Machine Learning Neurosciences

Michel Dojat

Machine Learning & Neurosciences

Michel Dojat

michel.dojat@inserm.fr

Artificial Intelligence

Intelligent Agent: an entity that takes the best possible action in a situation

- How to build an artificial intelligent agent?
- Test our models of natural intelligent agents?

Computer science

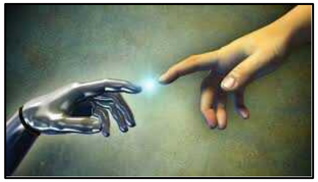

Neuroscience

Integration of heterogeneous datasets

Management of large repositories of data & knowledge

Knowledge discovery

Computer & Brain

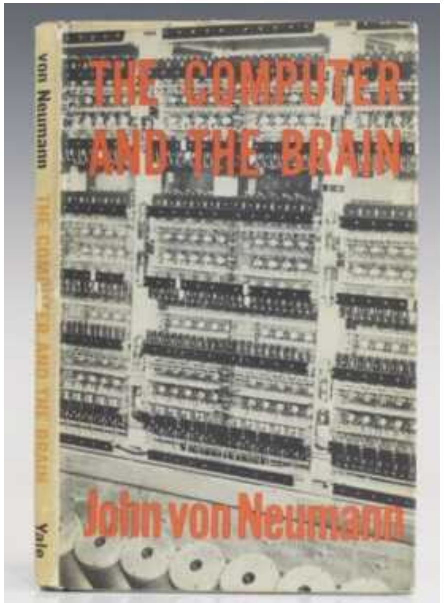

1958
Yale University Press, New Haven

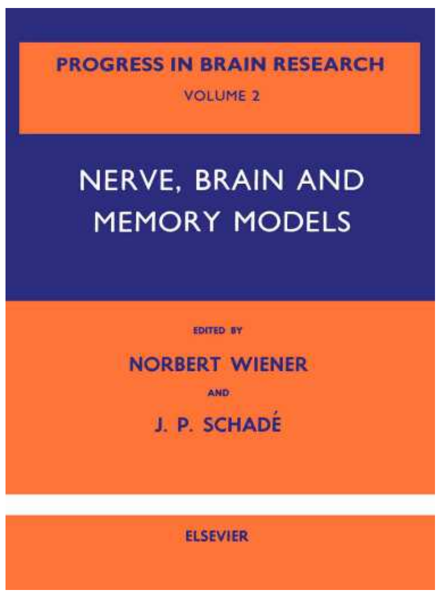

1963

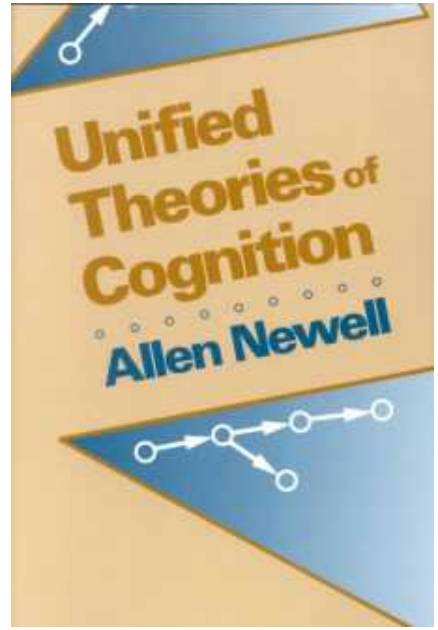

1982

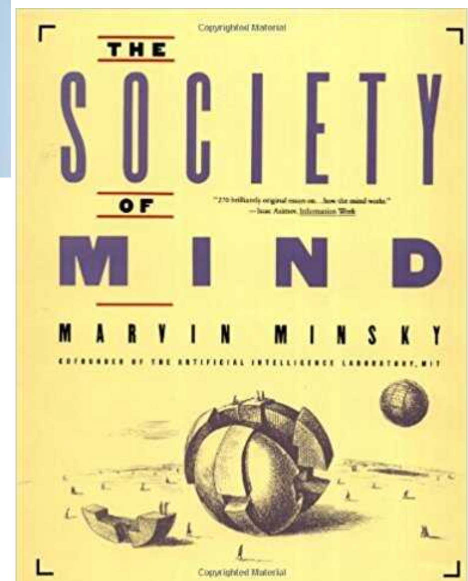

1988

Two main approaches

- Machine Learning
  - Bio-inspired
    - Artificial life
    - Neural Networks
      - WMcCulloch & W Pitts (1943)
        Artificial neurons
      - D Hebb (1949)
        Learning by modification of connections
      - F Rosenblatt (1963)
        Convergence theorem
      - M Minsky & S Paper
        Perceptrons (1969)
  - Classification (SVM,…)

Operations on large vectors

- Symbolic Processing
  - Problem-solving
  - Planning
  - Logic
  - Knowledge representation
    - Common knowledge
    - Meta-knowledge
    - Ontology
  - Multi-agents
  - Co-construction

Rule-based manipulation of symbols

But ...

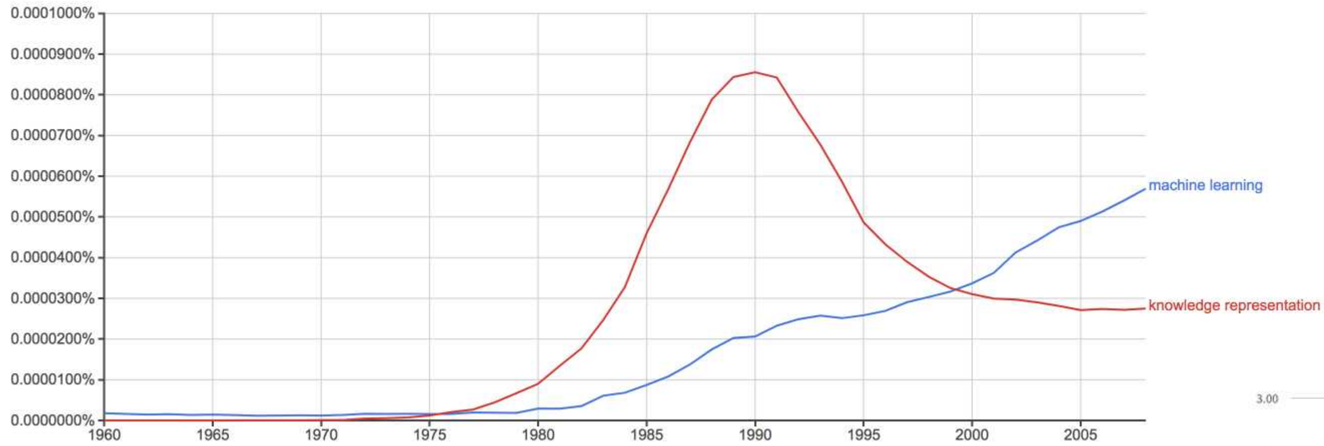

0.0001000%
0.0000900%
0.0000800%
0.0000700%
0.0000600%
0.0000500%
0.0000400%
0.0000300%
0.0000200%
0.0000100%
0.0000000%
1960
1965
1970
1975
1980
1985
1990
1995
2000
2005
machine learning
knowledge representation

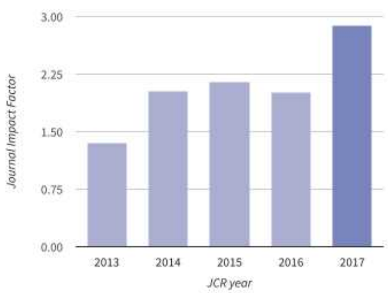

3.00
2.25
1.50
0.75
0.00
Journal Impact Factor
2013
2014
2015
2016
2017
JCR year

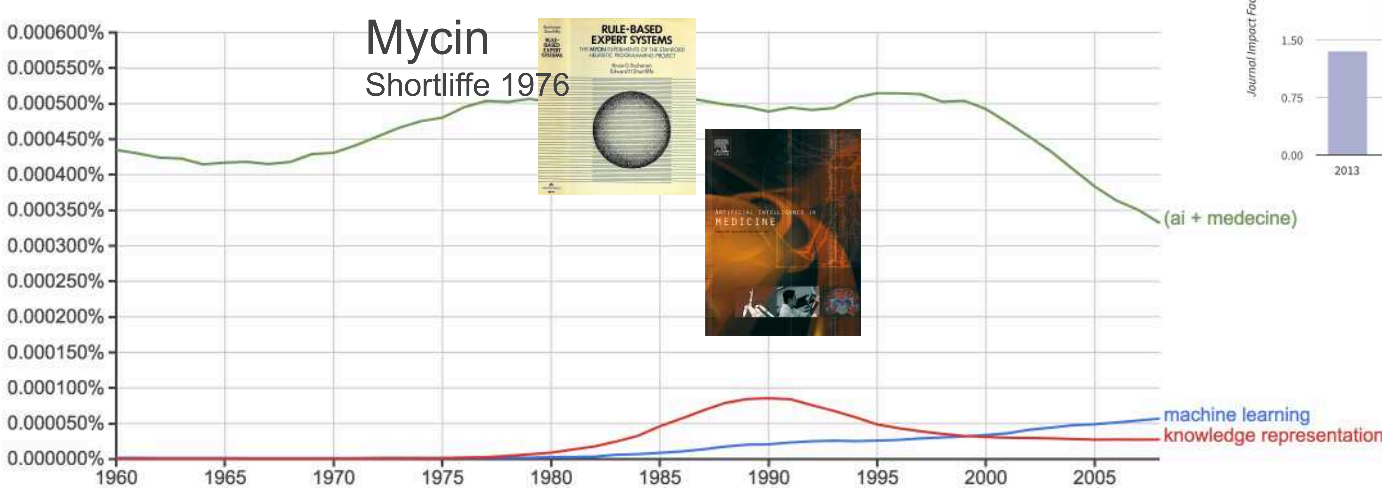

Mycin
Shortliffe 1976
RULE-BASED EXPERT SYSTEMS
ARTIFICIAL INTELLIGENCE IN MEDICINE
0.000600%
0.000550%
0.000500%
0.000450%
0.000400%
0.000350%
0.000300%
0.000250%
0.000200%
0.000150%
0.000100%
0.000050%
0.000000%
1960
1965
1970
1975
1980
1985
1990
1995
2000
2005
(ai + medecine)
machine learning
knowledge representation

ML: Natural extension of traditional statistical approaches

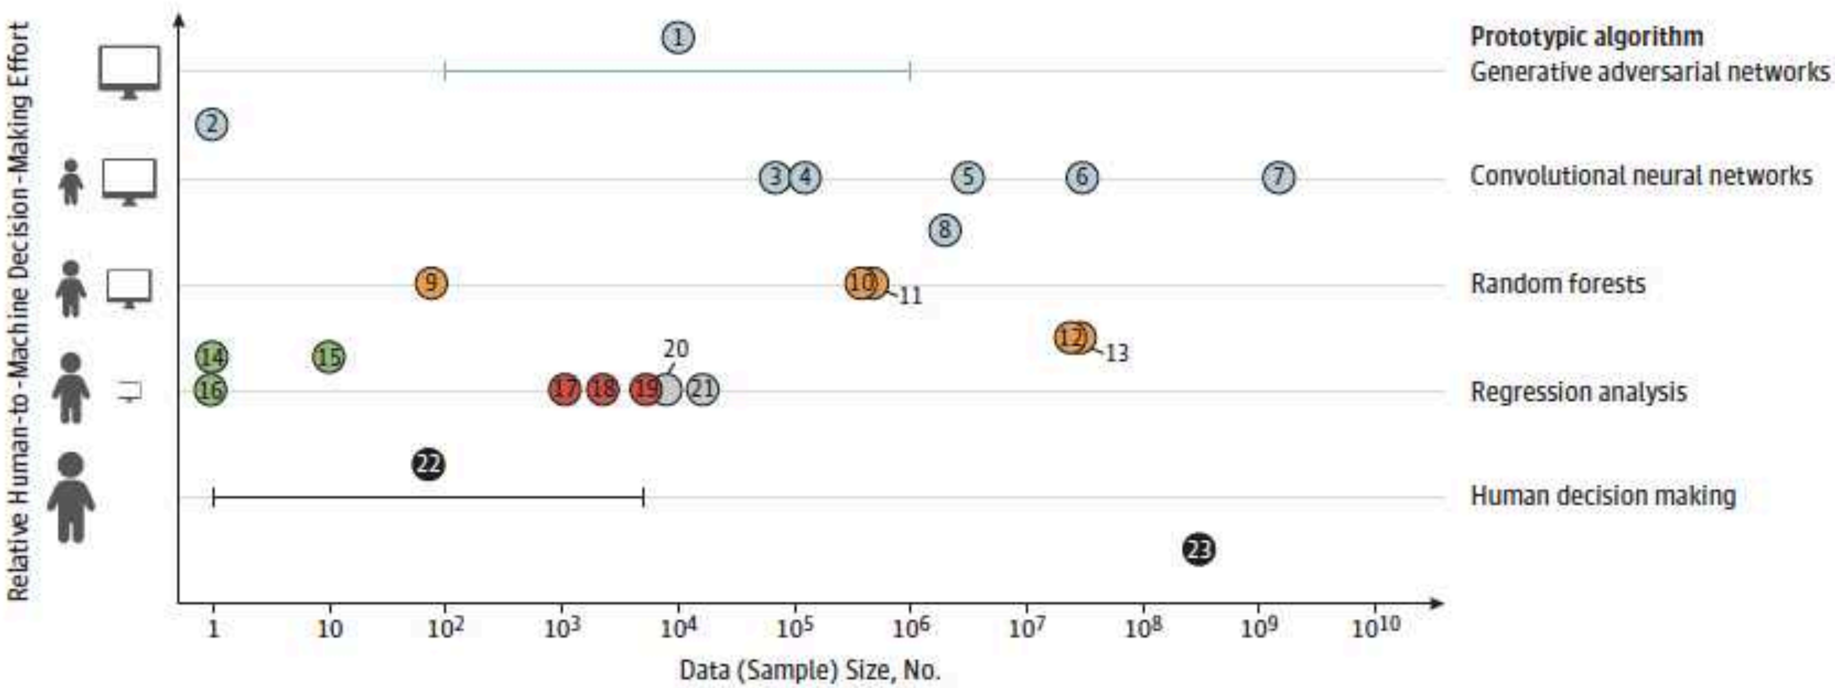

Deep learning

Classic machine learning

Risk calculators

Knowledge-Based systems

Other

Beam & Kohane Nature 2018

ML Approaches

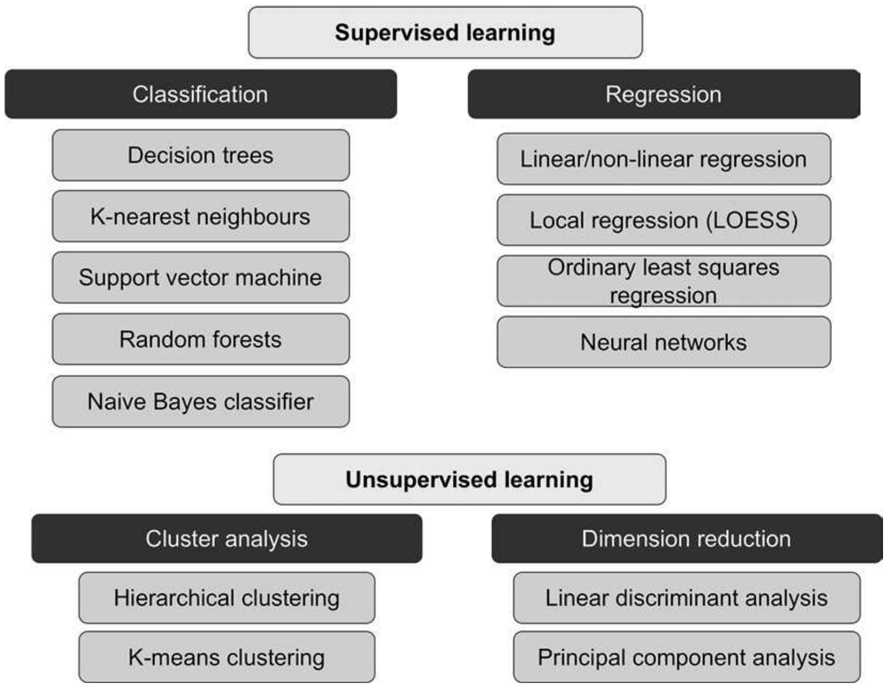

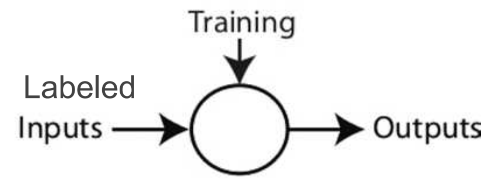

Learns known patterns
Predicts outcome

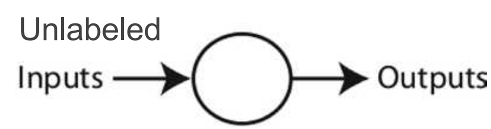

Learns unknown patterns
Find hidden patterns

Reinforcement learning

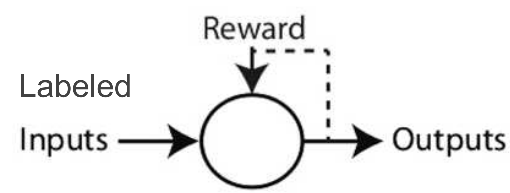

Generates data
Interacts with environment

Transfert learning al

Applies learned patterns to a different but related task

[From Choy et al. Radiology 2017]

ML Approaches

Supervised learning

- Classification
  - Decision trees
  - K-nearest neighbours
  - Support vector machine
  - Random forests
  - Naive Bayes classifier
- Regression
  - Linear/non-linear regression
  - Local regression (LOESS)
  - Ordinary least squares regression
  - Neural networks

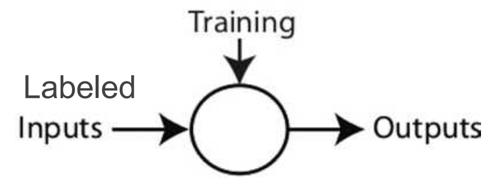

Learns known patterns
Predicts outcome

Unsupervised learning

- Cluster analysis
  - Hierarchical clustering
  - K-means clustering
- Dimension reduction
  - Linear discriminant analysis
  - Principal component analysis

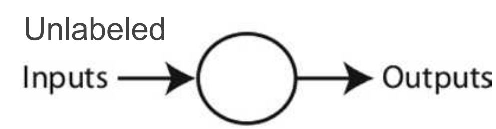

Learns unknown patterns
Find hidden patterns

Reinforcement learning

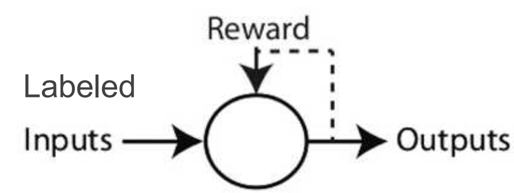

Generates data
Interacts with environment

Transfert learning

Applies learned patterns to a different but related task

[From Choy et al. Radiology 2017]

Neurosciences - ML Cross fertilization

- The Human Visual System
  - Encoding / Decoding
- Deep Learning in NN
  - CNN as Models of the Visual System
- Discussion - Perspectives

Mammalian visual system

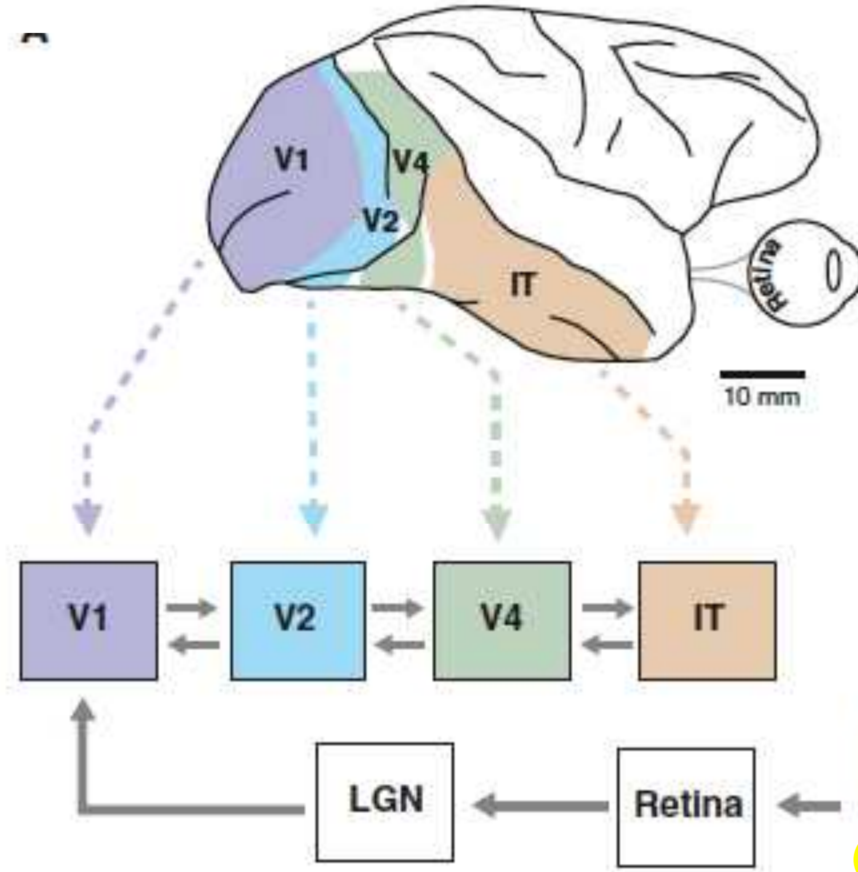

Simultaneous and parallel computing

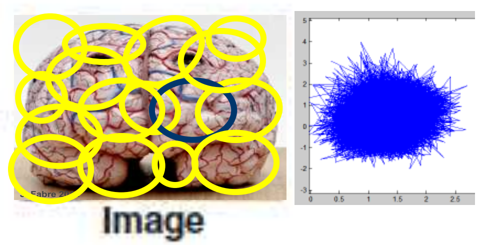

[adapted from Cox & Dean curr bio 2014]

Human visual system

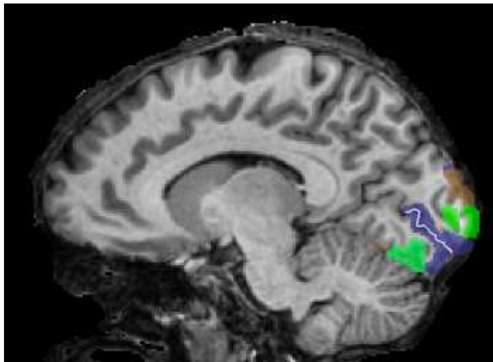

CORTEX

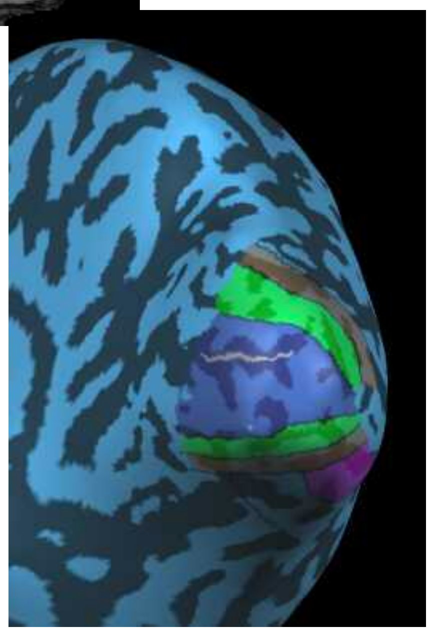

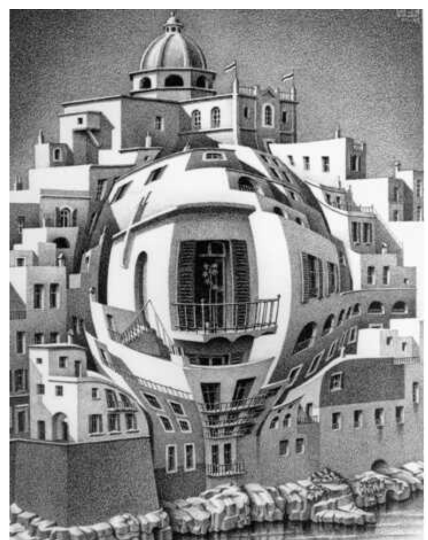

[Escher, Balcony 1945]

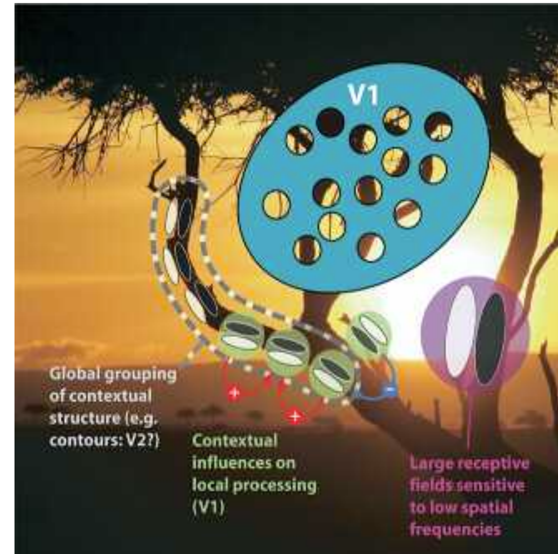

[Dakin & Frith Neuron 2005]

Brain connectivity

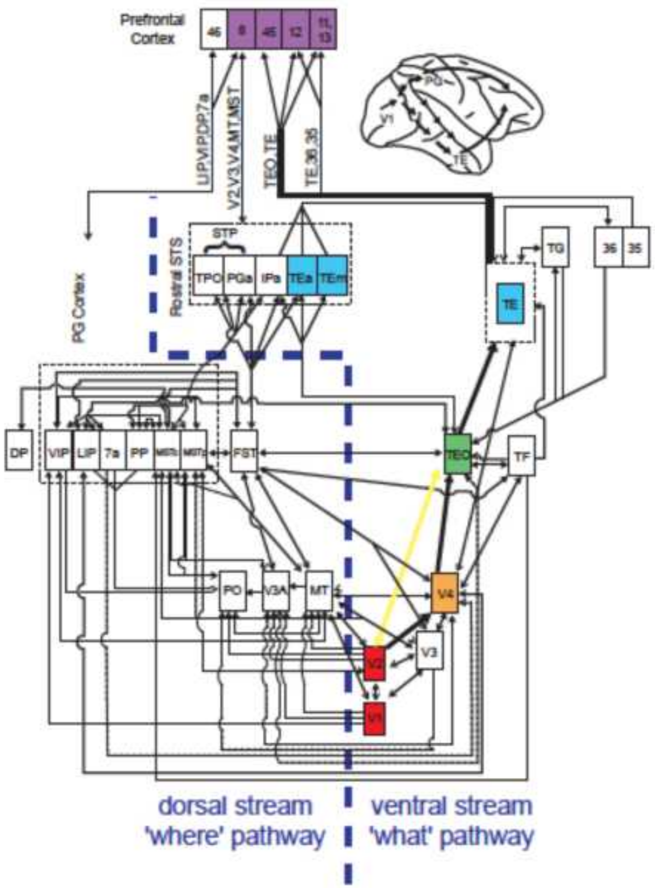

[Serre et al 2005 Tech Report]

- Hierarchical model
- Several areas interconnected (layers)
- Receptive field sizes increasing
  - Local features
  - Global grouping

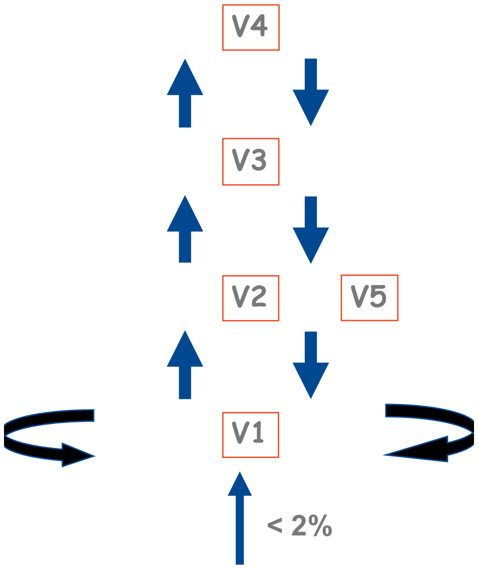

MAX pooling

MAX pooling: invariance to scale and translation; key mechanism for object recognition

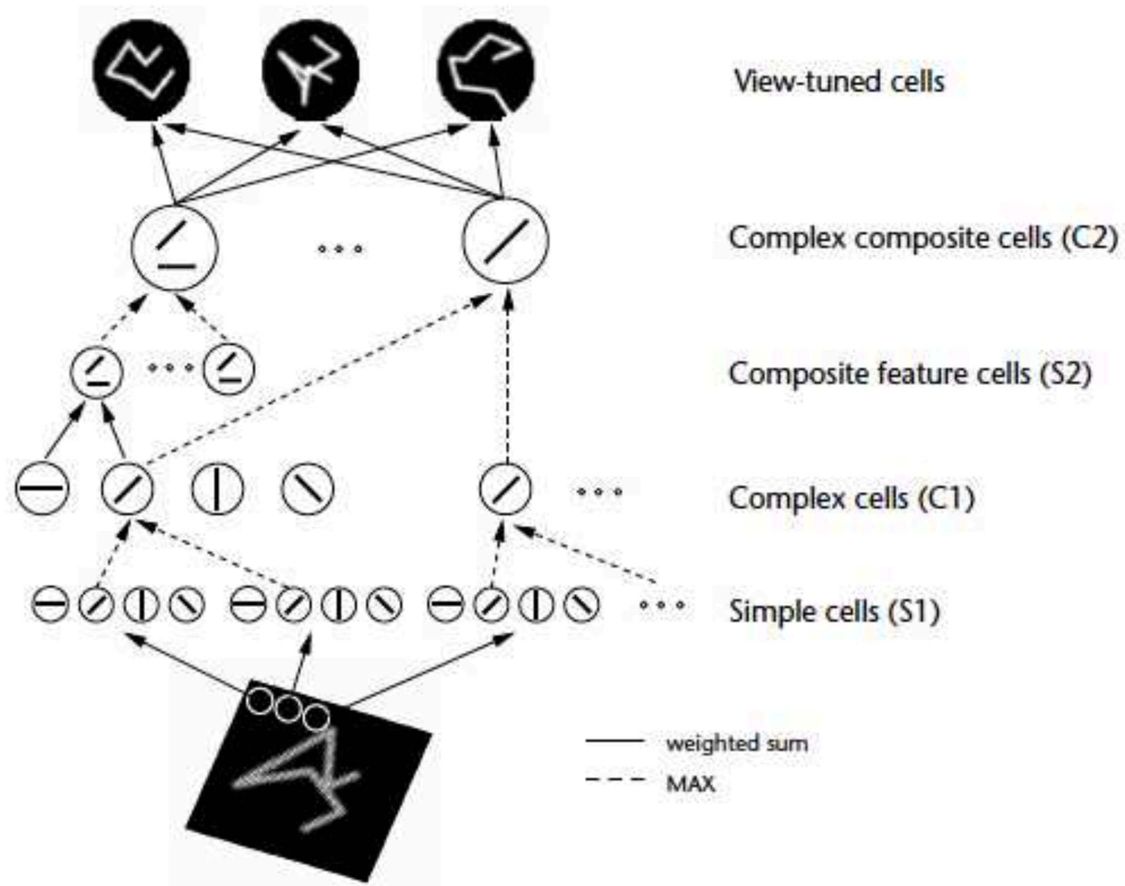

[Riesenhuber & Poggio Nature 1999]

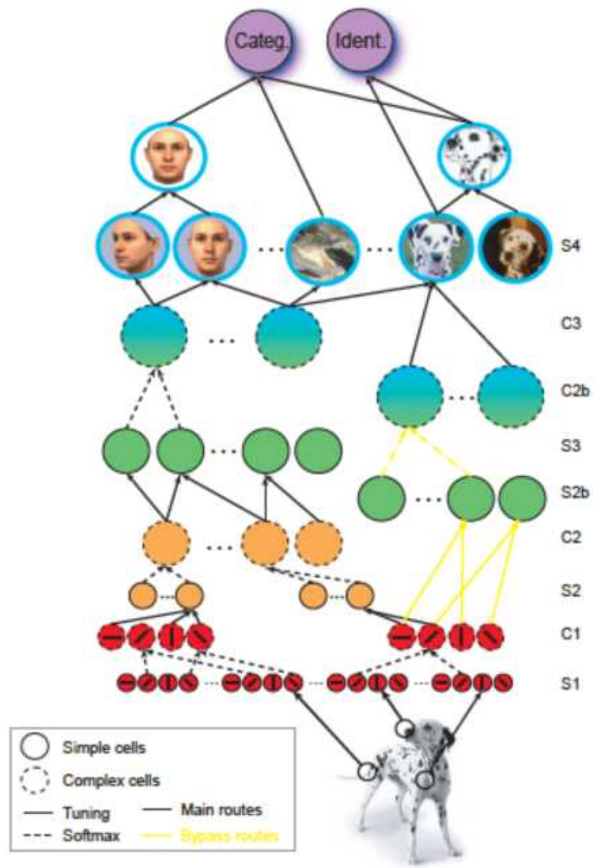

[Serre et al 2005 Tech Report]

Visual representation in the Human Brain

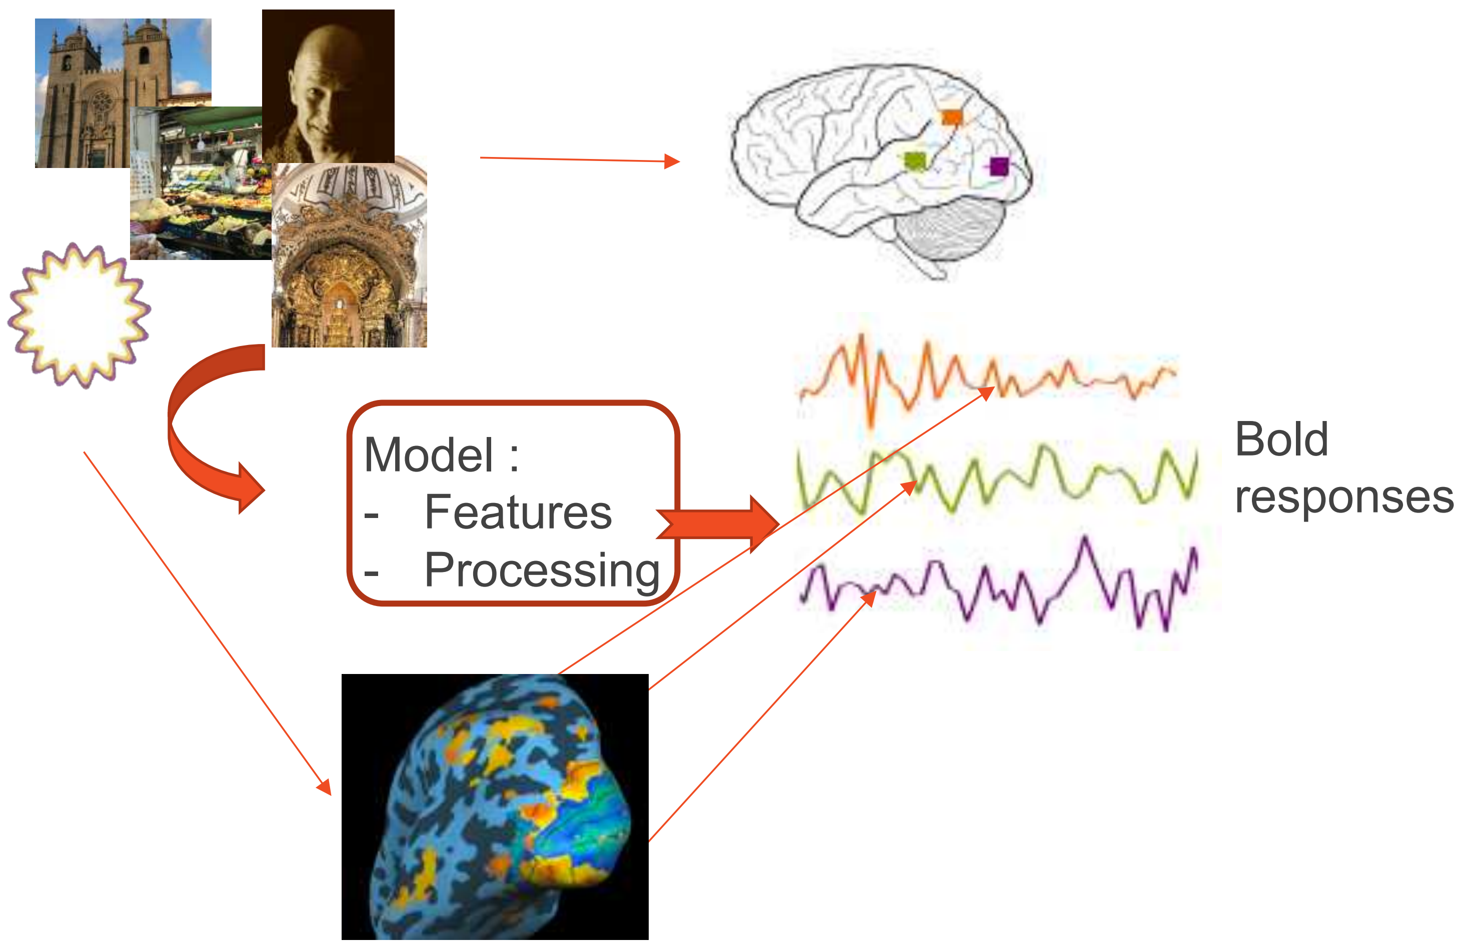

Functional imaging: Encoding models

Univariate Voxel Analysis

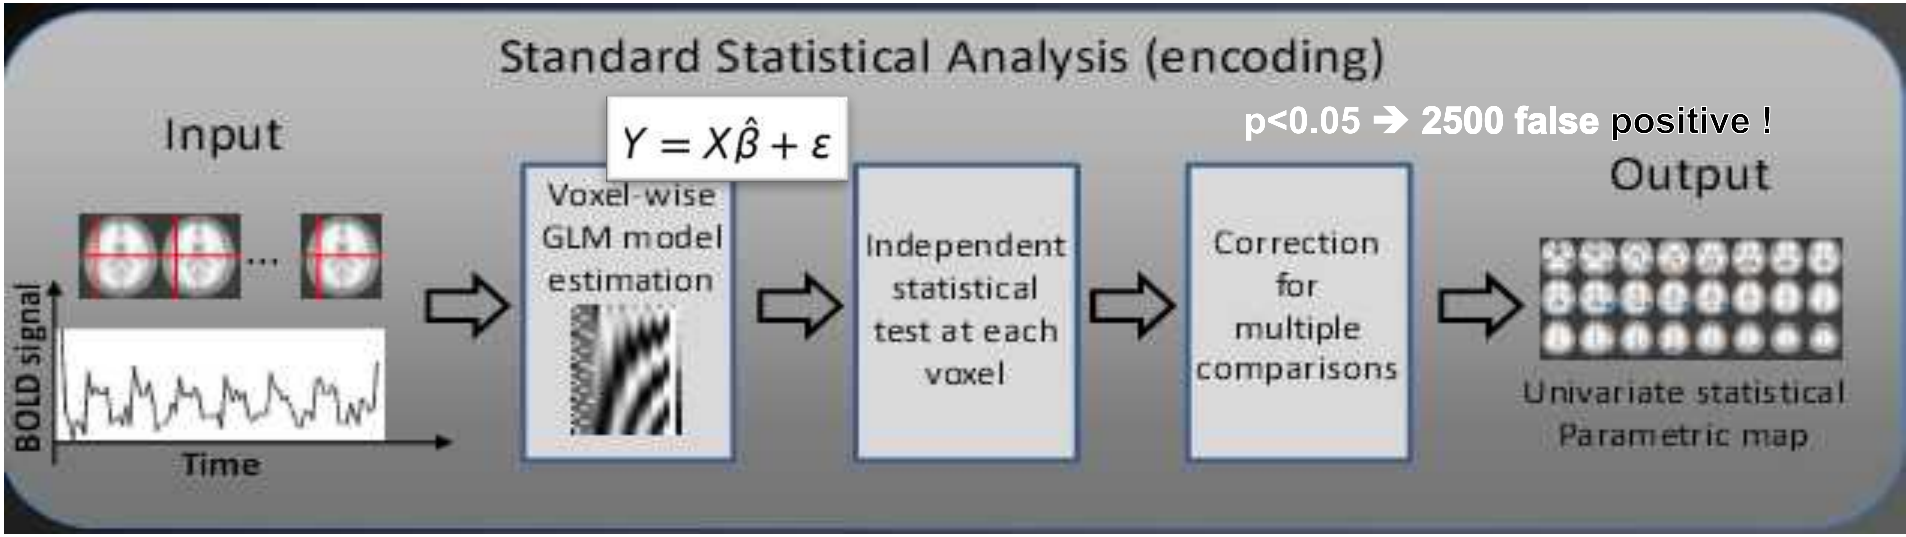

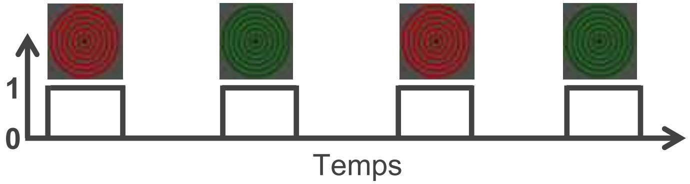

Some examples

Localizers

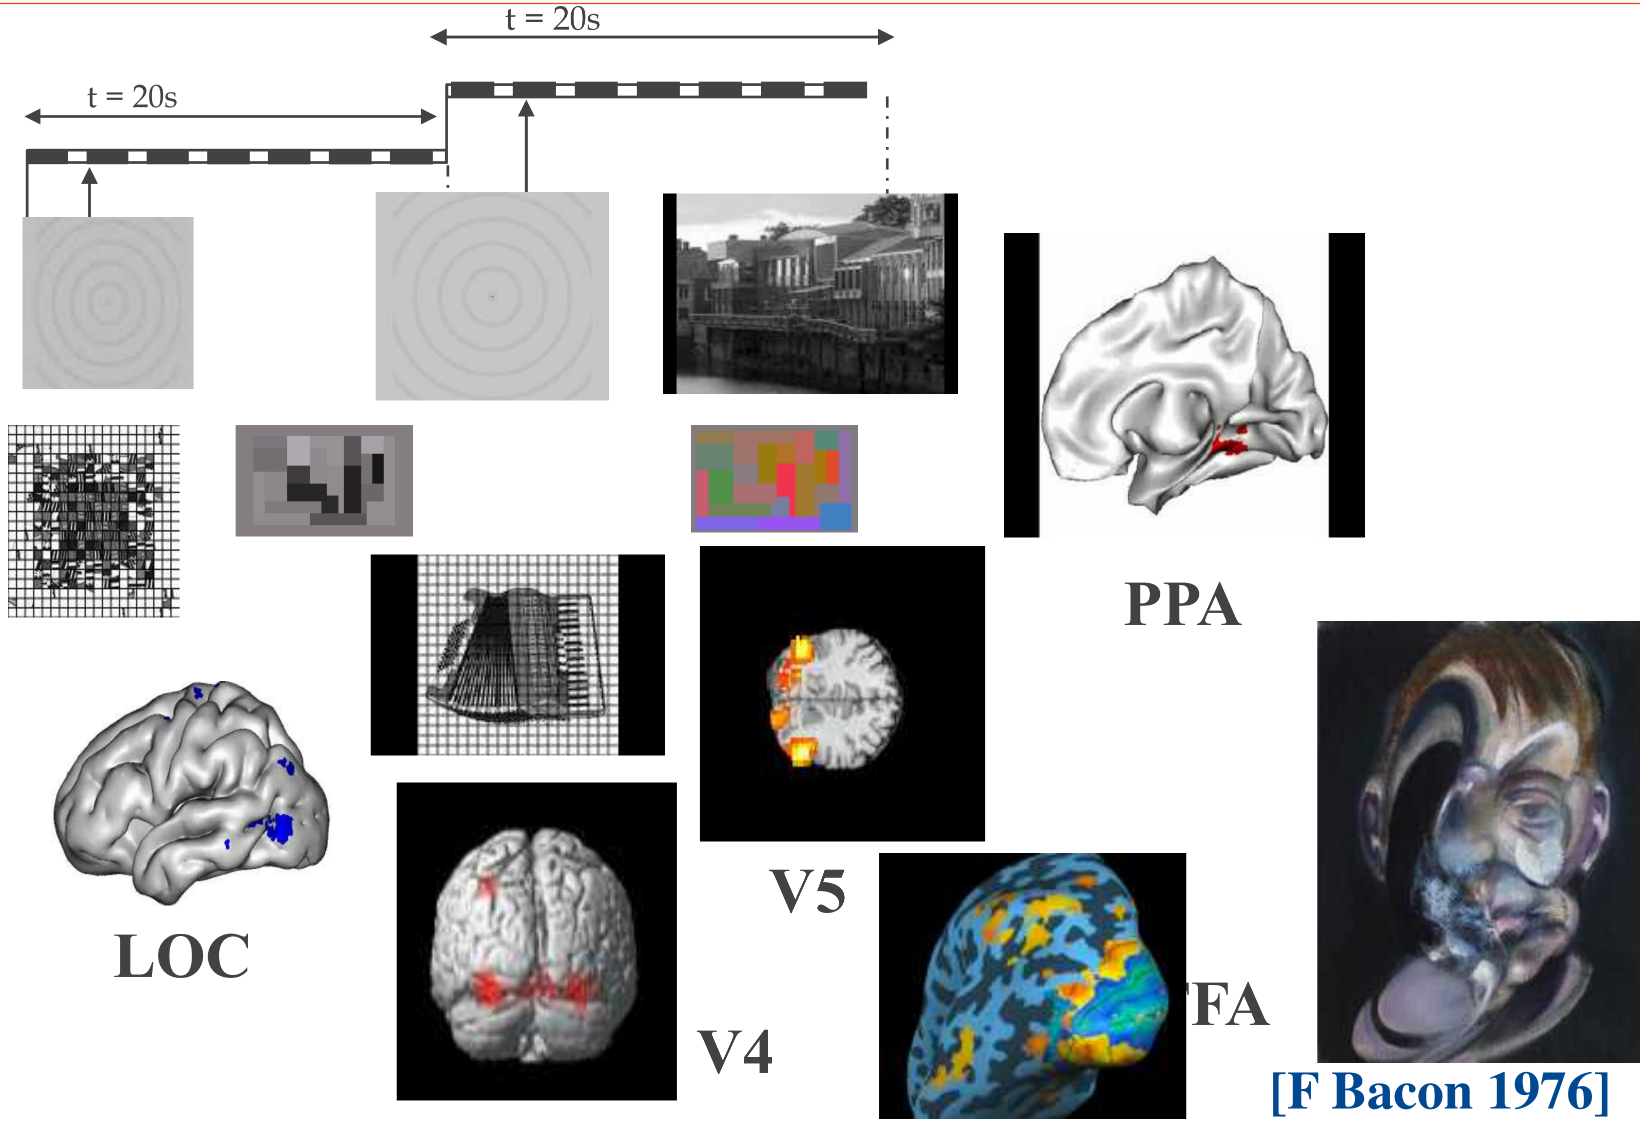

[F Bacon 1976]

Cortical activation – Space localisation

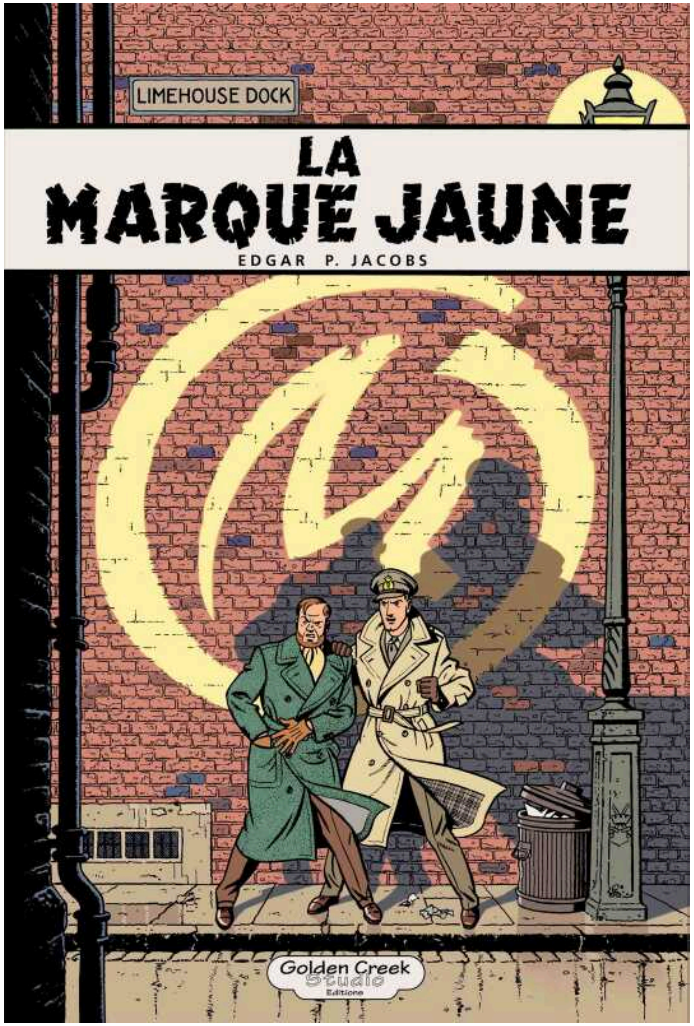

Cortical activation – Space localisation

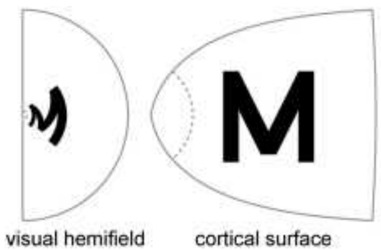

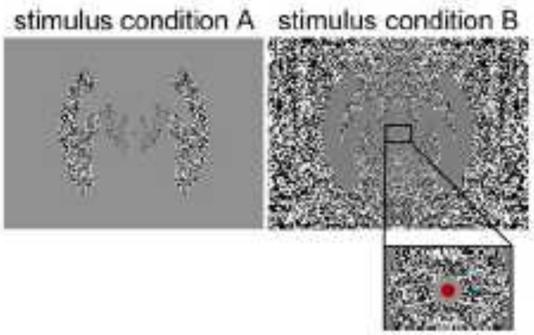

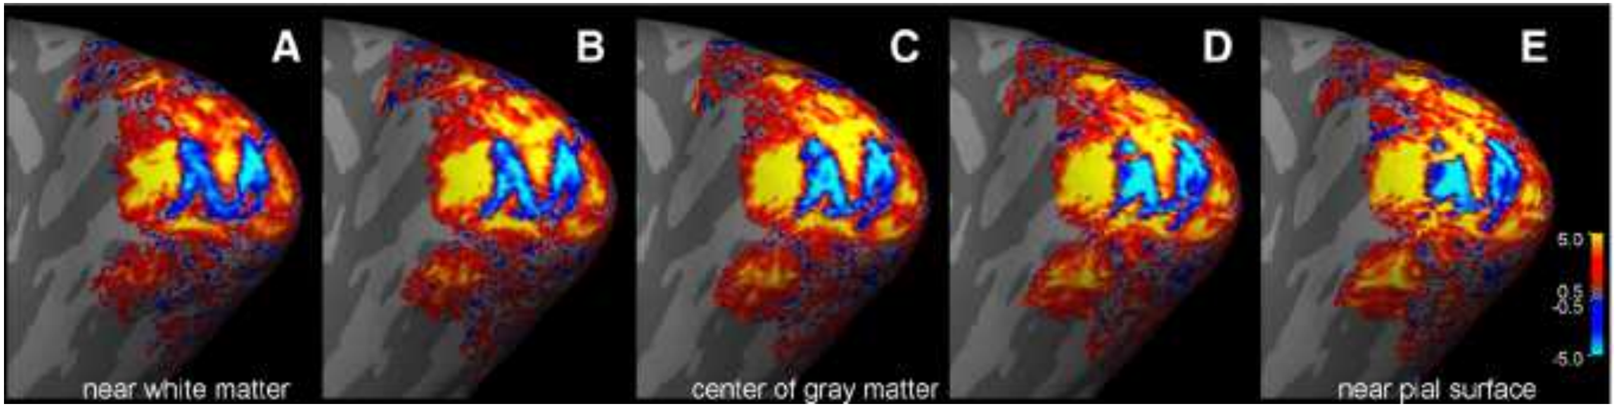

[Polimeni et al. NeuroIm 2010]

Pattern recognition

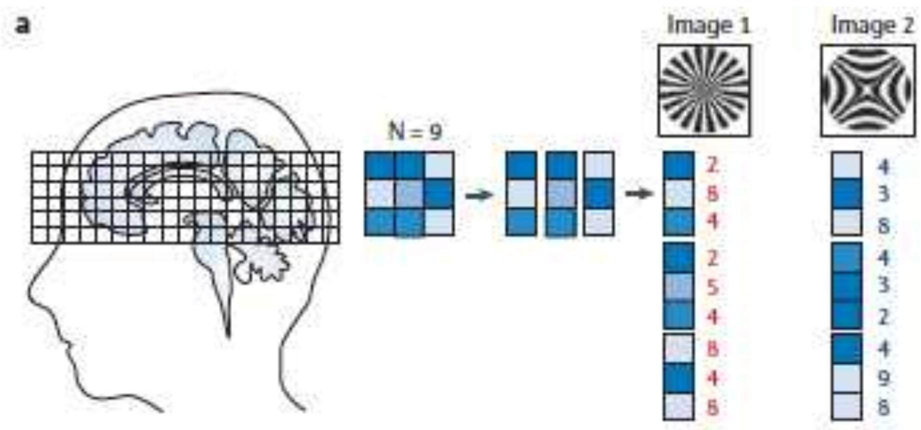

[Haynes and Rees Nat Neuro 2006]

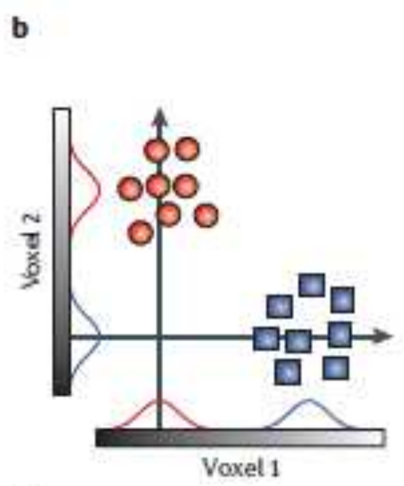

Pattern recognition

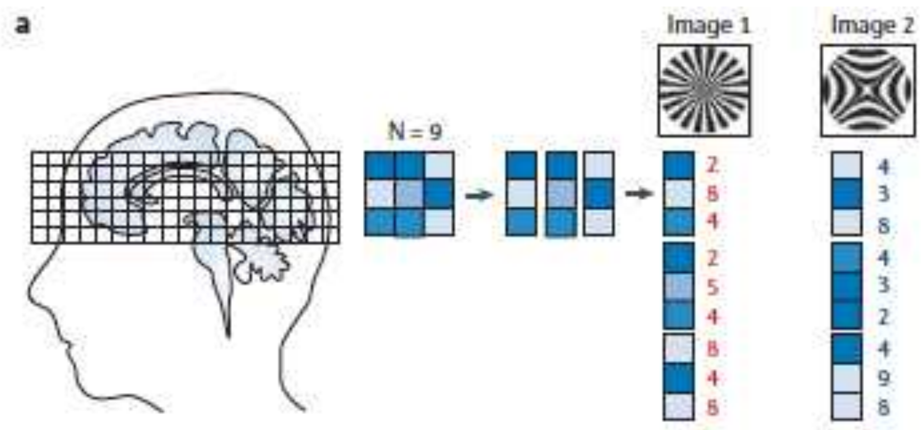

$$f(x) = w + b$$
$$f(x^*) = wx^* + B$$
$$\mathbf{w} = \sum_{i=1}^{N} \alpha_i \mathbf{x}_i$$

Risque d'overfitting

[Haynes and Rees Nat Neuro 2006]

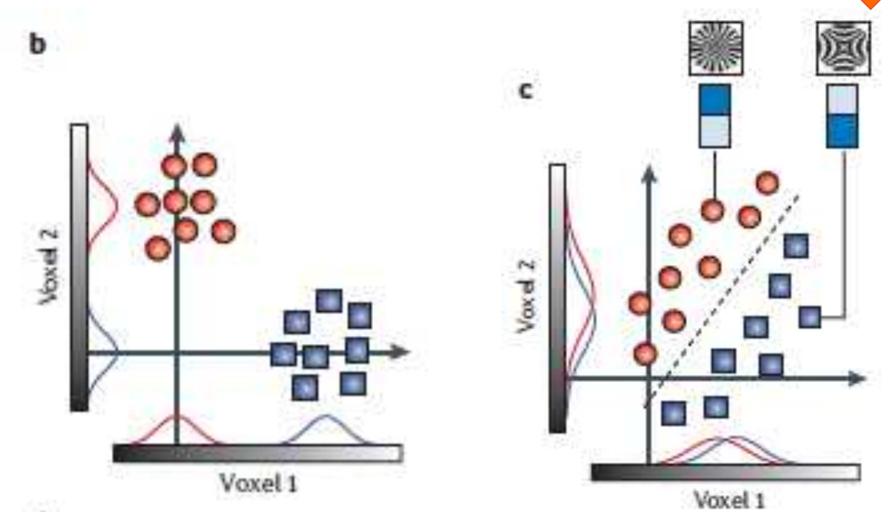

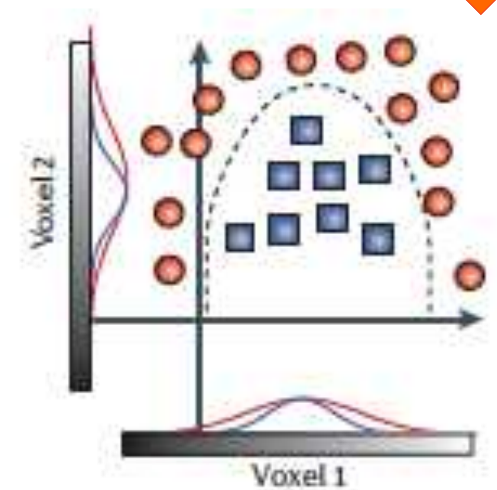

Mind Reading

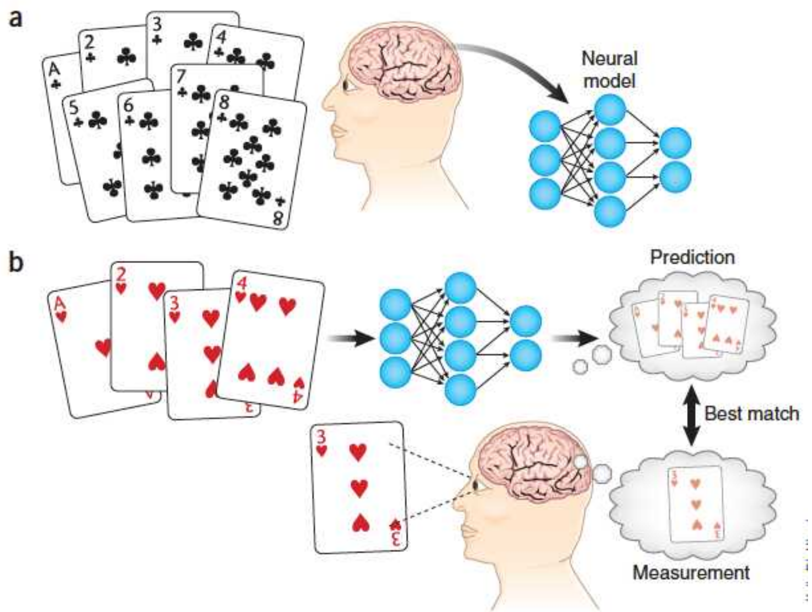

Identifying natural images from human brain activity
Kendrick N. Kay, Thomas Naselaris, Ryan J. Prenger & Jack L. Gallant

Nature Vol 452|20 March 2008

[Wandell Nature 2008]

A general brain reading device to access the visual contents of purely mental phenomena such as dreams and imagery

Water color effect

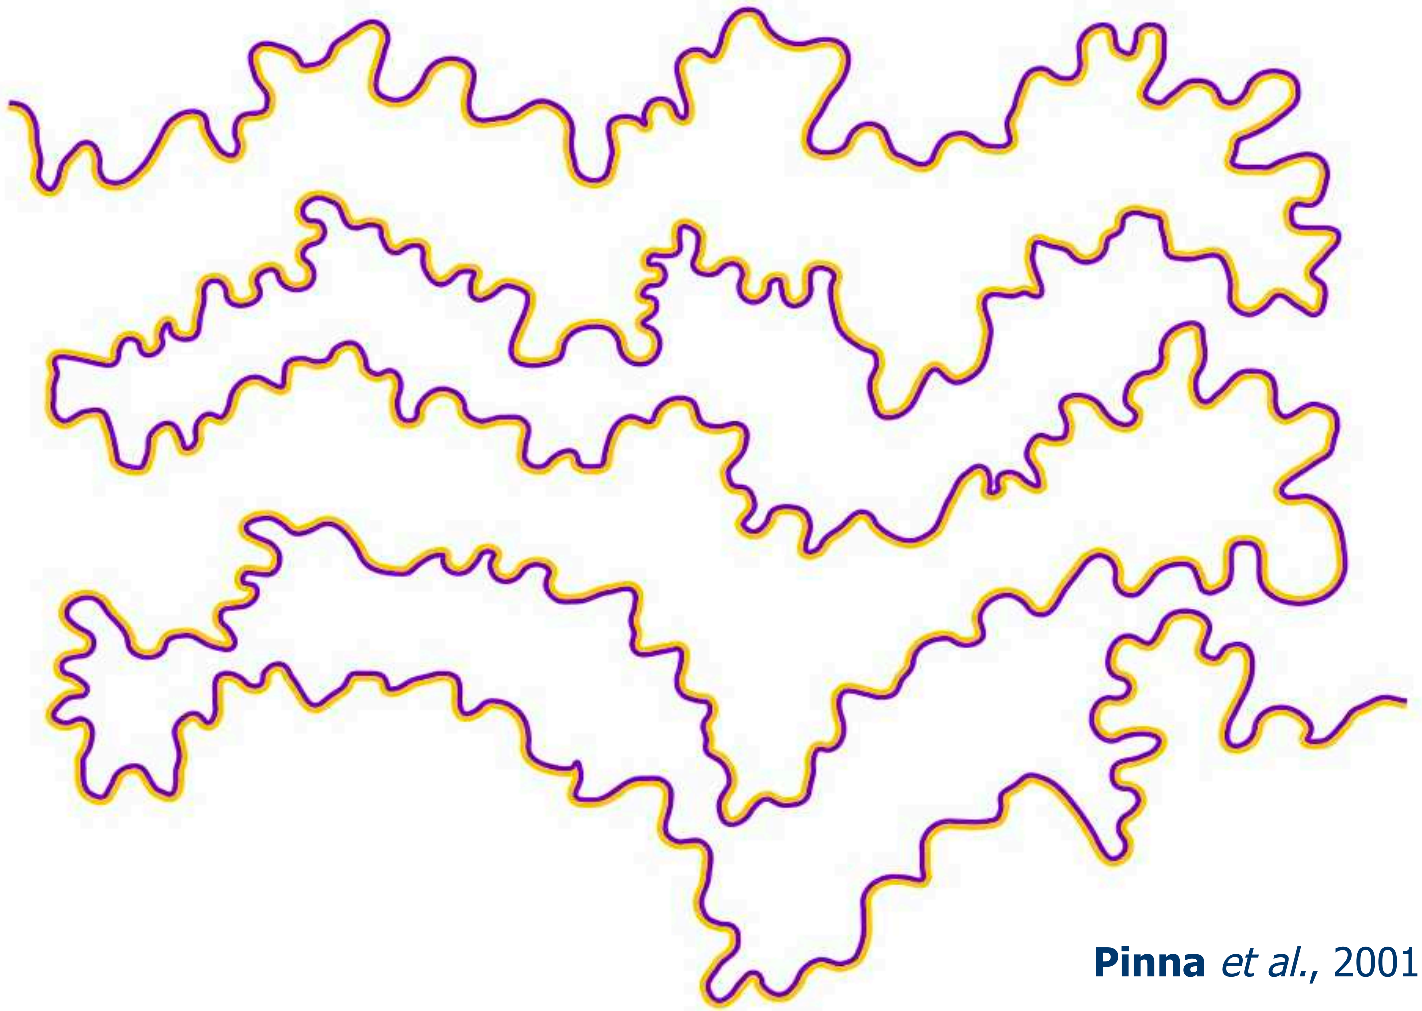

Pinna *et al.*, 2001.

Water color effect

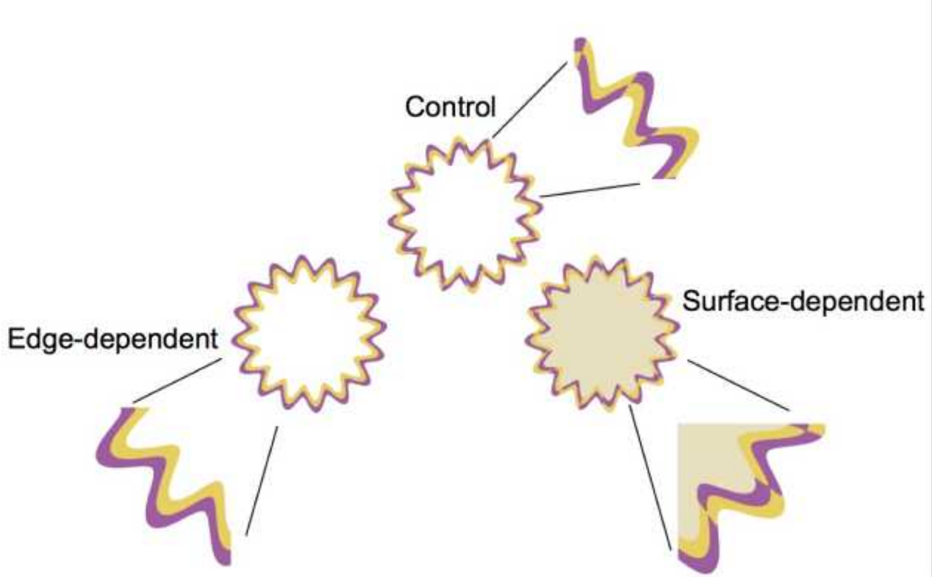

Stimulus and response differences tested by classifiers.

	Edge continuity	*Interior chromaticity*	*Interior perceived color*	
Edge-dependent vs Control	X		X	
Surface vs Control		X	X	
Surface vs Edge-dependent	X	X		

Gérardin *et al.*, Neuroimage 2018

Method

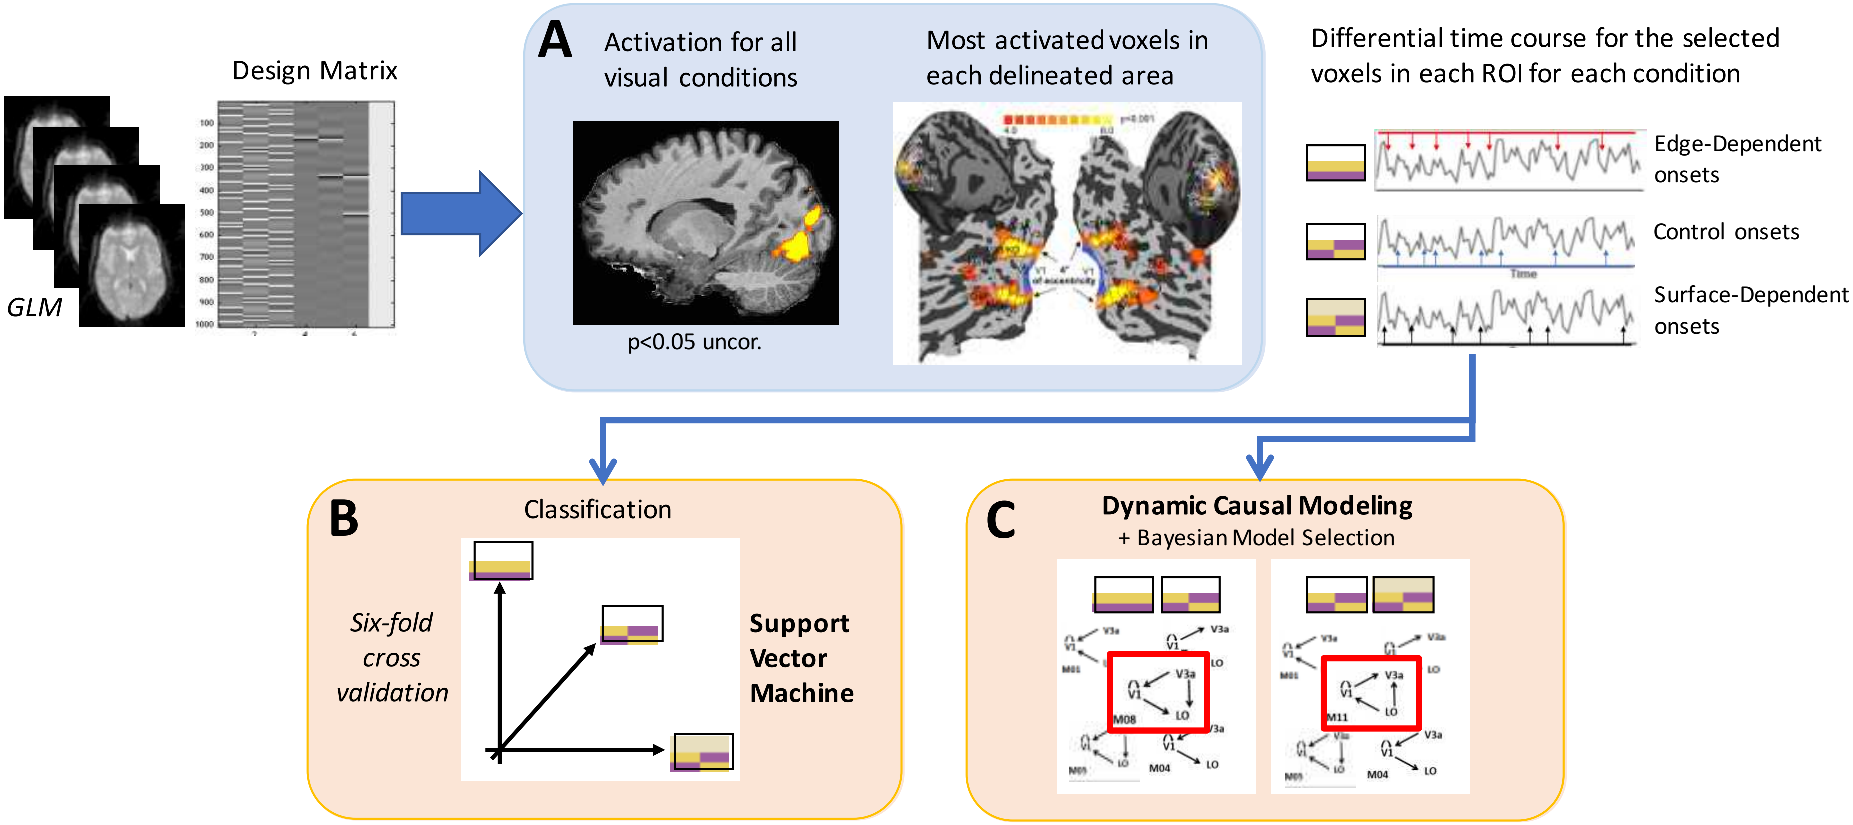

Design Matrix
GLM
A
Activation for all visual conditions
Most activated voxels in each delineated area
p<0.05 uncor.
Differential time course for the selected voxels in each ROI for each condition
Edge-Dependent onsets
Control onsets
Surface-Dependent onsets
B
Classification
Six-fold cross validation
Support Vector Machine
C
Dynamic Causal Modeling
+ Bayesian Model Selection

Results

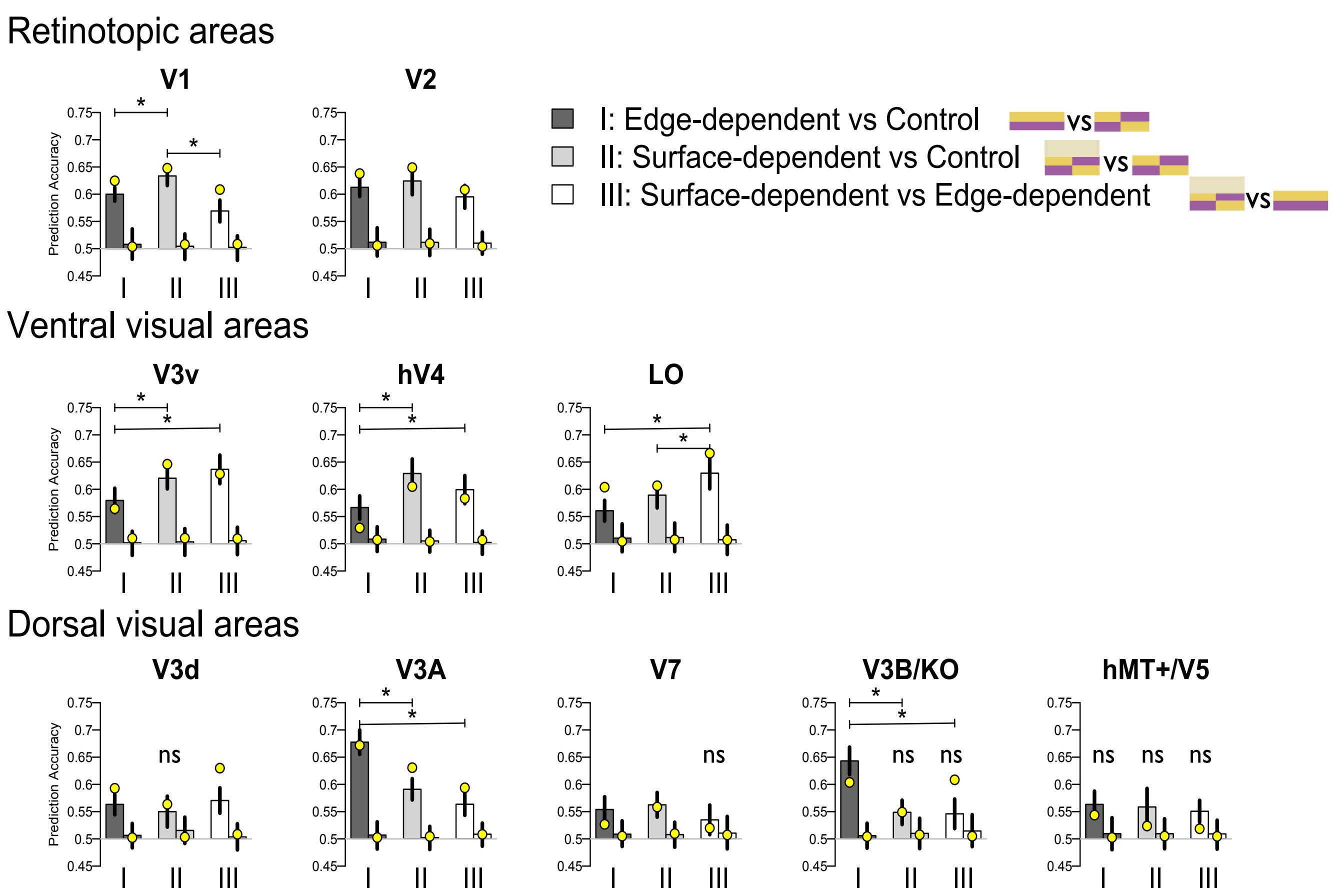

Multi-Voxel Pattern Analysis (MVPA) from fMRI data.

Conclusion

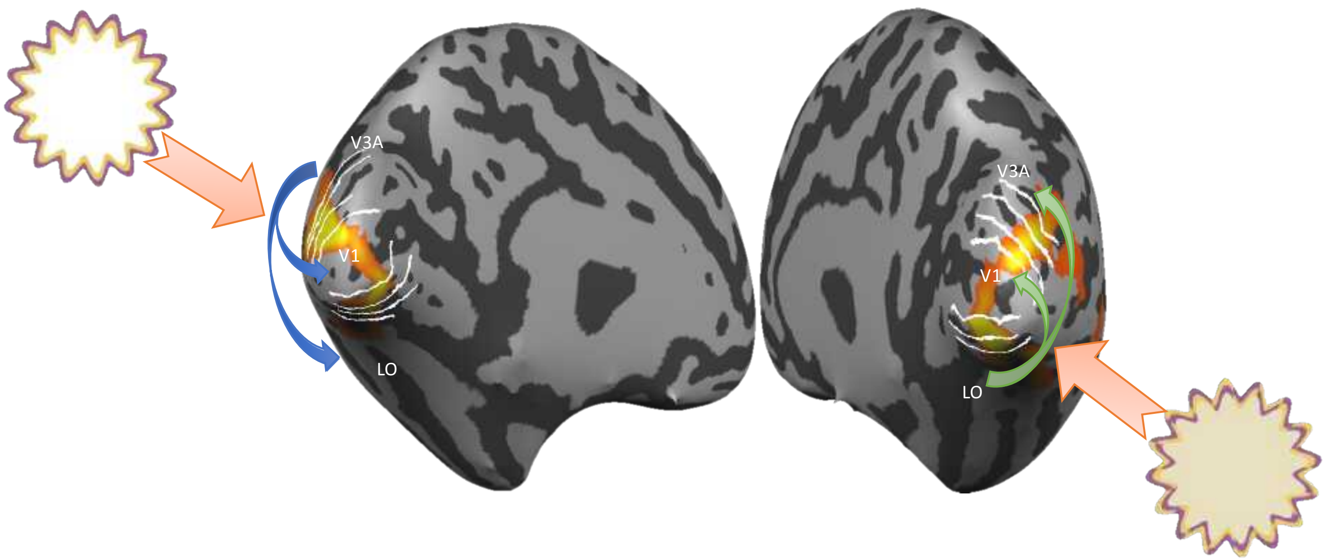

– Filling-in is best classified and best correlate with appearance by dorsal areas V3A & V3B/KO
– Uniform chromaticity by ventral areas hV4 & LO
– Feedback modulation from V3A to V1 and LO for filling-in
– Feedback from LO modulating V1 and V3A for uniform chromaticity

Color decoding

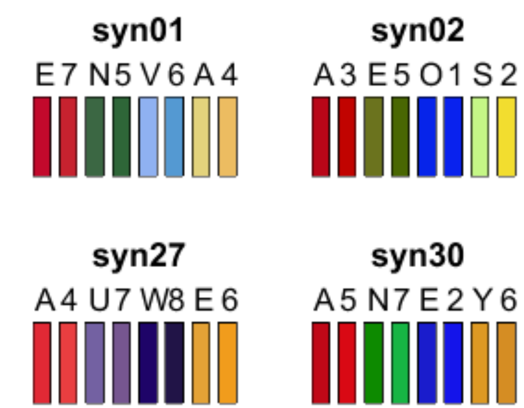

Multivariate pattern analysis of fMRI data for imaginary and real colours in grapheme-colour synaesthesia

Mathieu J. Ruiz, Michel Dojat, Jean-Michel Hupe

bioRxiv 214809; doi: https://doi.org/10.1101/214809

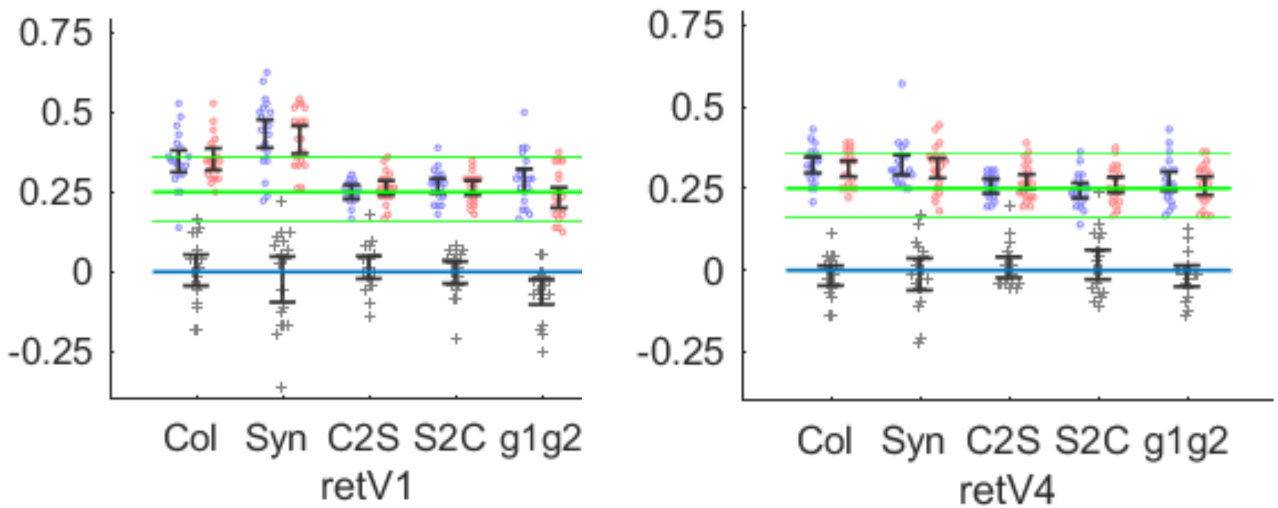

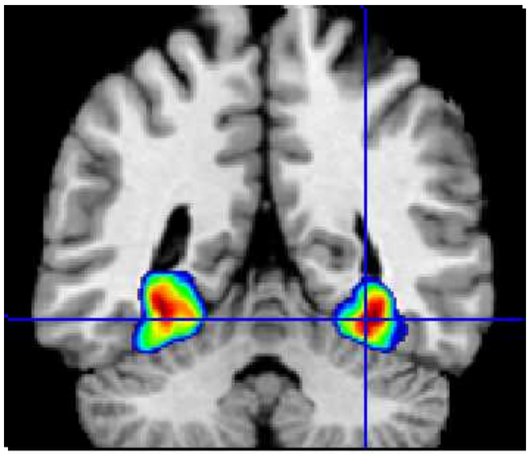

Neural Networks

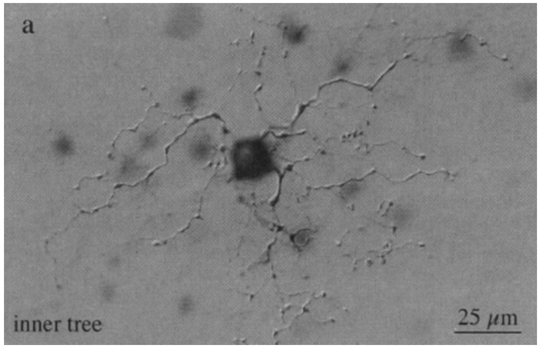

Artificial neurons (bio-inspirés)

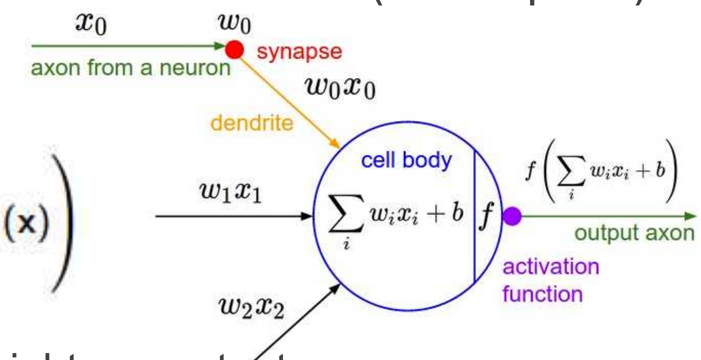

$$y(\mathbf{x}, \mathbf{w}) = f\left(\sum_{j=0}^{M} w_j \phi_j(\mathbf{x})\right)$$

x: observations, w: weights, y: output

ϕ: models basis or activation functions

$$t(\mathbf{x}) = y(\mathbf{x}, \mathbf{w}) + \epsilon(\mathbf{x})$$

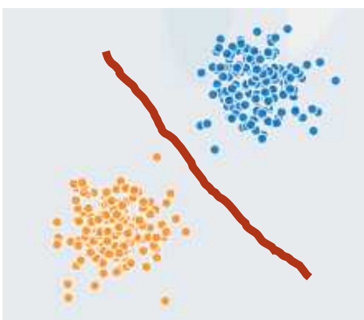

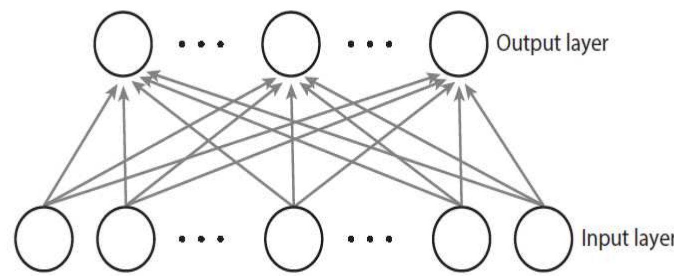

Multi-layer network

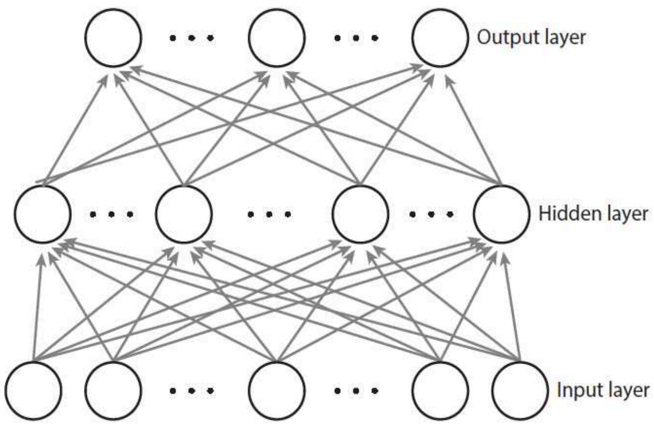

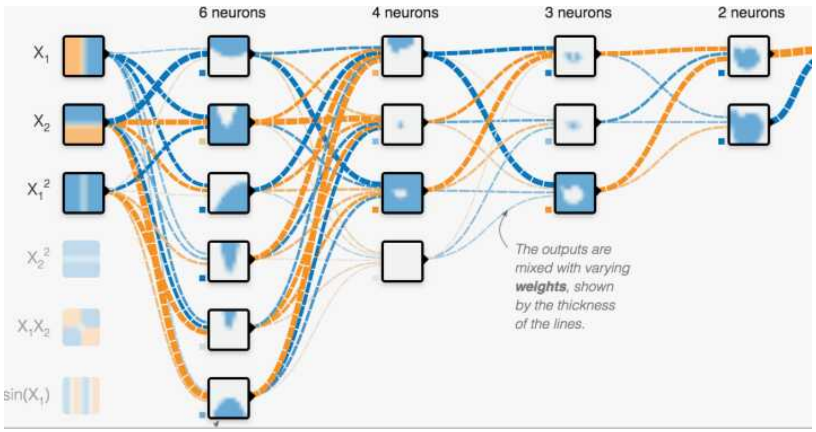

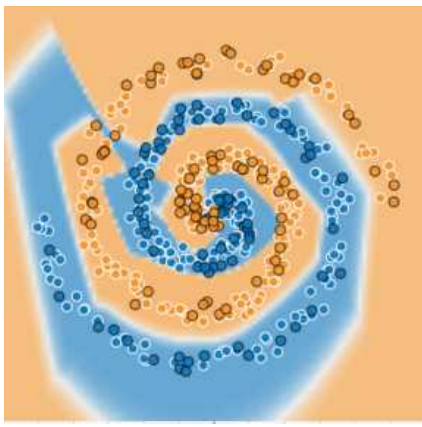

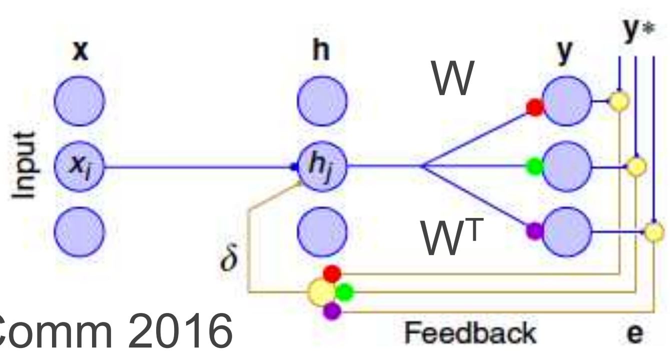

Lillicrap Nat Comm 2016

Convolutional Neural Network

Multi-scale representation
Hierarchical extraction of information
Features of filters learnt

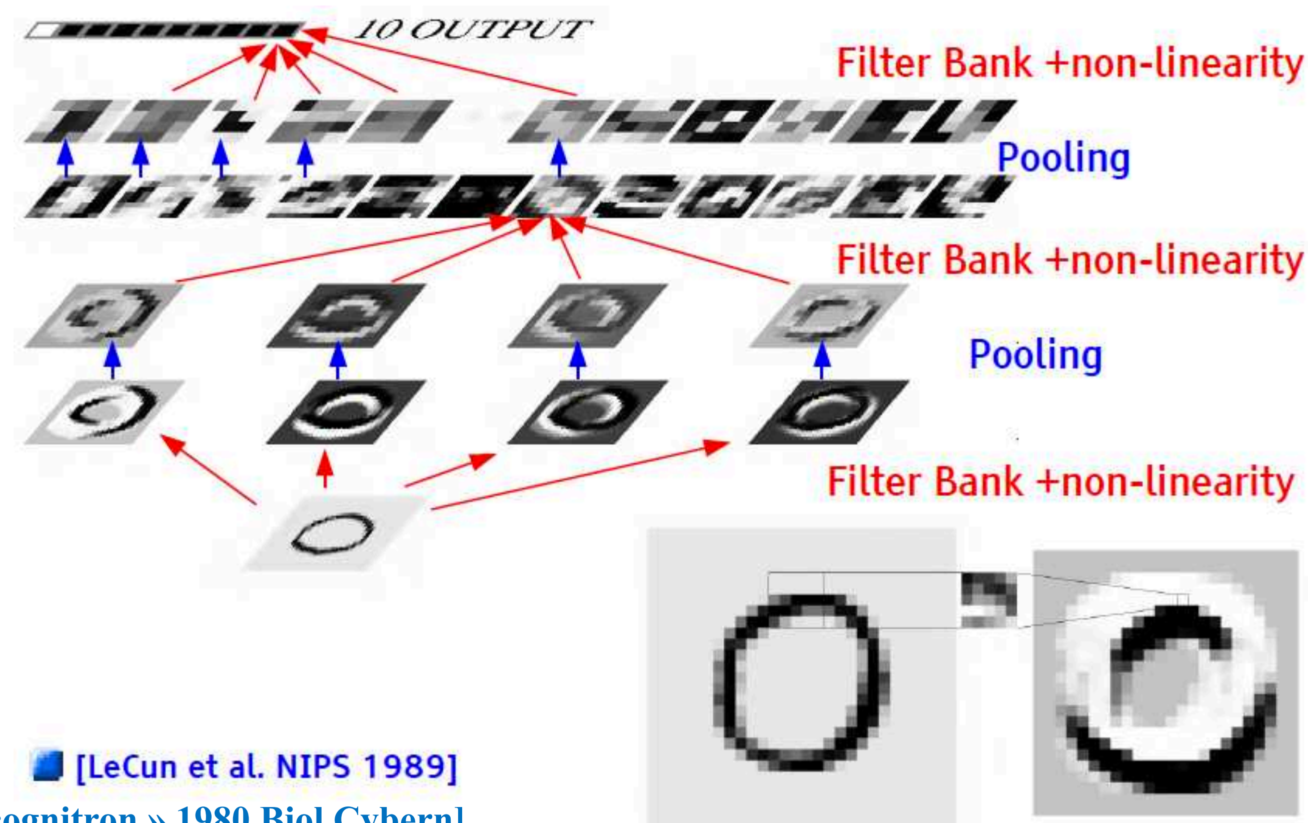

[Fukushima « Neocognitron » 1980 Biol Cybern]

CNN as a model of the visual system …

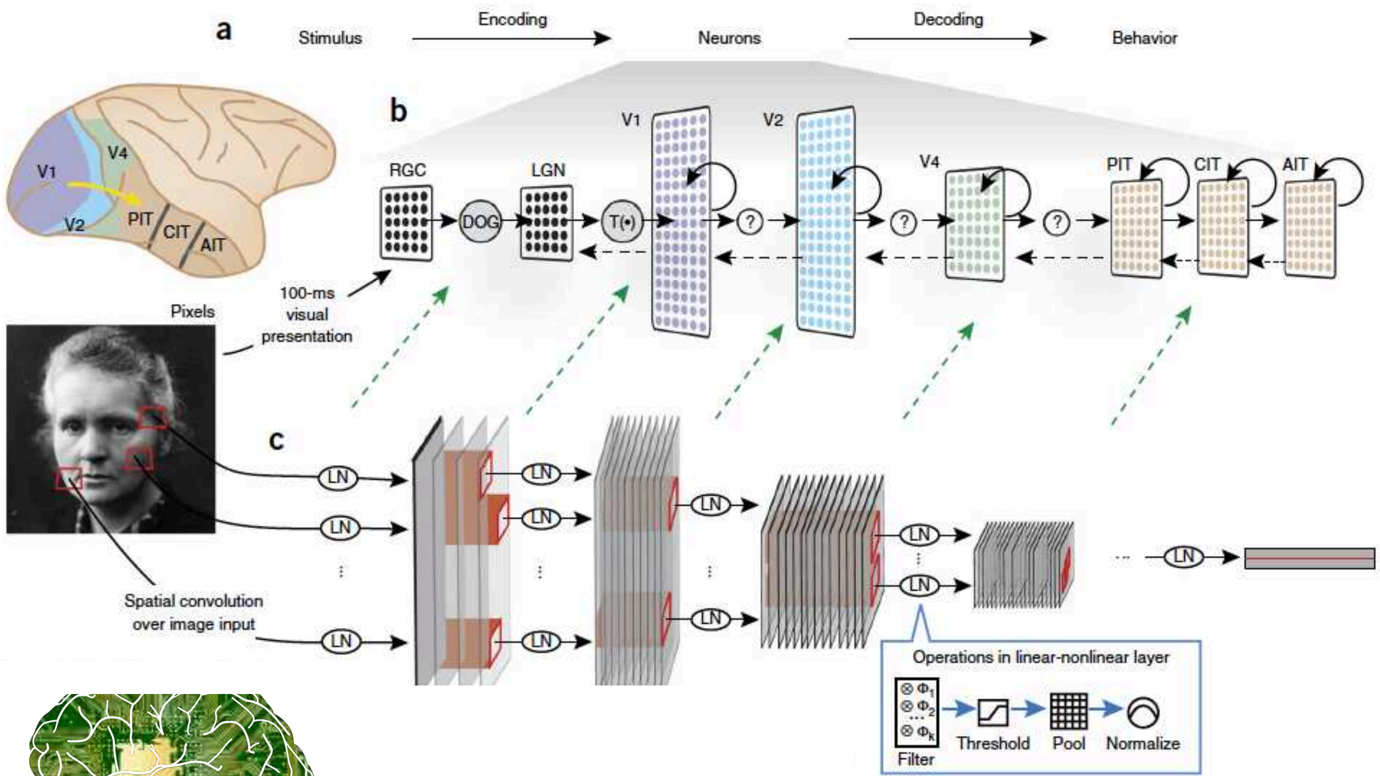

[Yamins and Dicarlo Nat Neuro 2016]

- ✓ Receptive field
- ✓ Hierarchical model
- ✓ Max pooling

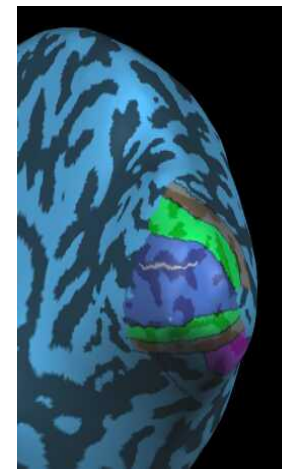

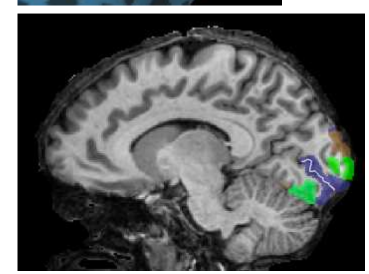

[Dojat Lavoisier ed. 2017]

Hierarchical models for neural responses prediction

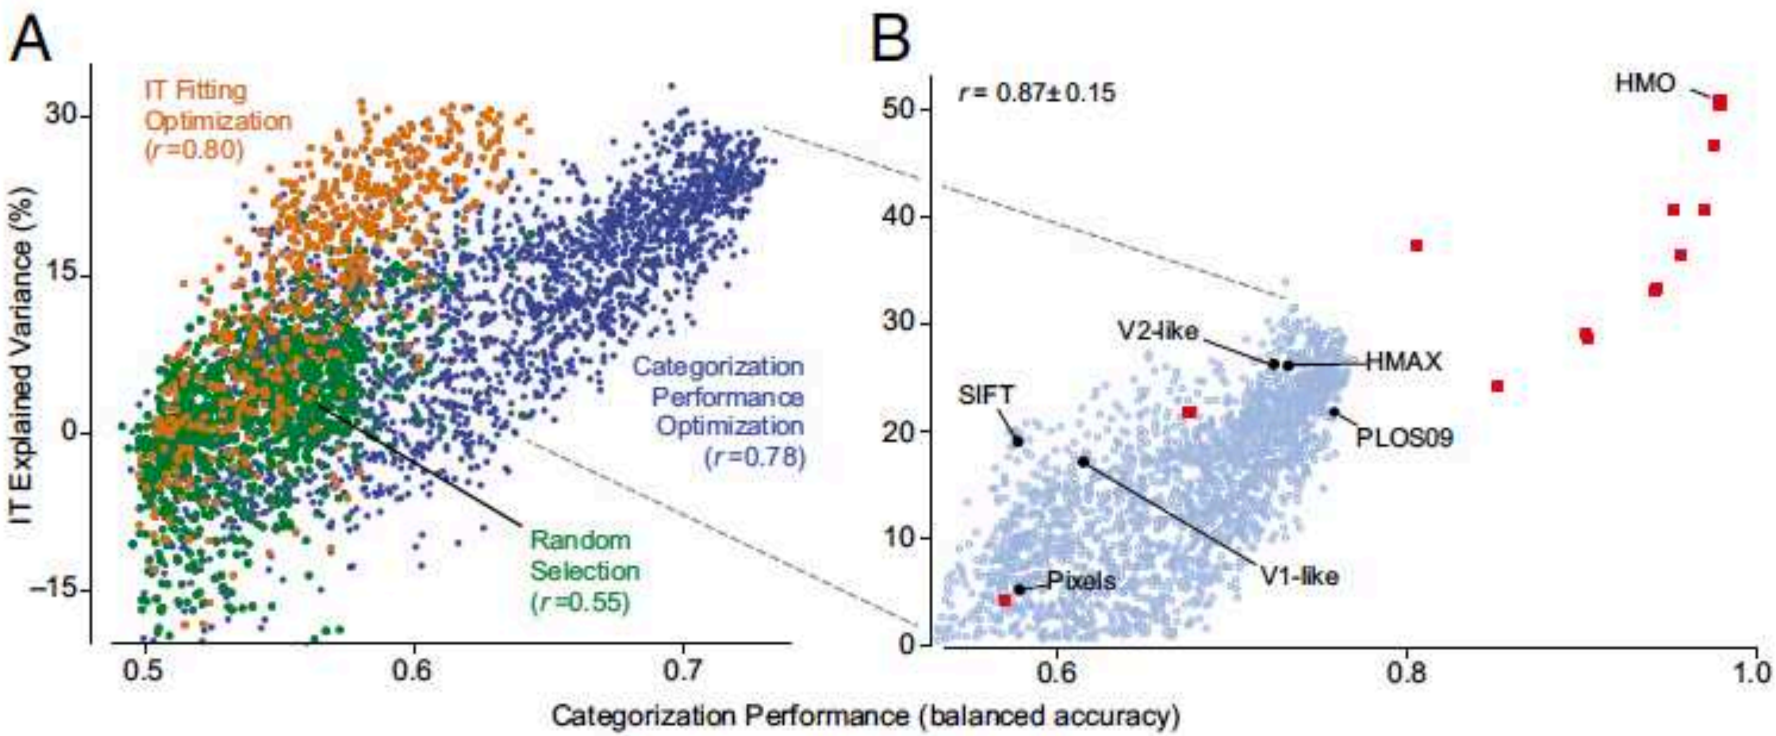

[Yamins et al 2014 PNAS]

Hierarchical models for neural responses prediction

[Yamins et al 2014 PNAS]

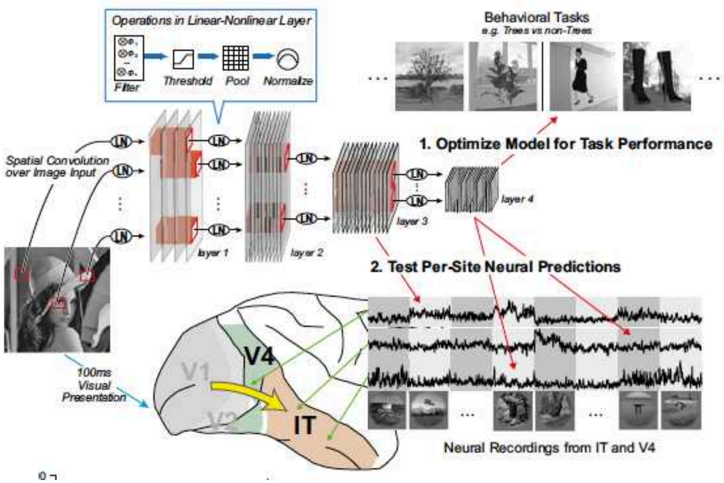

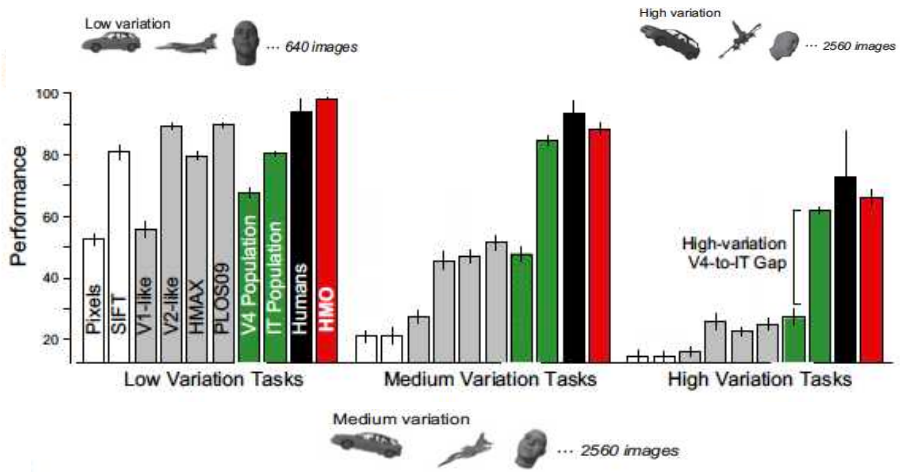

Representation dissimilarity matrices

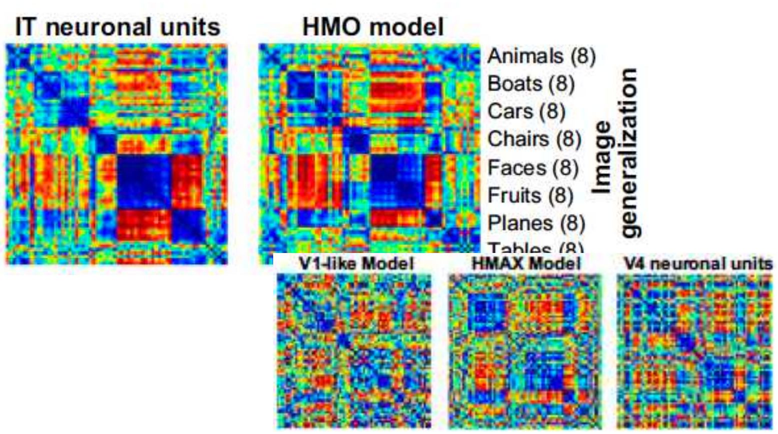

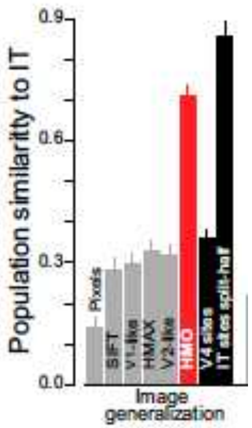

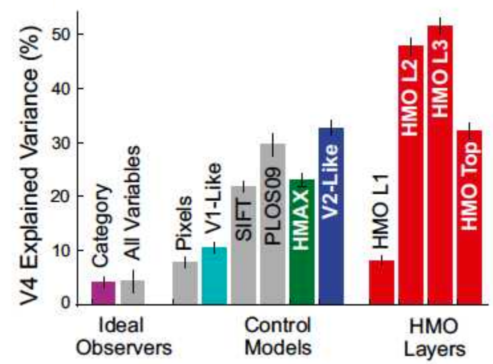

Hierarchical models for functional responses prediction

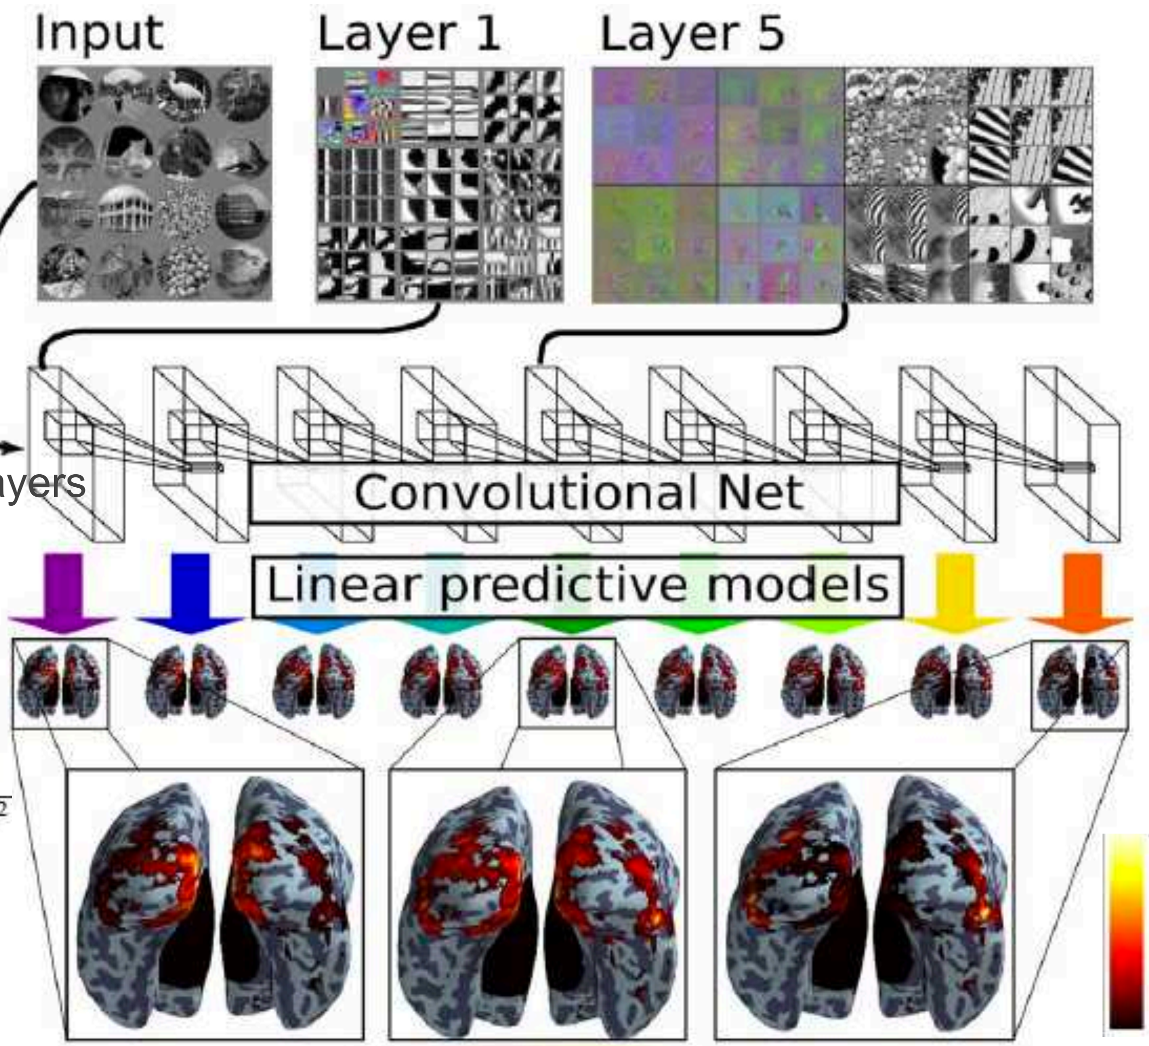

OverFeat
6 convolutional layers
3 fully connected

$$r_v^2 = 1 - \frac{\| y^v_{\mathrm{test}} - y^v_{\mathrm{pred}} \|^2}{\| y^v_{\mathrm{test}} - \mathrm{mean}(y^v_{\mathrm{test}}) \|^2}$$

Data:
Kay et al Nature 2008

[Eickenberg Neuroimage 2017]]

Hierarchical models for functional responses prediction

[Eickenberg Neuroimage 2017]

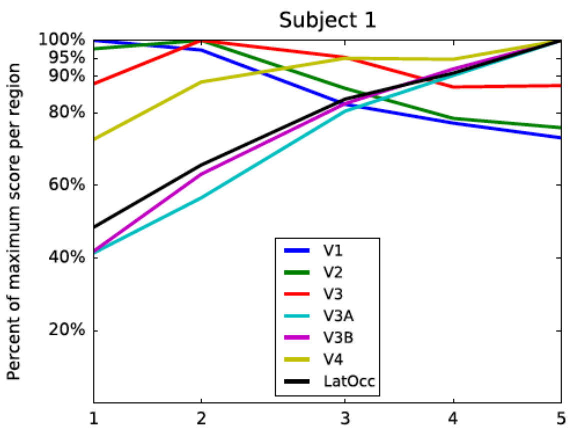

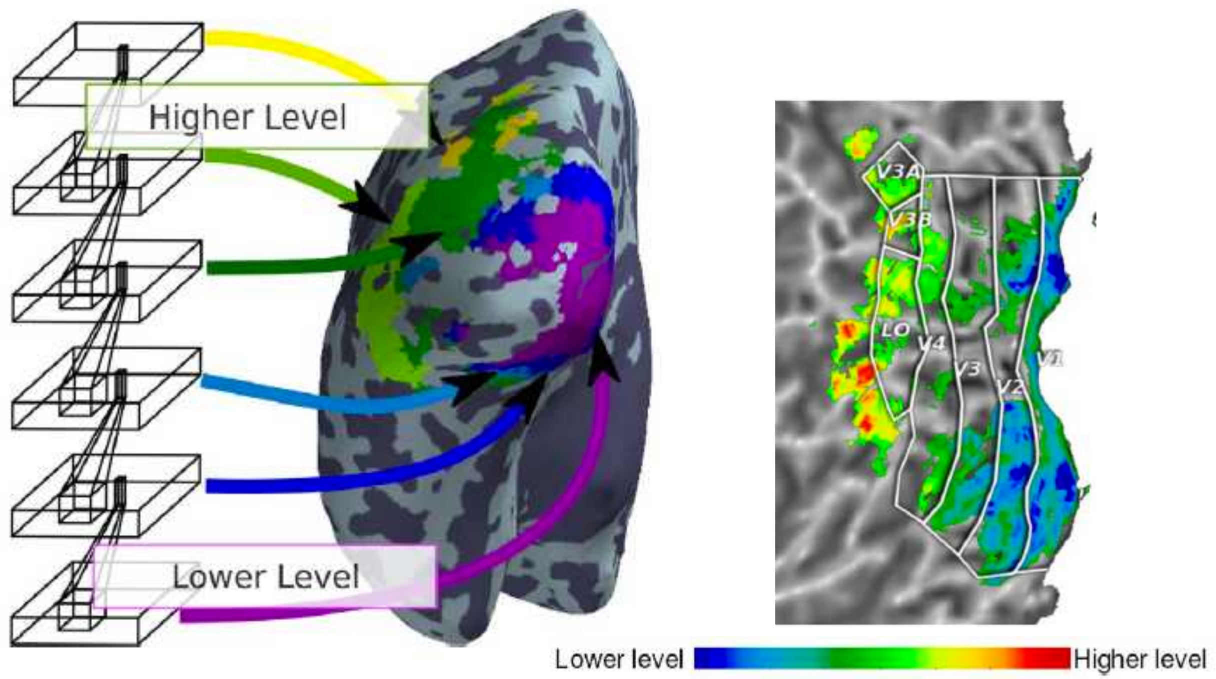

Data:
Naselaris et al Neuron 2009

[Güclü and van Gerven J Neuro 2015; Cadena et al 2017 bioRxiv; Cichy et al 2017 Scient Rep; Greene et al PLOS 2018; Seeliger et al . NeuroIm 2018; Wen et al Scient Rep 2018]

Hierarchical models for functional responses prediction

[Eickenberg Neuroimage 2017]

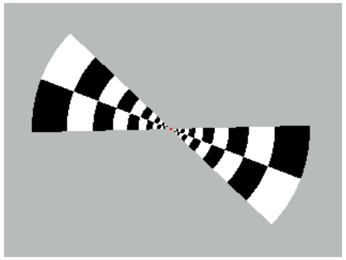

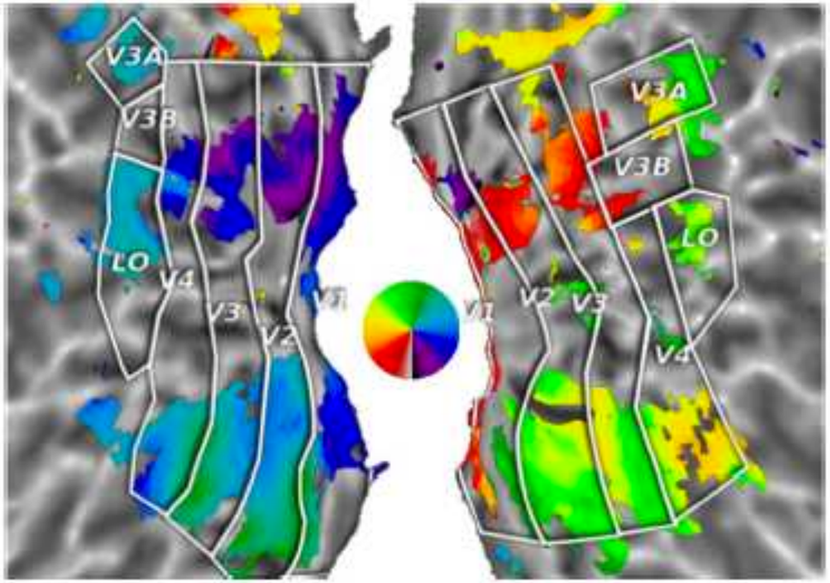

Seen/imagined object arbitrary categories

CNN1…8
HMAX1..3
GIST, SIFT+BoF

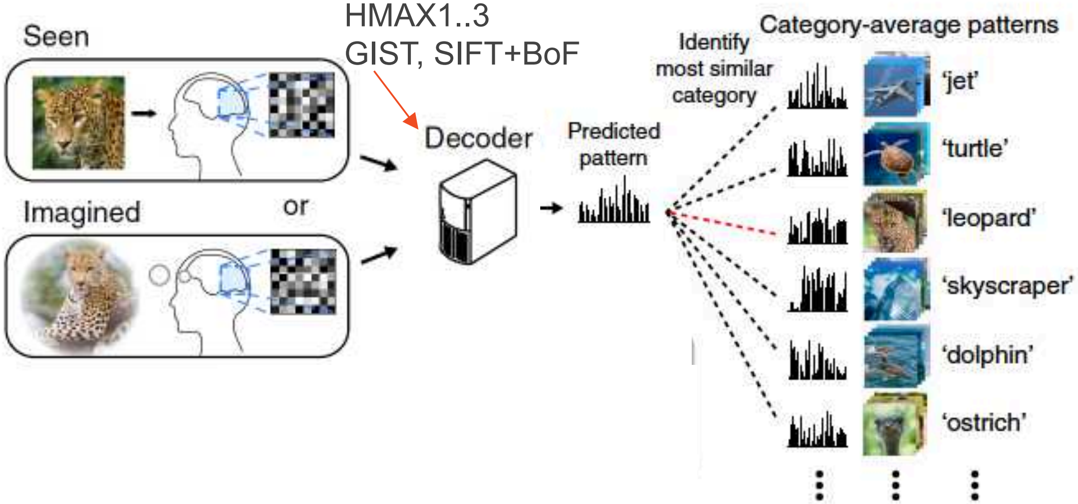

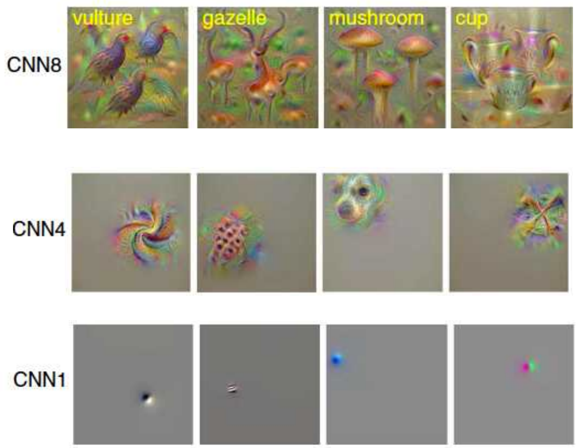

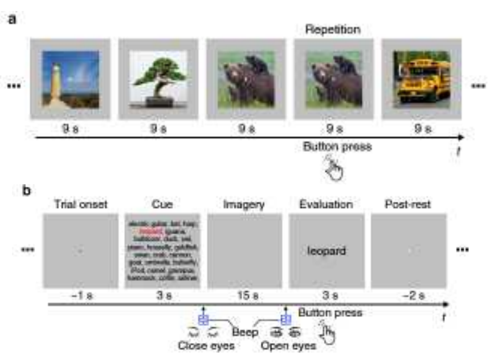

Train: 1200 im/50 obj cat (x24)
Test: 50 im/50 obj cat (x35)
9 min each->9h15

1im/50 obj cat (x10)
Idem test
10 min each ->3h33

Decoding accuracy

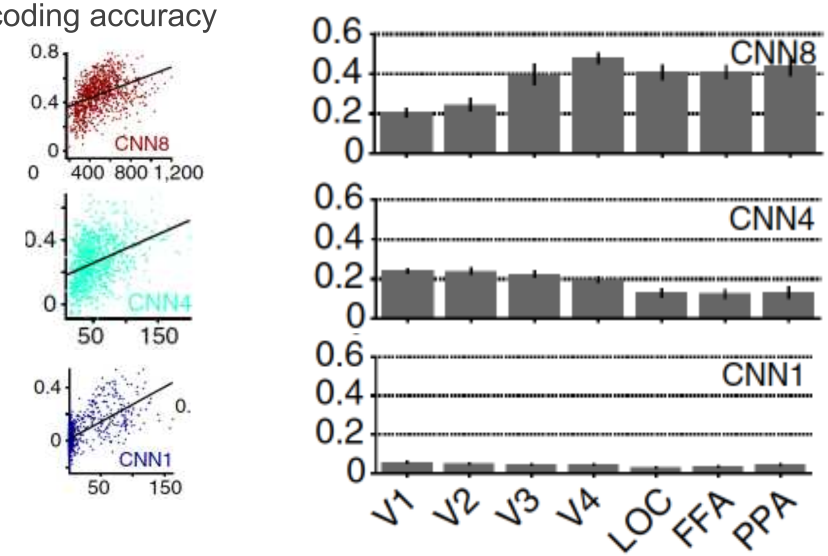

[Horikawa and Kamitani Nat Comm 2018]

Seen/imagined object arbitrary categories

CNN1…8
HMAX1..3
GIST, SIFT+BoF

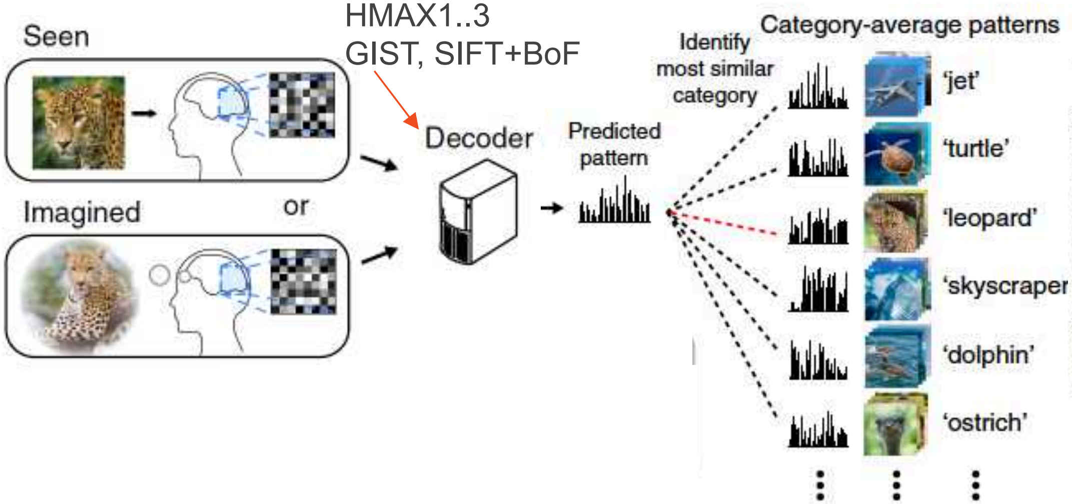

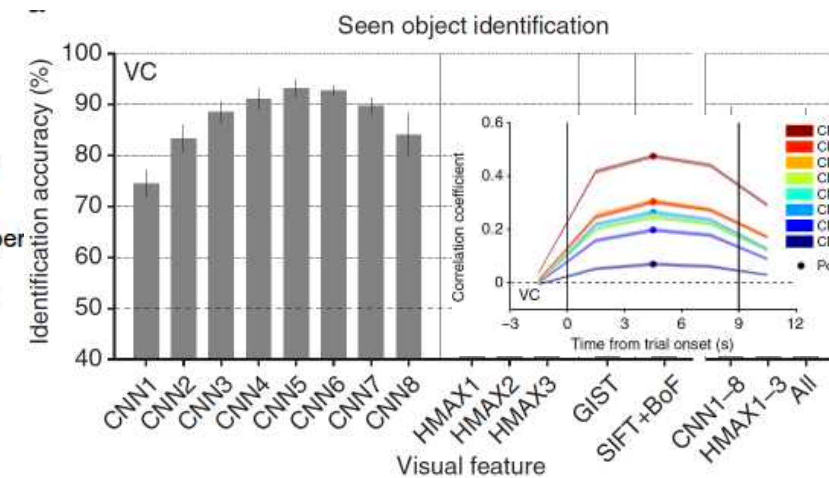

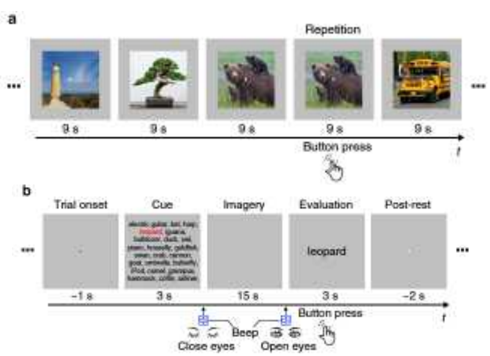

Train: 1200 im/50 obj cat (x24)
Test: 50 im/50 obj cat (x35)
9 min each->9h15

1im/50 obj cat (x10)
Idem test
10 min each ->3h33

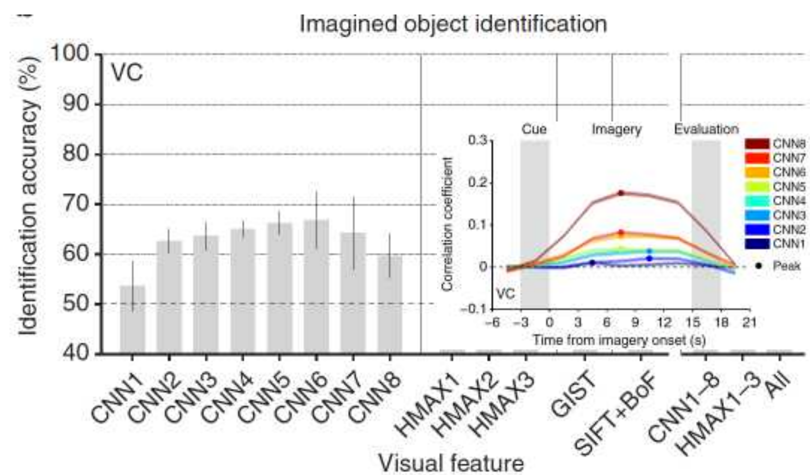

[Horikawa and Kamitani Nat Comm 2018]

Dynamic natural vision

[Wen el al CC 2017]

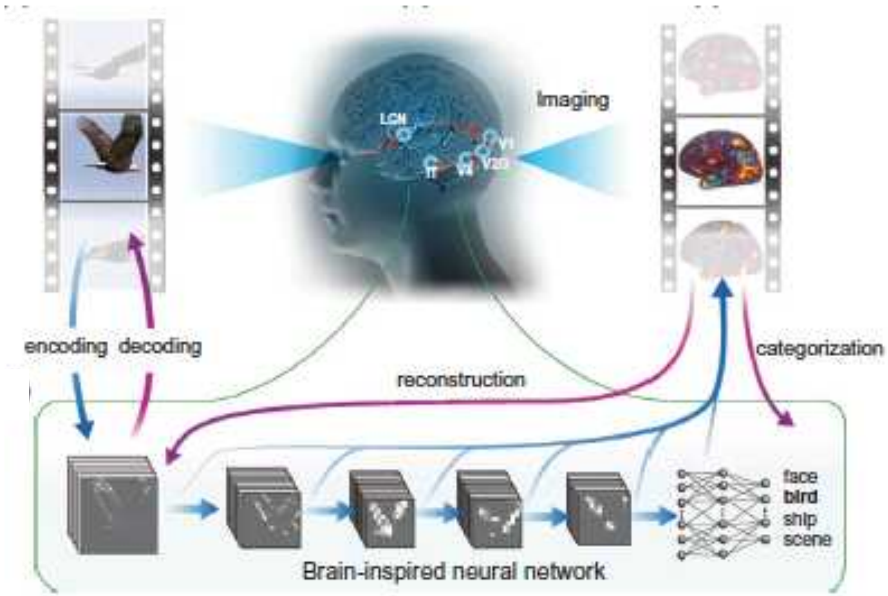

2.4h Training movie x 2
40 min Testing movie x10
3 subjects, 3T scanner
3.5mm^3, 220x220mm^2
AlexNet 8 layers

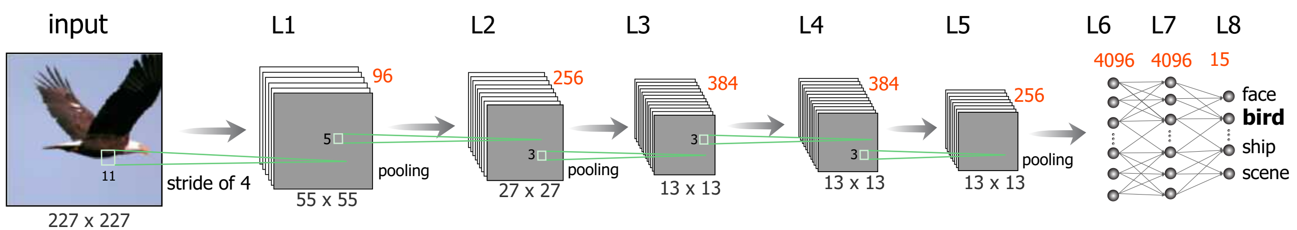

Exploring visual representation

2000 human faces

64 000 images
80 classes
800 categories

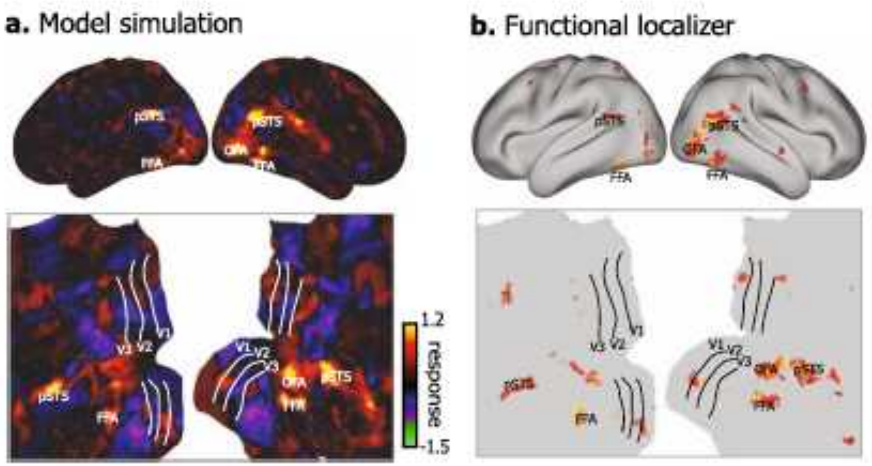

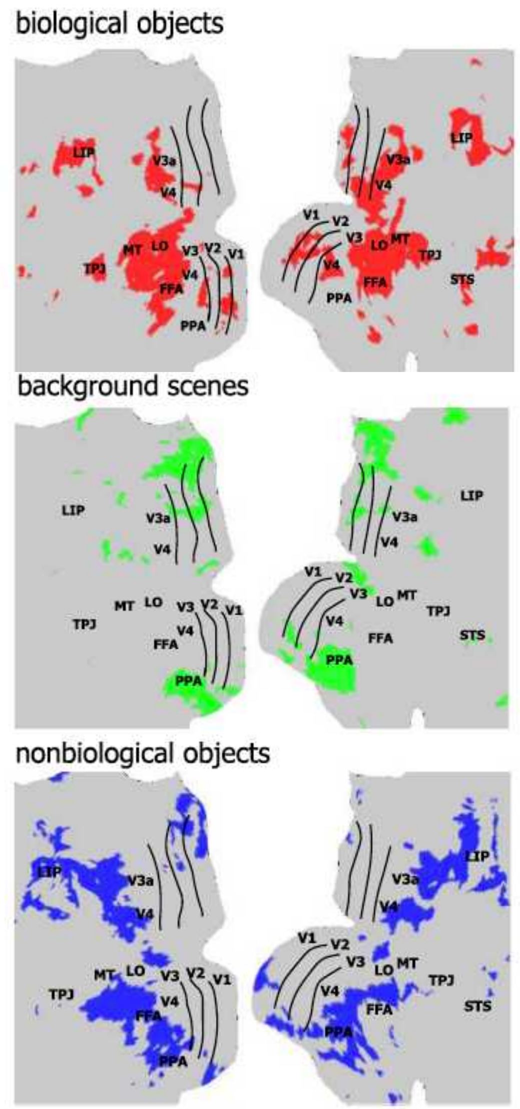

bear bird cat chicken
dog elephant goose horse
lion monkey sheep tiger

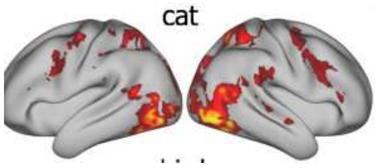

bedroom boat bridge building
classroom corridor door factory
house kitchen livingroom market

airplane bag ball bike bottle
bowl car cellphone watch
drink flag hat instrument

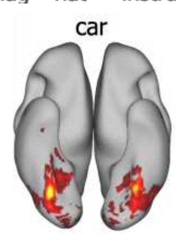

[Wen et al. Scient Rep 2018]

Nobody is perfect …

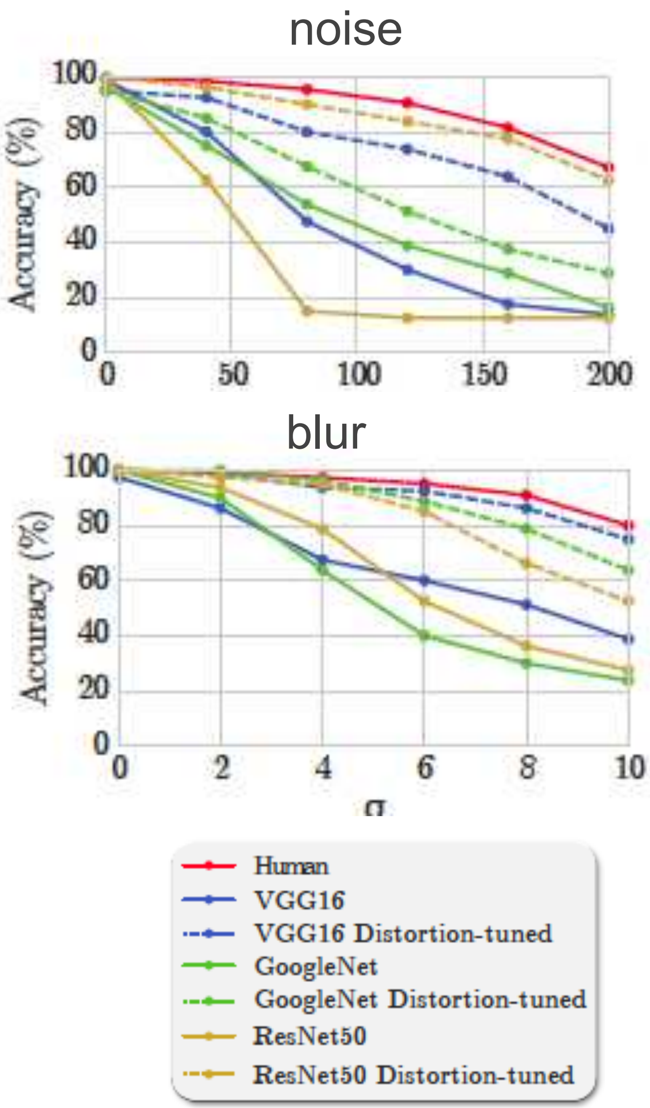

[Dodge and Karam ICCV 2017]

CNN as a model of the visual system …

- CNN driven for image recognition
  - Explains significant variance of neuroimaging data in ventral stream
  - Models a hierarchical representation of feedforward visual information processing (i.e. ventral stream)
  - Supports the generation of expected cortical activation
  - Supports decoding & then semantic categorisation
  - Validates the existing models

  BUT still incomplete …

CNN as a model of the visual system …

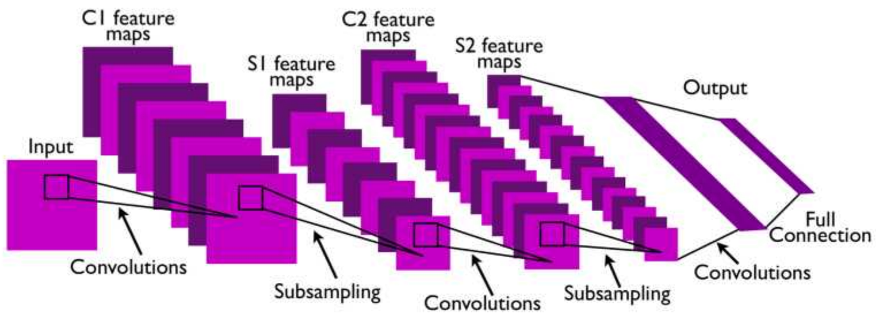

[Lecun et al 2010]

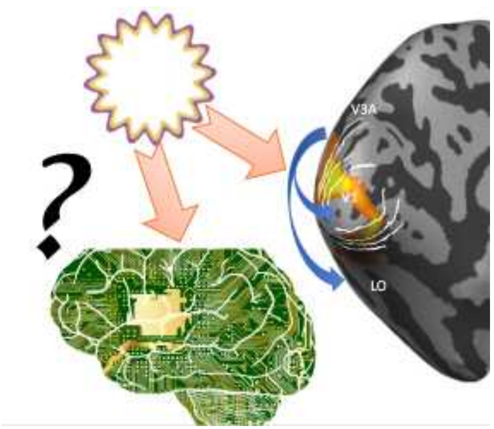

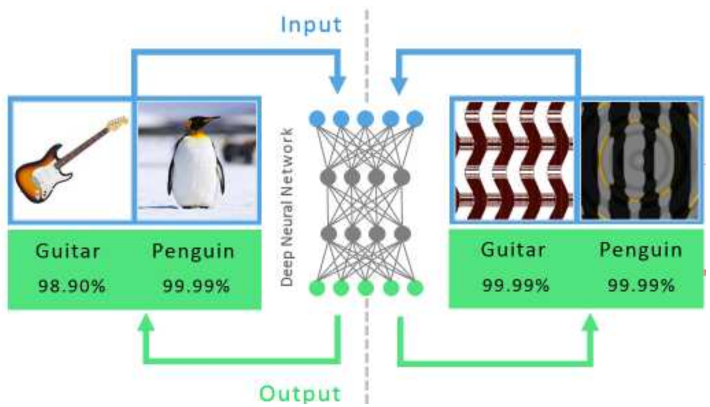

[Nguyen et al CVPR 2015]

- ✓ Receptive field
- ✓ Hierarchical model
- ✓ Max pooling
- ✓ Layer2layer conn.
- ✓ Magnitude factor
- ✓ Color
- ✓ Metamer
- ✓ Attention
- ✓ Local vs Global
- ✓ Perspective effects
- ✓ Learning
- ✓ Illusion
- ✓ Eye mvt

[Dali 1941]

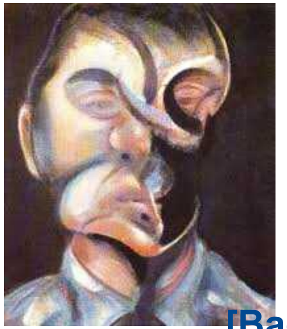

[Bacon 1972]

Illusions - Hallucinations

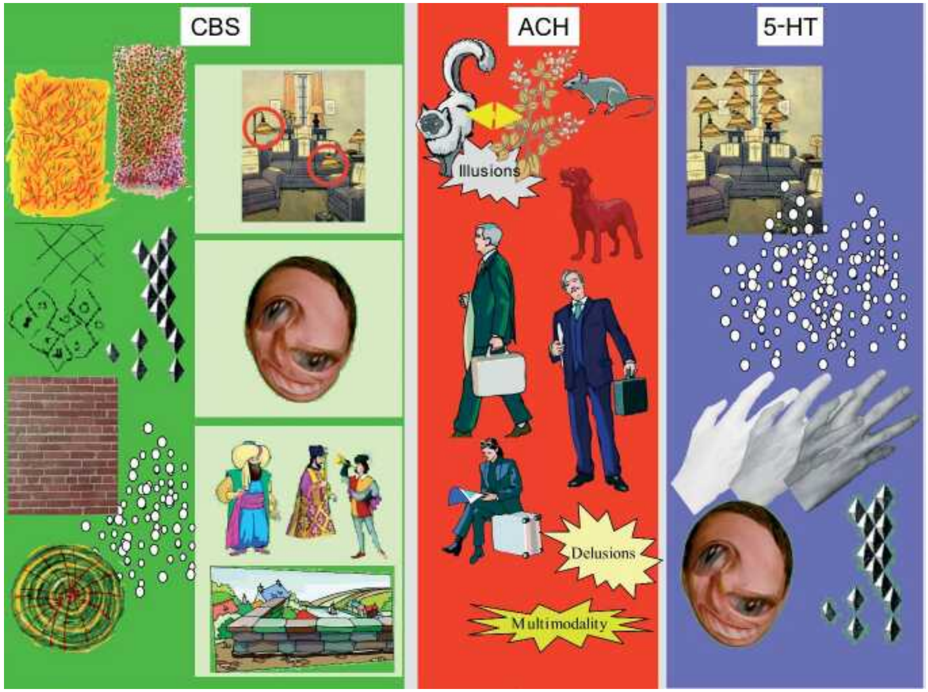

[Ffytche 2007 Dial Clin Neurosc]

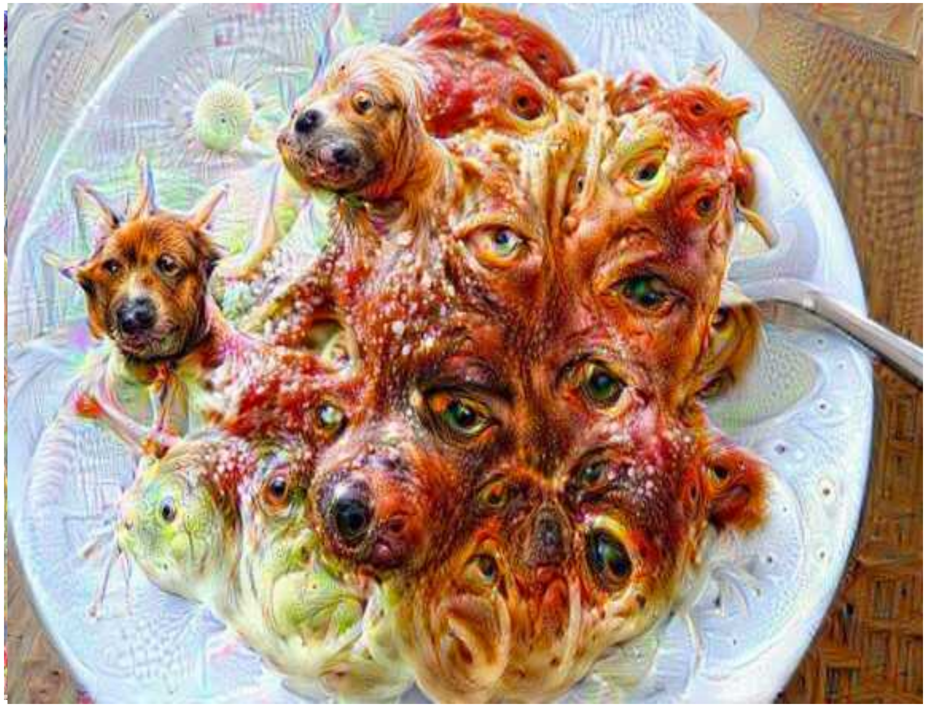

DeepDream.com

CNN a model for … auditory cortical responses

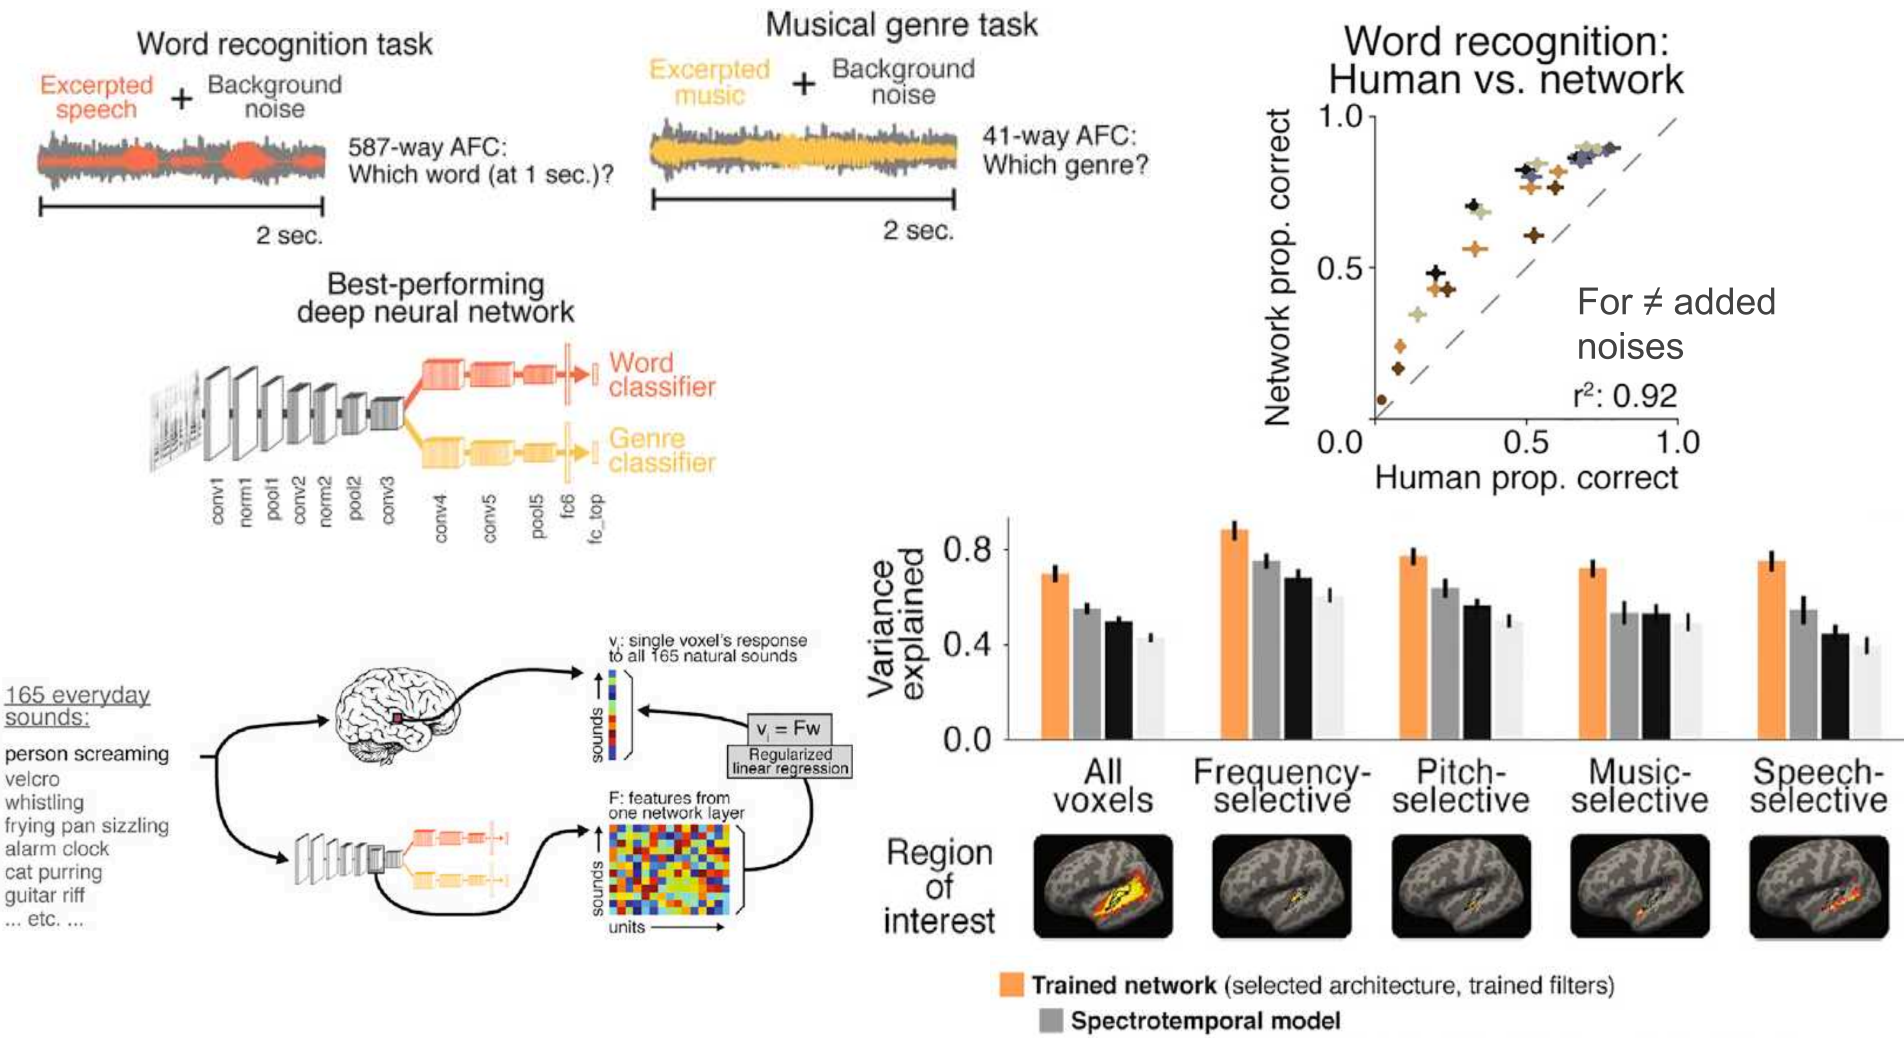

[Kell et al Neuron 2018]

AI for catching human errors

Medical error-the third leading cause of death in the US

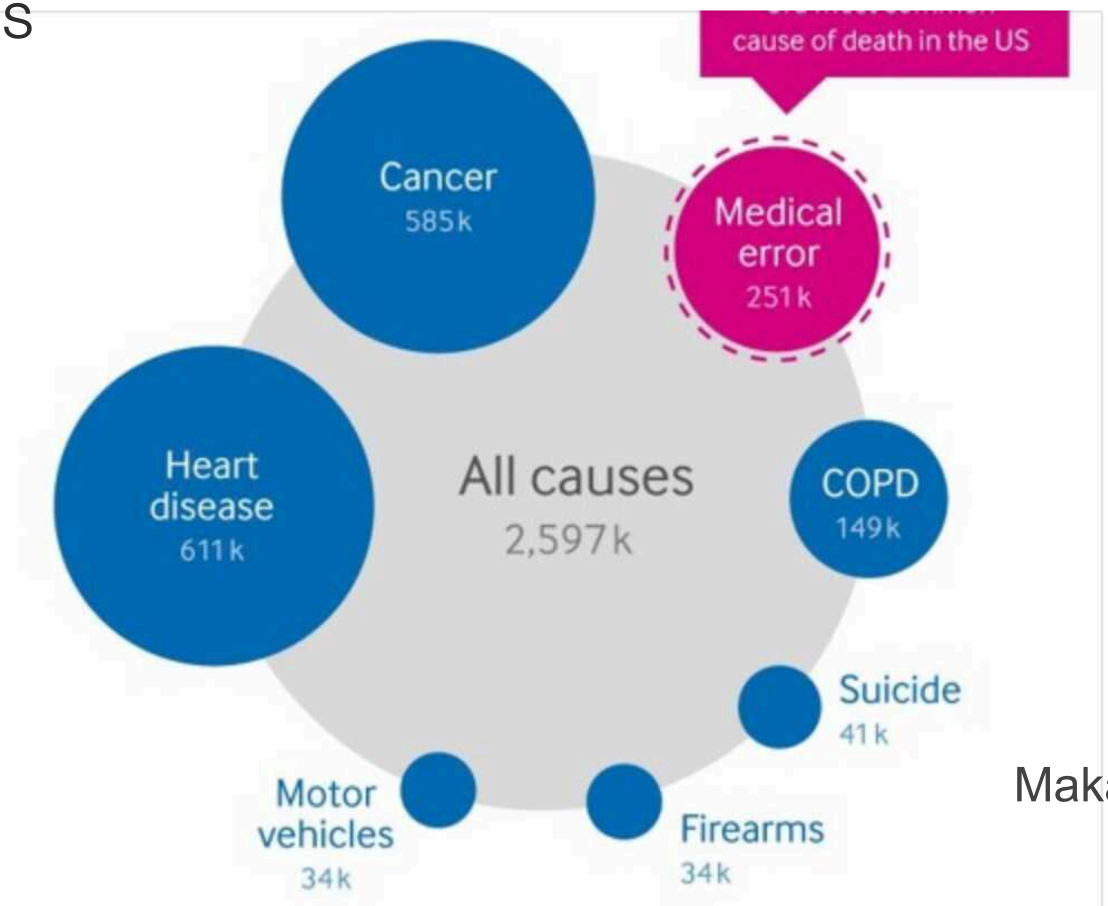

Makari et al BMJ 2016

New tools …

http://clickme.ai/

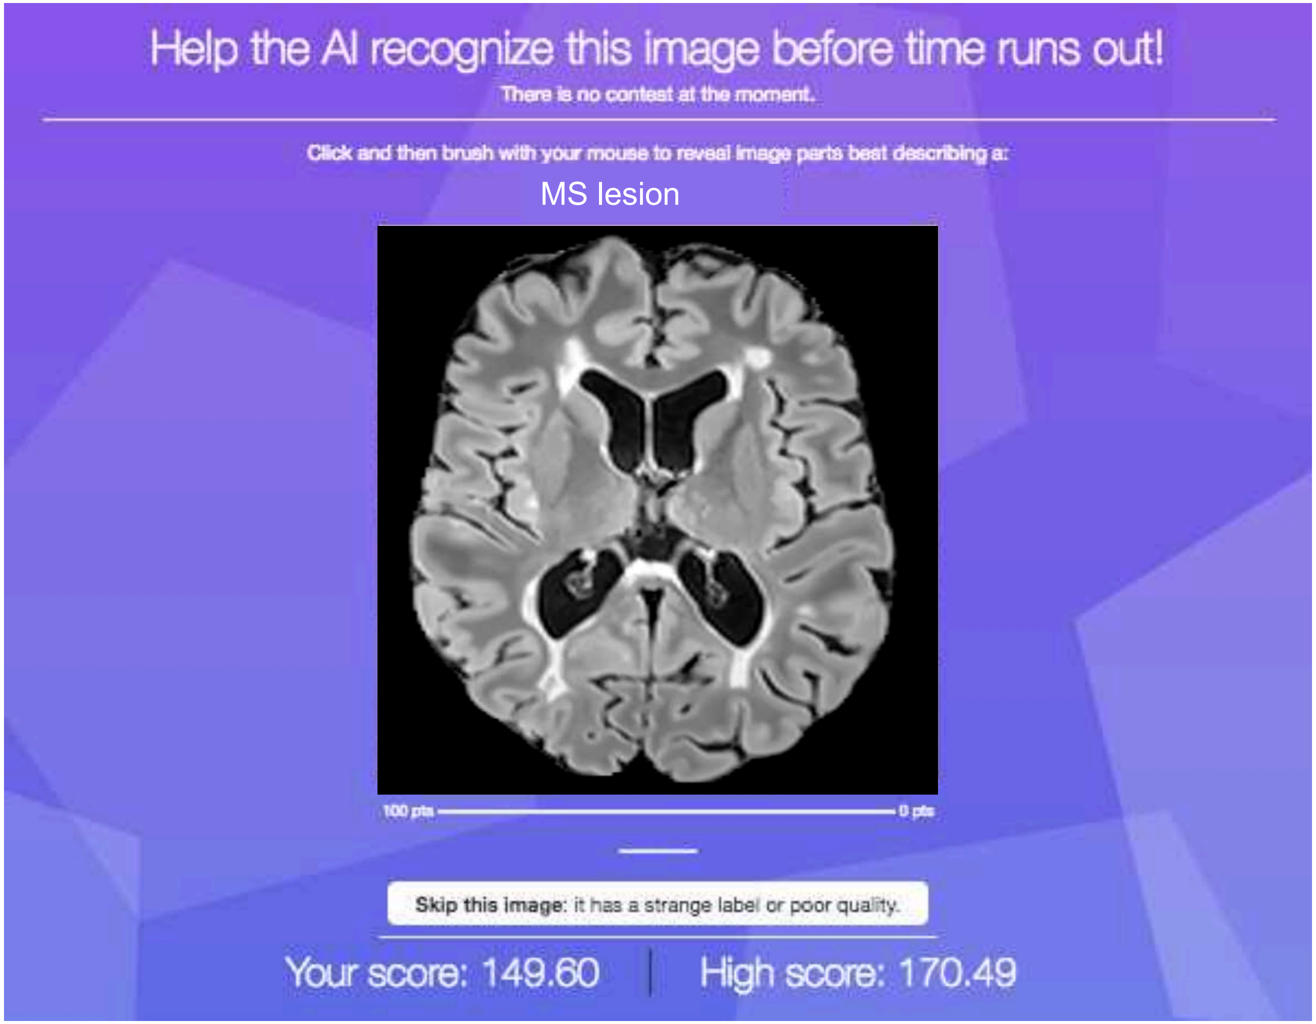

Art History

How characteristics of style are identified?
How the patterns evolve?

76921 paintings
Train(85%)
Val (9.5%)
Test (5.5%)

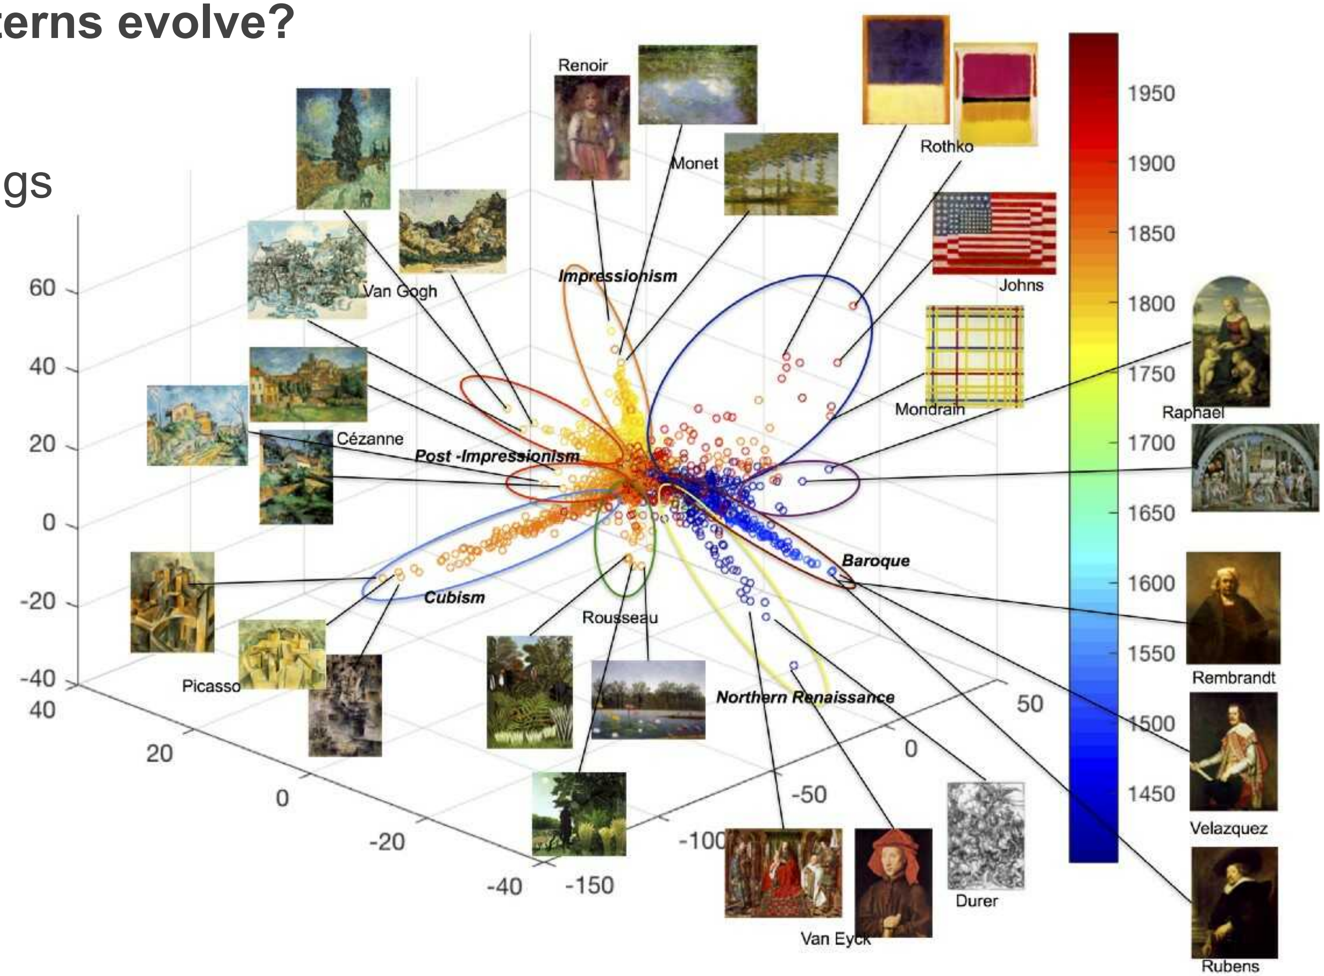

Elgammal et al. 2018 arxiv 1801.07729

e-Nosology

How characteristics of pathology are identified?
How the patterns evolve?

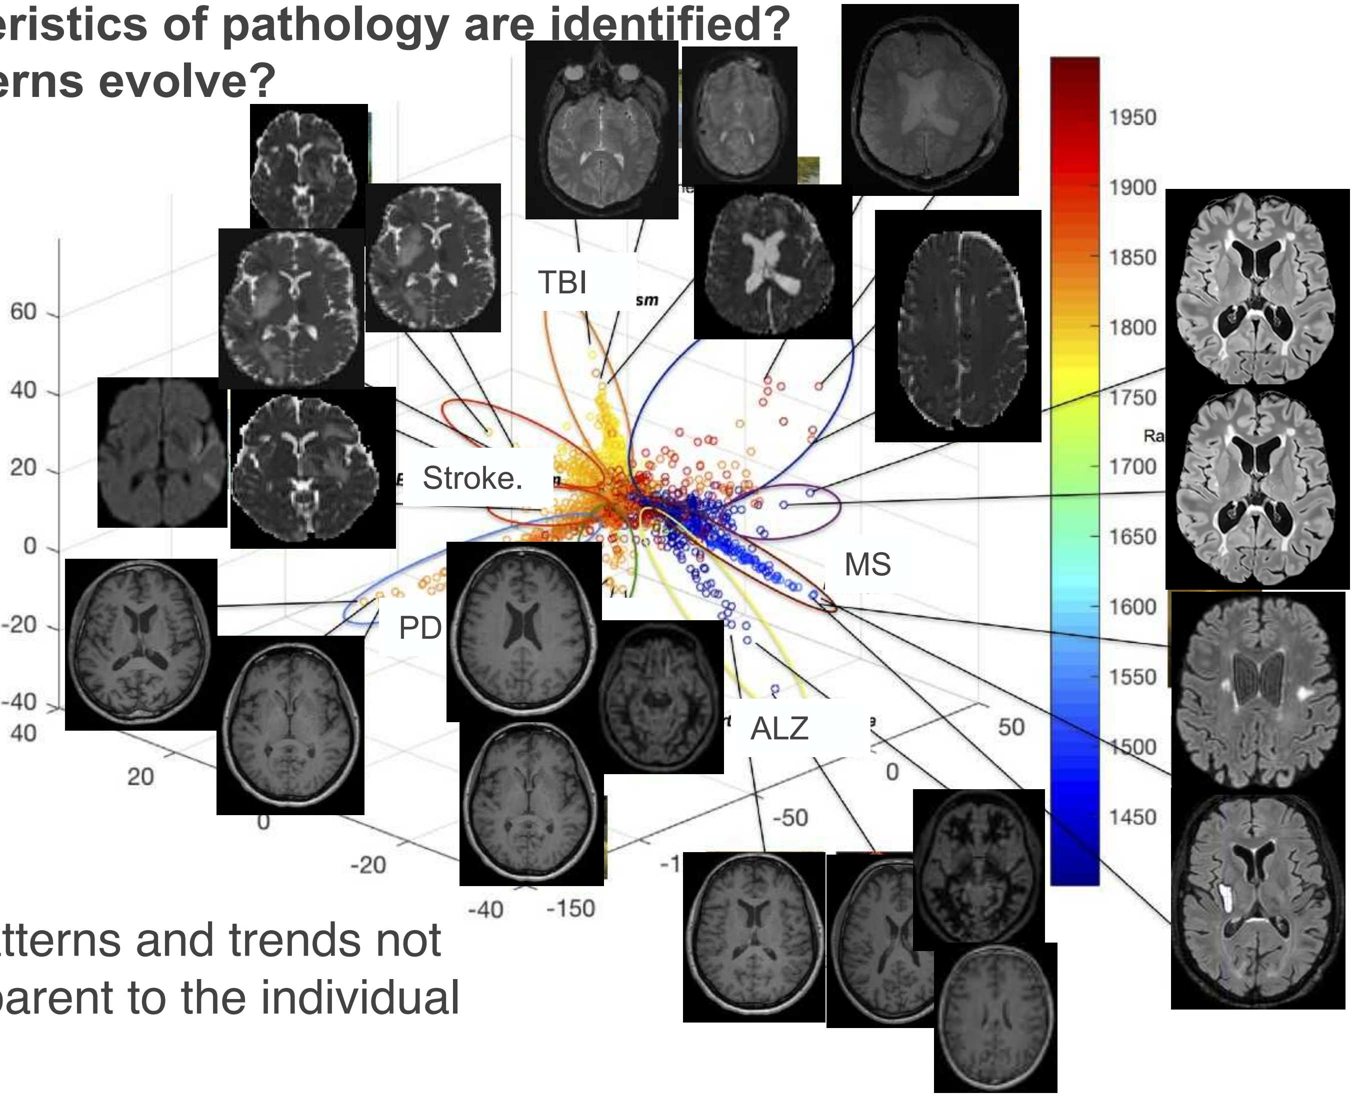

« to discover
fundamental patterns and trends not
necessarily apparent to the individual
human eye »

New models for neuroscience

How from the activity of millions of neurons distributed in several brain areas emerge simple percepts and how they are linked to emotion, motivation or action?

Experimentalists; e.g. extend or scale proof of concept finding

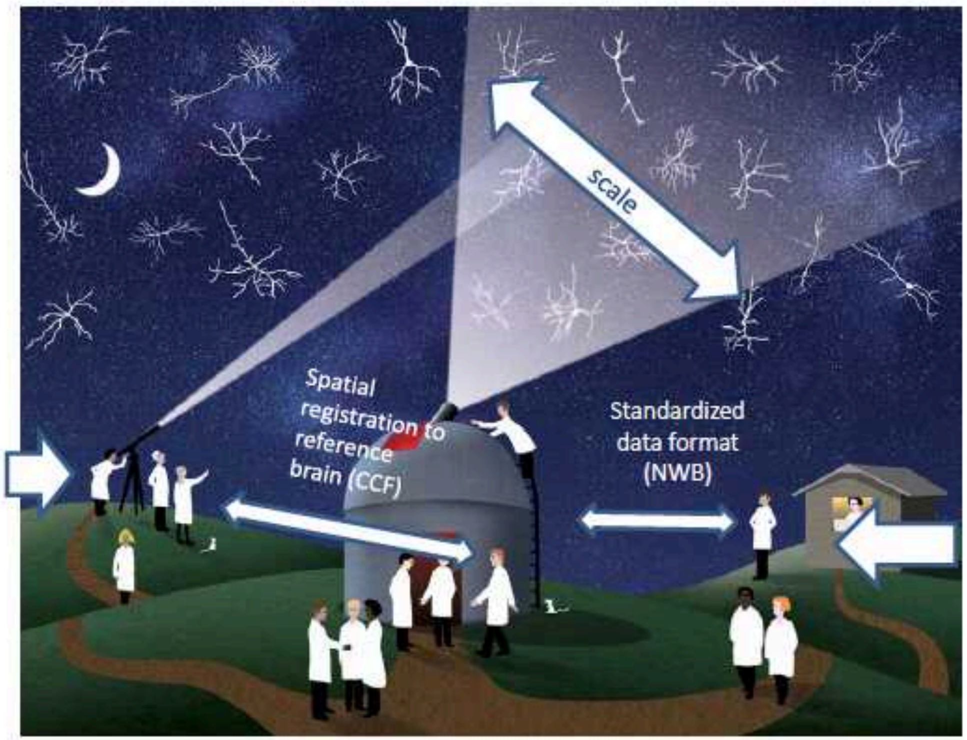

Big Data Analytics

Ontology Based Data Access

Researchers without access to sophisticated in vivo functional assays, Theoretical neuroscientists

from Mike Hawrylycz

Future

Mixing Hype AI (Supervised L+ **Renforcement L**) + Old fashion AI (Tree search)
[Silver et al Nature 2016]

Consciousness model

Soar:
[Newell 1983, 92]

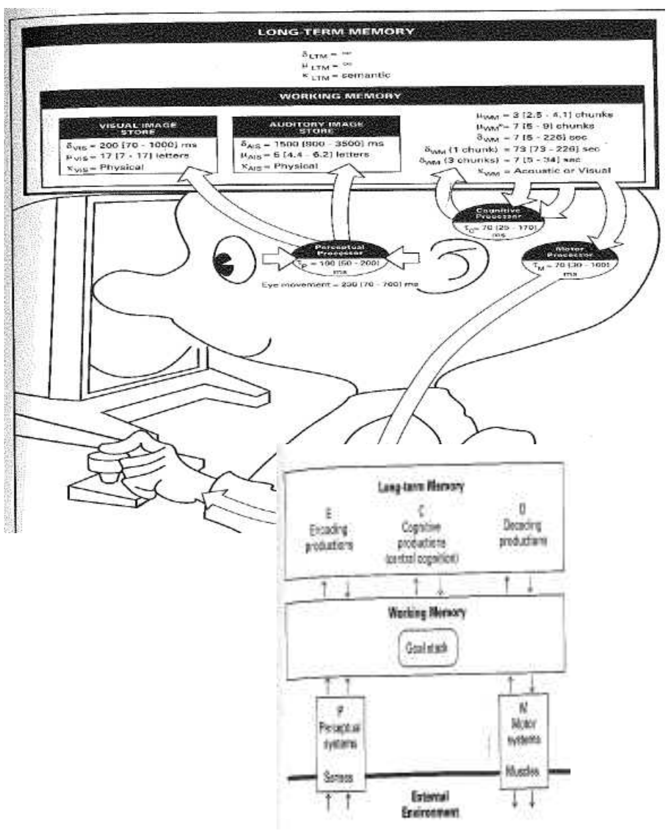

Global workspace theory [Baars 1988, 97, 2002]

Global ignition [Dehaene 2003]

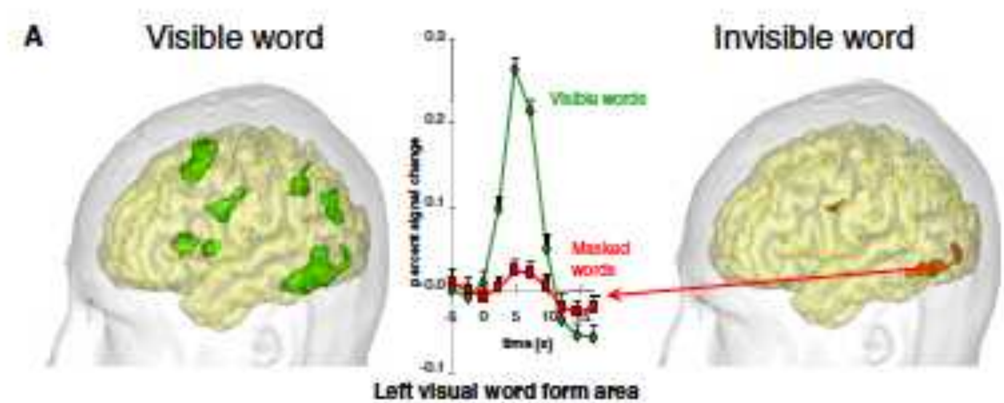

[Dehaene and Changeux Neuron 2011]

The Unreasonable Effectiveness of Data

- Classifiers based on millions of specific features perform better than elaborate models that try to discover general rules.

[Halevy et al IEEE Intell Syst 2009]

But …

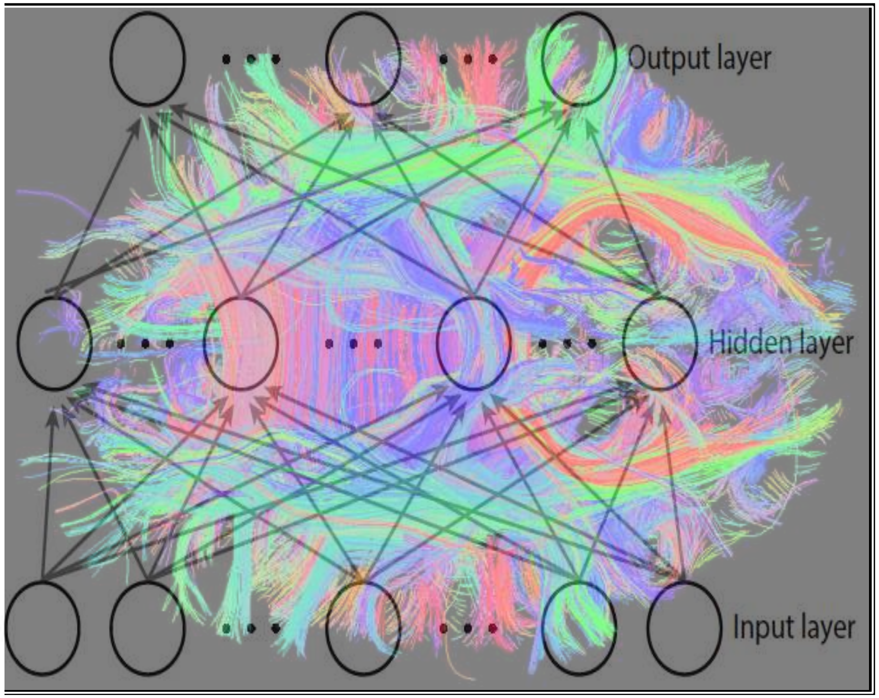

But ...

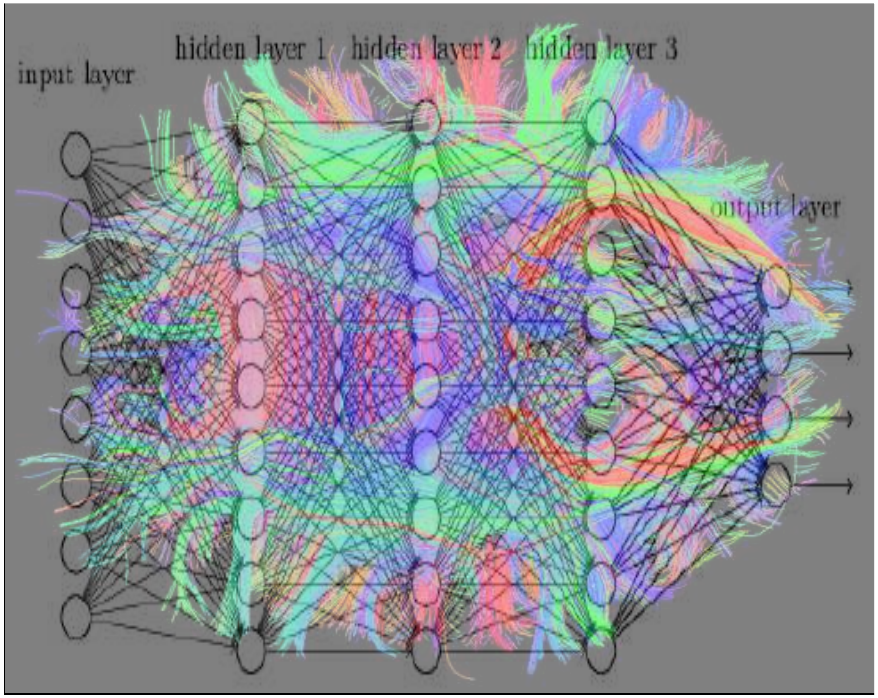

input layer
hidden layer 1
hidden layer 2
hidden layer 3
output layer


The End