See eye to eye

Ricardo Marroquim

www.lcg.ufrj.br/marroquim

Laboratório de Computação Gráfica

PESC
Programa de Engenharia de Sistemas e Computação

how do we see?

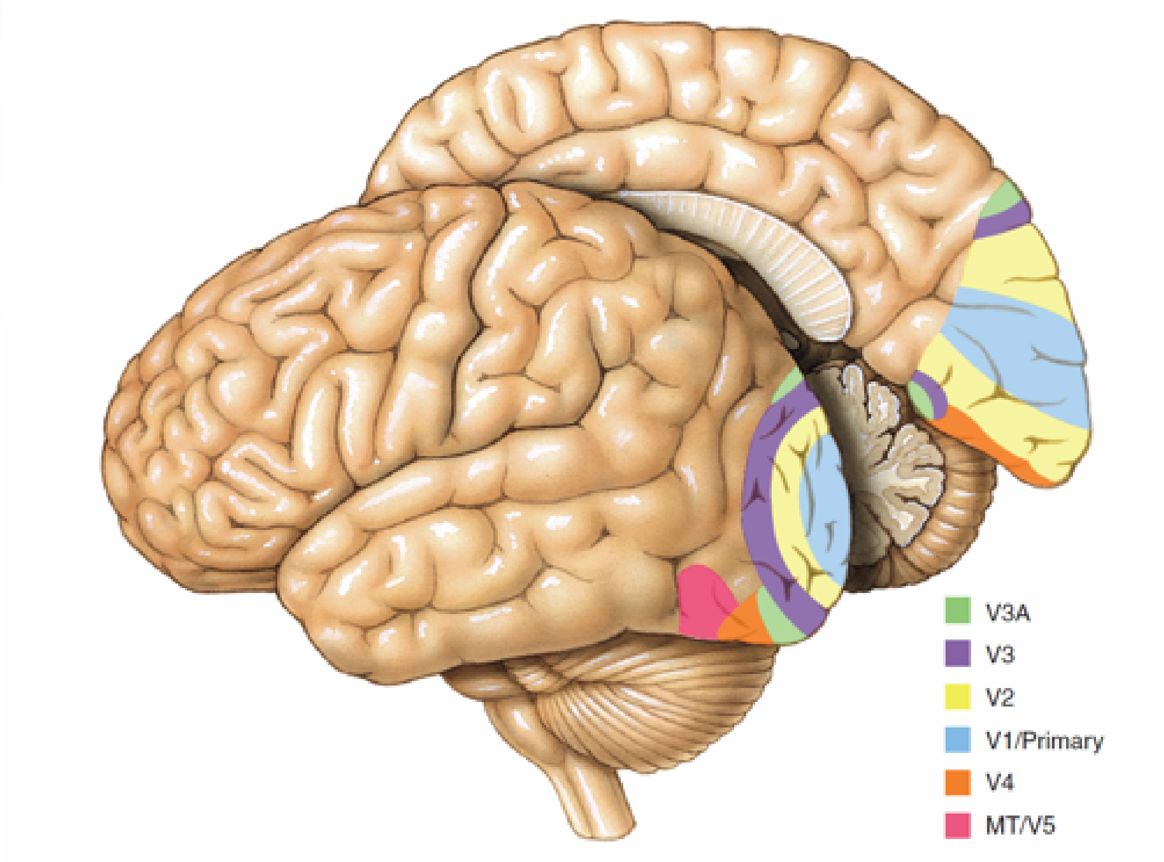

how computers see?

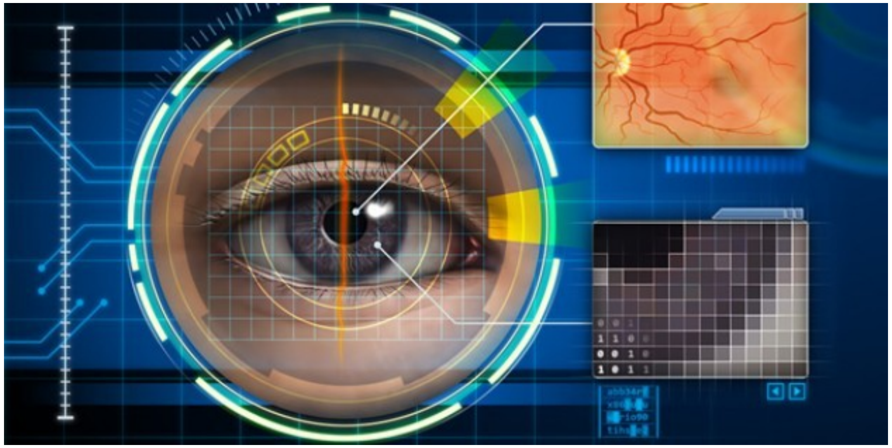

computer vision

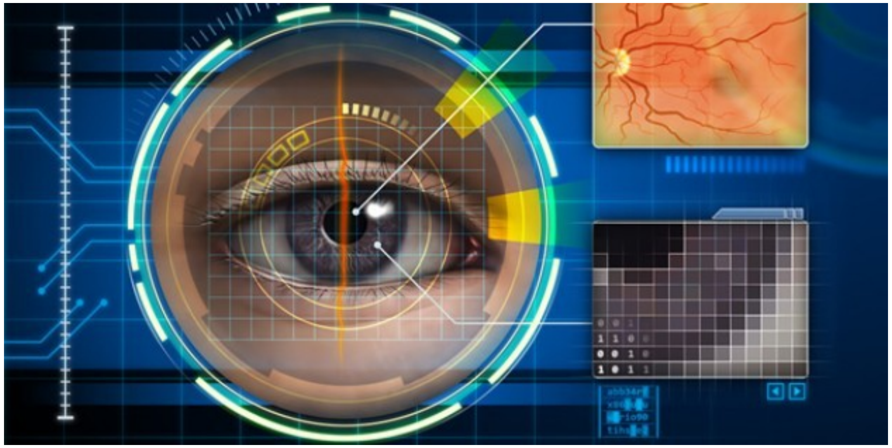

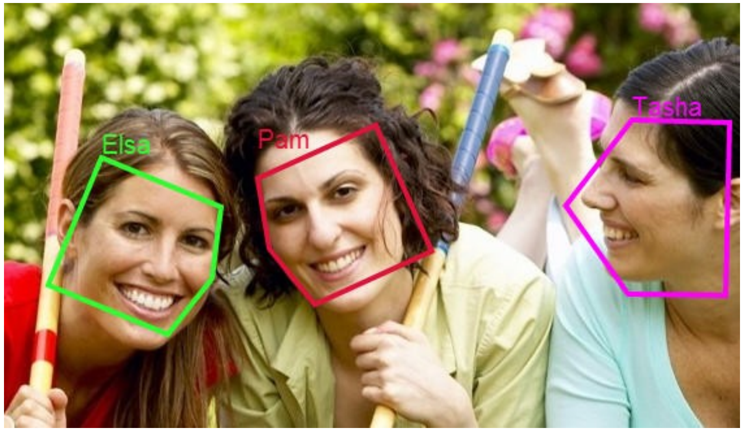

computer vision

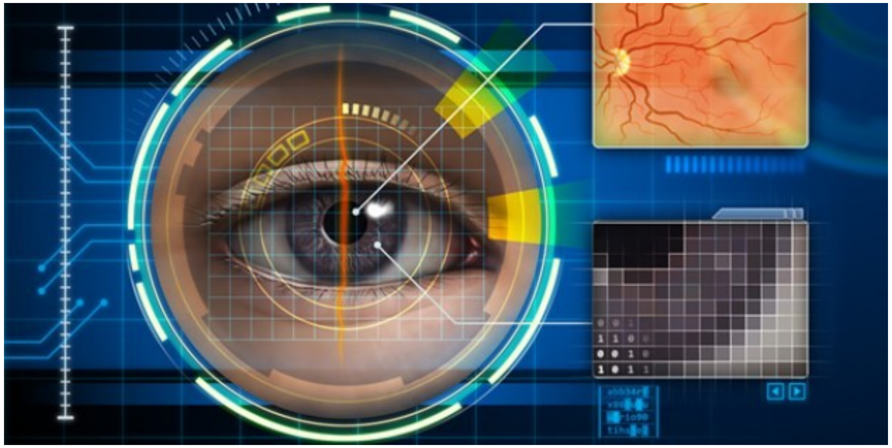

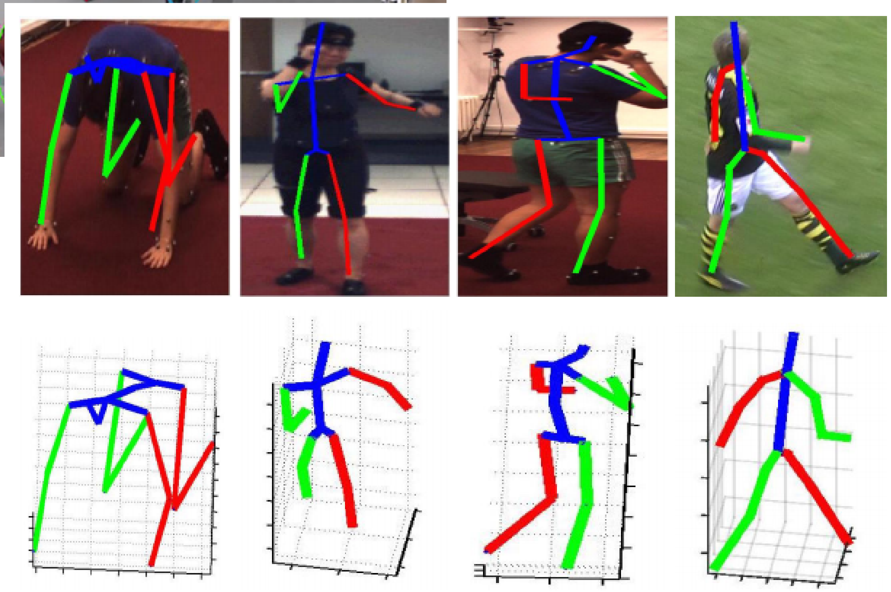

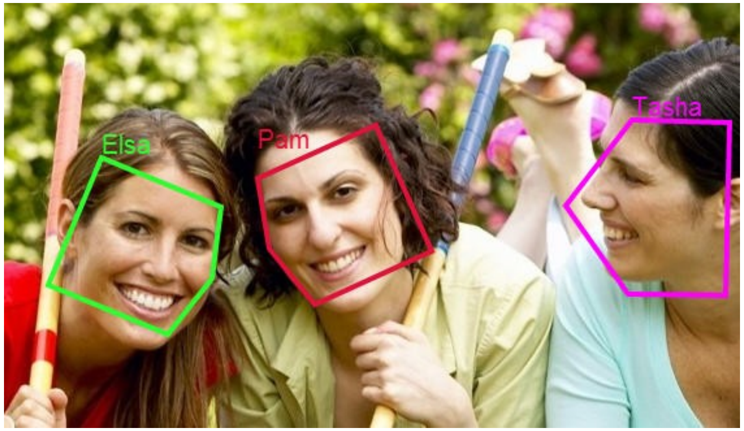

computer vision

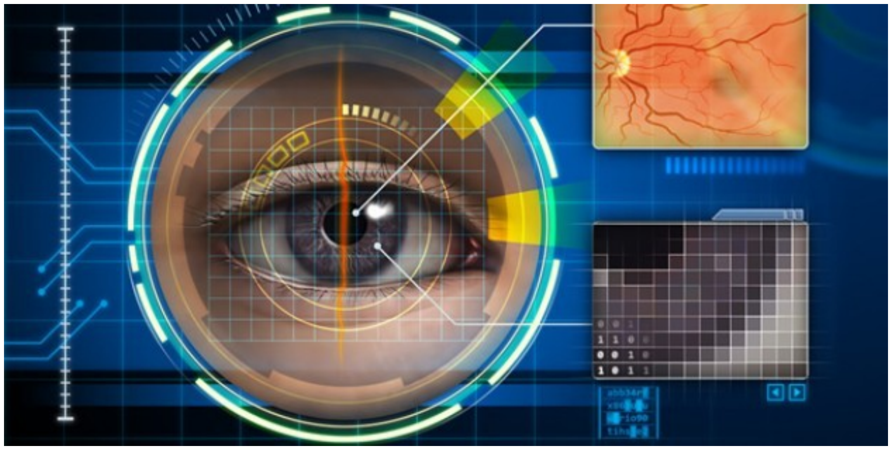

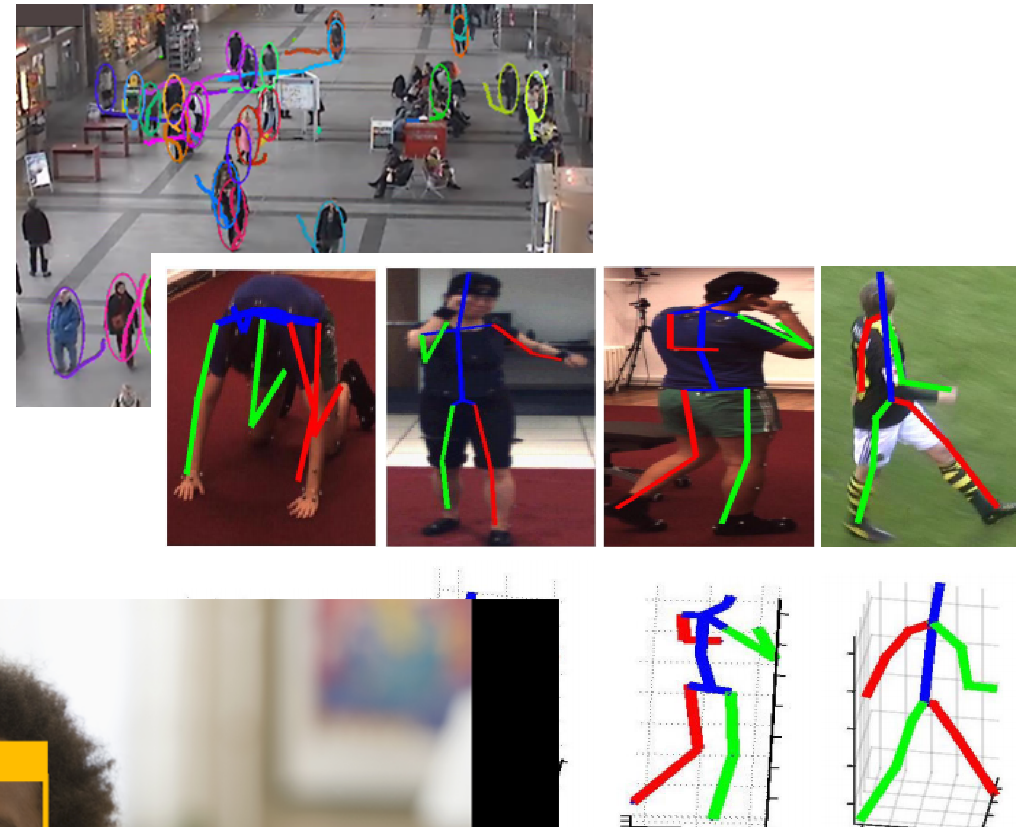

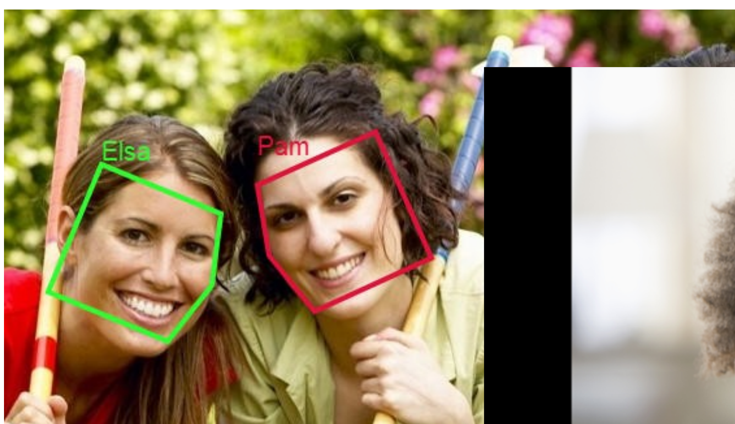

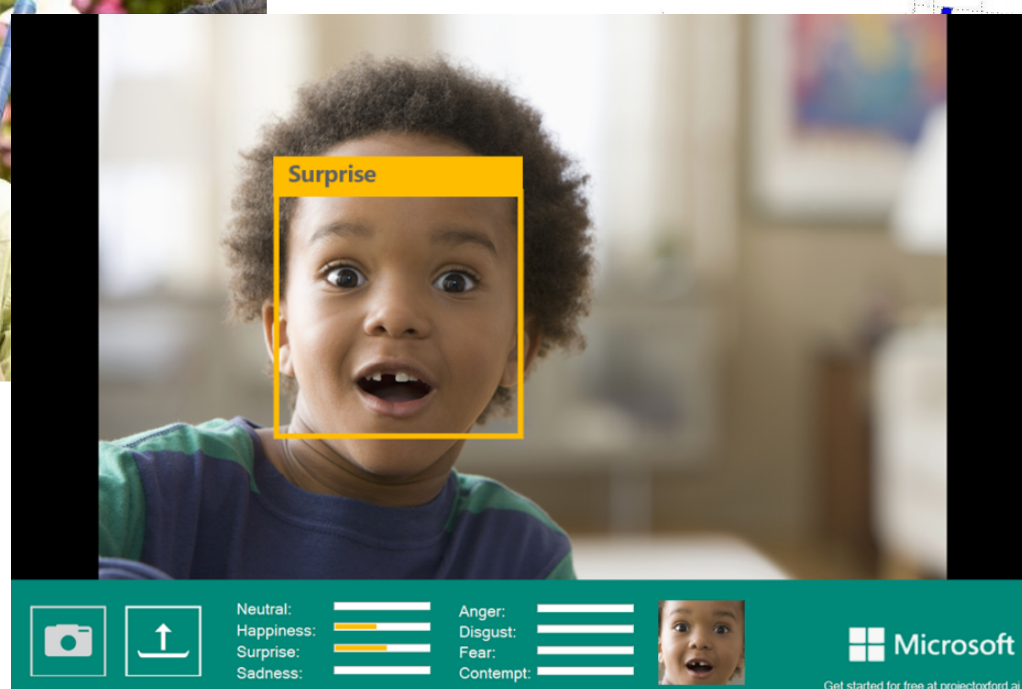

computer vision

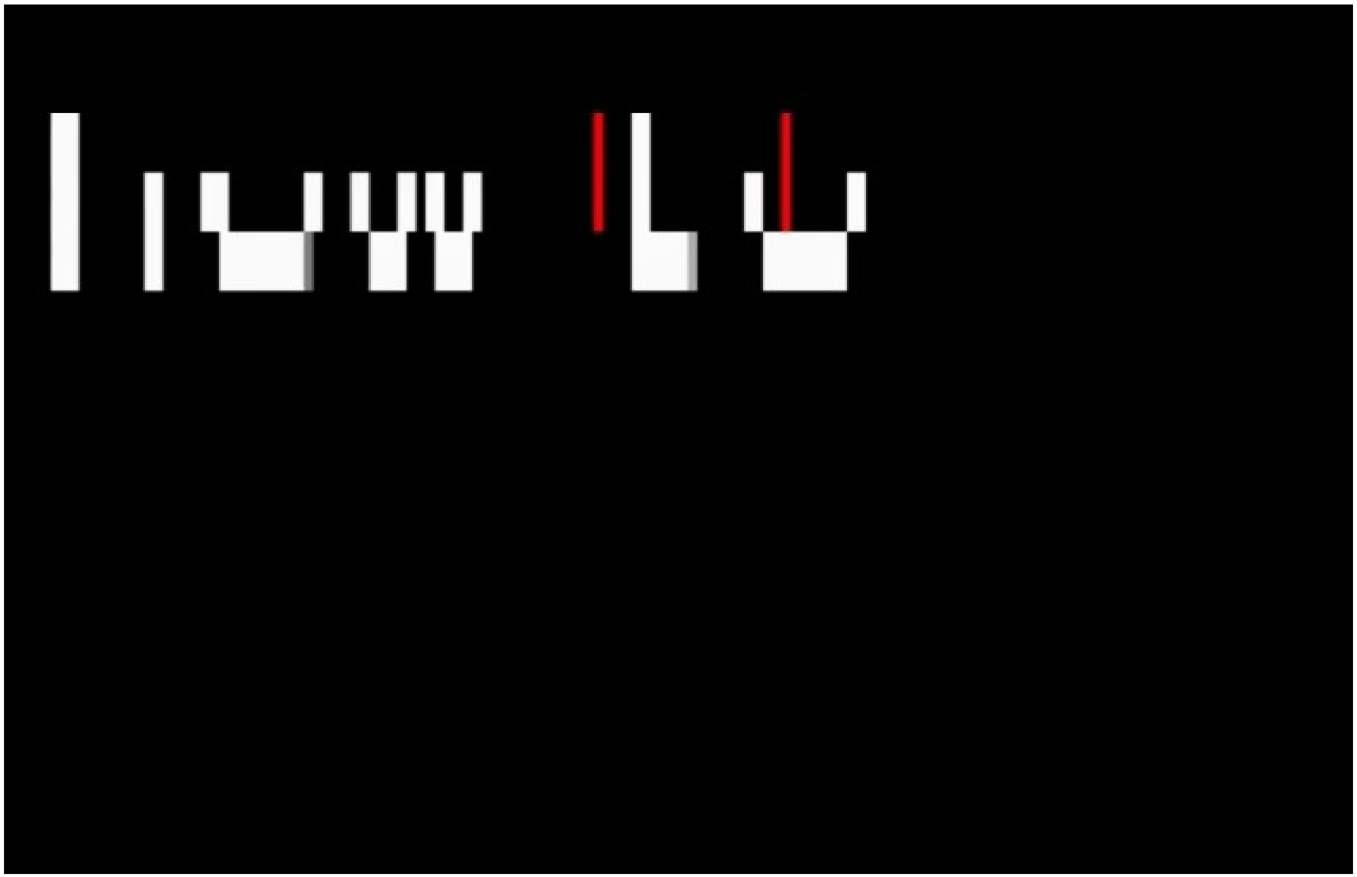

https://www.youtube.com/watch?v=eQLcDmfmGB0

humans vs computers

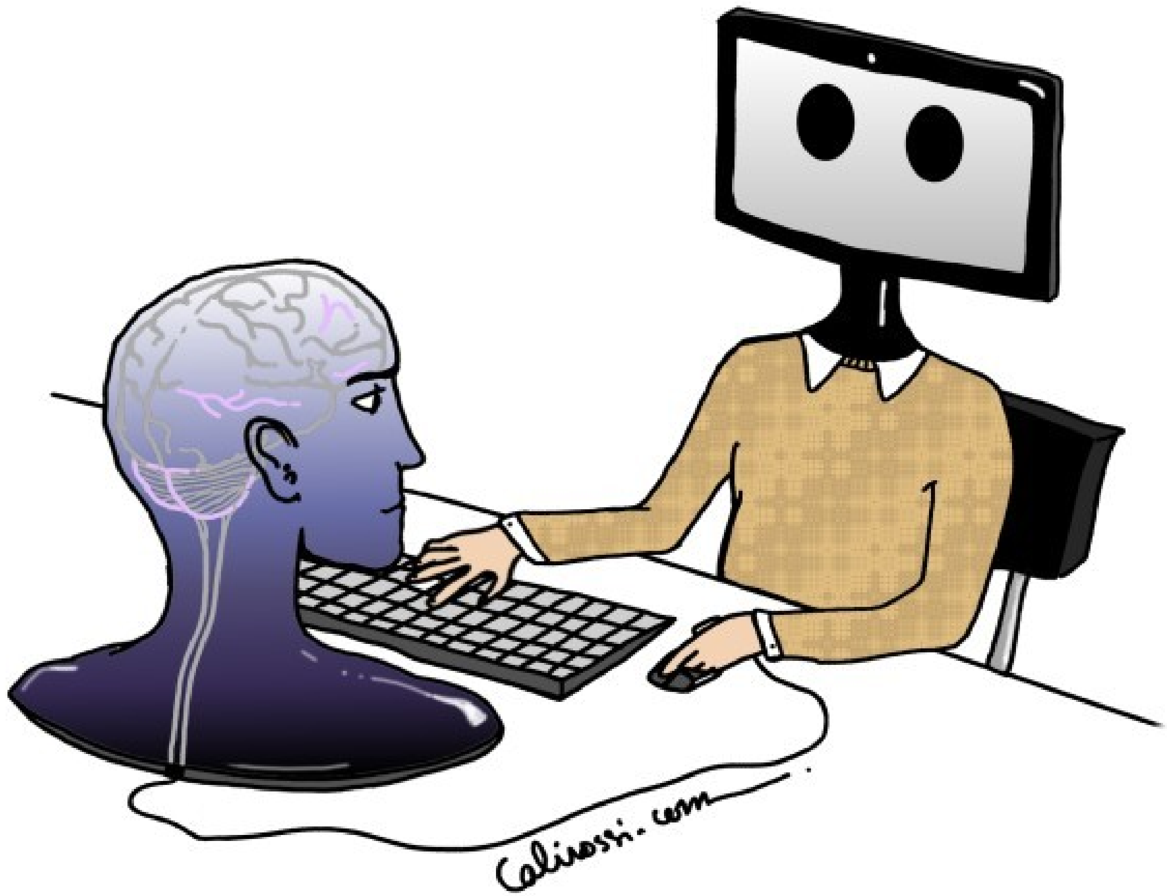

http://benedicterossi.com/

Marvin Minsky

- pioneer: Perceptrons, Logo turtle, Head-mounted display …
- 1969: Turing Award

IN CS, IT CAN BE HARD TO EXPLAIN THE DIFFERENCE BETWEEN THE EASY AND THE VIRTUALLY IMPOSSIBLE.

Larry Roberts

- 1963 - PhD Thesis: Machine Perception of Three-Dimensional Solids

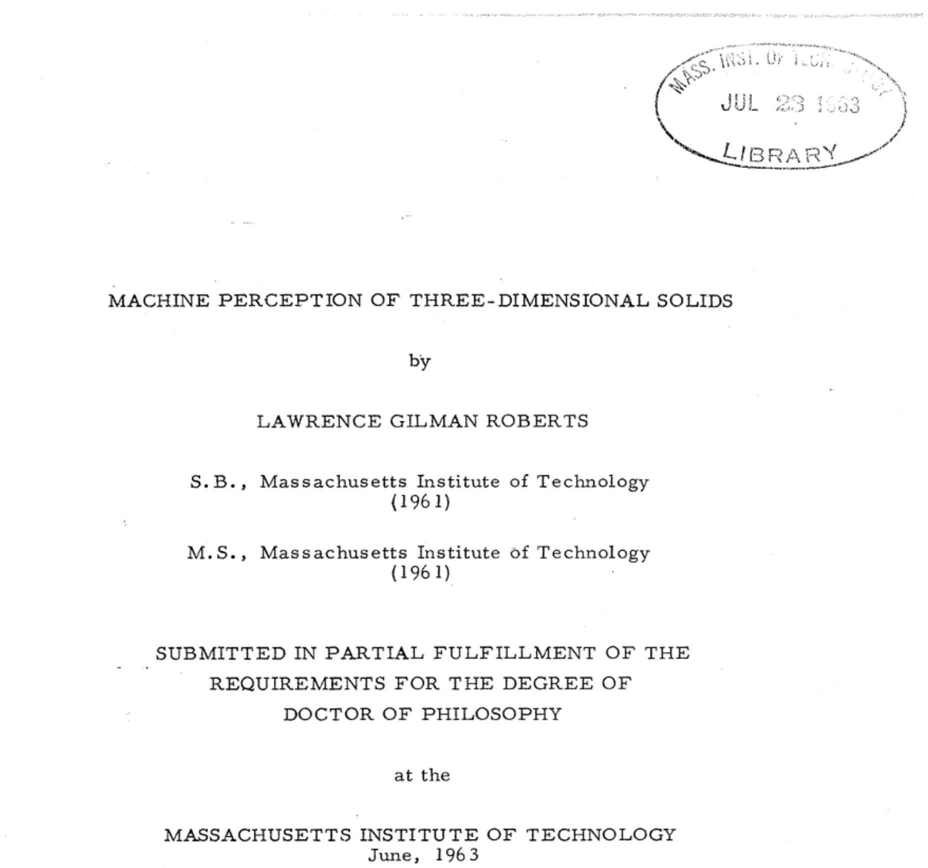

MACHINE PERCEPTION OF THREE-DIMENSIONAL SOLIDS

by

LAWRENCE GILMAN ROBERTS

S.B., Massachusetts Institute of Technology
(1961)

M.S., Massachusetts Institute of Technology
(1961)

SUBMITTED IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
June, 1963

David Marr

- 1982 - David Marr - Vision: A Computational Investigation into the Human Representation and Processing of Visual Information

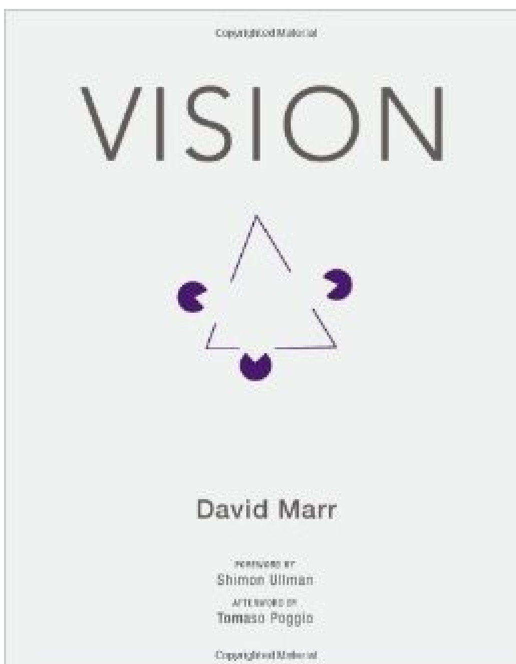

computer vision

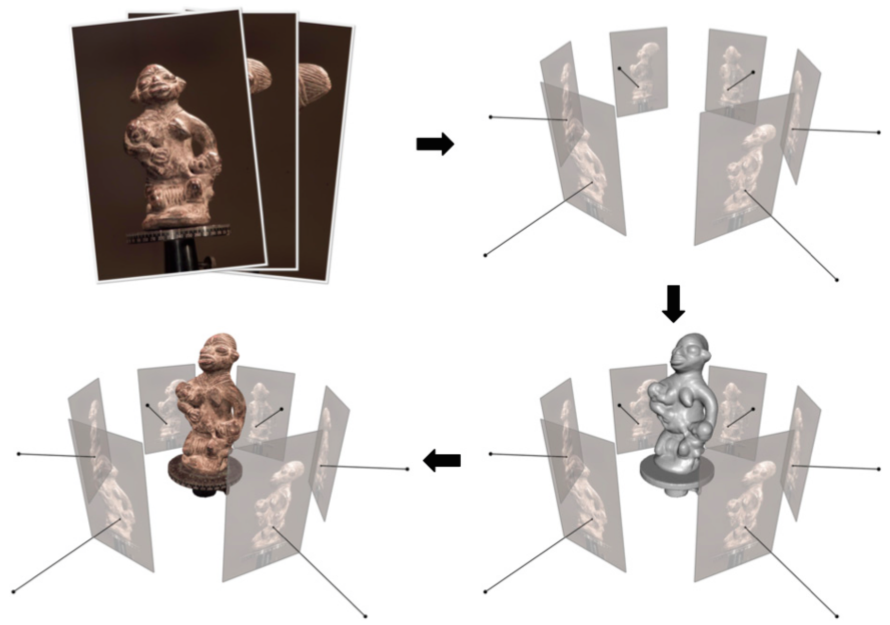

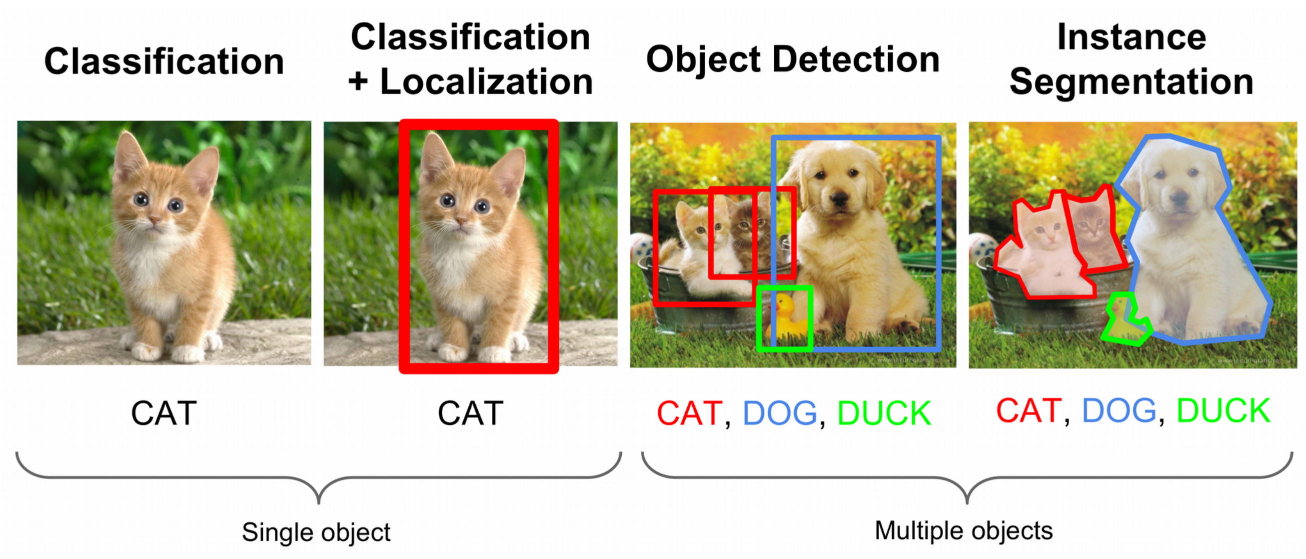

projective geometry

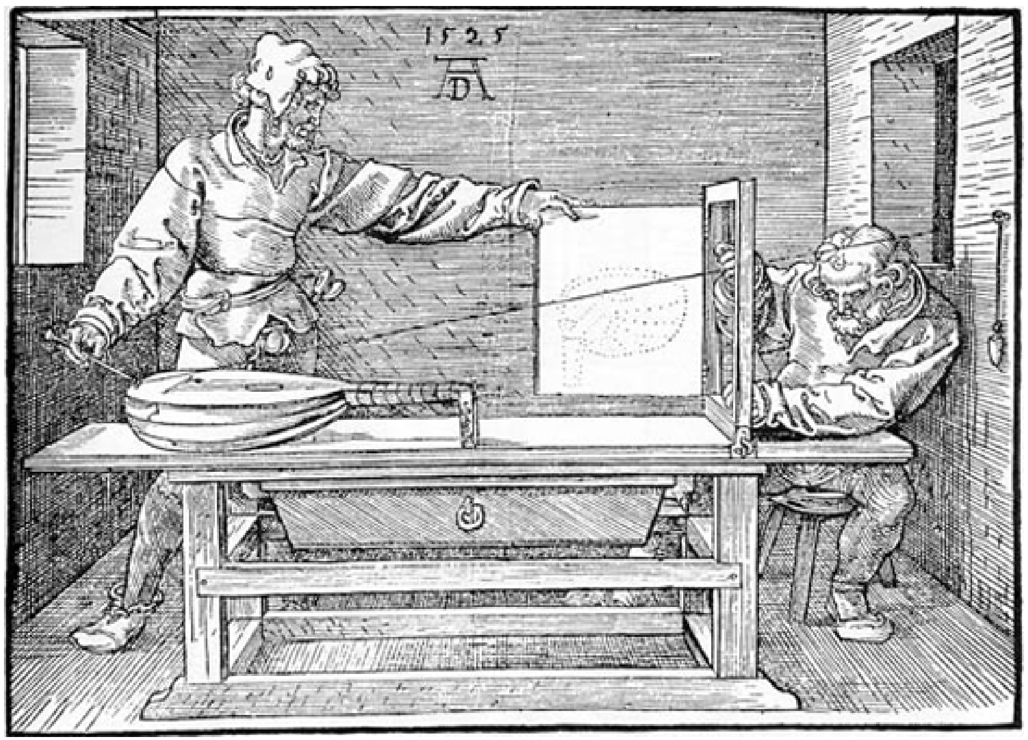

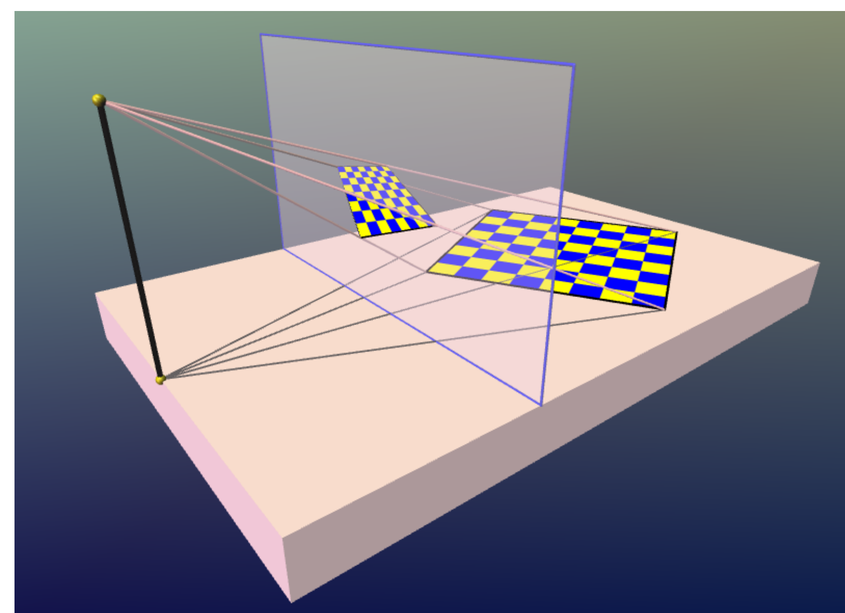

photo pop-up

http://dhoiem.cs.illinois.edu/projects/popup/

projective geometry

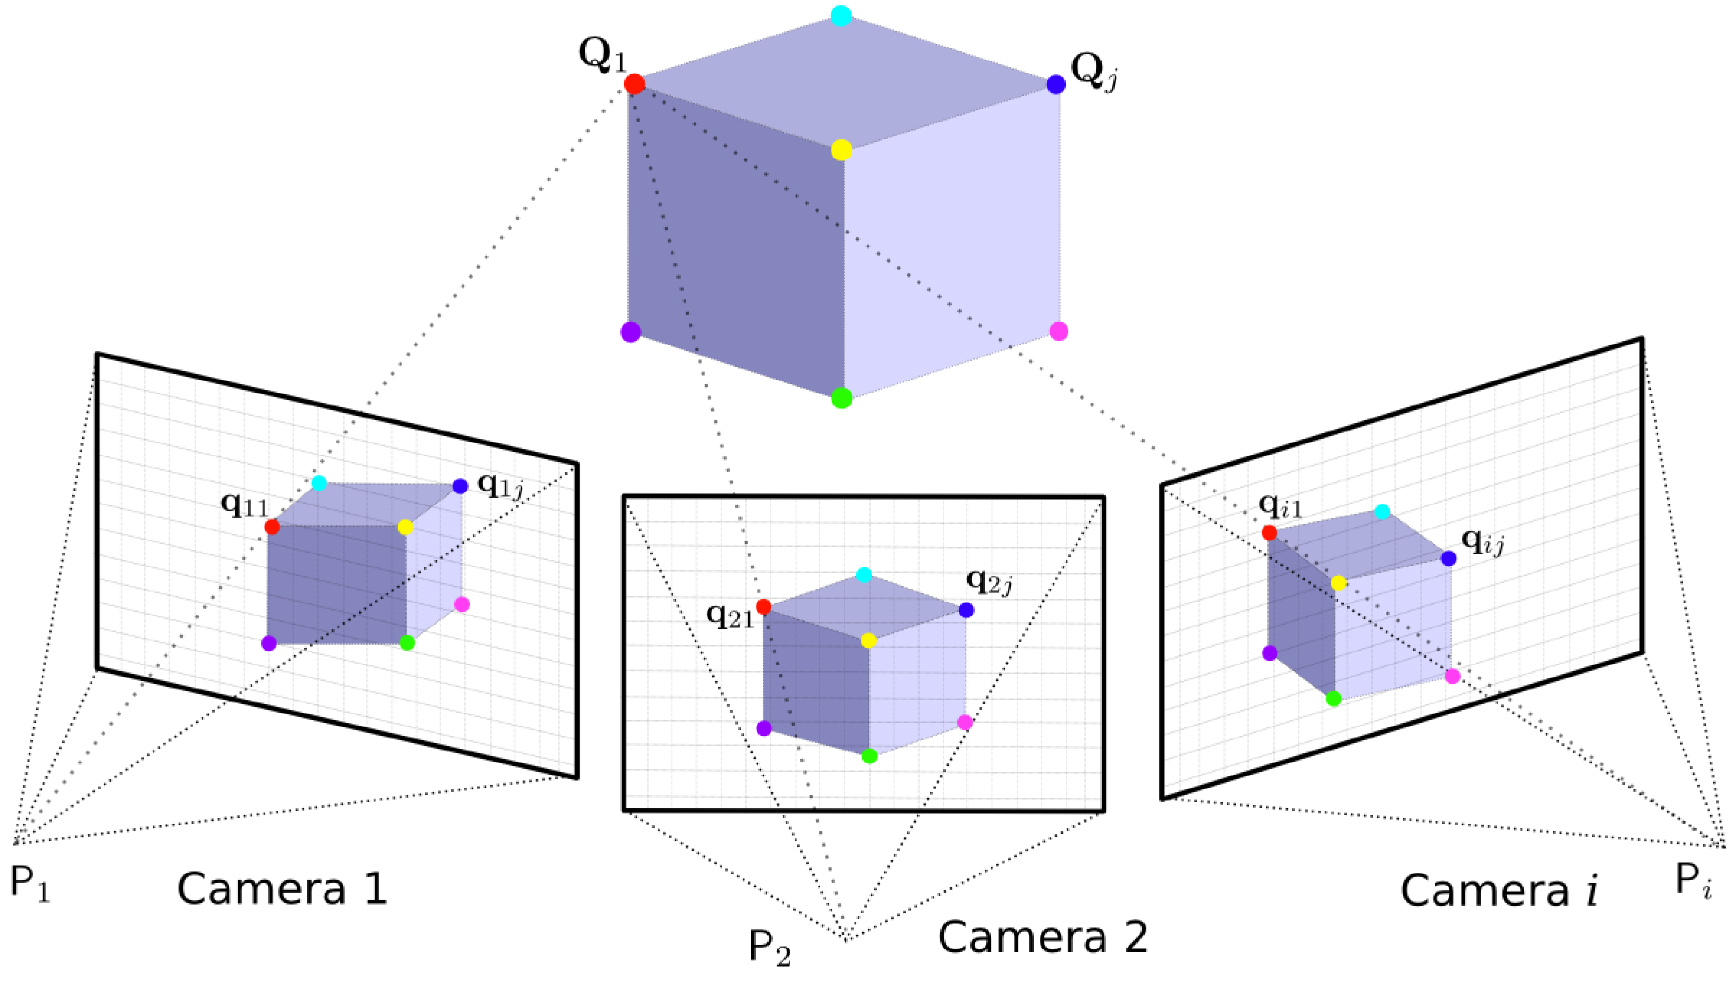

3D reconstruction

http://www.3dflow.net/

understanding

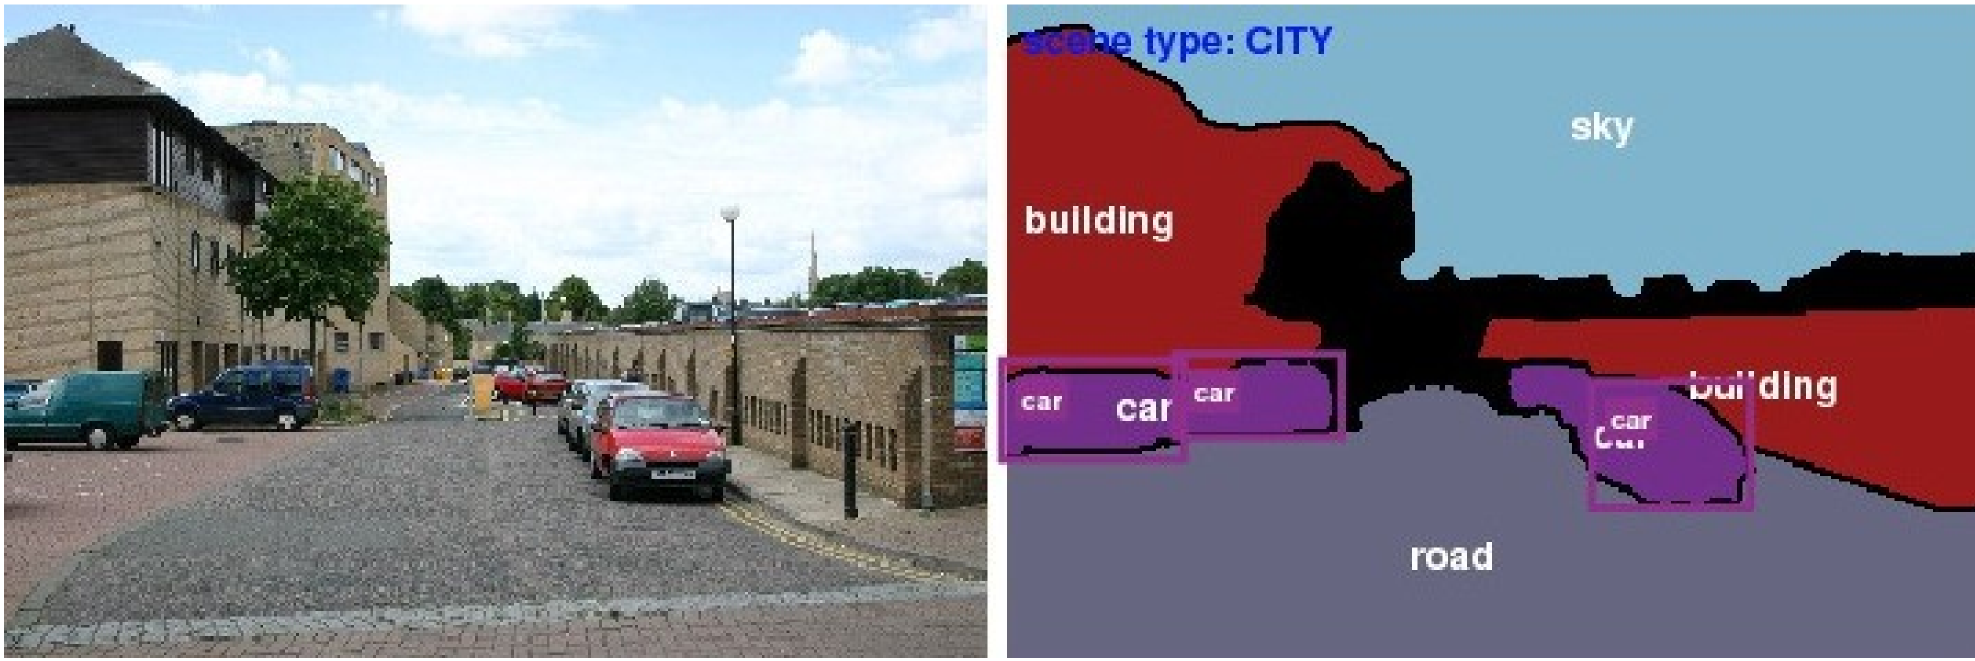

understanding

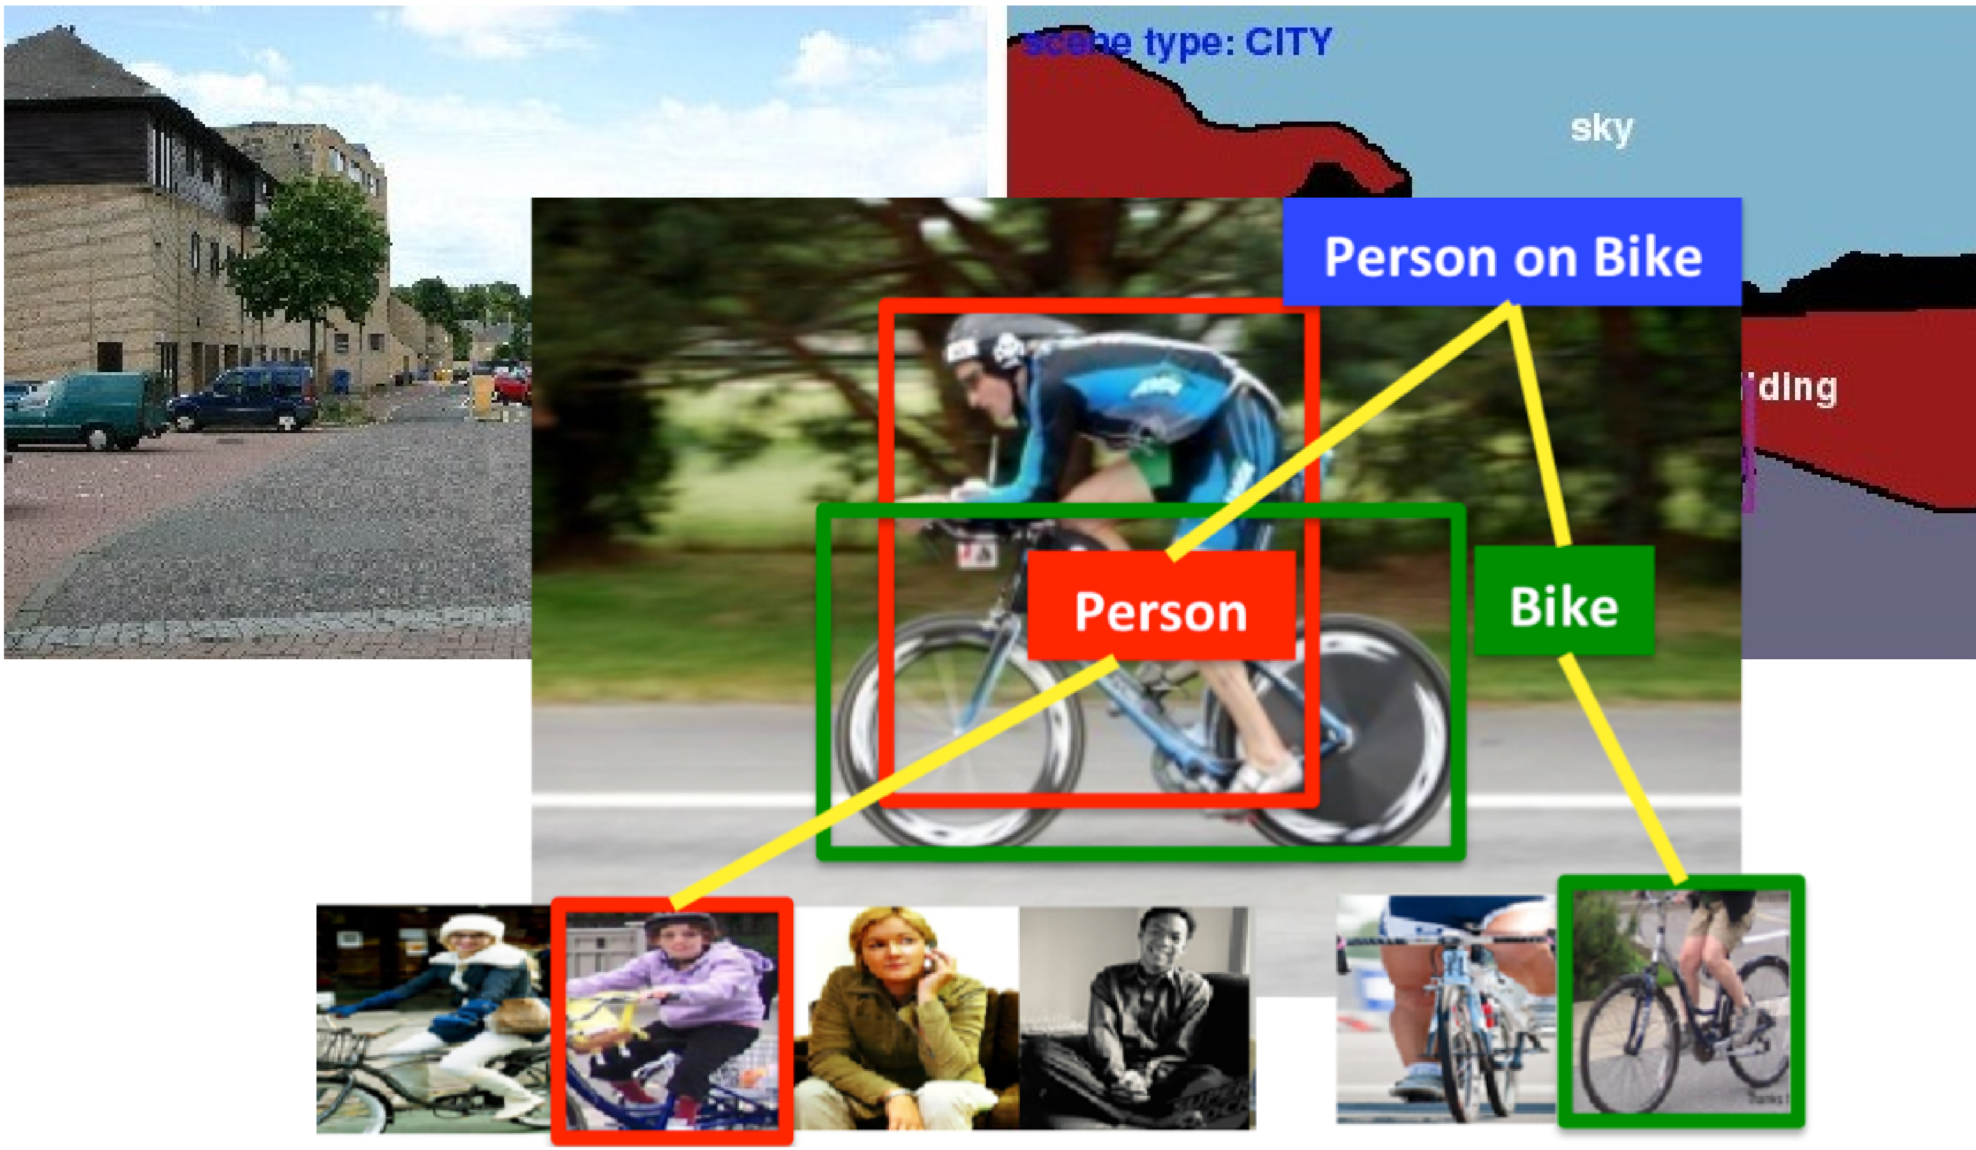

understanding

OPEN ACCESS Freely available online PLoS COMPUTATIONAL BIOLOGY

Why is Real-World Visual Object Recognition Hard?

Nicolas Pinto[1,2], David D. Cox[1,2,3], James J. DiCarlo[1,2*]

1 McGovern Institute for Brain Research, Massachusetts Institute of Technology, Cambridge, Massachusetts, United States
Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts, United States of America, 3 The Rowland Institu
States of America

**Progress in understanding the brain mechanisms underlying vision requires the cons
that not only emulate the brain's anatomy and physiology, but ultimately match it
recent years, "natural" images have become popular in the study of vision and ha
impressive progress in building such models. Here, we challenge the use of uncontr
that progress. In particular, we show that a simple V1-like model—a neuroscient
perform poorly at real-world visual object recognition tasks—outperforms state-of-t
(biologically inspired and otherwise) on a standard, ostensibly natural image recog
designed a "simpler" recognition test to better span the real-world variation in objec
show that this test correctly exposes the inadequacy of the V1-like model. Taken to
that tests based on uncontrolled natural images can be seriously misleading, potenti
direction. Instead, we reexamine what it means for images to be natural and argue
problem of object recognition—real-world image variation.**

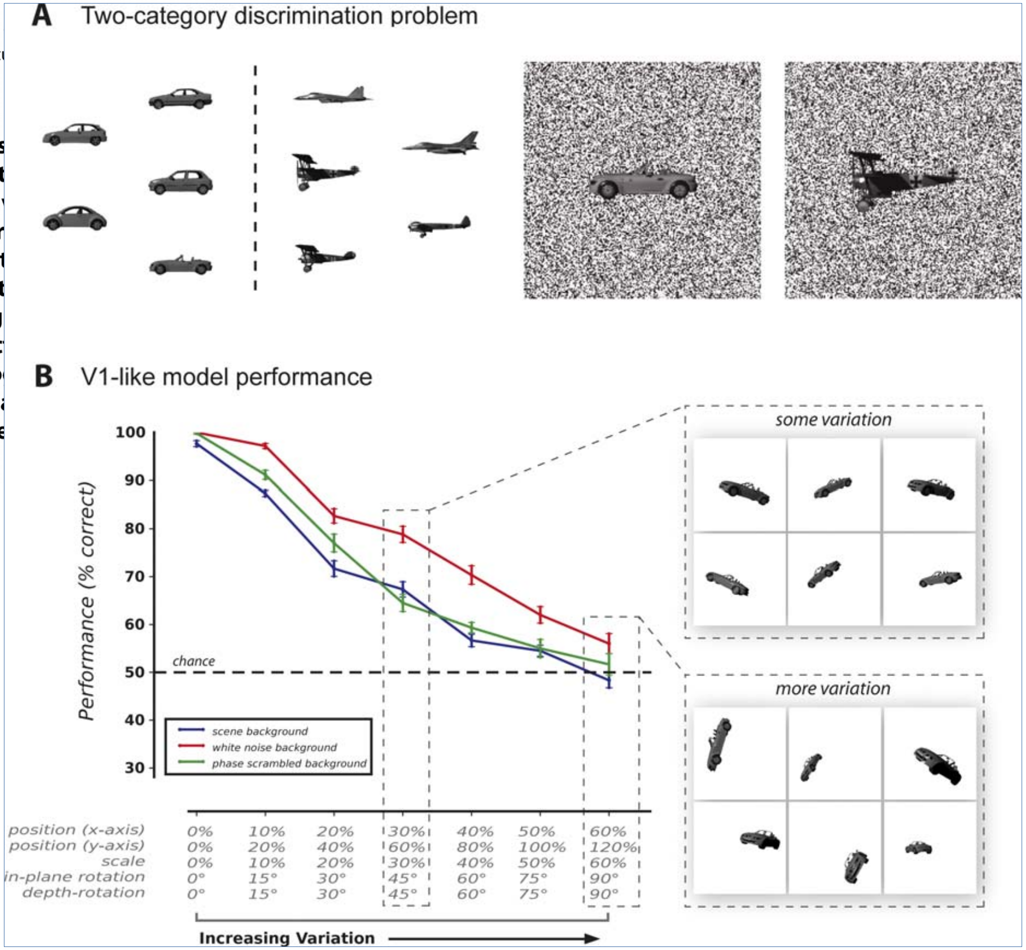

Fei Fei Li

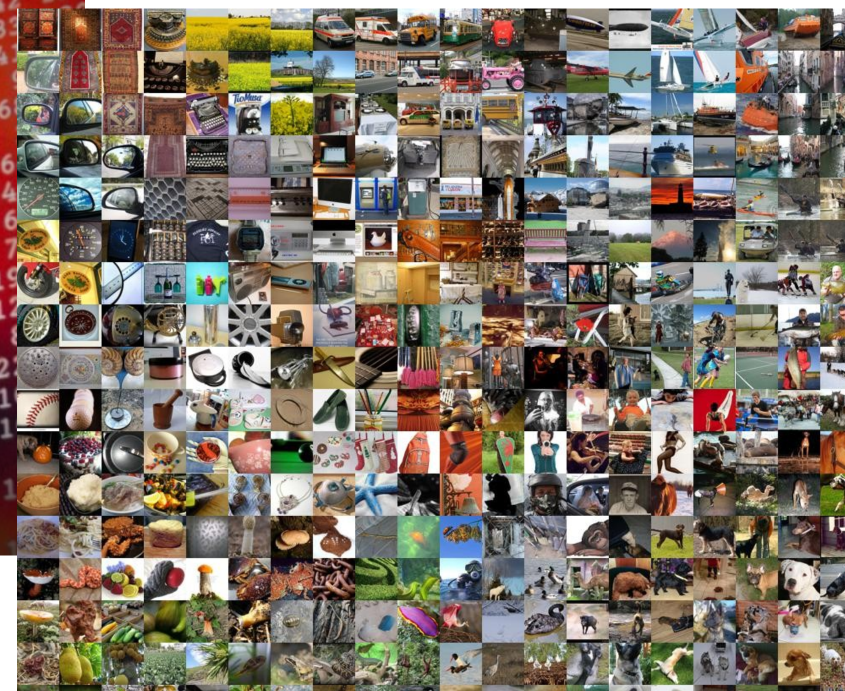

https://www.ted.com/talks/fei_fei_li_how_we_re_teaching_computers_to_understand_pictures

deep learning

ImageNet Classification with Deep Convolutional Neural Networks

Alex Krizhevsky
University of Toronto
kriz@cs.utoronto.ca

Ilya Sutskever
University of Toronto
ilya@cs.utoronto.ca

Geoffrey E. Hinton
University of Toronto
hinton@cs.utoronto.ca

Abstra

We trained a large, deep convolutional neu
high-resolution images in the ImageNet LS
ferent classes. On the test data, we achieve
and 17.0% which is considerably better th
neural network, which has 60 million para
of five convolutional layers, some of which
and three fully-connected layers with a fin
ing faster, we used non-saturating neurons
tation of the convolution operation. To red
layers we employed a recently-developed r
that proved to be very effective. We also
ILSVRC-2012 competition and achieved a
compared to 26.2% achieved by the second

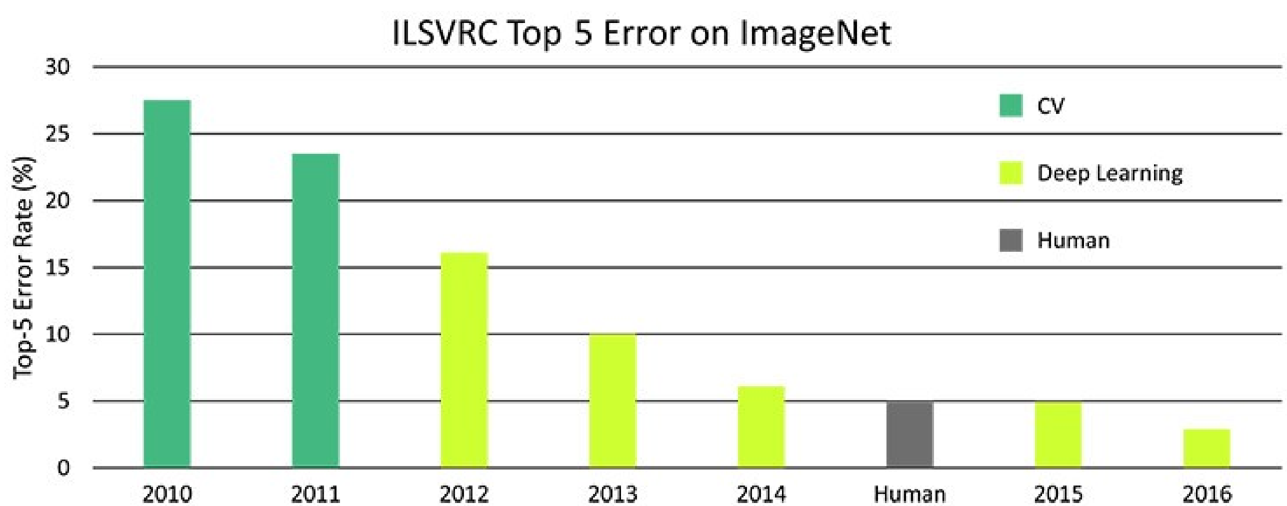

deep learning

DenseCap: Fully Convolutional Localization Networks for Dense Captioning

Justin Johnson* Andrej Karpathy* Li Fei-Fei
Department of Computer Science, Stanford University
{jcjohns,karpathy,feifeili}@cs.stanford.edu

Abstract

We introduce the dense captioning task, which requires a computer vision system to both localize and describe salient regions in images in natural language. The dense captioning task generalizes object detection when the descriptions consist of a single word, and Image Captioning when one predicted region covers the full image. To address the localization and description task jointly we propose a Fully Convolutional Localization Network (FCLN) architecture that processes an image with a single, efficient forward pass, requires no external regions proposals, and can be trained end-to-end with a single round of optimization. The architecture is composed of a Convolutional Network, a novel

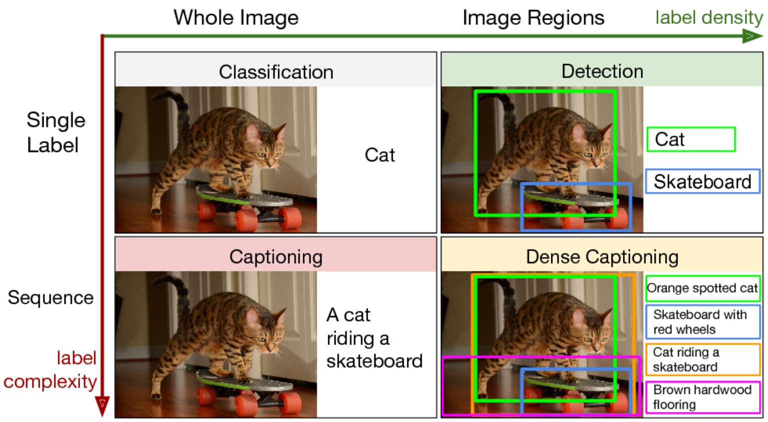

Figure 1. We address the Dense Captioning task (bottom right) with a model that jointly generates both dense and rich annotations in a single forward pass.

deep learning

DenseCap: Fully Convolutional Localization Networks for Dense Captioning

Justin Johnson* Andrej Karpathy* Li Fei-Fei
Department of Computer Science, Stanford University
{jcjohns,karpathy,feifeili}@cs.stanford.edu

Abstract

We introduce the dense captioning task, which requires a computer vision system to both localize and describe salient regions in images in natural language. The dense captioning task generalizes object detection when the descriptions consist of a single word, and Image Captioning when one predicted region covers the full image. To address the local-

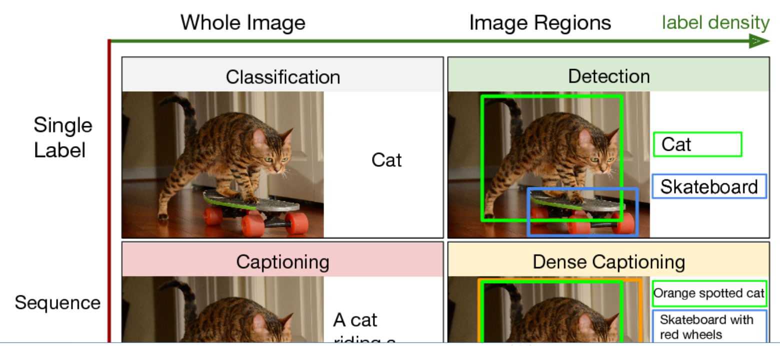

deep learning

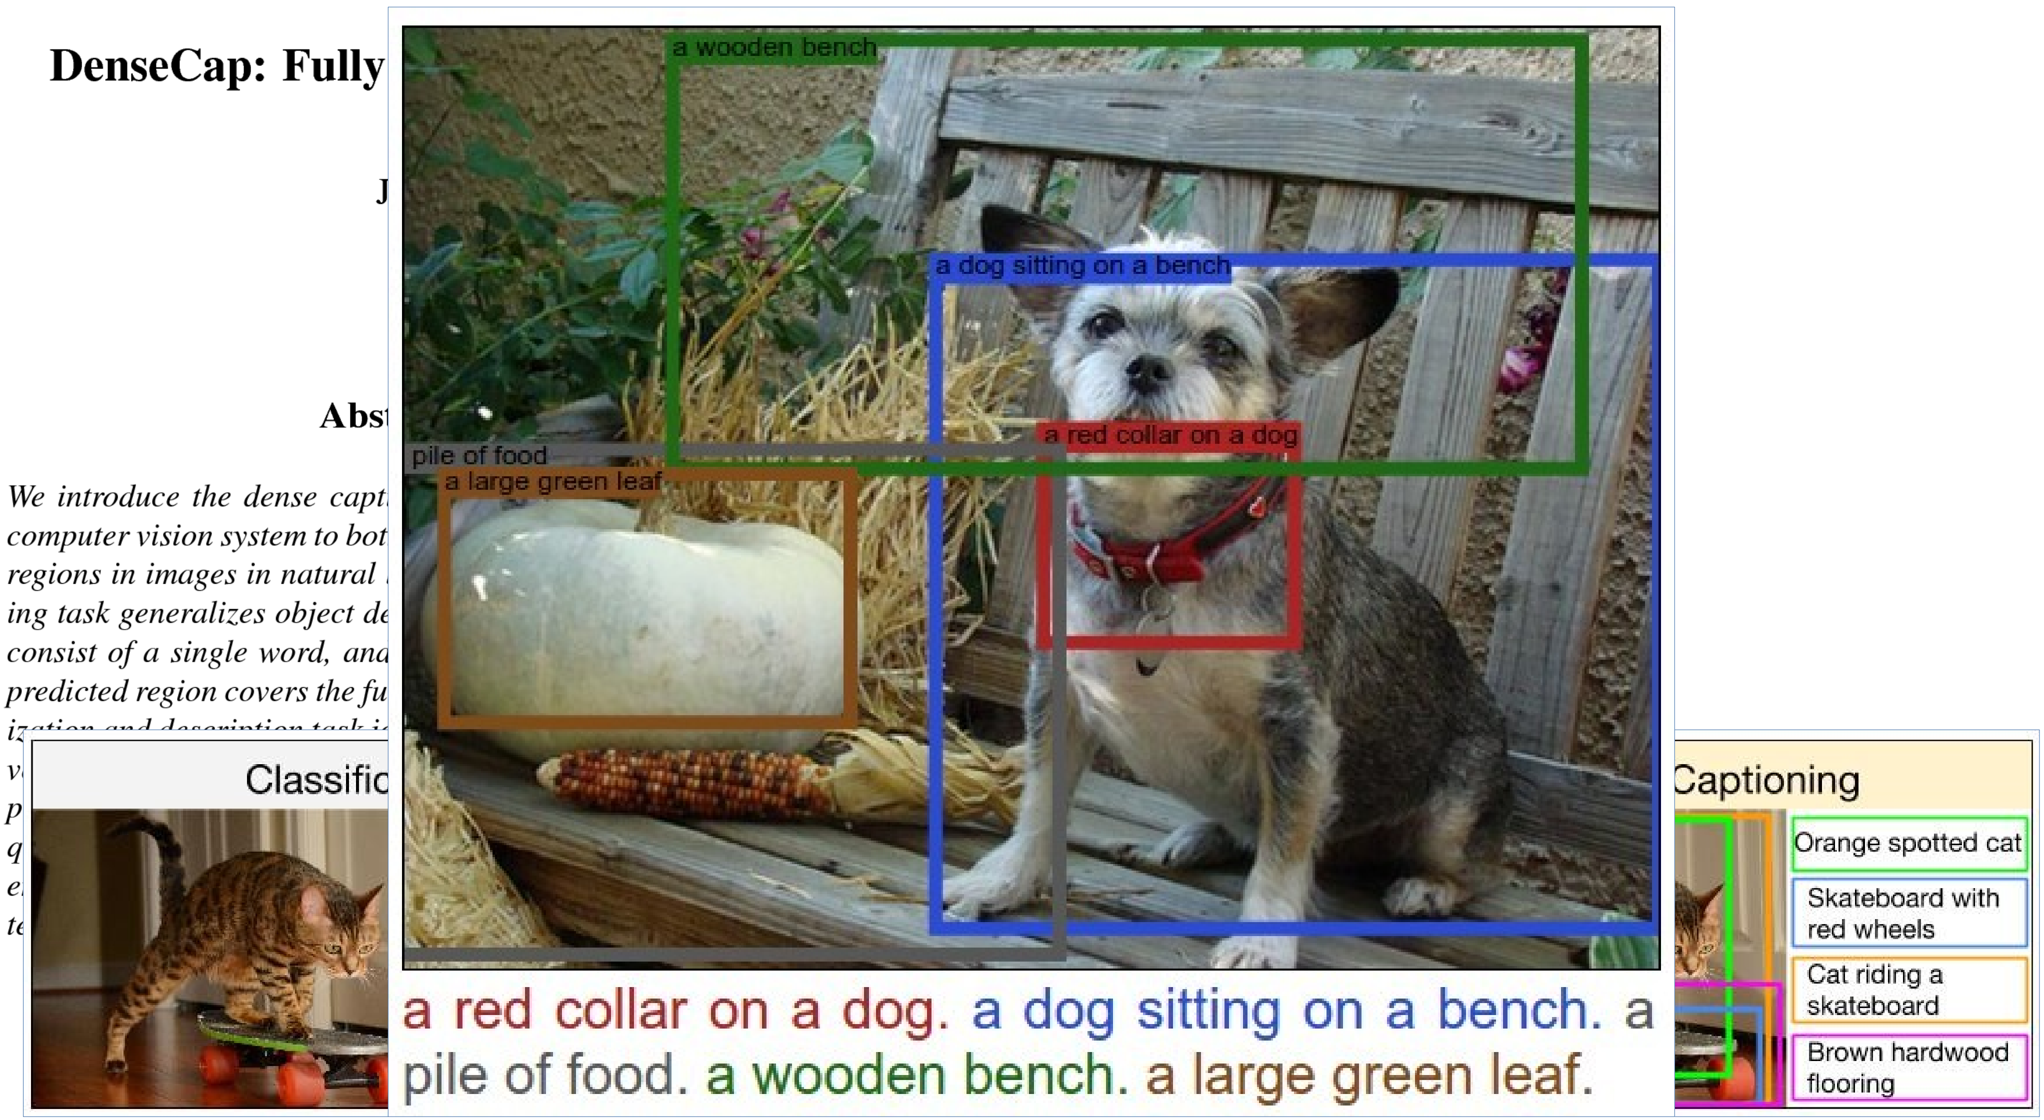

ACM Prize in Computing

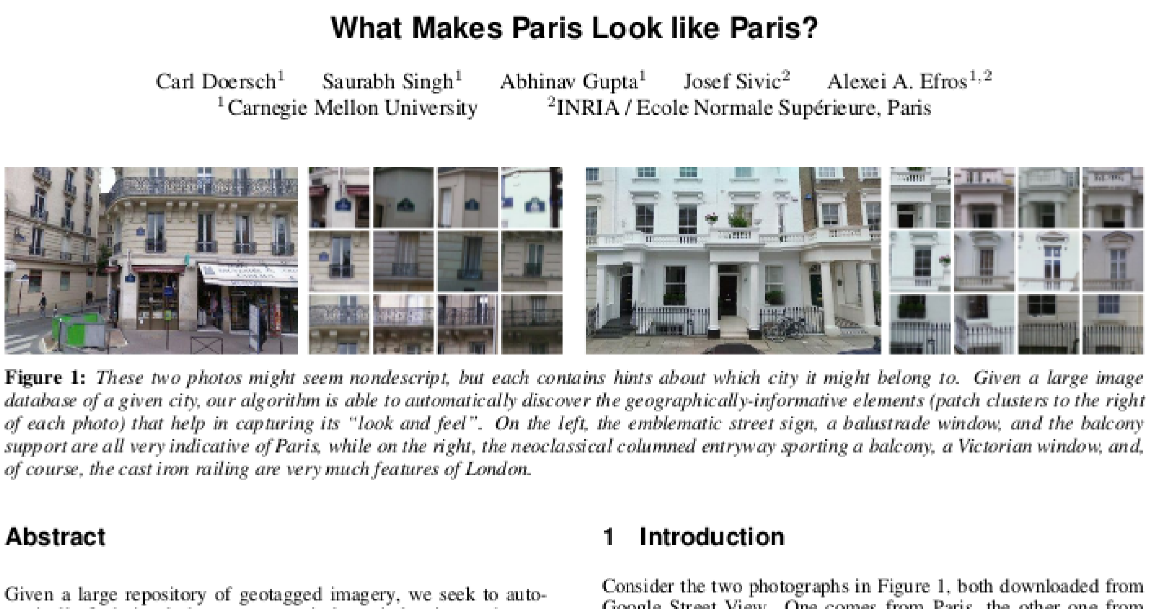

What Makes Paris Look like Paris?

Carl Doersch[1] Saurabh Singh[1] Abhinav Gupta[1] Josef Sivic[2] Alexei A. Efros[1,2]
[1]Carnegie Mellon University [2]INRIA / Ecole Normale Supérieure, Paris

Figure 1: *These two photos might seem nondescript, but each contains hints about which city it might belong to. Given a large image database of a given city, our algorithm is able to automatically discover the geographically-informative elements (patch clusters to the right of each photo) that help in capturing its "look and feel". On the left, the emblematic street sign, a balustrade window, and the balcony support are all very indicative of Paris, while on the right, the neoclassical columned entryway sporting a balcony, a Victorian window, and, of course, the cast iron railing are very much features of London.*

Abstract

Given a large repository of geotagged imagery, we seek to automatically find visual elements, e.g. windows, balconies, and street

ements at different geo-spatial scales, and geographically-informed image retrieval.

1 Introduction

Consider the two photographs in Figure 1, both downloaded from Google Street View. One comes from Paris, the other one from

windows with railings, the particular style of balconies, the distinctive doorways, the traditional blue/green/white street signs, etc. were particularly helpful. Finding those features can be difficult

Scene Completion Using Millions of Photographs

James Hays Alexei A. Efros
Carnegie Mellon University

Original Image Input Scene Matches Output

Figure 1: Given an input image with a missing region, we use matching scenes from a large collection of photographs to complete the image.

Abstract

What can you do with a million images? In this paper we present a new image completion algorithm powered by a huge database of photographs gathered from the Web. The algorithm patches up holes in images by finding similar image regions in the database that are not only seamless but also semantically valid. Our chief insight is that while the space of images is effectively infinite, the space of semantically differentiable scenes is actually not that large. For many image completion tasks we are able to find similar scenes which contain image fragments that will convincingly complete the image. Our algorithm is entirely data-driven, requiring no annotations or labelling by the user. Unlike existing image completion methods, our algorithm can generate a diverse set of results for each input image and we allow users to select among them. We demonstrate the superiority of our algorithm over existing image completion approaches.

Keywords: Image Completion, Image Database, Image Com-

There are two fundamentally different strategies for image completion. The first aims to reconstruct, as accurately as possible, the data that *should have been* there, but somehow got occluded or corrupted. Methods attempting an accurate reconstruction have to use some other source of data in addition to the input image, such as video (using various background stabilization techniques, e.g. [Irani et al. 1995]) or multiple photographs of the same physical scene [Agarwala et al. 2004; Snavely et al. 2006].

The alternative is to try finding a plausible way to fill in the missing pixels, hallucinating data that *could have been* there. This is a much less easily quantifiable endeavor, relying instead on the studies of human visual perception. The most successful existing methods [Criminisi et al. 2003; Drori et al. 2003; Wexler et al. 2004; Wilczkowiak et al. 2005; Komodakis 2006] operate by extending adjacent textures and contours into the unknown region. This idea is derived from example-based texture synthesis [Efros and Leung 1999; Efros and Freeman 2001; Kwatra et al. 2003; Kwatra et al.

Scene Completion Using Millions of Photographs

James Hays Alexei A. Efros
Carnegie Mellon University

What can you do with a million images? In this paper we present a new image completion algorithm powered by a huge database of photographs gathered from the Web. The algorithm patches up

pletion. The first aims to reconstruct, as accurately as possible, the data that *should have been* there, but somehow got occluded or corrupted. Methods attempting an accurate reconstruction have

Colorful Image Colorization

Richard Zhang, Phillip Isola, Alexei A. Efros
{rich.zhang,isola,efros}@eecs.berkeley.edu

University of California, Berkeley

Abstract. Given a grayscale photograph as input, this paper attacks the problem of hallucinating a *plausible* color version of the photograph. This problem is clearly underconstrained, so previous approaches have either relied on significant user interaction or resulted in desaturated colorizations. We propose a fully automatic approach that produces vibrant and realistic colorizations. We embrace the underlying uncertainty of the problem by posing it as a classification task and use class-rebalancing at training time to increase the diversity of colors in the result. The system is implemented as a feed-forward pass in a CNN at test time and is trained on over a million color images. We evaluate our algorithm using a "colorization Turing test," asking human participants to choose between a generated and ground truth color image. Our method successfully fools humans on 32% of the trials, significantly higher than previous methods. Moreover, we show that colorization can be a powerful pretext task for

Colorful Image Colorization

Richard Zhang, Phillip Isola, Alexei A. Efros

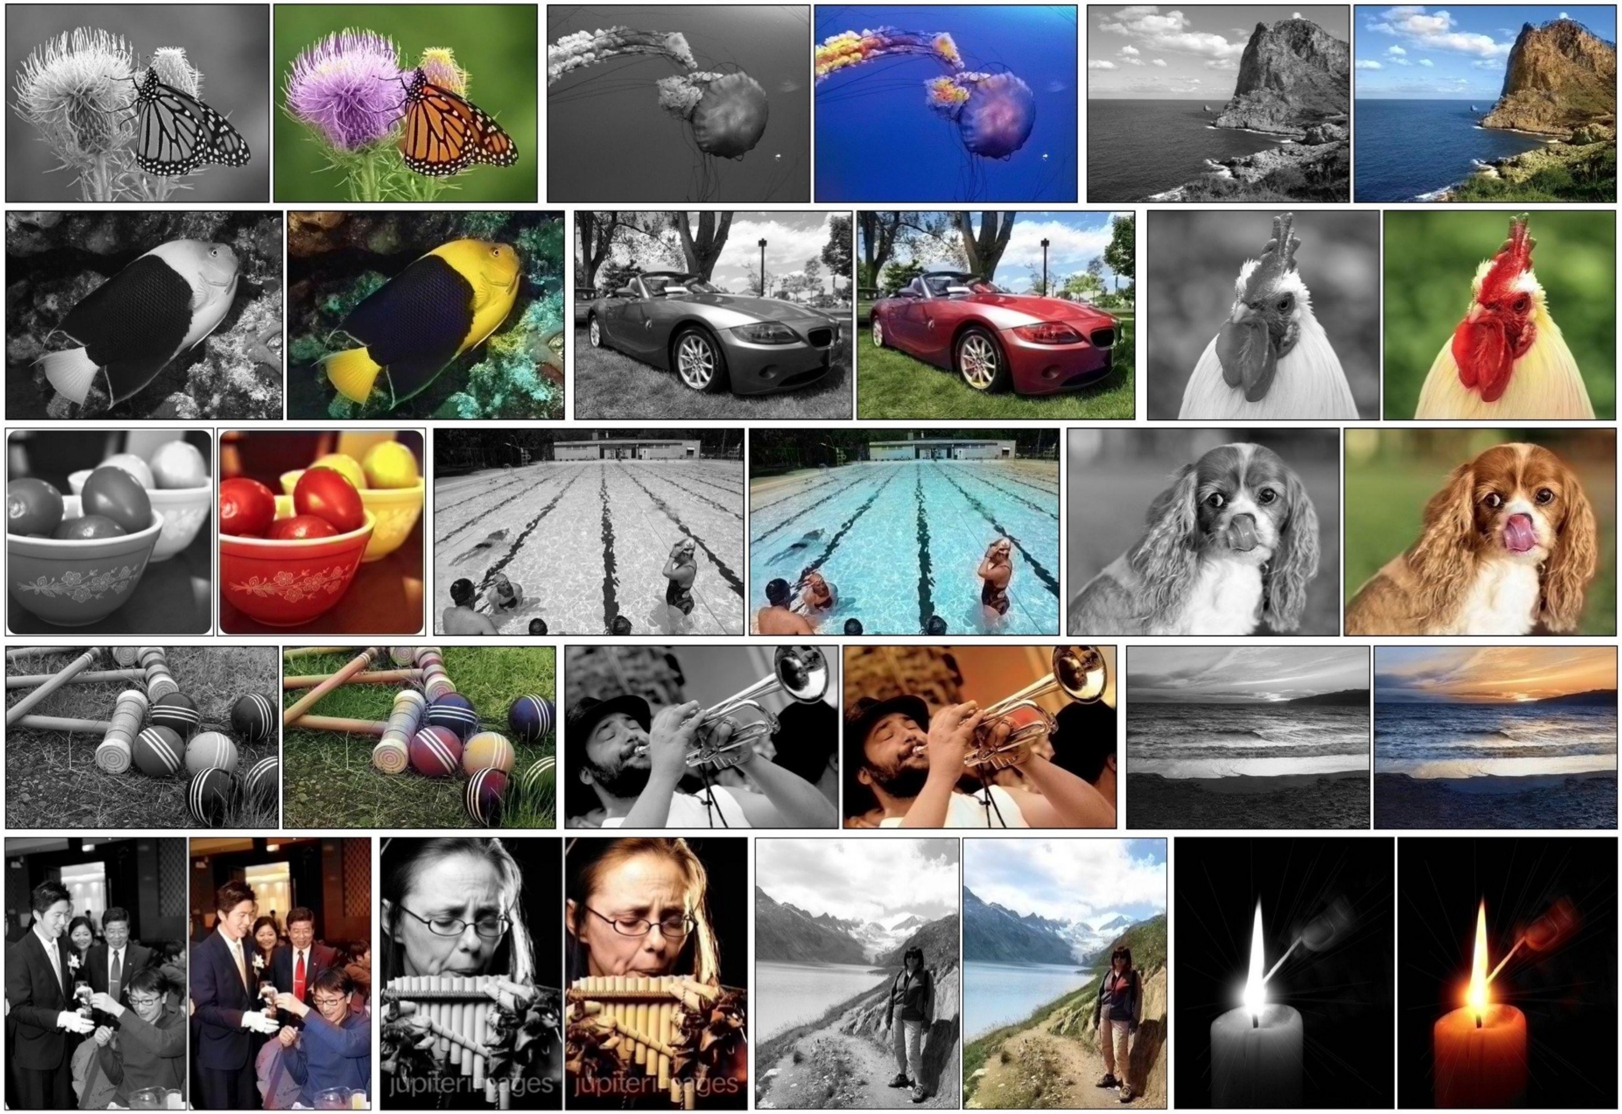

deep learning

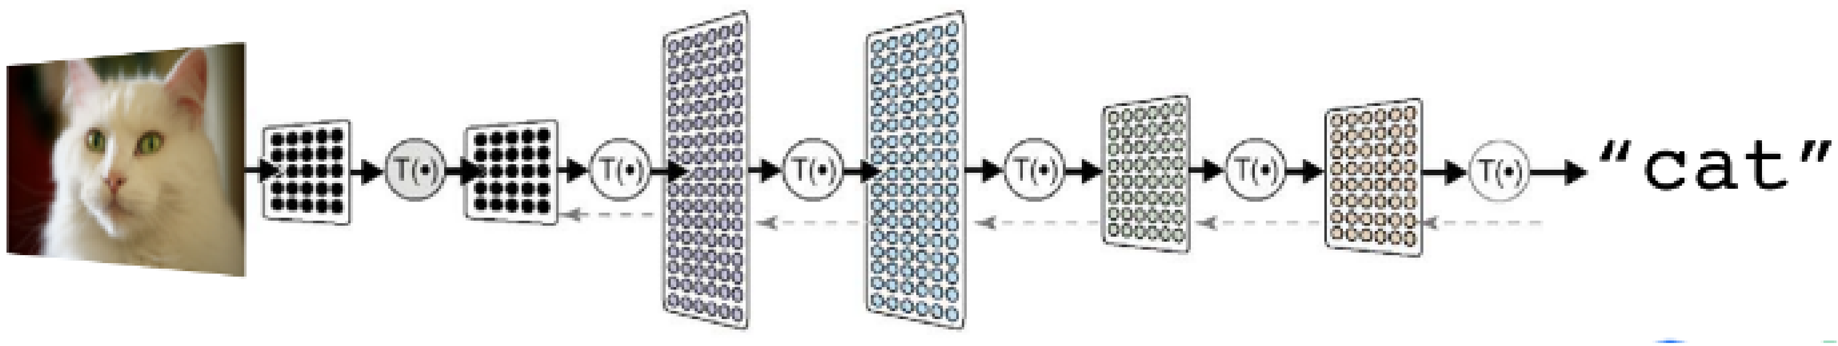

Alchemy or Chemistry?

deepvis

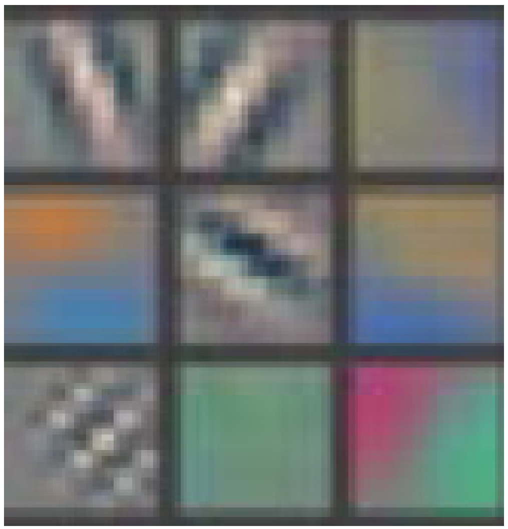

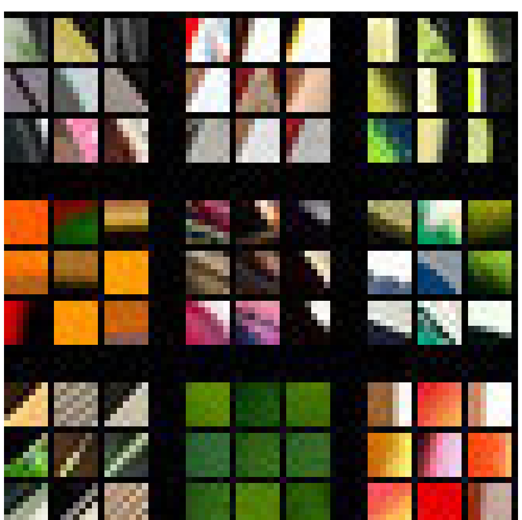

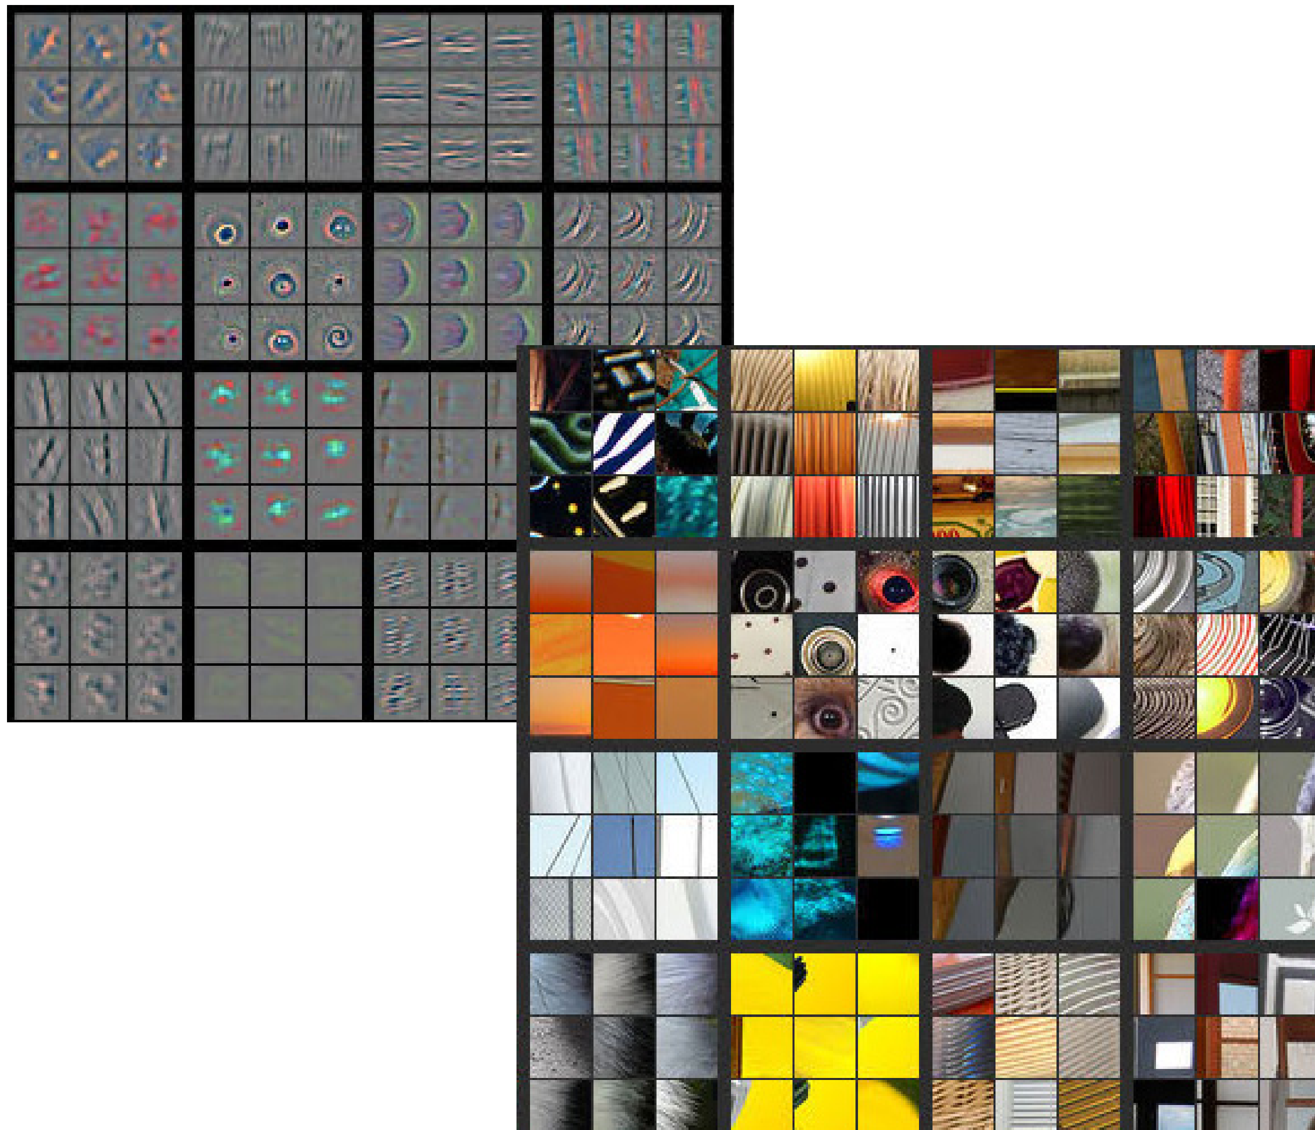

deepvis

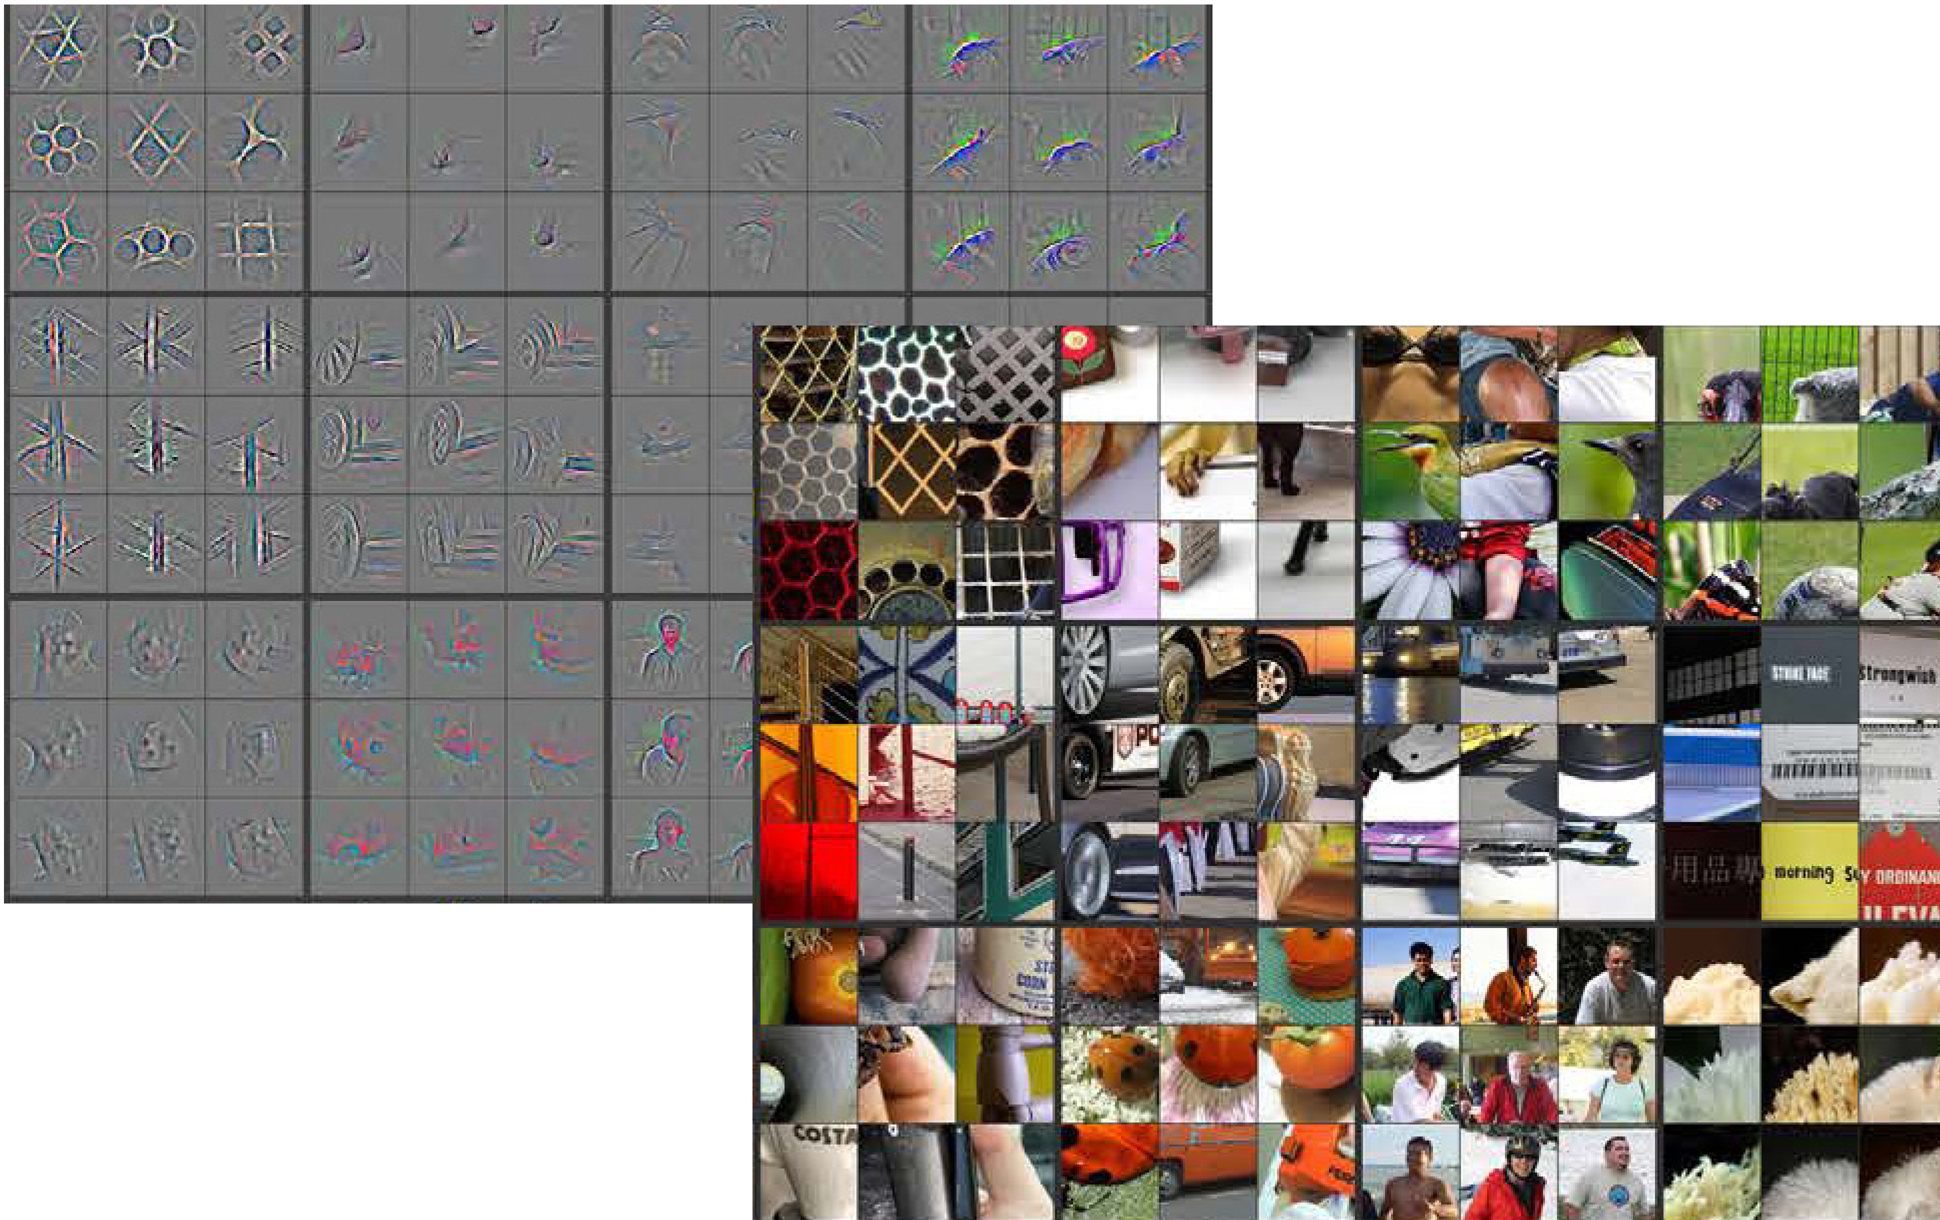

deepvis

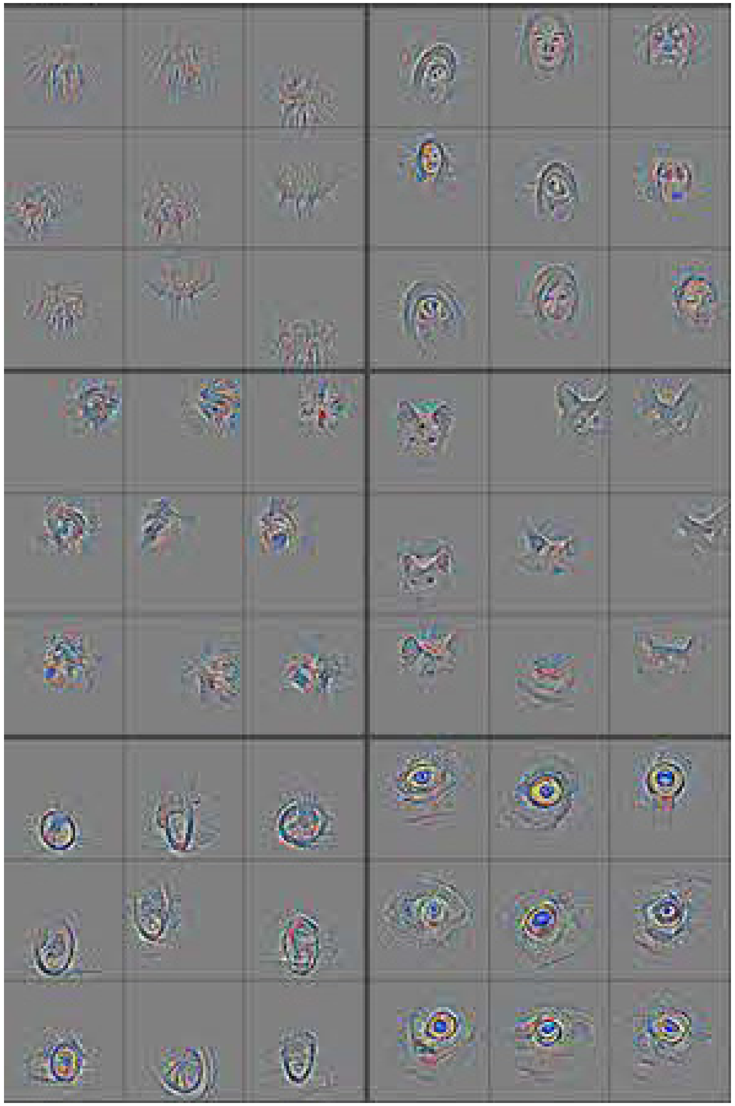

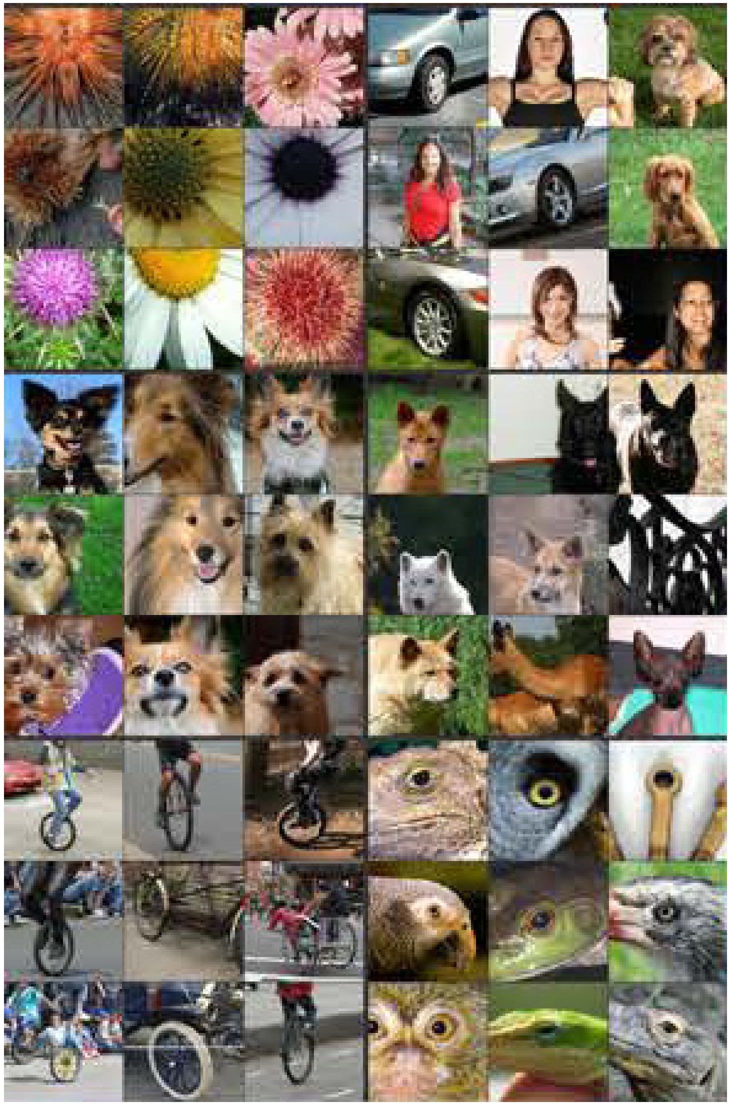

Colorful Image Colorization

Richard Zhang, Phillip Isola, Alexei A. Efros
{rich.zhang,isola,efros}@eecs.berkeley.edu

University of California, Berkeley

Abstra
the pro
This pr
either r
orizatio

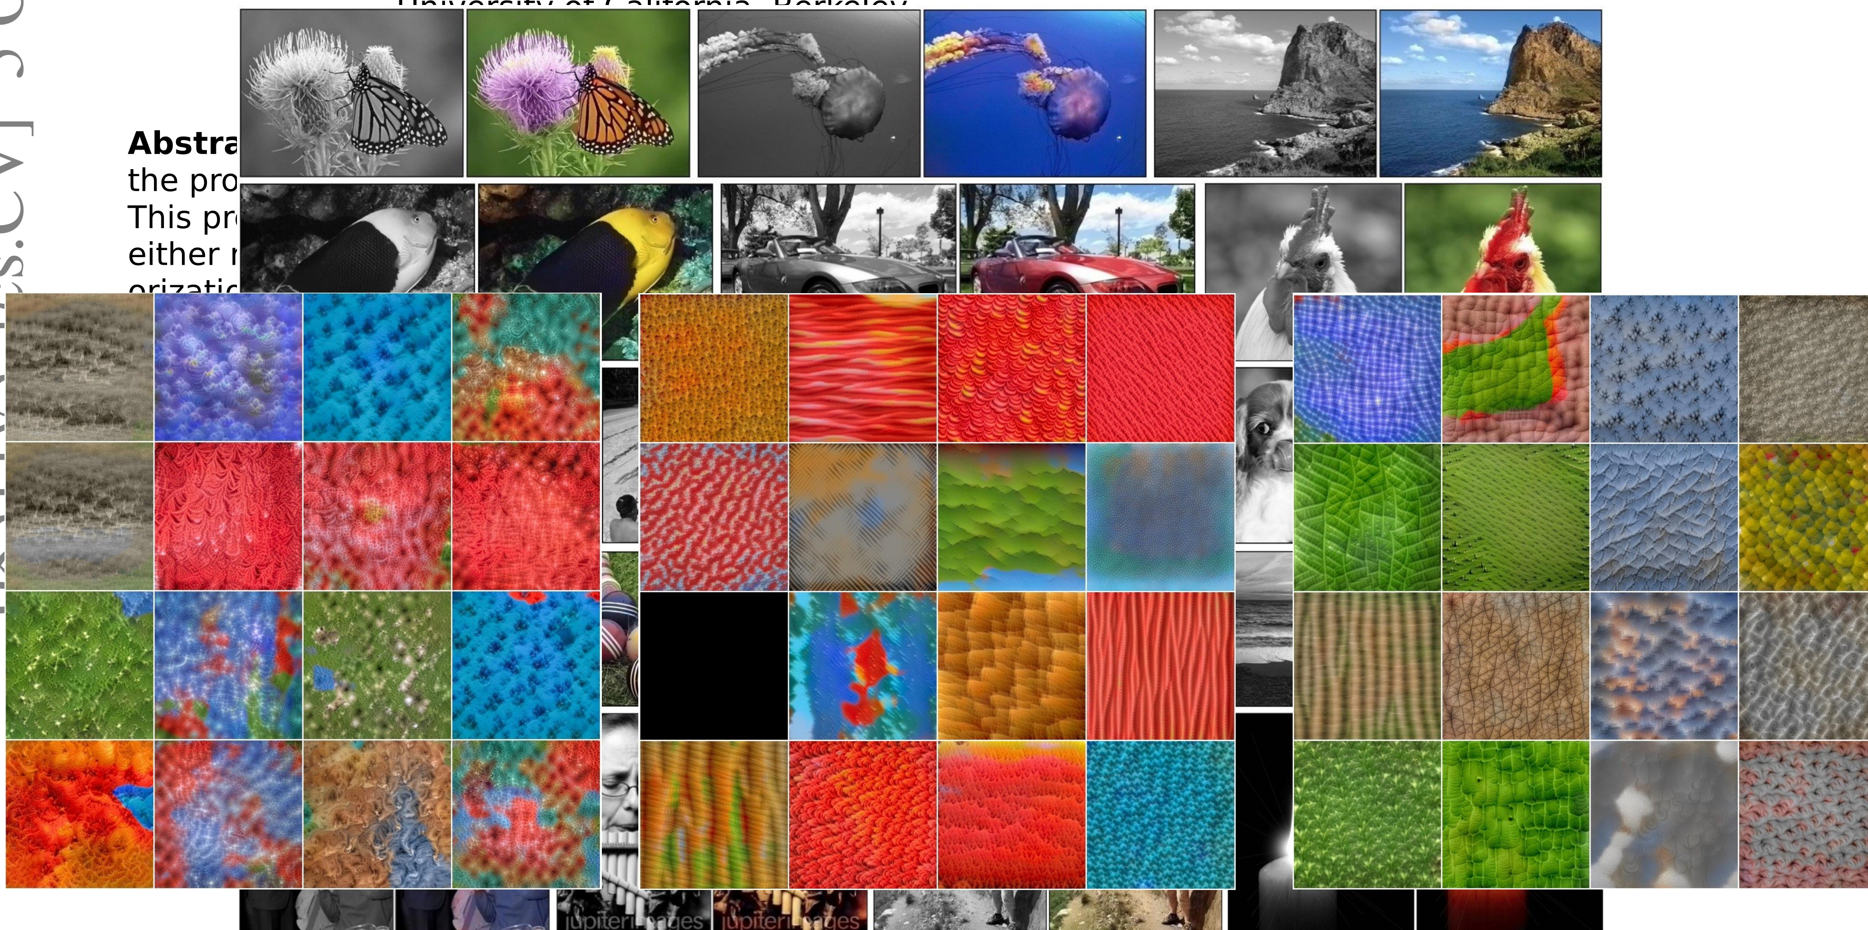

Distill Journal

The Building Blocks of Interpretability

Interpretability techniques are normally studied in isolation.
We explore the powerful interfaces that arise when you combine them —
and the rich structure of this combinatorial space.

For instance, by combining feature visualization (*what is a neuron looking for?*) with attribution (*how does it affect the output?*), we can explore how the network decides between labels like **Labrador retriever** and **tiger cat**.

Several floppy ear detectors seem to be important when distinguishing dogs, whereas pointy ears are used to classify "tiger cat".

CHANNELS THAT MOST SUPPORT ... LABRADOR RETRIEVER ▾ ⟷ TIGER CAT ▾

feature visualization of channel

hover for attribution maps →

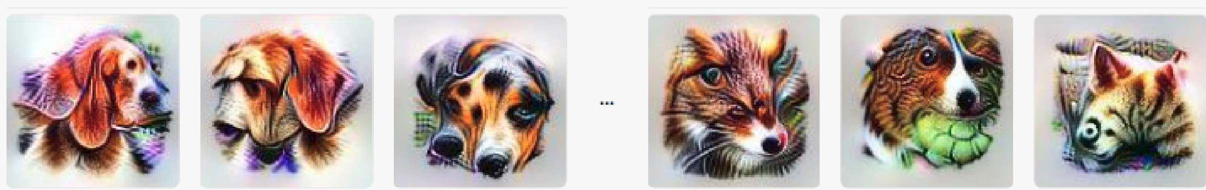

				...			
net evidence	1.63	1.51	1.19		1.32	1.54	1.72
for "Labrador retriever"	1.22	1.24	1.32		-0.70	-1.24	-0.43
for "tiger cat"	-0.40	-0.27	0.13		0.62	0.30	1.29

deep dreams

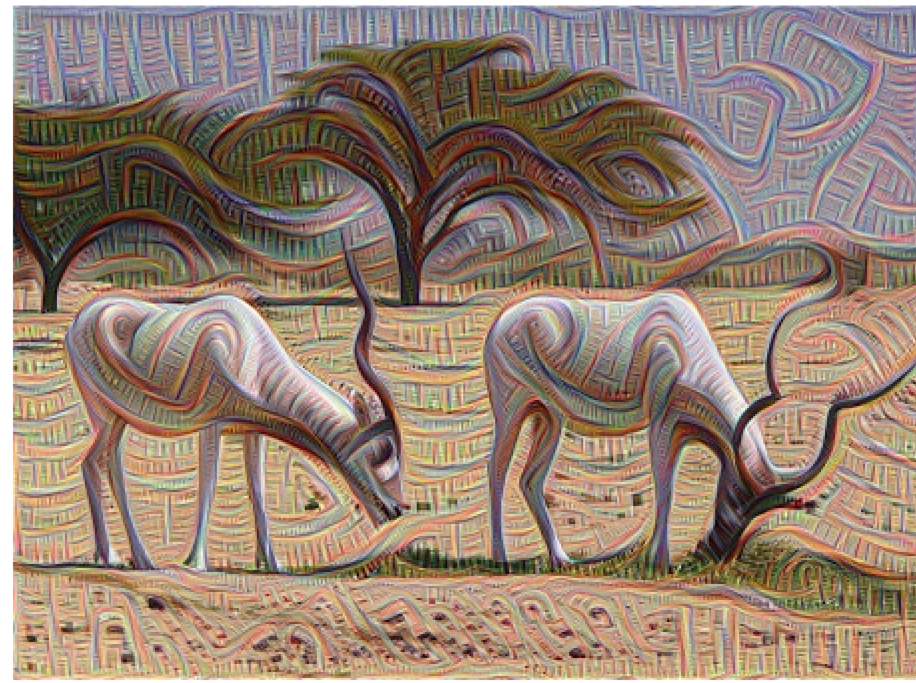

deep dreams

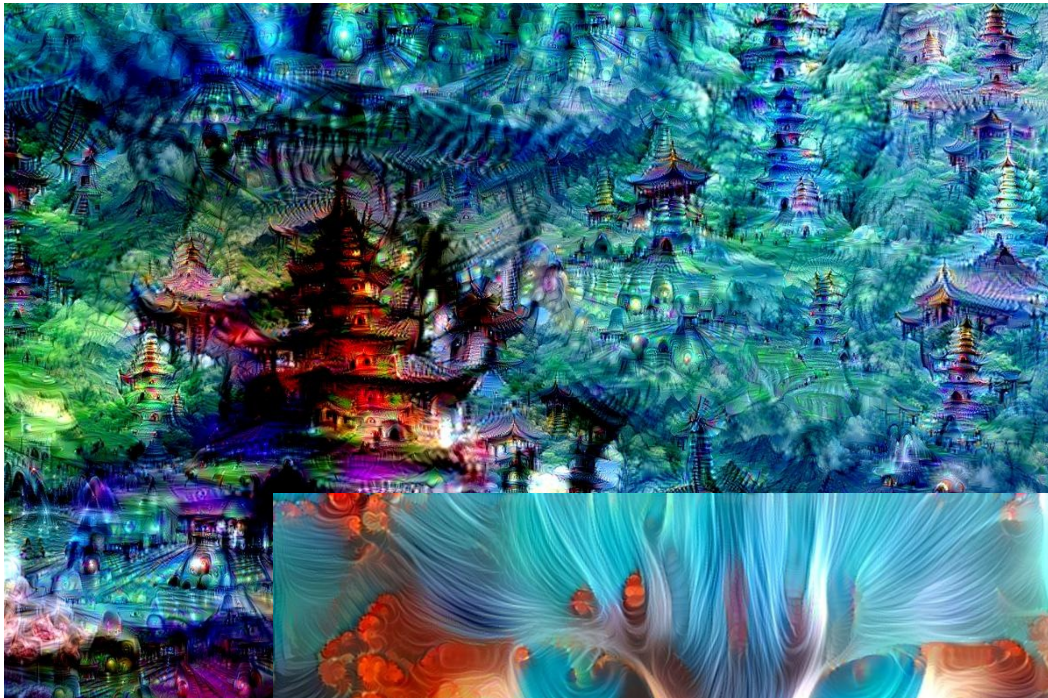
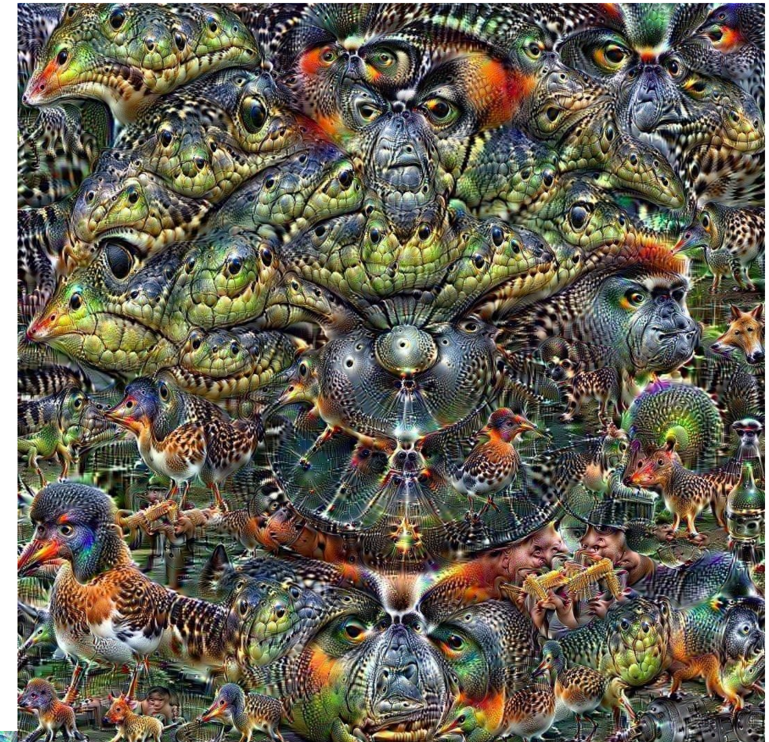

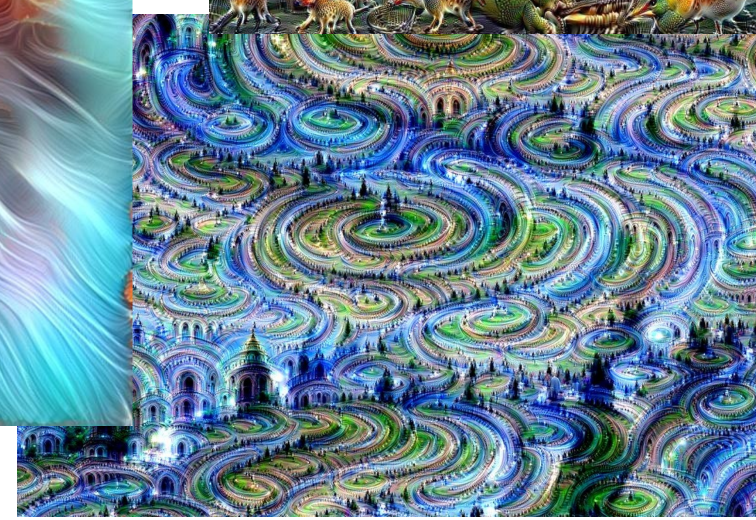

real deep dreams?

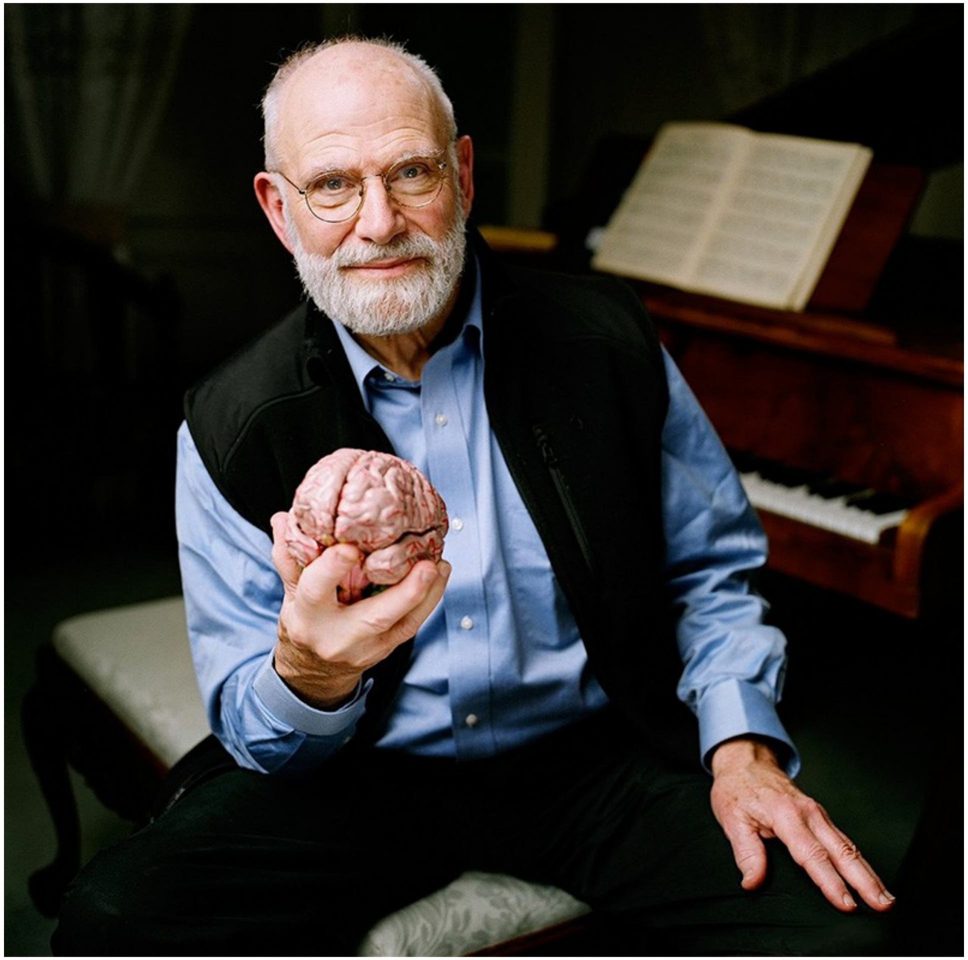

https://www.ted.com/talks/oliver_sacks_what_hallucination_reveals_about_our_minds

human vs machine

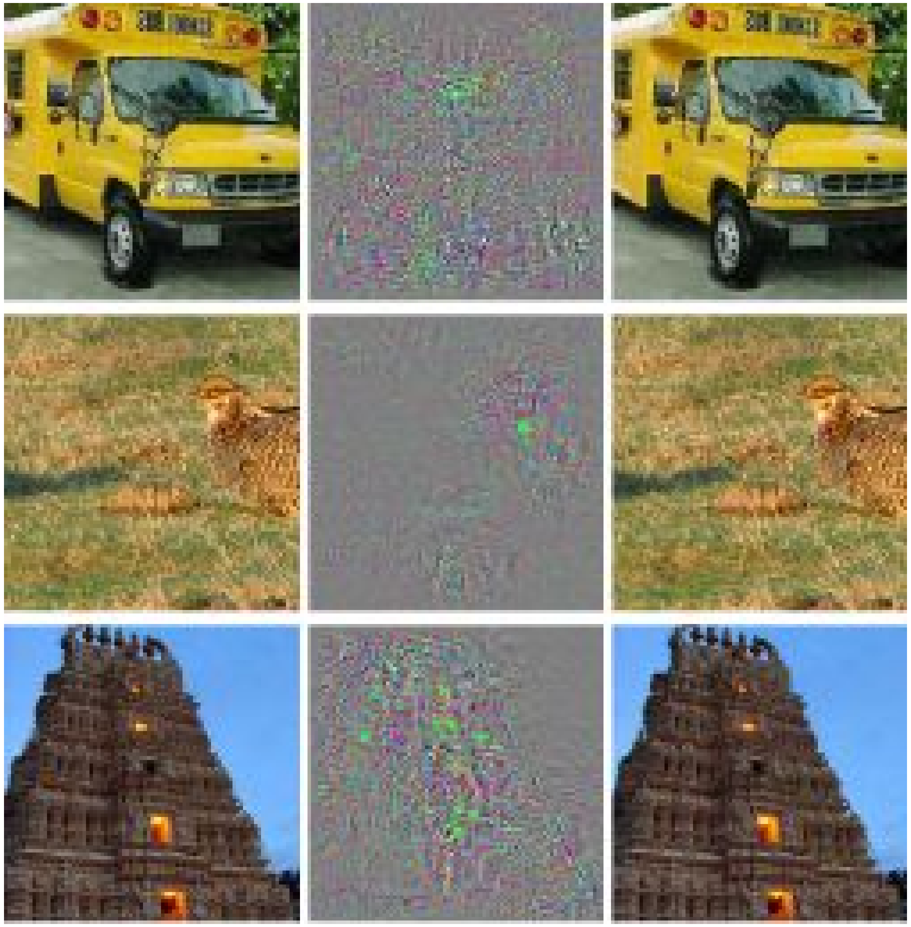

human vs machine

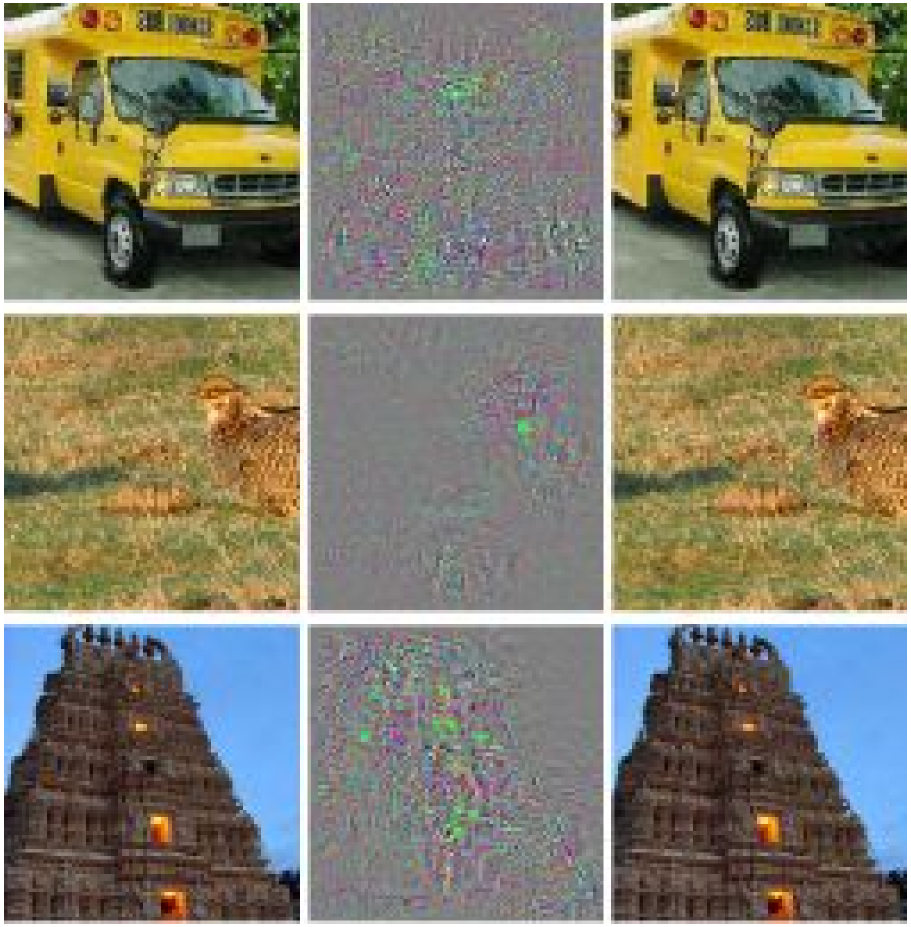

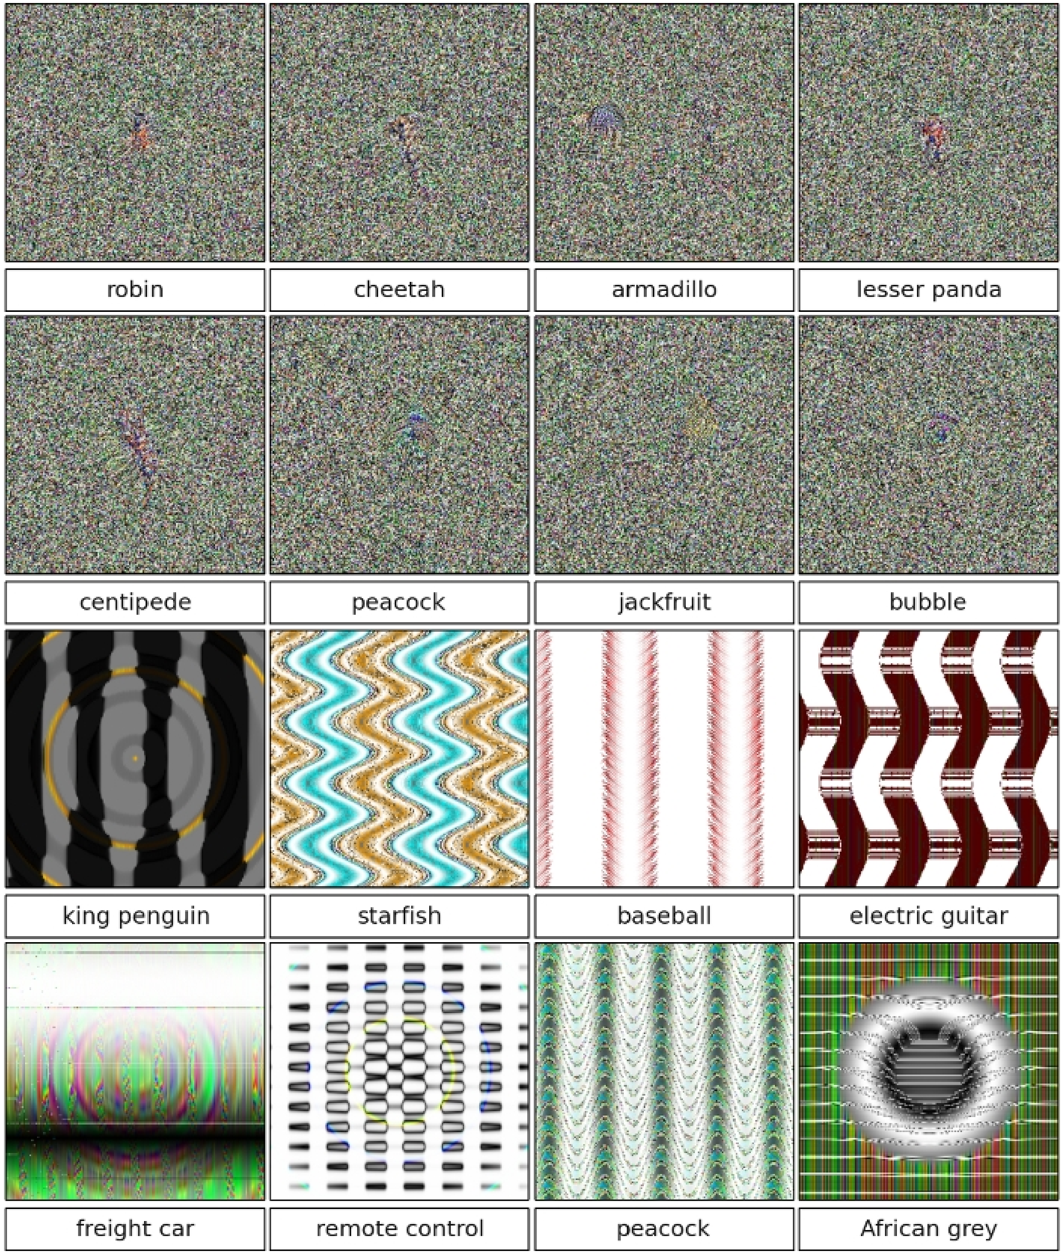

human vs machine

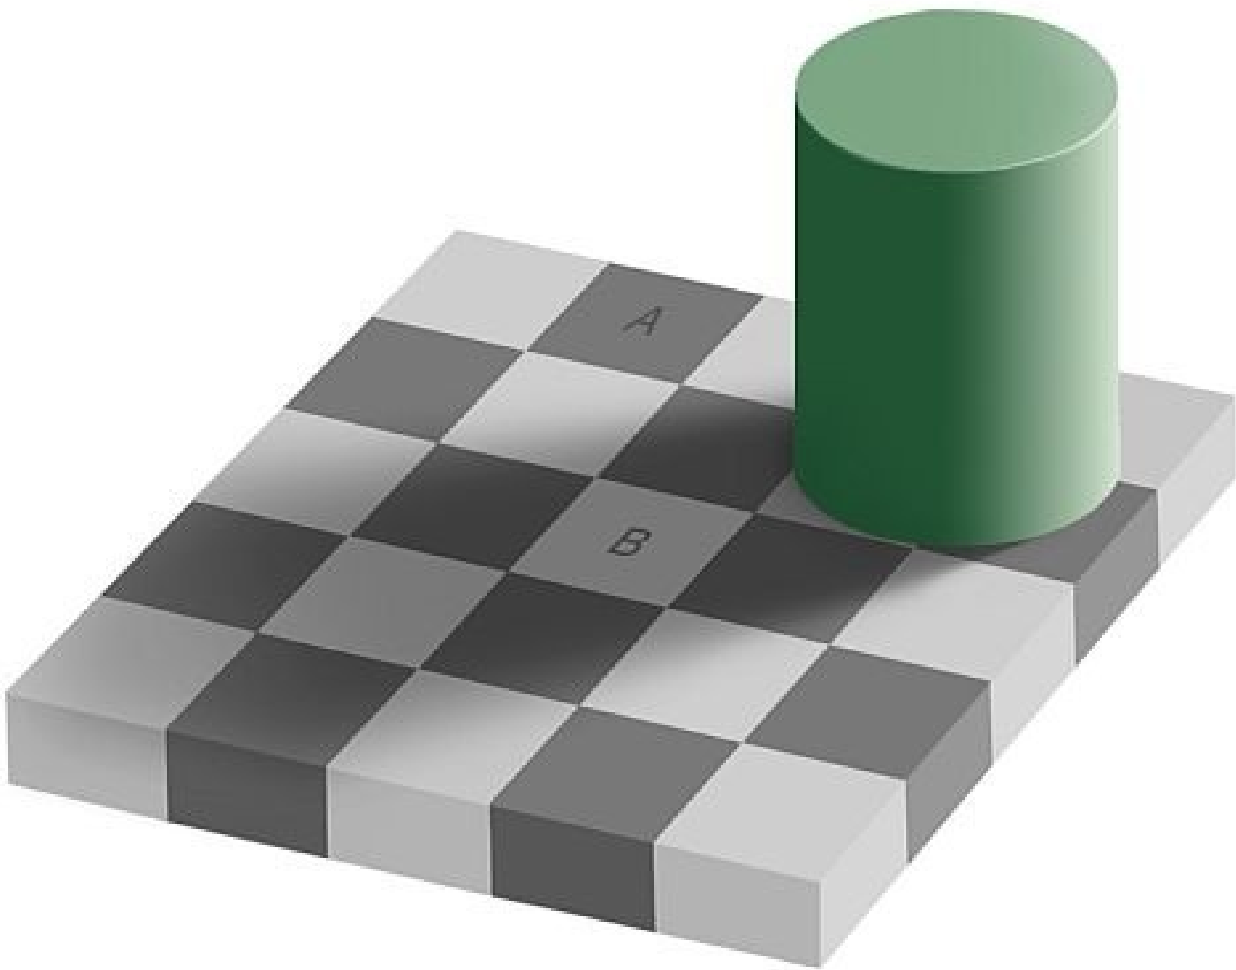

human vs machine

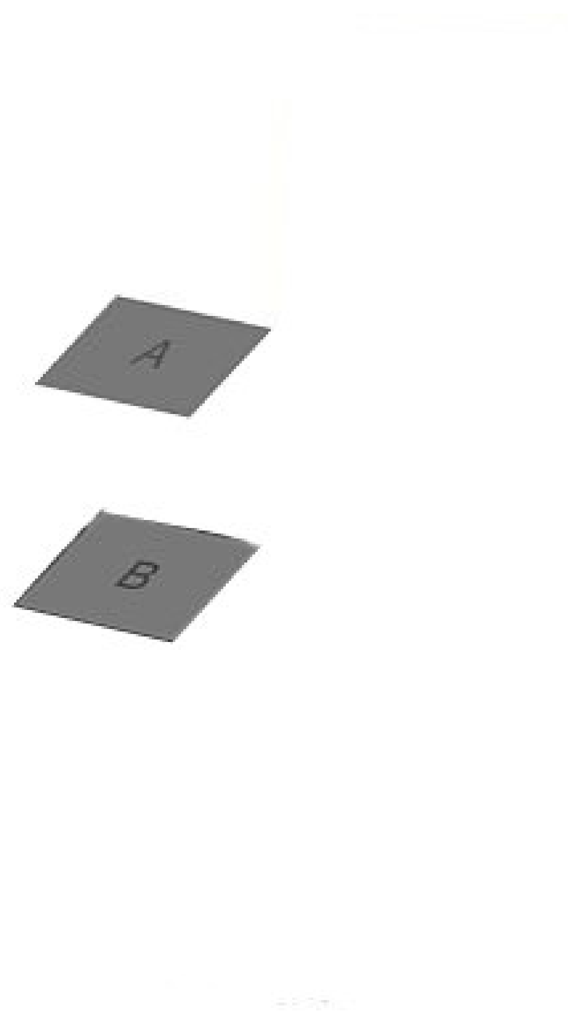

learning to see

https://www.ted.com/talks/pawan_sinha_on_how_brains_learn_to_see

learning to see

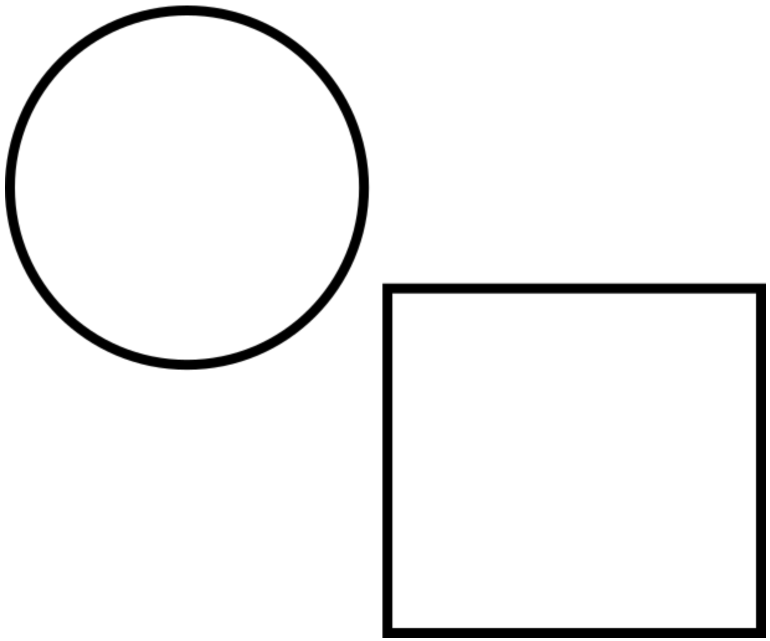

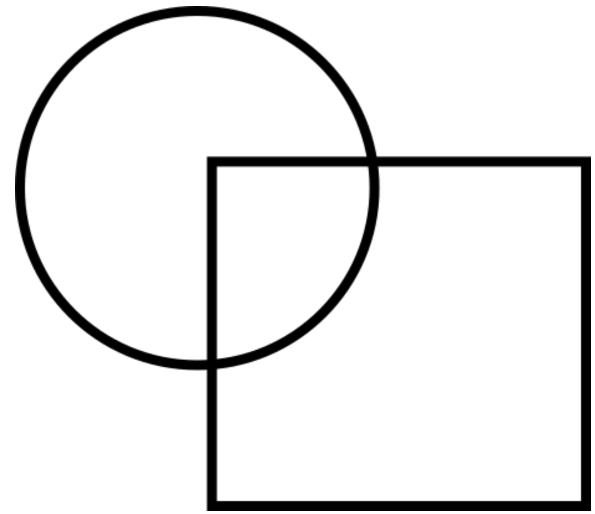

https://www.ted.com/talks/pawan_sinha_on_how_brains_learn_to_see

learning to see

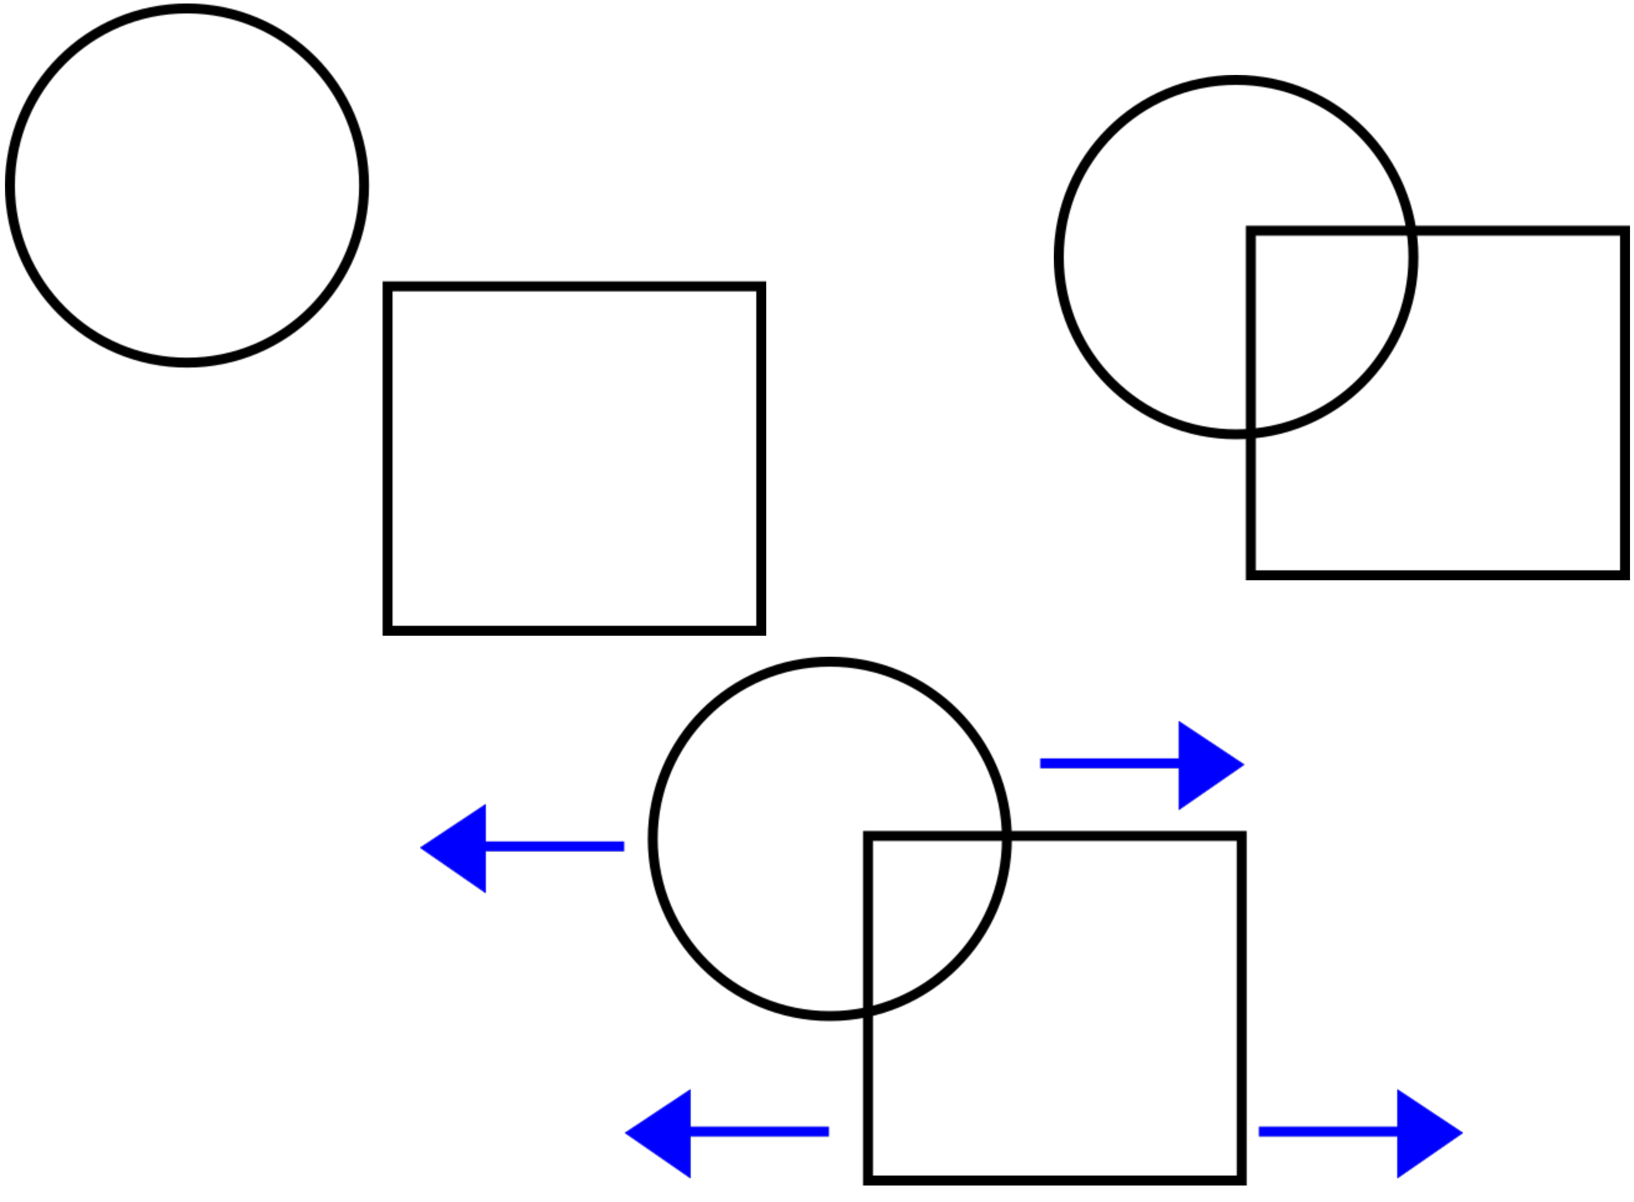

https://www.ted.com/talks/pawan_sinha_on_how_brains_learn_to_see

connectcome

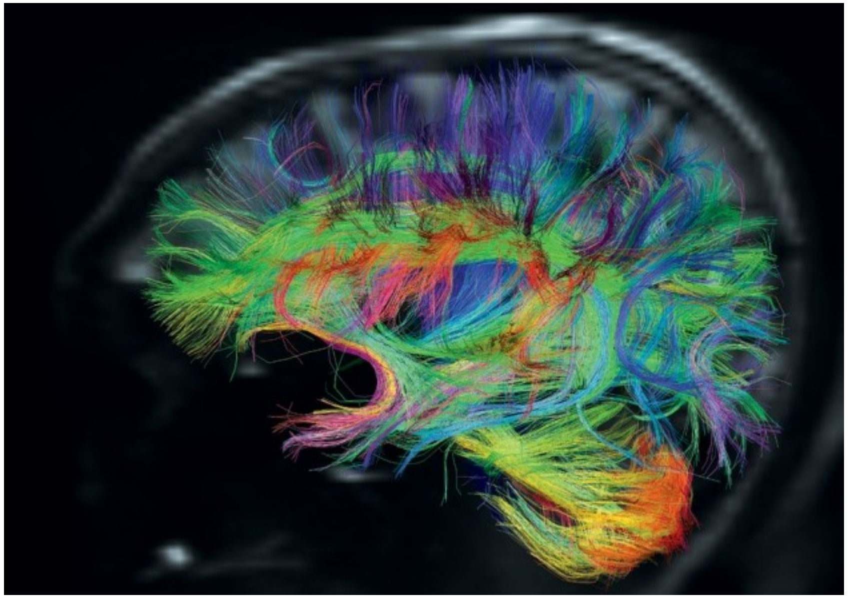

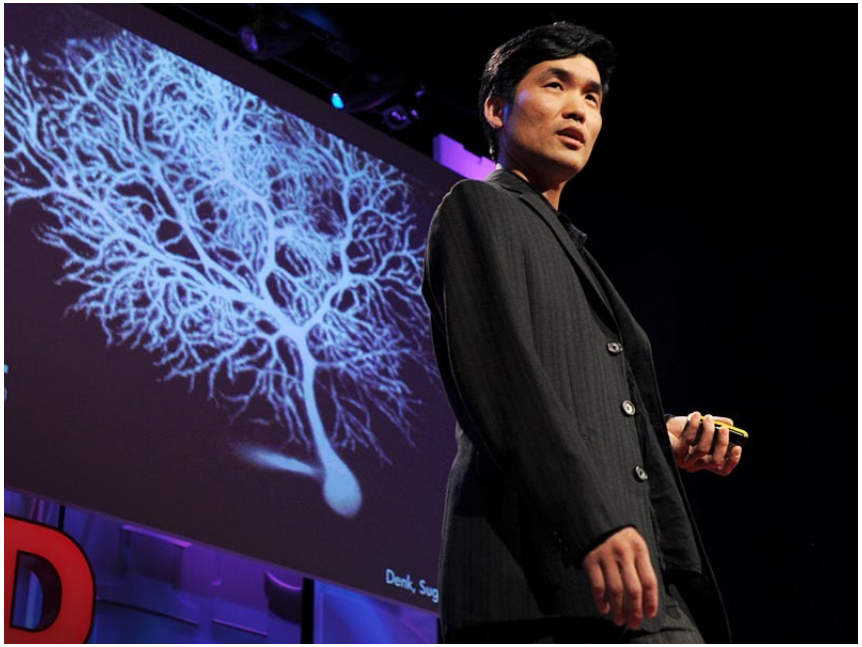

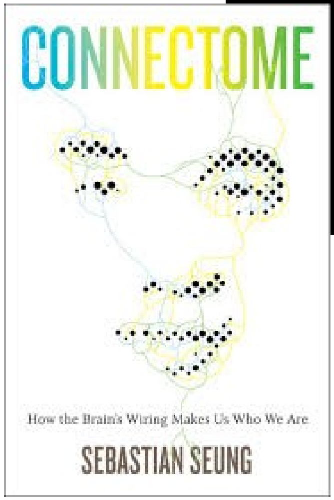

Blue Brain Project

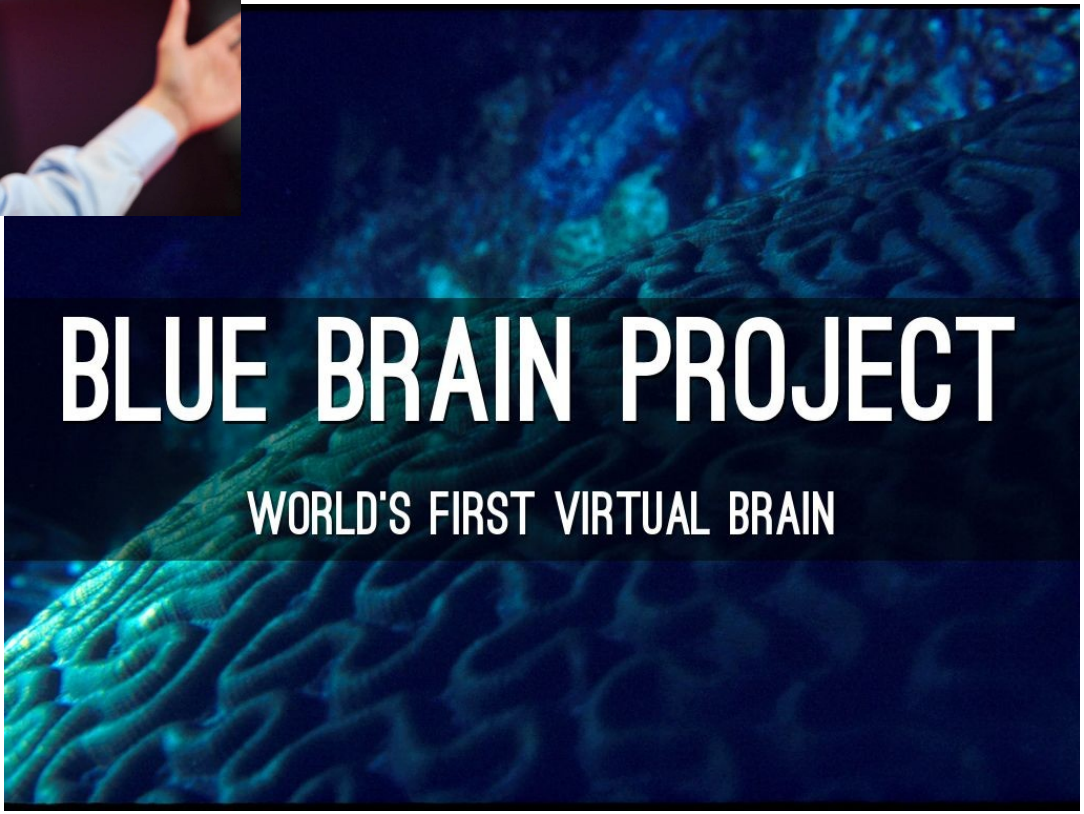

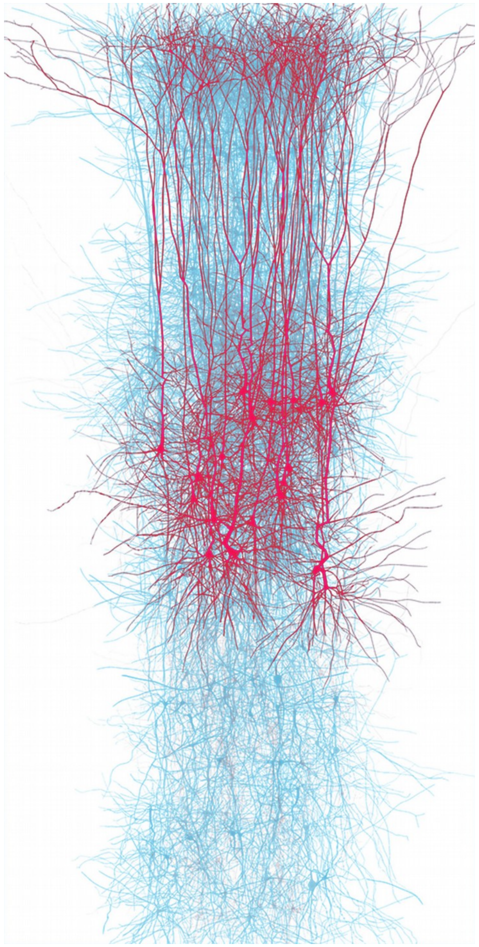

https://www.ted.com/talks/henry_markram_supercomputing_the_brain_s_secrets
https://www.youtube.com/watch?time_continue=2&v=2qTuZlMvFgY

Reading Minds

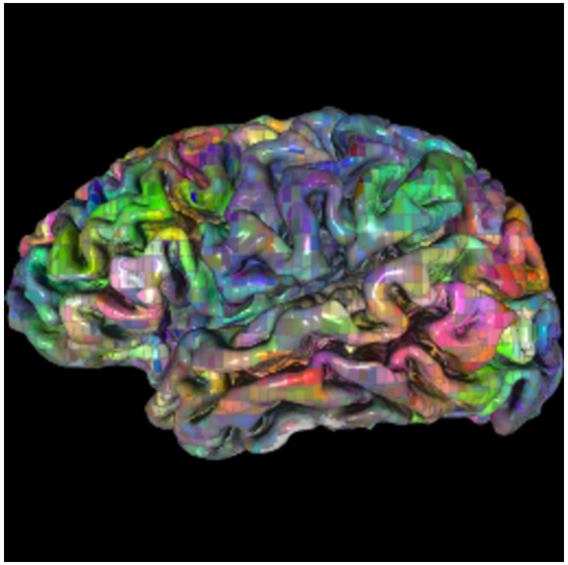

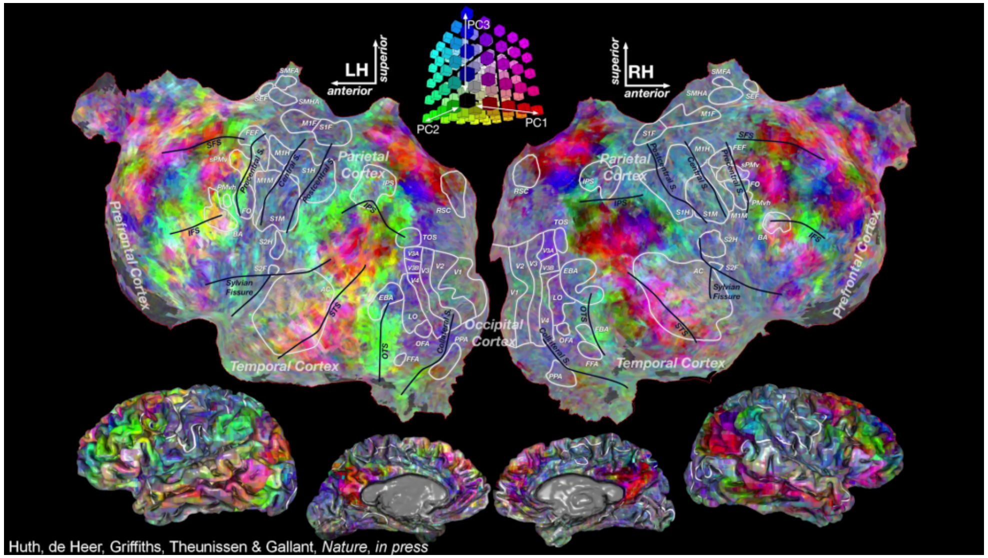

https://www.youtube.com/watch?v=Ecvv-EvOj8M

The Big and the Small: Challenges of Imaging the Brain's Circuits

Jeff W. Lichtman[1]* and Winfried Denk[2]*

The relation between the structure of the nervous system and its function is more poorly understood than the relation between structure and function in any other organ system. We explore why bridging the structure-function divide is uniquely difficult in the brain. These difficulties also explain the thrust behind the enormous amount of innovation centered on microscopy in neuroscience. We highlight some recent progress and the challenges that remain.

A central theme of biology is the relation between the structure and function of things. By structure, we mean the physical form of something, a property that humans can apprehend by touch (if the object is big enough) or by sight. Right now, the leading edge of this effort is the field known by the general name "structural biology" but is focused

tem, where much progress has been made at the molecular and functional level. But notwithstanding the extraordinary insights of neurobiology's foremost structural biologist, Cajal, our understanding of the relation between the structure and function of the brain remains primitive, especially when compared to other organ systems. There is no other organ system where so

research to determine the full extent of cell-type diversity in this small part of the nervous system, because the range of cell types continues to grow as the analysis becomes more refined. Moreover, neuronal cellu
one region to
brain serves a
itself. Few be
of retinal cell
in trying to u
bral cortex.
Researche
diversity by
ities to categ
basis of the m
molecular bio
generating tis
cent markers
of neuronal ce
is clear: The
thesis of inhi
only in inhib
molecular ma
derstood (4).

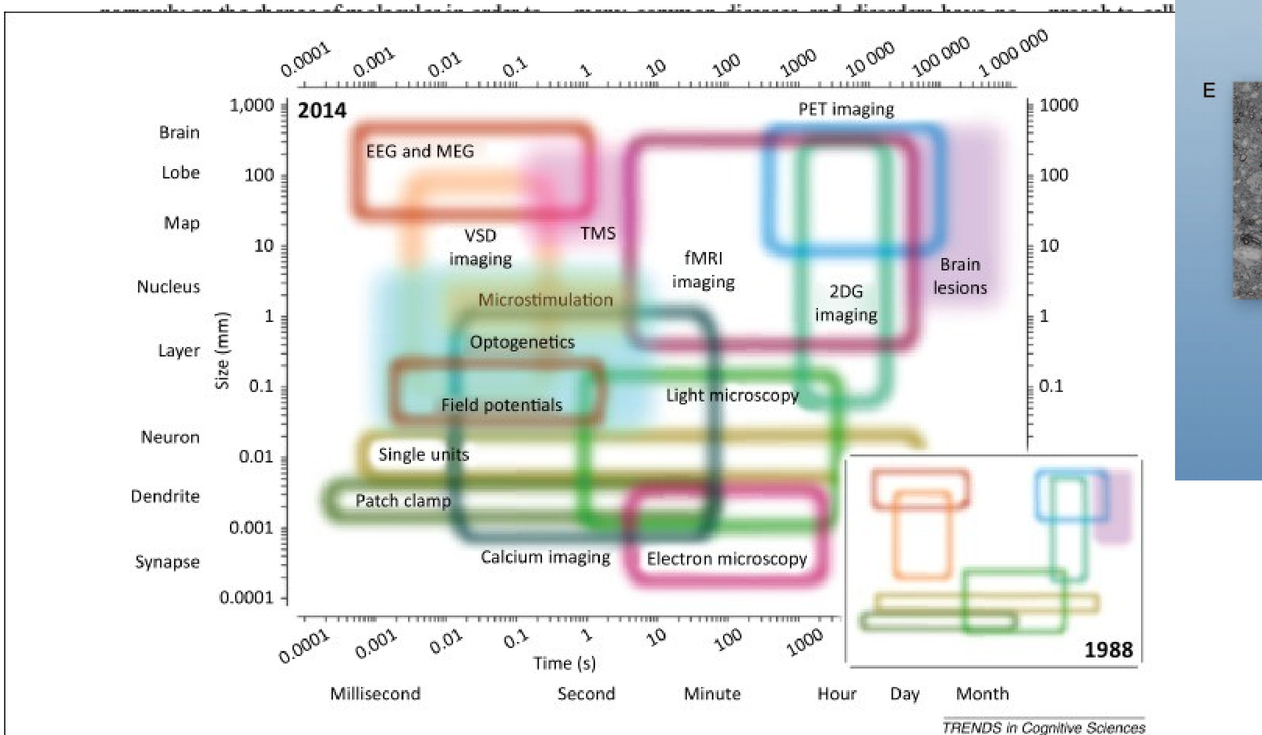

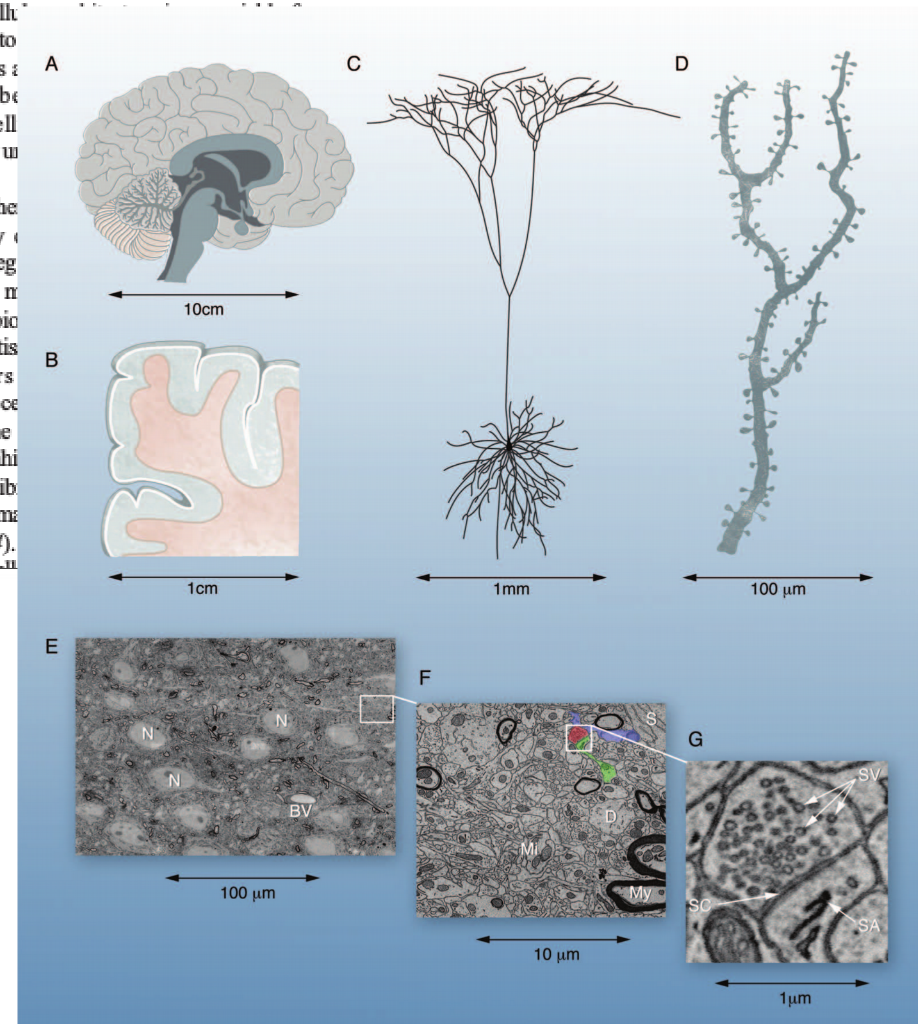

Network Neuroscience

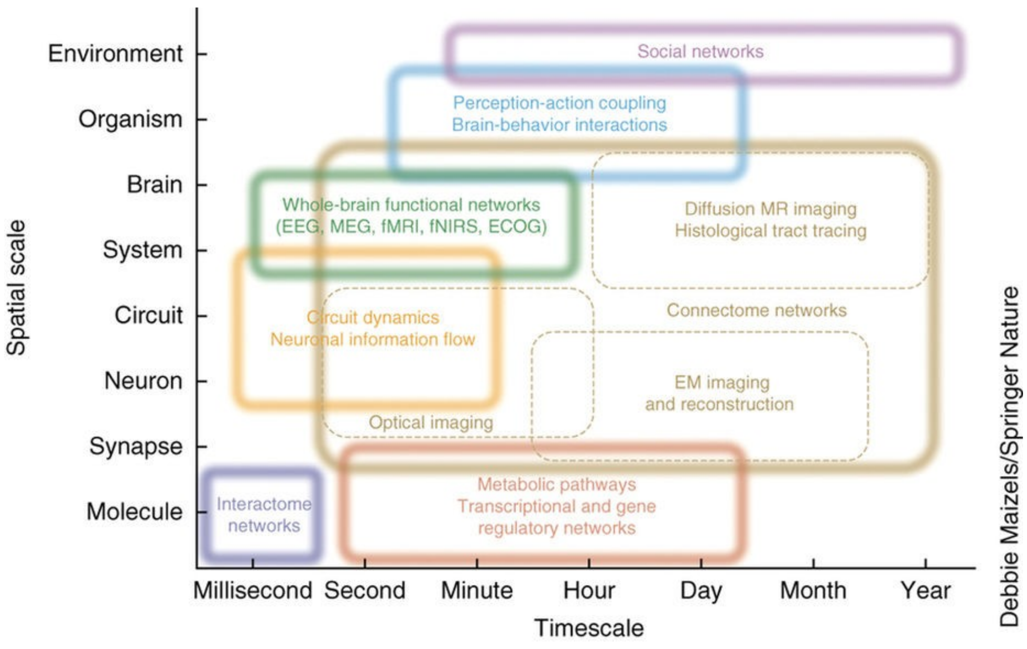

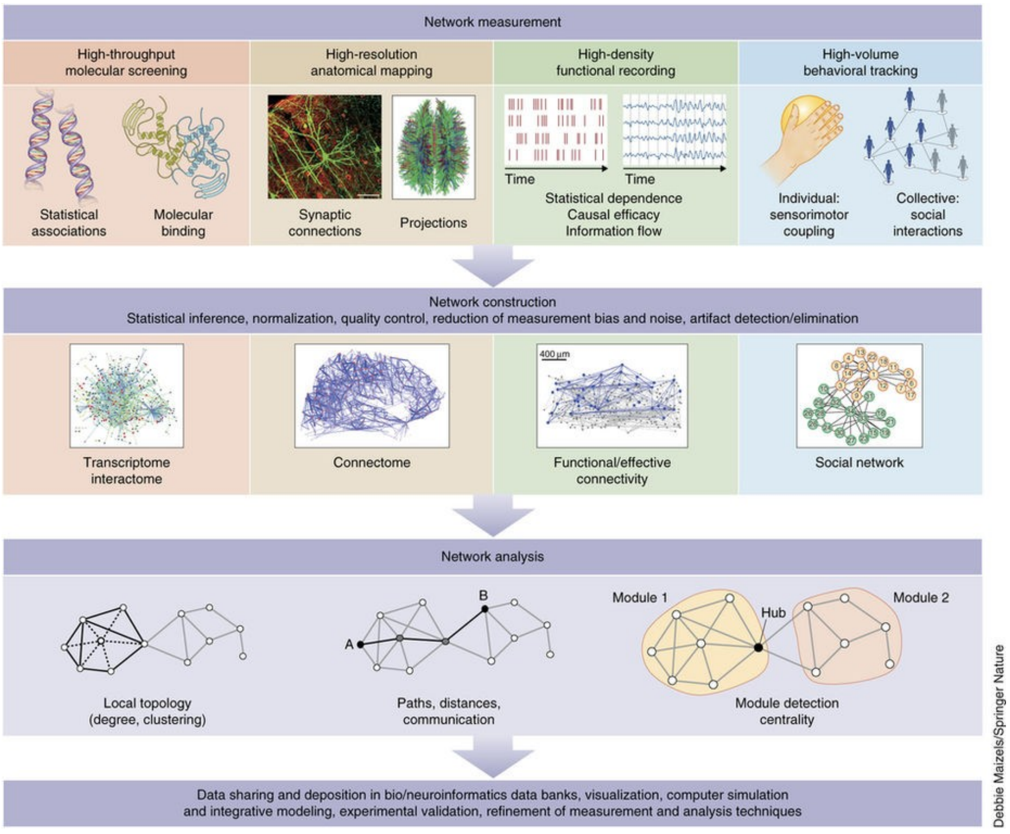

applications

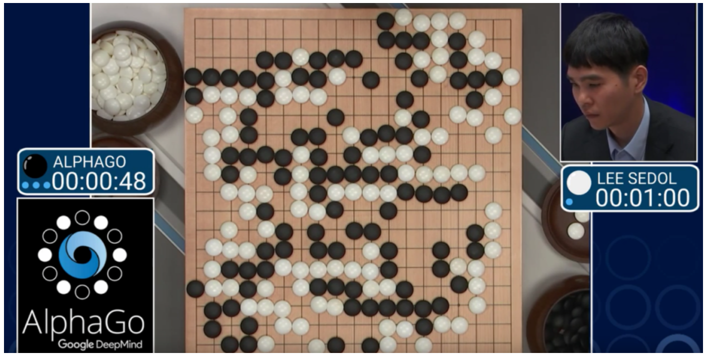

https://www.alphagomovie.com/

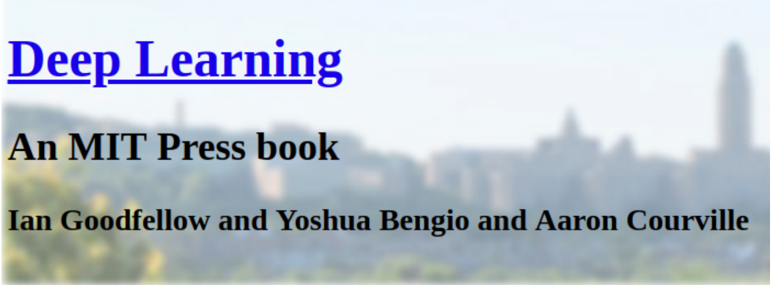

http://www.deeplearningbook.org/

Google's Neural Machine Translation

...reduced translation errors by an average of 60% when compared to the prior Google Translate technology

https://www.ted.com/talks/jeremy_howard_the_wonderful_and_terrifying_implications_of_computers_that_can_learn

http://www.fast.ai/

Moshe Vardi

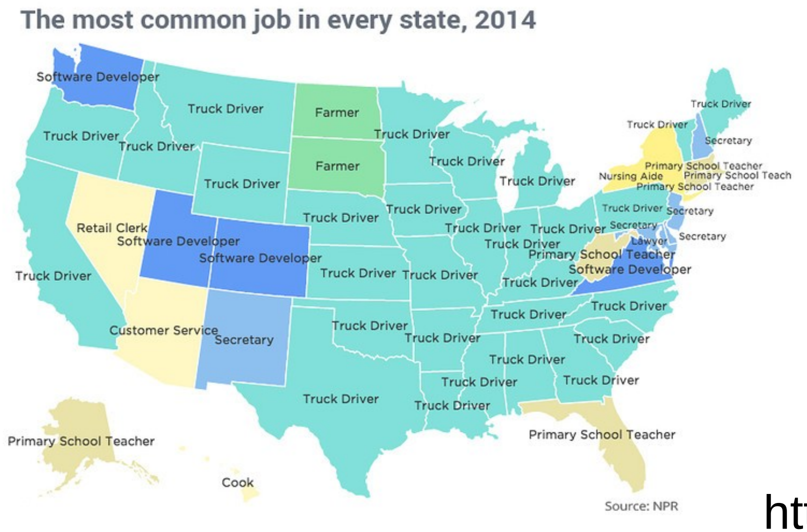

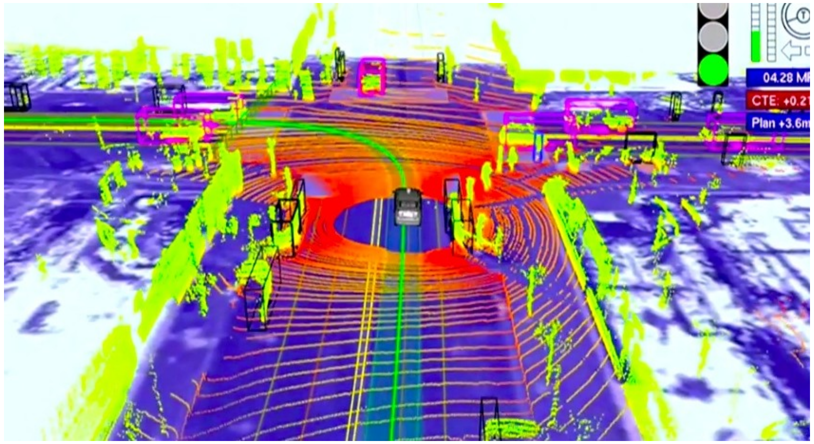

https://www.youtube.com/watch?v=O1hjwb7edlI

Jeannette Wing

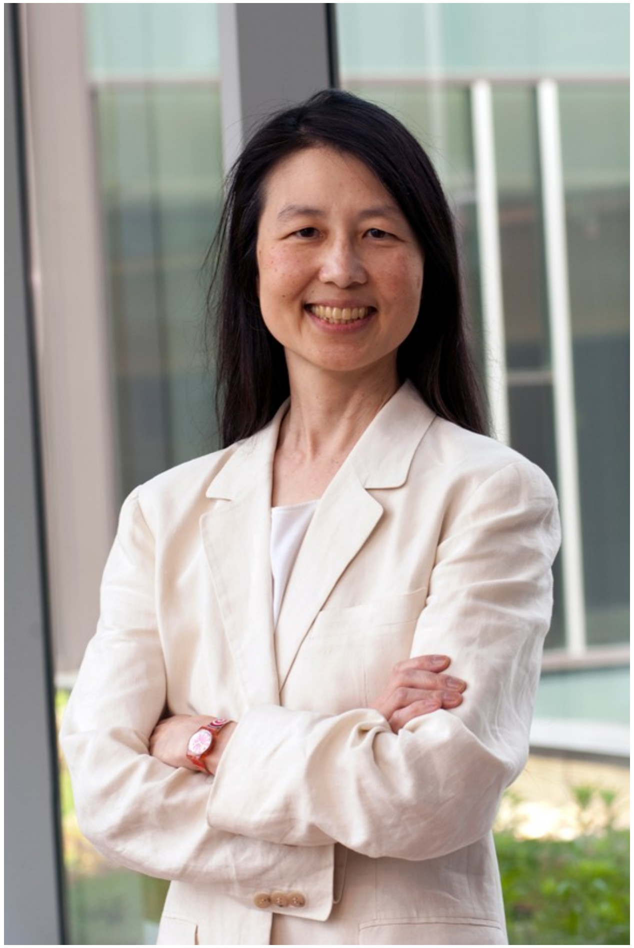

https://www.facebook.com/abciencias/videos/vb.191135830918527/1883125591719534/?type=2&theater

Nick Bostrom

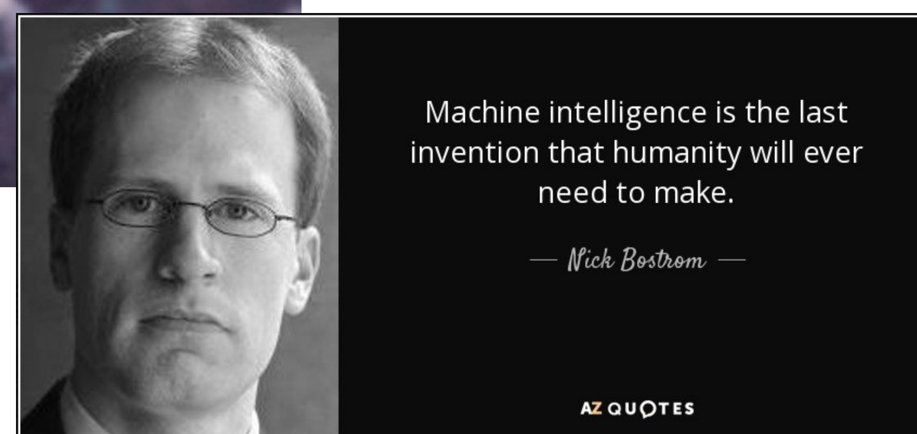

https://www.ted.com/talks/nick_bostrom_what_happens_when_our_computers_get_smarter_than_we_are

René Descartes

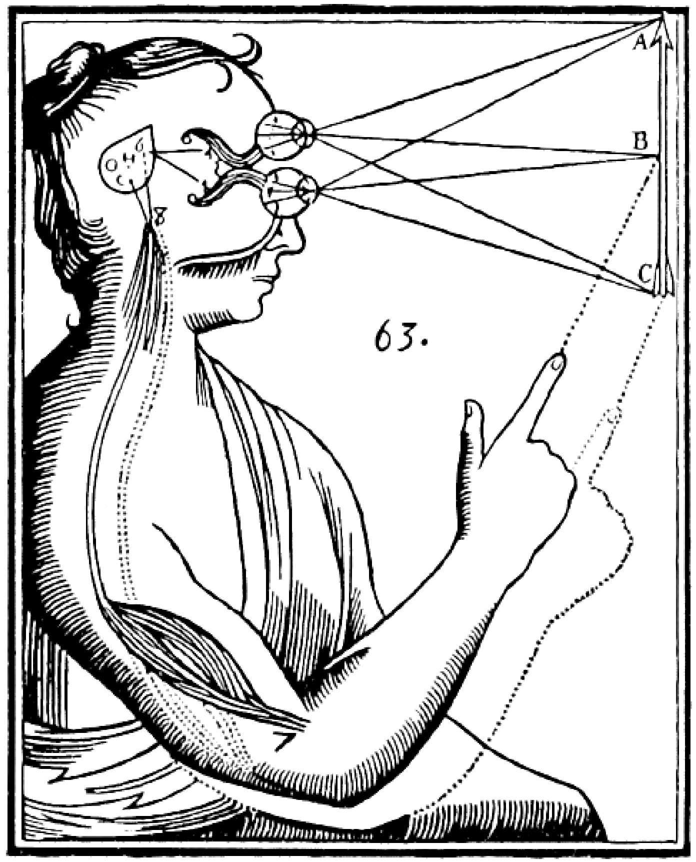

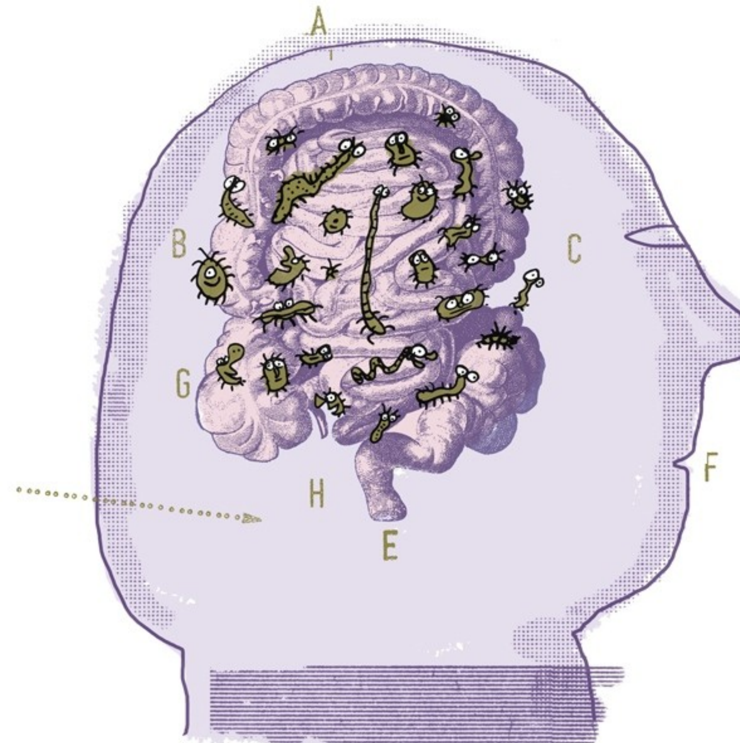

René Descartes

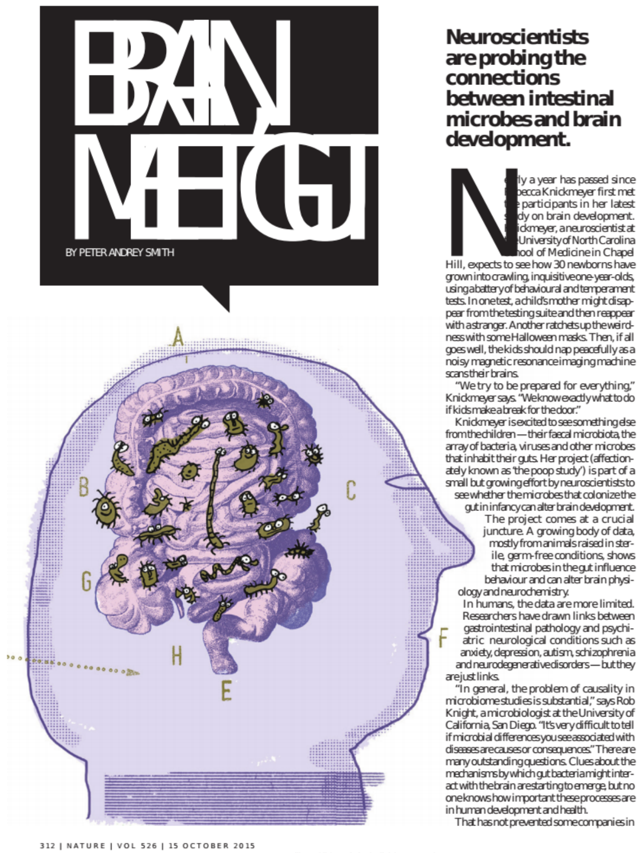

BY PETER ANDREY SMITH

Neuroscientists are probing the connections between intestinal microbes and brain development.

Nearly a year has passed since Rebecca Knickmeyer first met the participants in her latest study on brain development. Knickmeyer, a neuroscientist at the University of North Carolina School of Medicine in Chapel Hill, expects to see how 30 newborns have grown into crawling, inquisitive one-year-olds, using a battery of behavioural and temperament tests. In one test, a child's mother might disappear from the testing suite and then reappear with a stranger. Another ratchets up the weirdness with some Halloween masks. Then, if all goes well, the kids should nap peacefully as a noisy magnetic resonance imaging machine scans their brains.

"We try to be prepared for everything," Knickmeyer says. "We know exactly what to do if kids make a break for the door."

Knickmeyer is excited to see something else from the children — their faecal microbiota, the array of bacteria, viruses and other microbes that inhabit their guts. Her project (affectionately known as 'the poop study') is part of a small but growing effort by neuroscientists to see whether the microbes that colonize the gut in infancy can alter brain development.

The project comes at a crucial juncture. A growing body of data, mostly from animals raised in sterile, germ-free conditions, shows that microbes in the gut influence behaviour and can alter brain physiology and neurochemistry.

In humans, the data are more limited. Researchers have drawn links between gastrointestinal pathology and psychiatric neurological conditions such as anxiety, depression, autism, schizophrenia and neurodegenerative disorders — but they are just links.

"In general, the problem of causality in microbiome studies is substantial," says Rob Knight, a microbiologist at the University of California, San Diego. "It's very difficult to tell if microbial differences you see associated with diseases are causes or consequences." There are many outstanding questions. Clues about the mechanisms by which gut bacteria might interact with the brain are starting to emerge, but no one knows how important these processes are in human development and health.

That has not prevented some companies in

ILLUSTRATION BY SERGE BLOCH

The tantalizing links between gut microbes and the brain

Nature 2015, Vol 526

René Descartes

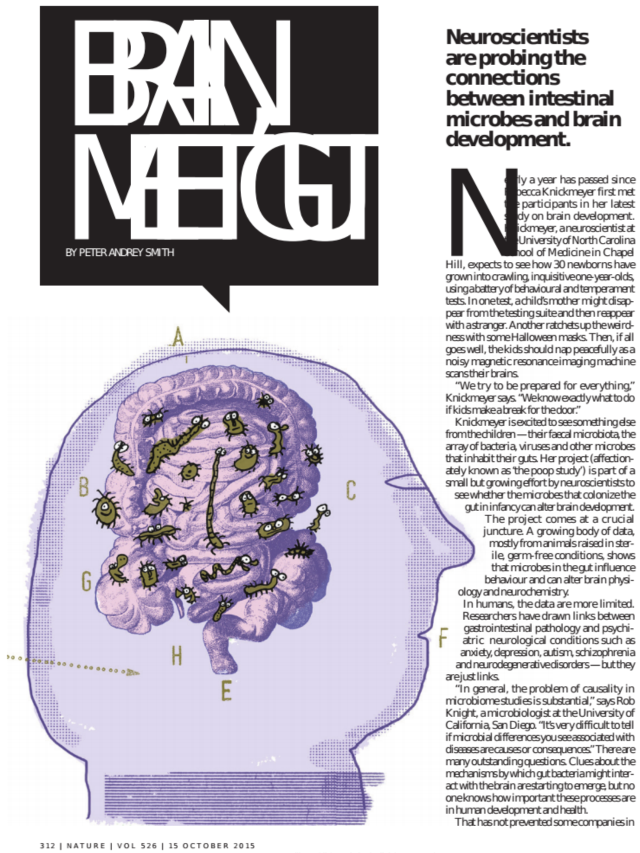

BY PETER ANDREY SMITH

Neuroscientists are probing the connections between intestinal microbes and brain development.

Nearly a year has passed since Rebecca Knickmeyer first met the participants in her latest study on brain development. Knickmeyer, a neuroscientist at the University of North Carolina School of Medicine in Chapel Hill, expects to see how 30 newborns have grown into crawling, inquisitive one-year-olds, using a battery of behavioural and temperament tests. In one test, a child's mother might disappear from the testing suite and then reappear with a stranger. Another ratchets up the weirdness with some Halloween masks. Then, if all goes well, the kids should nap peacefully as a noisy magnetic resonance imaging machine scans their brains.

"We try to be prepared for everything," Knickmeyer says. "We know exactly what to do if kids make a break for the door."

Knickmeyer is excited to see something else from the children — their faecal microbiota, the array of bacteria, viruses and other microbes that inhabit their guts. Her project (affectionately known as 'the poop study') is part of a small but growing effort by neuroscientists to see whether the microbes that colonize the gut in infancy can alter brain development.

The project comes at a crucial juncture. A growing body of data, mostly from animals raised in sterile, germ-free conditions, shows that microbes in the gut influence behaviour and can alter brain physiology and neurochemistry.

In humans, the data are more limited. Researchers have drawn links between gastrointestinal pathology and psychiatric neurological conditions such as anxiety, depression, autism, schizophrenia and neurodegenerative disorders — but they are just links.

"In general, the problem of causality in microbiome studies is substantial," says Rob Knight, a microbiologist at the University of California, San Diego. "It's very difficult to tell if microbial differences you see associated with diseases are causes or consequences." There are many outstanding questions. Clues about the mechanisms by which gut bacteria might interact with the brain are starting to emerge, but no one knows how important these processes are in human development and health.

That has not prevented some companies in

ILLUSTRATION BY SERGE BLOCH

312 | NATURE | VOL 526 | 15 OCTOBER 2015

© 2015 Macmillan Publishers Limited. All rights reserved

The tantalizing links between gut microbes and the brain

Nature 2015, Vol 526

See eye to eye!

Ricardo Marroquim

www.lcg.ufrj.br/marroquim

Laboratório de Computação Gráfica

PESC
Programa de Engenharia de Sistemas e Computação

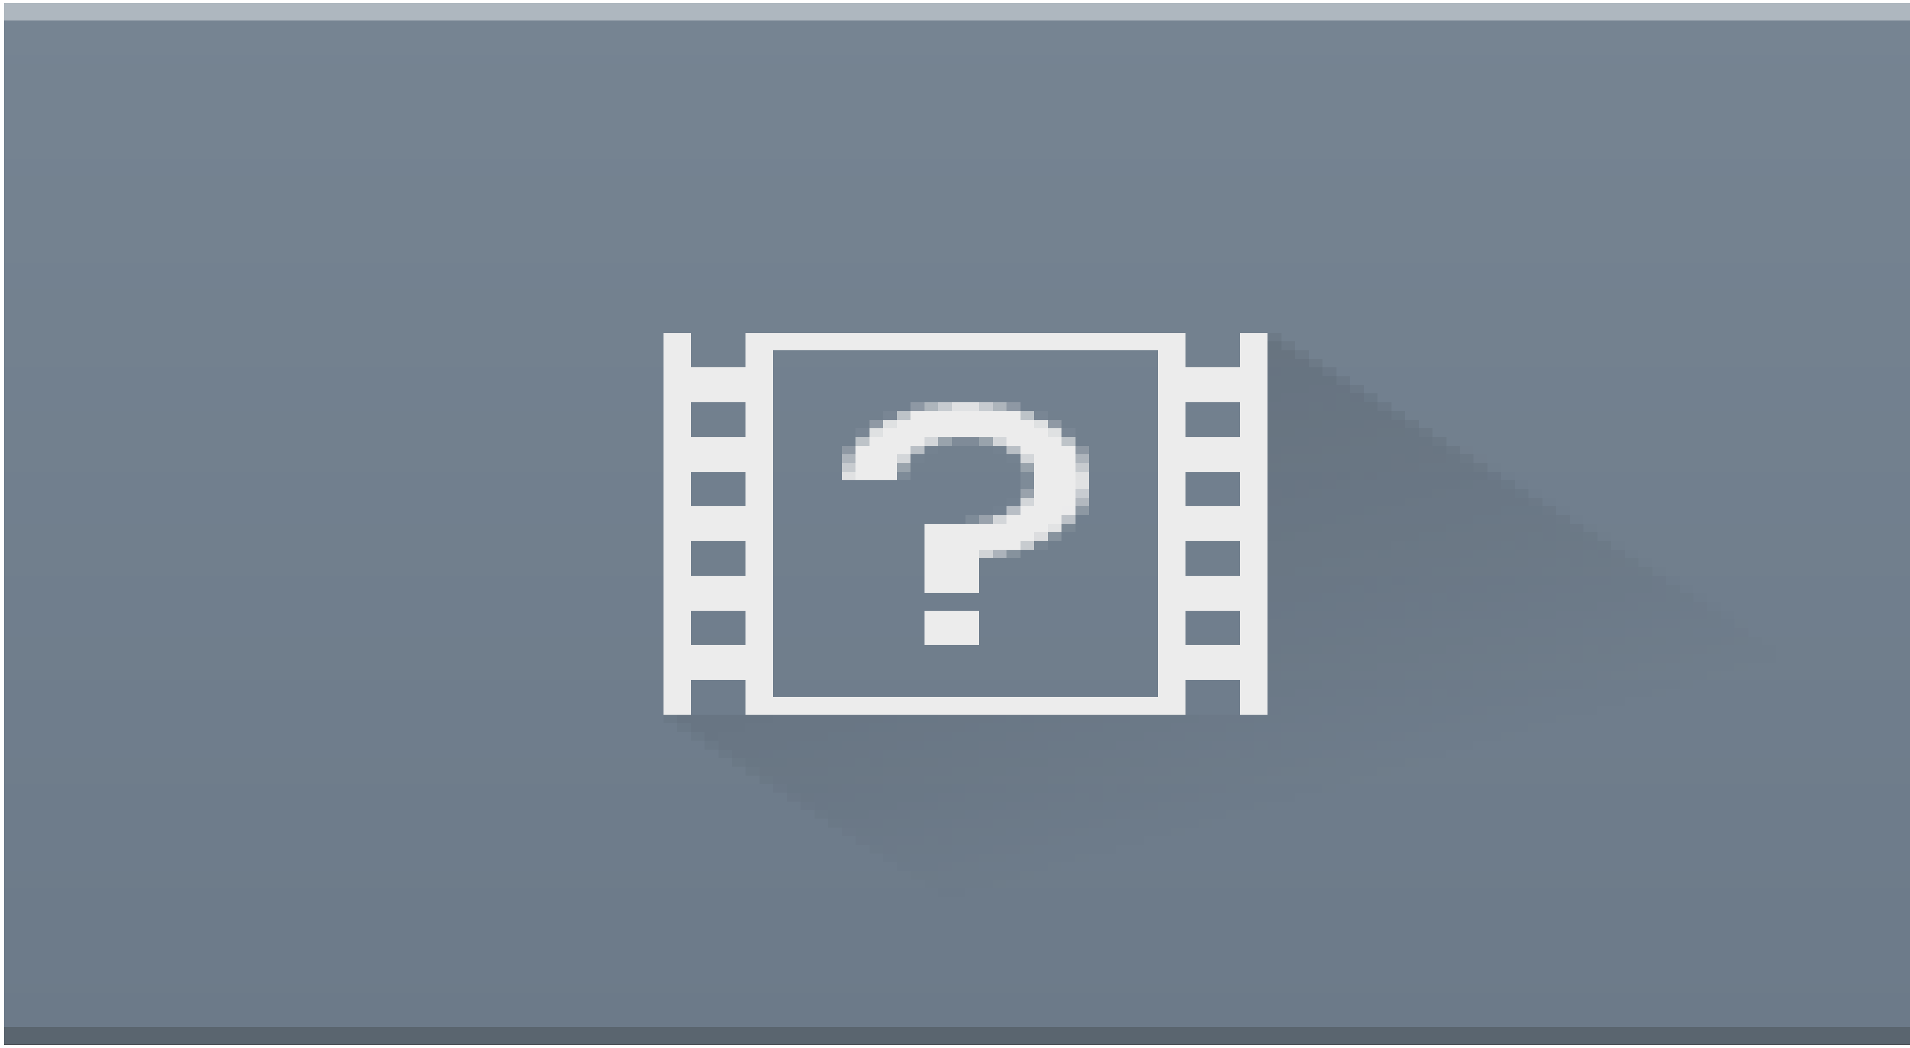

https://vimeo.com/132700334

image references 1/9

http://www.orosend.com/keep-your-eyes-healthy-nmw15/

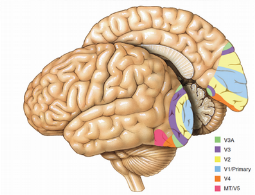

https://www.studyblue.com/notes/note/n/the-visual-system/deck/6955406

http://www.kdnuggets.com/2016/08/seven-steps-understanding-computer-vision.html

http://semiengineering.com/seeing-the-future-of-vision/

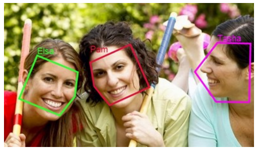

http://www.rcrwireless.com/20110723/wireless/google-buys-facial-recognition-firm-despite-privacy-concerns#prettyPhoto

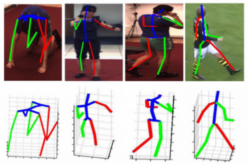

http://cvlab.epfl.ch/research/surv/human-pose-estimation

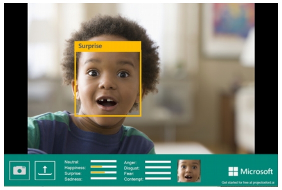

http://venturebeat.com/2015/11/11/microsoft-launches-project-oxford-apis-for-face-tracking-emotion-speaker-recognition-spell-checking/

image references 2/9

http://www.2001italia.it/2014/04/a-full-cast-list-for-2001-part-4.html

http://www.masswerk.at/minskytro/

http://xkcd.com/1425/

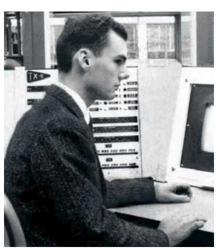

http://webdesignpi.tripod.com/roberts.htm

https://blogs.royalsociety.org/publishing/350-anniversary-issue-author-q-a-richard-morris/

image references 3/9

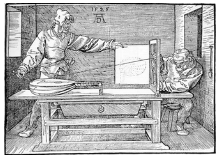

http://xahlee.info/3d/tech_drawing.html

http://b3ck.blogspot.com.br/

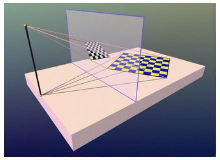

http://cgunn3.blogspot.com.br/

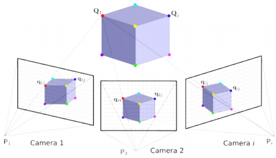

http://www.cs.cornell.edu/courses/cs4670/2013fa/lectures/lectures.html

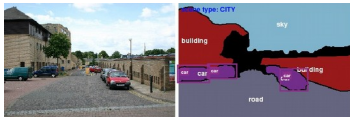

http://ttic.uchicago.edu/~yaojian/HolisticSceneUnderstanding.html

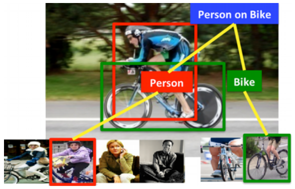

http://cs.stanford.edu/~taranlan/

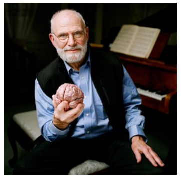

http://www.oliversacks.com/about-oliver-sacks/

image references 4/9

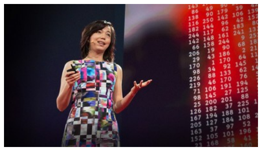

https://www.ted.com/talks/fei_fei_li_how_we_re_teaching_computers_to_understand_pictures

https://www.wired.com/2015/01/karpathy/

http://cs.stanford.edu/people/karpathy/

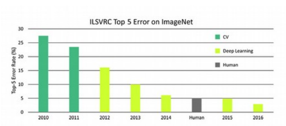

https://www.dsiac.org/resources/journals/dsiac/winter-2017-volume-4-number-1/real-time-situ-intelligent-video-analytics

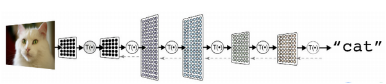

http://redcatlabs.com/2014-12-18_DeepLearning.js/img/img-to-cat_700x131.png

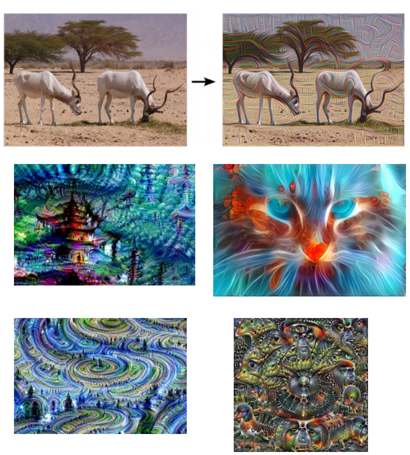

https://photos.google.com/share/AF1QipPX0SCl7OzWilt9LnuQliattX4OUCj_8EP65_cTVnBmS1jnYgsGQAieQUc1VQWdgQ?key=aVBxWjhwSzg2RjJWLWRuVFBBZEN1d205bUdEMnhB

image references 5/9

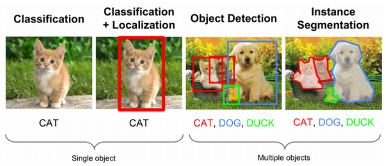

https://chaosmail.github.io/images/deep-learning/localizationVsDetection.png

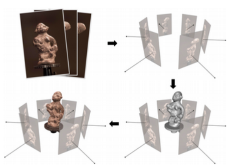

https://www.cse.wustl.edu/~furukawa/

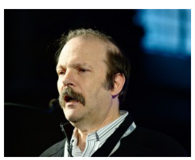

https://en.wikipedia.org/wiki/Moshe_Vardi

http://blog.caranddriver.com/driving-not-optional-google-shows-off-prototype-autonomous-car-with-no-steering-wheel-or-pedals/

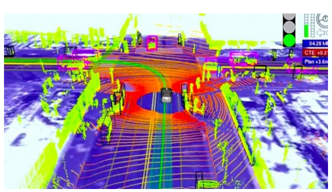

http://spectrum.ieee.org/automaton/robotics/artificial-intelligence/how-google-self-driving-car-works

http://www.marketwatch.com/story/keep-on-truckin-in-a-majority-of-states-its-the-most-popular-job-2015-02-09

image references 6/9

https://awards.acm.org/acm-prize

https://www.ted.com/talks/henry_markram_supercomputing_the_brain_s_secrets

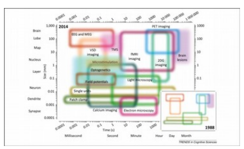

Self-portraits of the brain: cognitive science, data visualization, and communicating brain structure and function

http://www.brainybehavior.com/blog/tag/neuroimagin/

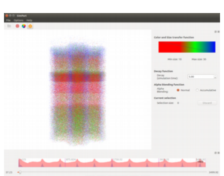

http://gmrv.es/visimp/

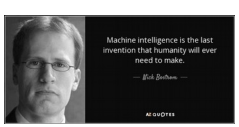

http://www.azquotes.com/author/44436-Nick_Bostrom

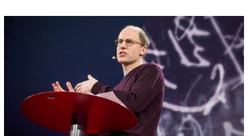

https://www.ted.com/talks/nick_bostrom_what_happens_when_our_computers_get_smarter_than_we_are

image references 7/9

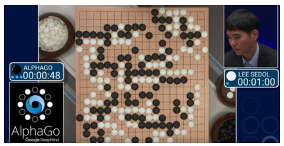

http://www.popularmechanics.com/technology/a19863/googles-alphago-ai-wins-second-game-go/

https://www.ted.com/talks/jeremy_howard_the_wonderful_and_terrifying_implications_of_computers_that_can_learn

https://cdn.slidesharecdn.com/ss_thumbnails/fastaitalk-170518062514-thumbnail-4.jpg?cb=1495091926

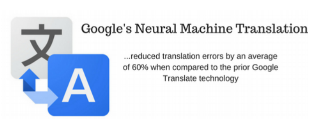

http://www.darlukman.com/wp-content/uploads/Googles-Neural-Machine-Translation.png

https://www.tensorflow.org/_static/images/tensorflow/logo.png

http://mobilemarketingmagazine.com/wp-content/uploads/2017/01/IBM-Watson-Logo.png

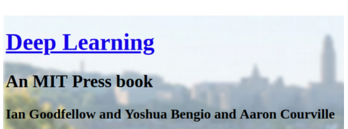

http://www.deeplearningbook.org/

image references 8/9

https://arts.mit.edu/artists/pawan-sinha/

https://industry.datascience.columbia.edu/profile/jeannette-wing

https://www.techlicious.com/blog/uber-self-driving-car-fatality-video/

By André Hatala [e.a.] (1997) De eeuw van Rembrandt, Bruxelles: Crédit communal de Belgique, ISBN 2-908388-32-4., Public Domain, https://commons.wikimedia.org/w/index.php?curid=2774313

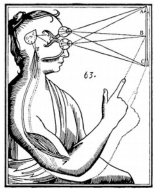

Public Domain, https://commons.wikimedia.org/w/index.php?curid=1918592

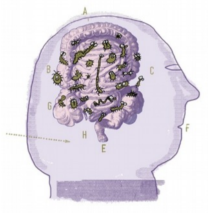

https://www.nature.com/news/the-tantalizing-links-between-gut-microbes-and-the-brain-1.18557

image references 9/9

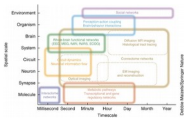 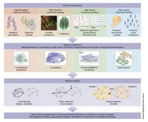

https://www.nature.com/articles/nn.4502/figures/1

https://hiveminer.com/Tags/herbsttheater/Timeline

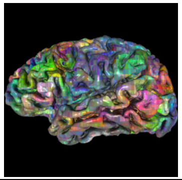

https://twitter.com/gallantlab

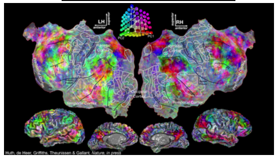

http://opentranscripts.org/transcript/engineering-thoughts-memories/

https://www.popsci.com/science/article/2011-09/mind-reading-tech-reconstructs-videos-brain-images

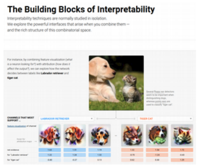

https://distill.pub/2018/building-blocks/

https://9to5google.com/guides/deepmind/

https://www.scientificamerican.com/magazine/sa/2018/03-01/

bibliography 1/2

Gray, C. M., & Singer, W. (1989). Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. Proceedings of the National Academy of Sciences of the United States of America, 86(5), 1698–1702. http://doi.org/10.1073/pnas.86.5.1698

Rauber, P. E., Fadel, S., Falcao, A., & Telea, A. (2016). Visualizing the Hidden Activity of Artificial Neural Networks. IEEE Transactions on Visualization and Computer Graphics, 1, 1–1. http://doi.org/10.1109/TVCG.2016.2598838

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, R. F. (2014). Intriguing properties of neural networks. In International Conference on Learning Representations. http://doi.org/10.1021/ct2009208

Nguyen, A., Yosinski, J., & Clune, J. (2015). Deep Neural Networks are Easily Fooled. In Computer Vision and Pattern Recognition, 2015 IEEE Conference on (pp. 427–436). http://doi.org/10.1109/CVPR.2015.7298640

Fiorani, M., Azzi, J. C. B., Soares, J. G. M., & Gattass, R. (2014). Automatic mapping of visual cortex receptive fields: A fast and precise algorithm. Journal of Neuroscience Methods, 221, 112–126. http://doi.org/10.1016/j.jneumeth.2013.09.012

Rauber, P. E., Silva, R. R. O., Feringa, S., Celebi, M. E., Falcão, A. X., & Telea, A. C. (2015). Interactive Image Feature Selection Aided by Dimensionality Reduction. In EuroVis Workshop on Visual Analytics (pp. 2–6). http://doi.org/10.2312/eurova.20151098

Silva, R. R. O., Rauber, P. E., Martins, R. M., Minghim, R., & Telea, A. C. (2015). Attribute-based Visual Explanation of Multidimensional Projections. In EuroVis Workshop on Visual Analytics. http://doi.org/10.2312/eurova.20151100

Ostrovsky, Y., Meyers, E., Ganesh, S., Mathur, U., & Sinha, P. (2009). Visual parsing after recovery from blindness. Psychological Science, 20(12), 1484–1491. http://doi.org/10.1111/j.1467-9280.2009.02471.x

Pinto, N., Cox, D. D., & DiCarlo, J. J. (2008). Why is real-world visual object recognition hard? PLoS Computational Biology, 4(1), 0151–0156. http://doi.org/10.1371/journal.pcbi.0040027

Cox, D. D., & Dean, T. (2014). Neural networks and neuroscience-inspired computer vision. Current Biology, 24(18), R921–R929. http://doi.org/10.1016/j.cub.2014.08.026

Roberts, L. Gi. (1965). Machine perception of three-dimensional solids. PhD Thesis, (November), 159–197. Retrieved from http://oai.dtic.mil/oai/oai?verb=getRecord&metadataPrefix=html&identifier=AD0413529

Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., … Fei-Fei, L. (2015). ImageNet Large Scale Visual Recognition Challenge. International Journal of Computer Vision, 115(3), 211–252. http://doi.org/10.1007/s11263-015-0816-y

Johnson, J., Karpathy, A., & Fei-Fei, L. (2016). DenseCap: Fully Convolutional Localization Networks for Dense Captioning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Pattern Recognition. http://doi.org/10.1109/CVPR.2016.494

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. In Advances In Neural Information Processing Systems (pp. 1–9). http://doi.org/http://dx.doi.org/10.1016/j.protcy.2014.09.007

bibliography 2/2

Doersch, Carl and Singh, Saurabh and Gupta, Abhinav and Sivic, Josef and Efros, Alexei A. (2012),What Makes Paris Look Like Paris?. ACM Transactions on Graphics, vol 31 num 4, http://doi.acm.org/10.1145/2185520.2185597

James Hays, Alexei A. Efros. (2007). Scene Completion Using Millions of Photographs. ACM Transactions on Graphics (SIGGRAPH 2007), vol. 26, No. 3. http://doi.acm.org/10.1145/1276377.1276382

Richard Zhang, Phillip Isola, Alexei A. Efros. (2016). European Conference on Computer Vision. https://link.springer.com/chapter/10.1007%2F978-3-319-46487-9_40

Lichtman, Jeff W. and Denk, Winfried. (2011). The Big and the Small: Challenges of Imaging the Brain's Circuits. Science. vol. 334, No. 6056. http://science.sciencemag.org/content/334/6056/618

Jeremy Kawahara, Colin J. Brown, Steven P. Miller, Brian G. Booth, Vann Chau, Ruth E. Grunau, Jill G. Zwicker, Ghassan Hamarneh. (2017). BrainNetCNN: Convolutional neural networks for brain networks; towards predicting neurodevelopment. NeuroImage, vol 146. https://doi.org/10.1016/j.neuroimage.2016.09.046

Robert L. Goldstone, Franco Pestilli, and Katy Börner. (2015). Self-portraits of the brain: cognitive science, data visualization, and communicating brain structure and function. Trends in Cognitive Science, vol 19, No 8. http://dx.doi.org/10.1016/j.tics.2015.05.012

Kexin Pei, Yinzhi Cao, Junfeng Yang, Suman Jana (2017), DeepXplore: Automated Whitebox Testing of Deep Learning Systems, SOSP '17 Proceedings of the 26th Symposium on Operating Systems Principles

Yuri Ostrovsky, Ethan Meyers, Suma Ganesh, Umang Mathur, Pawan Sinha, (2009) Visual Parsing After Recovery From Blindness, Psychol Sci. 2009 Dec;20(12)

Danielle S Bassett, Olaf Sporns (2017), Network neuroscience, Nature Neuroscience vol 20, pp 353–364