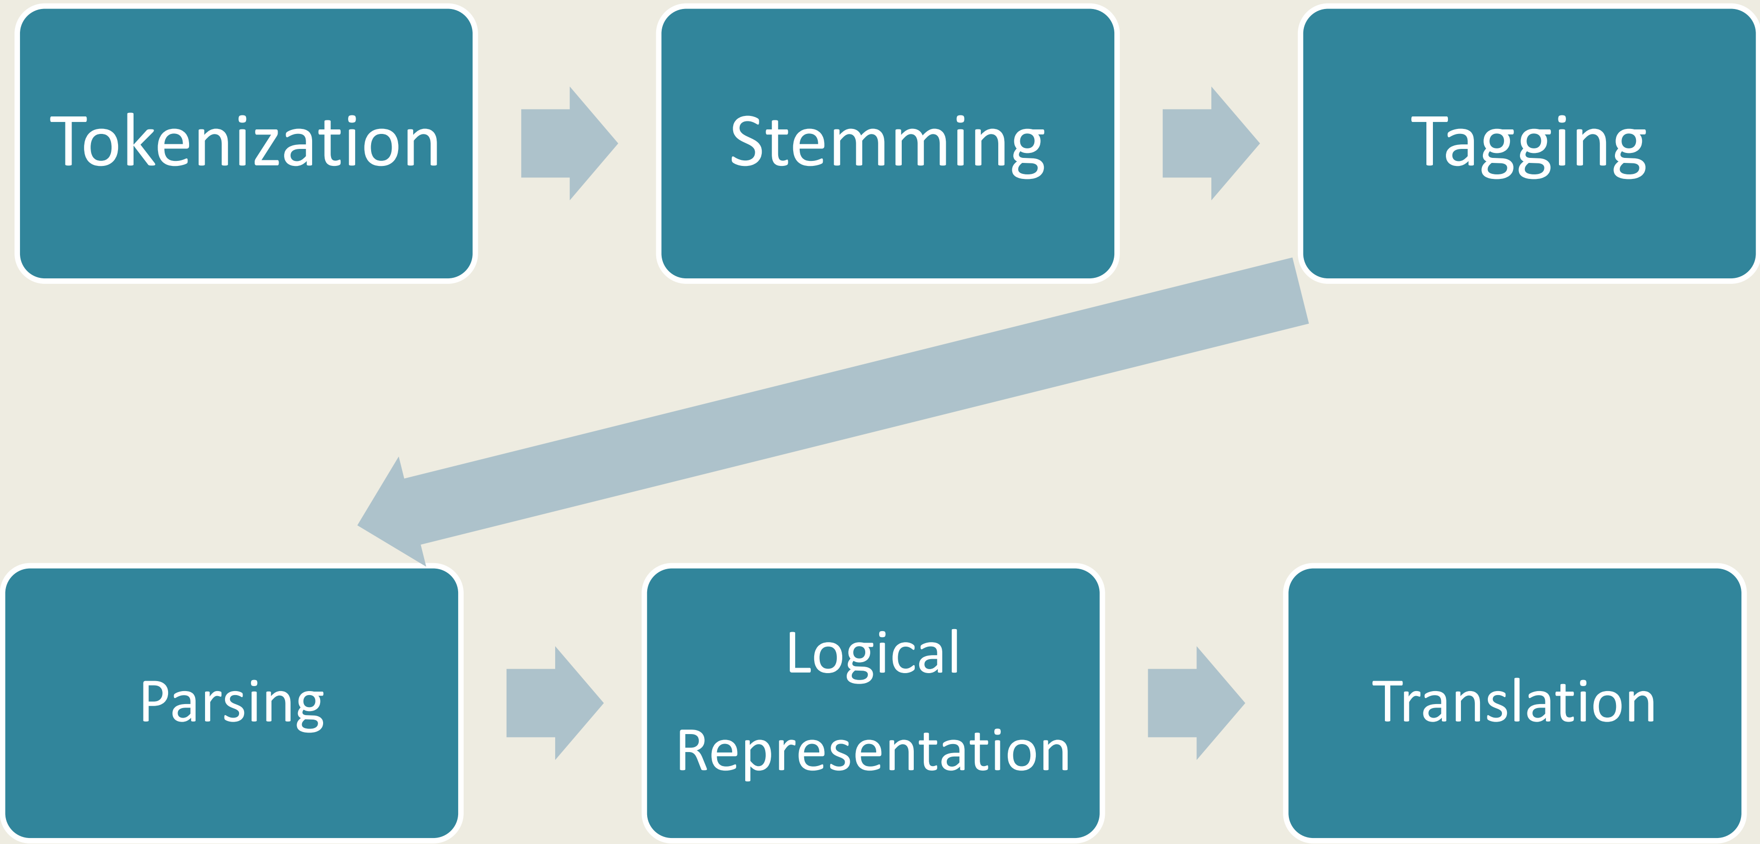

1957 Chomsky Grammar

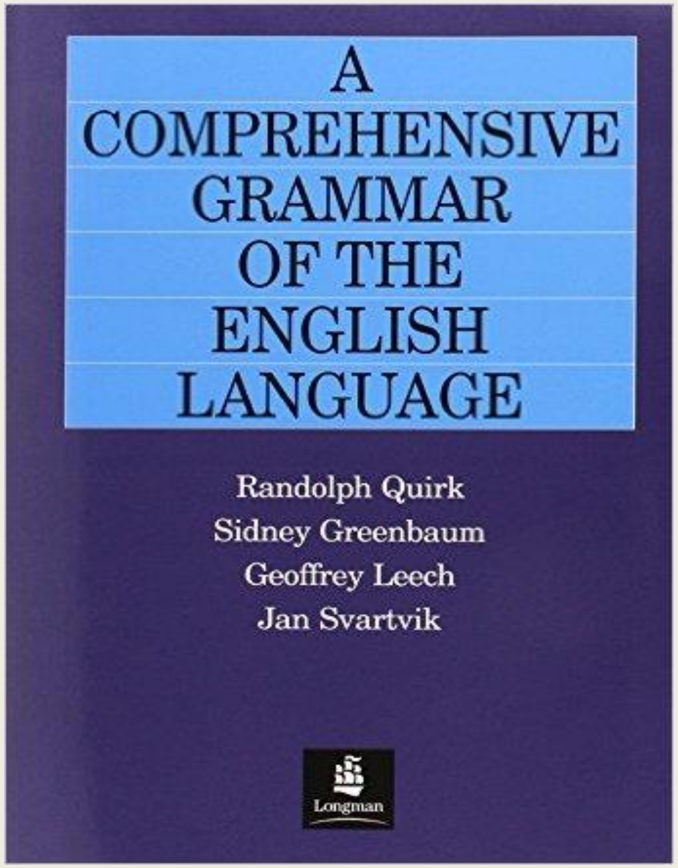

1970 Montague Semantics

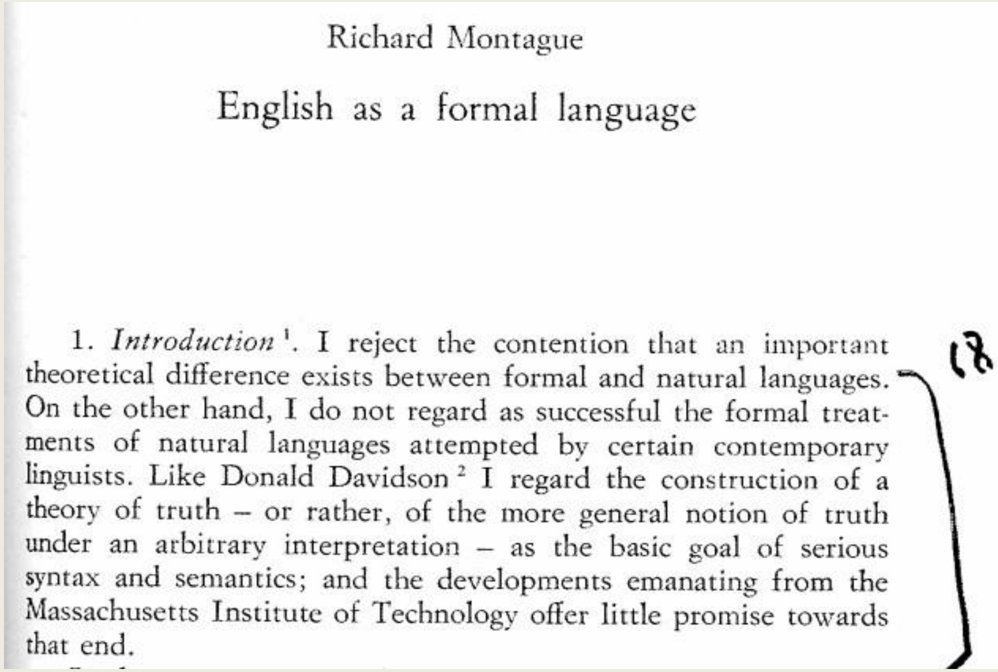

Richard Montague

English as a formal language

1. *Introduction*[1]. I reject the contention that an important theoretical difference exists between formal and natural languages. On the other hand, I do not regard as successful the formal treatments of natural languages attempted by certain contemporary linguists. Like Donald Davidson[2] I regard the construction of a theory of truth – or rather, of the more general notion of truth under an arbitrary interpretation – as the basic goal of serious syntax and semantics; and the developments emanating from the Massachusetts Institute of Technology offer little promise towards that end.

"English as a Formal Language". In: Bruno Visentini (ed.): *Linguaggi nella società e nella tecnica*. Mailand 1970, 189–223.

1792 Pages!

http://www.amazon.com/Comprehensive-Grammar-English-Language-General/dp/0582517346/ref=sr_1_2

Linguistic Rules are Complicated!

- A velvet new comfortable dress – **INCORRECT**
- A comfortable new velvet dress – **CORRECT**

1	2	3	4	5	6	7	8
Opinion	Size	Shape	Age	Color	Nationality/Origin	Material	Purpose

Note: *Not everyone agrees on this order, and there may be exceptions*

http://www.espressoenglish.net/order-of-adjectives-in-english/

Computational Linguistics

December 2006, Vol. 32, No. 4, Pages 527-549
Posted Online November 21, 2006.
(doi:10.1162/coli.2006.32.4.527)

N-gram-based Machine Translation

José B. Mariño[*] Rafael E. Banchs[*] Josep M. Crego[*] Adrià de Gispert[*] Patrik Lambert[*] José A. R. Fonollosa[*] Marta R. Costa-jussà[*]
[*] Department of Signal Theory and Communications, Campus Nord, Barcelona 08034, Spain.

PDF (284.888 KB) | **PDF Plus** (350.623 KB)

This article describes in detail an n-gram approach to statistical machine translation. This approach consists of a log-linear combination of a translation model based on n-grams of bilingual units, which are referred to as tuples, along with four specific feature functions. Translation performance, which happens to be in the state of the art, is demonstrated with Spanish-to-English and English-to-Spanish translations of the European Parliament Plenary Sessions (EPPS).

2006: Simple n-gram models with lots of data beat complicated hand built linguistic models!

http://www.mitpressjournals.org/doi/abs/10.1162/coli.2006.32.4.527#.VjWP0_mfM-U

EXPERT OPINION

Contact Editor: **Brian Brannon**, bbrannon@computer.org

The Unreasonable Effectiveness of Data

Alon Halevy, Peter Norvig, and Fernando Pereira, *Google*

2009: And data is cheap and plentiful!

Much cheaper than linguists or programmers!

http://www.computer.org/csdl/mags/ex/2009/02/mex2009020008-abs.html

1962: Roger Shepard Cognitive Geometry

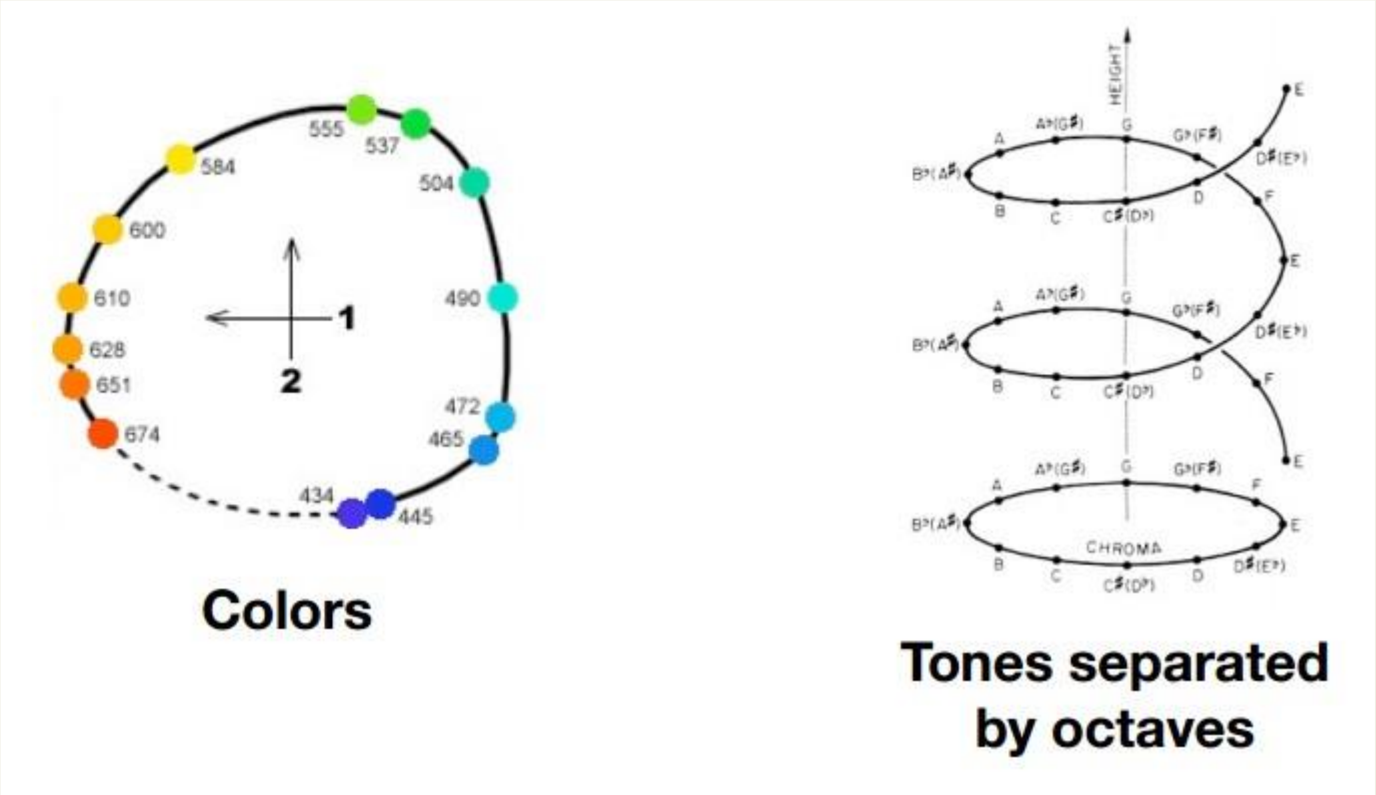

http://link.springer.com/article/10.1007/BF02289630

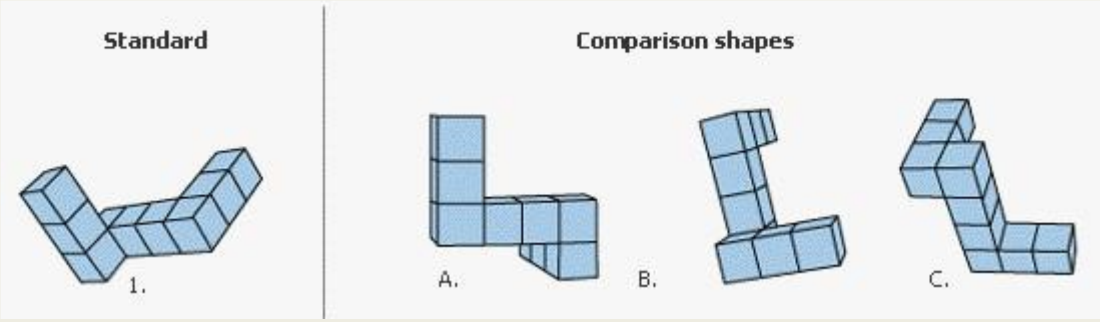

https://psychlopedia.wikispaces.com/mental+rotation

Word2Vec – Mikolov 2013

- Distributional Semantics – Firth 1957
- Represent words by vectors
- Close vectors represent similar contexts
- Certain relations represented by translation:

King – Man + Woman = Queen

- Also tense, temperature, location, plurals,…

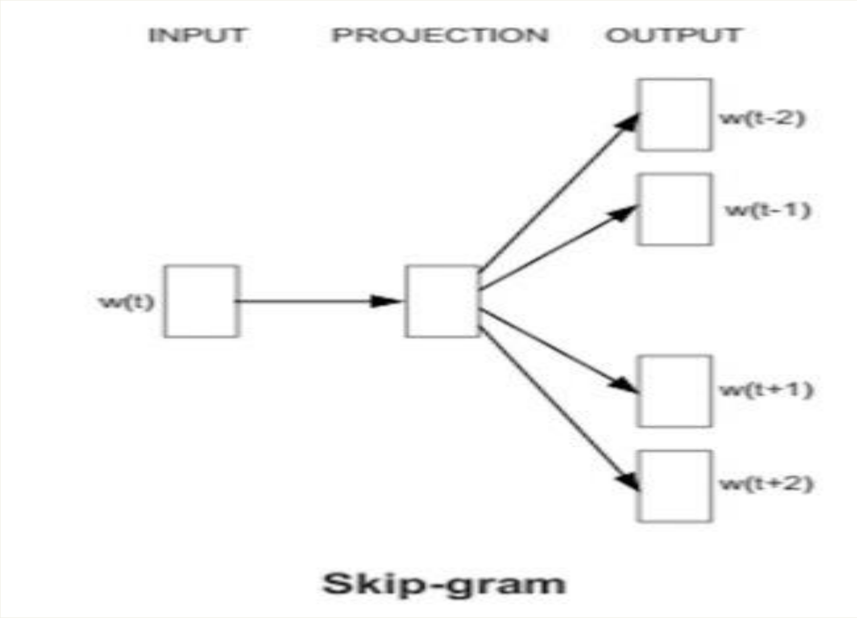

http://deeplearning4j.org/word2vec.html

2013 Mikolov:

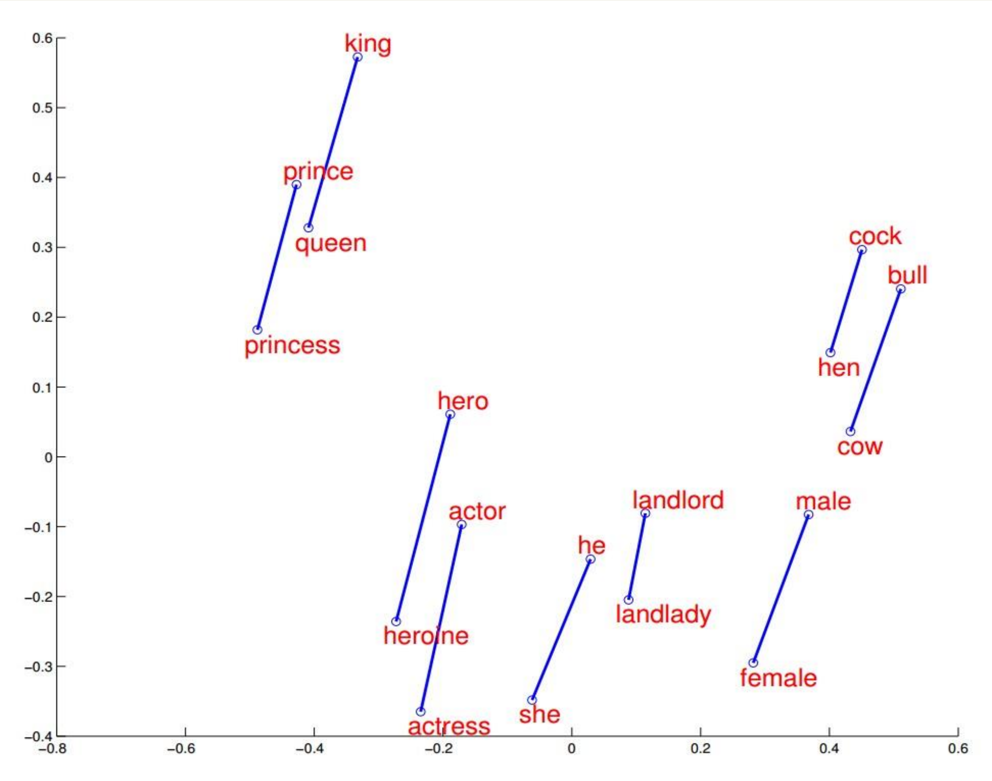

https://drive.google.com/file/d/0B7XkCwpI5KDYRWRnd1RzWXQ2TWc/edit?usp=sharing

Why? Same context shift for all male -> female

The man ate *his* lunch.

The king ate *his* lunch.

The woman at *her* lunch.

The queen ate *her* lunch.

More Semantic Relations

- Paris – France + Italy = Rome https://code.google.com/p/word2vec/
- Human – Animal = Ethics http://byterot.blogspot.com/2015/06/five-crazy-abstractions-my-deep-learning-word2doc-model-just-did-NLP-gensim.html
- Obama – USA + Russia = Putin
- Library – Books = Hall
- Biggest – Big + Small = Smallest http://arxiv.org/pdf/1301.3781.pdf
- Ethical – Possibly + Impossibly = Unethical
- Picasso – Einstein + Scientist = Painter
- Forearm – Leg + Knee = Elbow http://deeplearning4j.org/word2vec.html
- Architect – Building + Software = Programmer

Marr’s “Neat” Vision Pipeline

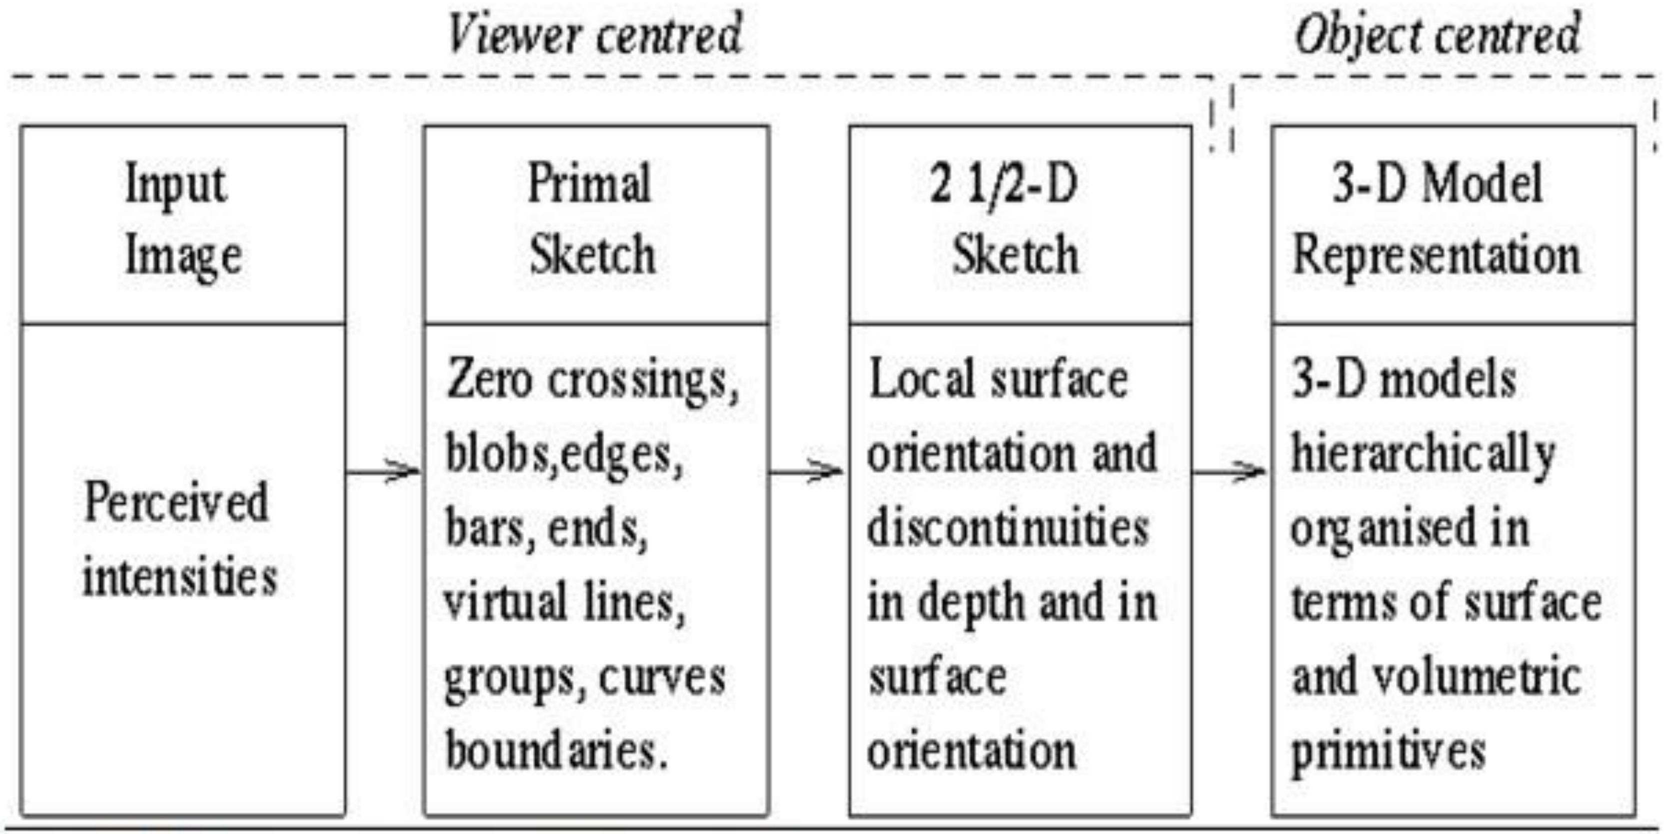

Figure 1: Marr's representational framework

http://homepages.inf.ed.ac.uk/rbf/CVonline/LOCAL_COPIES/GOMES1/marr.html

http://www.amazon.com/Vision-Computational-Investigation-Representation-Information/dp/0262514621/ref=sr_1_2

Deep Neural Net Face Recognition

http://arxiv.org/abs/1503.03832

Google FaceNet, June 2015
Record accuracy 99.63% on Labeled Faces in the Wild dataset
Cuts best previous error rate by 30%
22 layer feedforward net, 140M weights, 1.6 GFLOP/image, conv/pool/norm
Trained on triples pushing same faces together, different apart

CMU OpenFace, Oct. 2015
Open Source version of FaceNet
84.83% accuracy, <.1 training faces

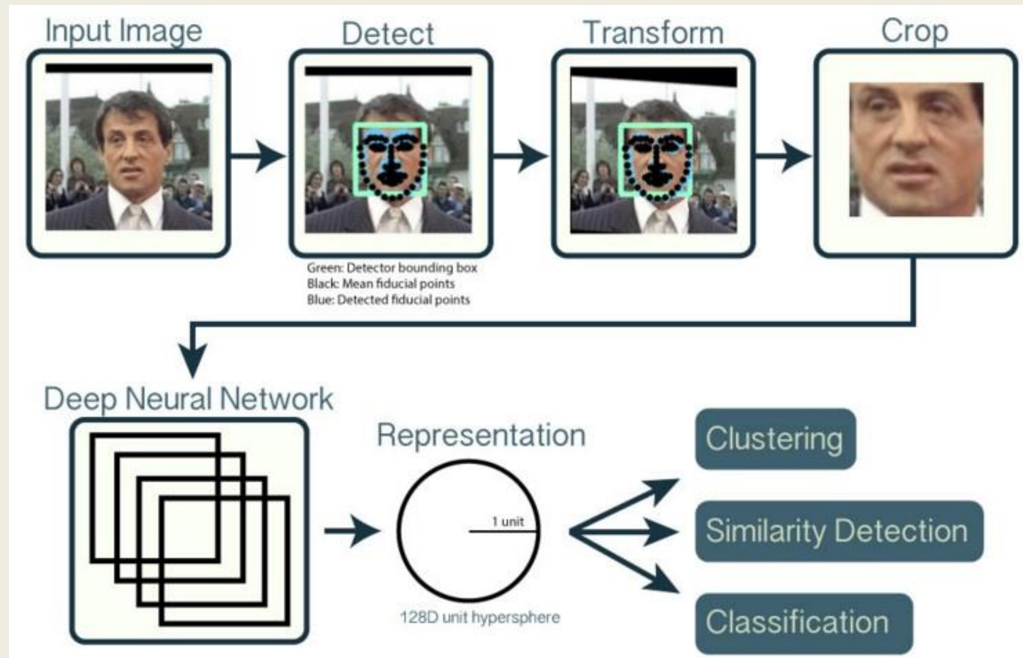

https://github.com/cmusatyalab/openface

Please use responsibly!

We do not support the use of this project in applications that violate privacy and security. We are using this to help cognitively impaired users to sense and understand the world around them.

http://www.aliexpress.com/cheap/cheap-image-sensor-module.html

http://www.ebay.com/sch/i.html?_nkw=cmos+image+sensor&_sop=15

Cheap Cameras
+
Face Recognition
+
Body Recognition

600TVL 1/4" CMOS image sensor board Pixelplus PC70...
US $3.20 / piece
Shipping: US $2.56 / piece
Min. Order: 1 piece
14 | 16 Orders

MT9V143M05STC DIGITAL IMAGE SENSOR 1/4 INCH VGA CMOS ACTIVE PIXEL Camera Chip
$2.95 From Israel
or Best Offer
+$3.95 shipping

$3.20 on Alibaba $2.95 on ebay

Watch this open-source program recognize faces in real time

by LAUREN HOCKENSON Tweet — 17d ago in DESIGN & DEV

http://thenextweb.com/dd/2015/10/15/watch-this-open-source-program-recognize-faces-in-real-time/

Facebook Can Now Recognize You in Photos Without Even Seeing Your Face

BY ADARSH VERMA · JUNE 23, 2015

http://fossbytes.com/facebook-can-now-recognize-you-in-photos-without-even-seeing-your-face/

Cameras know you by your walk

Improvements in gait analysis mean your characteristic way of walking could soon be used to identify you – wherever you are

https://www.newscientist.com/article/mg21528835-600-cameras-know-you-by-your-walk/

=

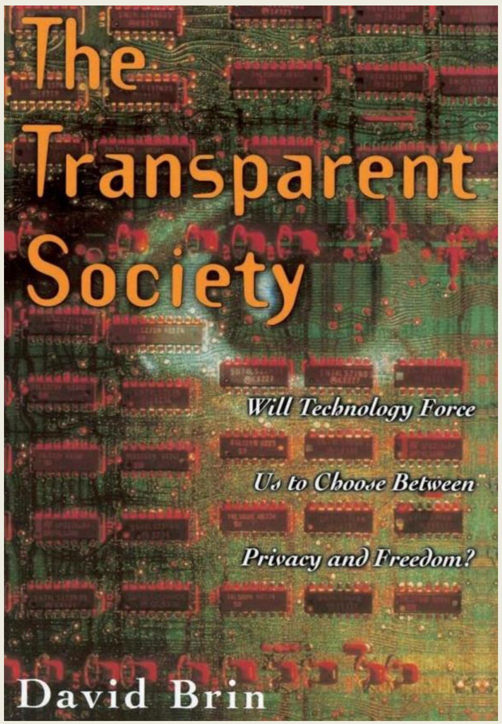

Brin's "Transparent Society"

http://www.amazon.com/Transparent-Society-Technology-Between-Privacy/dp/0738201448/ref=sr_1_1

Biological Networks are Recurrent

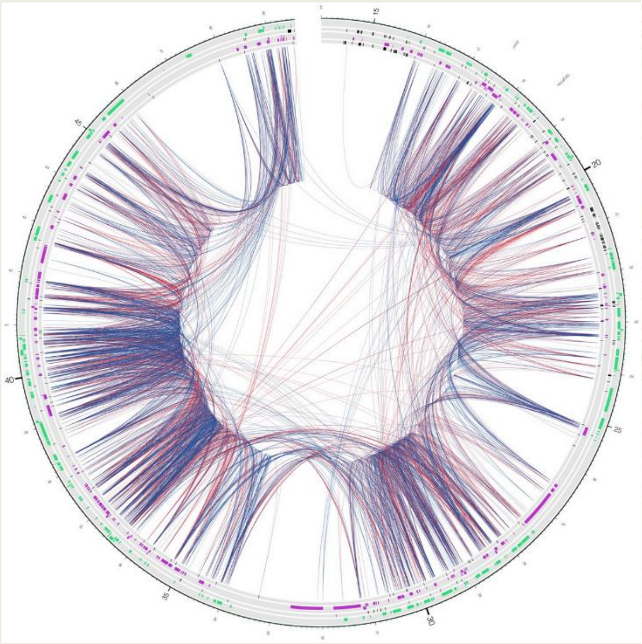

Gene network
Chromosome 22

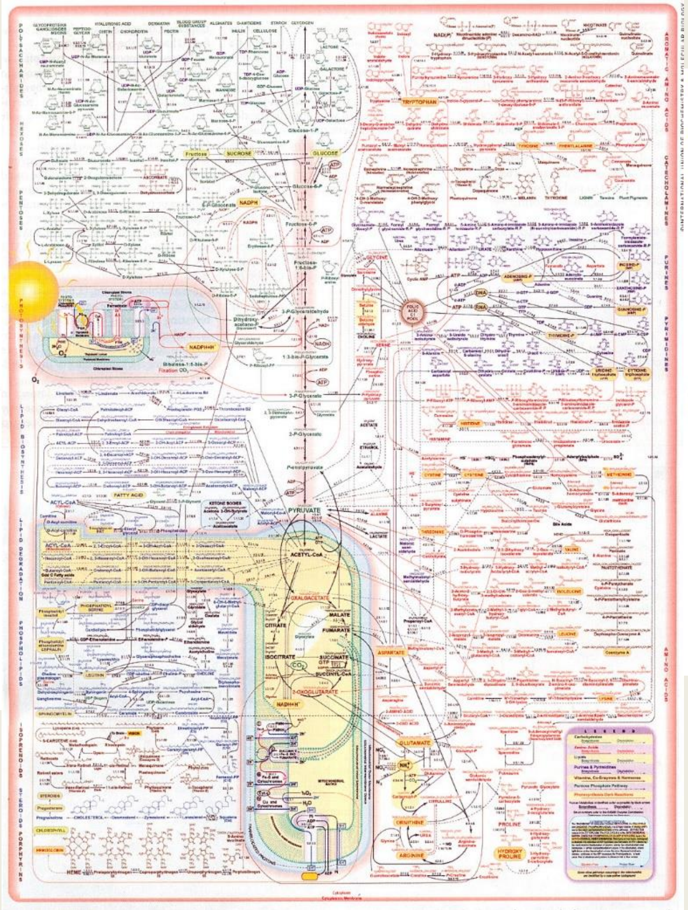

Human Metabolome

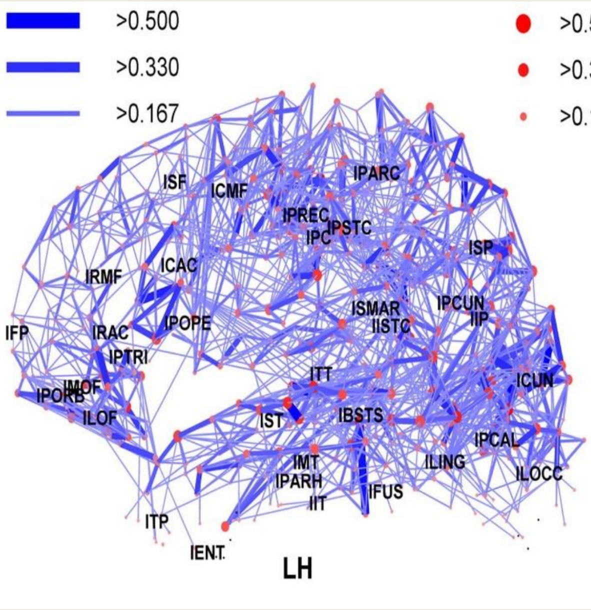

Brain Connectome

https://en.wikipedia.org/wiki/Hub_(network_science_concept)

https://41.media.tumblr.com/tumblr_m5l6rzlqwc1r1171mo1_1280.jpg

http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0028213#pone-0028213-g010 http://karpathy.github.io/2015/05/21/rnn-effectiveness/

Andrej Karpathy blog

About · Hacker's guide to Neural Networks

The Unreasonable Effectiveness of Recurrent Neural Networks

May 21, 2015

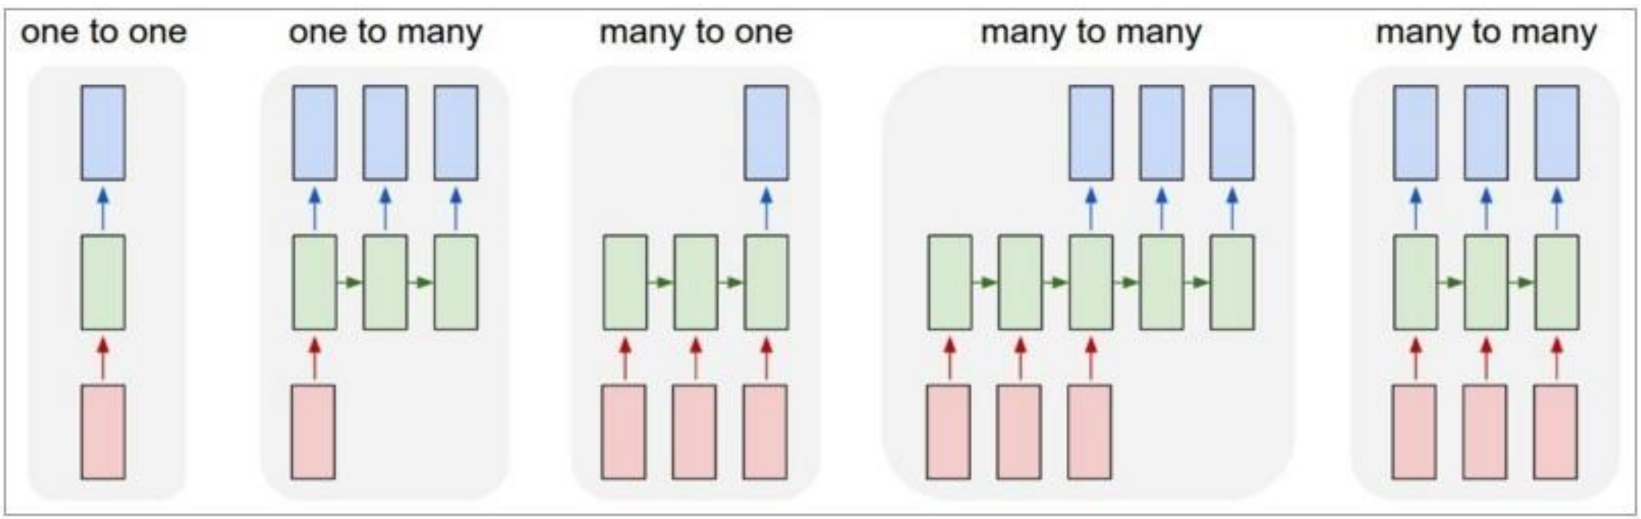

Image Classification

Image Captioning

Sentence Sentiment

English->French Translation

Real-time Video Frame Classification

Recurrent Net Hallucinates C Code

Karpathy: 464MB of C code, 3 layer LSTM, 10 million parameters

http://karpathy.github.io/2015/05/21/rnn-effectiveness/

```
/*
 * Increment the size file of the new incorrect UI_FILTER group information
 * of the size generatively.
 */
static int indicate_policy(void)
{
  int error;
  if (fd == MARN_EPT) {
    /*
     * The kernel blank will coeld it to userspace.
     */
    if (ss->segment < mem_total)
      unblock_graph_and_set_blocked();
    else
      ret = 1;
    goto bail;
  }
  segaddr = in_SB(in.addr);
  selector = seg / 16;
  setup_works = true;
  for (i = 0; i < blocks; i++) {
    seq = buf[i++];
    bpf = bd->bd.next + i * search;
    if (fd) {
      current = blocked;
    }
  }
  rw->name = "Getjbbregs";
  bprm_self_clearl(&iv->version);
  regs->new = blocks[(BPF_STATS << info->historidac)] | PFMR_CLOBATHINC_SECONDS << 12;
  return segtable;
}
```

http://arlingtonva.s3.amazonaws.com/wp-content/uploads/sites/25/2013/12/rat.jpg

The rat escaped.

The rat the cat attacked escaped.

The rat the cat the dog chased attacked escaped.

DeepMind Deep-Q Networks

http://www.nature.com/nature/journal/v518/n7540/full/nature14236.html

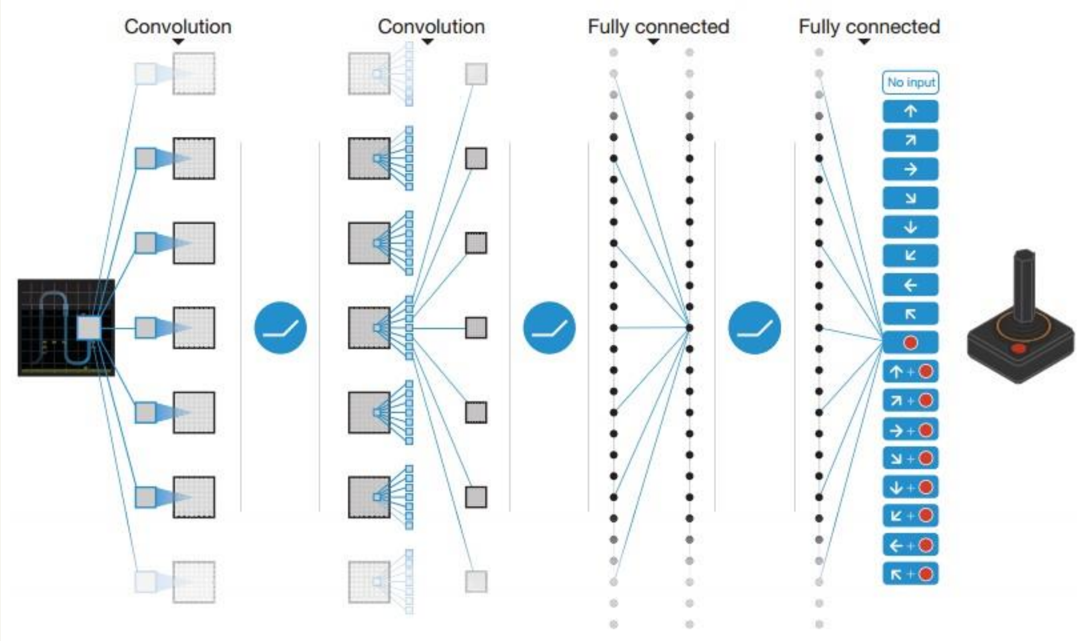

Feb. 2015:
49 Atari 2600 Games
Raw pixels
Same net all games
Beat previous Ais
Beat humans on half

May 2015:
100’s of games

https://www.youtube.com/watch?v=08Cl7ii6viY&feature=youtu.be&t=15m31s

May 2015:
3D games
TORCS racing
Beat Ais from pixels

Aerial Drones: $98 Billion by 2025

http://www.businessinsider.com/the-market-for-commercial-drones-2014-2

Delivery, Surveillance, Agriculture, Military, Police

http://www.flybestdrones.com/best-5-drones-with-camera-under-50-dollars/

http://mint-tek.com/wp-content/uploads/2015/08/commercialdronesforhire.jpg

Deep Learning Has Blindspots

Full Citation: Nguyen A, Yosinski J, Clune J. *Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images*. In Computer Vision and Pattern Recognition (CVPR '15), IEEE, 2015.

Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images

Anh Nguyen
University of Wyoming
anguyen8@uwyo.edu

Jason Yosinski
Cornell University
yosinski@cs.cornell.edu

Jeff Clune
University of Wyoming
jeffclune@uwyo.edu

arXiv:1412.1897v4 [cs.CV] 2 Apr 2015

Abstract

Deep neural networks (DNNs) have recently been achieving state-of-the-art performance on a variety of pattern-recognition tasks, most notably visual classification problems. Given that DNNs are now able to classify objects in images with near-human-level performance, questions naturally arise as to what differences remain between computer and human vision. A recent study [30] revealed that changing an image (e.g. of a lion) in a way imperceptible to humans can cause a DNN to label the image as something else entirely (e.g. mislabeling a lion a library). Here we show a related result: it is easy to produce images that are completely unrecognizable to humans, but that state-of-the-art DNNs believe to be recognizable objects with 99.99% confidence (e.g. labeling with certainty that white noise static is a lion). Specifically, we take convolutional neural networks trained to perform well on either the ImageNet or MNIST datasets and then find images with evolutionary algorithms or gradient ascent that DNNs label with high confidence as belonging to each dataset class. It is possible to produce images totally unrecognizable to human eyes that DNNs believe with near certainty are familiar objects, which we call "fooling images" (more generally, fooling examples). Our results shed light on interesting differences between human vision and current DNNs, and raise questions about the generality of DNN computer vision.

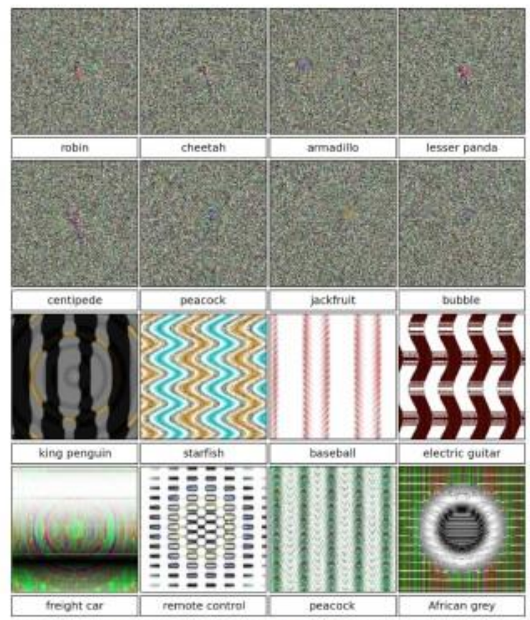

Figure 1. Evolved images that are unrecognizable to humans, but that state-of-the-art DNNs trained on ImageNet believe with $\geq$ 99.6% certainty to be a familiar object. This result highlights differences between how DNNs and humans recognize objects. Images are either directly (*top*) or indirectly (*bottom*) encoded.

http://arxiv.org/abs/1412.1897

Other Issues

- Typically have *problems to solve* rather than reinforcement signals
- Want *confidence* that system solves problem
- Want confidence in *no unintended behaviors*
- Systems often have to obey legal, corporate, or design *constraints*

http://78813809ba6486e732cd-642fac701798512a2848affc62d0ffb0.r60.cf2.rackcdn.com/465DAB1D-1F8E-4164-8D18-3BFD150E02F4.jpg

Rational Decision Making

http://commons.wikimedia.org/wiki/File:John_von_Neumann.jpg

1. *Have utility function*
2. *Have a model of the world*
3. *Choose the action with highest expected utility*
4. *Update the model based on what happens*

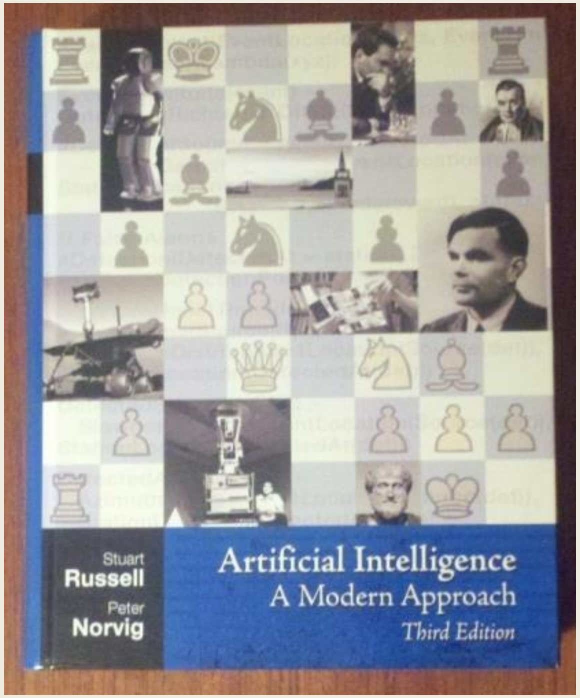

http://aima.cs.berkeley.edu/

- Von Neumann and Morgenstern, 1944
- Savage, 1954
- Anscombe and Aumann, 1963

Modern Approach to AI

Fully Rational Systems

Sensations

$S_1, \ldots, S_N$

Actions

$A_1, \ldots, A_N$

Utility function: $U(S_1, \ldots, S_N)$ Prior Probability: $P(S_1, \ldots, S_N \mid A_1, \ldots, A_N)$

Rational Action at time t:

$$A_t^R(S_1, A_1, \ldots, A_{t-1}, S_t) =$$

$$\underset{A_t^R}{\operatorname{argmax}} \sum_{S_{t+1}, \ldots, S_N} U(S_1, \ldots, S_N) P(S_1, \ldots, S_N \mid A_1, \ldots, A_{t-1}, A_t^R, \ldots, A_N^R)$$

The Formula for Intelligence!

It includes Bayesian Inference, Search, and Deliberation.

But it requires $O(NS^N A^N)$ computational steps.

Approximately Rational Architectures

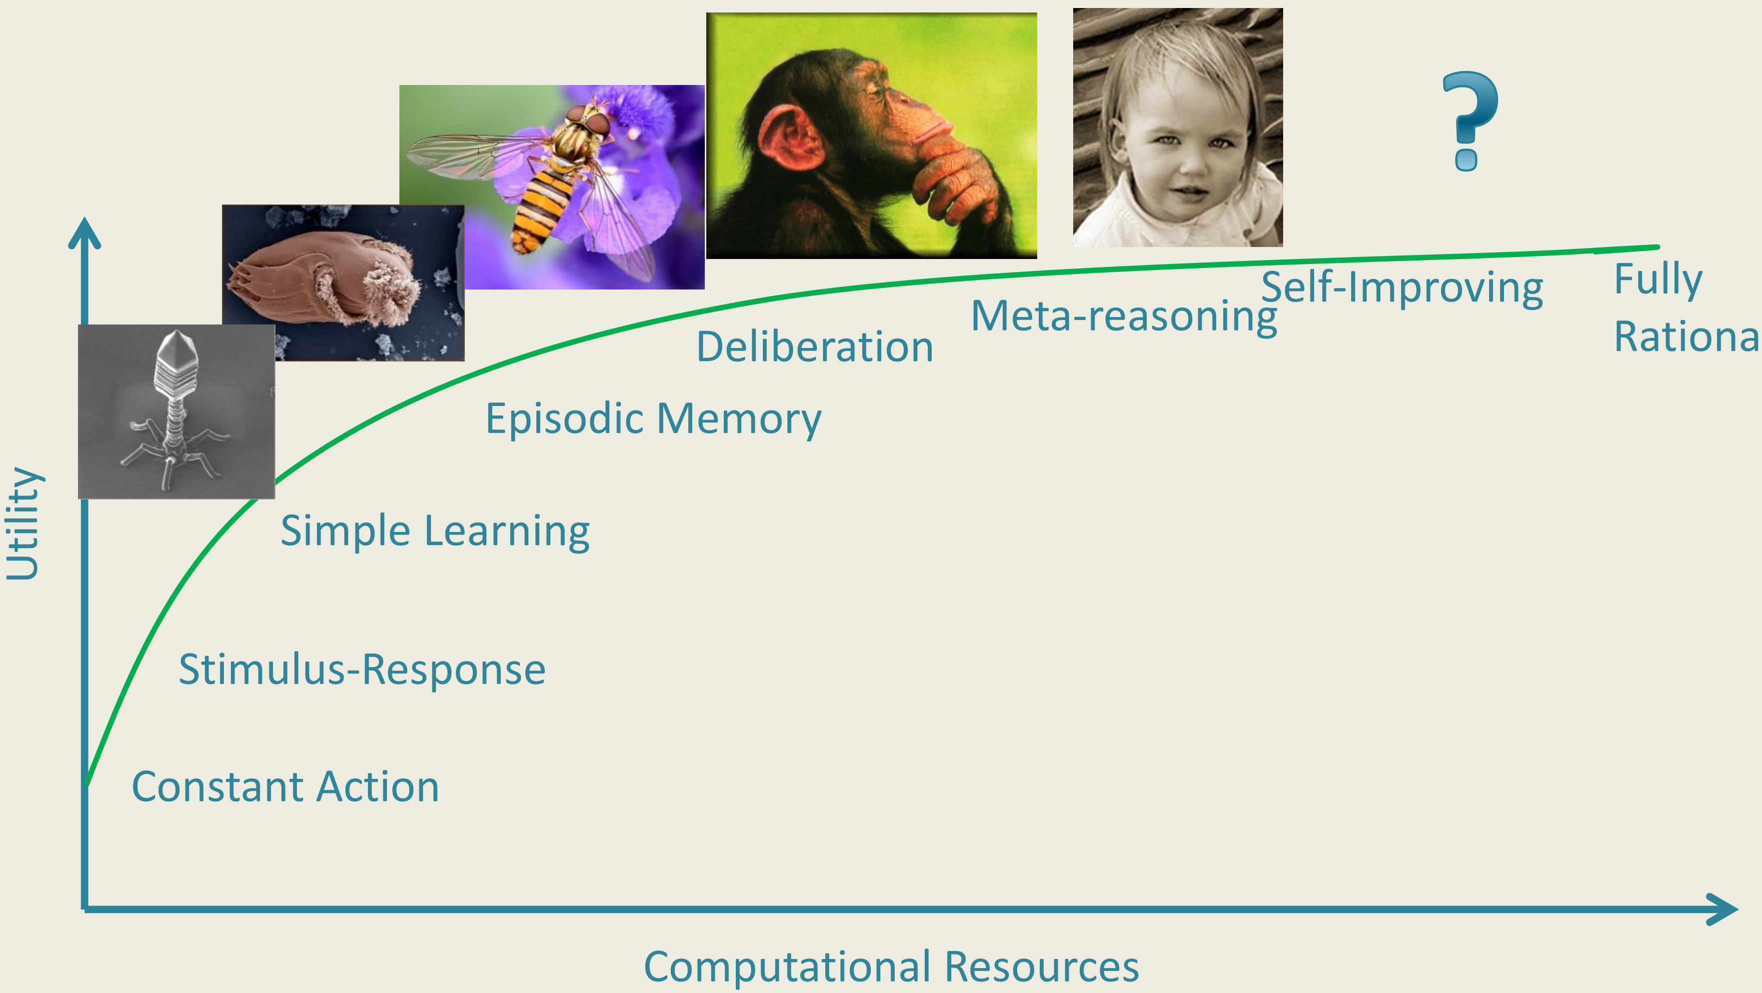

Technology Needs Semantics!

- Analyzing camera, sensor, weather data
- Better search, question answering, info
- Analysis and optimization of business processes
- Health monitoring, medical diagnosis
- Financial markets trading, stabilization
- Autonomous cars, trucks, boats, subs, planes
- Pollution monitoring and cleanup
- Improved robotic manufacturing
- Software and Hardware design

Approaches to Semantics

Representation, Encoding, Learning, Communication, Reasoning

- Montague – map into Typed Lambda Calculus
- Denotational – map into CS Domains
- Mathematical – map into Set Theory
- Categorical – map into Category Theory
- Distributional – Statistics of Contexts

http://engineering.missouri.edu/wp-content/uploads/australian-cloudy-sky.jpg

http://flinttown.com/wp-content/uploads/2015/07/Fireworks.jpg