PSIORI

$\Pr(E|H_p) \times \Pr(H_p)$

Siegen Lernen - Neuronale Netze in der Wirtschaft

Dr. Sascha Lange (sascha@psiori.com)

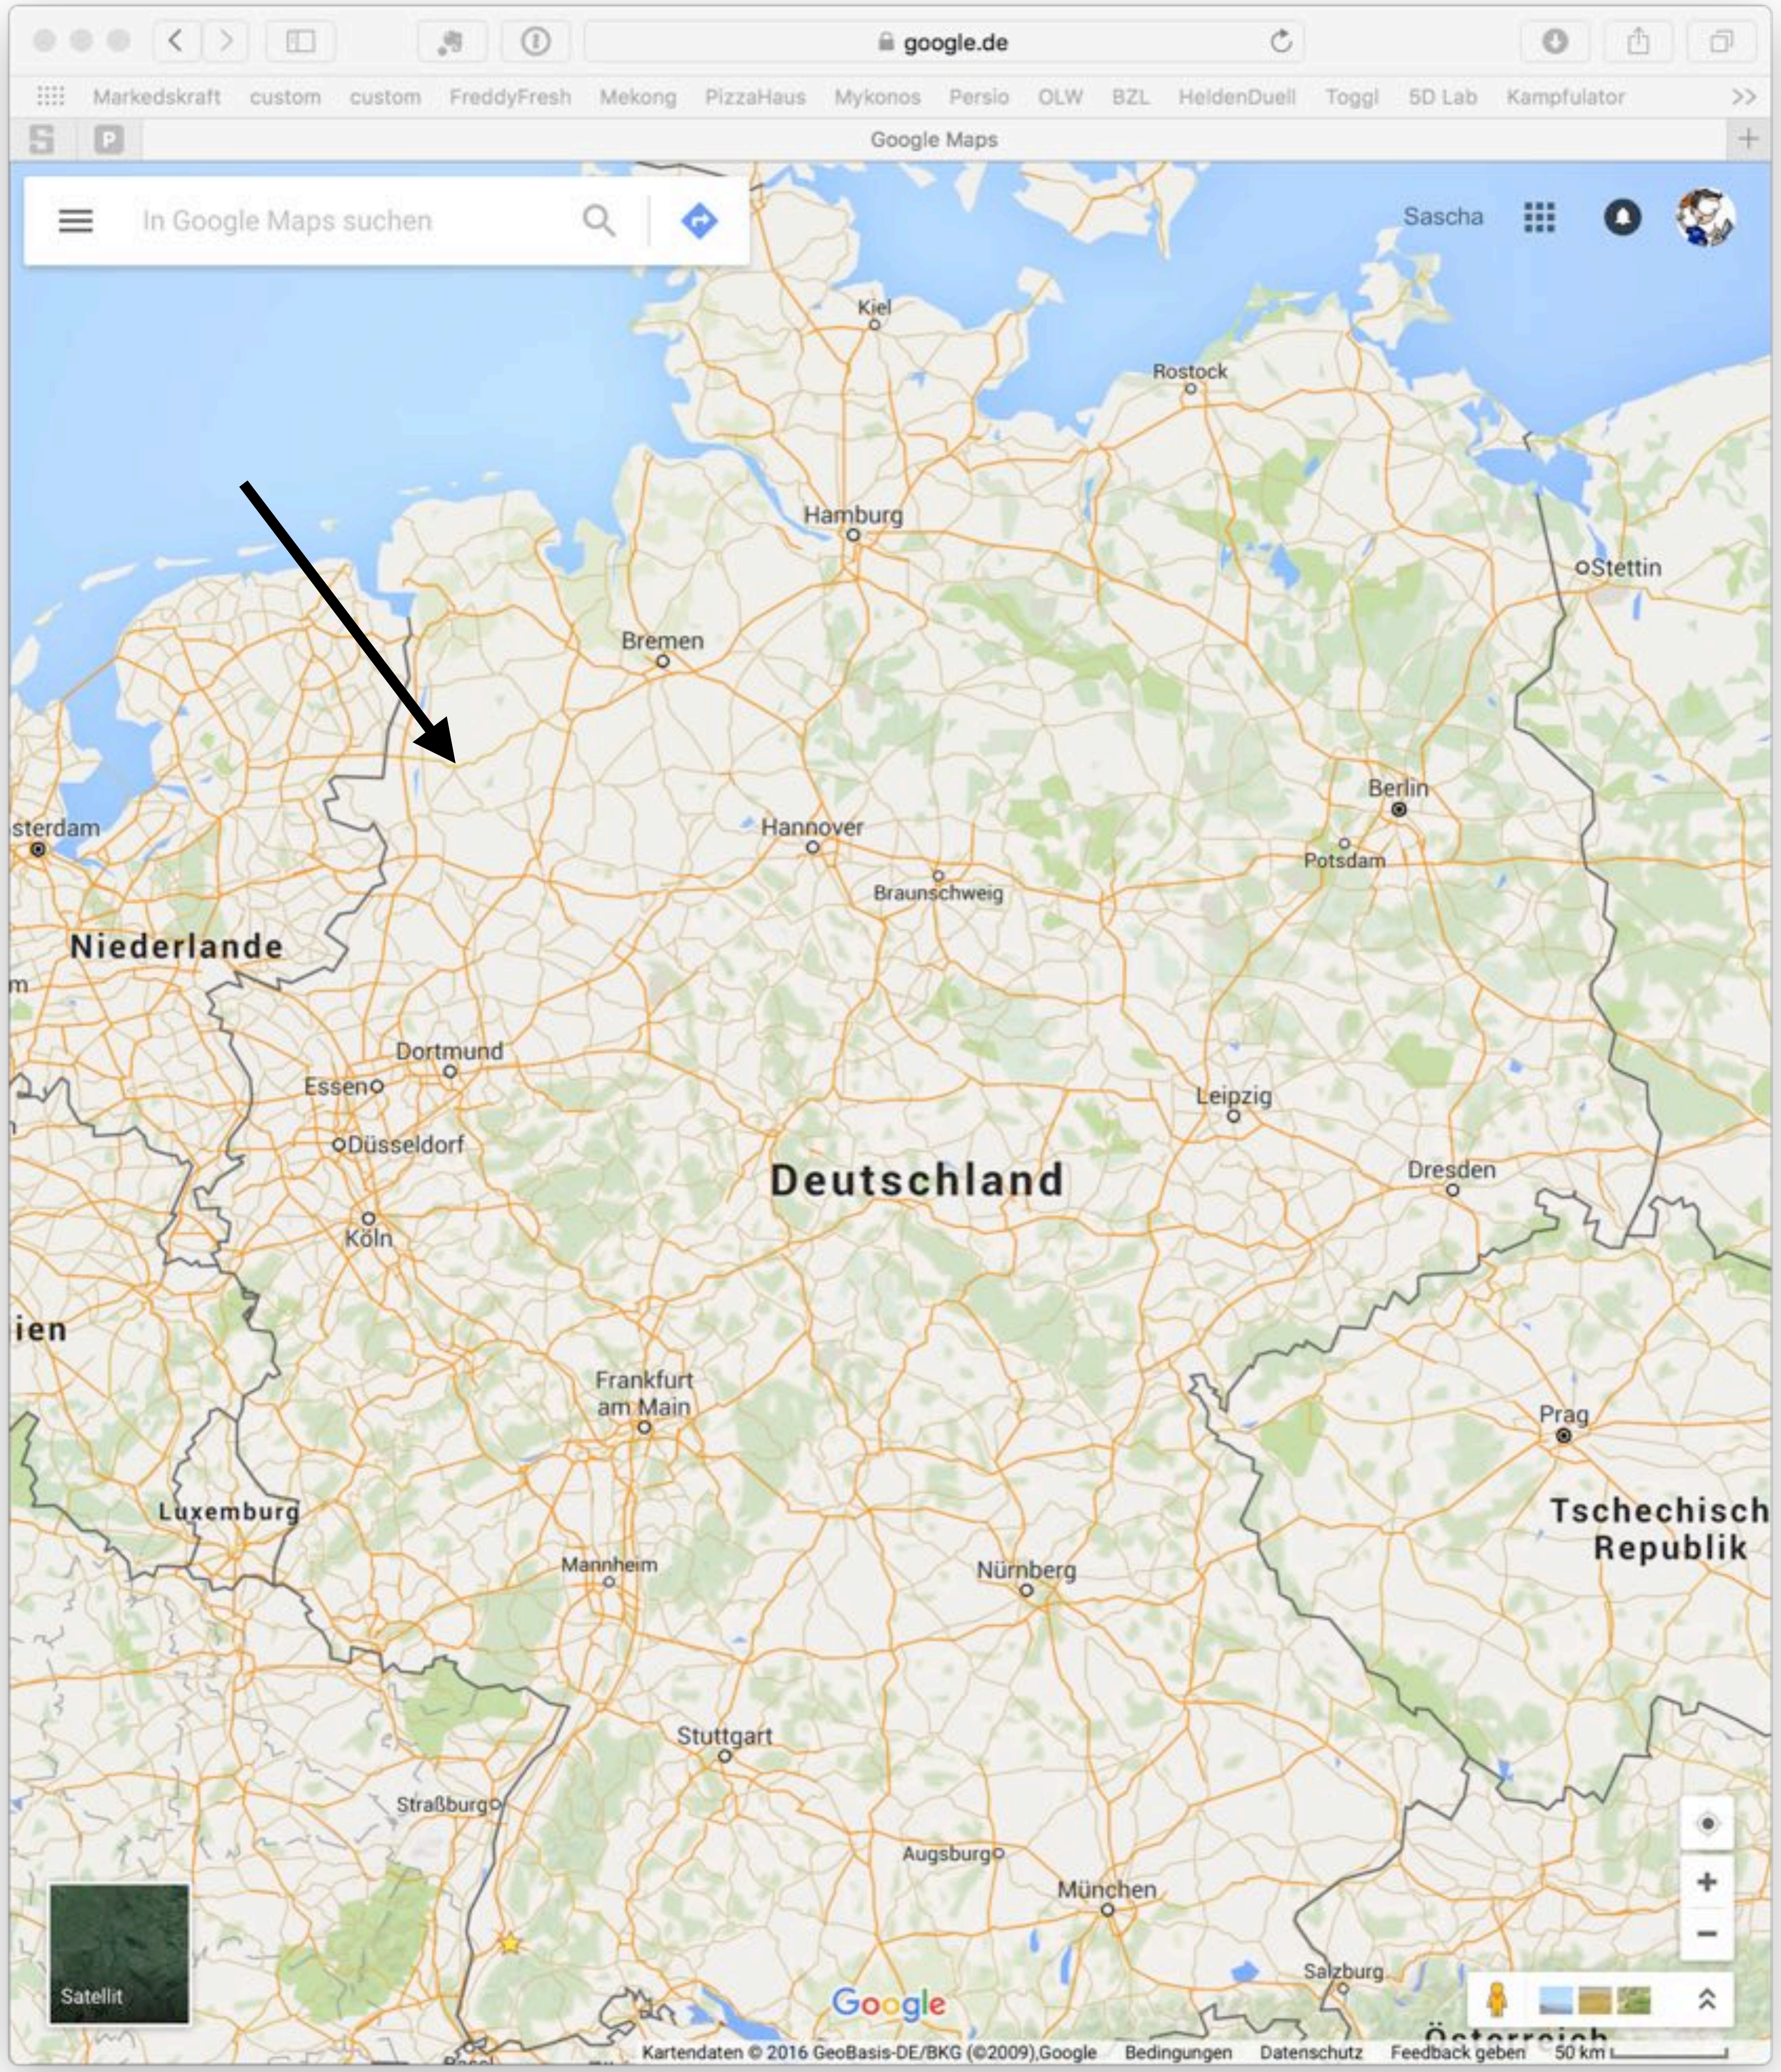
google.de
Markedskraft
custom
custom
FreddyFresh
Mekong
PizzaHaus
Mykonos
Persio
OLW
BZL
HeldenDuell
Toggl
5D Lab
Kampfulator
Google Maps
In Google Maps suchen
Sascha
Kiel
Rostock
Hamburg
Stettin
Bremen
Berlin
Hannover
Potsdam
Braunschweig
sterdam
Niederlande
Dortmund
Essen
Düsseldorf
Leipzig
Deutschland
Dresden
Köln
ien
Frankfurt am Main
Prag
Luxemburg
Tschechische Republik
Mannheim
Nürnberg
Stuttgart
Straßburg
Augsburg
München
Salzburg
Satellit
Google
Kartendaten © 2016 GeoBasis-DE/BKG (©2009),Google
Bedingungen
Datenschutz
Feedback geben
50 km

Machine Learning Lab
Founder: Prof. Dr. Martin Riedmiller

UNI FREIBURG

Resilient Propagation (1992)
Neural Fitted Q-Iteration (NN based Reinforcement Learning, 2005)
Deep Reinforcement Learning (2010)

FESTO · axel springer · KARMANN · BOSCH · DAIMLER · DEEPMIND · HARTING Pushing Performance · Fraunhofer · DFKI Deutsches Forschungszentrum für Künstliche Intelligenz GmbH · BrainLinks BrainTools acting.thoughts University of Freiburg

"Involved" @DeepMind, London, now Google / Alphabet daughter

bytroLabs
Felix Faber: BSc Thesis + Team Member
Tobias Kringe, Christopher Lörken: HiWi 2004

Kolja Hegelich
Team Captain Tribots
Head of Development
Twistbox Games

Stefan Welker
Arne Voigtländer
Christian Müller
Students & Tribots Team

Robot Soccer

2004-2008 DFG-Project &
Team: Brainstormers Tribots

World Champion
2006 and 2007

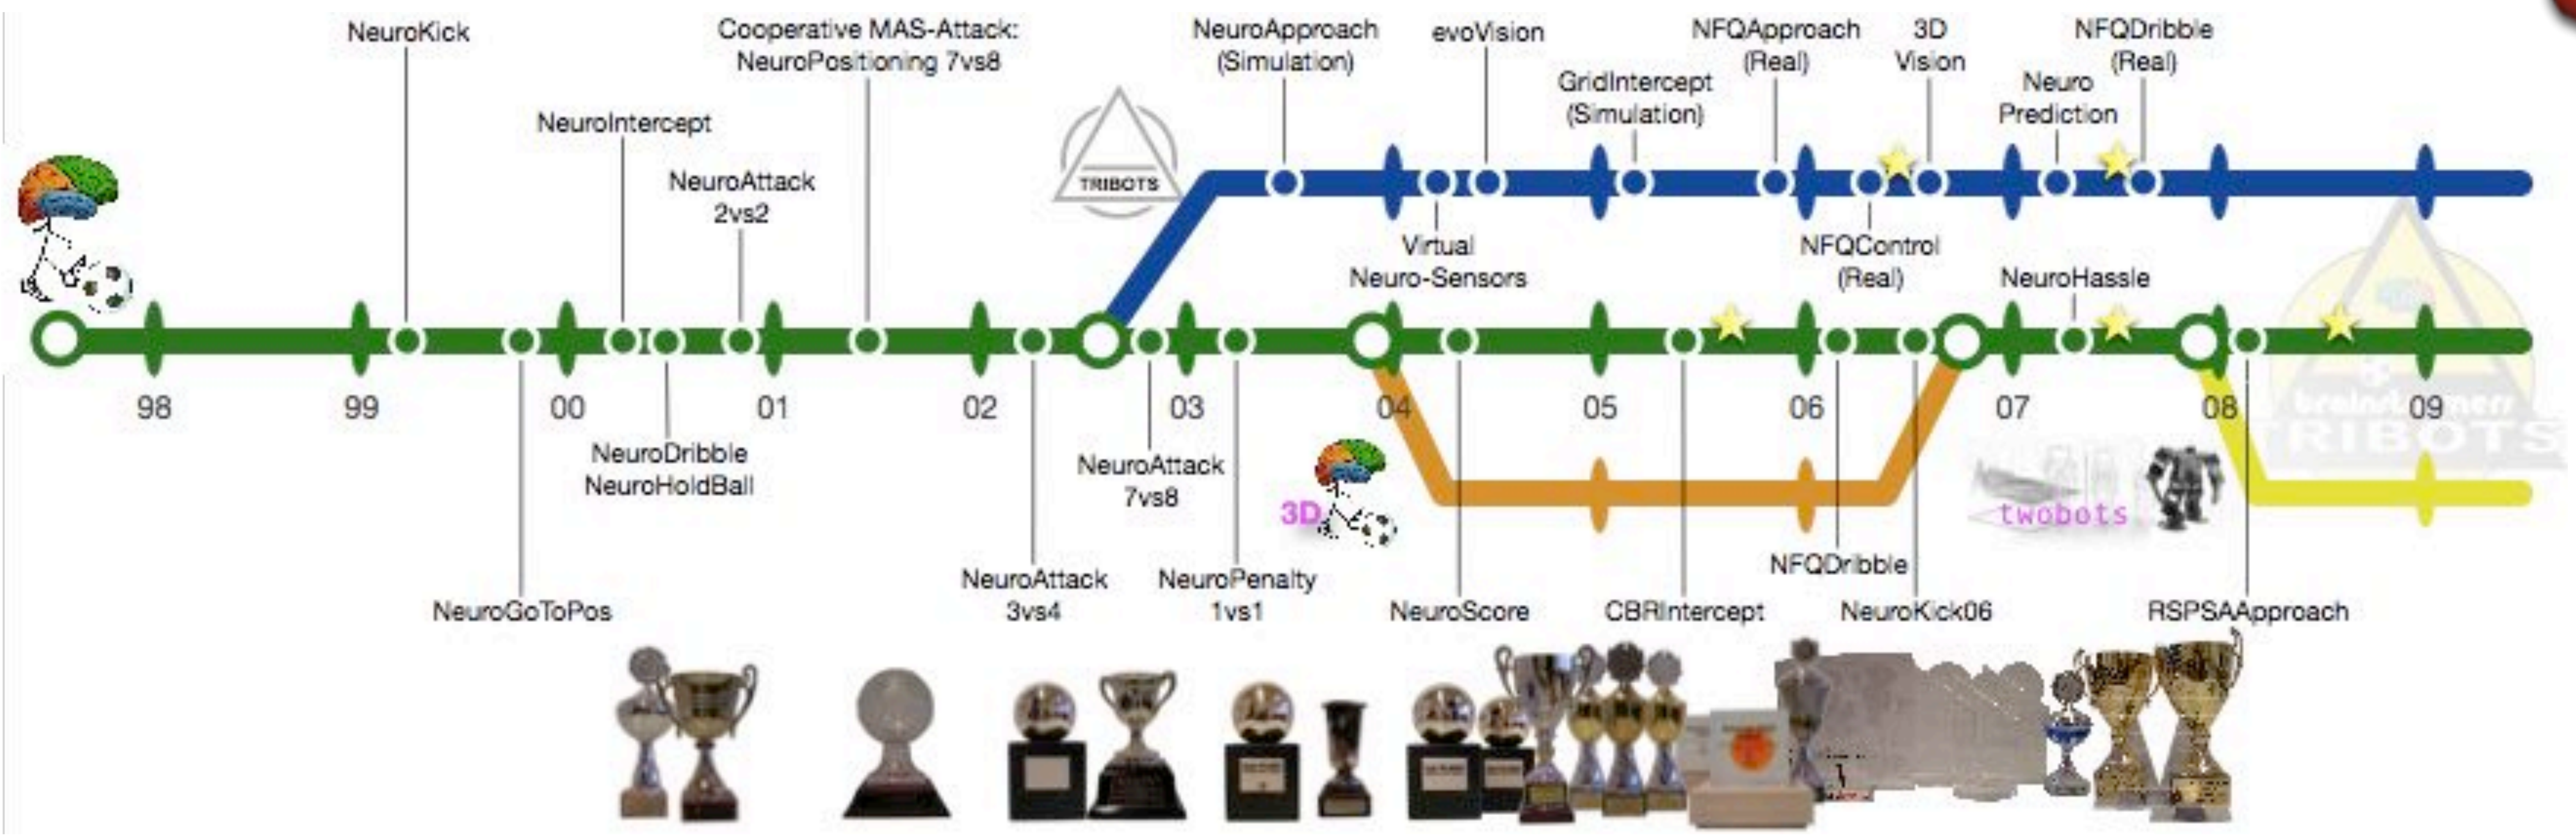

Deep Reinforcement Learning

RL directly on raw image data using Deep Neural Networks for learning optimal policies in latent feature spaces

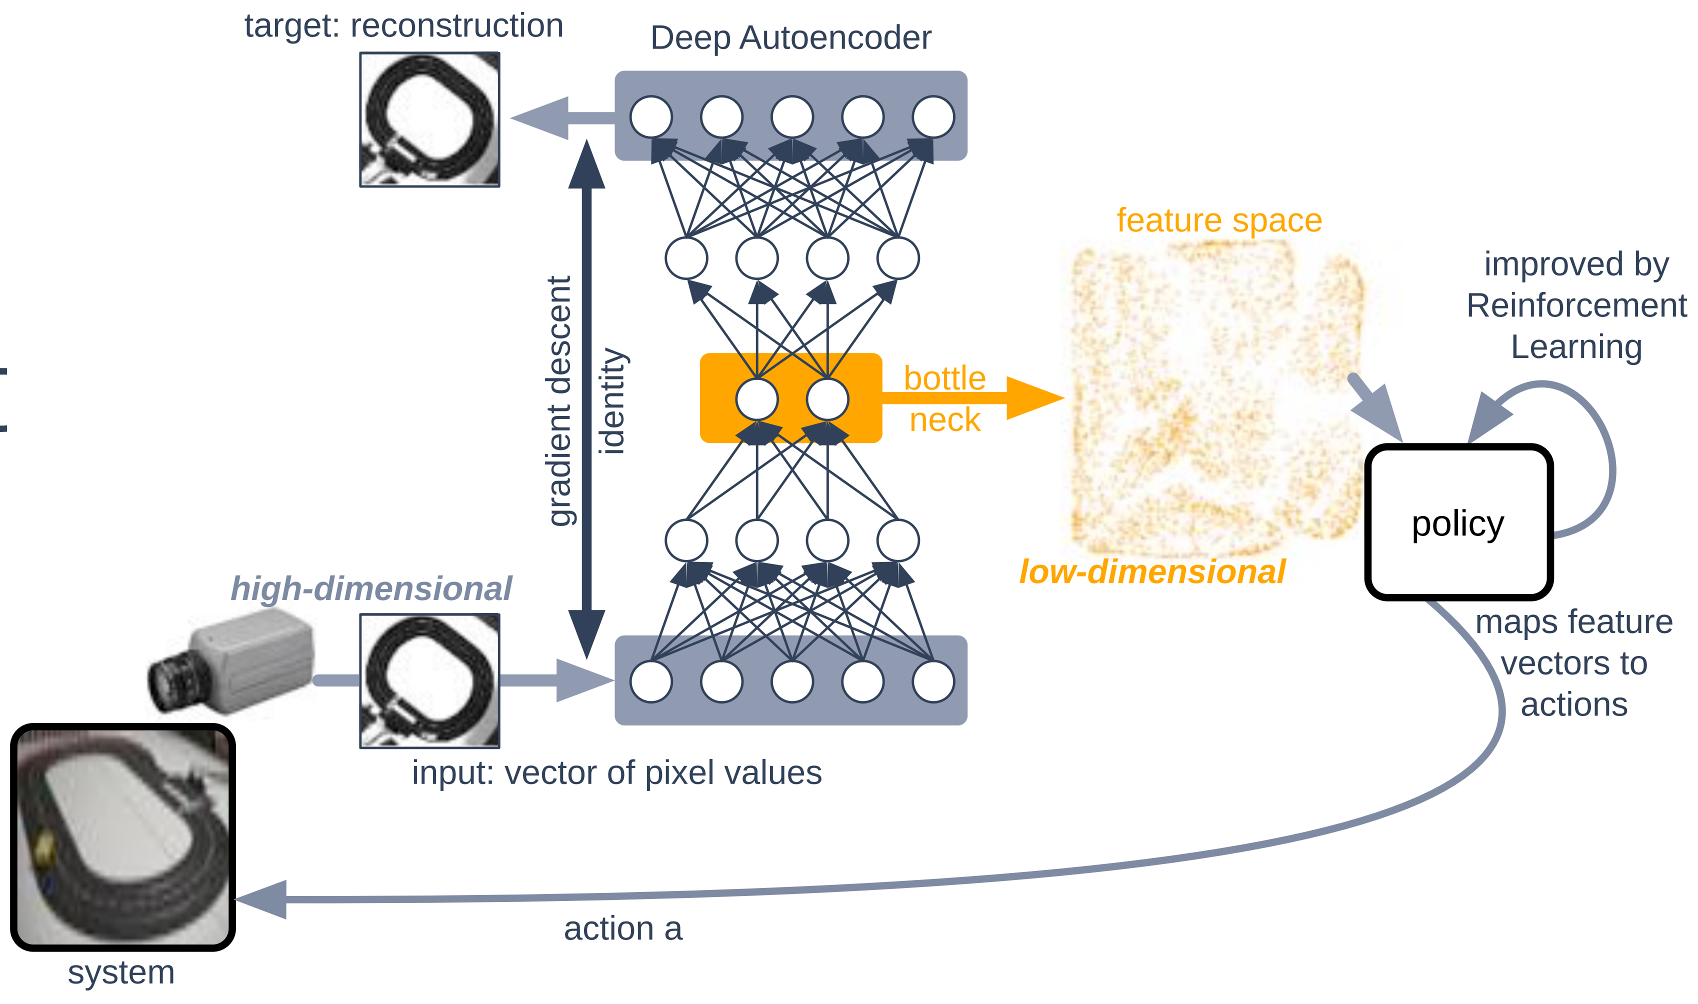

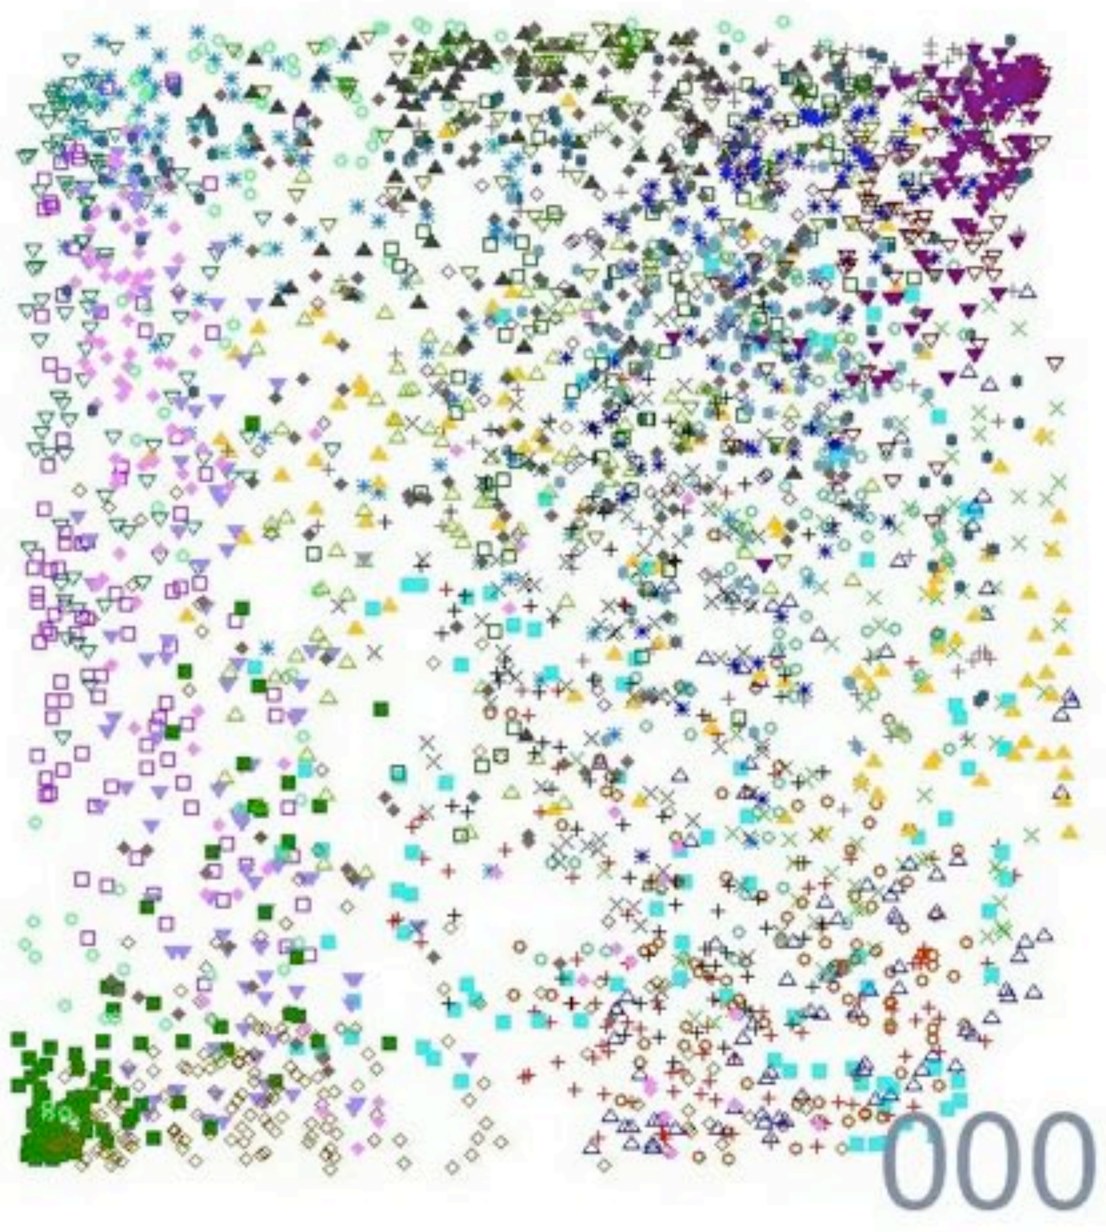

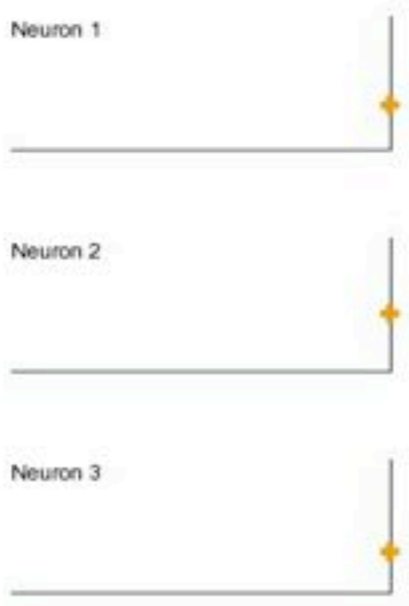

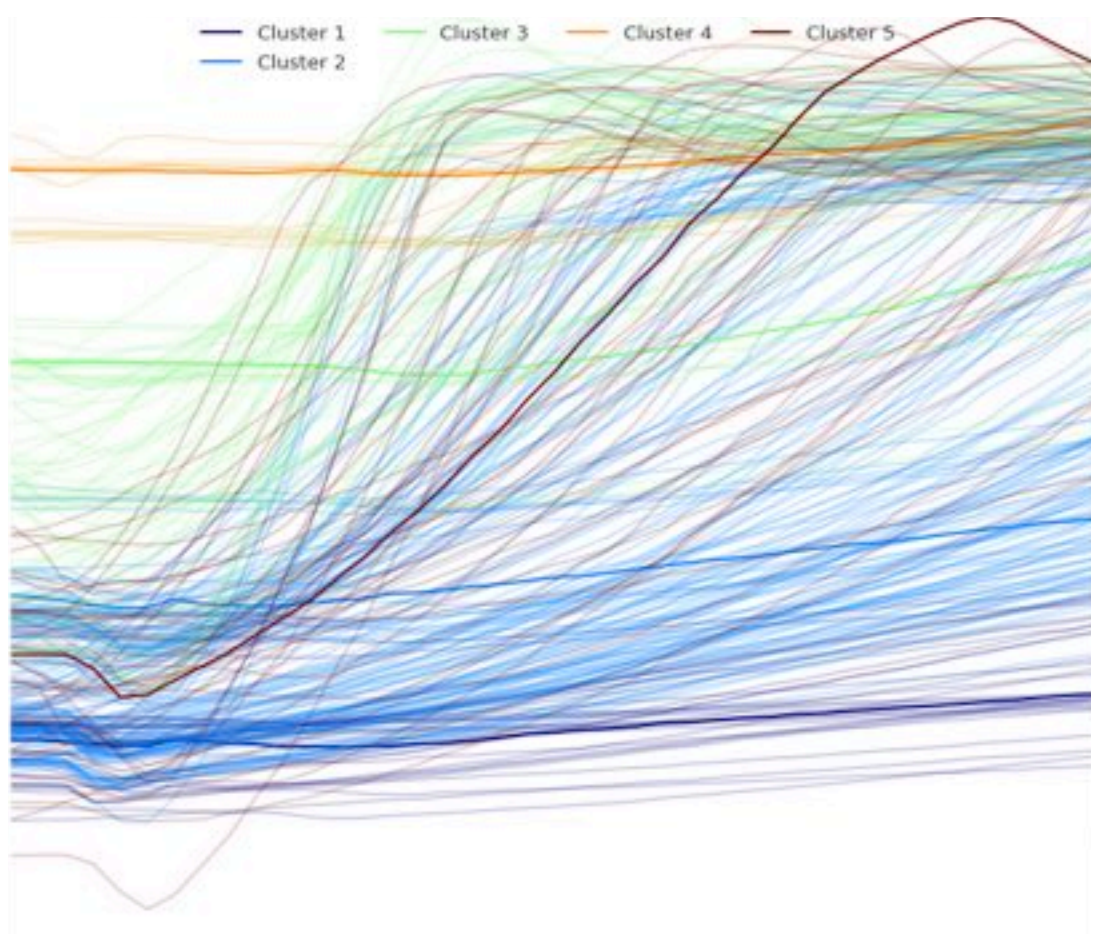

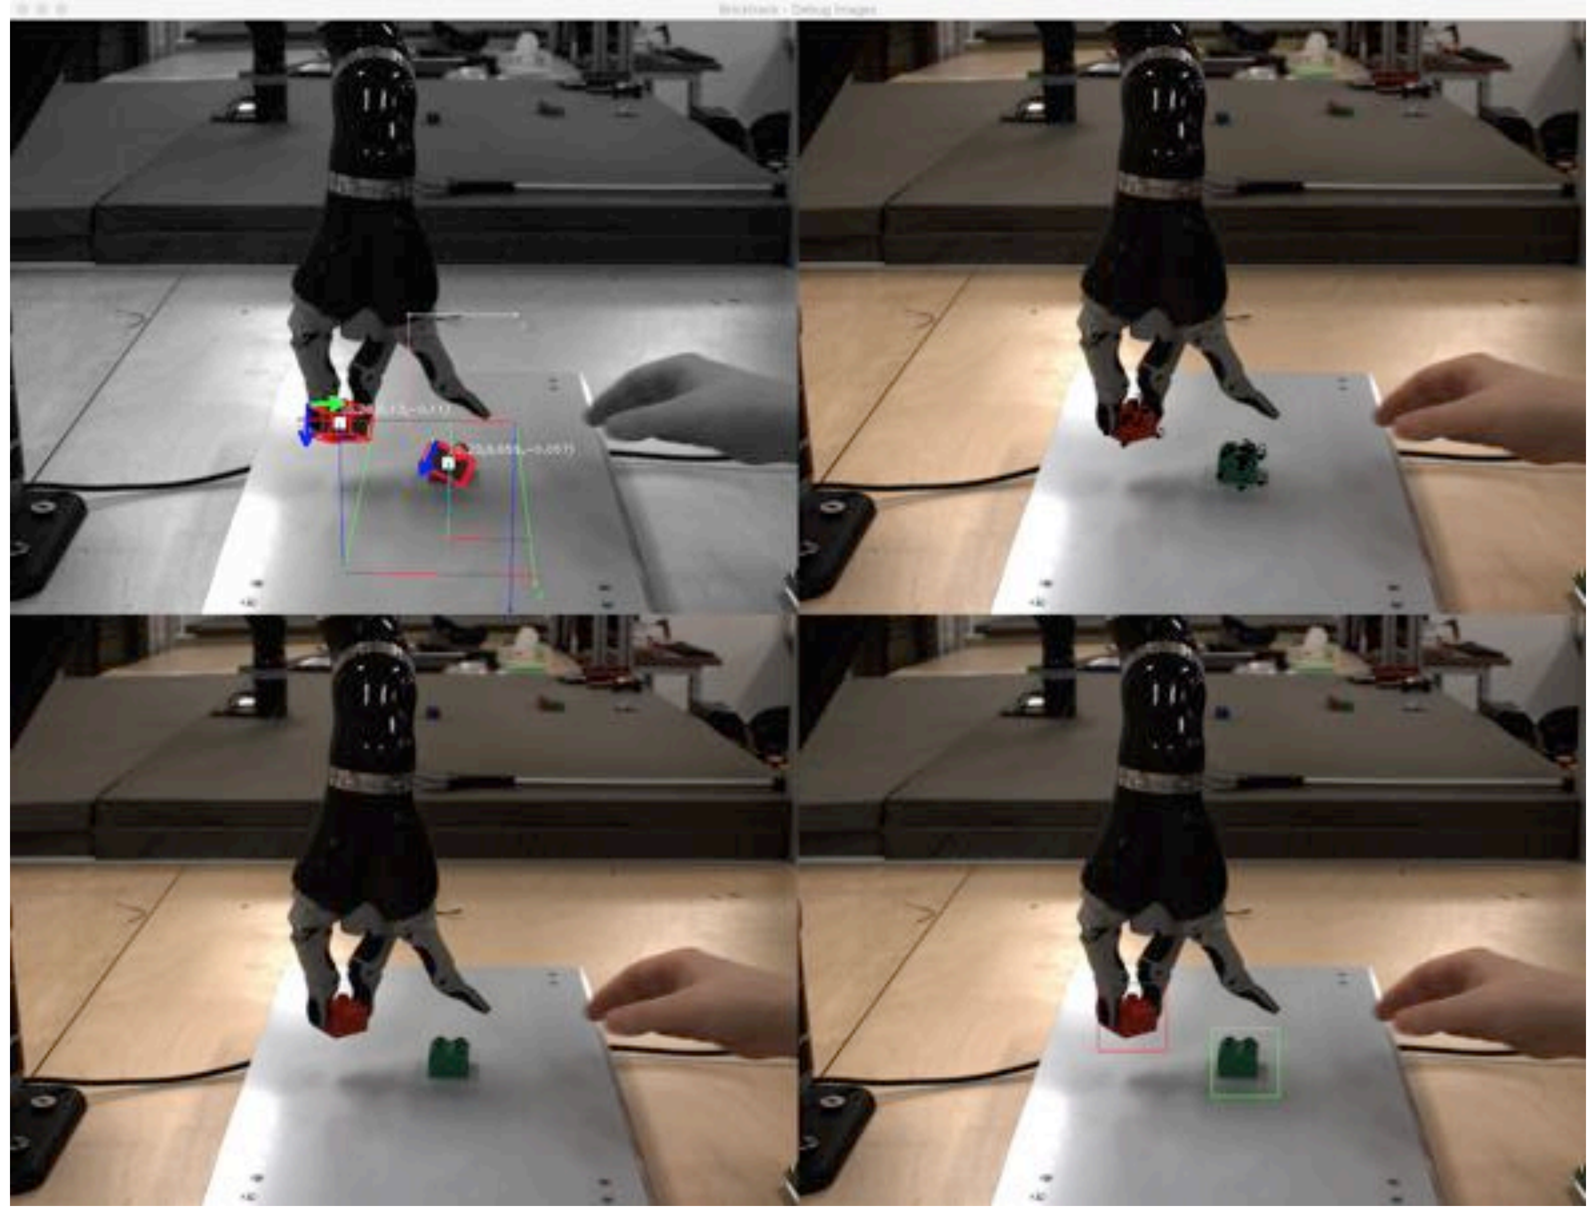

2014: "5d lab research", now: PSIORI

$\Pr(E|H_p) \times \Pr(H_p)$


ENIT SYSTEMS
CAFE &
BISTRO

pmOne
HOMAG
ProSiebenSat.1 Media AG
arotop food & environment GmbH
DUNKIN' DONUTS
SMART DATAONE HIGH QUALITY PRODUCT INFORMATION
NETCHECK
kr3m
HTTC GMBH & CO. KG
PAWISDA SYSTEMS
meinestadt.de
arvato BERTELSMANN
Imperia
Google DeepMind
Audi
ThyssenKrupp
metricminds
DÜRR
Kennedy Space Center VISITOR COMPLEX
Moguntia
nexum Consulting & Design
spilgames
automotive engineering iav
BARCO
fluxguide
igus
WORLD SYNC

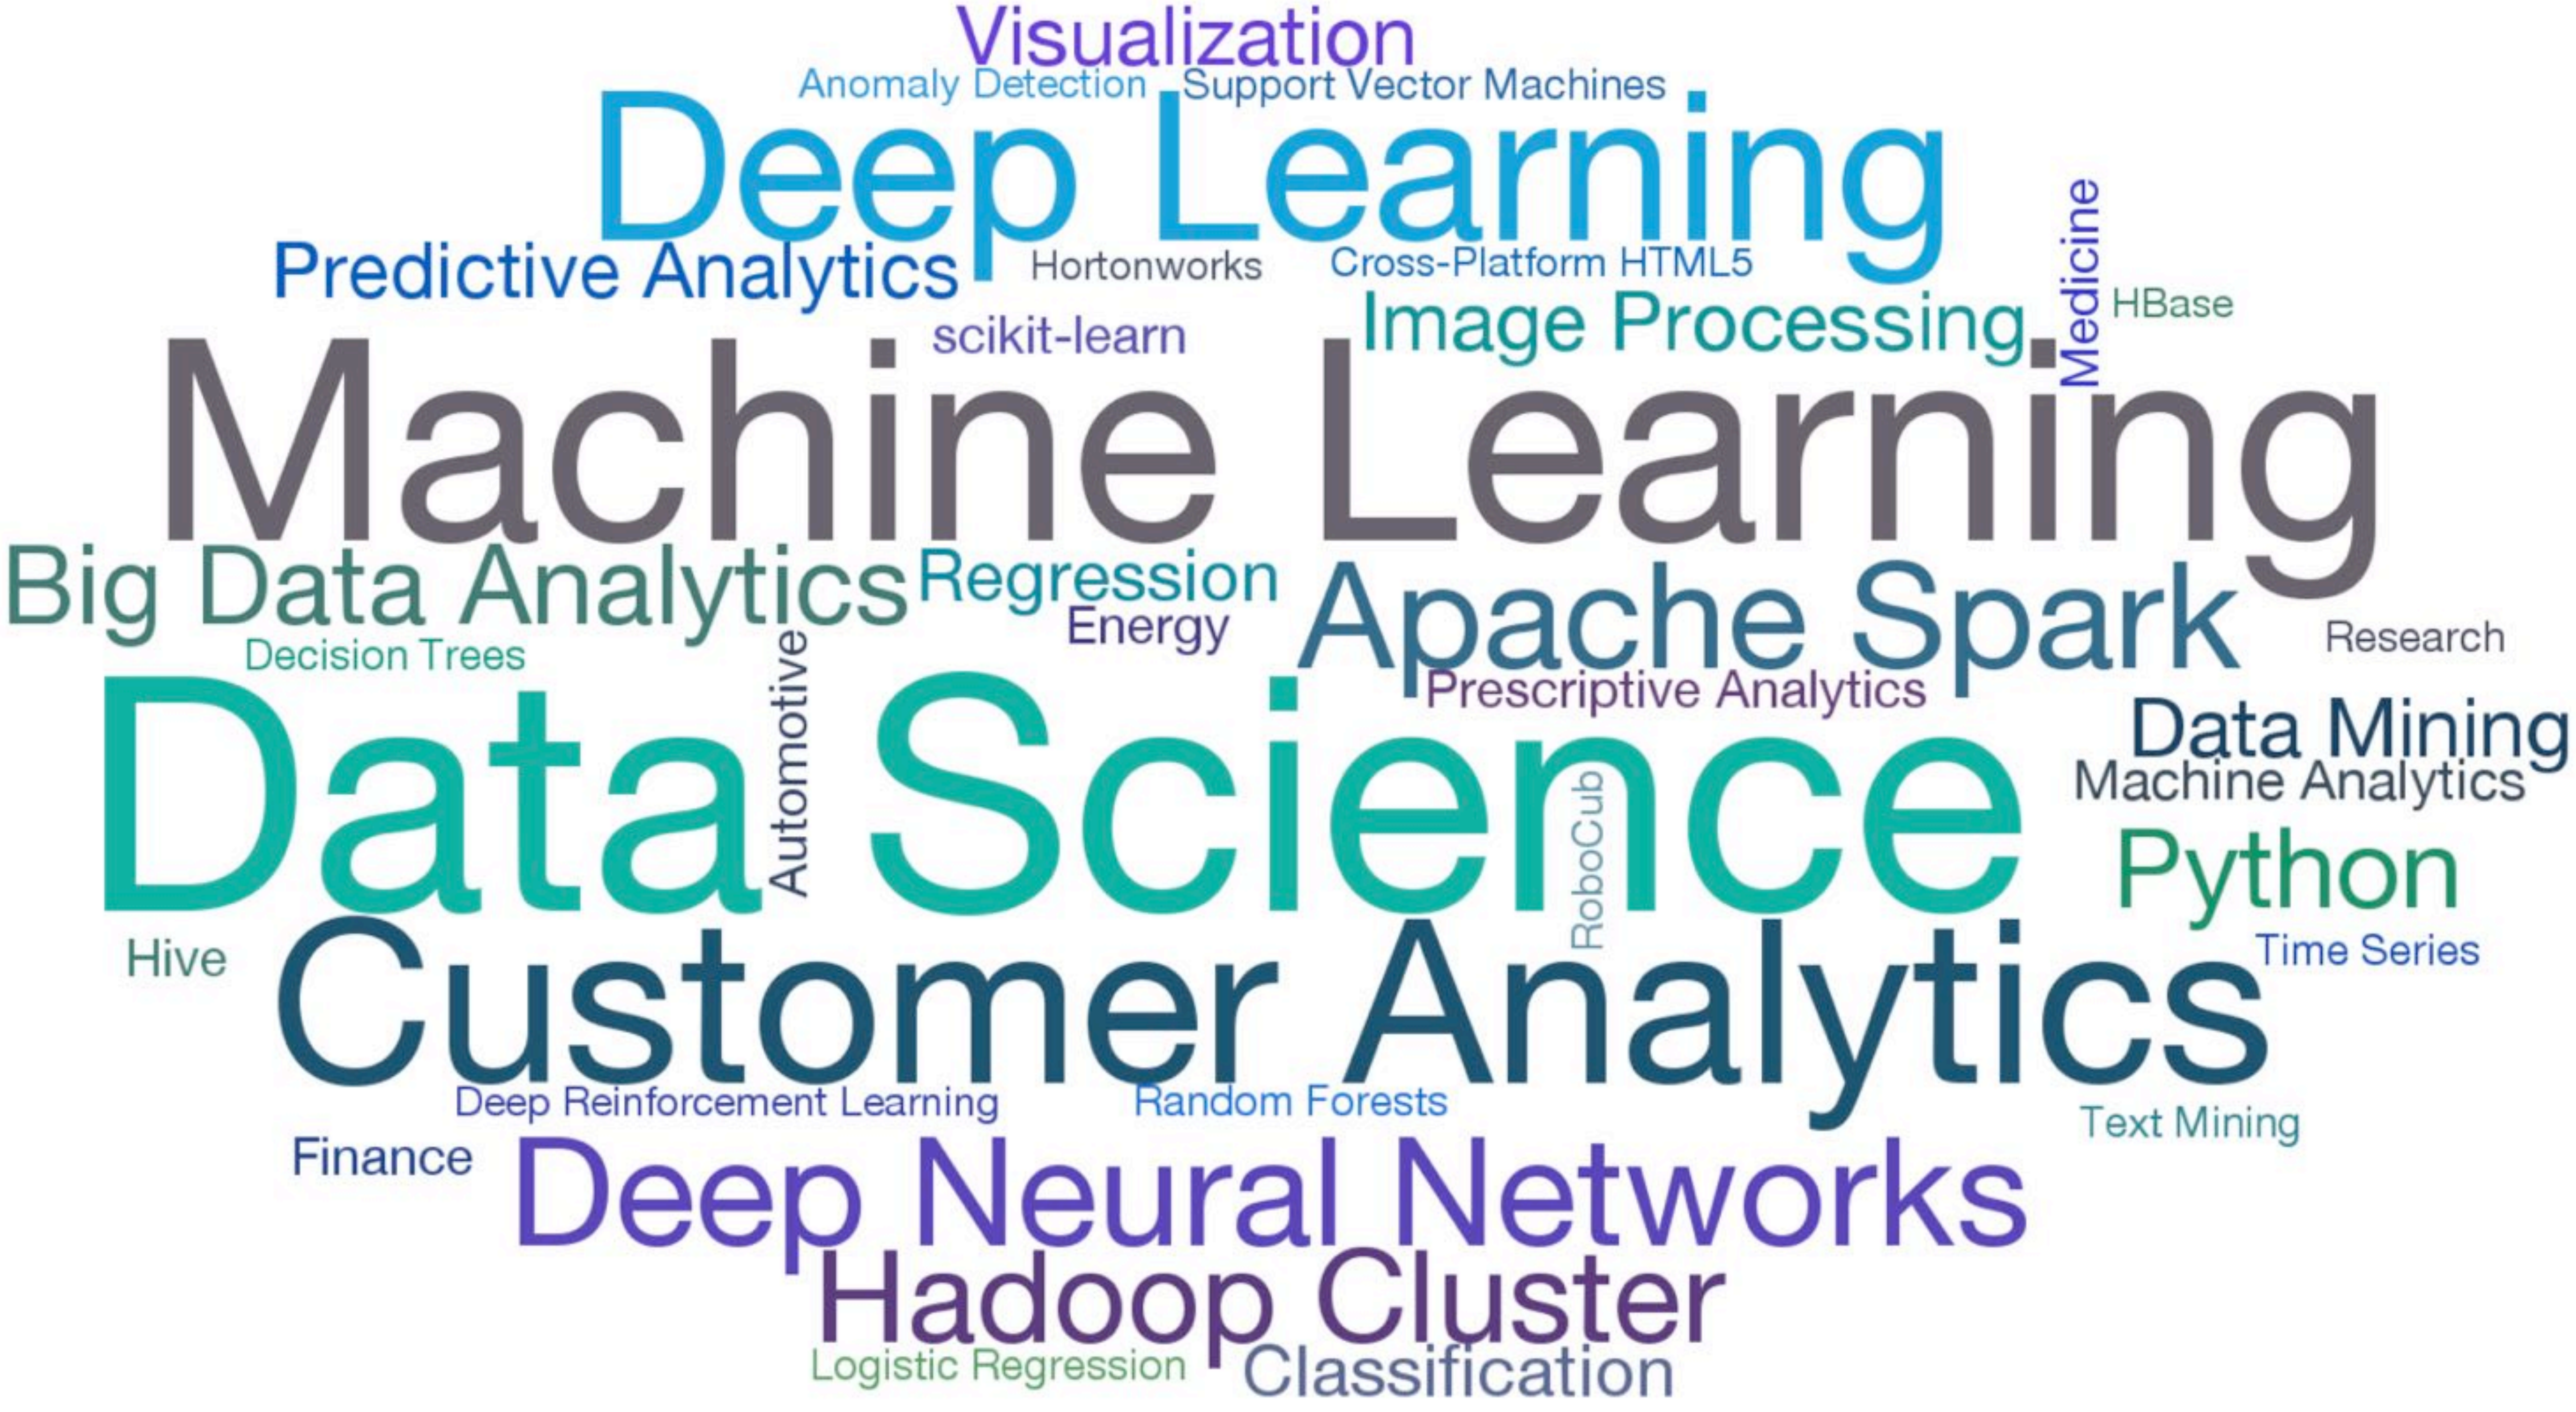
Visualization
Anomaly Detection
Support Vector Machines
Deep Learning
Predictive Analytics
Hortonworks
Cross-Platform HTML5
Medicine
HBase
scikit-learn
Image Processing
Machine Learning
Big Data Analytics
Regression
Apache Spark
Energy
Research
Decision Trees
Automotive
Prescriptive Analytics
Data Mining
Data Science
Machine Analytics
RoboCub
Python
Hive
Customer Analytics
Time Series
Deep Reinforcement Learning
Random Forests
Text Mining
Finance
Deep Neural Networks
Hadoop Cluster
Logistic Regression
Classification

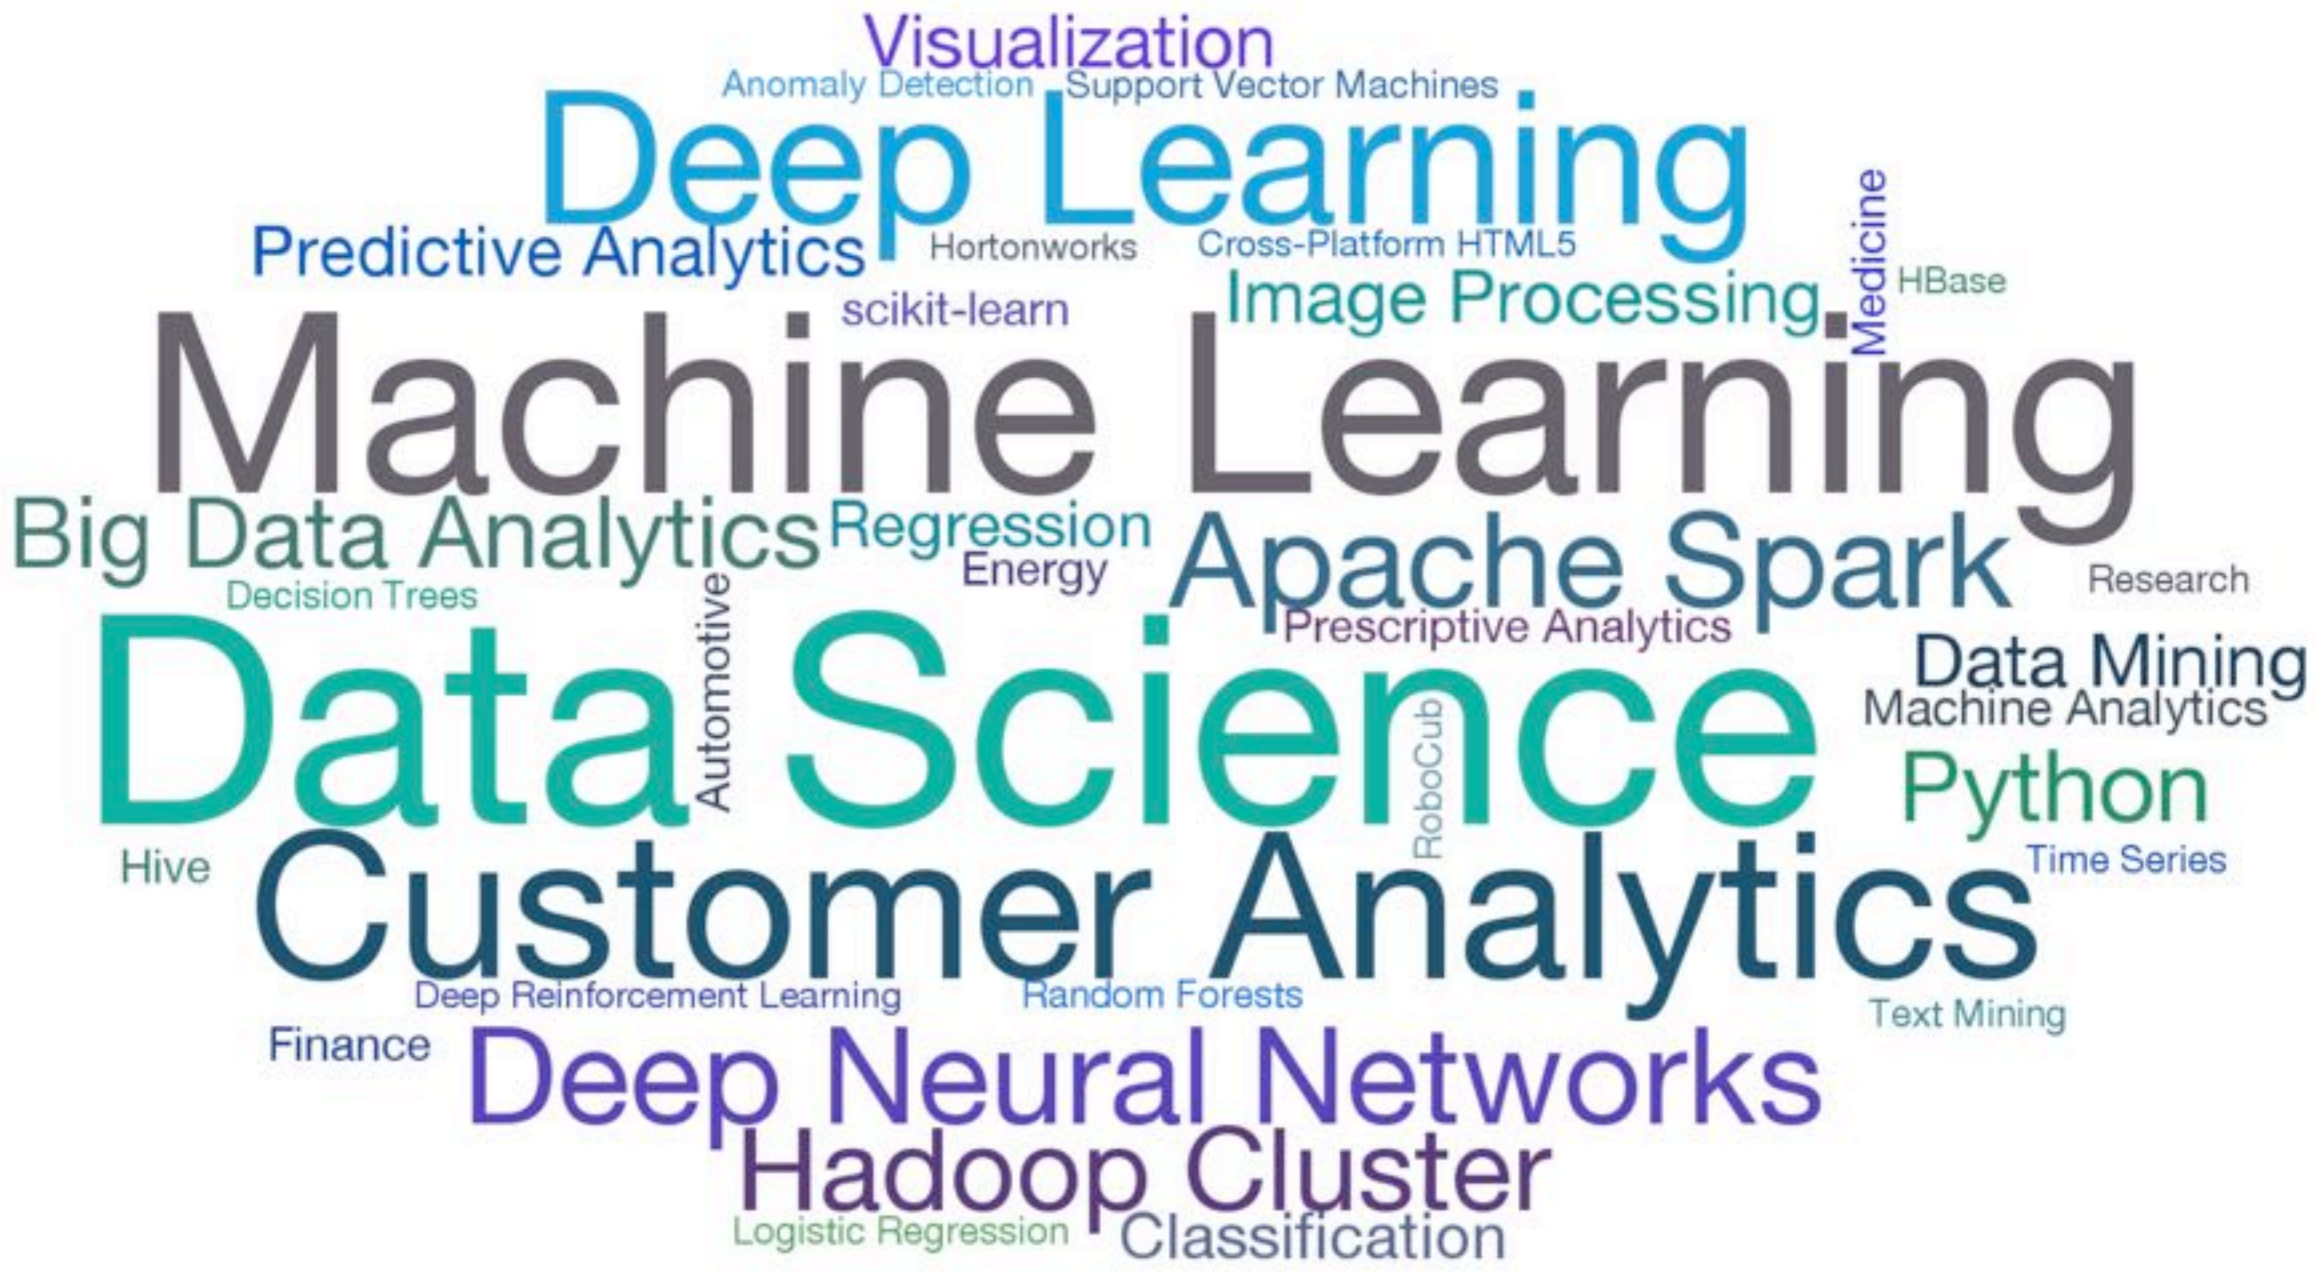

Outline

- Part I: The Research Perspective:

  Why is Artificial Intelligence hyped **right now**?
  What is Machine Learning?
  What is Deep Learning?
  What is Deep Reinforcement Learning?

- Part 2: The Business Perspective:

  Applications:
  Visualizing
  Predicting
  (Acting)

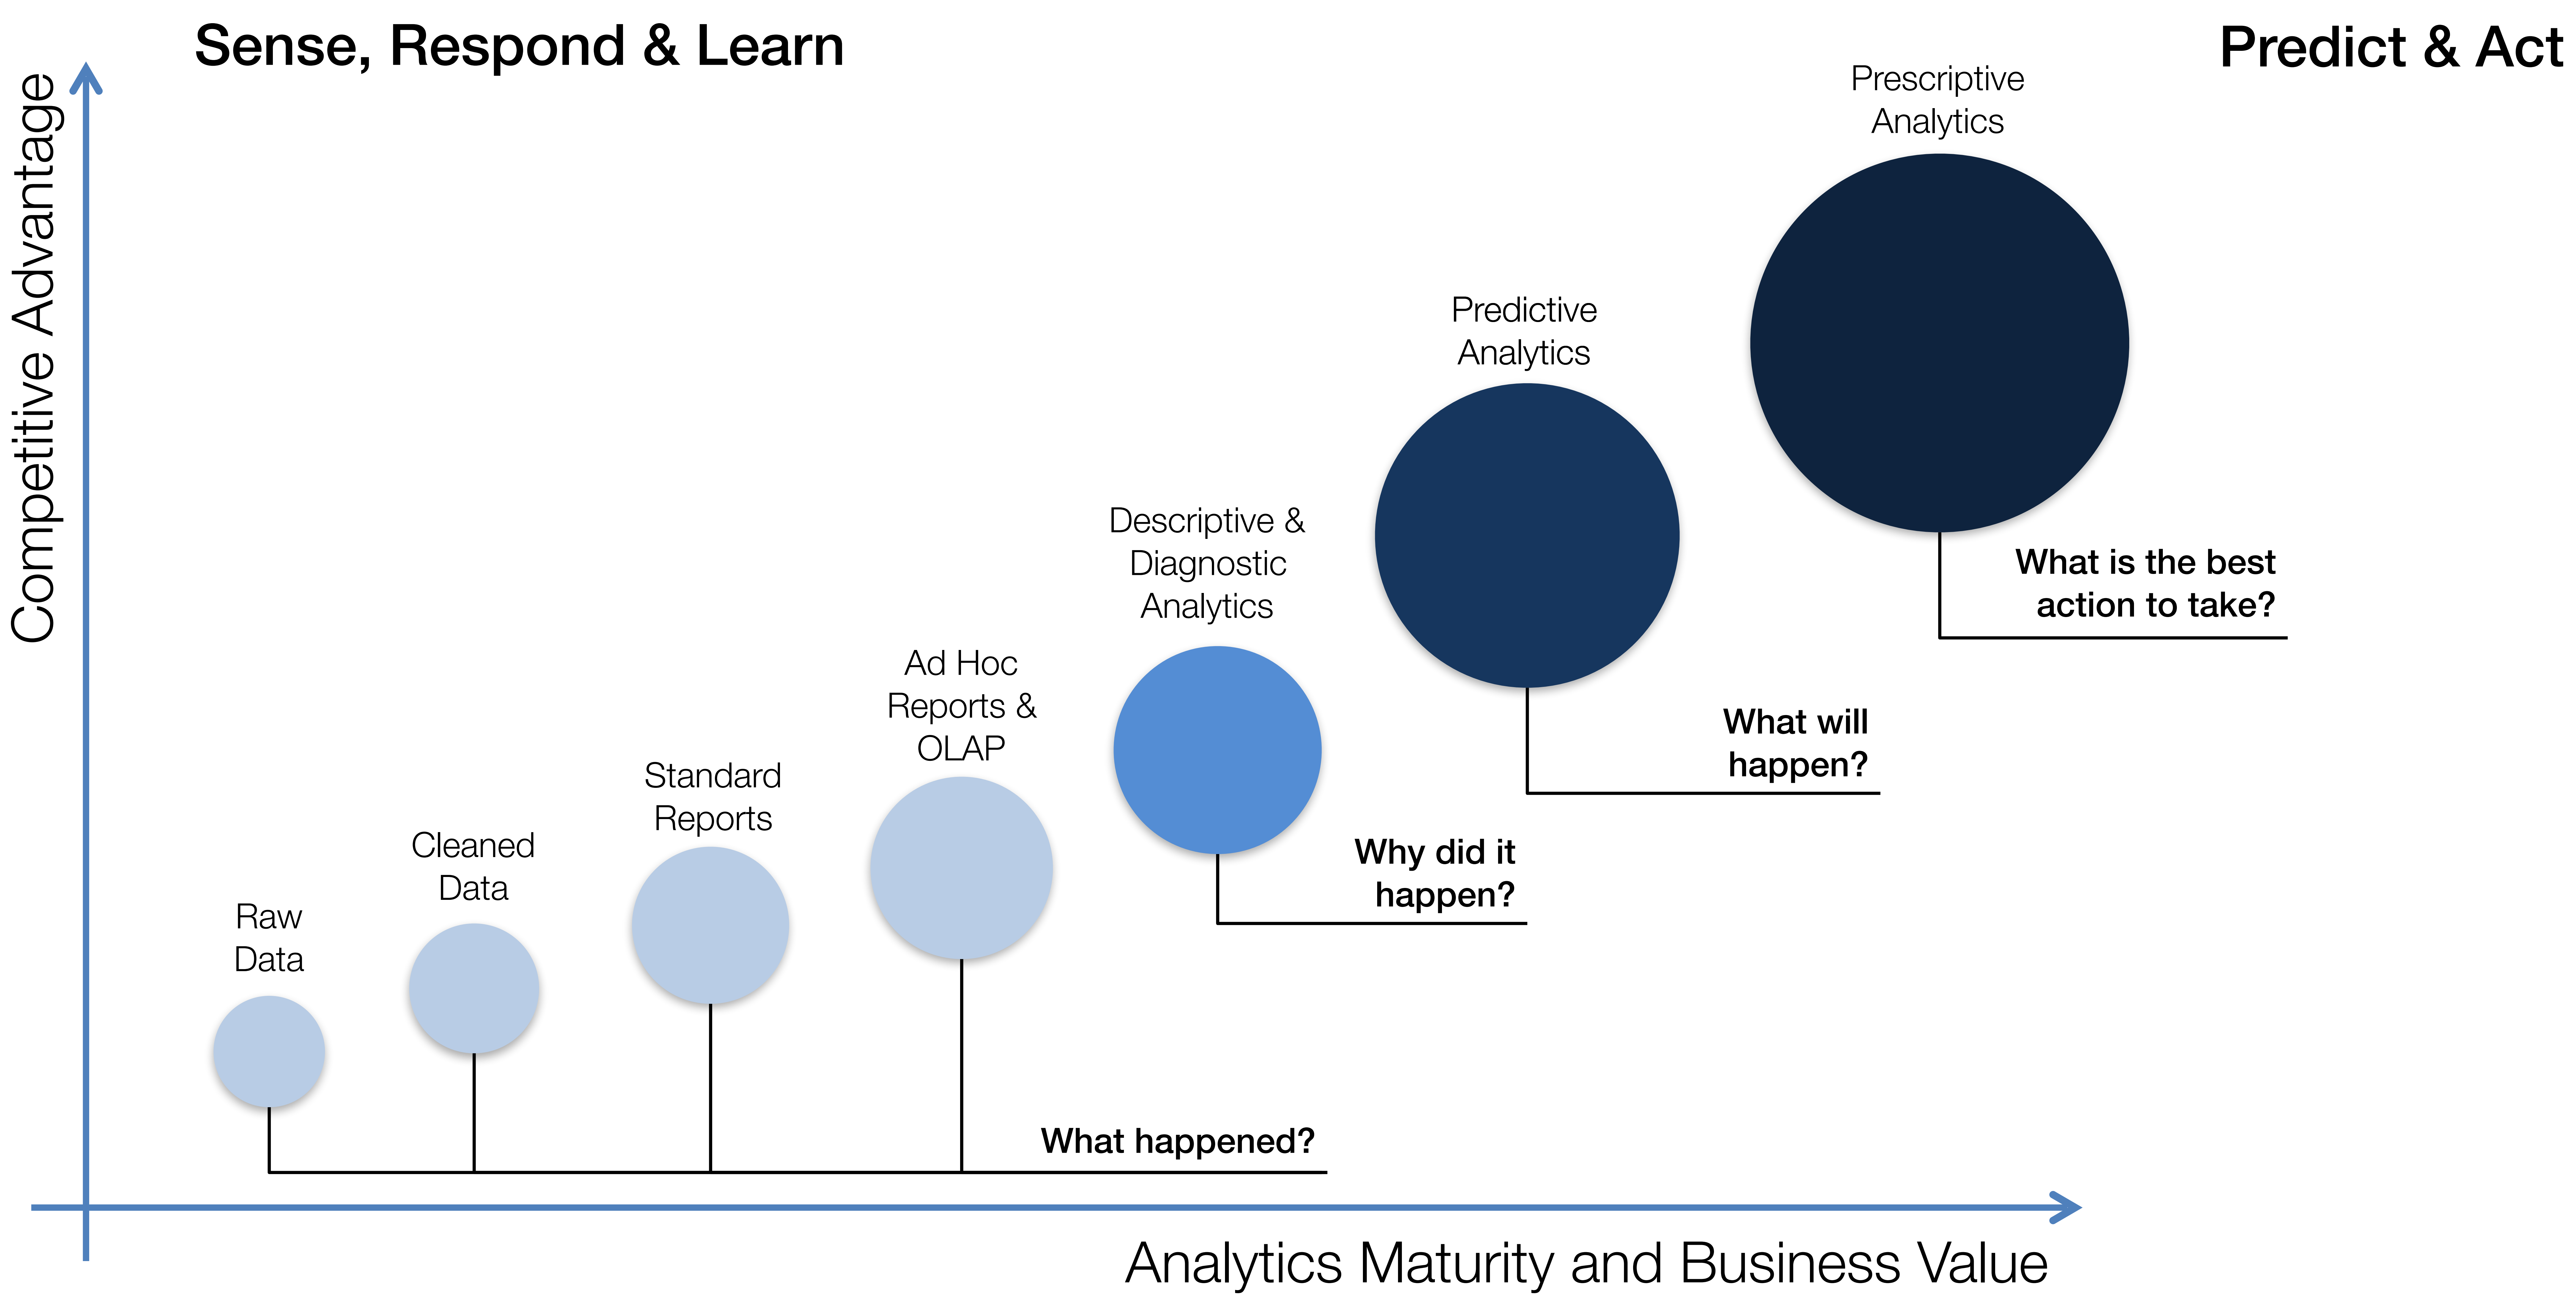
Sense, Respond & Learn
Predict & Act
Competitive Advantage
Analytics Maturity and Business Value
Raw Data
Cleaned Data
Standard Reports
Ad Hoc Reports & OLAP
What happened?
Descriptive & Diagnostic Analytics
Why did it happen?
Predictive Analytics
What will happen?
Prescriptive Analytics
What is the best action to take?

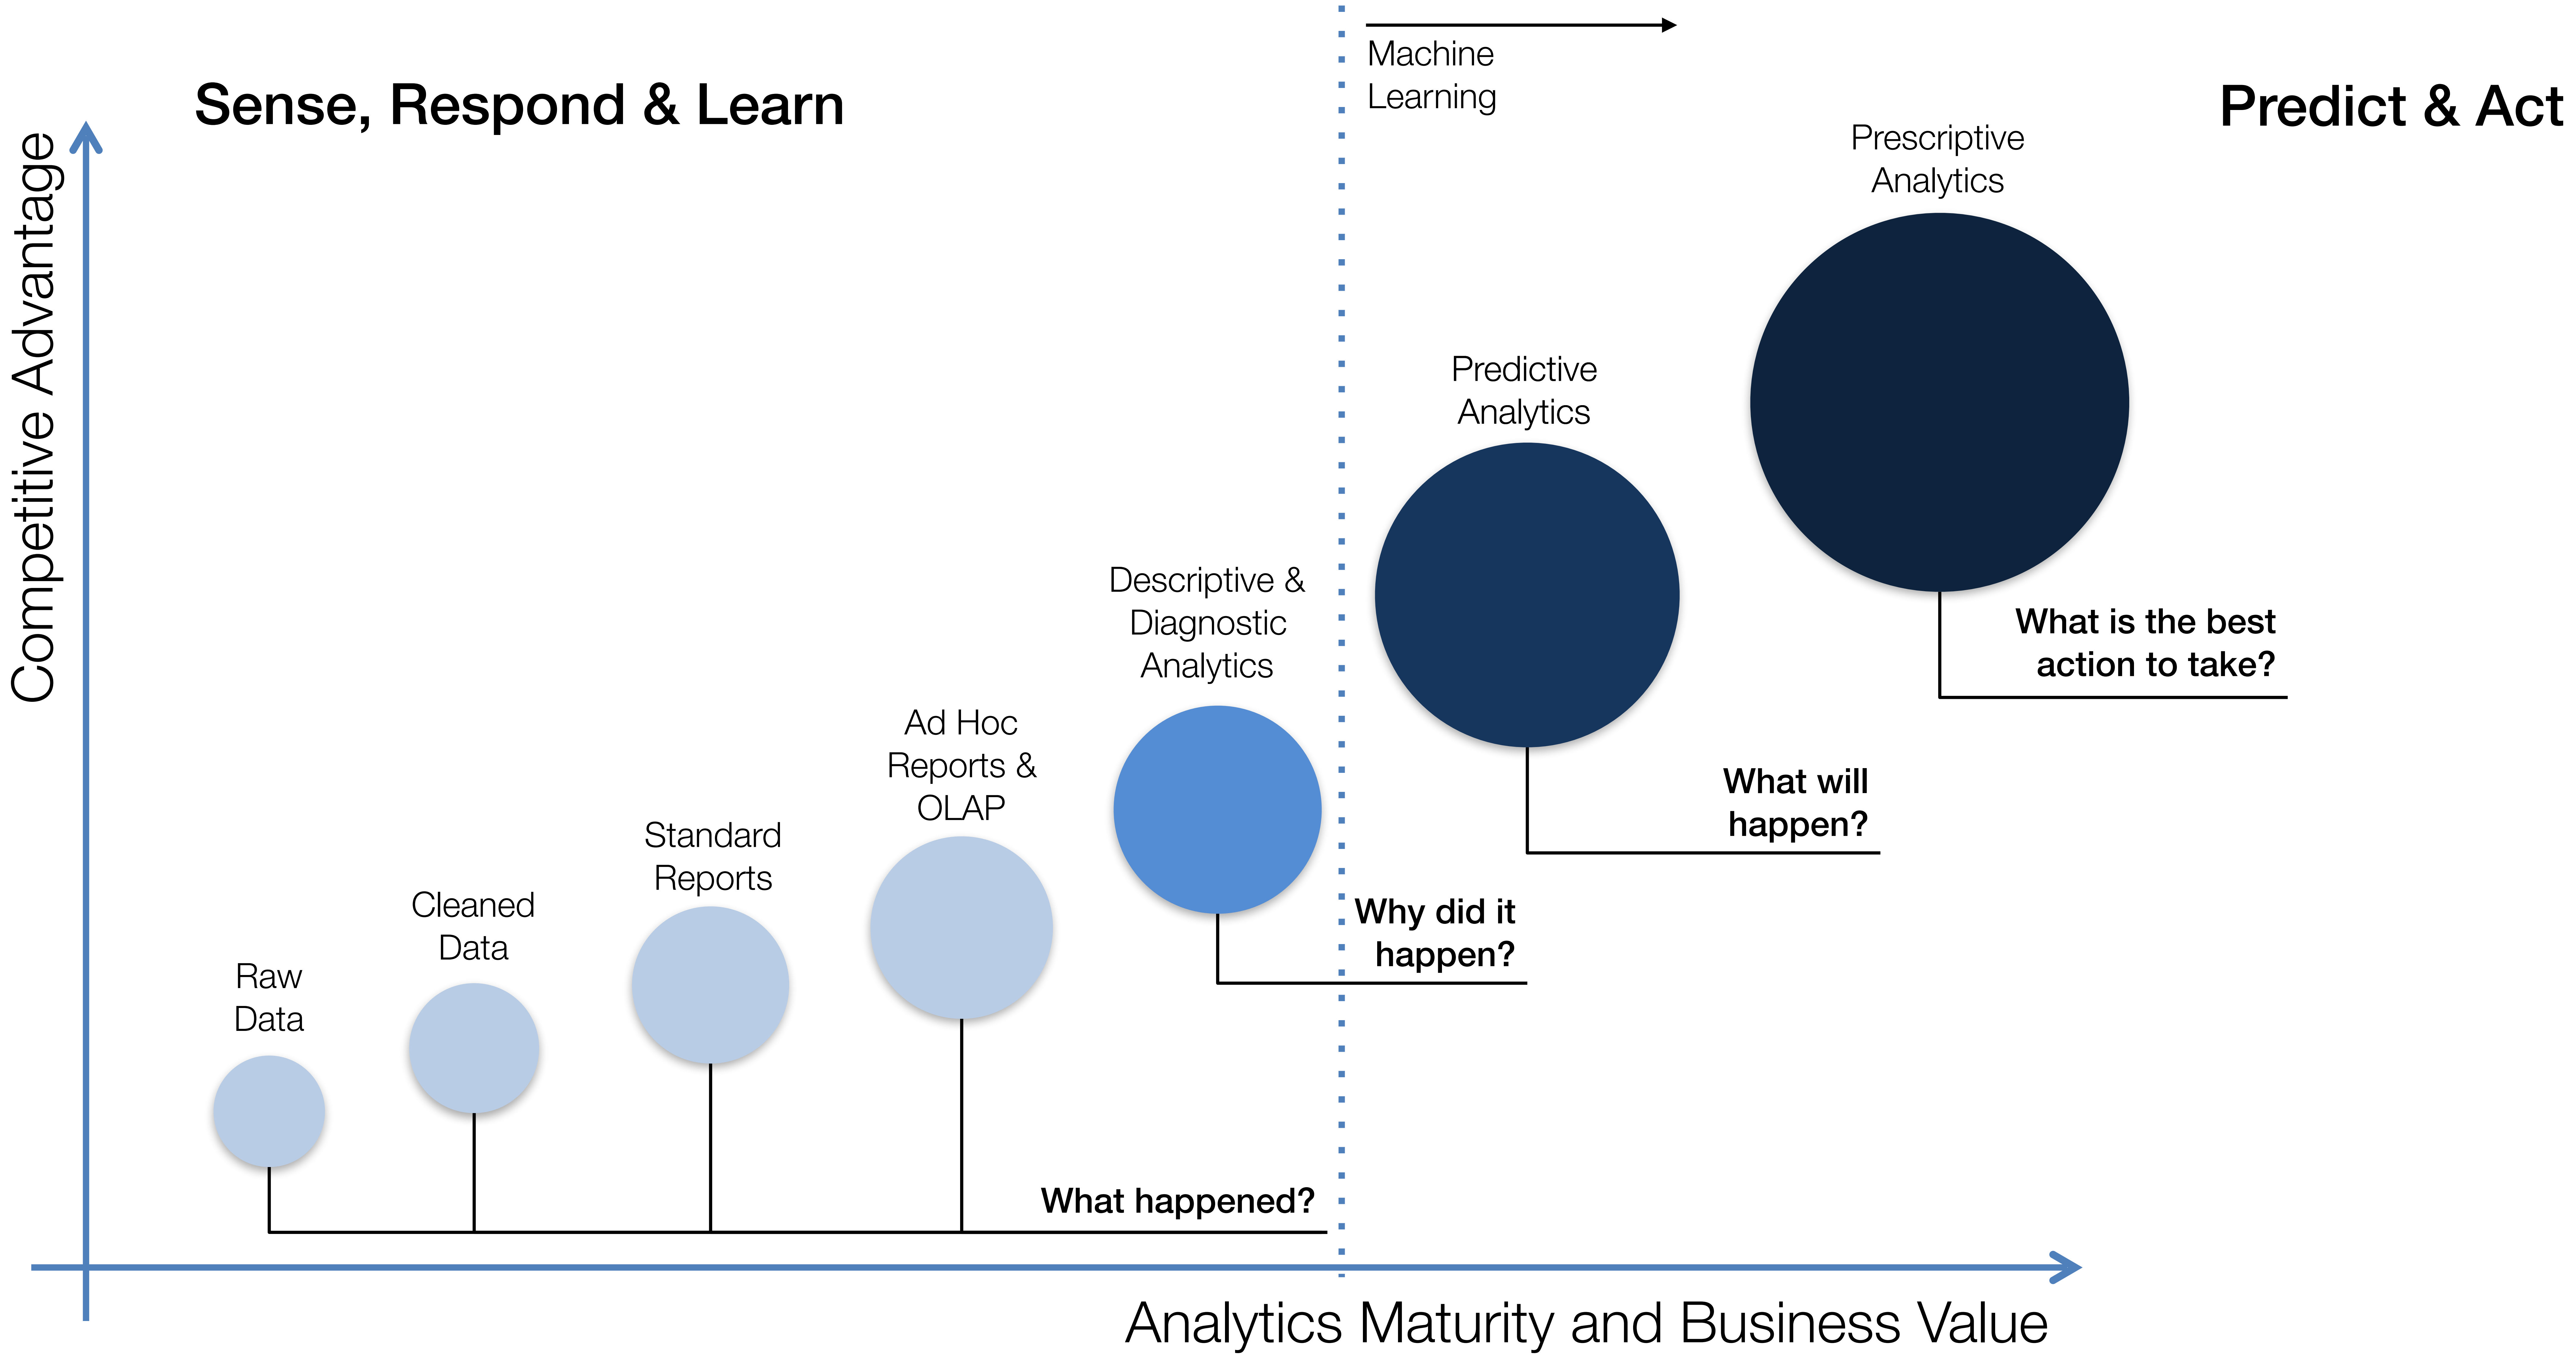
Sense, Respond & Learn
Predict & Act
Machine Learning
Competitive Advantage
Raw Data
Cleaned Data
Standard Reports
Ad Hoc Reports & OLAP
Descriptive & Diagnostic Analytics
Predictive Analytics
Prescriptive Analytics
What happened?
Why did it happen?
What will happen?
What is the best action to take?
Analytics Maturity and Business Value

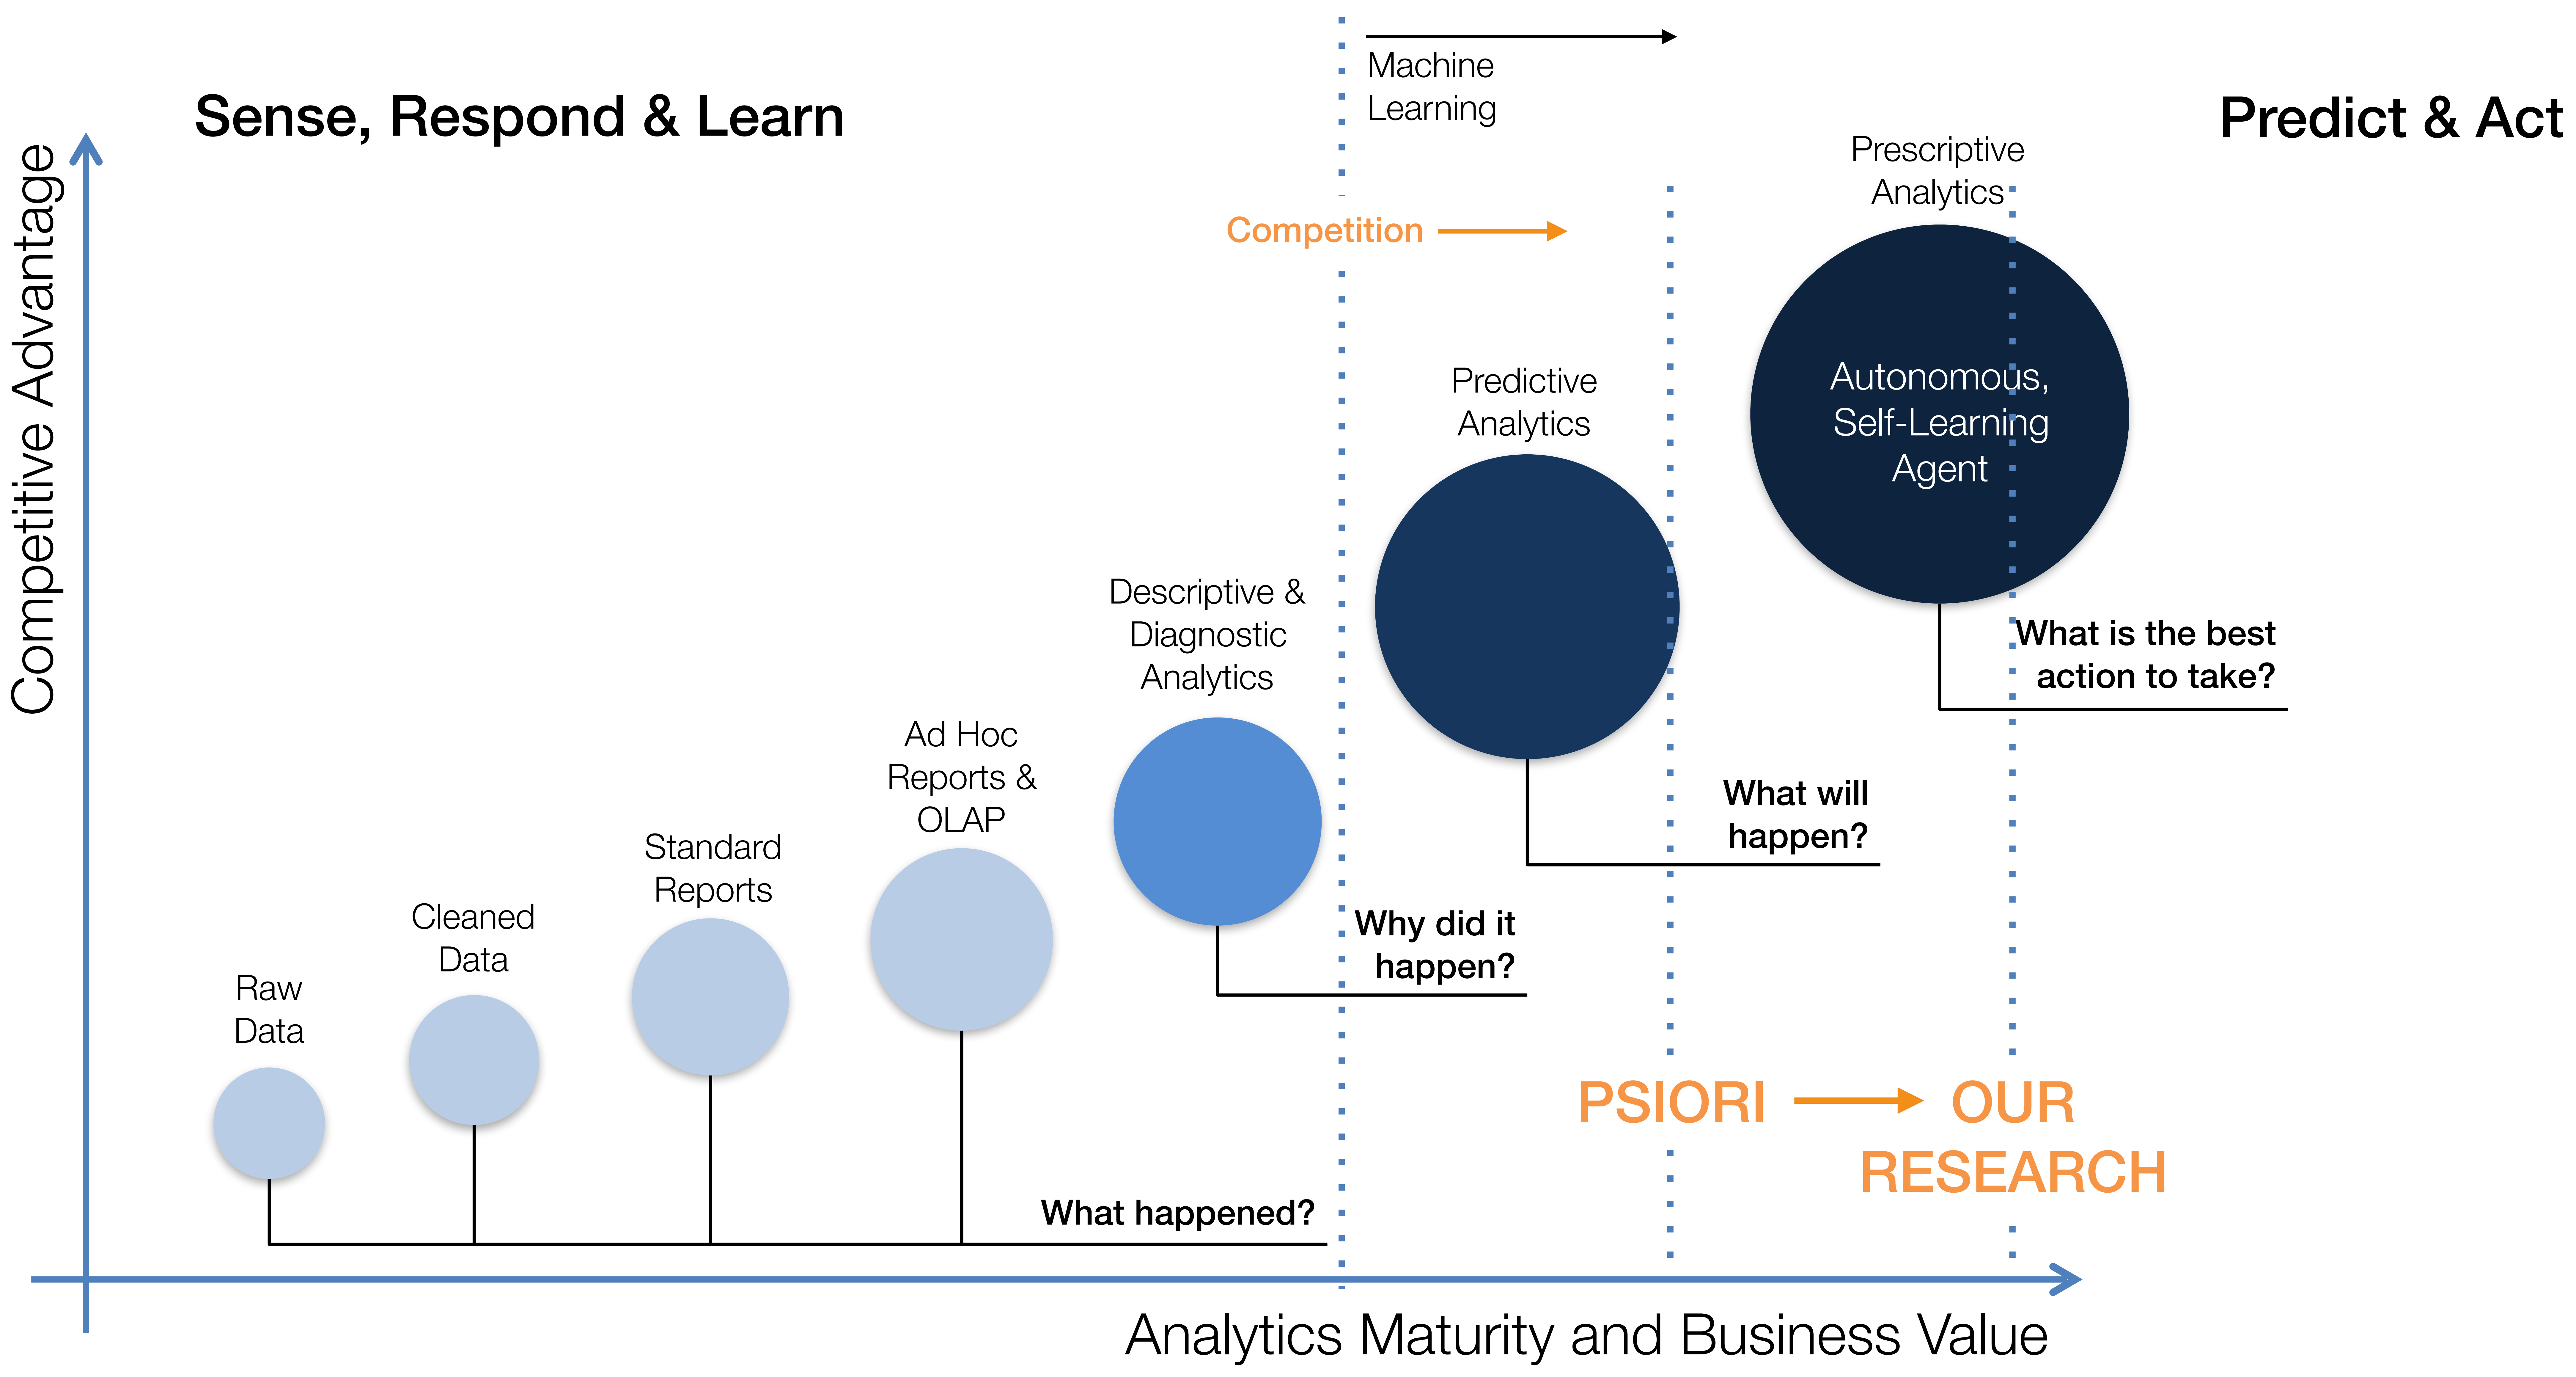
Sense, Respond & Learn
Predict & Act
Competitive Advantage
Machine Learning
Competition
Prescriptive Analytics
Autonomous, Self-Learning Agent
Predictive Analytics
Descriptive & Diagnostic Analytics
Ad Hoc Reports & OLAP
Standard Reports
Cleaned Data
Raw Data
What is the best action to take?
What will happen?
Why did it happen?
What happened?
PSIORI
OUR RESEARCH
Analytics Maturity and Business Value

Let’s talk about Machine Learning!

What is Machine Learning?

> [..] Machine learning explores the study and construction of algorithms that can learn from and make predictions on data[3] [..]
>
> — Wikipedia, March 2017

Artificial Intelligence

Knowledge Representation | … | Planning | Computer Vision | Machine Learning

Different Types of Machine Learning

- Supervised Learning
- Unsupervised Learning
- Reinforcement Learning

Supervised Learning

- External teacher provides "samples" of input and desired output.
  Example: Data from the past, existing users.

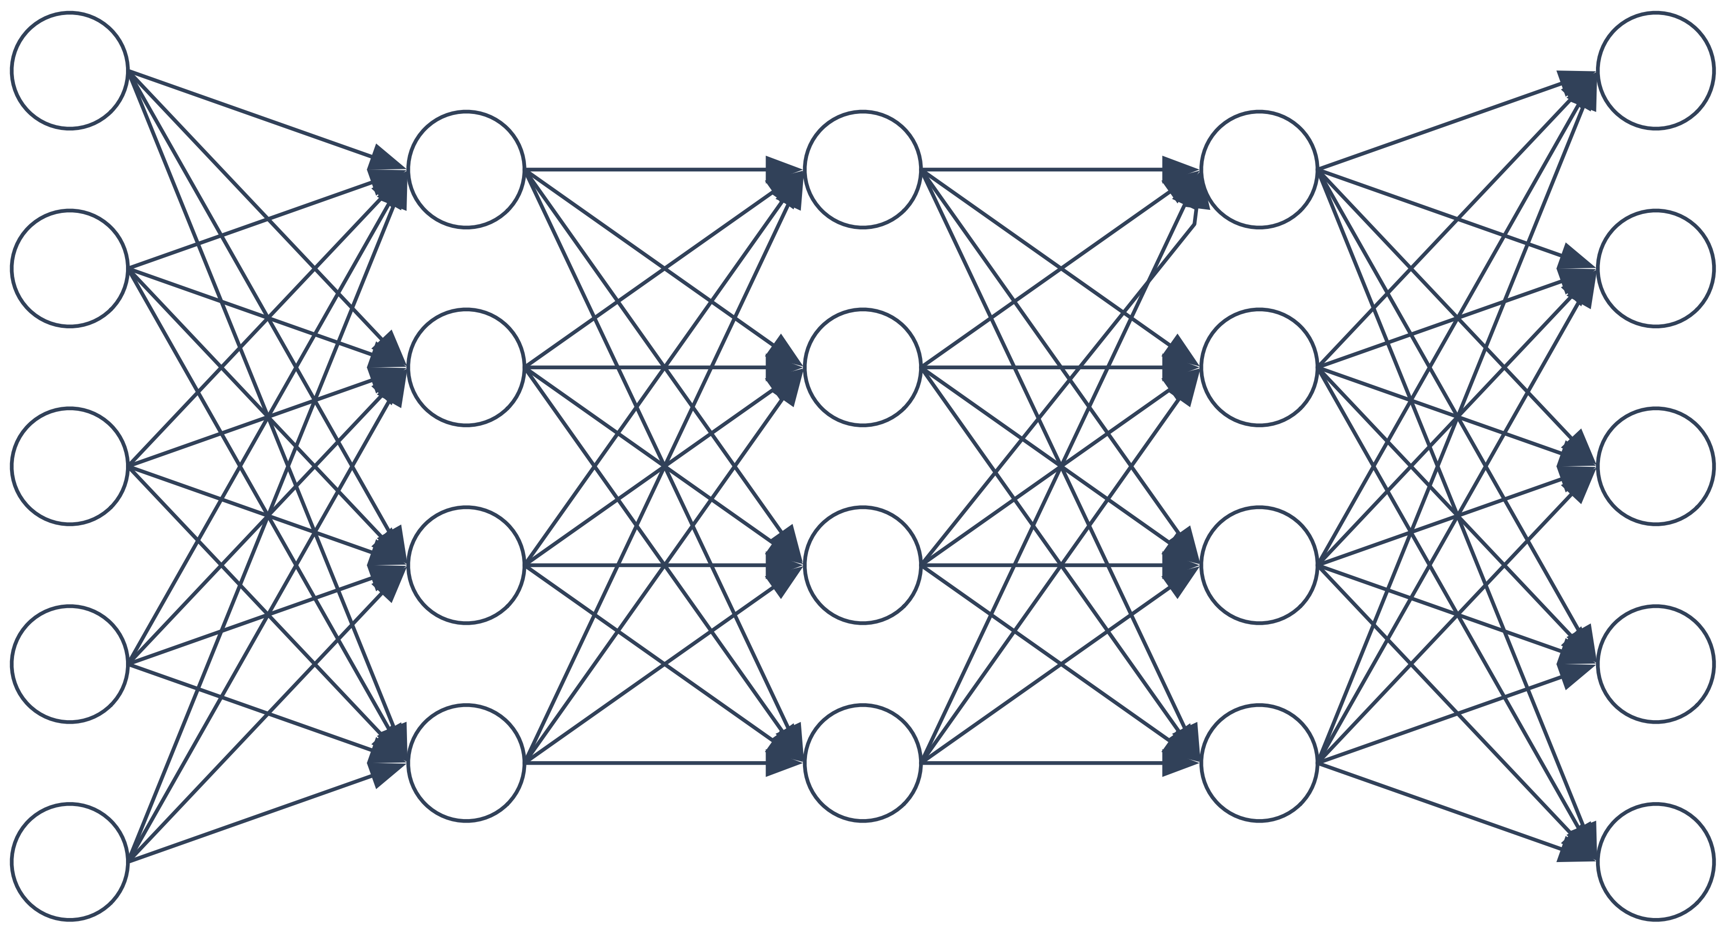

Layers: Input Hidden Output

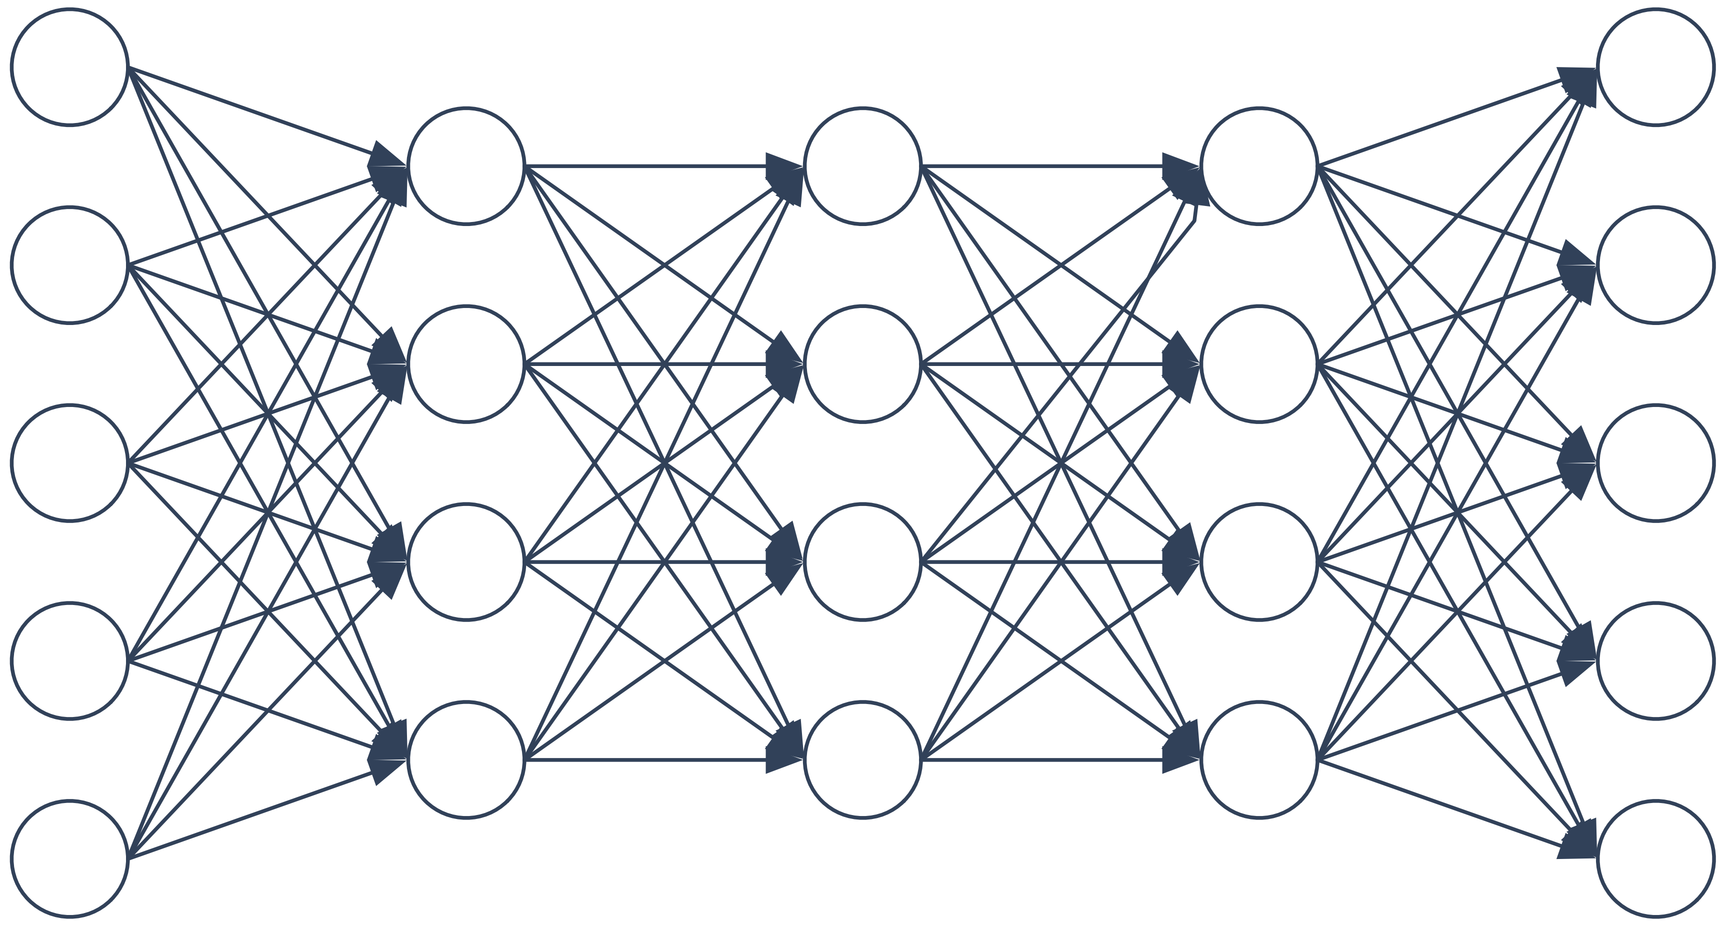

artificial neural network = directed acyclic graph

i = (0.2, 0.3, 0.0, -0.1, 0.9) → **o** = (0.0, 1.0, 0.0, 0.0, 0.0)

(forward pass / propagation)

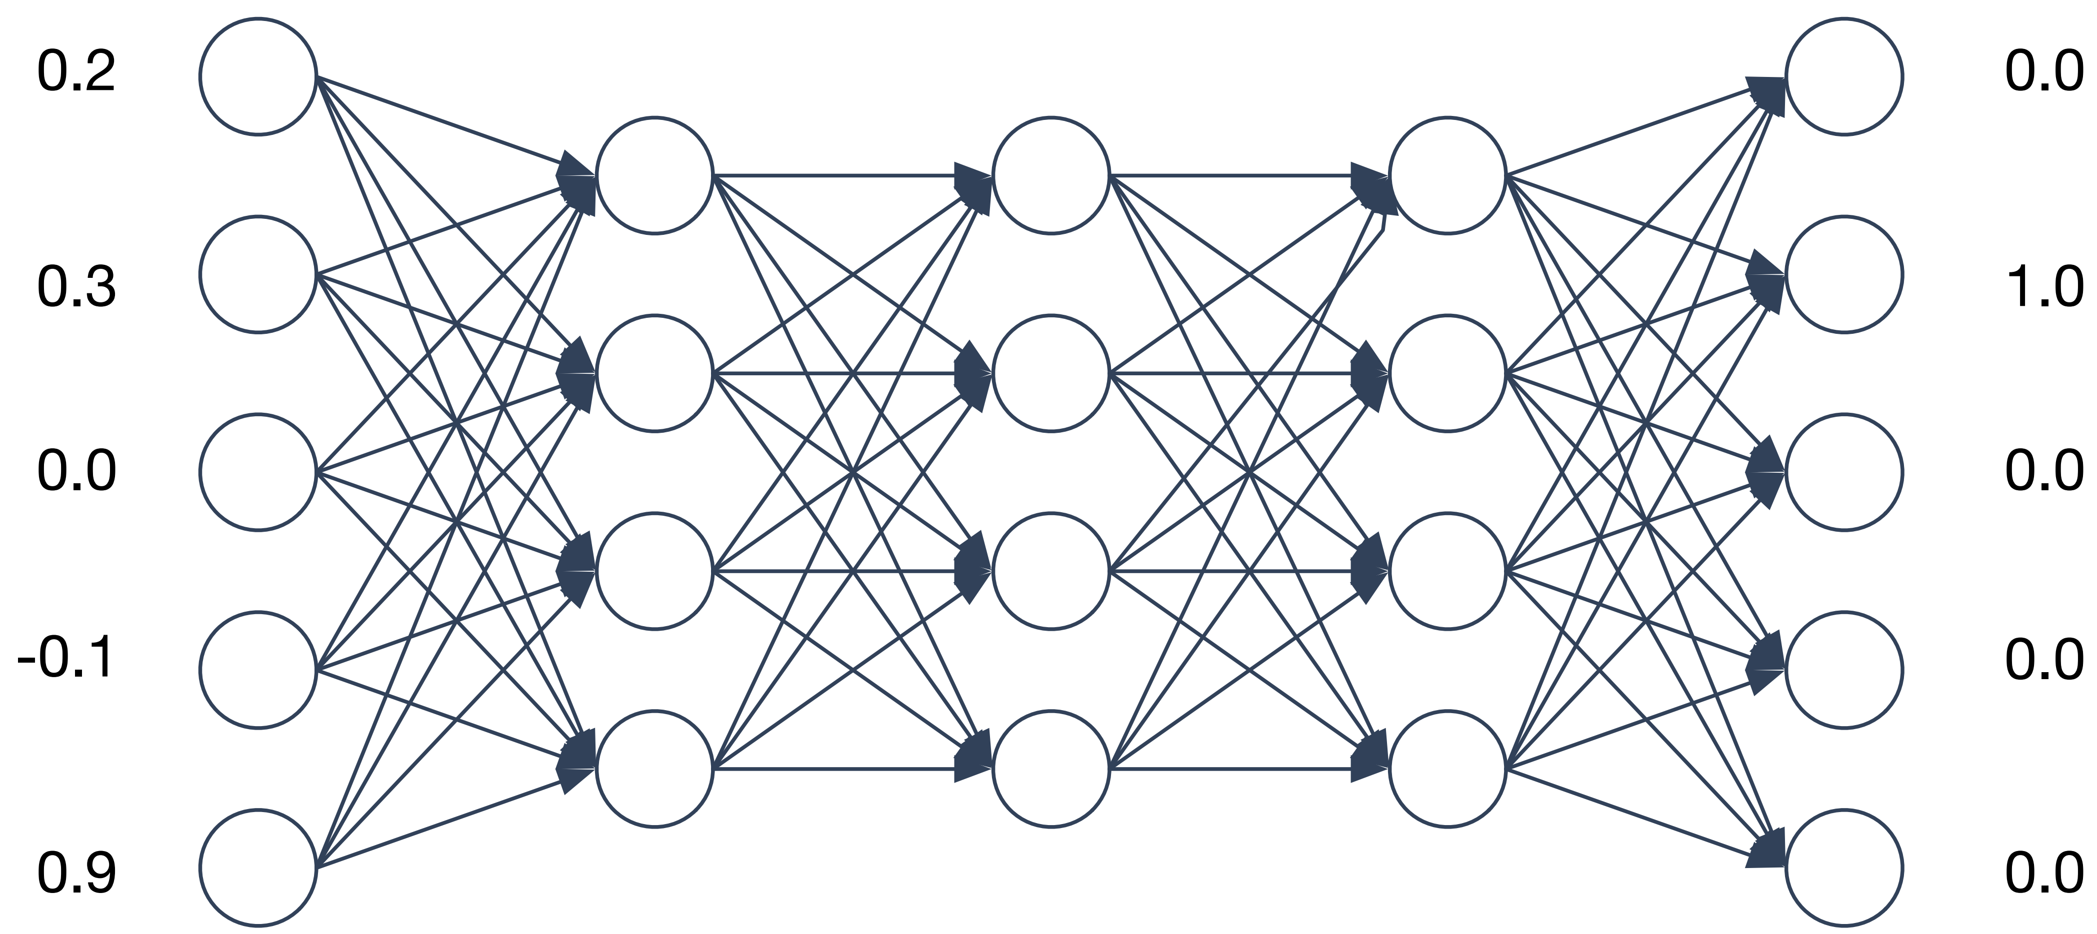

artificial neural network = directed acyclic graph

$$\frac{\partial E}{\partial w_{ji}}$$

(backpropagation algorithm)

o = (0.0, 1.0, 0.0, 0.0, 0.0)

t = (0.0, 0.0, 1.0, 0.0, 0.0)

$$E = \frac{1}{2}\sum_{i=1}^{|O|}(y_i^p - o_i^p)^2$$

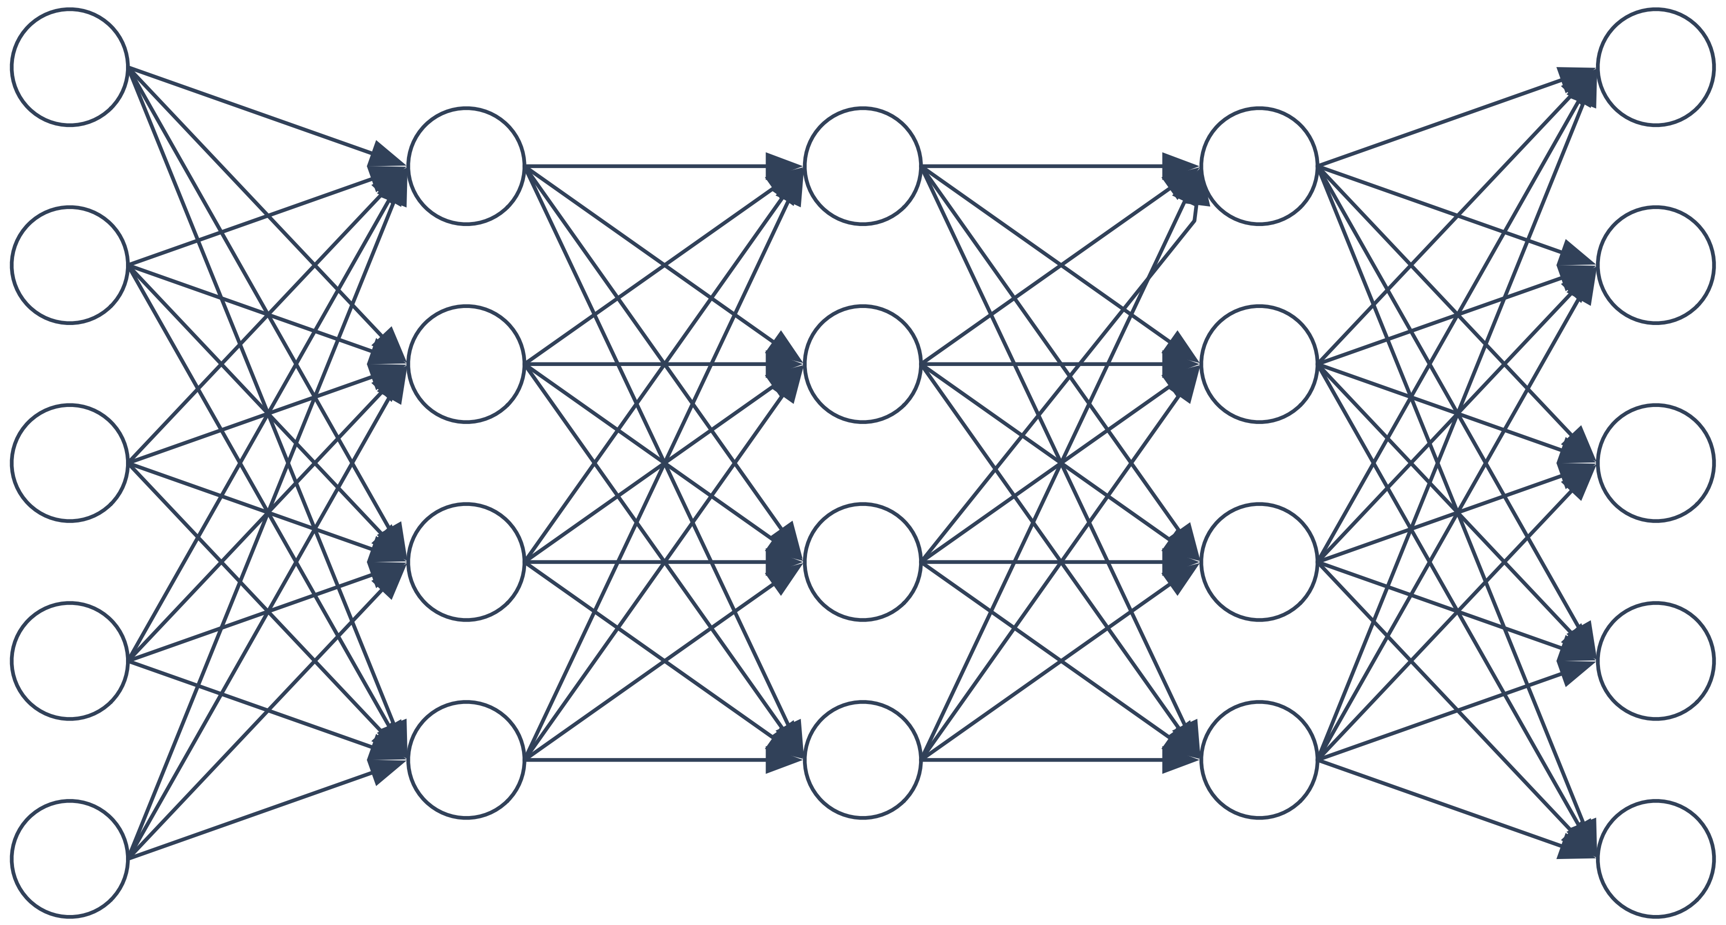

0.0

1.0

0.0

0.0

0.0

$$\frac{\partial E}{\partial w_{ji}}$$

(backpropagation algorithm)

o = (0.0, 1.0, 0.0, 0.0, 0.0)

t = (0.0, 0.0, 1.0, 0.0, 0.0)

$$E = \frac{1}{2} \sum_{i=1}^{|O|} (y_i^p - o_i^p)^2$$

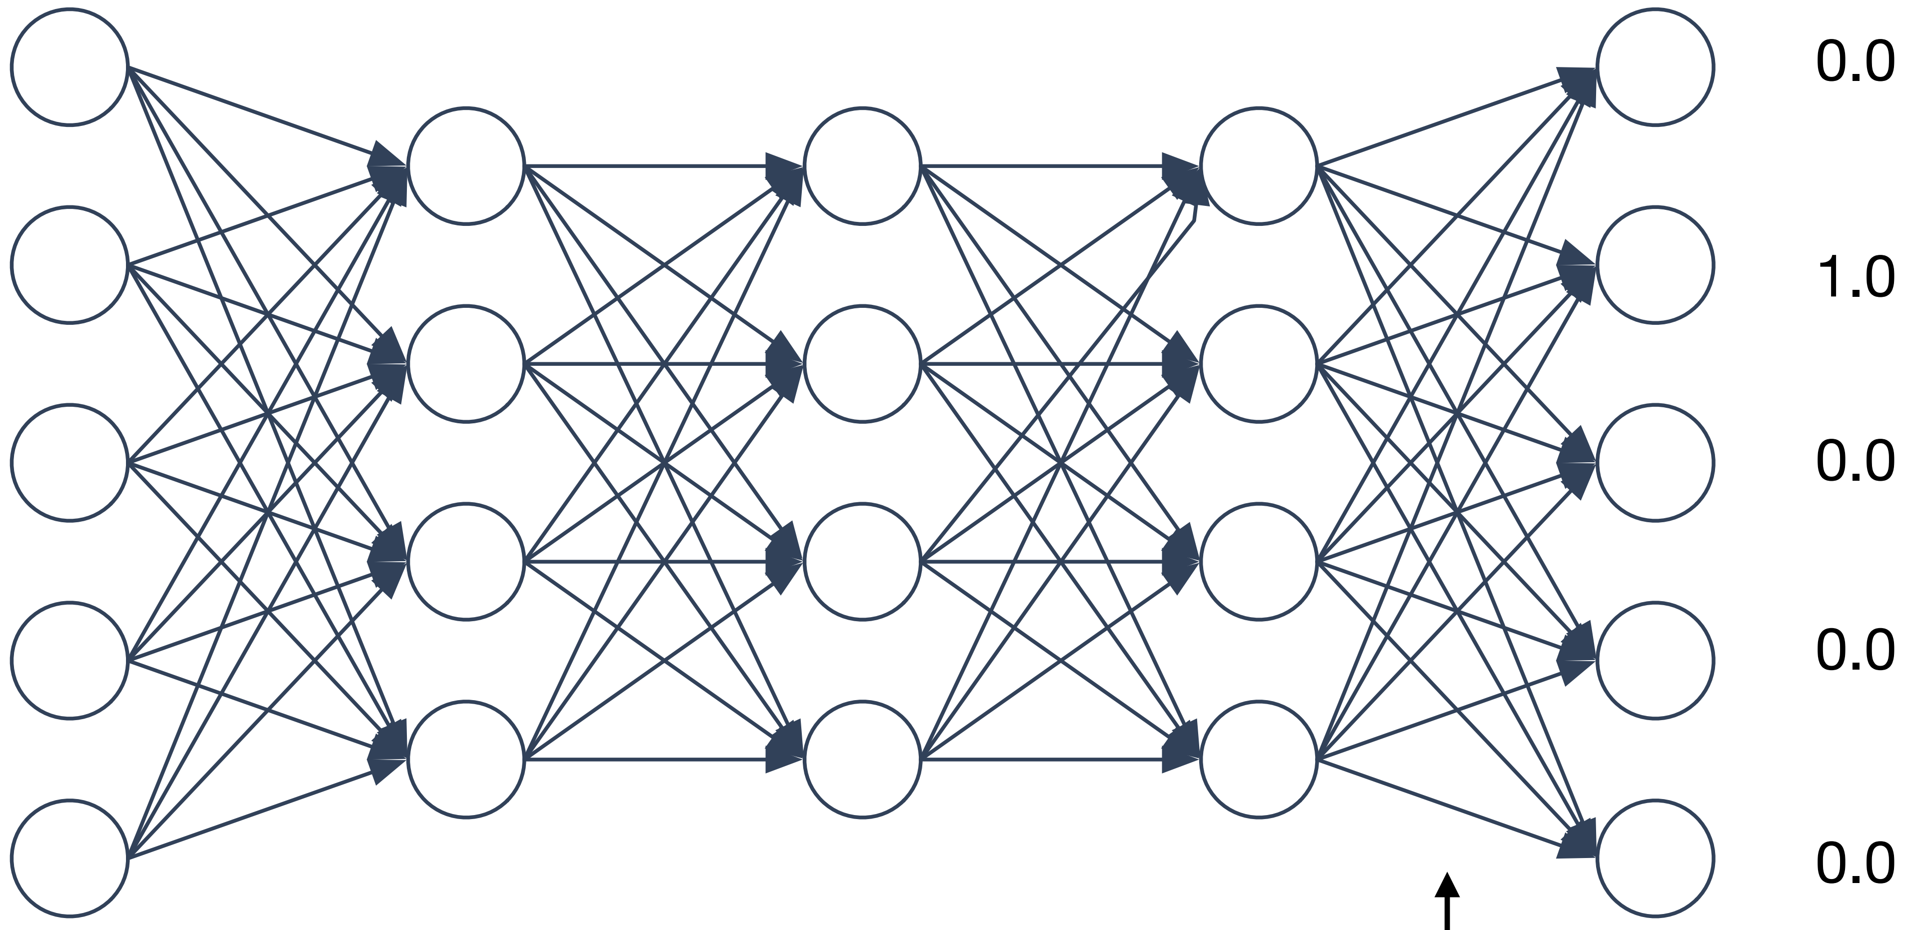

$$\frac{\partial E}{\partial w_{ji}}$$

$$\frac{\partial E}{\partial w_{ji}}$$

(backpropagation algorithm)

o = (0.0, 1.0, 0.0, 0.0, 0.0)

t = (0.0, 0.0, 1.0, 0.0, 0.0)

$$E = \frac{1}{2} \sum_{i=1}^{|O|} (y_i^p - o_i^p)^2$$

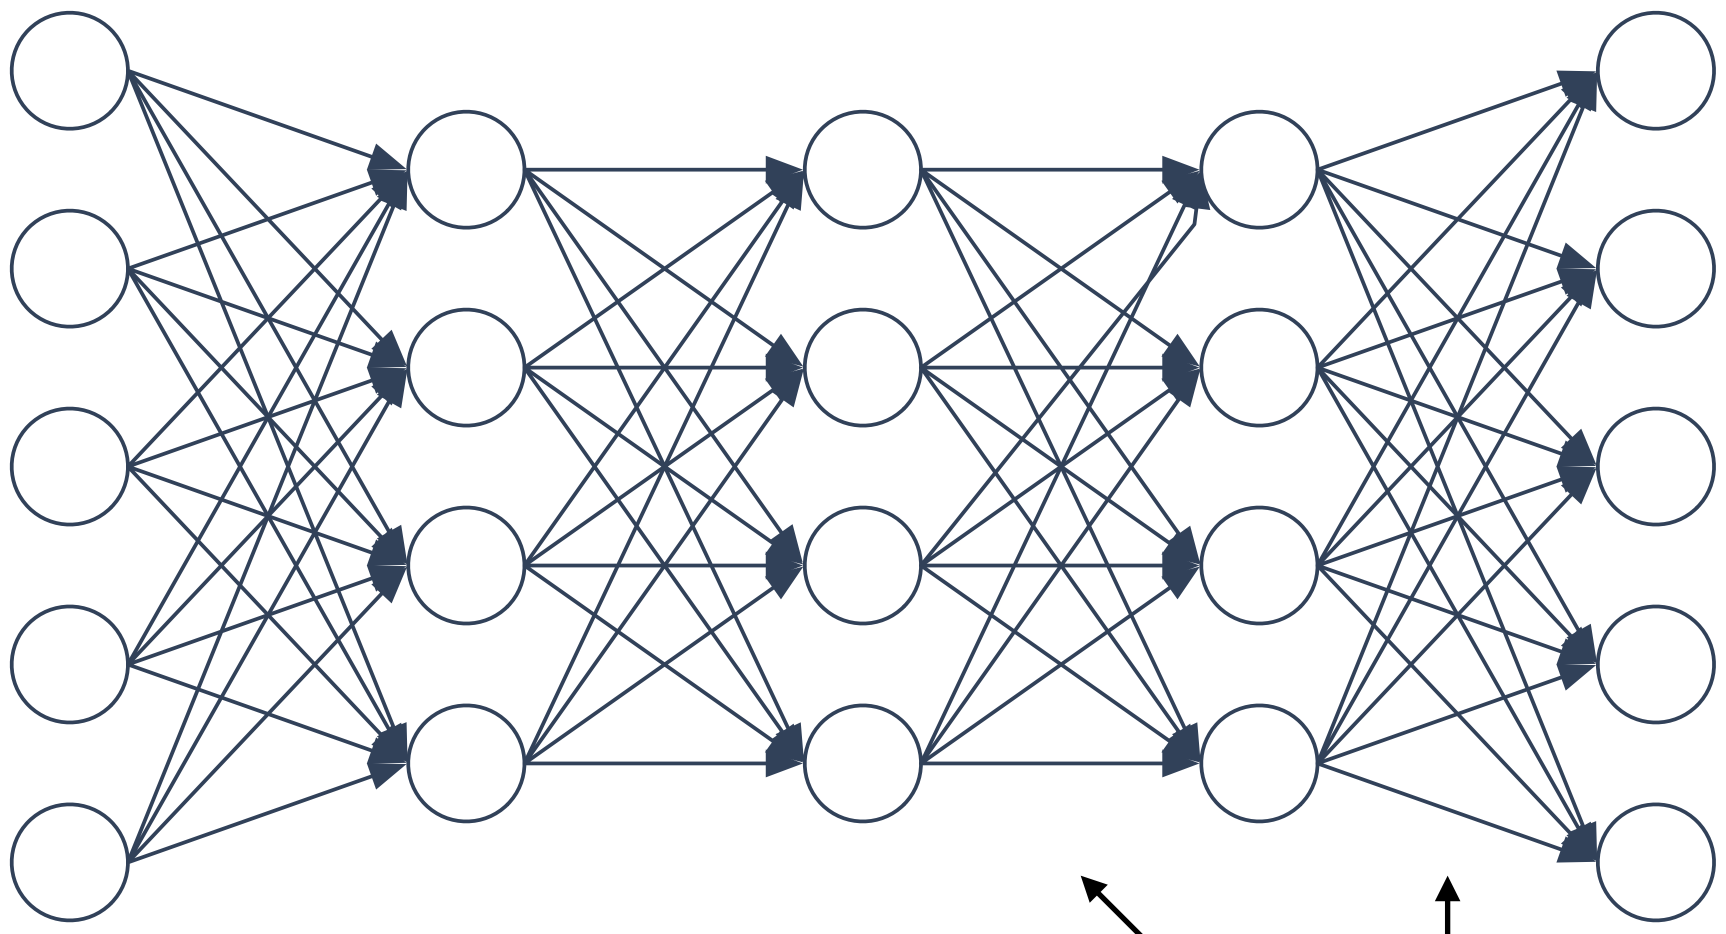

0.0

1.0

0.0

0.0

0.0

$$\frac{\partial E}{\partial w_{ji}}$$

$$\frac{\partial E}{\partial w_{ji}}$$

(backpropagation algorithm)

o = (0.0, 1.0, 0.0, 0.0, 0.0)

t = (0.0, 0.0, 1.0, 0.0, 0.0)

$$E = \frac{1}{2} \sum_{i=1}^{|O|} (y_i^p - o_i^p)^2$$

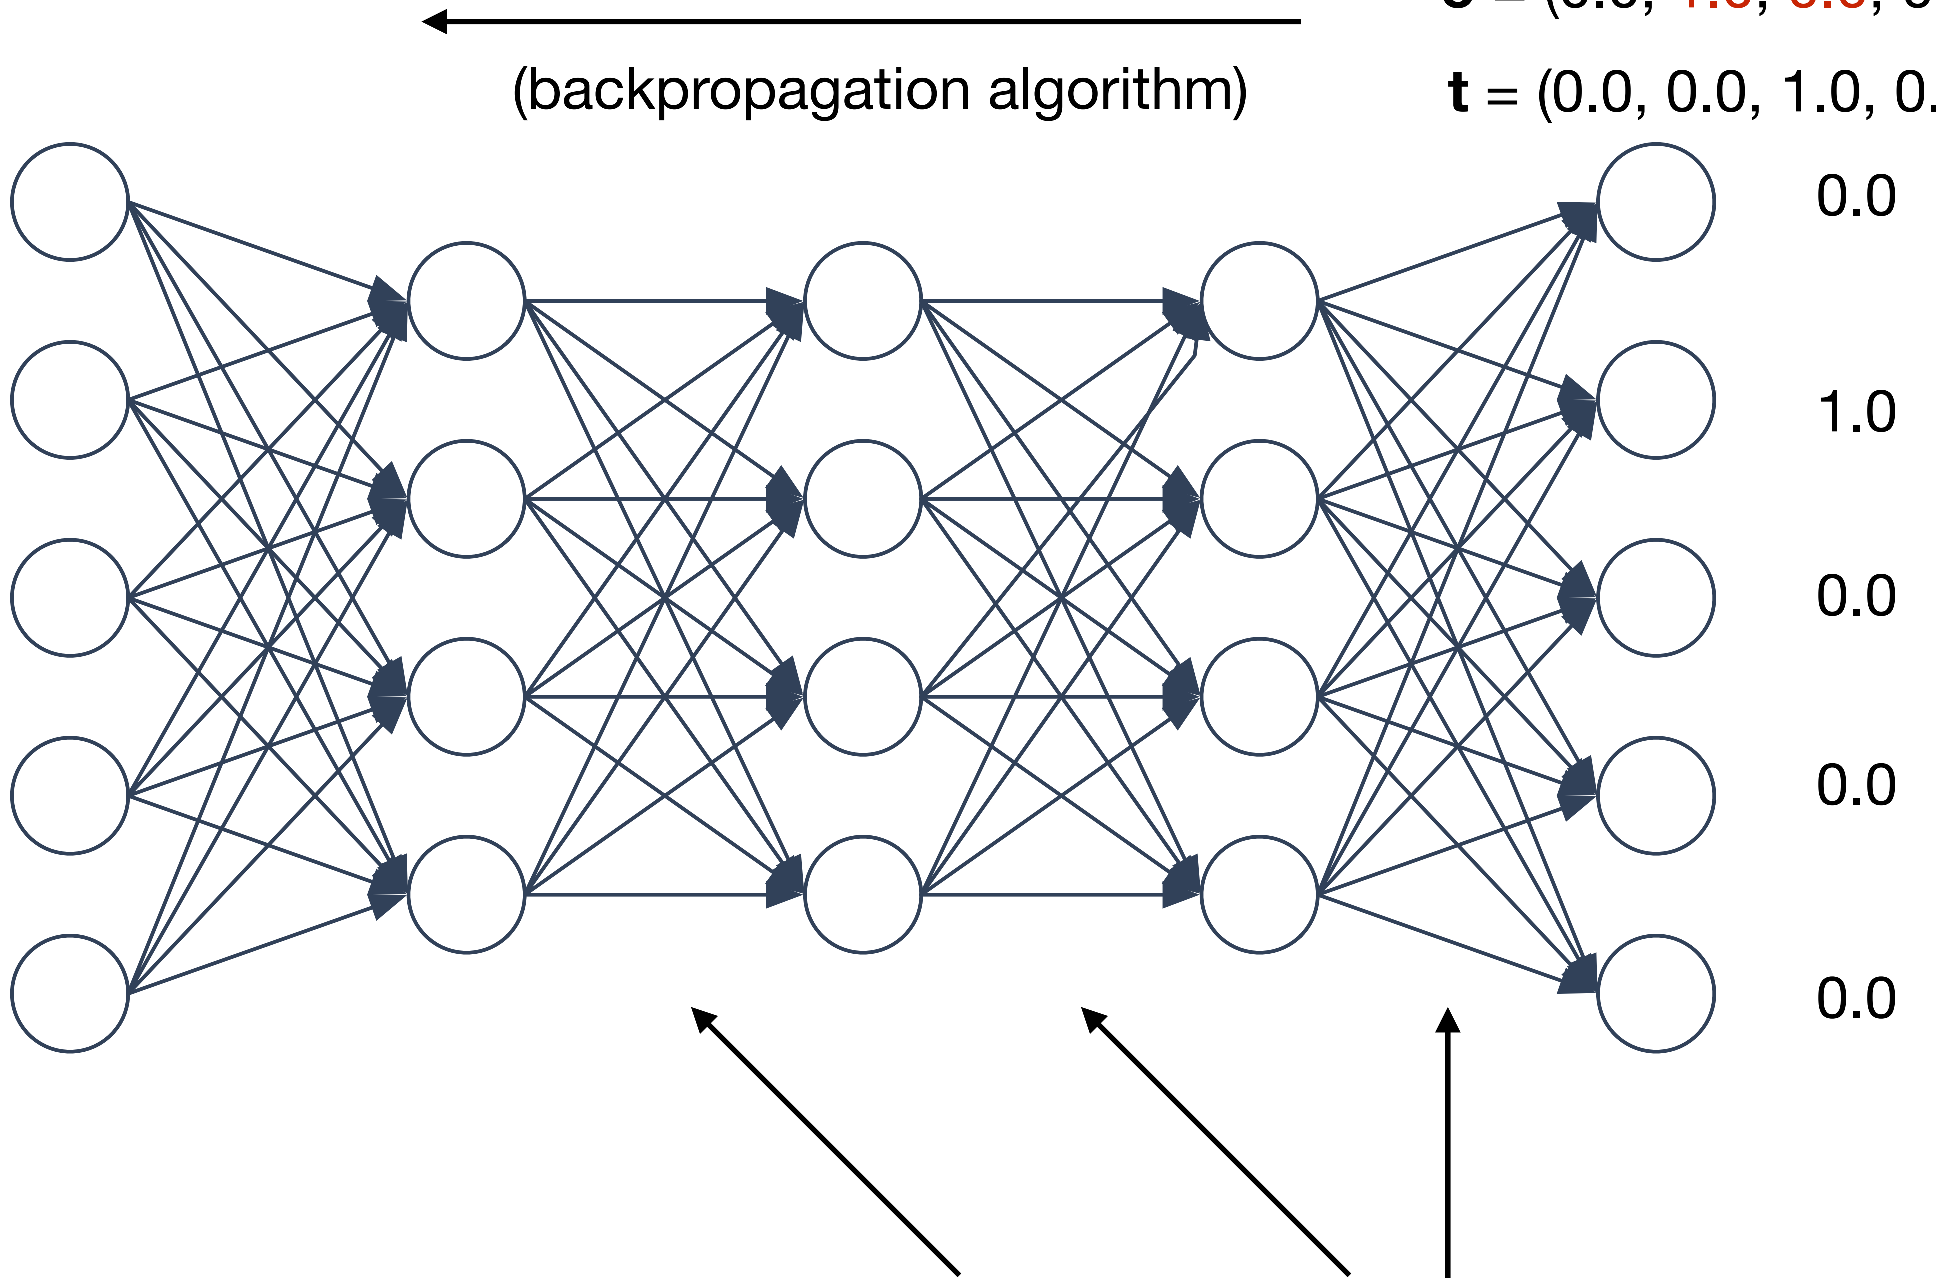

$$\frac{\partial E}{\partial w_{ji}}$$

$$\frac{\partial E}{\partial w_{ji}}$$

(backpropagation algorithm)

o = (0.0, 1.0, 0.0, 0.0, 0.0)

t = (0.0, 0.0, 1.0, 0.0, 0.0)

$$E = \frac{1}{2} \sum_{i=1}^{|O|} (y_i^p - o_i^p)^2$$

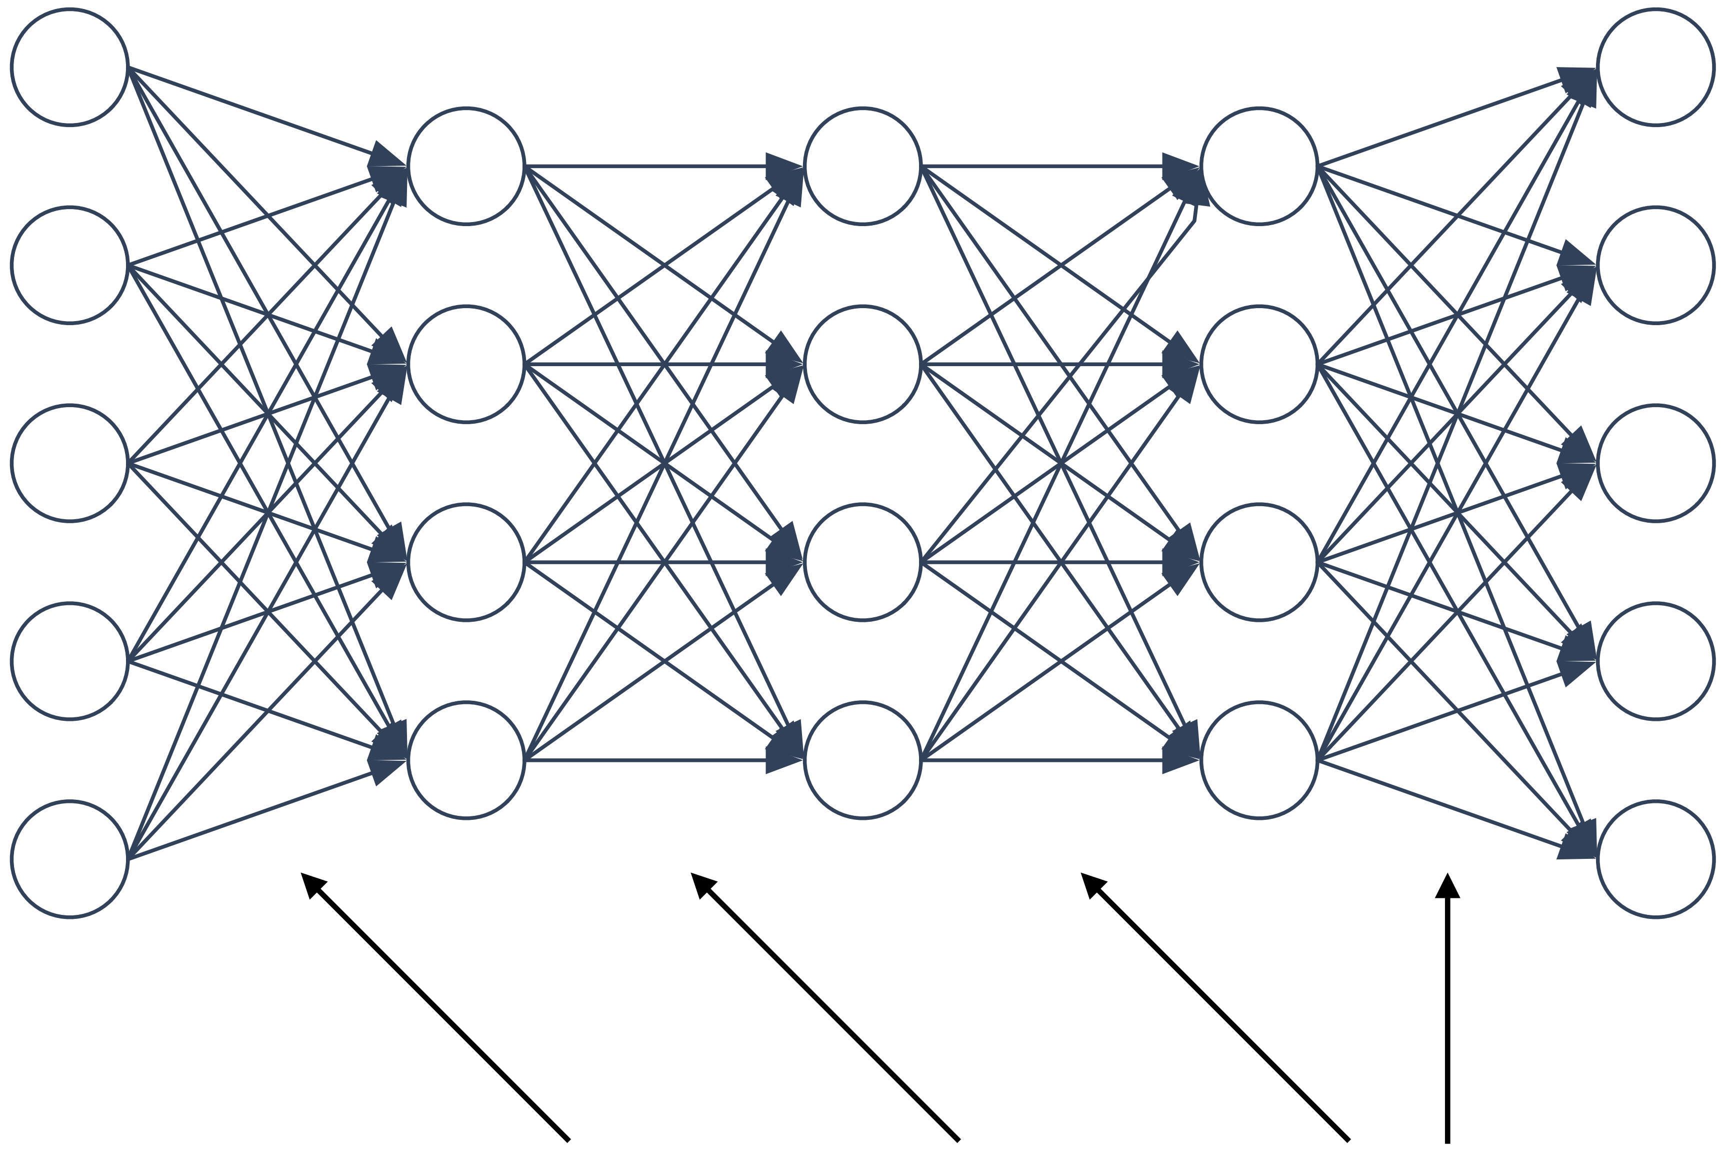

0.0

1.0

0.0

0.0

0.0

$$\frac{\partial E}{\partial w_{ji}}$$

$$\frac{\partial E}{\partial w_{ji}}$$

(backpropagation algorithm)

o = (0.0, 0.1, 0.9, 0.0, 0.0)

t = (0.0, 0.0, 1.0, 0.0, 0.0)

$$E = \frac{1}{2}\sum_{i=1}^{|O|}(y_i^p - o_i^p)^2$$

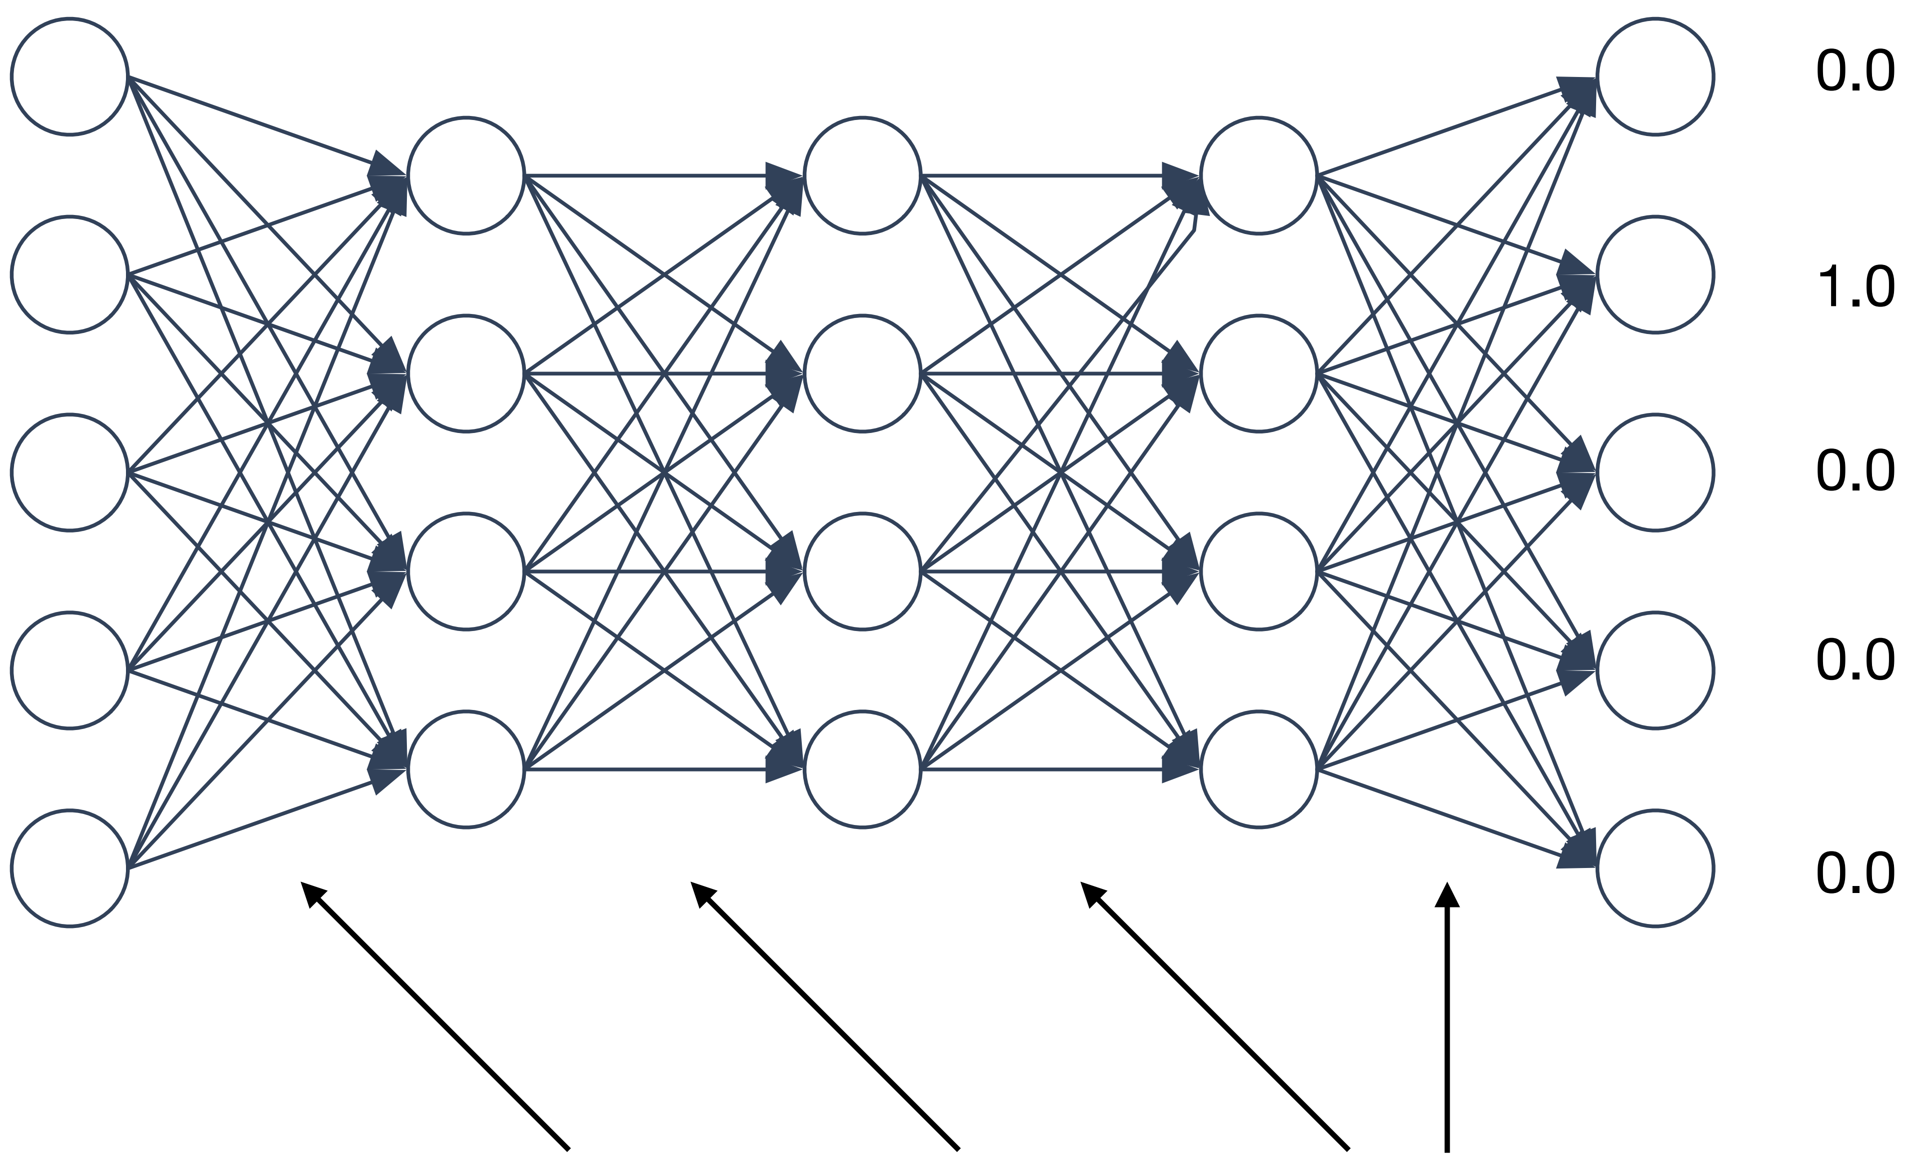

$$\frac{\partial E}{\partial w_{ji}}$$

$$\frac{\partial E}{\partial w_{ji}}$$

(backpropagation algorithm)

o = (0.0, 0.1, 0.9, 0.0, 0.0)

t = (0.0, 0.0, 1.0, 0.0, 0.0)

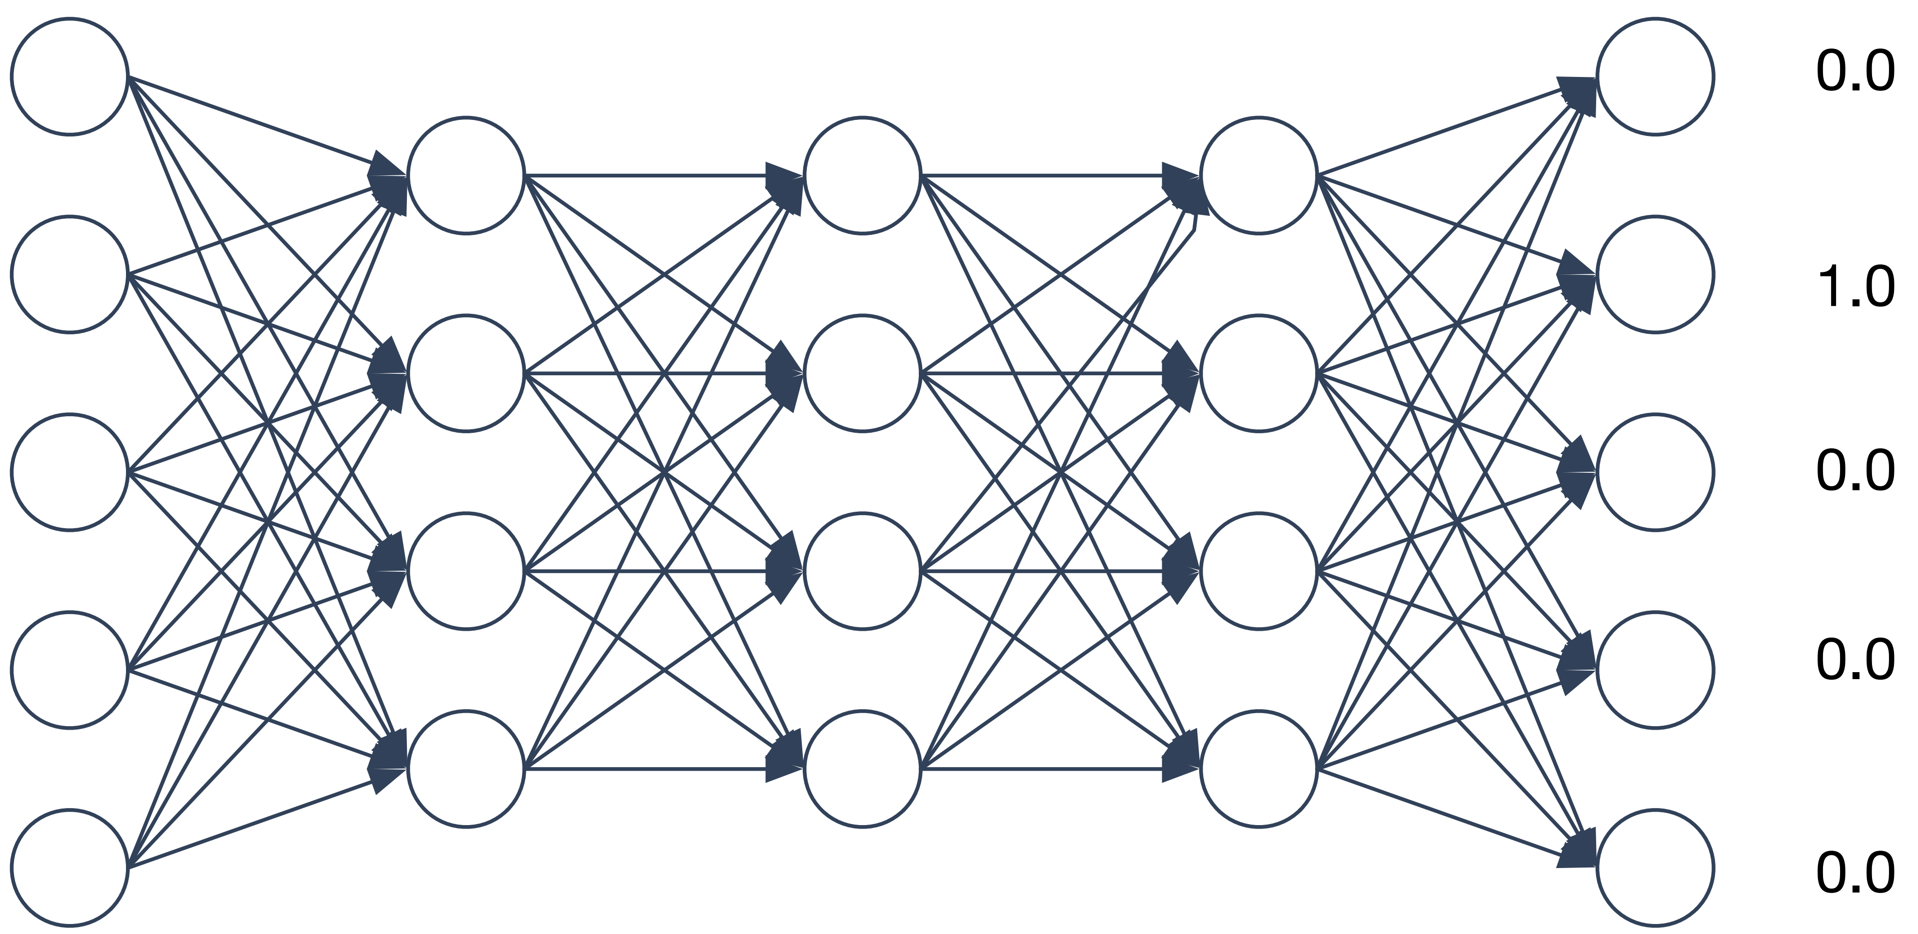

"Shallow" nets: fine!

$$\frac{\partial E}{\partial w_{ji}}$$

(backpropagation algorithm)

o = (0.0, 0.1, 0.9, 0.0, 0.0)

t = (0.0, 0.0, 1.0, 0.0, 0.0)

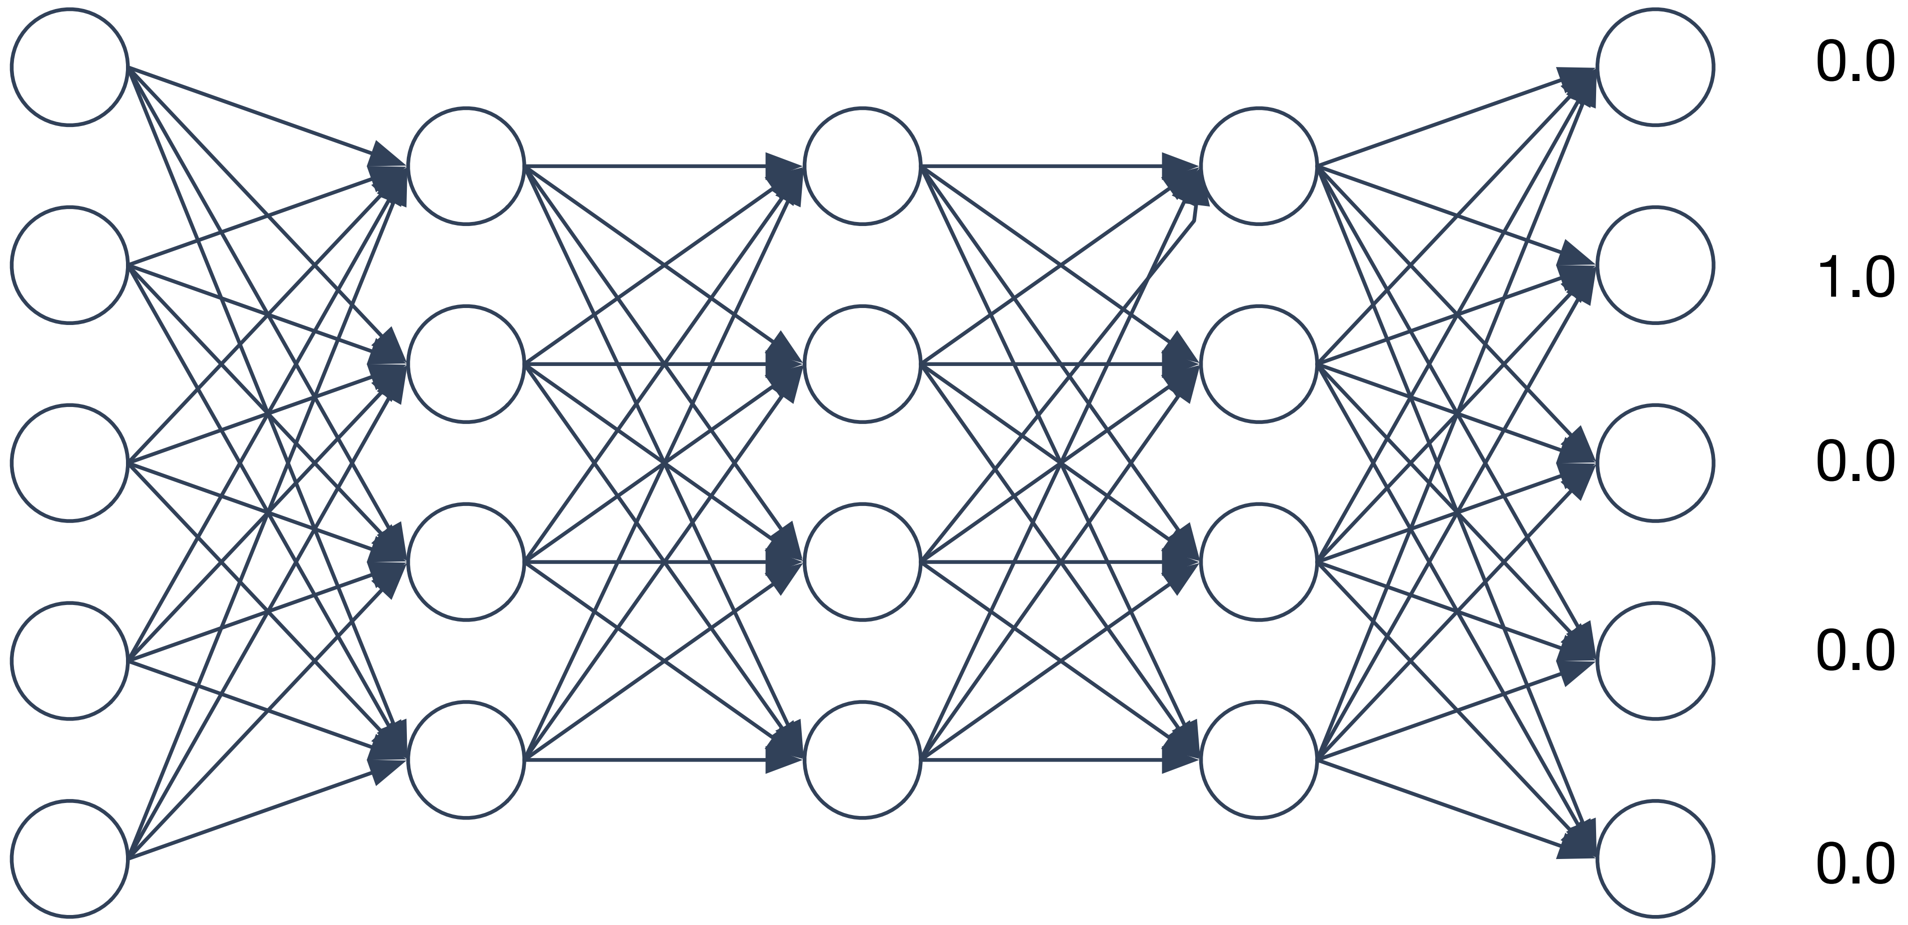

"Shallow" nets: fine!

"Shallow" nets = 1 hidden layer

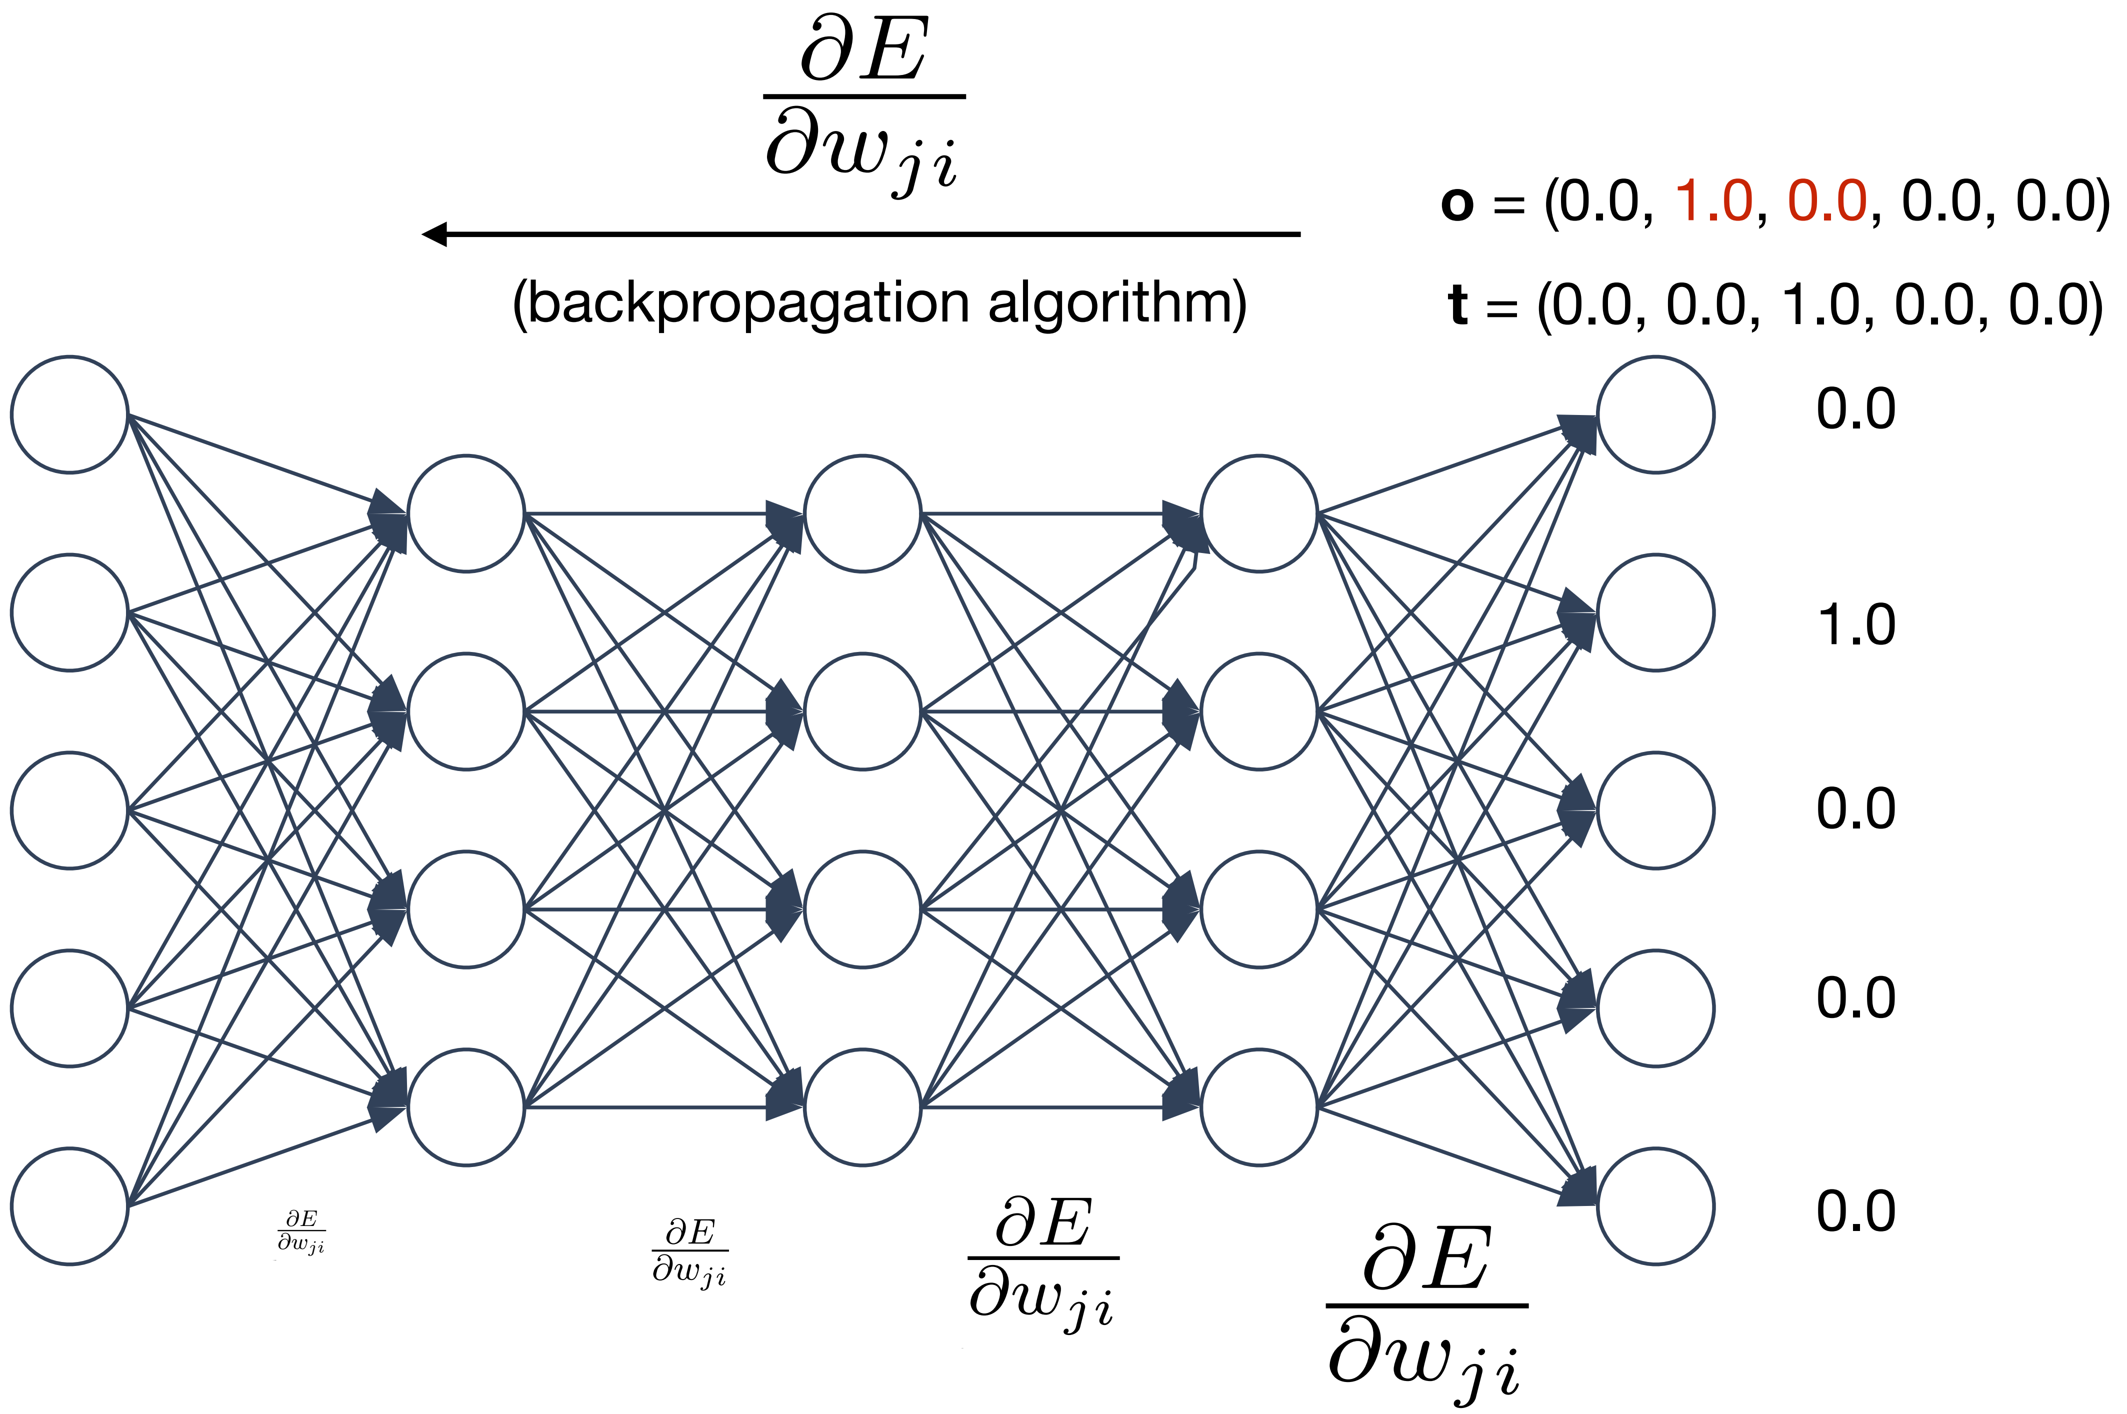
$\frac{\partial E}{\partial w_{ji}}$
(backpropagation algorithm)
o = (0.0, 1.0, 0.0, 0.0, 0.0)
t = (0.0, 0.0, 1.0, 0.0, 0.0)
0.0
1.0
0.0
0.0
0.0
$\frac{\partial E}{\partial w_{ji}}$
$\frac{\partial E}{\partial w_{ji}}$
$\frac{\partial E}{\partial w_{ji}}$
$\frac{\partial E}{\partial w_{ji}}$

"deep" nets: oh....

Vanishing Gradient

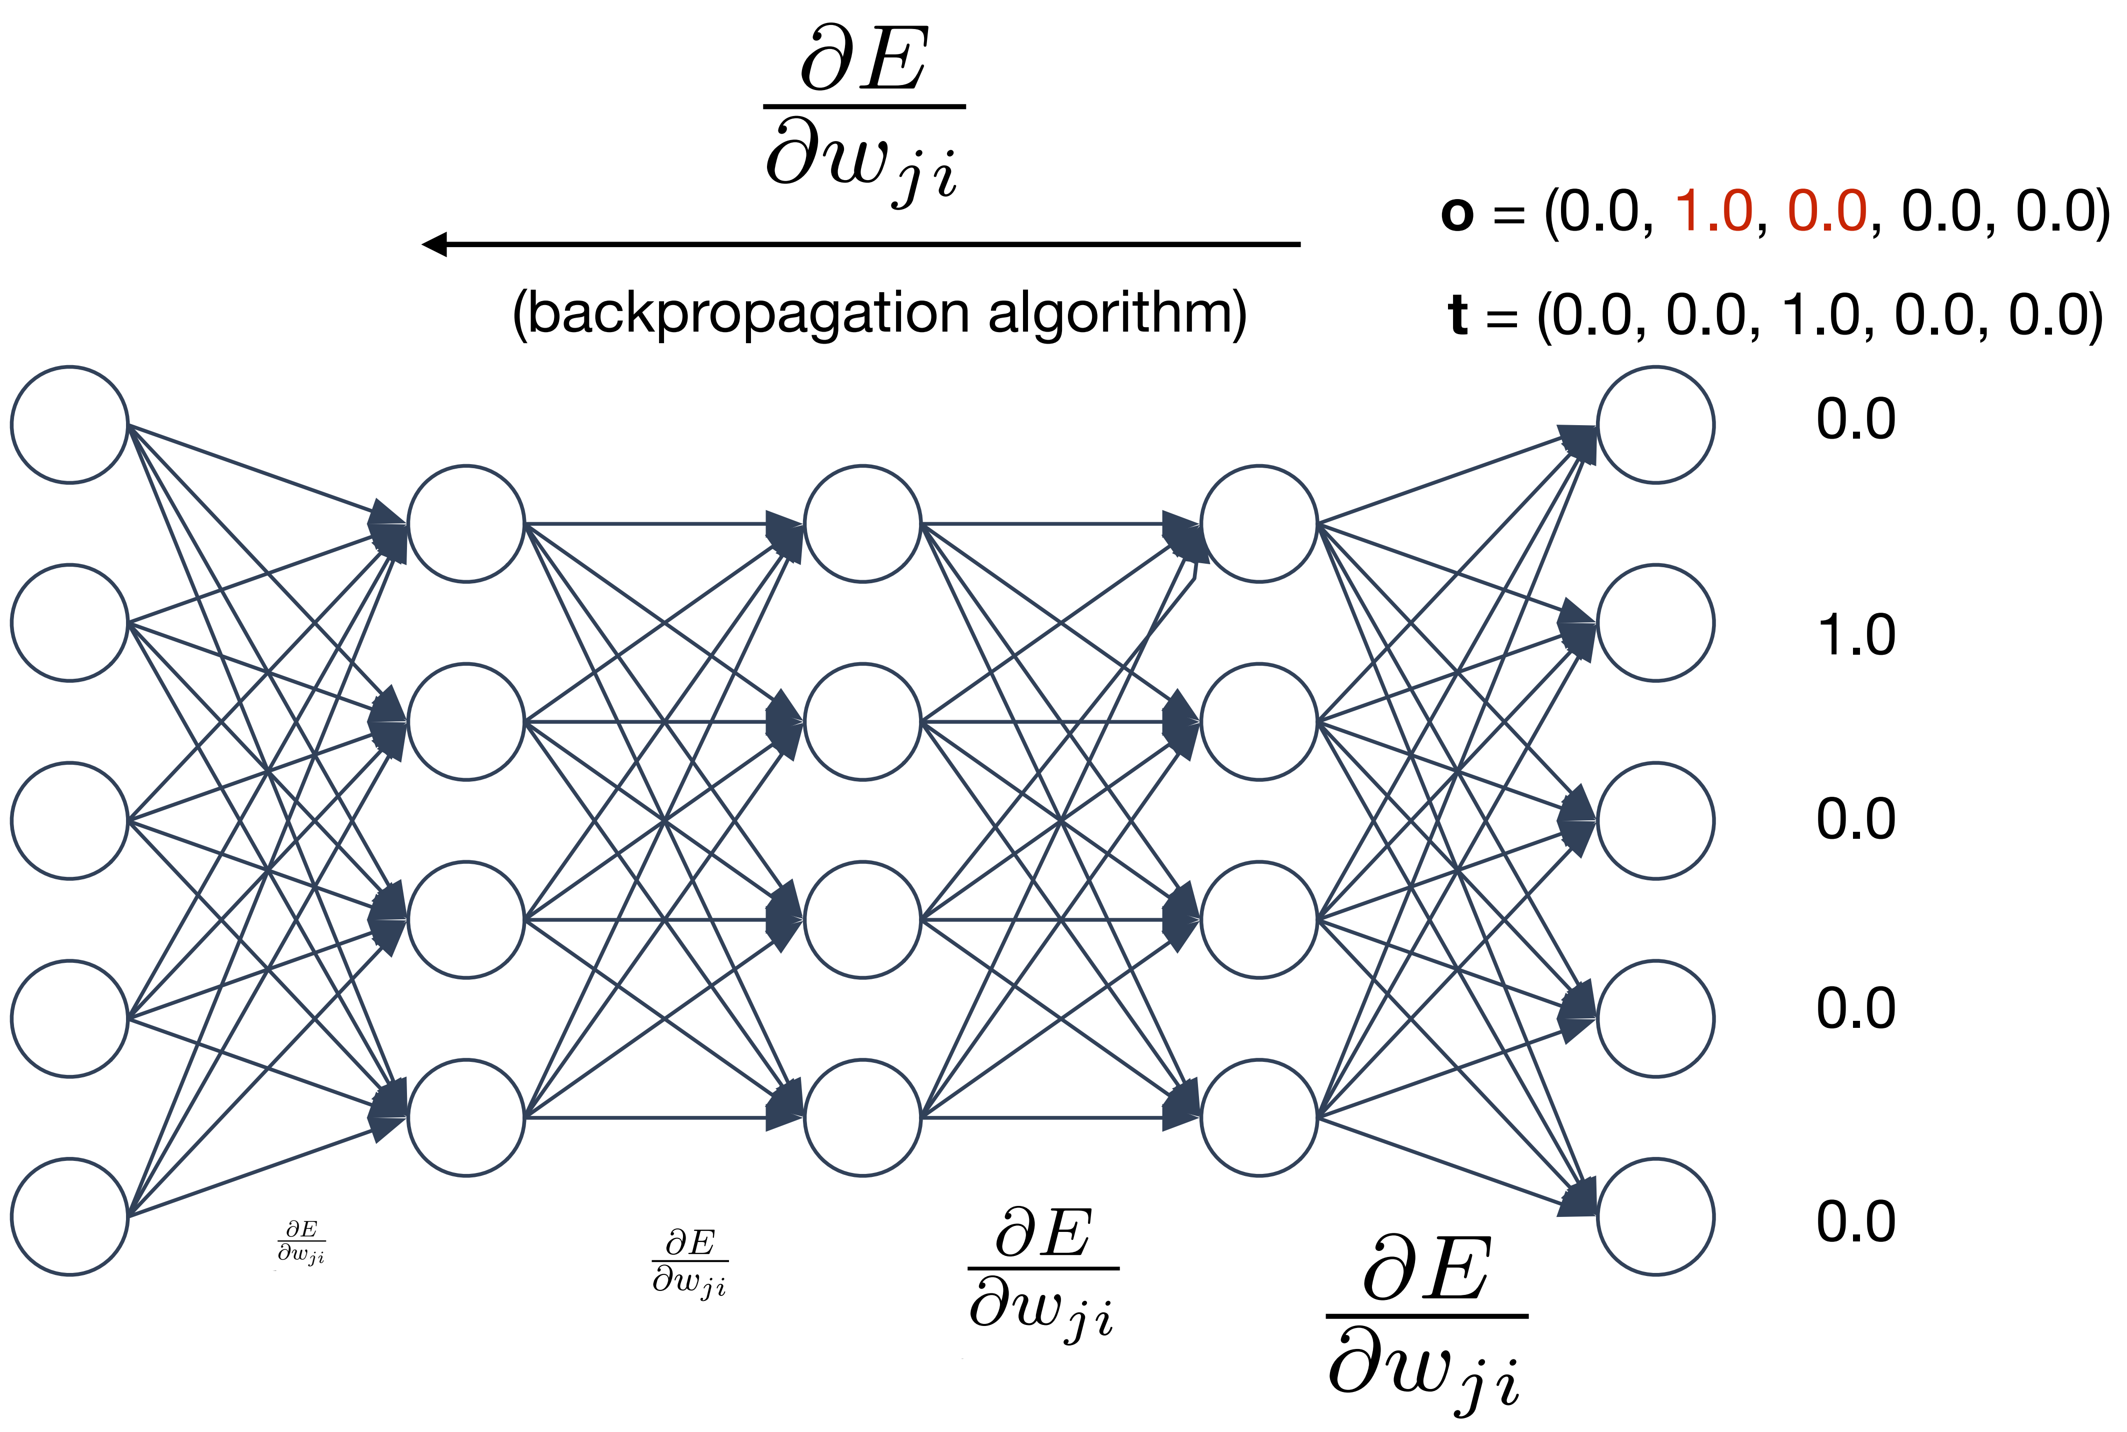

S. Hochreiter. Untersuchungen zu dynamischen neuronalen Netzen. Diploma thesis, Institut f. Informatik, Technische Univ. Munich, 1991. Advisor: J. Schmidhuber

Time passes …

THE MNIST DATABASE

of handwritten digits

Yann LeCun, Courant Institute, NYU
Corinna Cortes, Google Labs, New York
Christopher J.C. Burges, Microsoft Research, Redmond

Neural Nets			
2-layer NN, 300 hidden units, mean square error	none	4.7	LeCun et al. 1998
2-layer NN, 300 HU, MSE, [distortions]	none	3.6	LeCun et al. 1998
2-layer NN, 300 HU	deskewing	1.6	LeCun et al. 1998
2-layer NN, 1000 hidden units	none	4.5	LeCun et al. 1998
2-layer NN, 1000 HU, [distortions]	none	3.8	LeCun et al. 1998
3-layer NN, 300+100 hidden units	none	3.05	LeCun et al. 1998
3-layer NN, 300+100 HU [distortions]	none	2.5	LeCun et al. 1998
3-layer NN, 500+150 hidden units	none	2.95	LeCun et al. 1998
3-layer NN, 500+150 HU [distortions]	none	2.45	LeCun et al. 1998
3-layer NN, 500+300 HU, softmax, cross entropy, weight decay	none	1.53	Hinton, unpublished, 2005
2-layer NN, 800 HU, Cross-Entropy Loss	none	1.6	Simard et al., ICDAR 2003
2-layer NN, 800 HU, cross-entropy [affine distortions]	none	1.1	Simard et al., ICDAR 2003
2-layer NN, 800 HU, MSE [elastic distortions]	none	0.9	Simard et al., ICDAR 2003
2-layer NN, 800 HU, cross-entropy [elastic distortions]	none	0.7	Simard et al., ICDAR 2003
NN, 784-500-500-2000-30 + nearest neighbor, RBM + NCA training [no distortions]	none	1.0	Salakhutdinov and Hinton, AI-Stats 2007

THE MNIST DATABASE

of handwritten digits

Yann LeCun, Courant Institute, NYU
Corinna Cortes, Google Labs, New York
Christopher J.C. Burges, Microsoft Research, Redmond

Neural Nets			
2-layer NN, 300 hidden units, mean square error	none	4.7	LeCun et al. 1998
2-layer NN, 300 HU, MSE, [distortions]	none	3.6	LeCun et al. 1998
2-layer NN, 300 HU	deskewing	1.6	LeCun et al. 1998
2-layer NN, 1000 hidden units	none	4.5	LeCun et al. 1998
2-layer NN, 1000 HU, [distortions]	none	3.8	LeCun et al. 1998
3-layer NN, 300+100 hidden units	none	3.05	LeCun et al. 1998
3-layer NN, 300+100 HU [distortions]	none	2.5	LeCun et al. 1998
3-layer NN, 500+150 hidden units	none	2.95	LeCun et al. 1998
3-layer NN, 500+150 HU [distortions]	none	2.45	LeCun et al. 1998
3-layer NN, 500+300 HU, softmax, cross entropy, weight decay	none	1.53	Hinton, unpublished, 2005
2-layer NN, 800 HU, Cross-Entropy Loss	none	1.6	Simard et al., ICDAR 2003
2-layer NN, 800 HU, cross-entropy [affine distortions]	none	1.1	Simard et al., ICDAR 2003
2-layer NN, 800 HU, MSE [elastic distortions]	none	0.9	Simard et al., ICDAR 2003
2-layer NN, 800 HU, cross-entropy [elastic distortions]	none	0.7	Simard et al., ICDAR 2003
NN, 784-500-500-2000-30 + nearest neighbor, RBM + NCA training [no distortions]	none	1.0	Salakhutdinov and Hinton, AI-Stats 2007

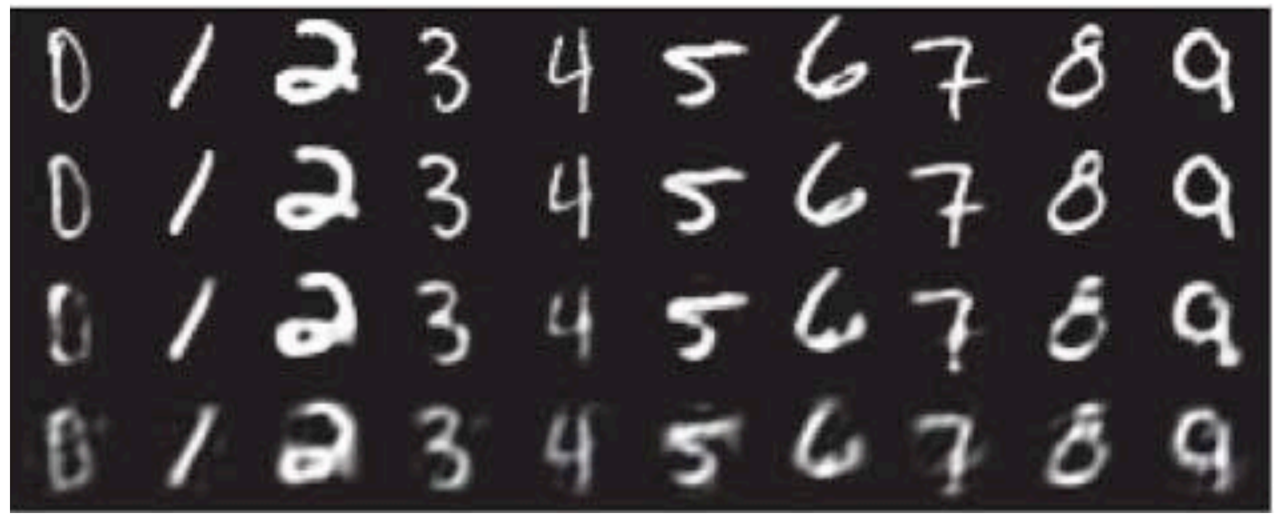

THE MNIST DATABASE

of handwritten digits

Yann LeCun, Courant Institute, NYU
Corinna Cortes, Google Labs, New York
Christopher J.C. Burges, Microsoft Research, Redmond

Neural Nets			
2-layer NN, 300 hidden units, mean square error	none	4.7	LeCun et al. 1998
2-layer NN, 300 HU, MSE, [distortions]	none	3.6	LeCun et al. 1998
2-layer NN, 300 HU	deskewing	1.6	LeCun et al. 1998
2-layer NN, 1000 hidden units	none	4.5	LeCun et al. 1998
2-layer NN, 1000 HU, [distortions]	none	3.8	LeCun et al. 1998
3-layer NN, 300+100 hidden units	none	3.05	LeCun et al. 1998
3-layer NN, 300+100 HU [distortions]	none	2.5	LeCun et al. 1998
3-layer NN, 500+150 hidden units	none	2.95	LeCun et al. 1998
3-layer NN, 500+150 HU [distortions]	none	2.45	LeCun et al. 1998
3-layer NN, 500+300 HU, softmax, cross entropy, weight decay	none	1.53	Hinton, unpublished, 2005
2-layer NN, 800 HU, Cross-Entropy Loss	none	1.6	Simard et al., ICDAR 2003
2-layer NN, 800 HU, cross-entropy [affine distortions]	none	1.1	Simard et al., ICDAR 2003
2-layer NN, 800 HU, MSE [elastic distortions]	none	0.9	Simard et al., ICDAR 2003
2-layer NN, 800 HU, cross-entropy [elastic distortions]	none	0.7	Simard et al., ICDAR 2003
NN, 784-500-500-2000-30 + nearest neighbor, RBM + NCA training [no distortions]	none	1.0	Salakhutdinov and Hinton, AI-Stats 2007

Reducing the Dimensionality of Data with Neural Networks
G. E. Hinton, *et al.*
Science **313**, 504 (2006);
DOI: 10.1126/science.1127647

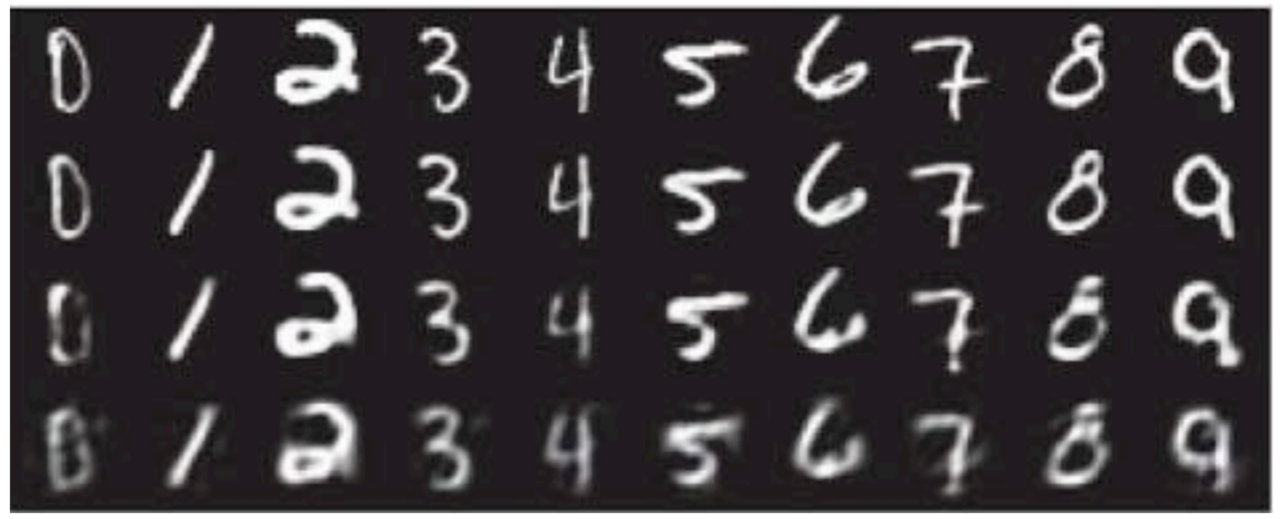

THE MNIST DATABASE

of handwritten digits

Yann LeCun, Courant Institute, NYU
Corinna Cortes, Google Labs, New York
Christopher J.C. Burges, Microsoft Research, Redmond

Neural Nets			
2-layer NN, 300 hidden units, mean square error	none	4.7	LeCun et al. 1998
2-layer NN, 300 HU, MSE, [distortions]	none	3.6	LeCun et al. 1998
2-layer NN, 300 HU	deskewing	1.6	LeCun et al. 1998
2-layer NN, 1000 hidden units	none	4.5	LeCun et al. 1998
2-layer NN, 1000 HU, [distortions]	none	3.8	LeCun et al. 1998
3-layer NN, 300+100 hidden units	none	3.05	LeCun et al. 1998
3-layer NN, 300+100 HU [distortions]	none	2.5	LeCun et al. 1998
3-layer NN, 500+150 hidden units	none	2.95	LeCun et al. 1998
3-layer NN, 500+150 HU [distortions]	none	2.45	LeCun et al. 1998
3-layer NN, 500+300 HU, softmax, cross entropy, weight decay	none	1.53	Hinton, unpublished, 2005
2-layer NN, 800 HU, Cross-Entropy Loss	none	1.6	Simard et al., ICDAR 2003
2-layer NN, 800 HU, cross-entropy [affine distortions]	none	1.1	Simard et al., ICDAR 2003
2-layer NN, 800 HU, MSE [elastic distortions]	none	0.9	Simard et al., ICDAR 2003
2-layer NN, 800 HU, cross-entropy [elastic distortions]	none	0.7	Simard et al., ICDAR 2003
NN, 784-500-500-2000-30 + nearest neighbor, RBM + NCA training [no distortions]	none	1.0	Salakhutdinov and Hinton, AI-Stats 2007

deep networks!?!

Reducing the Dimensionality of Data with Neural Networks
G. E. Hinton, *et al.*
Science **313**, 504 (2006);
DOI: 10.1126/science.1127647

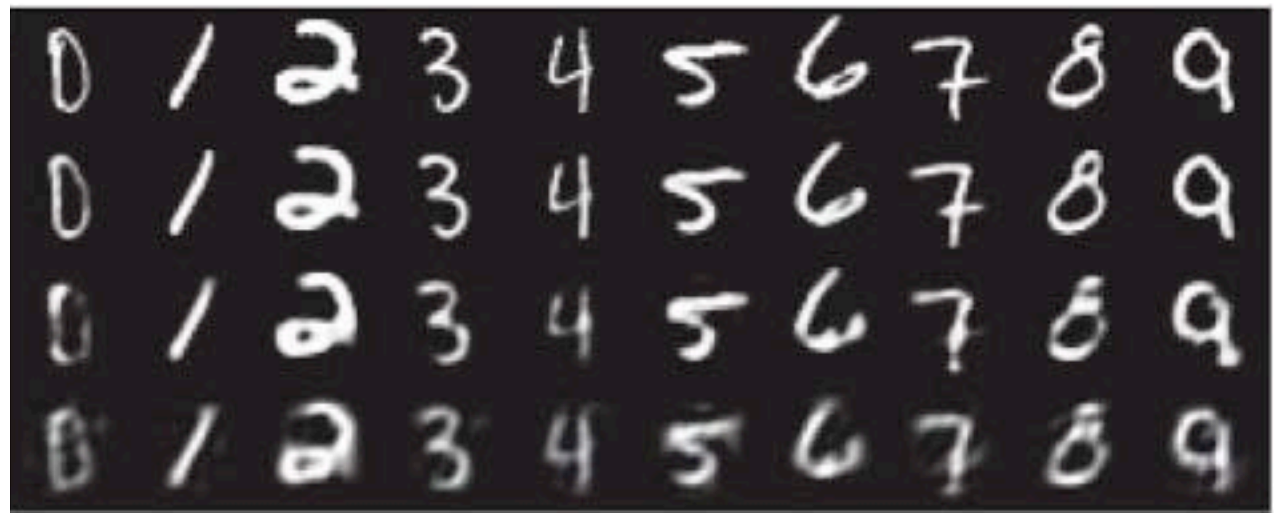

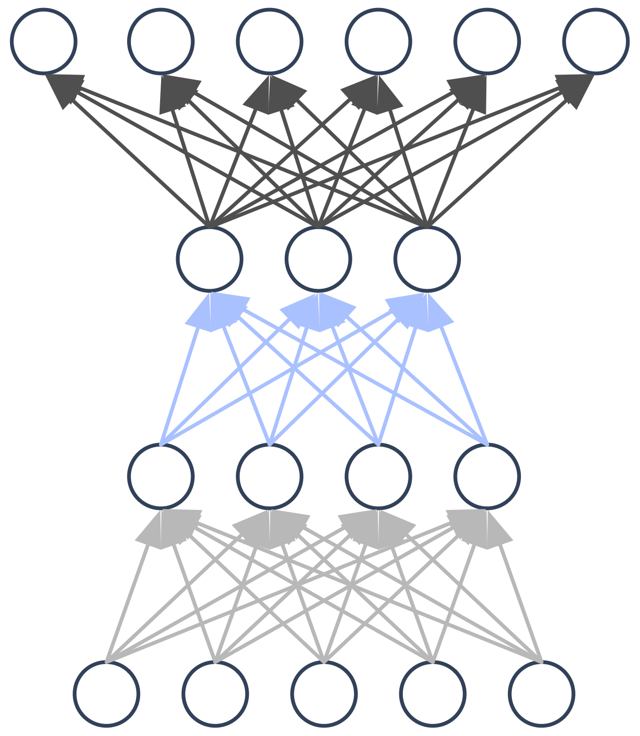

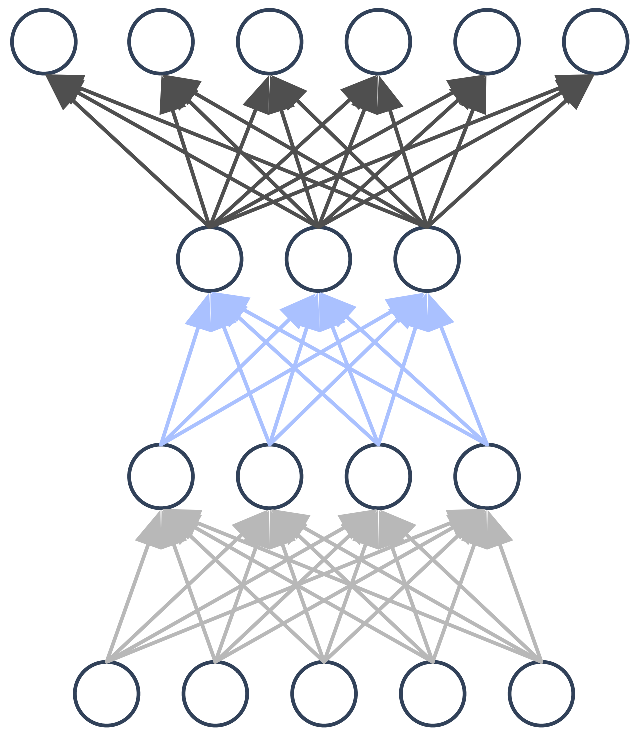

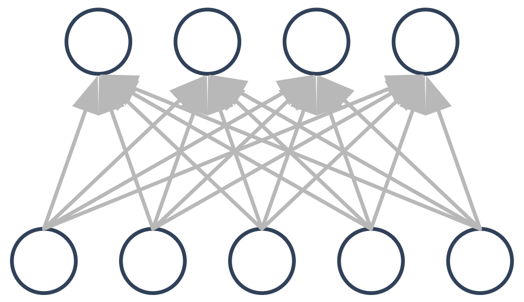

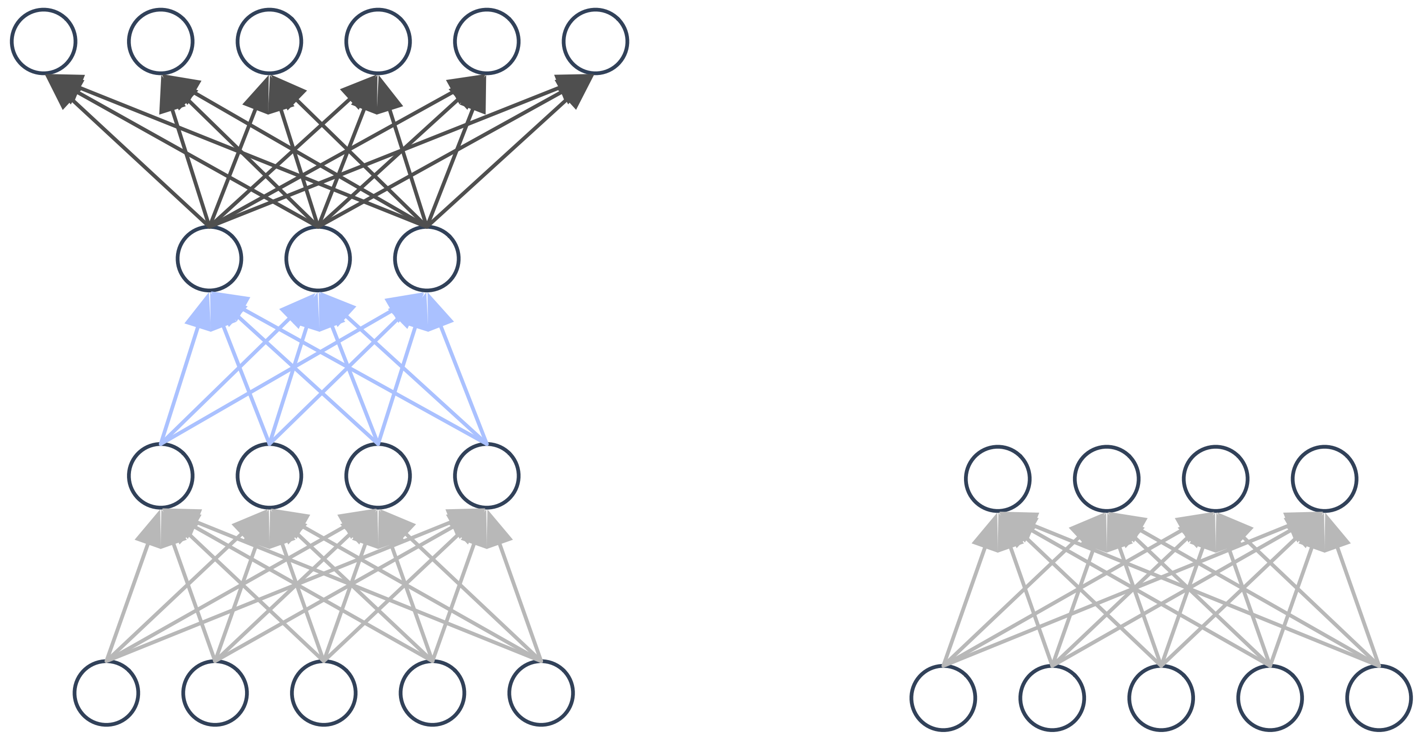

t = ?

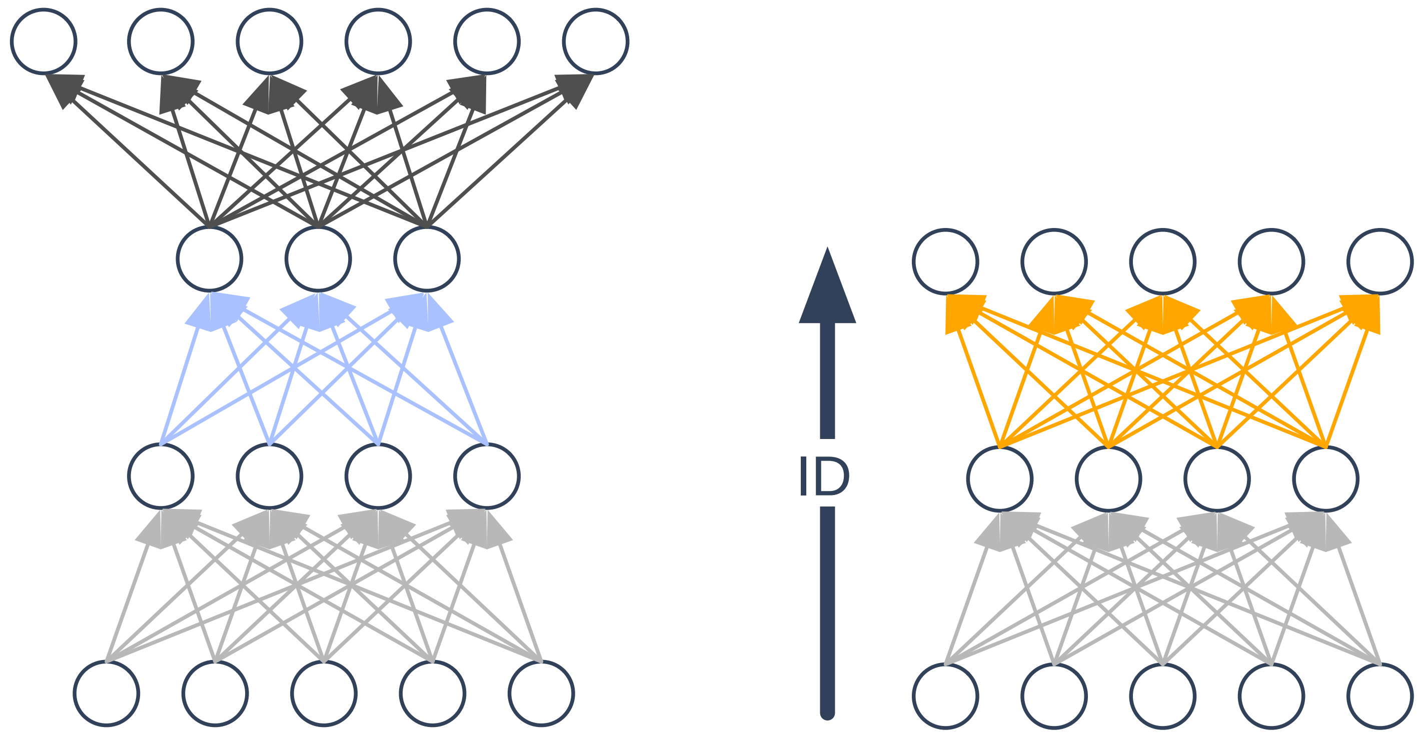

t = i !

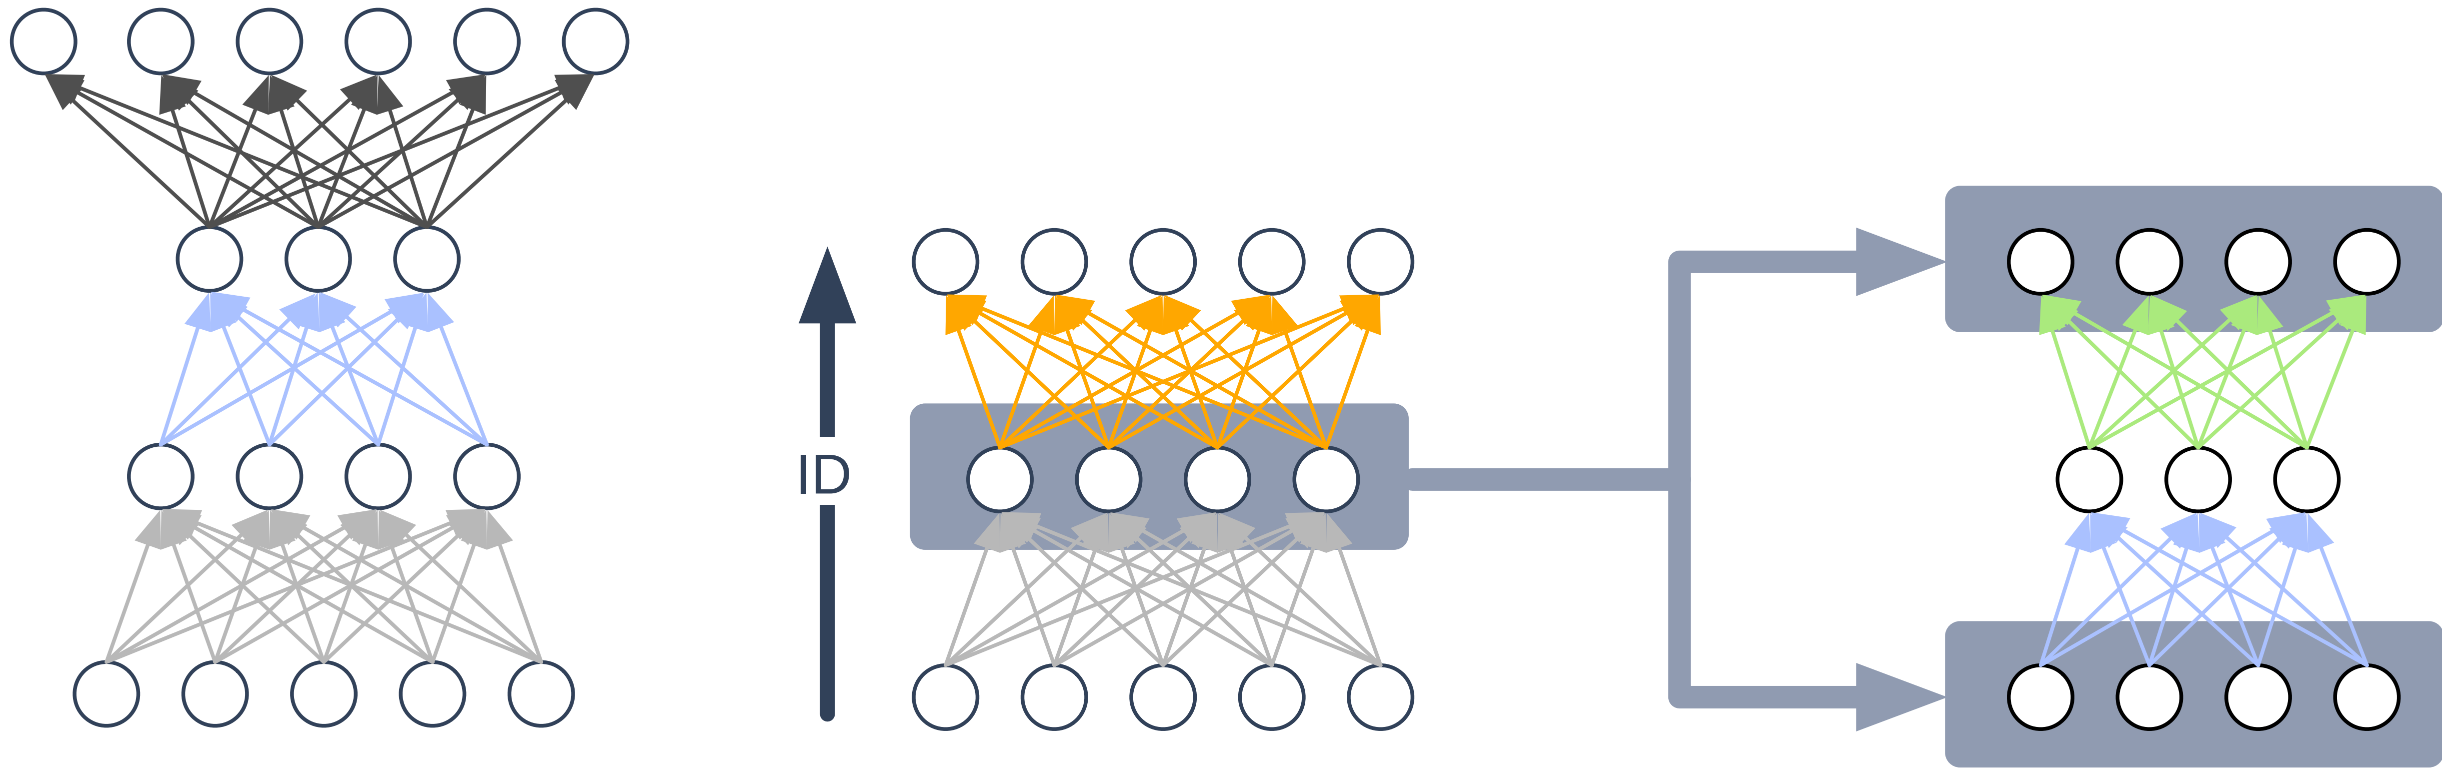
ID

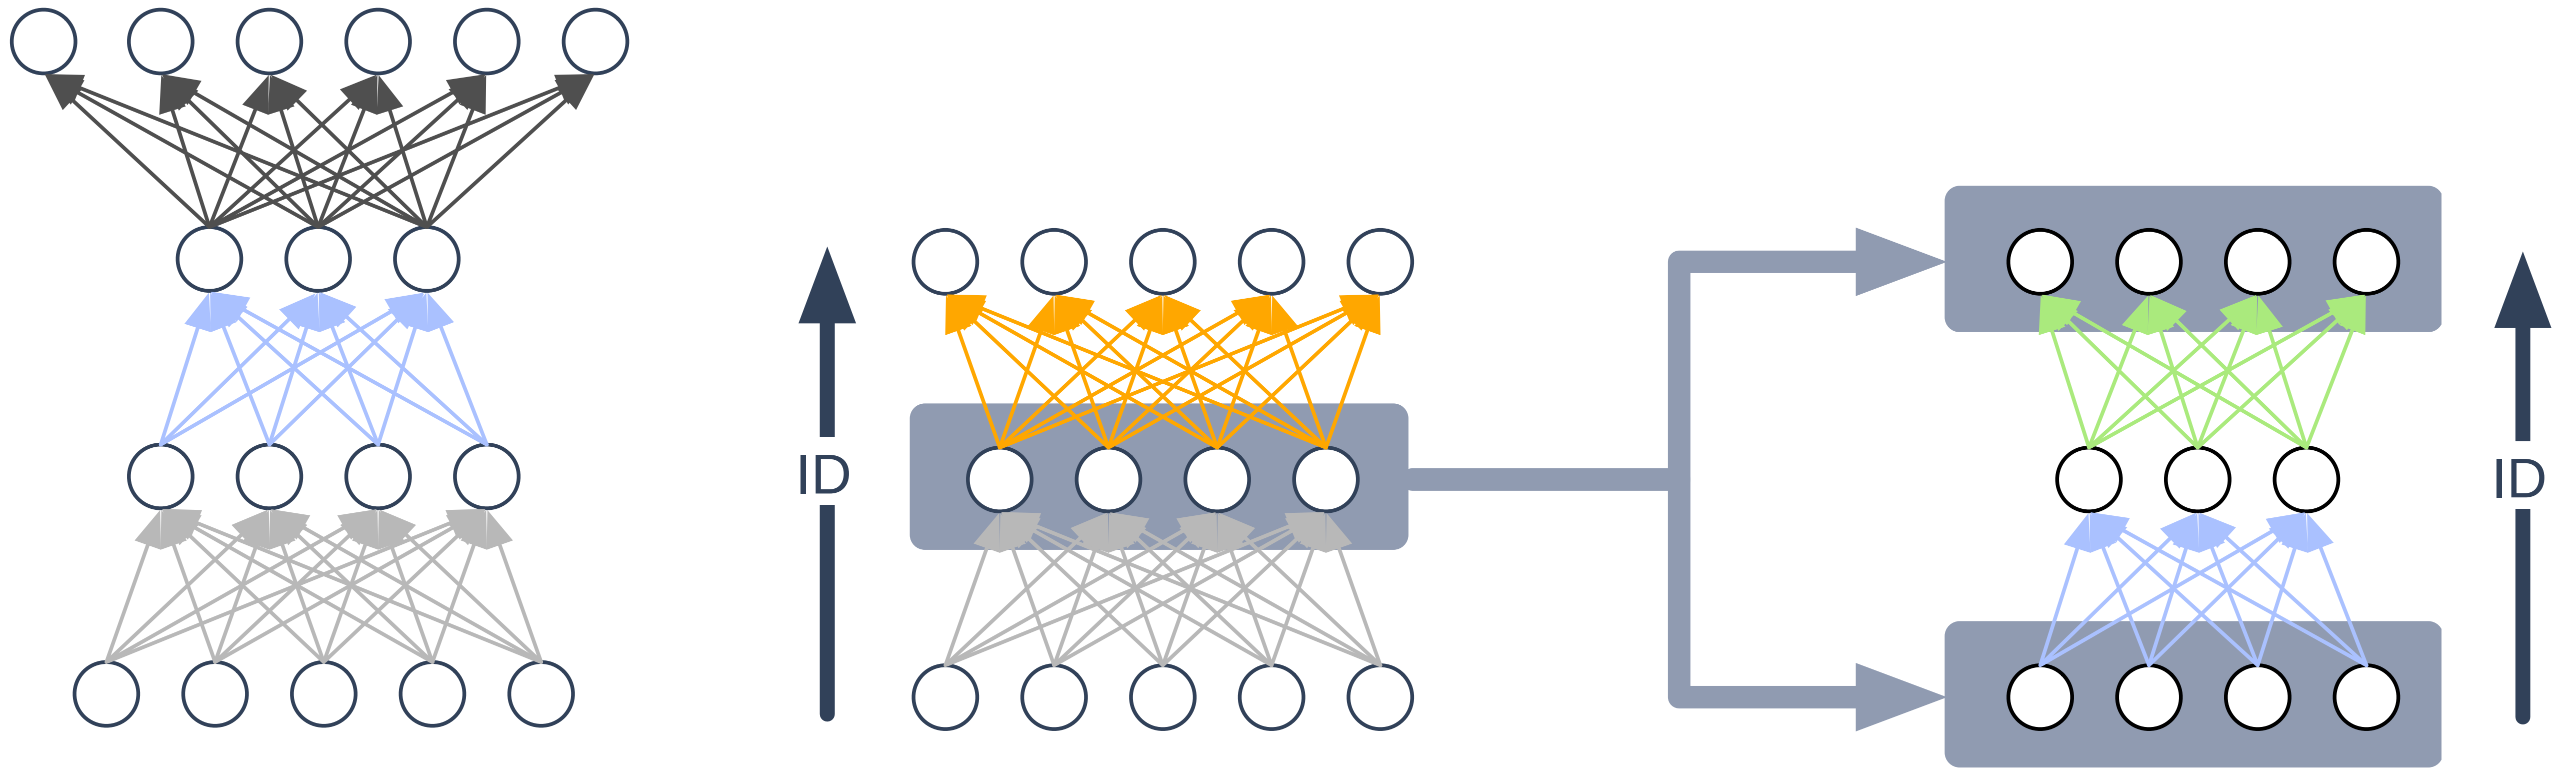
ID
ID

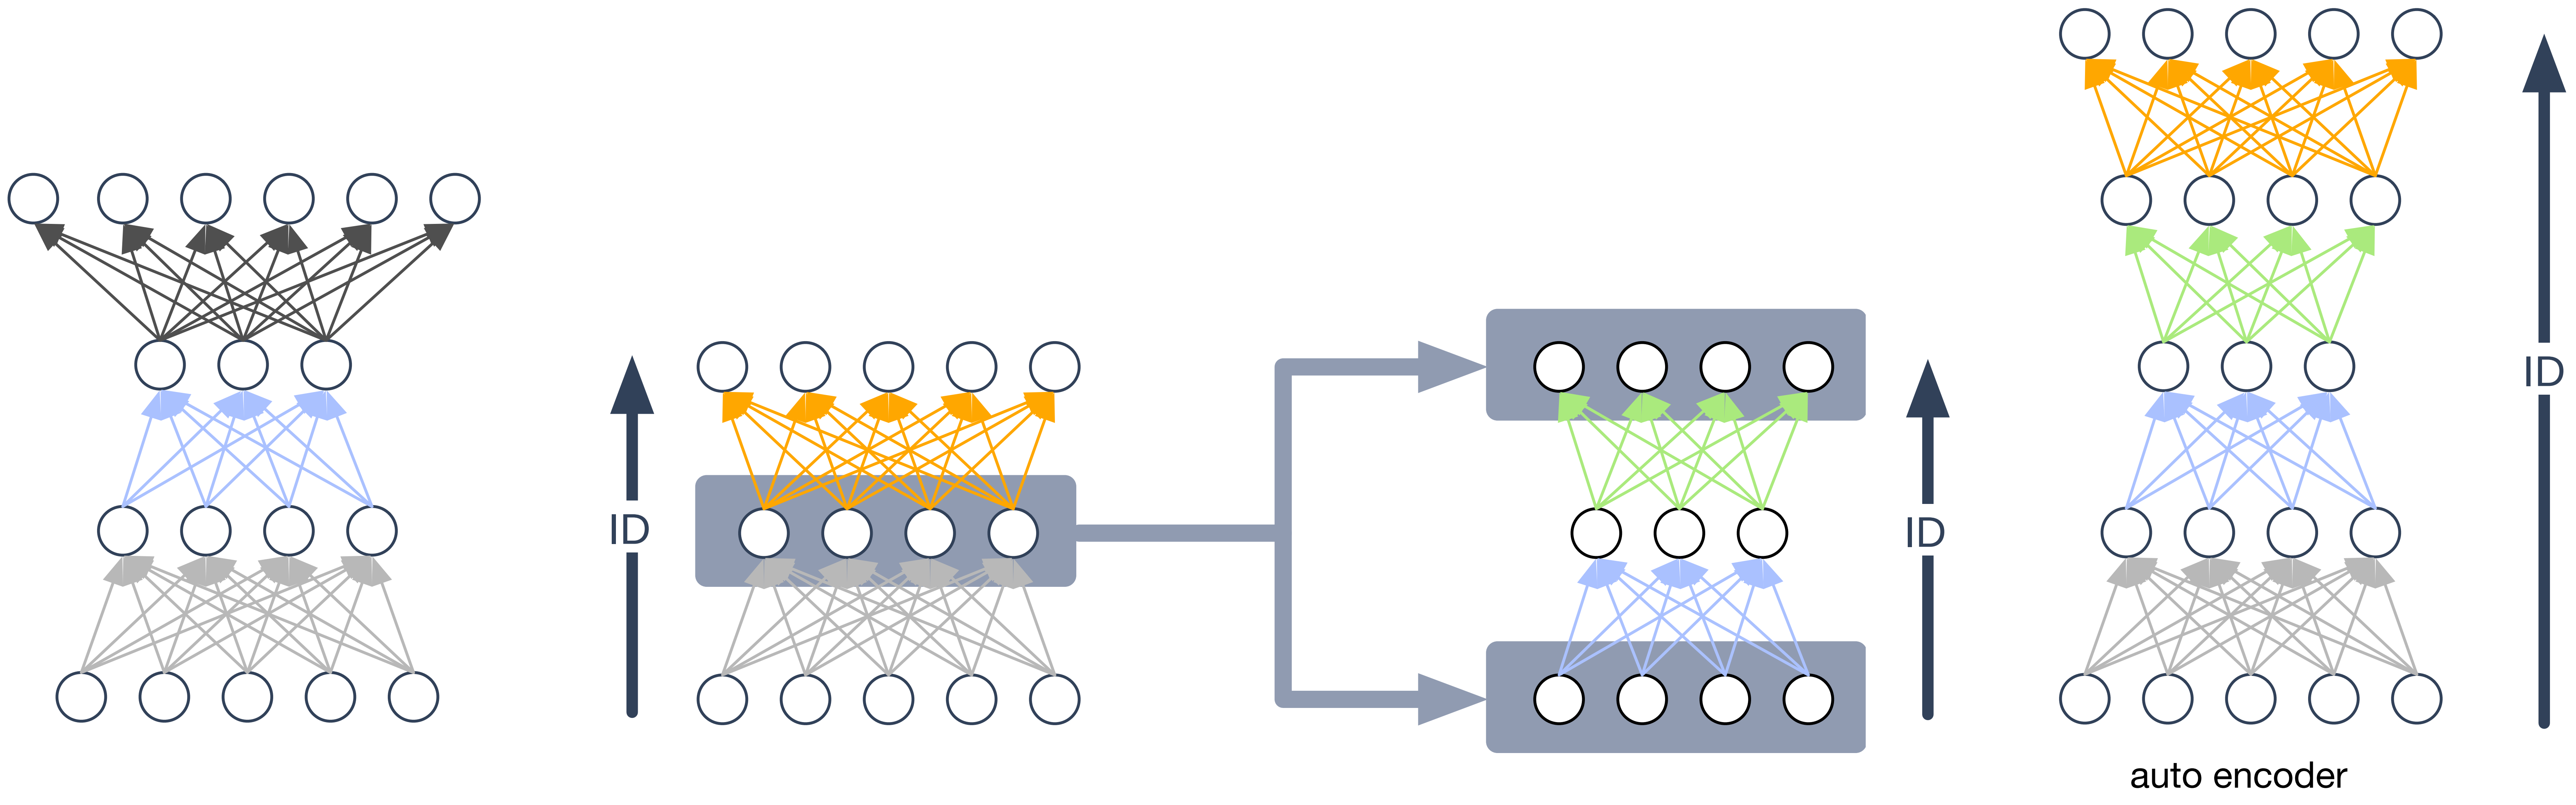
ID
ID
ID
auto encoder

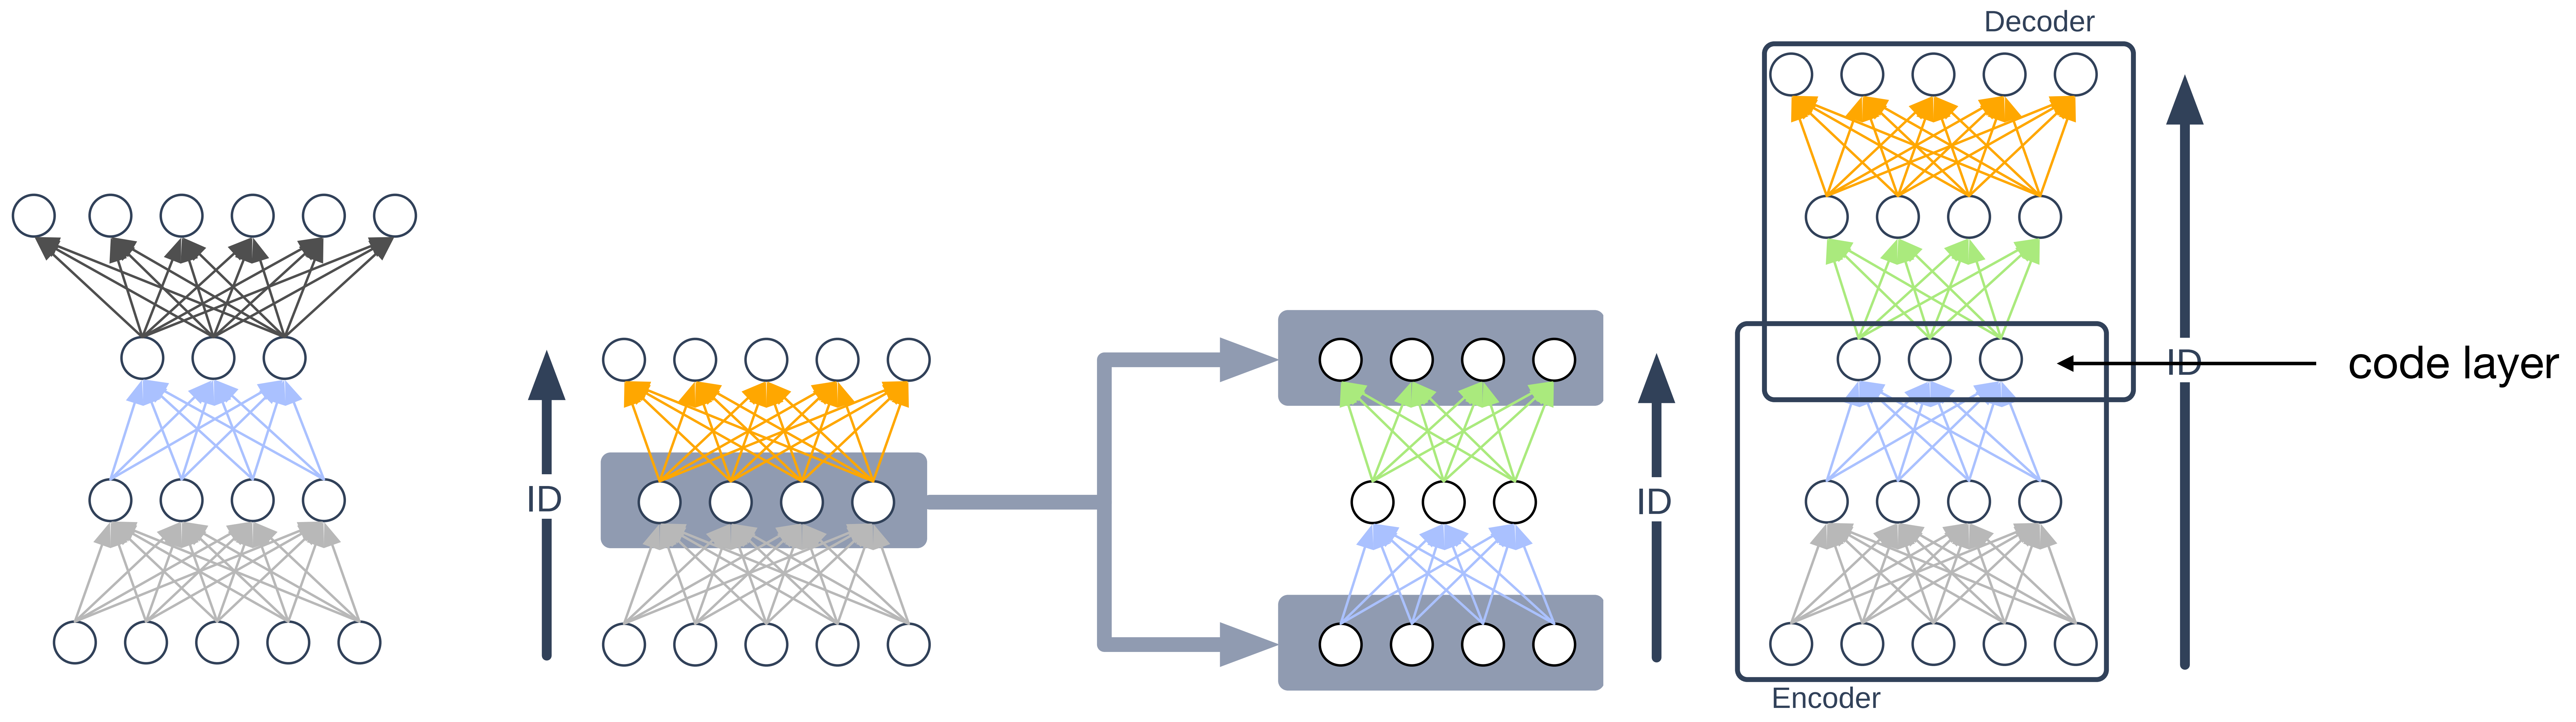

ID
ID
Decoder
ID
code layer
Encoder

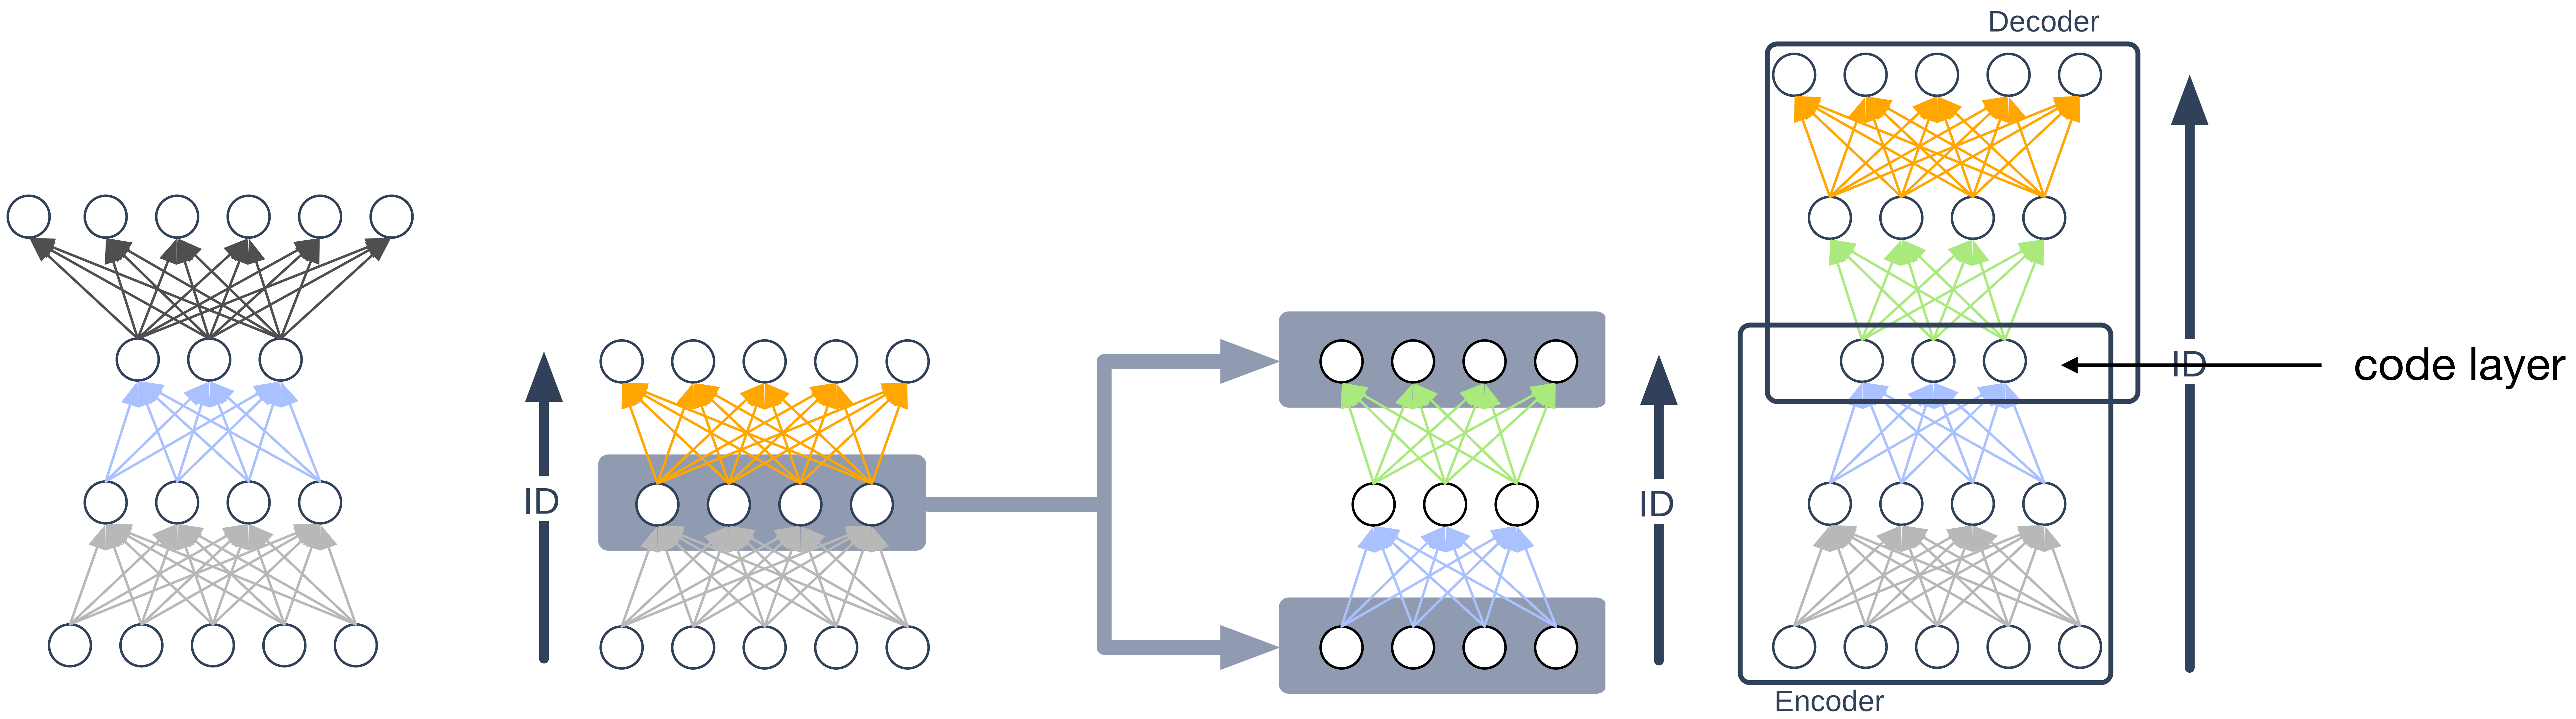

How does the learned "code" look like?

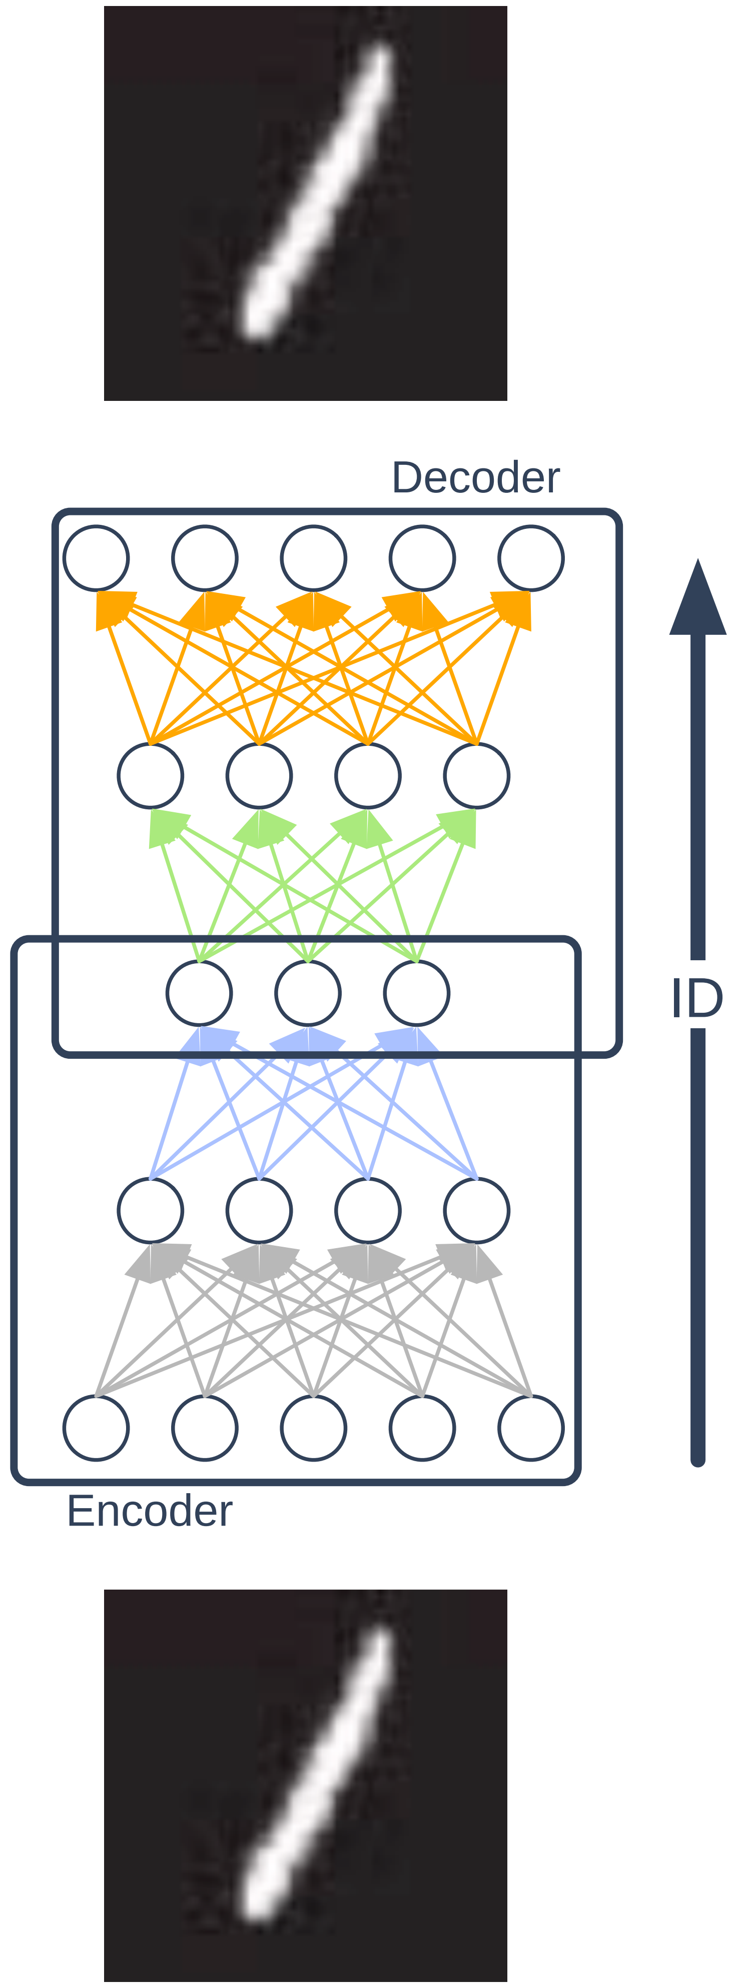

Reducing the Dimensionality of Data with Neural Networks
G. E. Hinton, *et al.*
Science **313**, 504 (2006);
DOI: 10.1126/science.1127647

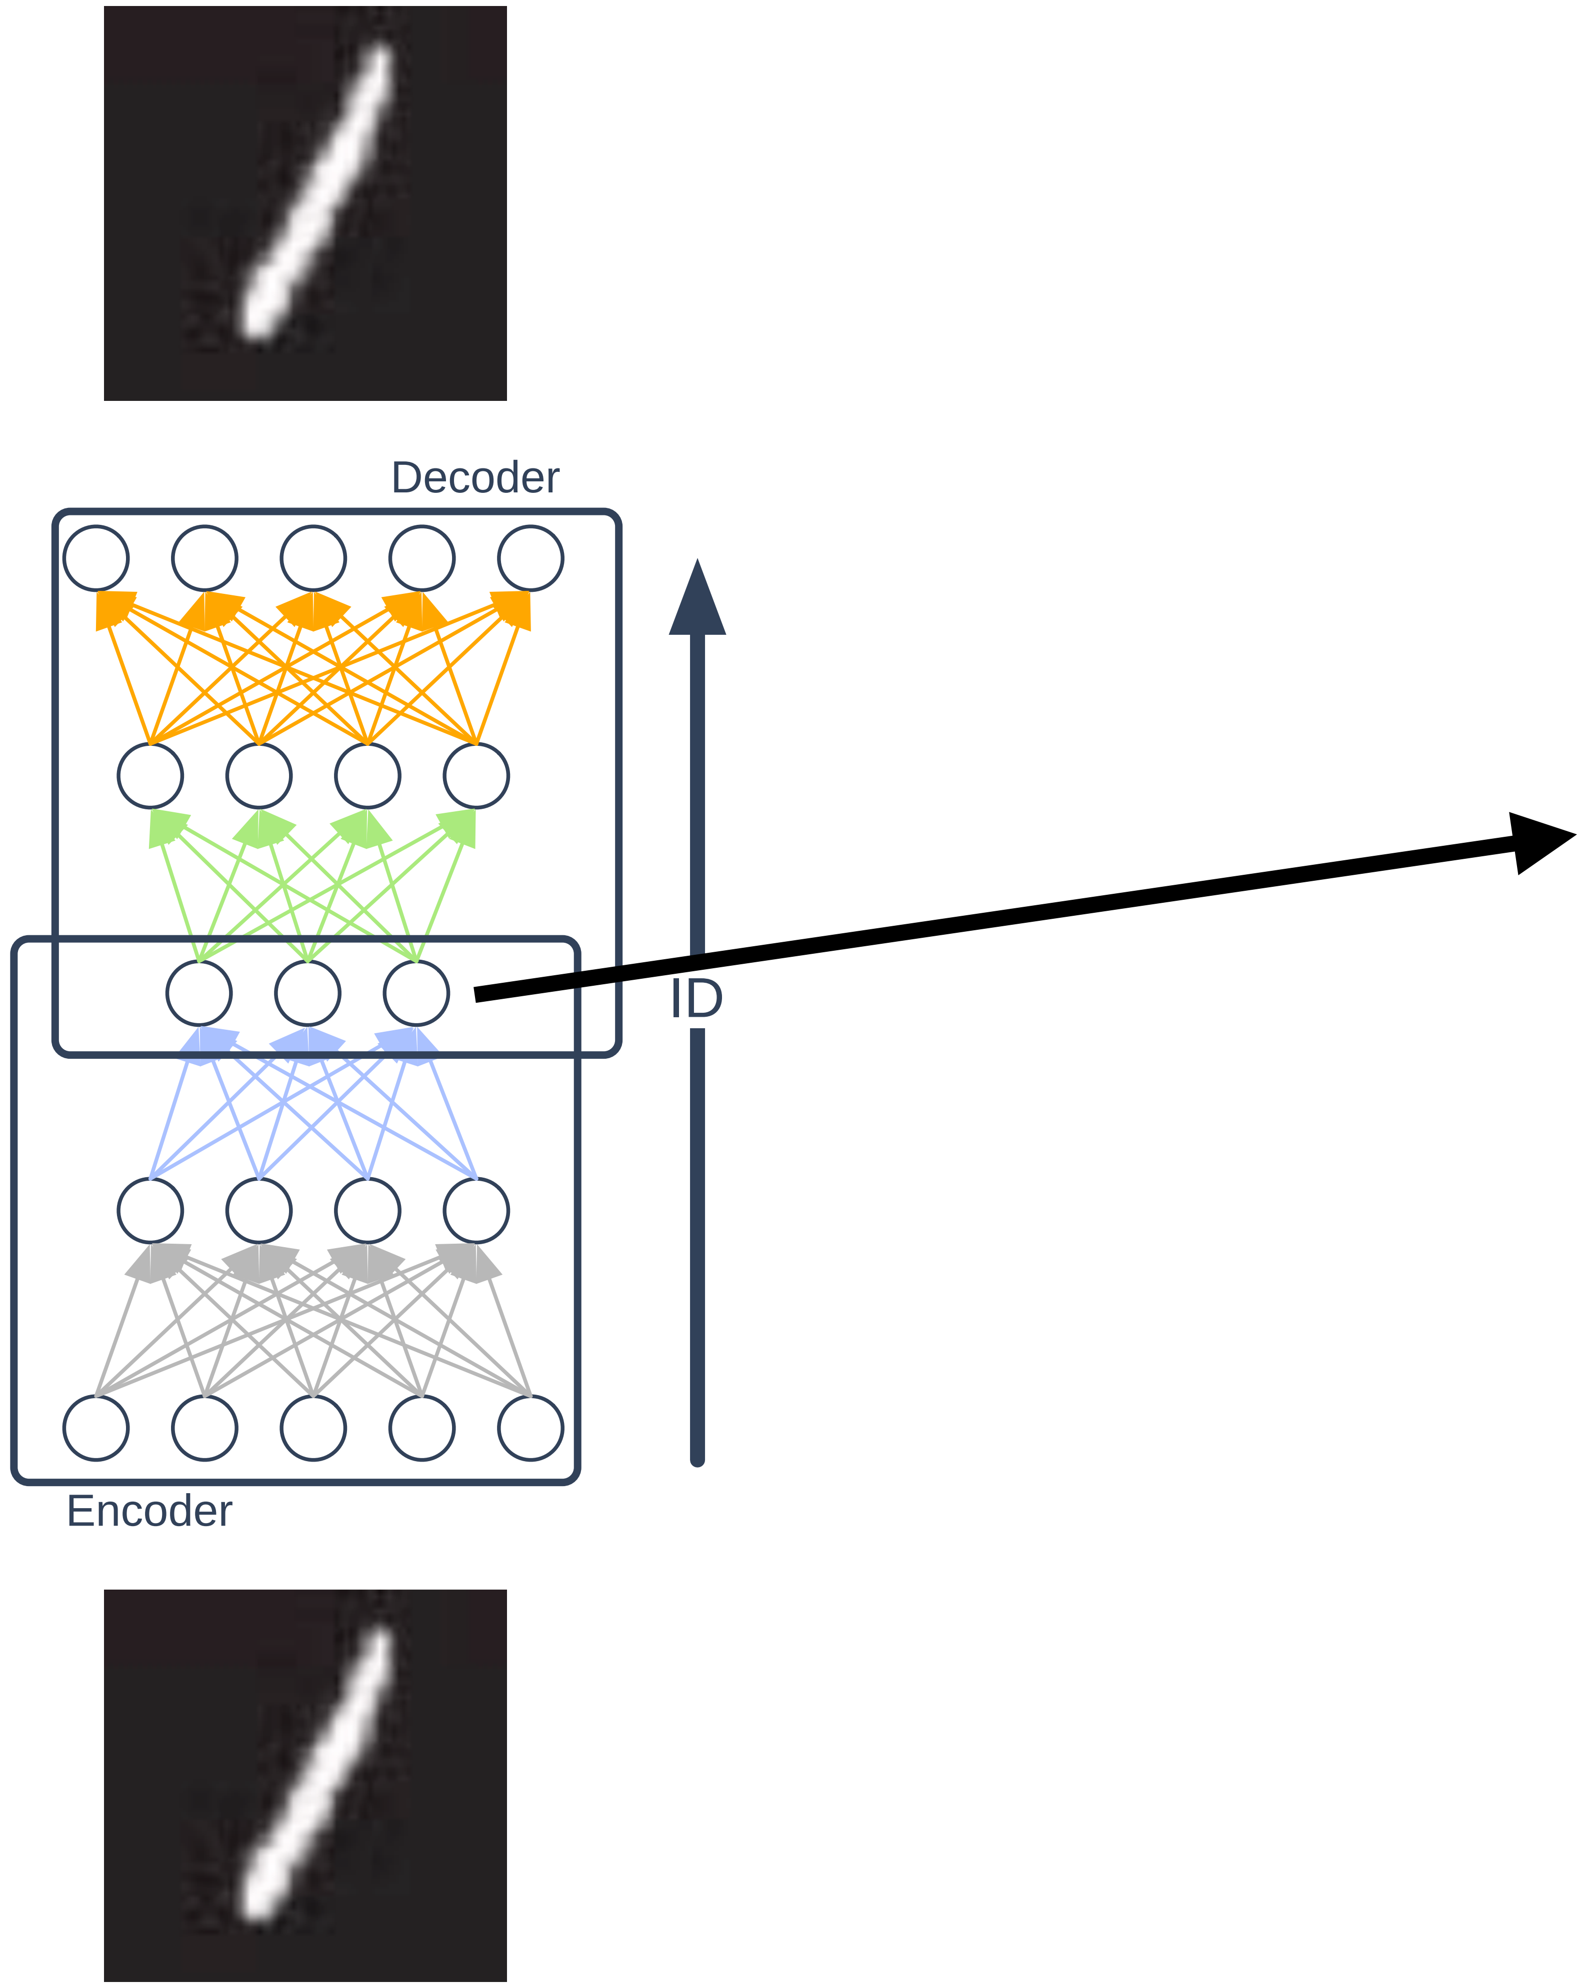

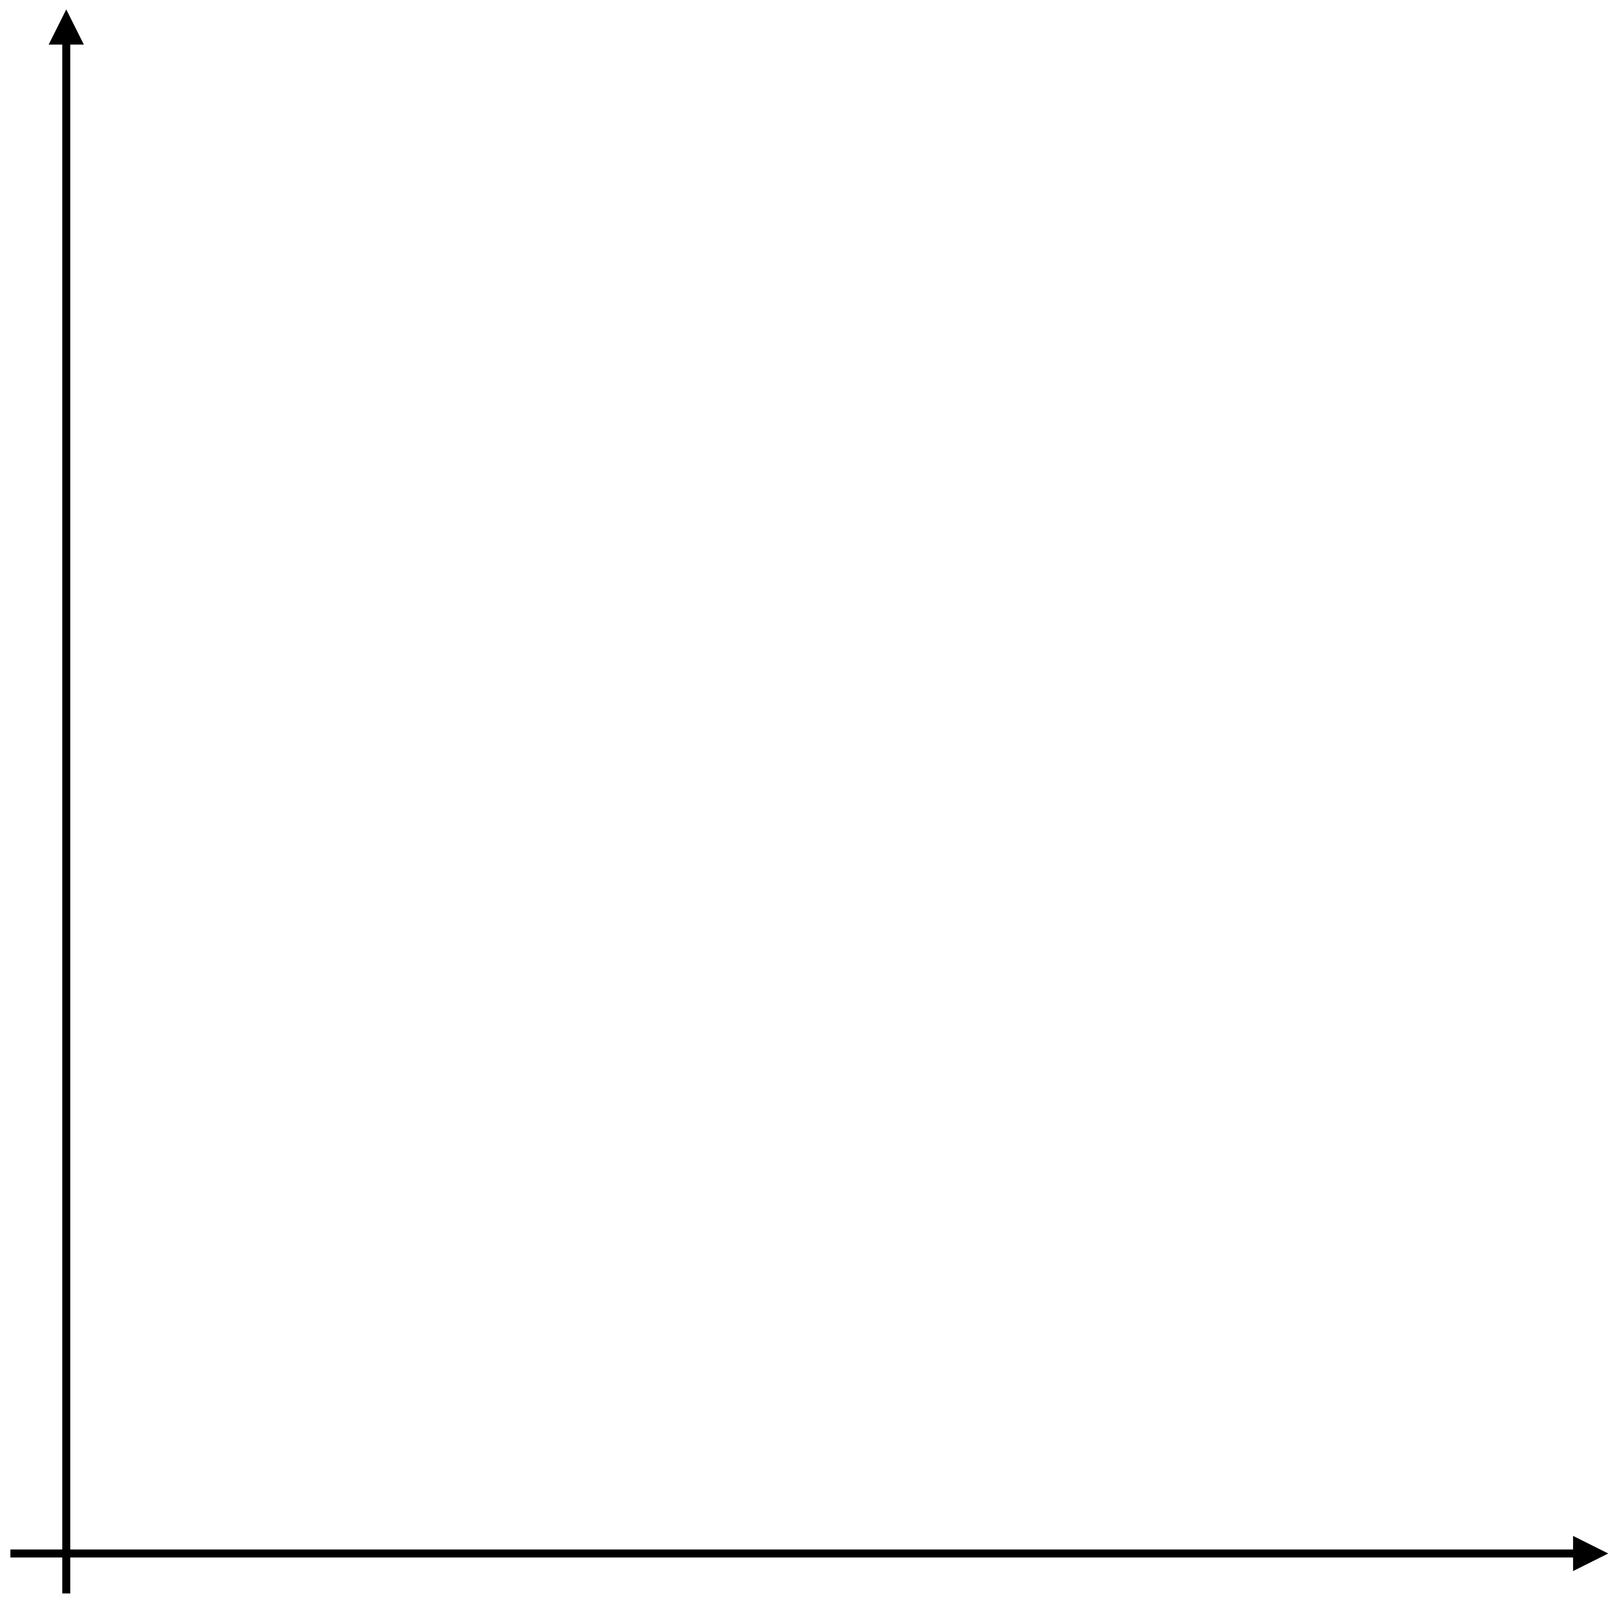

Reducing the Dimensionality of Data with Neural Networks
G. E. Hinton, *et al.*
Science **313**, 504 (2006);
DOI: 10.1126/science.1127647

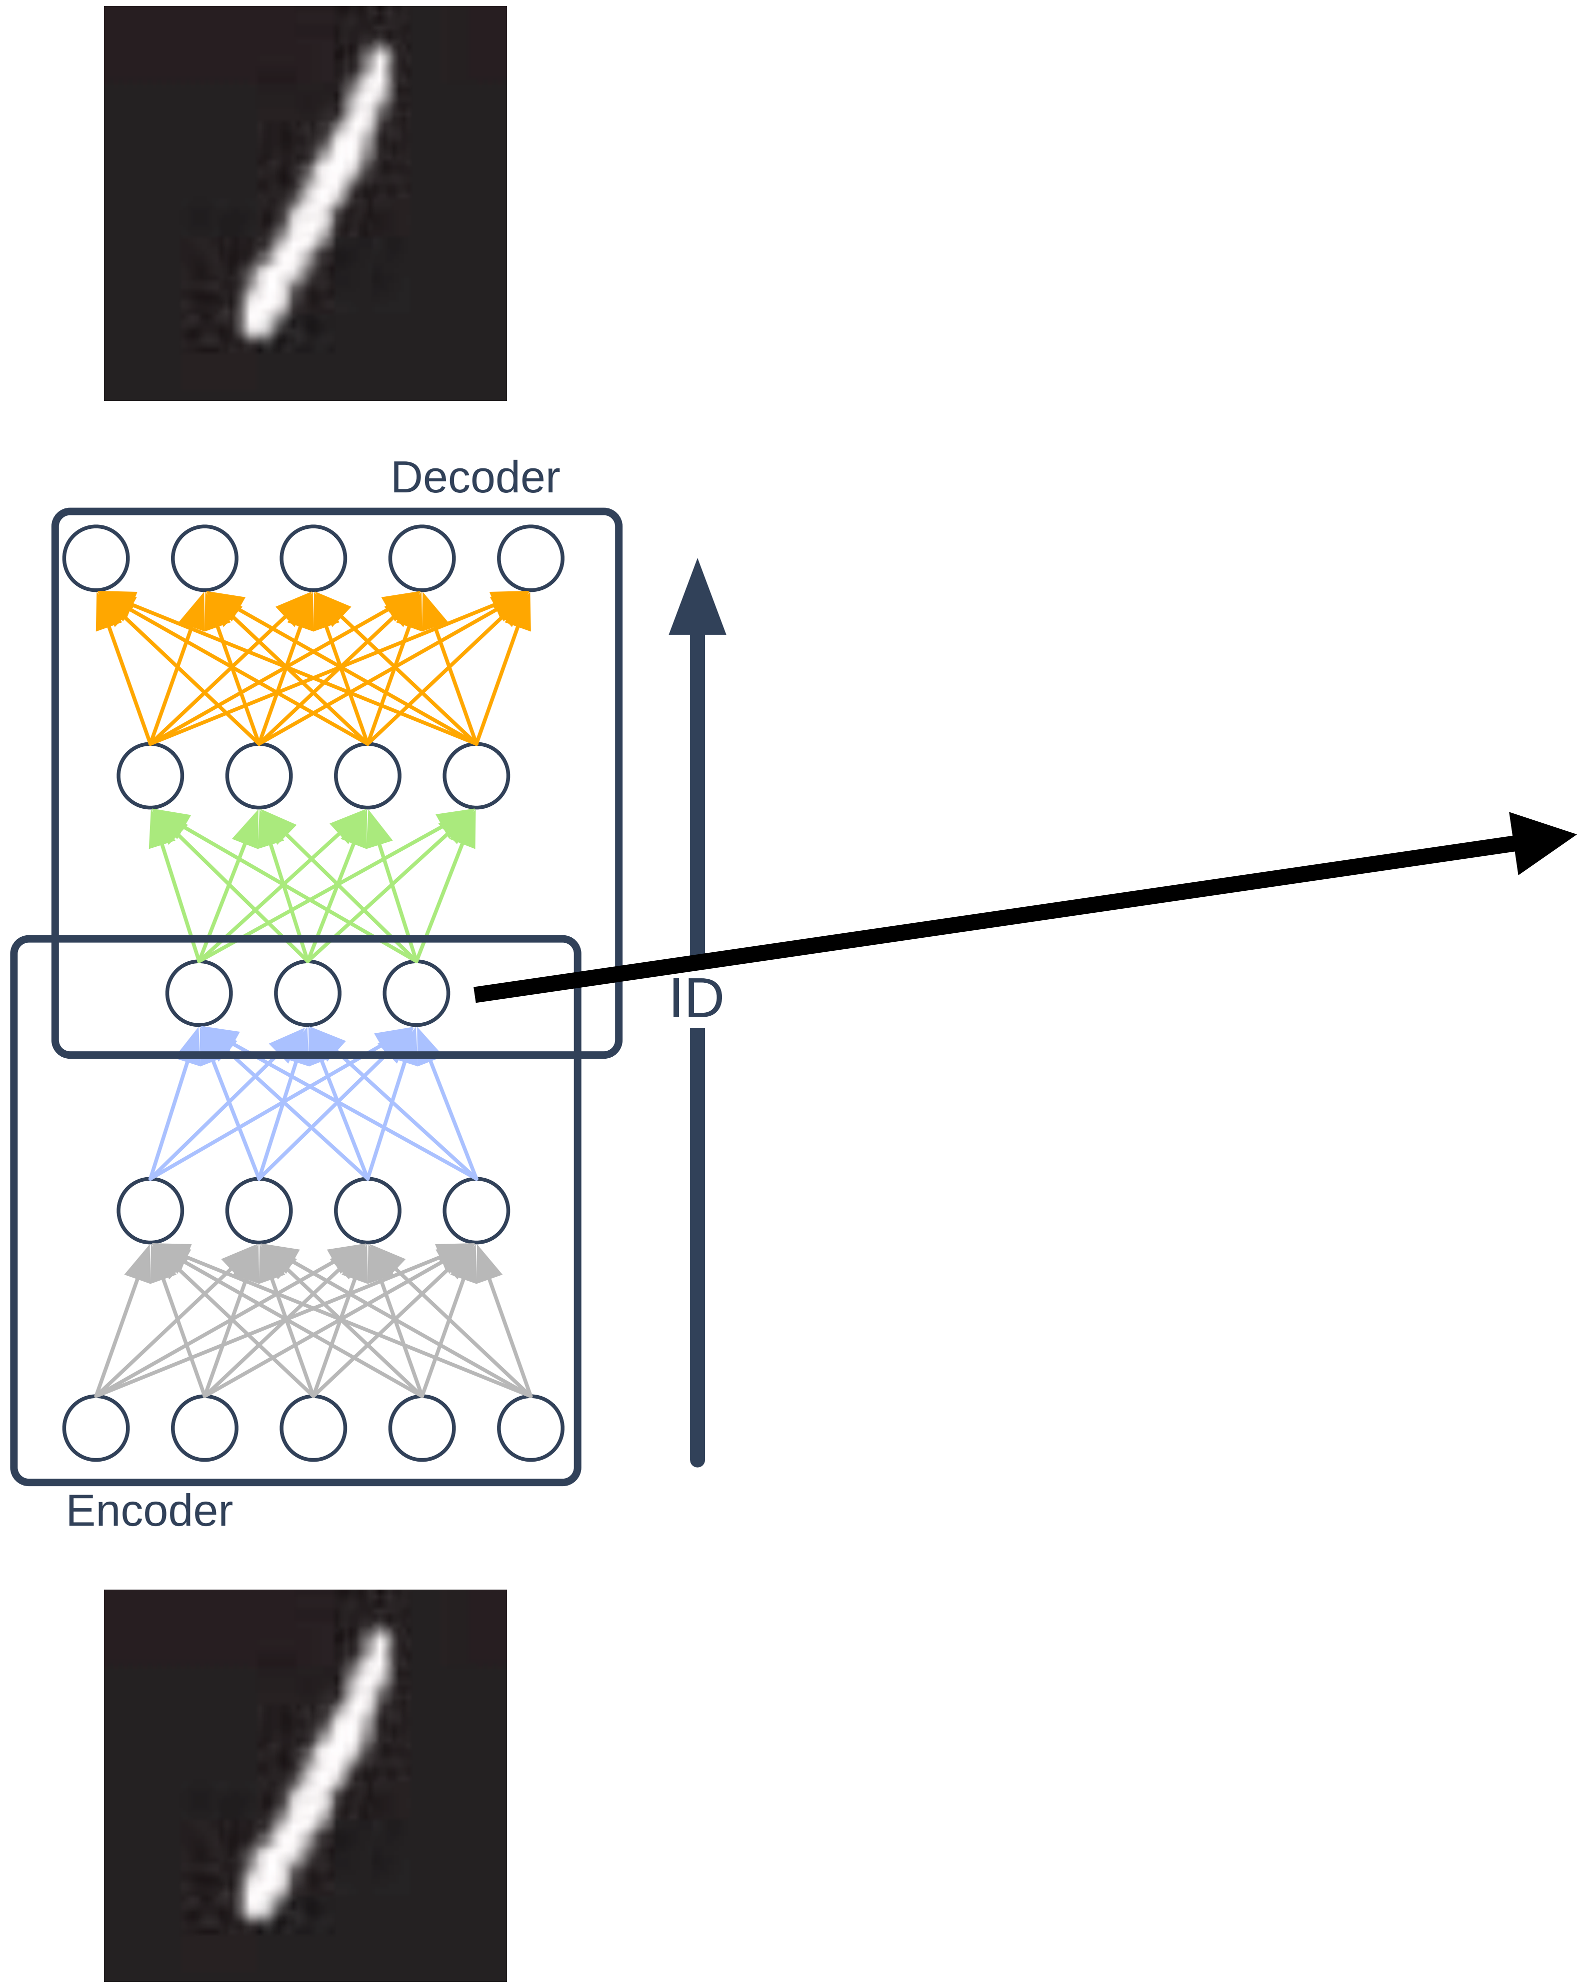

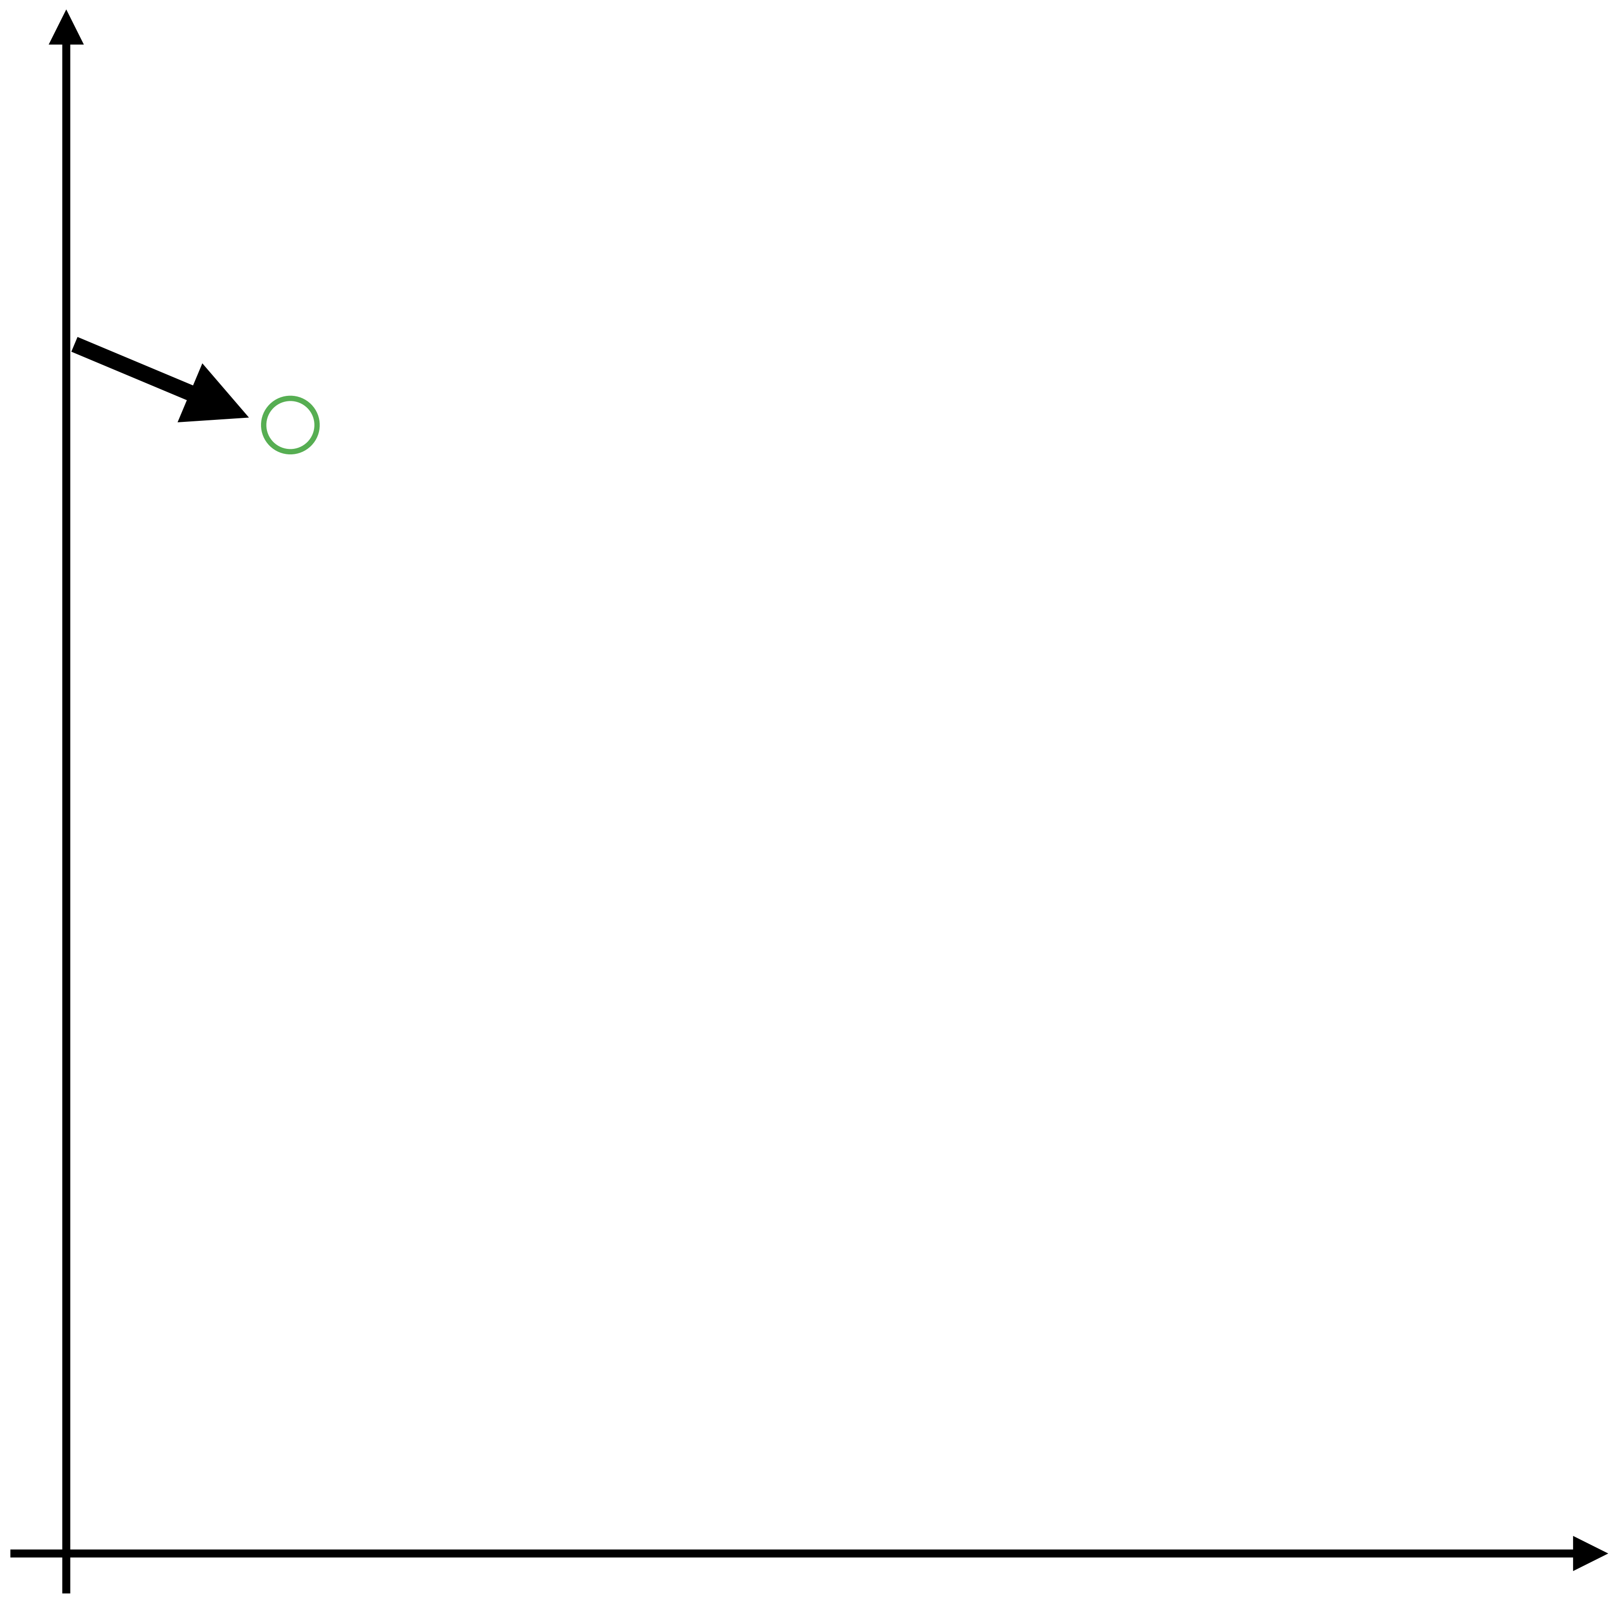

Science AAAS

Reducing the Dimensionality of Data with Neural Networks
G. E. Hinton, *et al.*
Science **313**, 504 (2006);
DOI: 10.1126/science.1127647

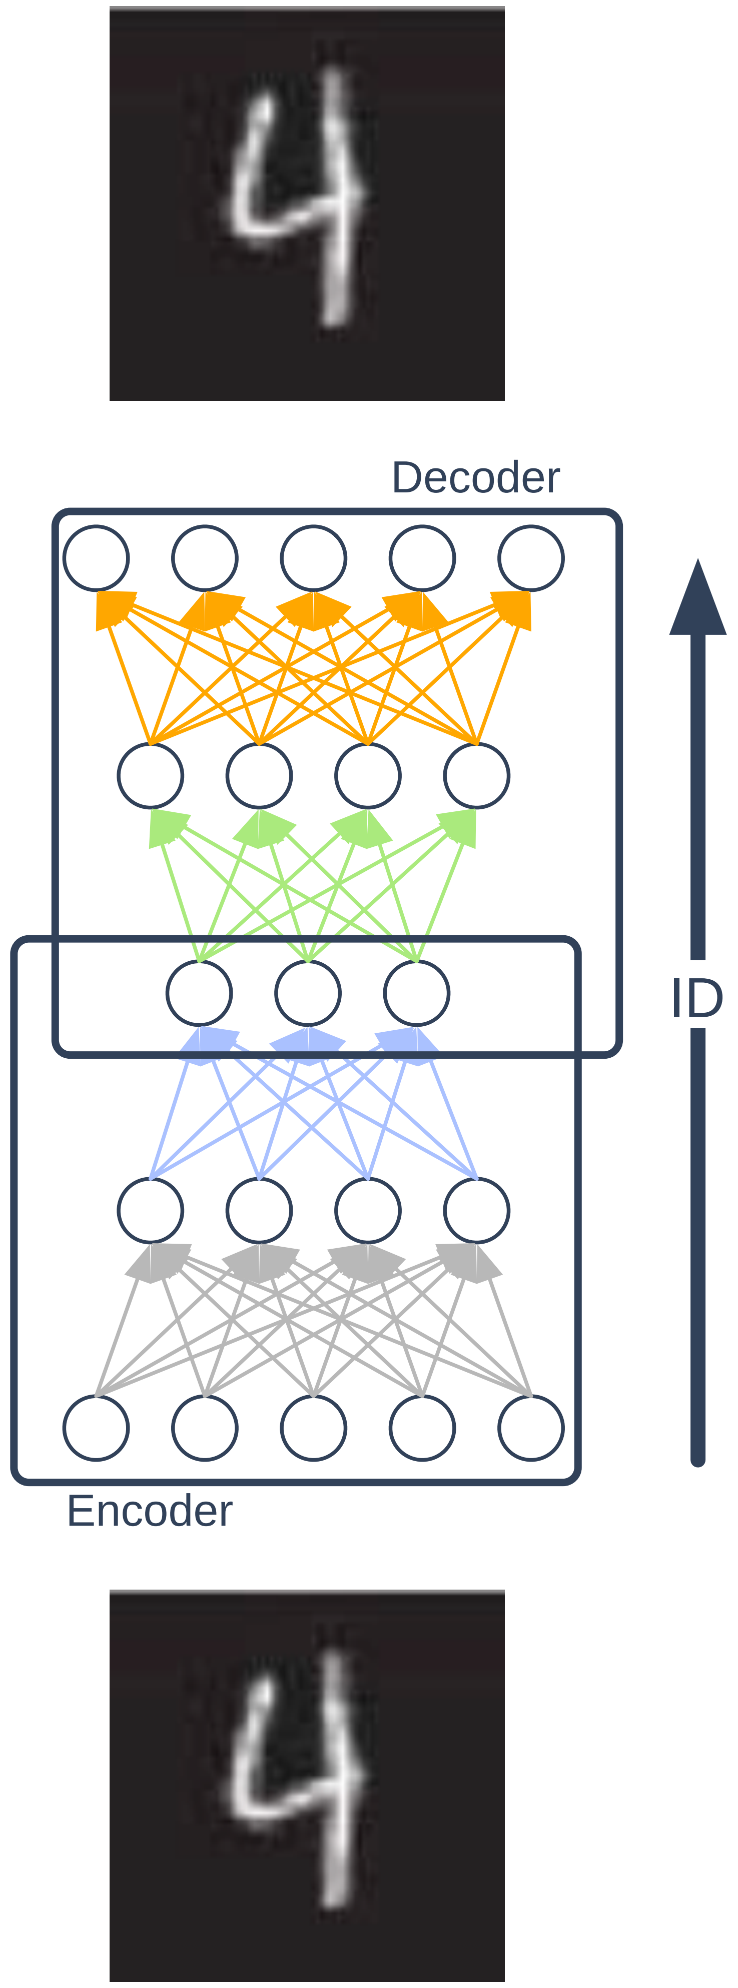

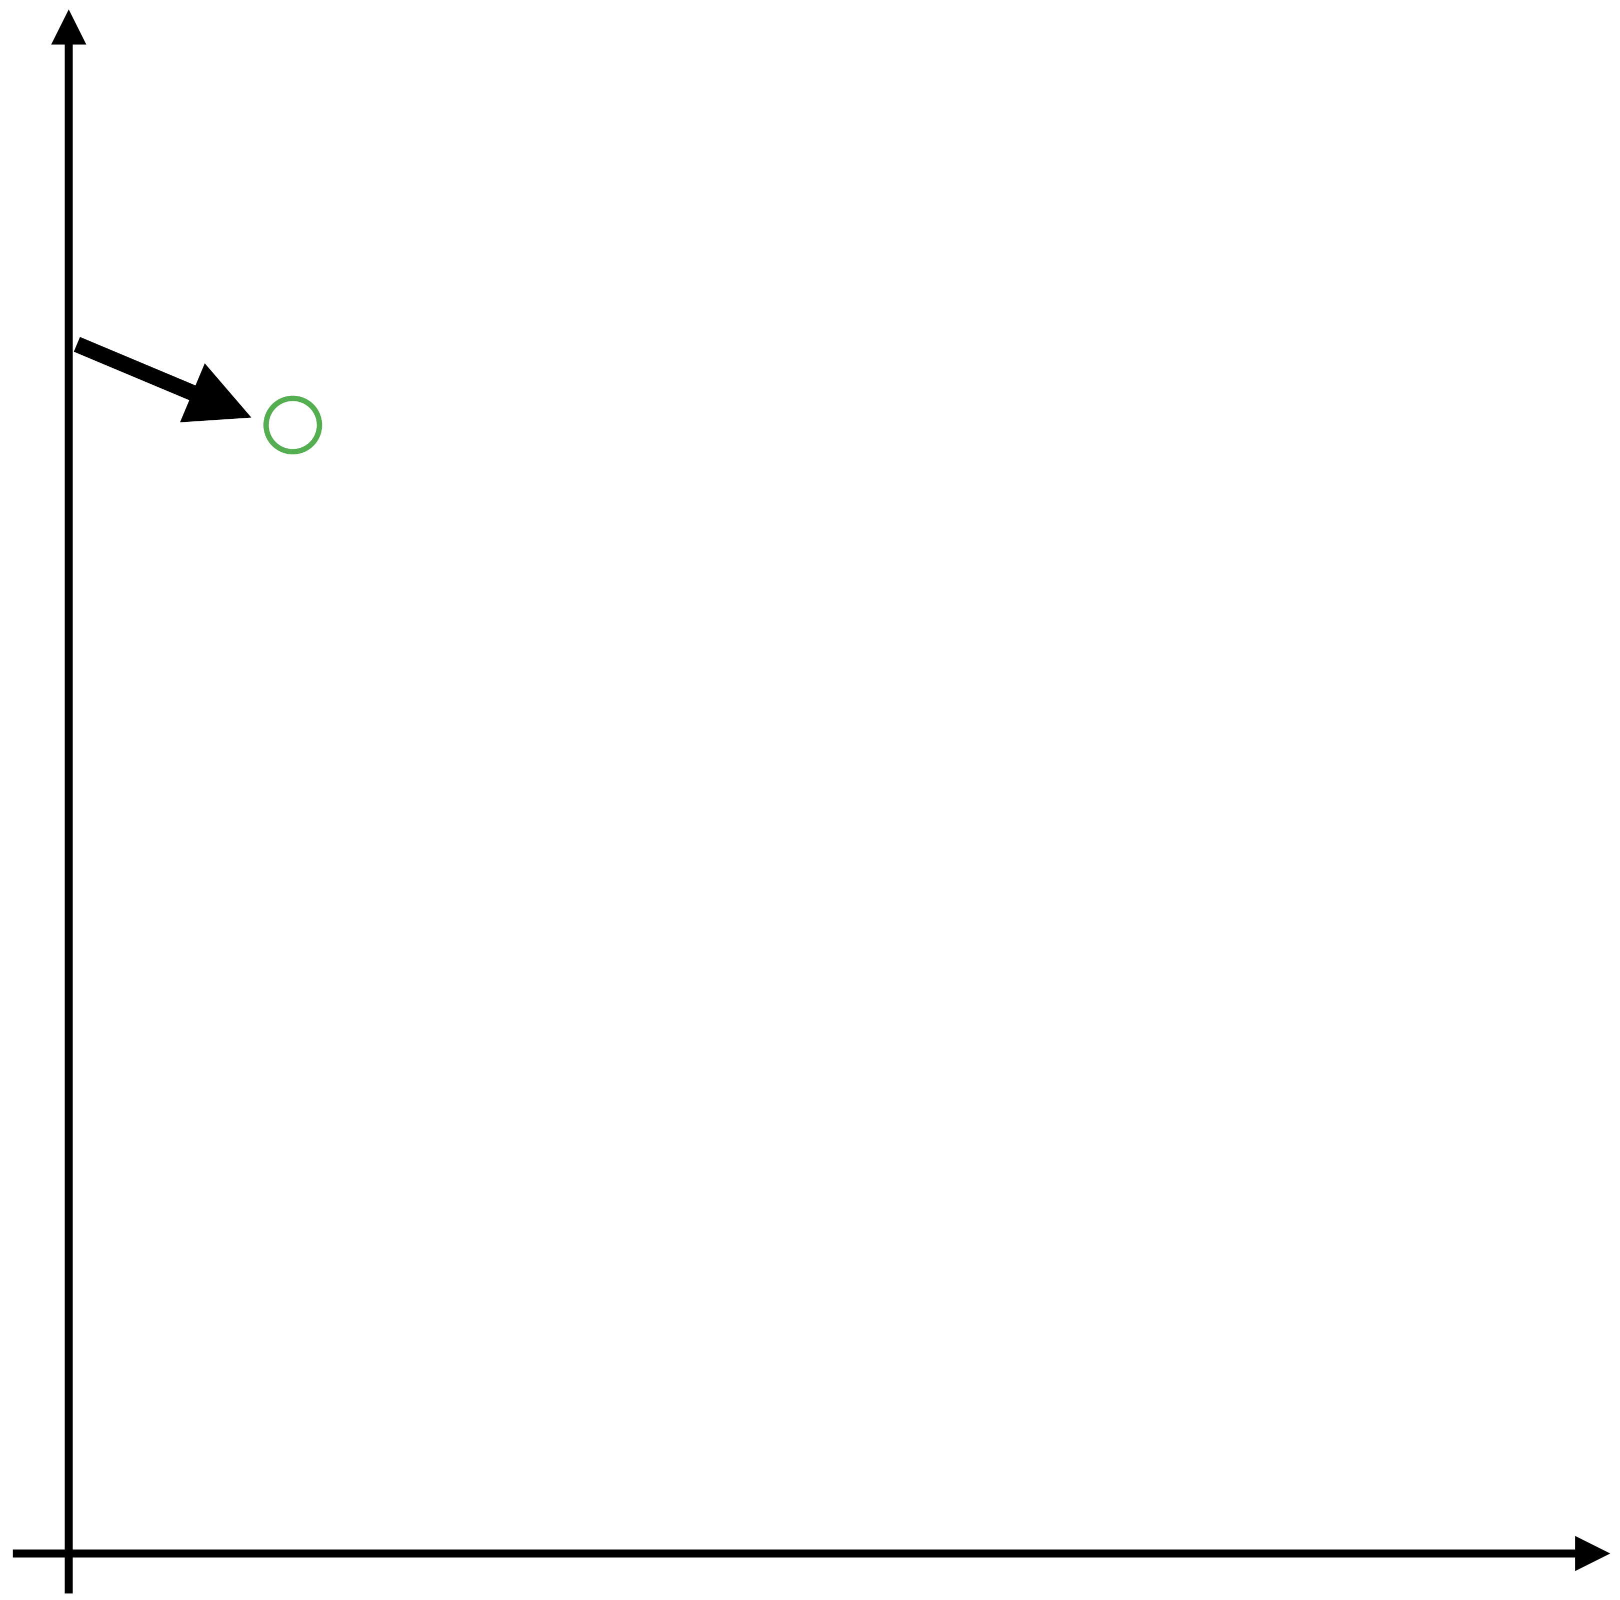

Reducing the Dimensionality of Data with Neural Networks
G. E. Hinton, *et al.*
Science **313**, 504 (2006);
DOI: 10.1126/science.1127647

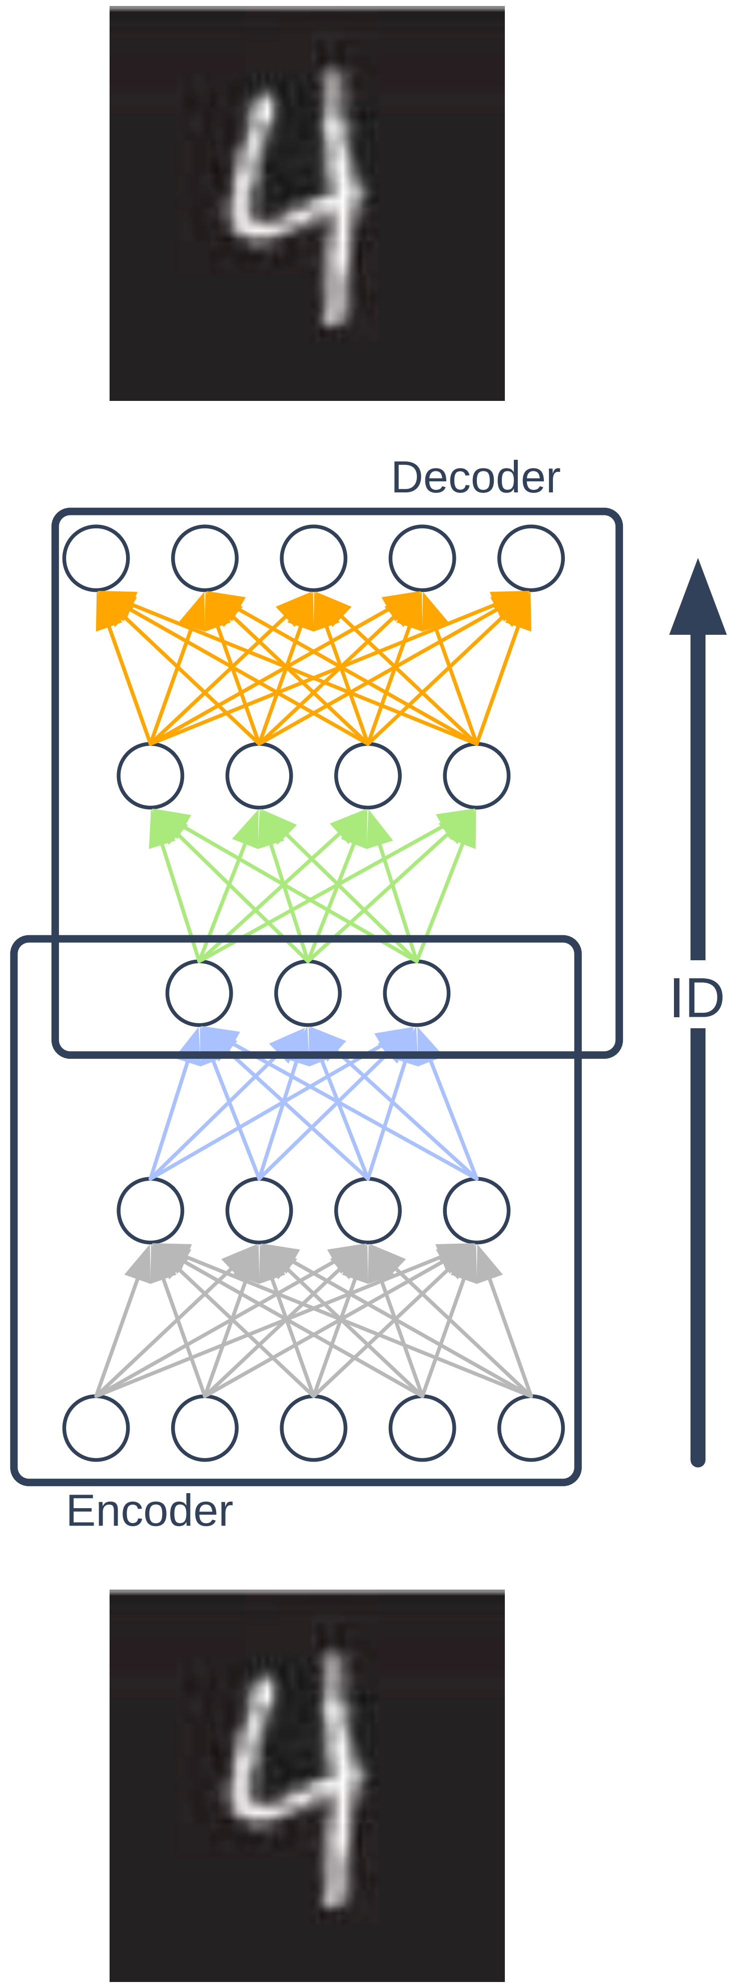

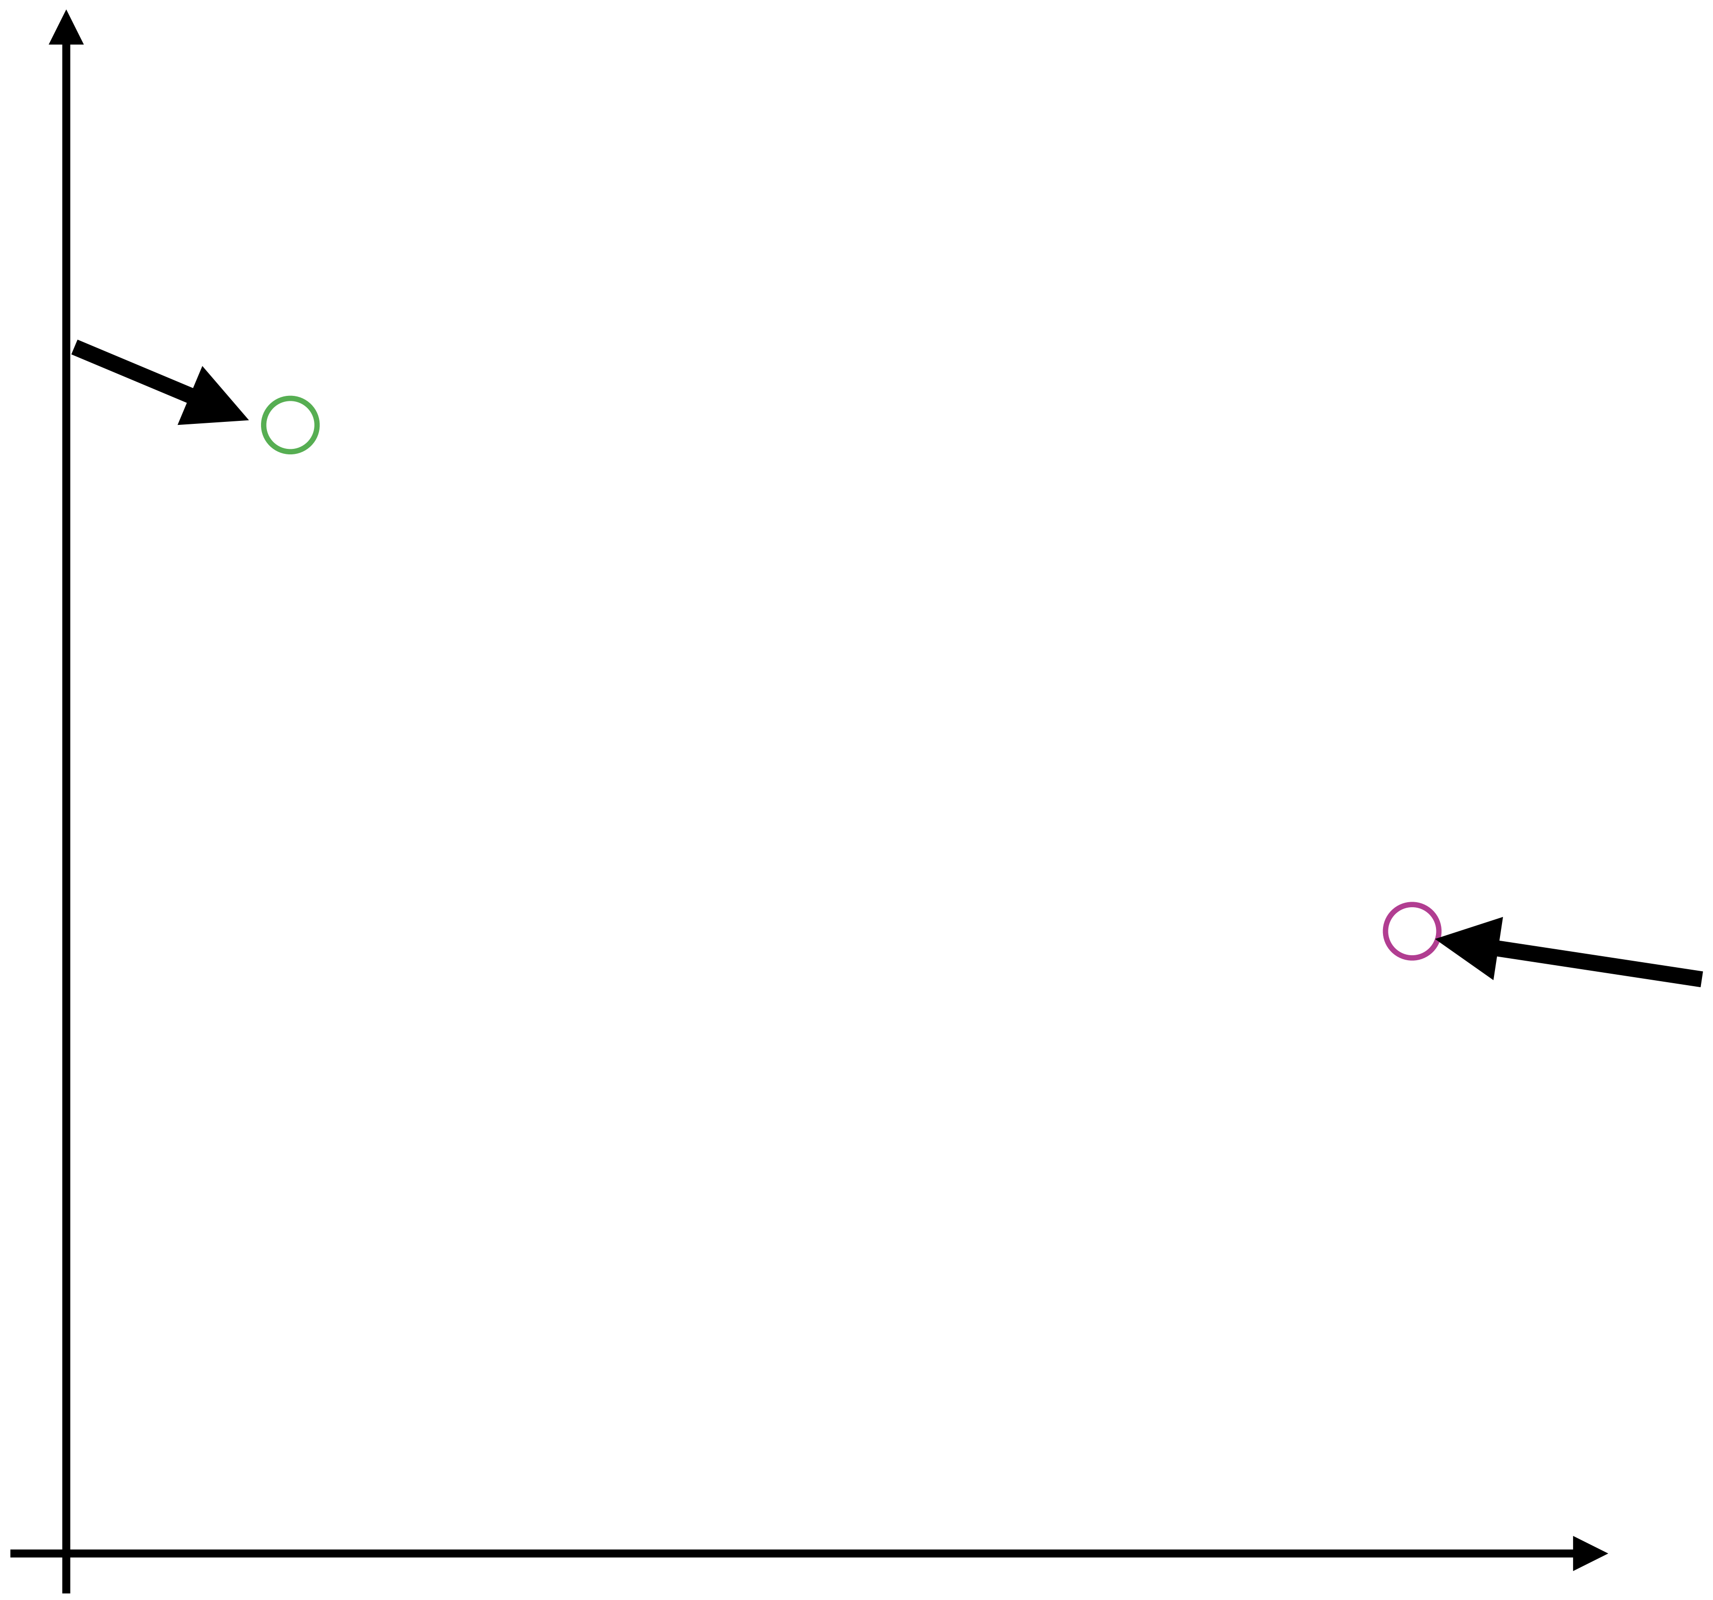

Reducing the Dimensionality of Data with Neural Networks
G. E. Hinton, *et al.*
Science **313**, 504 (2006);
DOI: 10.1126/science.1127647

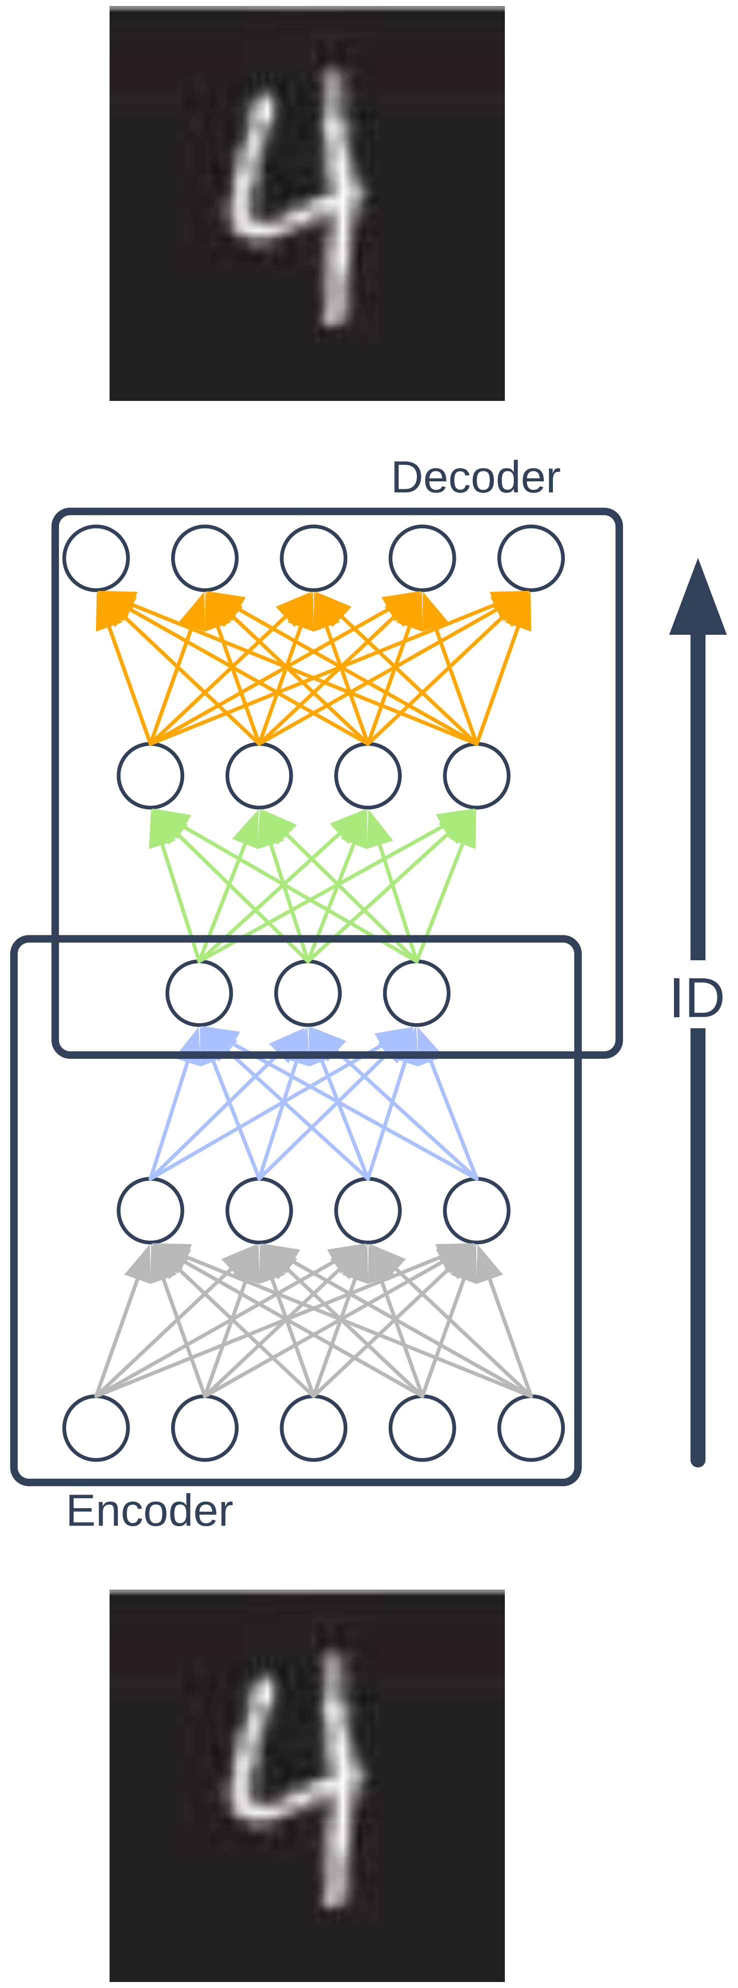

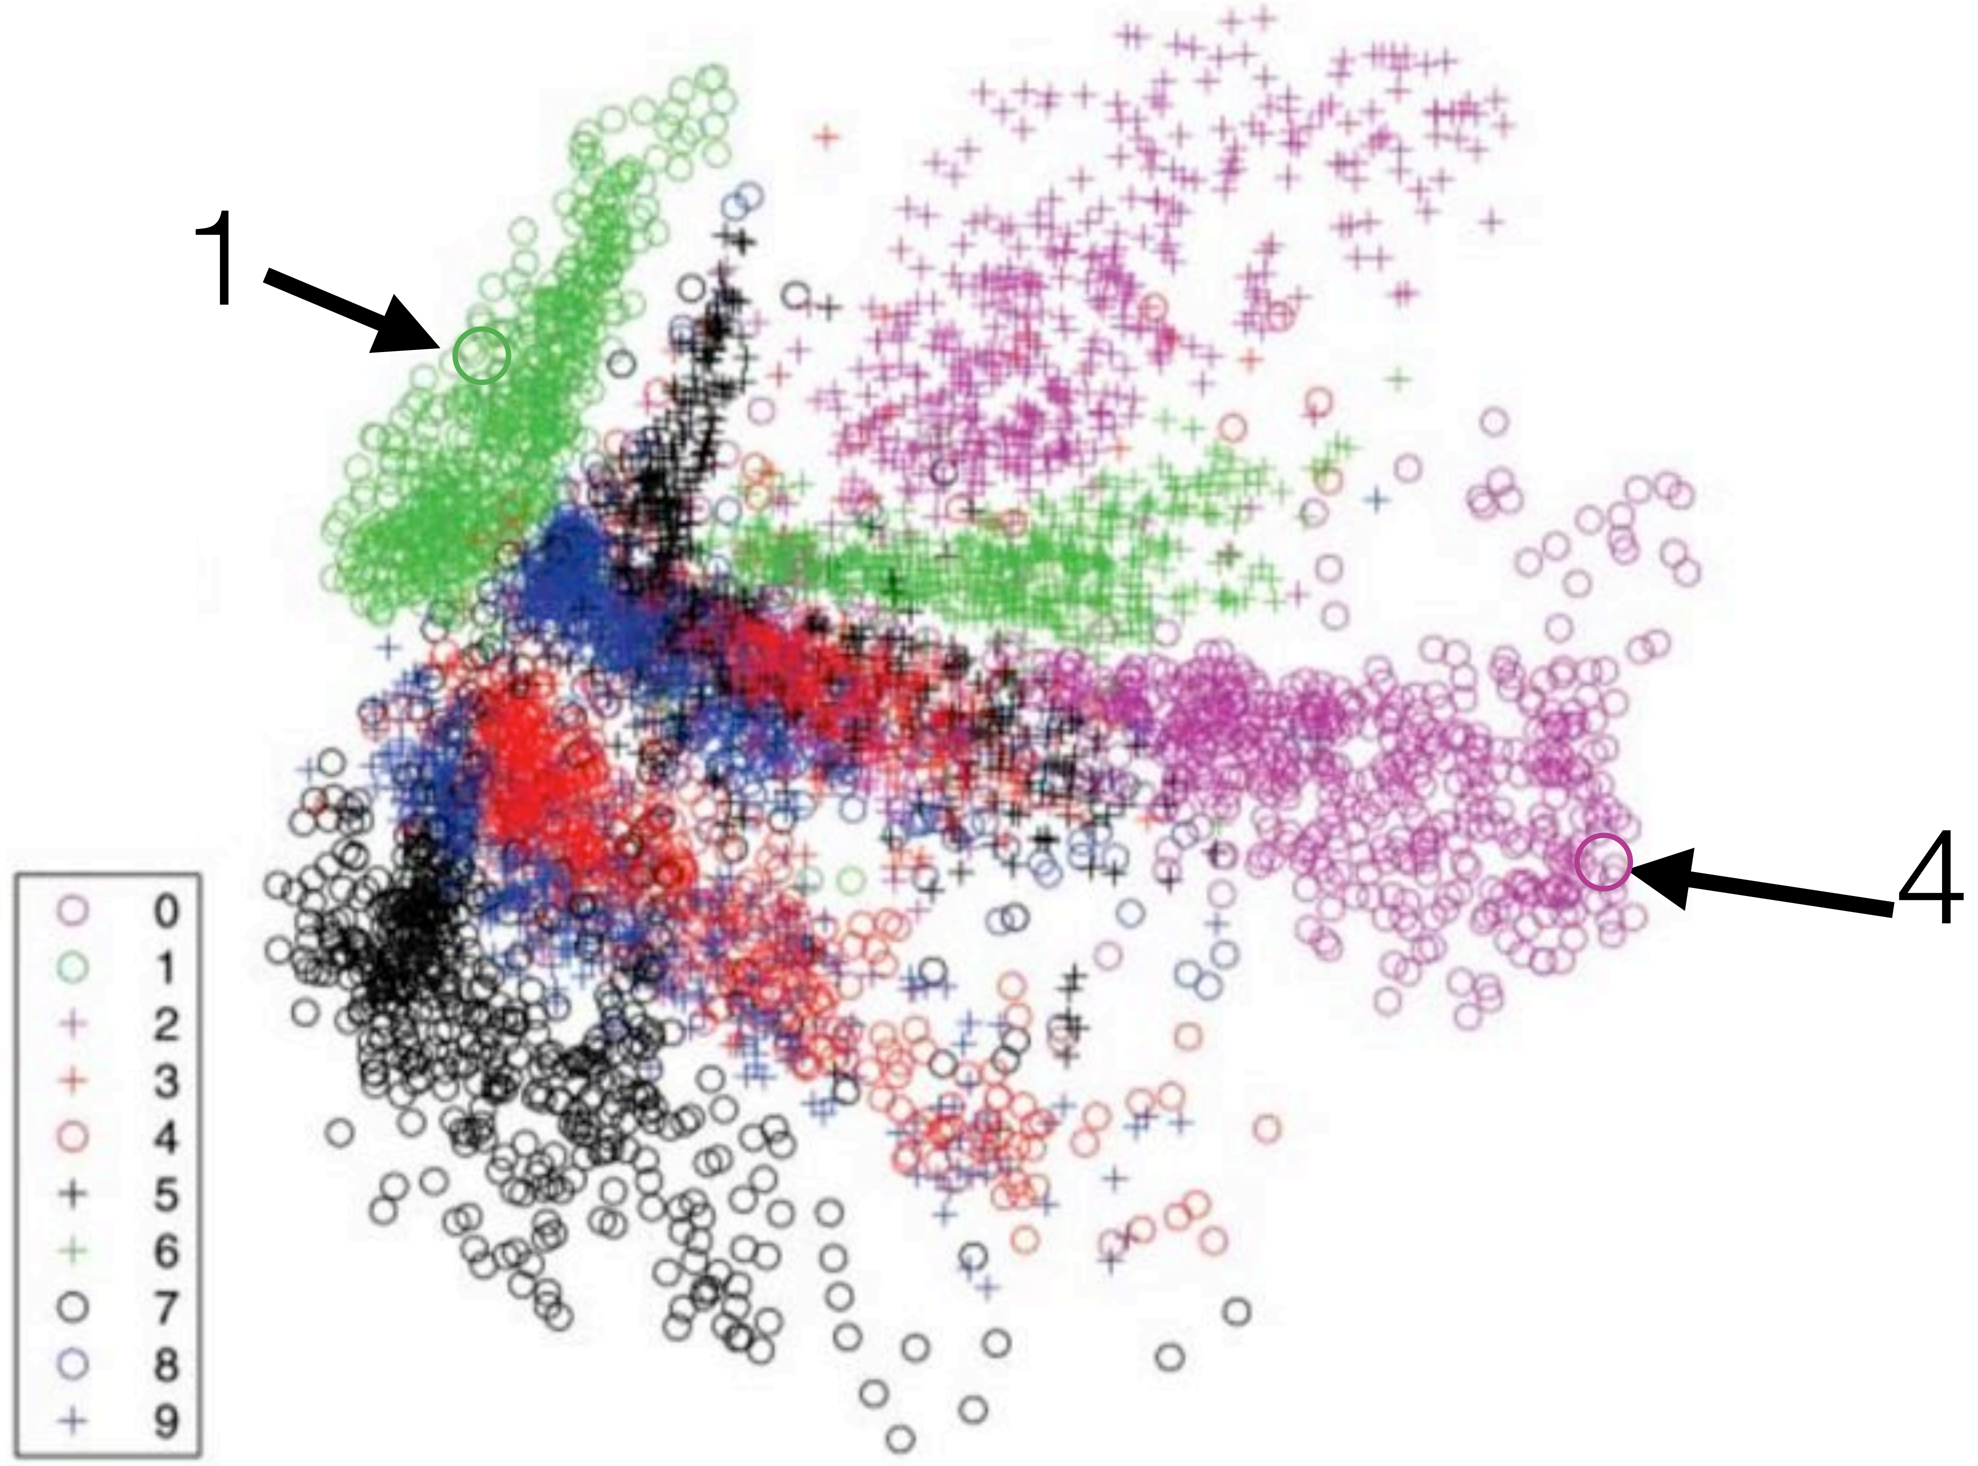

Science AAAS

Reducing the Dimensionality of Data with Neural Networks
G. E. Hinton, *et al.*
Science **313**, 504 (2006);
DOI: 10.1126/science.1127647

Unsupervised Learning

- There is just the data, no target, no feedback.
  Example: Probability Density Estimation, Clustering, **Mannifold Learning, Feature Learning**.

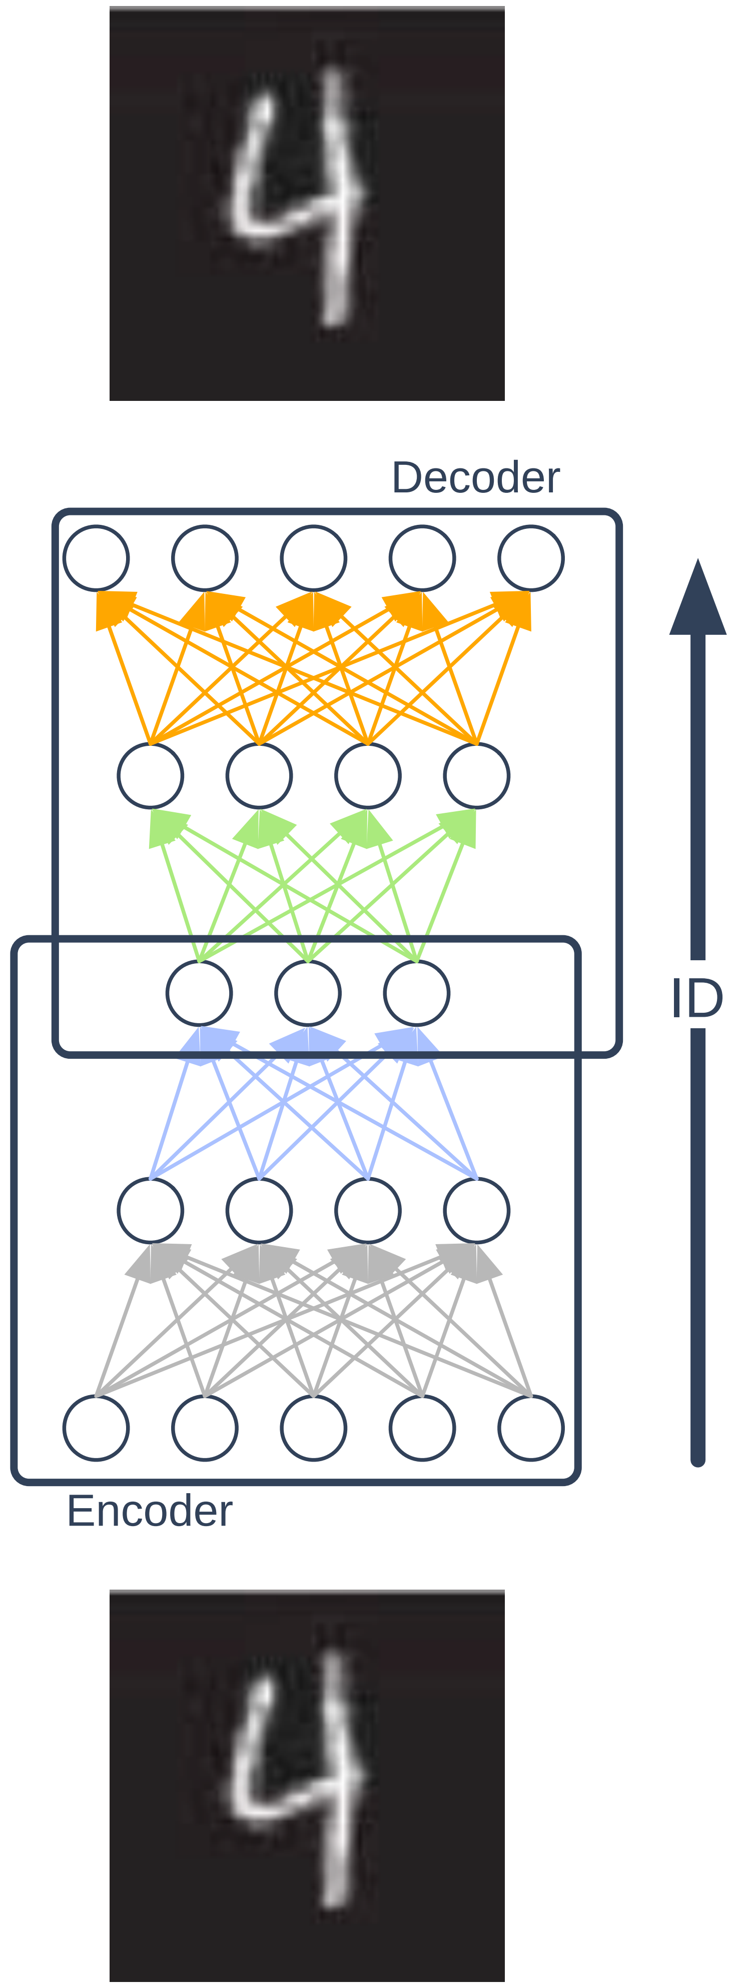

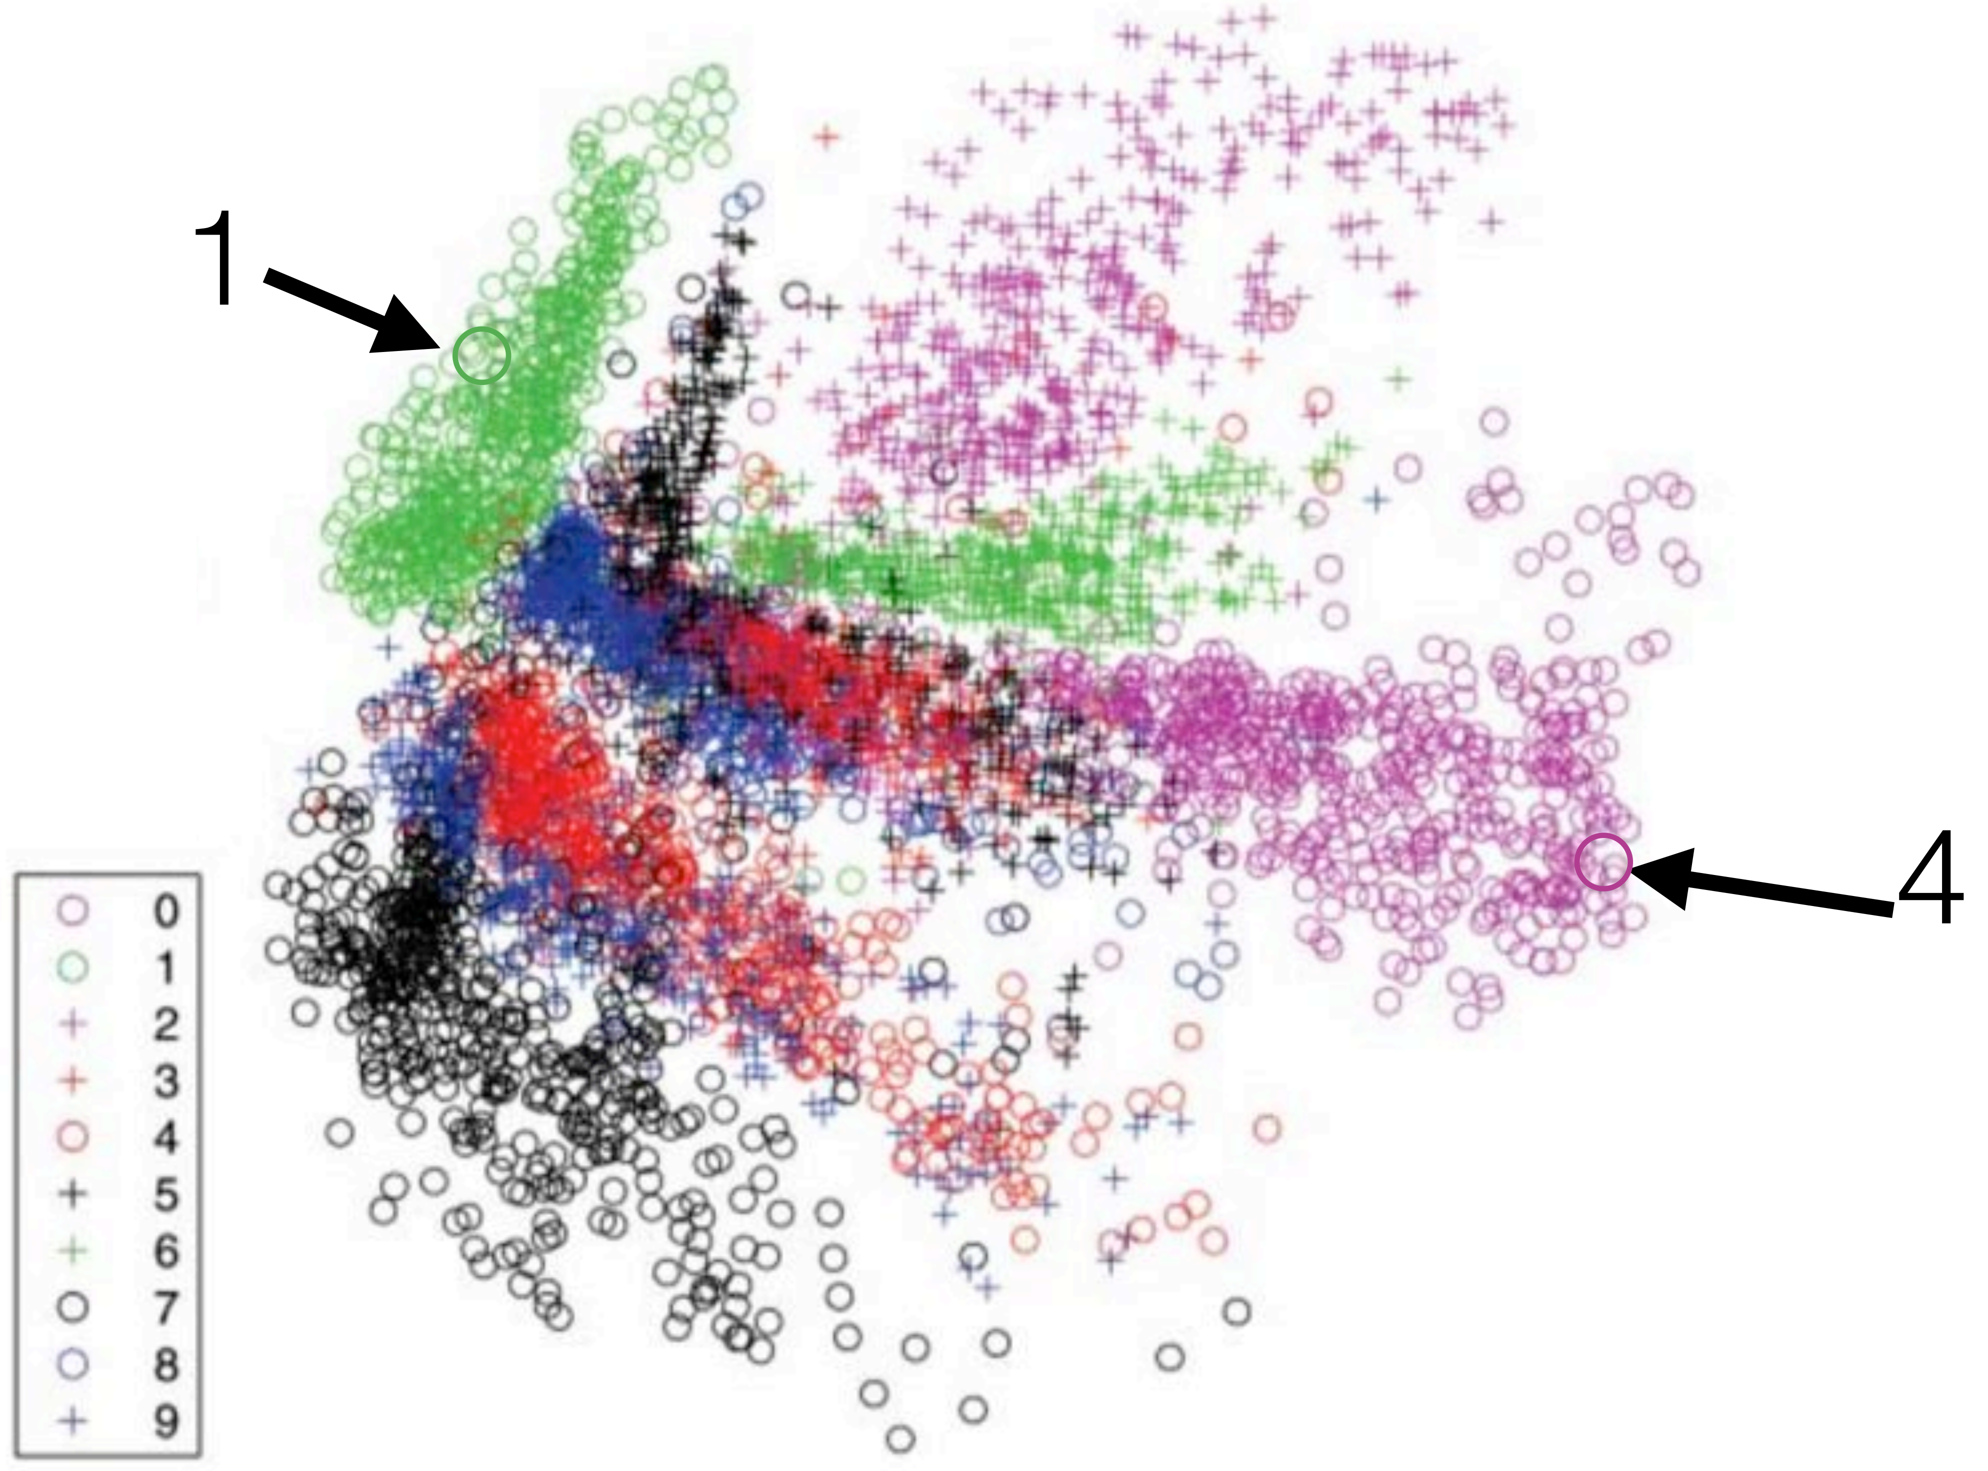

Science AAAS

Reducing the Dimensionality of Data with Neural Networks
G. E. Hinton, *et al.*
Science **313**, 504 (2006);
DOI: 10.1126/science.1127647

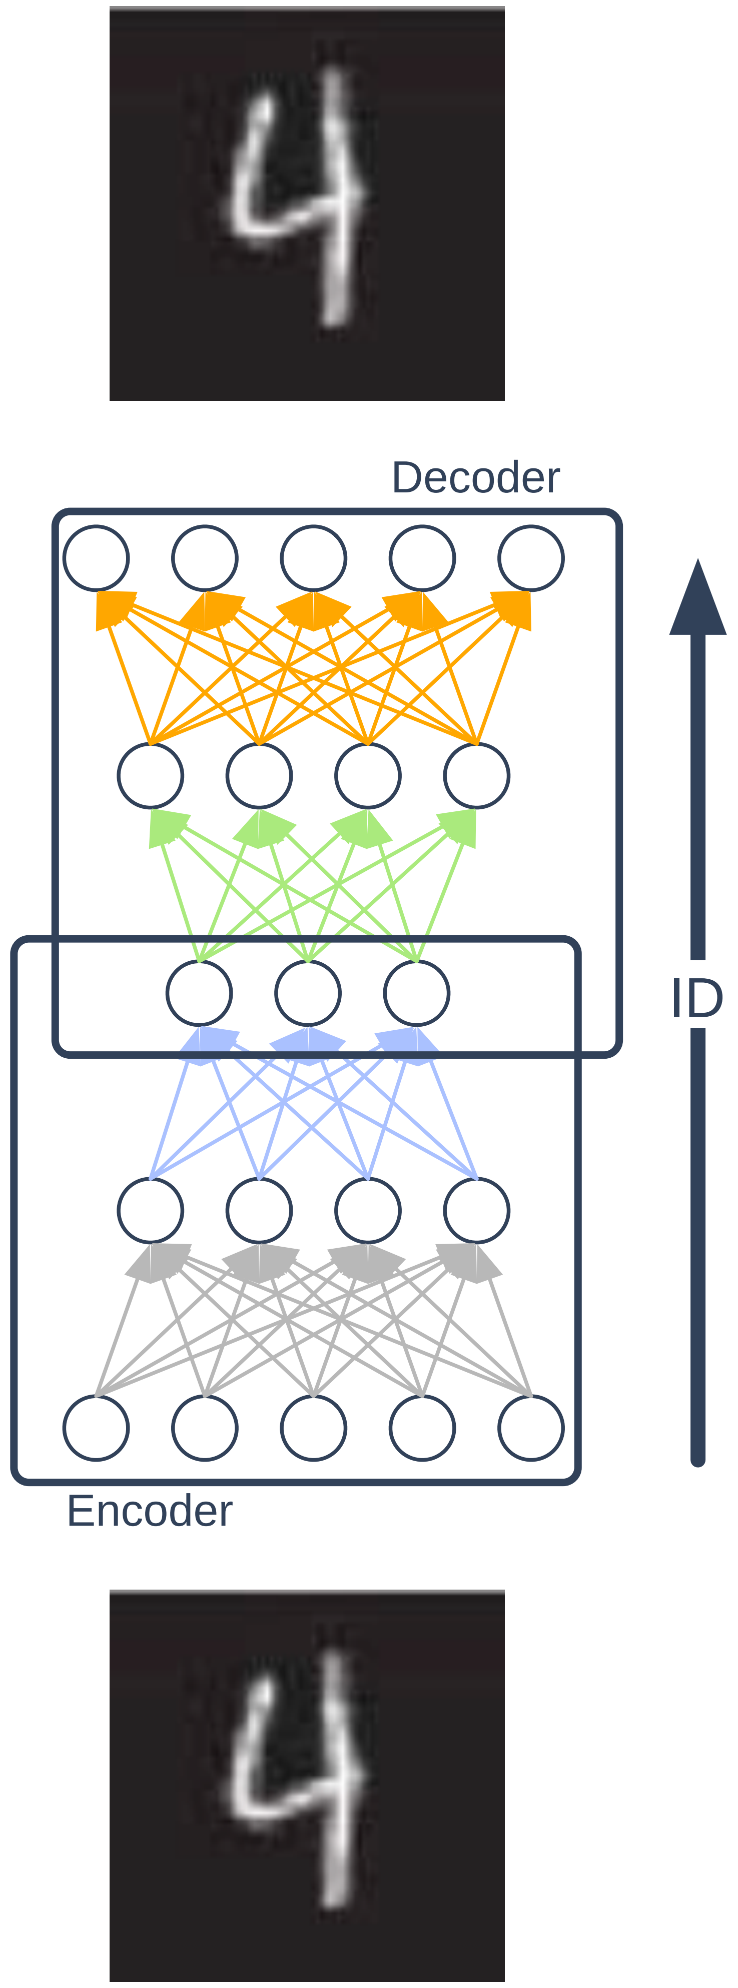

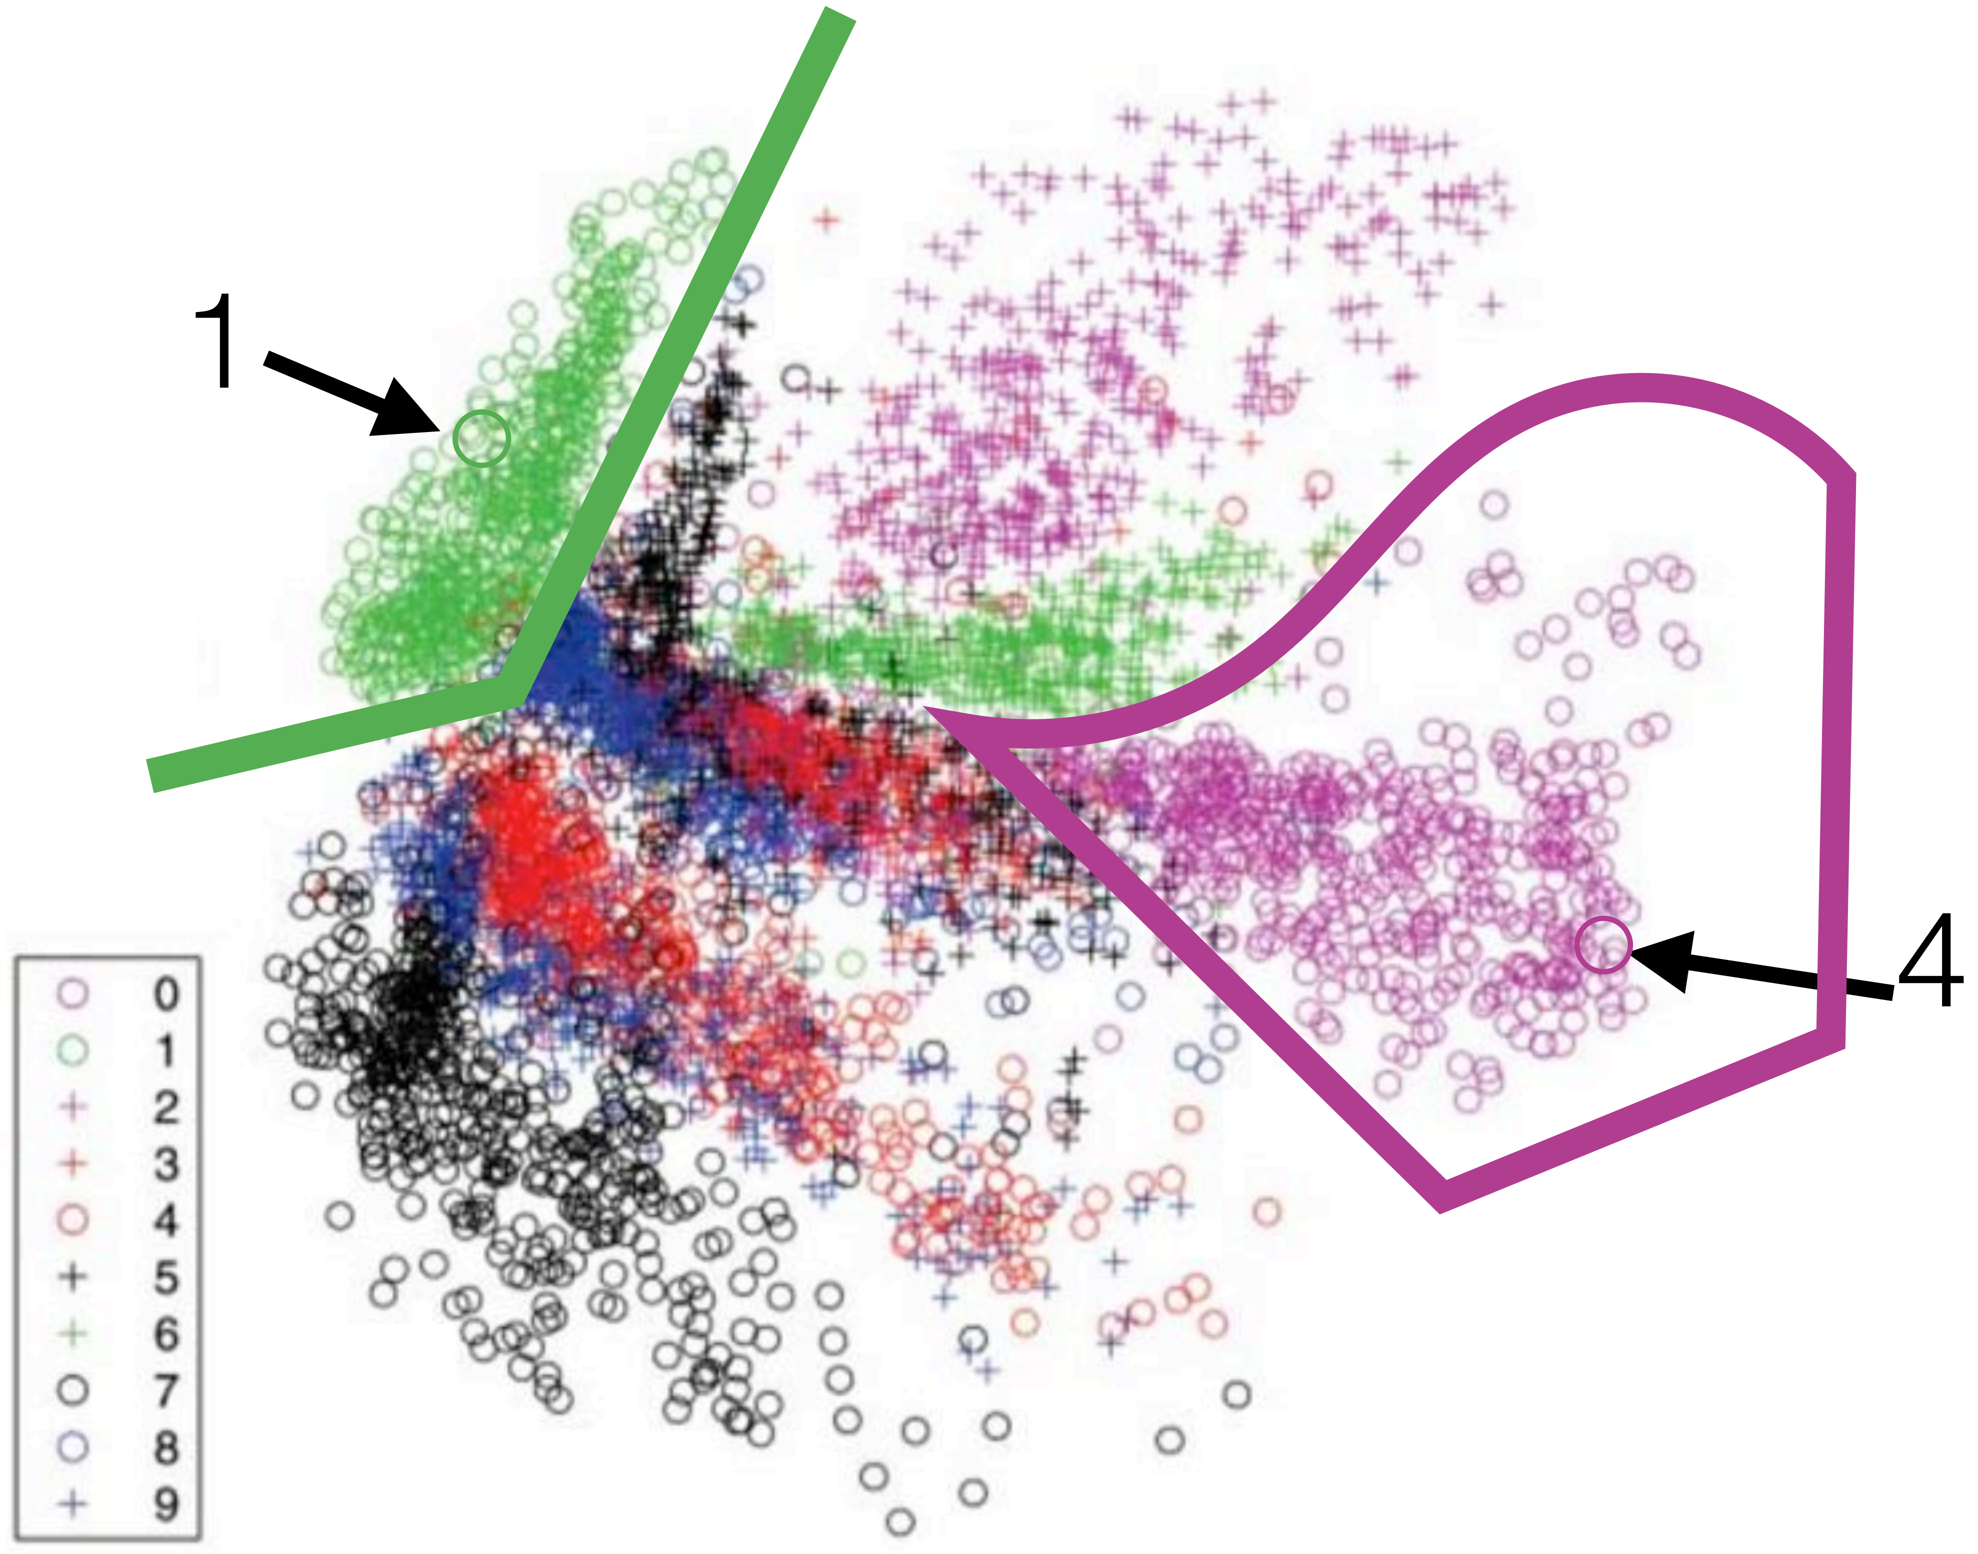

Science AAAS

Reducing the Dimensionality of Data with Neural Networks
G. E. Hinton, *et al.*
Science **313**, 504 (2006);
DOI: 10.1126/science.1127647

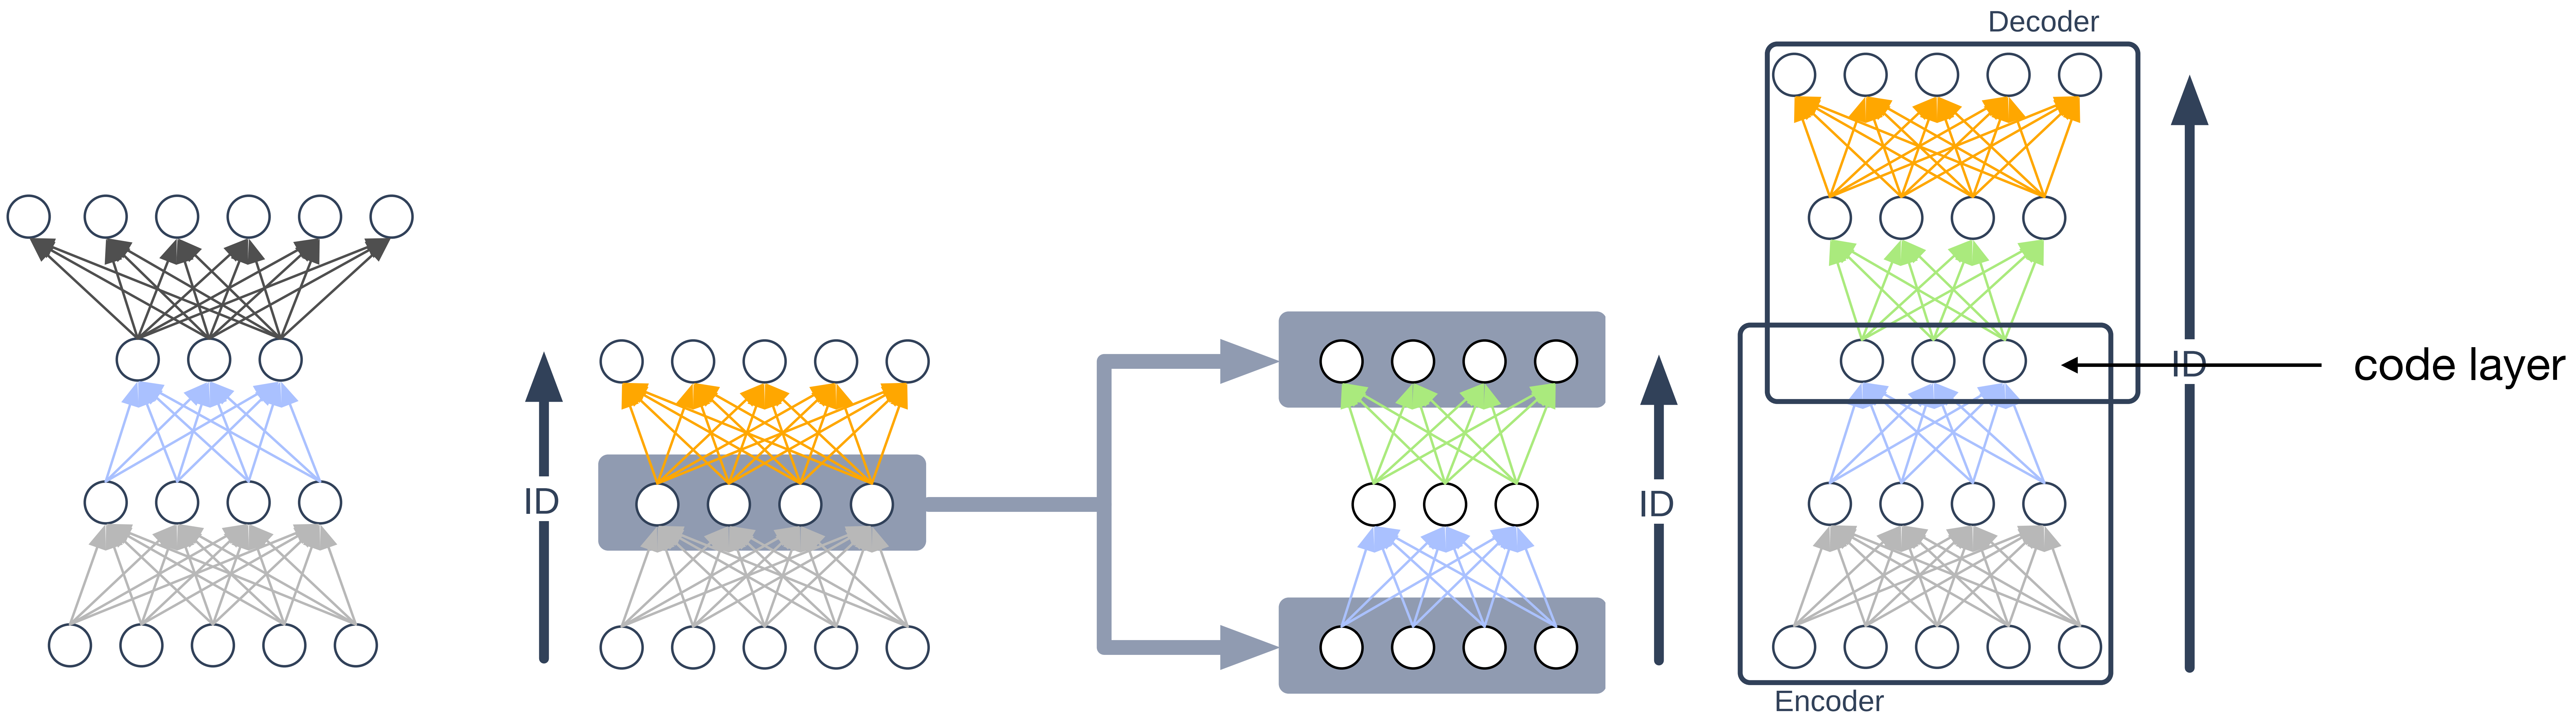
ID
ID
Decoder
ID
code layer
Encoder

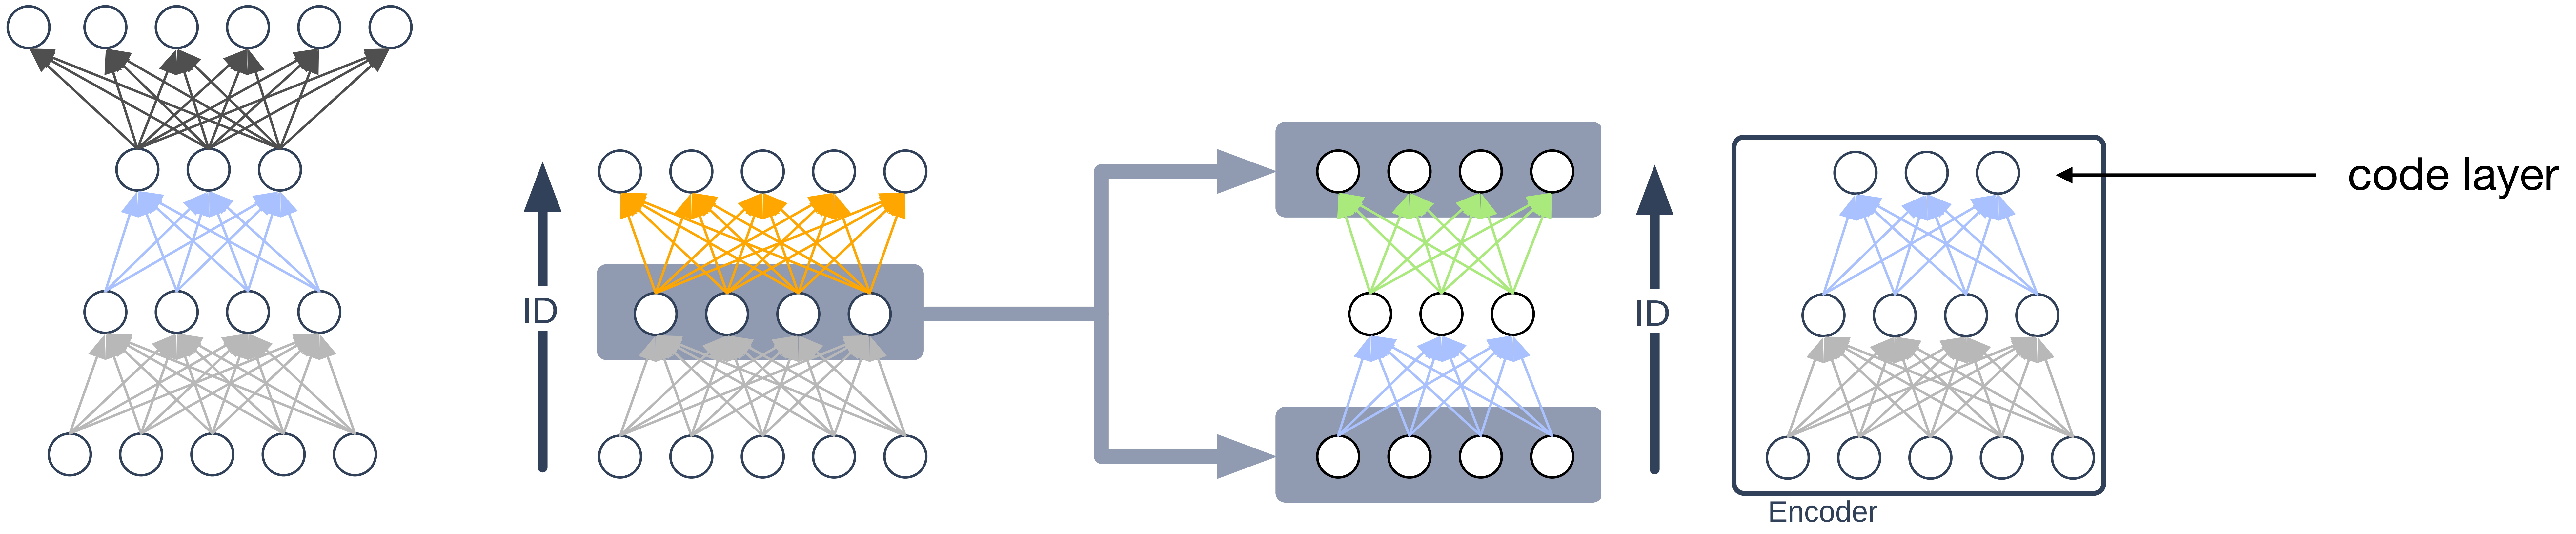

ID
ID
code layer
Encoder

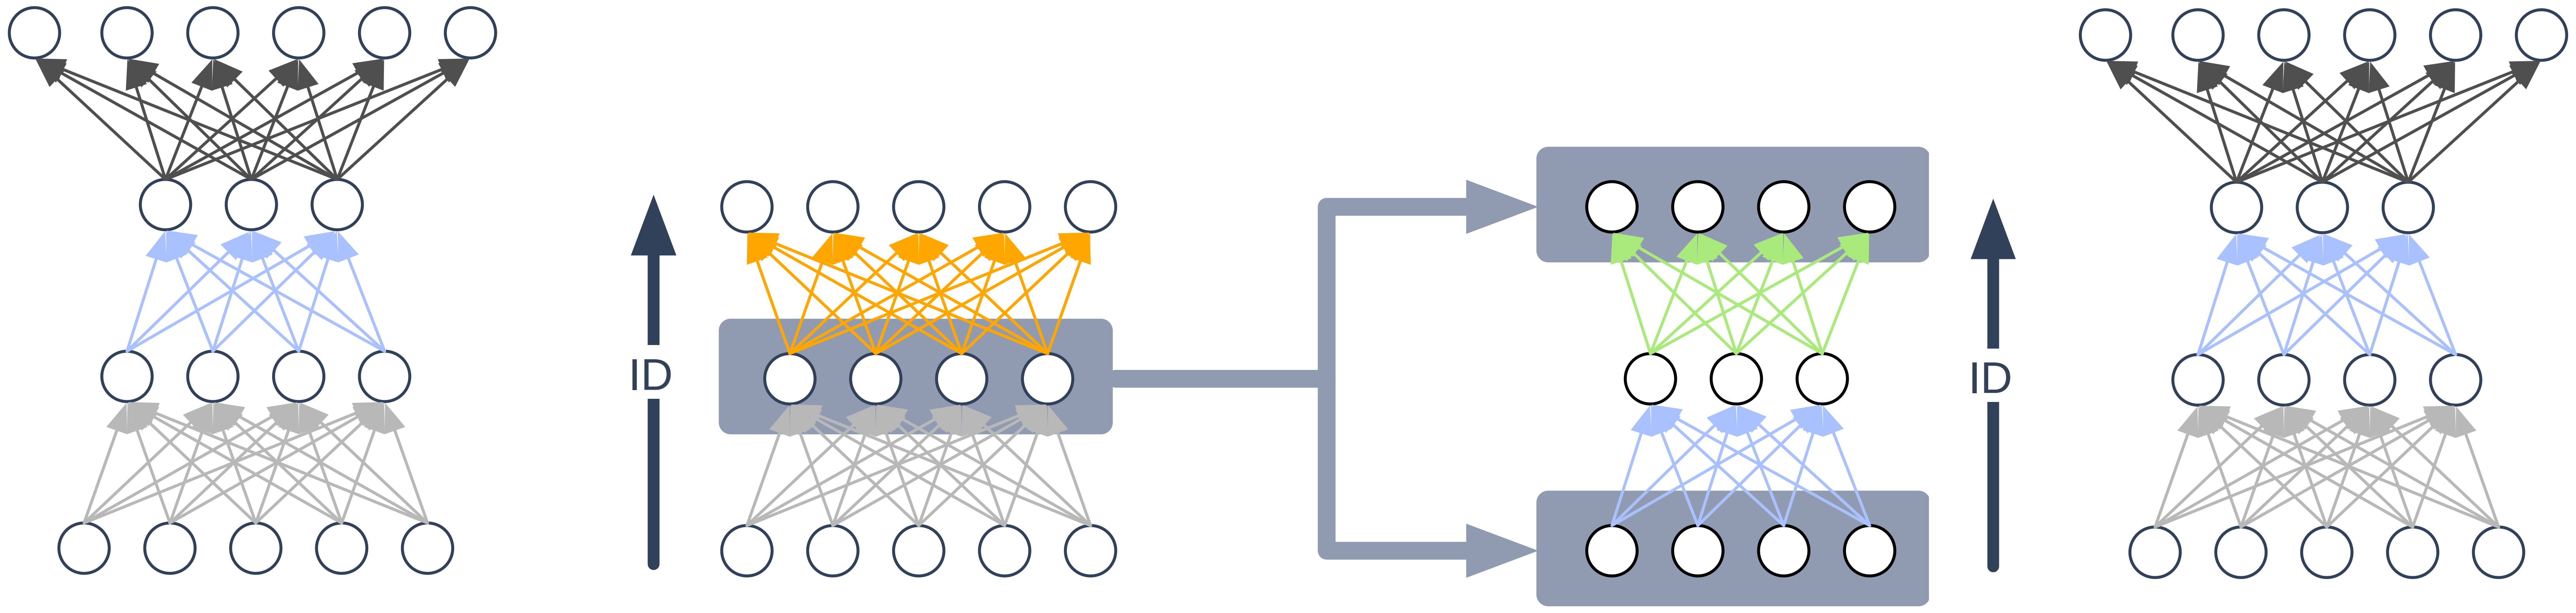
ID
ID

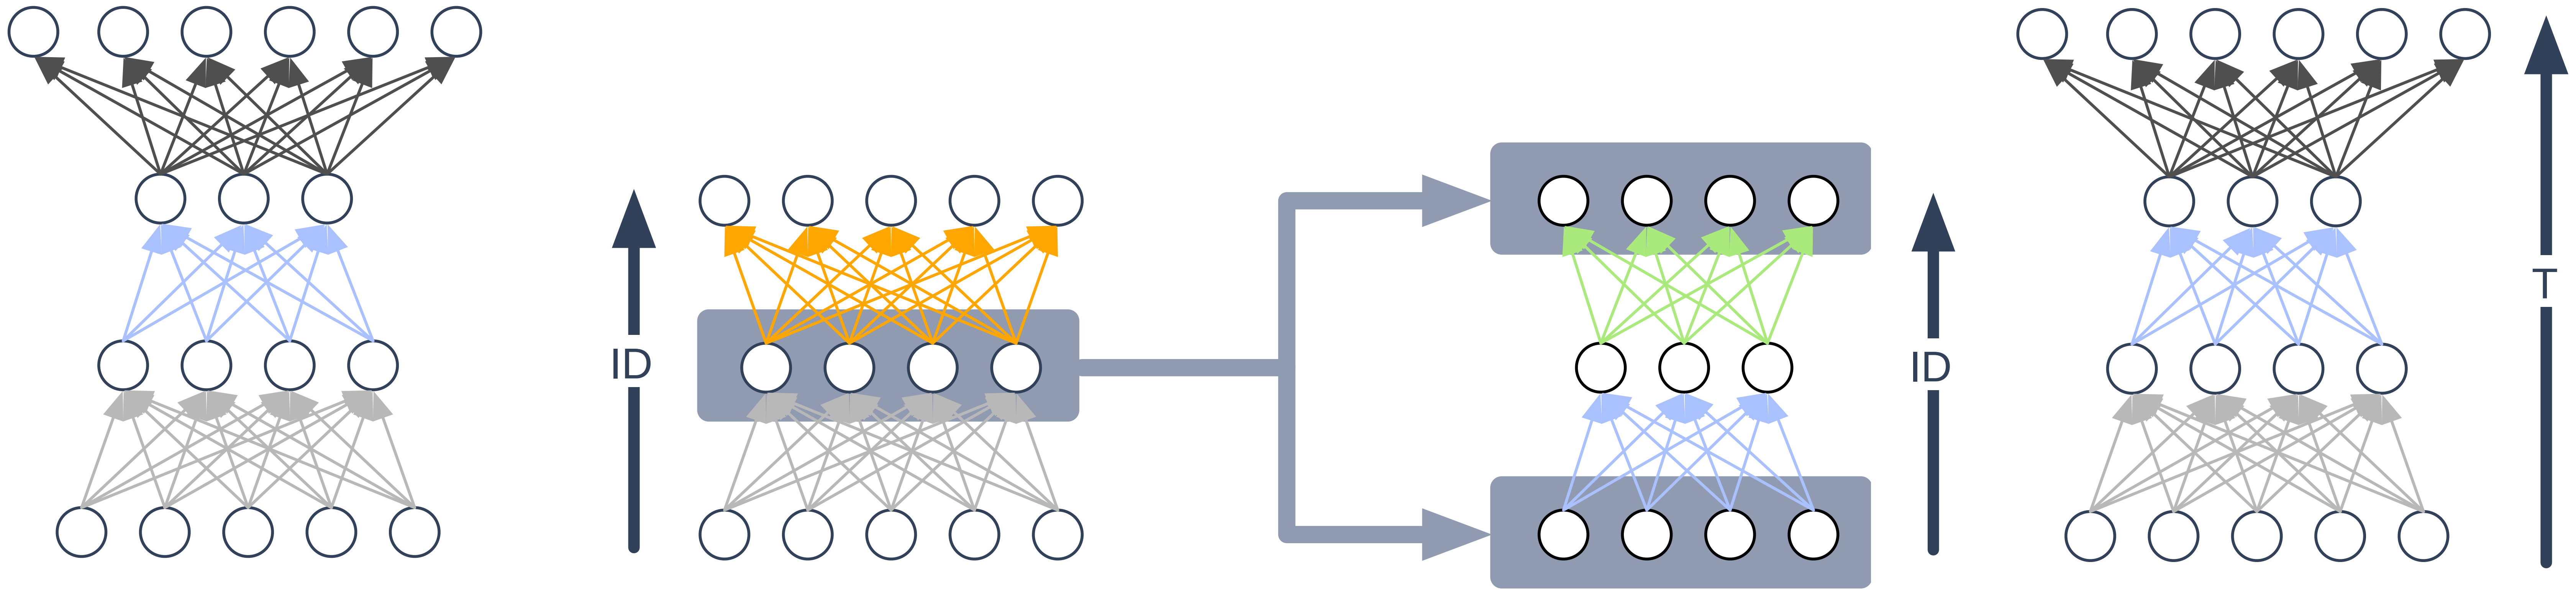
ID
ID
T

THE MNIST DATABASE

of handwritten digits

Yann LeCun, Courant Institute, NYU
Corinna Cortes, Google Labs, New York
Christopher J.C. Burges, Microsoft Research, Redmond

Neural Nets			
2-layer NN, 300 hidden units, mean square error	none	4.7	LeCun et al. 1998
2-layer NN, 300 HU, MSE, [distortions]	none	3.6	LeCun et al. 1998
2-layer NN, 300 HU	deskewing	1.6	LeCun et al. 1998
2-layer NN, 1000 hidden units	none	4.5	LeCun et al. 1998
2-layer NN, 1000 HU, [distortions]	none	3.8	LeCun et al. 1998
3-layer NN, 300+100 hidden units	none	3.05	LeCun et al. 1998
3-layer NN, 300+100 HU [distortions]	none	2.5	LeCun et al. 1998
3-layer NN, 500+150 hidden units	none	2.95	LeCun et al. 1998
3-layer NN, 500+150 HU [distortions]	none	2.45	LeCun et al. 1998
3-layer NN, 500+300 HU, softmax, cross entropy, weight decay	none	1.53	Hinton, unpublished, 2005
2-layer NN, 800 HU, Cross-Entropy Loss	none	1.6	Simard et al., ICDAR 2003
2-layer NN, 800 HU, cross-entropy [affine distortions]	none	1.1	Simard et al., ICDAR 2003
2-layer NN, 800 HU, MSE [elastic distortions]	none	0.9	Simard et al., ICDAR 2003
2-layer NN, 800 HU, cross-entropy [elastic distortions]	none	0.7	Simard et al., ICDAR 2003
NN, 784-500-500-2000-30 + nearest neighbor, RBM + NCA training [no distortions]	none	1.0	Salakhutdinov and Hinton, AI-Stats 2007

Reducing the Dimensionality of Data with Neural Networks
G. E. Hinton, *et al.*
Science **313**, 504 (2006);
DOI: 10.1126/science.1127647

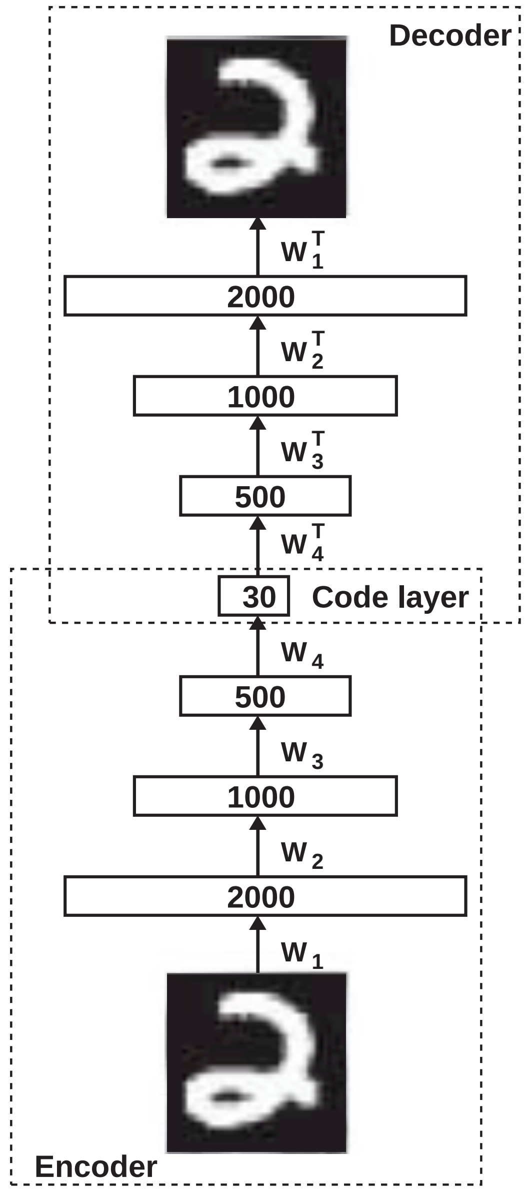

Unrolling

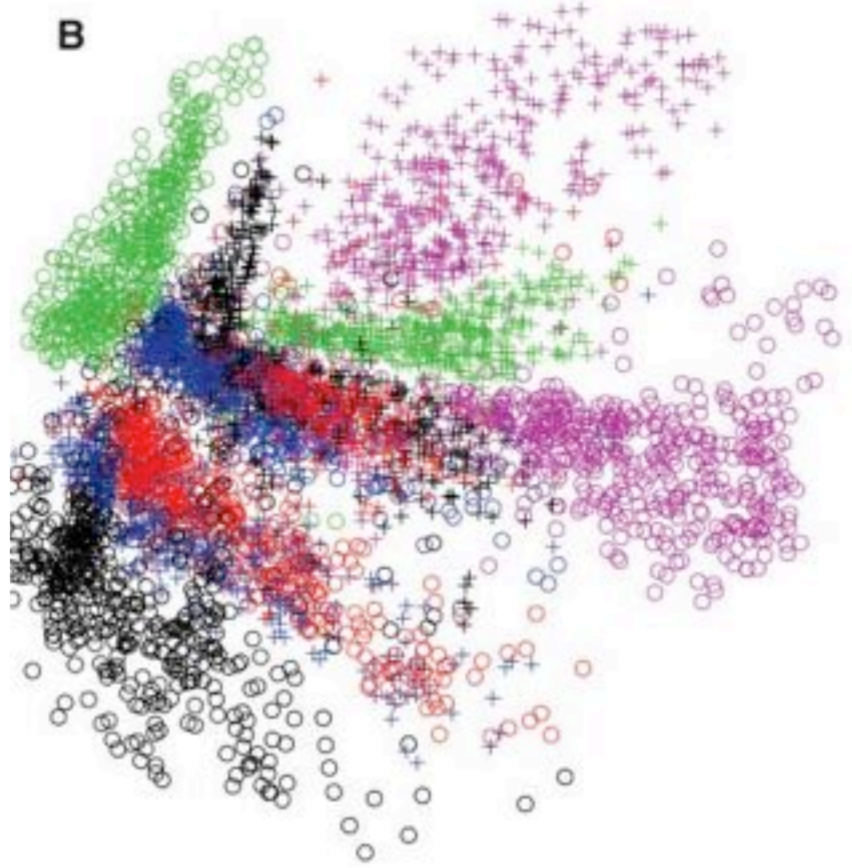

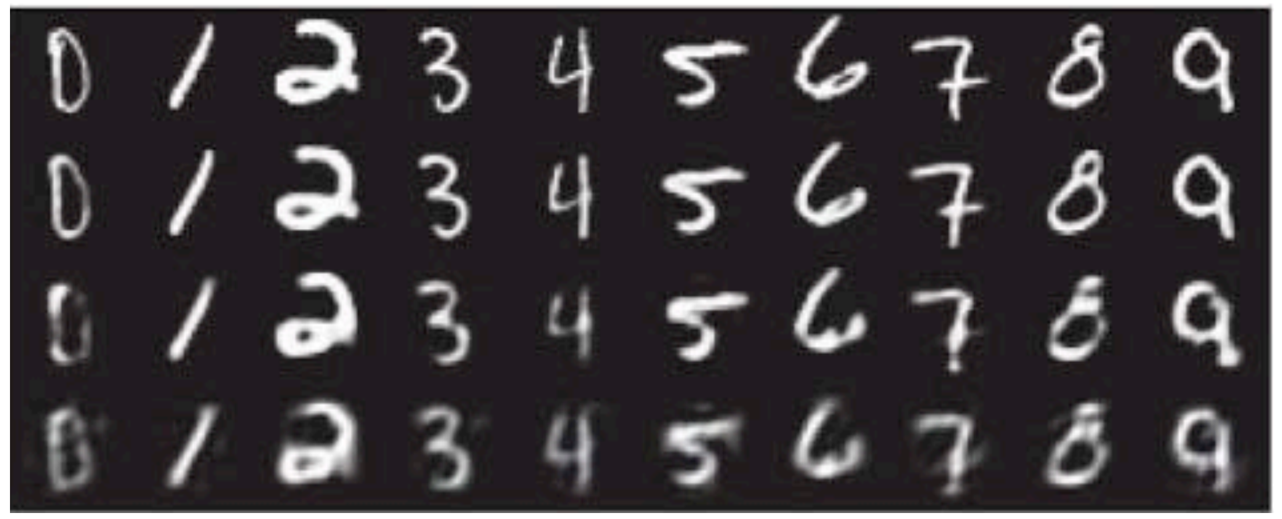

What is Deep Learning?

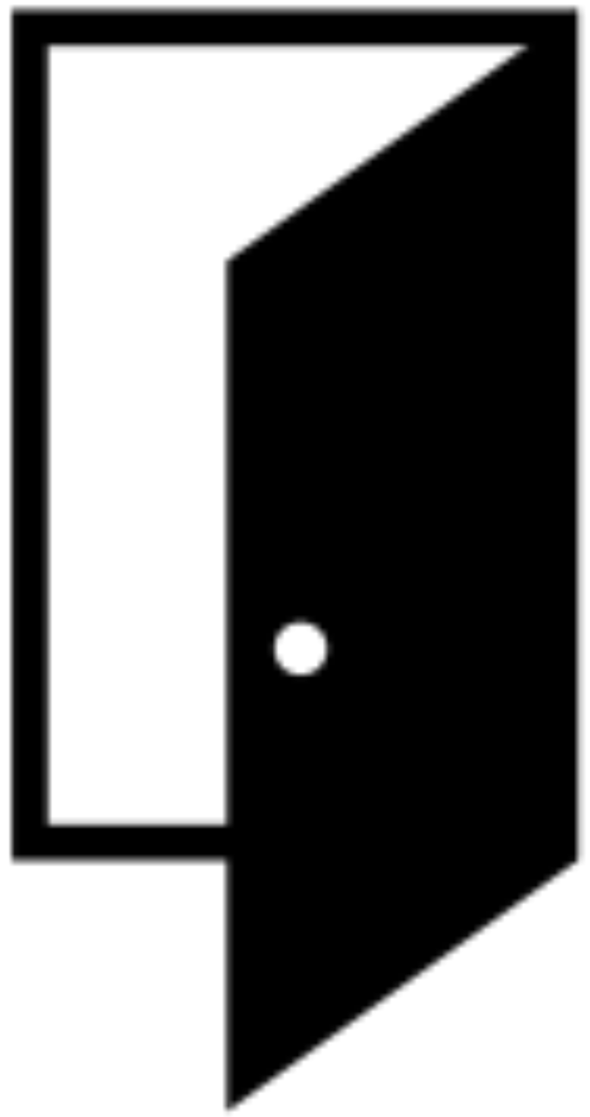

It started as clever idea for

training neural networks with
multiple hidden layers

DeepFace

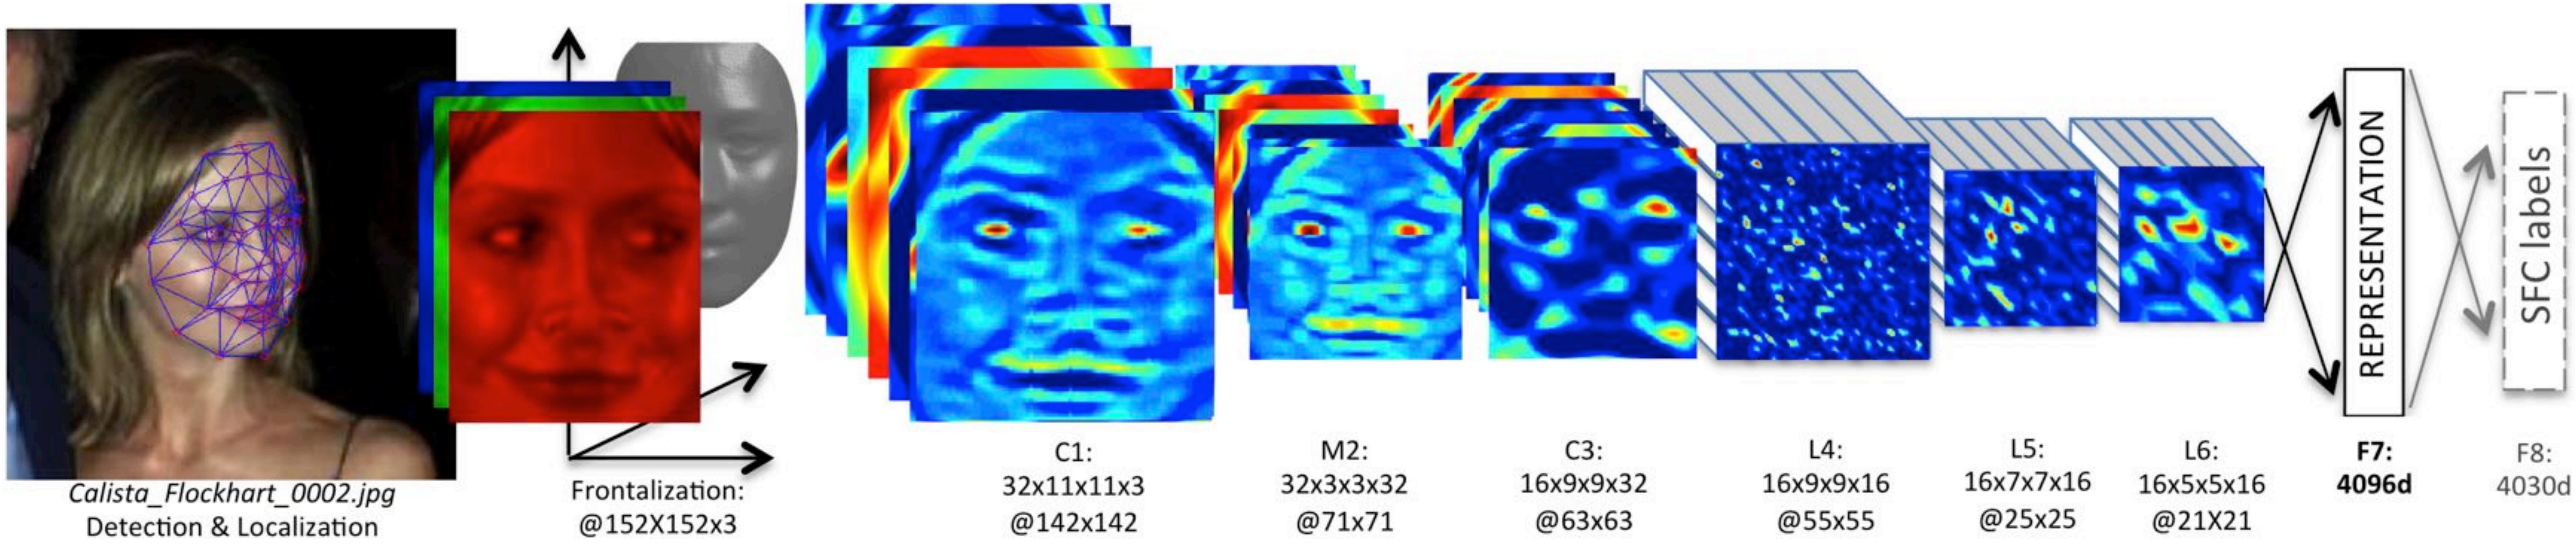

Taigman, Yang, Ranzato, Wolf, DeepFace: Closing the Gap to Human-Level Performance in Face Verification. CVPR 2014

Google LeNet 2014

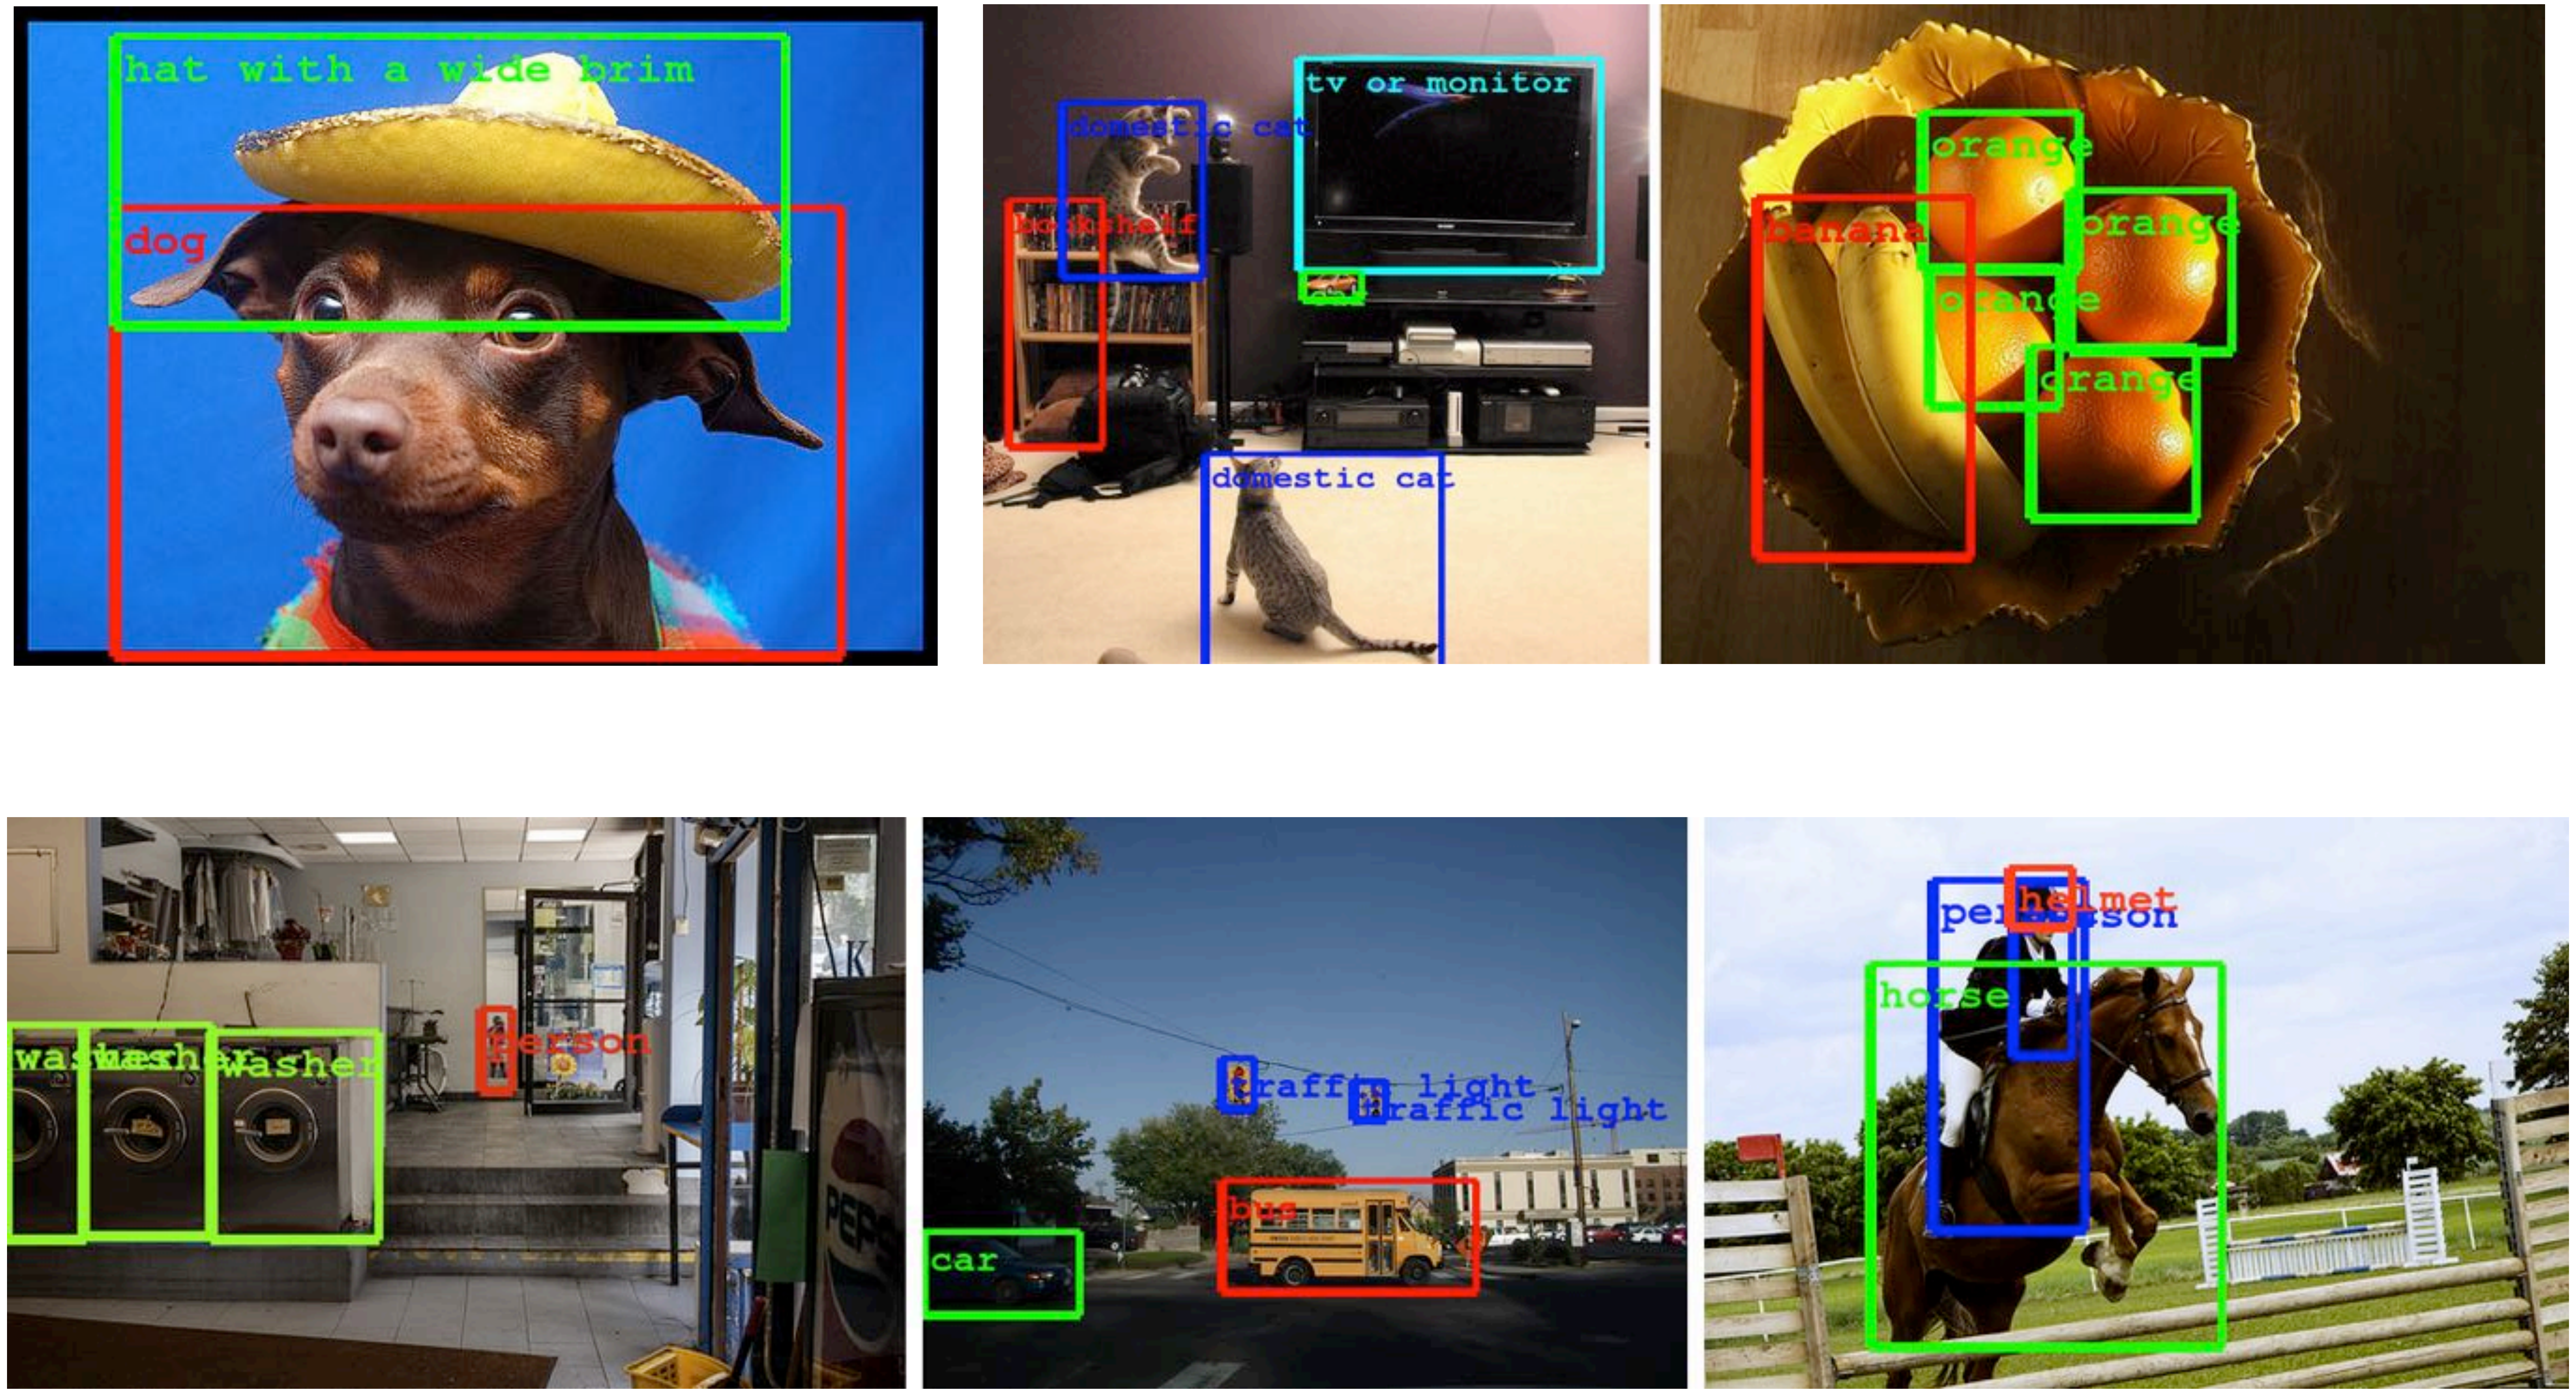

http://googleresearch.blogspot.de/2014/09/building-deeper-understanding-of-images.html

Image: Google
See e.g. http://arxiv.org/pdf/1412.0035v1.pdf

Image: Google

Gatys, Leon A., Alexander S. Ecker, and Matthias Bethge. "A neural algorithm of artistic style." arXiv preprint arXiv:1508.06576 (2015).

CVPR 2007: Minneapolis, Minnesota, USA

2007 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2007), 18-23 June 2007, Minneapolis, Minnesota, USA. IEEE Computer Society 2007, ISBN 1-4244-1179-3

Oral session 1A: Matching and Features

- Eric Nowak, Frédéric Jurie:
  Learning Visual Similarity Measures for Comparing Never Seen Objects.
- Herve Jegou, Hedi Harzallah, Cordelia Schmid:
  A contextual dissimilarity measure for accurate and efficient image search.
- Simon A. J. Winder, Matthew Brown:
  Learning Local Image Descriptors.
- Hongli Deng, Wei Zhang, Eric N. Mortensen, Thomas G. Dietterich, Linda G. Shapiro:
  Principal Curvature-Based Region Detector for Object Recognition.
- Alexander Toshev, Jianbo Shi, Kostas Daniilidis:
  Image Matching via Saliency Region Correspondences.

Oral session 1B: Motion Segmentation and Tracking

- Roberto Tron, René Vidal:
  A Benchmark for the Comparison of 3-D Motion Segmentation Algorithms.
- Hongdong Li:
  Two-View Motion Segmentation from Linear Programming Relaxation.
- Le Lu, Gregory D. Hager:
  A Nonparametric Treatment for Location/Segmentation Based Visual Tracking.

resistance is futile

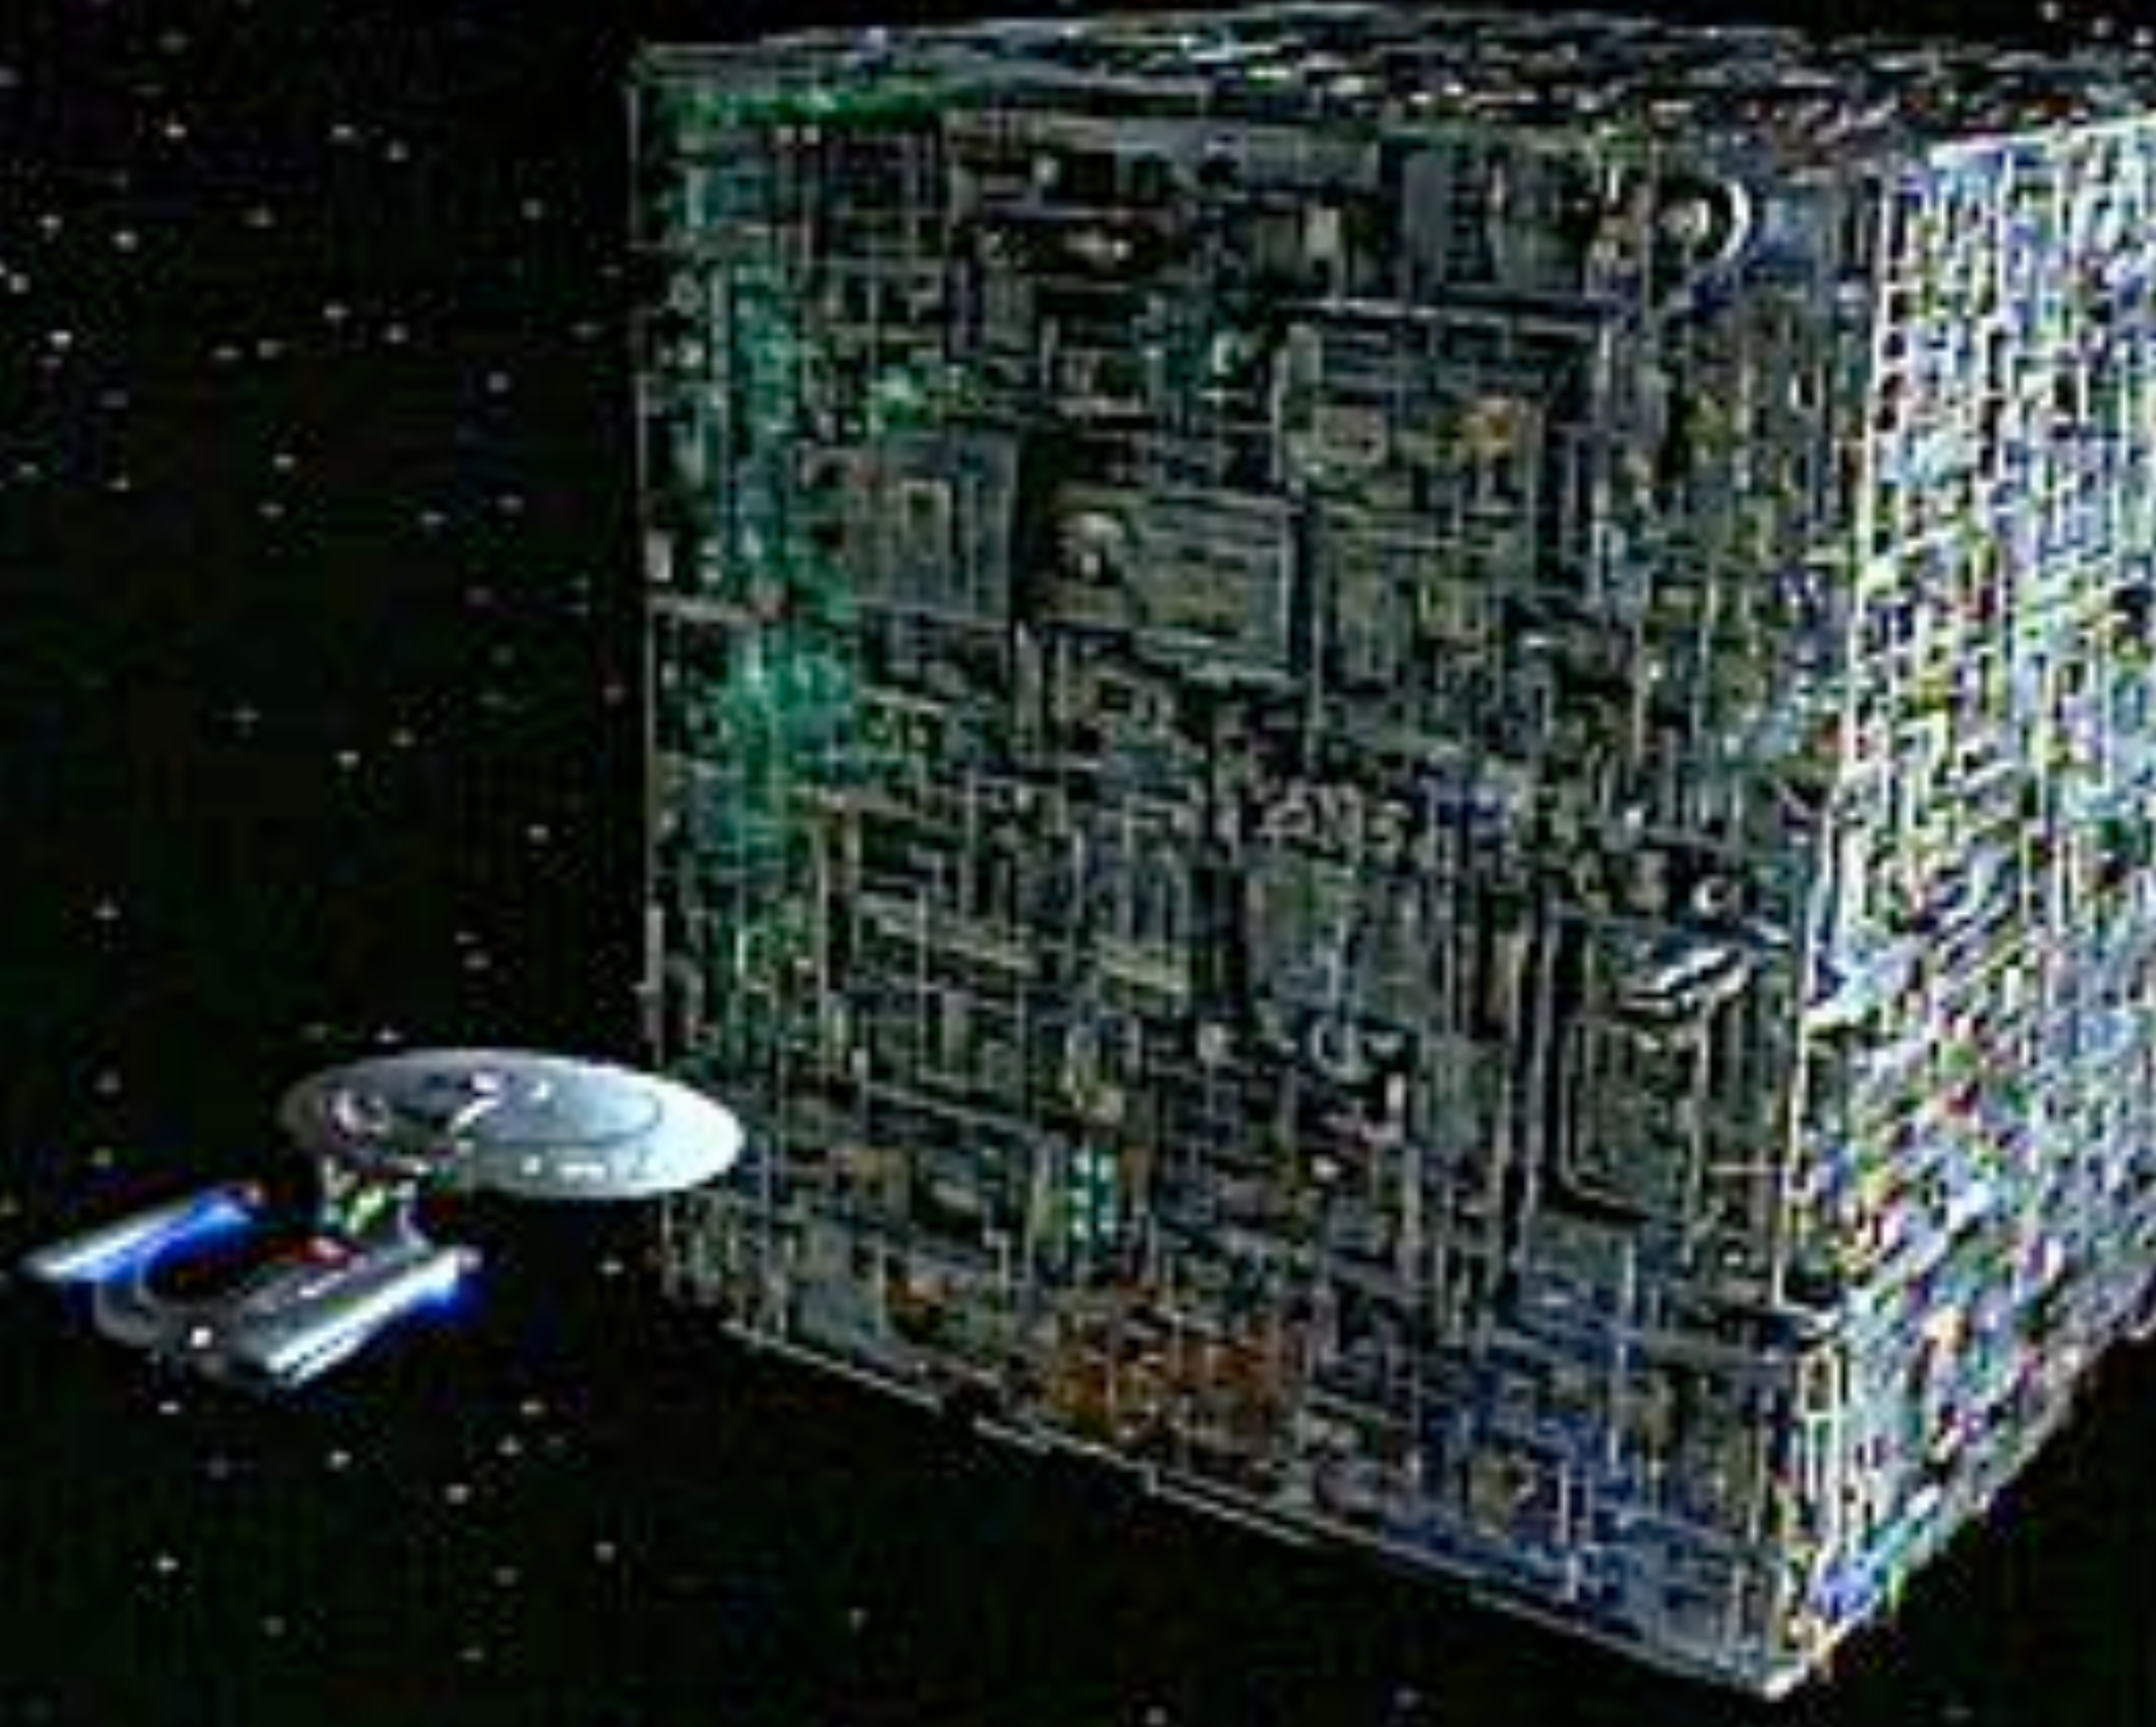

CVPR 2015

Google

DEEPMIND

Martin Riedmiller

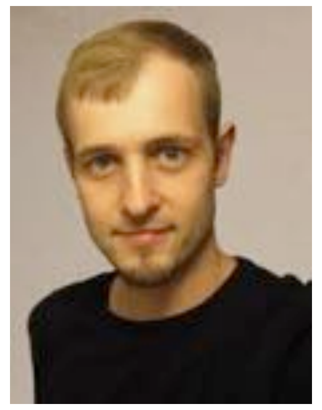

Geoffrey Hinton

Yann LeCun

Baidu 百度

Andrew Ng

facebook

1: Deep Learning

2: Reinforcement Learning!

Reinforcement Learning

- Learning through interaction, by means of trial and error. **Examples:** Closed Loop Control, Gas Turbine, Engine .

Batch Reinforcement Learning: NFQ

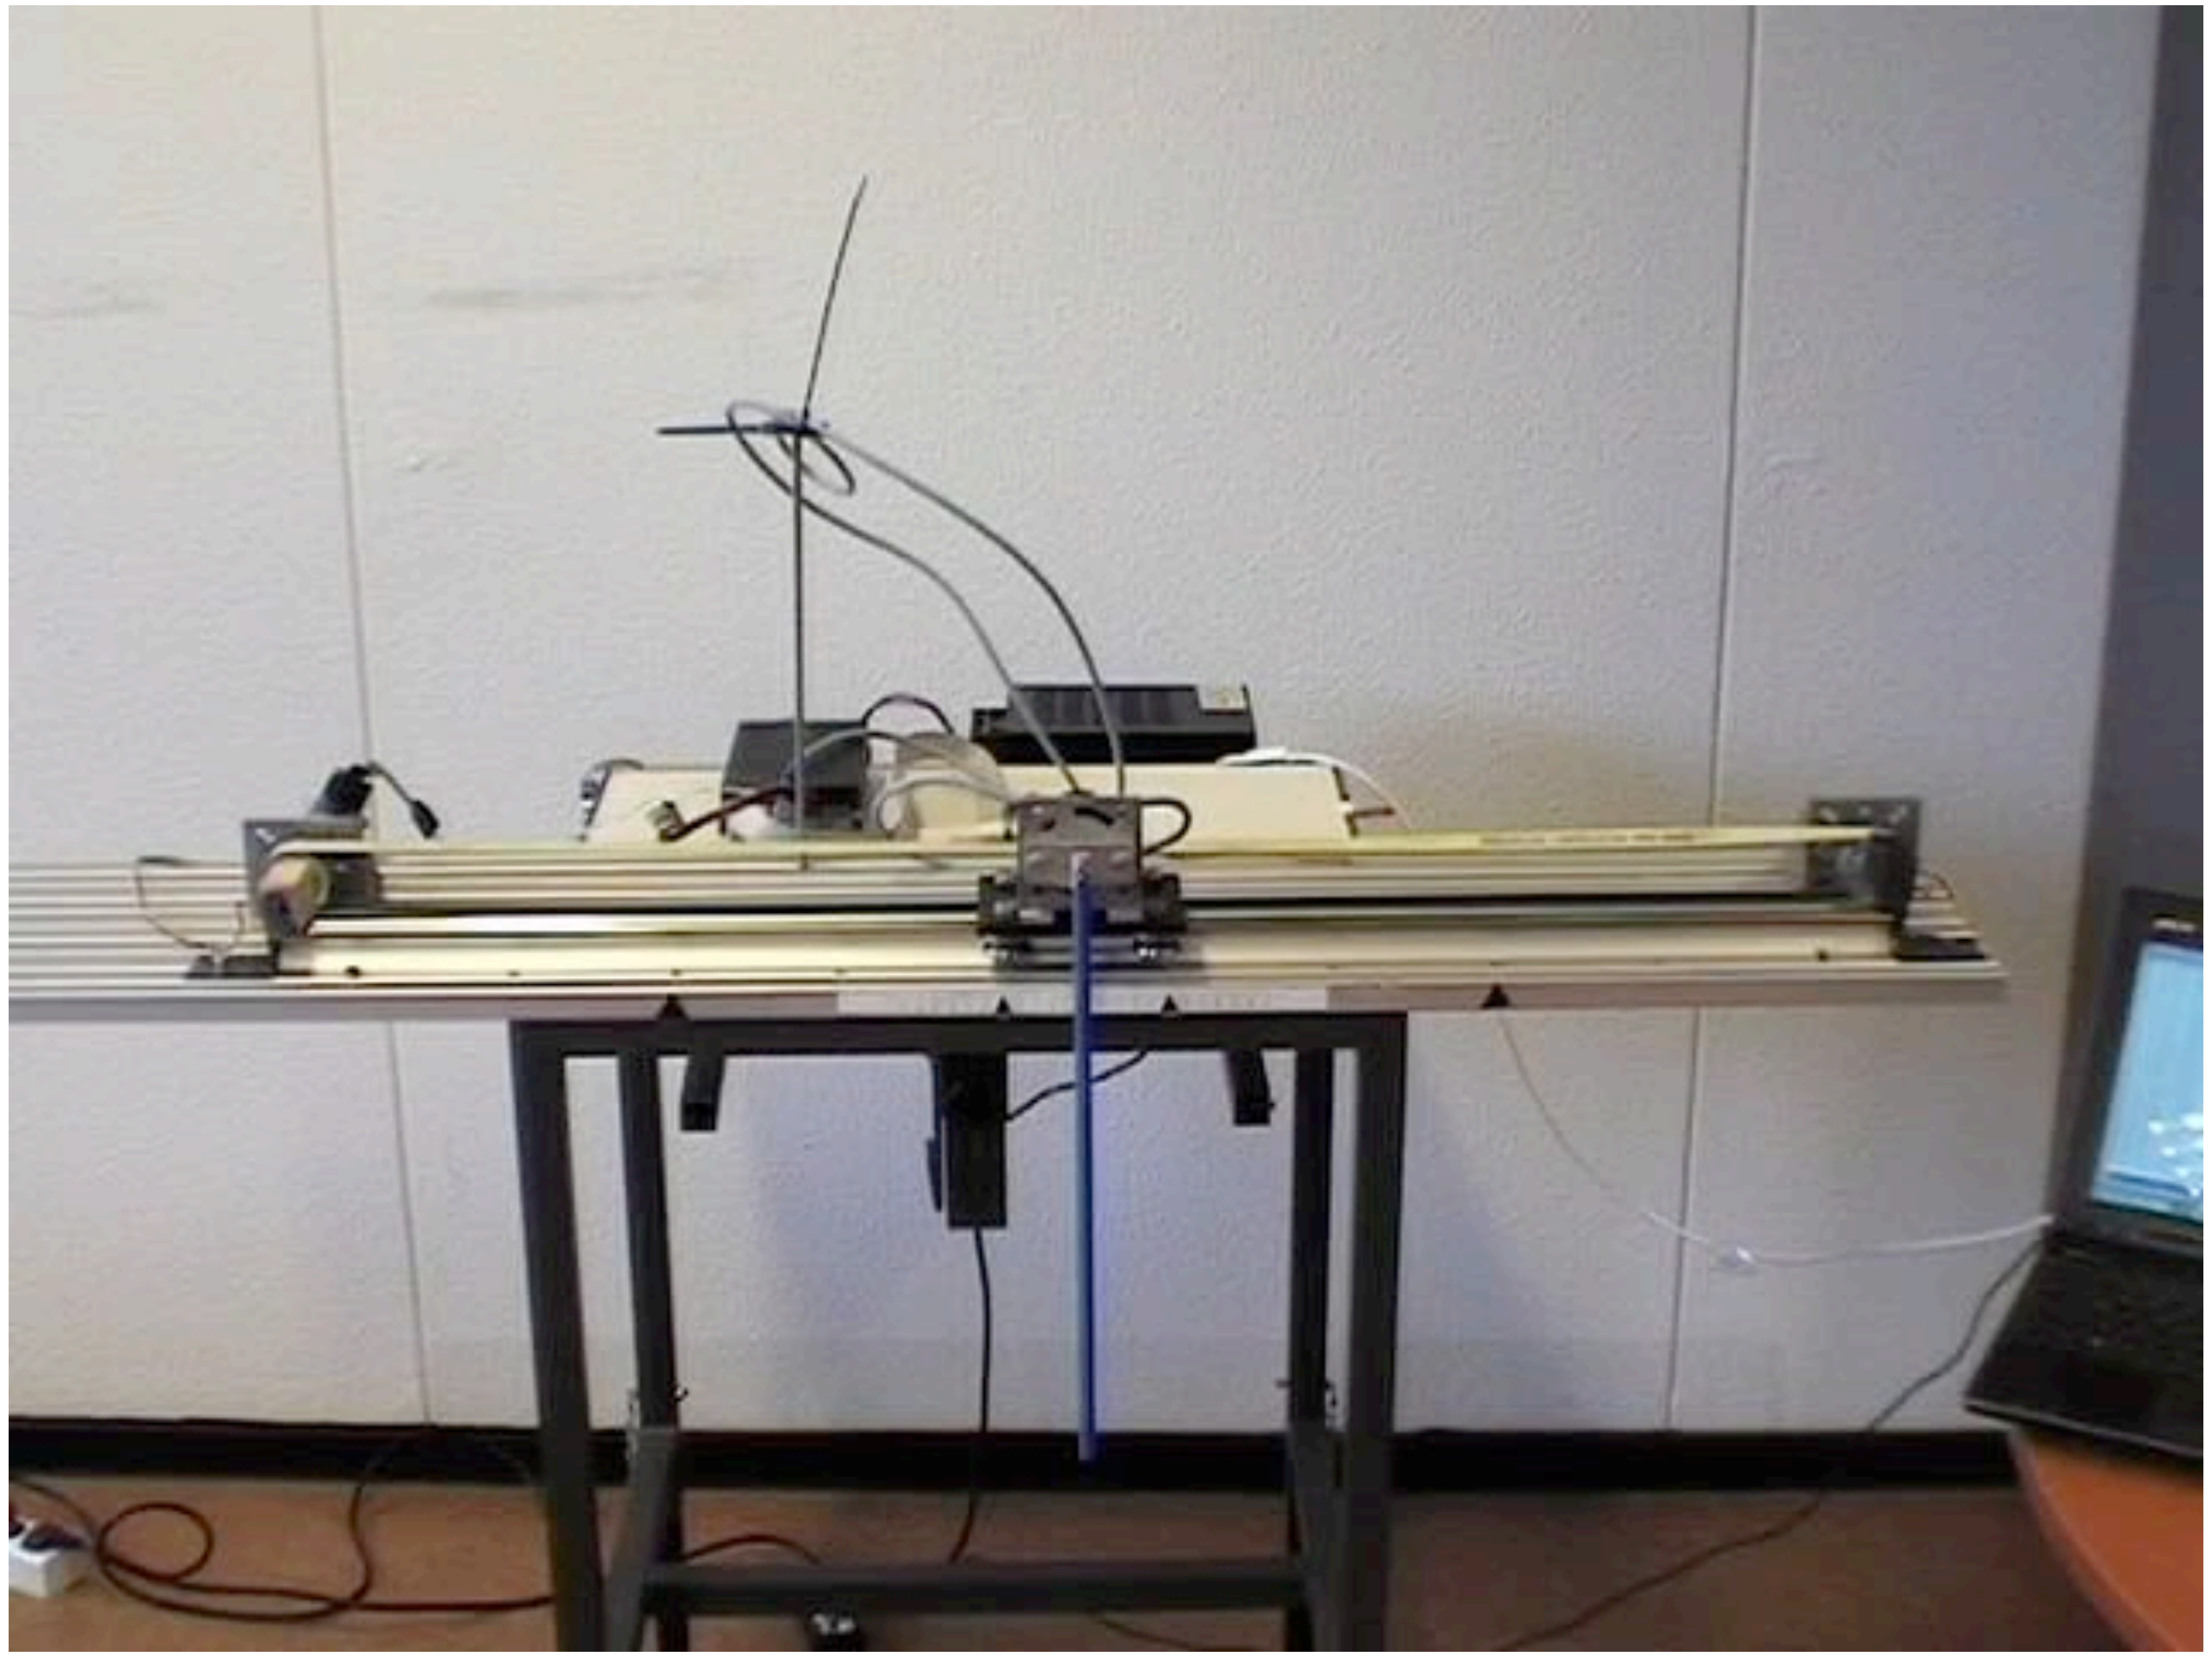

Martin Riedmiller. "Neural fitted Q iteration–first experiences with a data efficient neural reinforcement learning method." ECML 2005.

Break Through:
Sufficient stability with Neural Network
Data efficiency: Batch Reinforcement Learning

Batch RL on Soccer Robots

Riedmiller, Martin, Thomas Gabel, Roland Hafner, and Sascha Lange. "Reinforcement learning for robot soccer." *Autonomous Robots* 27, no. 1 (2009): 55-73.

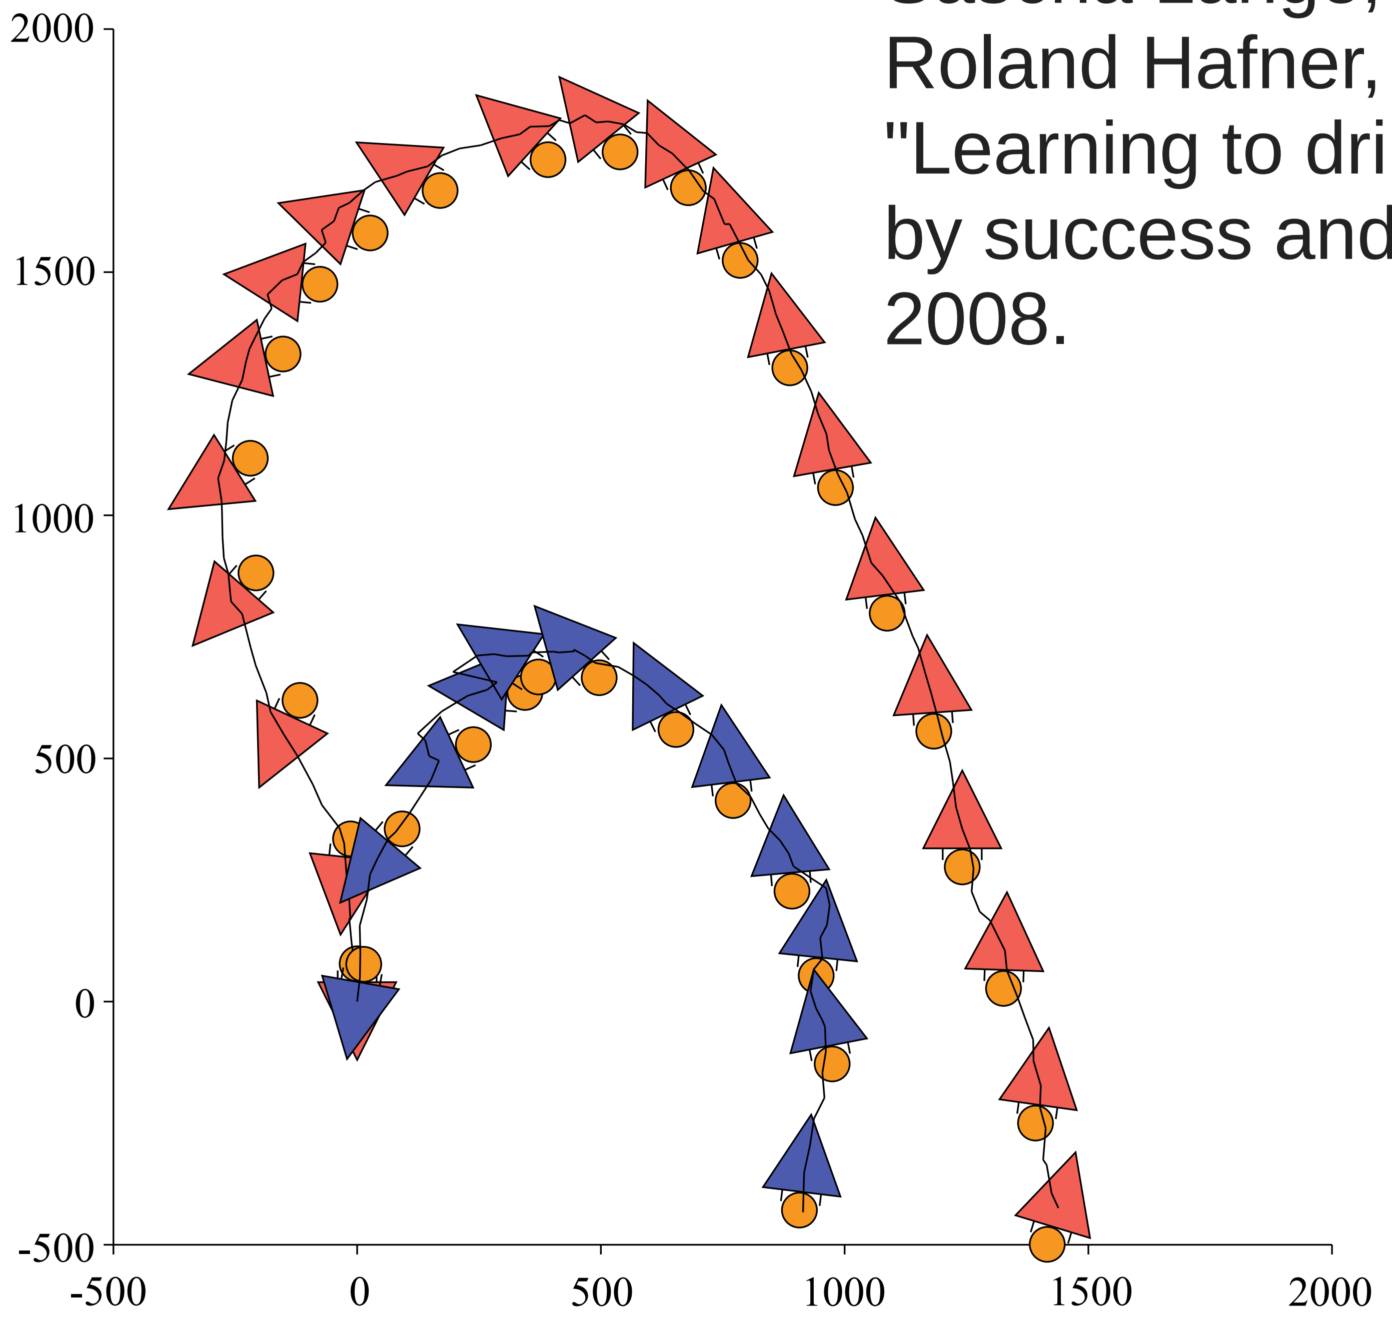

Riedmiller, Martin, Roland Hafner, Sascha Lange, Martin Lauer, Roland Hafner, and Sascha Lange. "Learning to dribble on a real robot by success and failure." ICRA 2008.

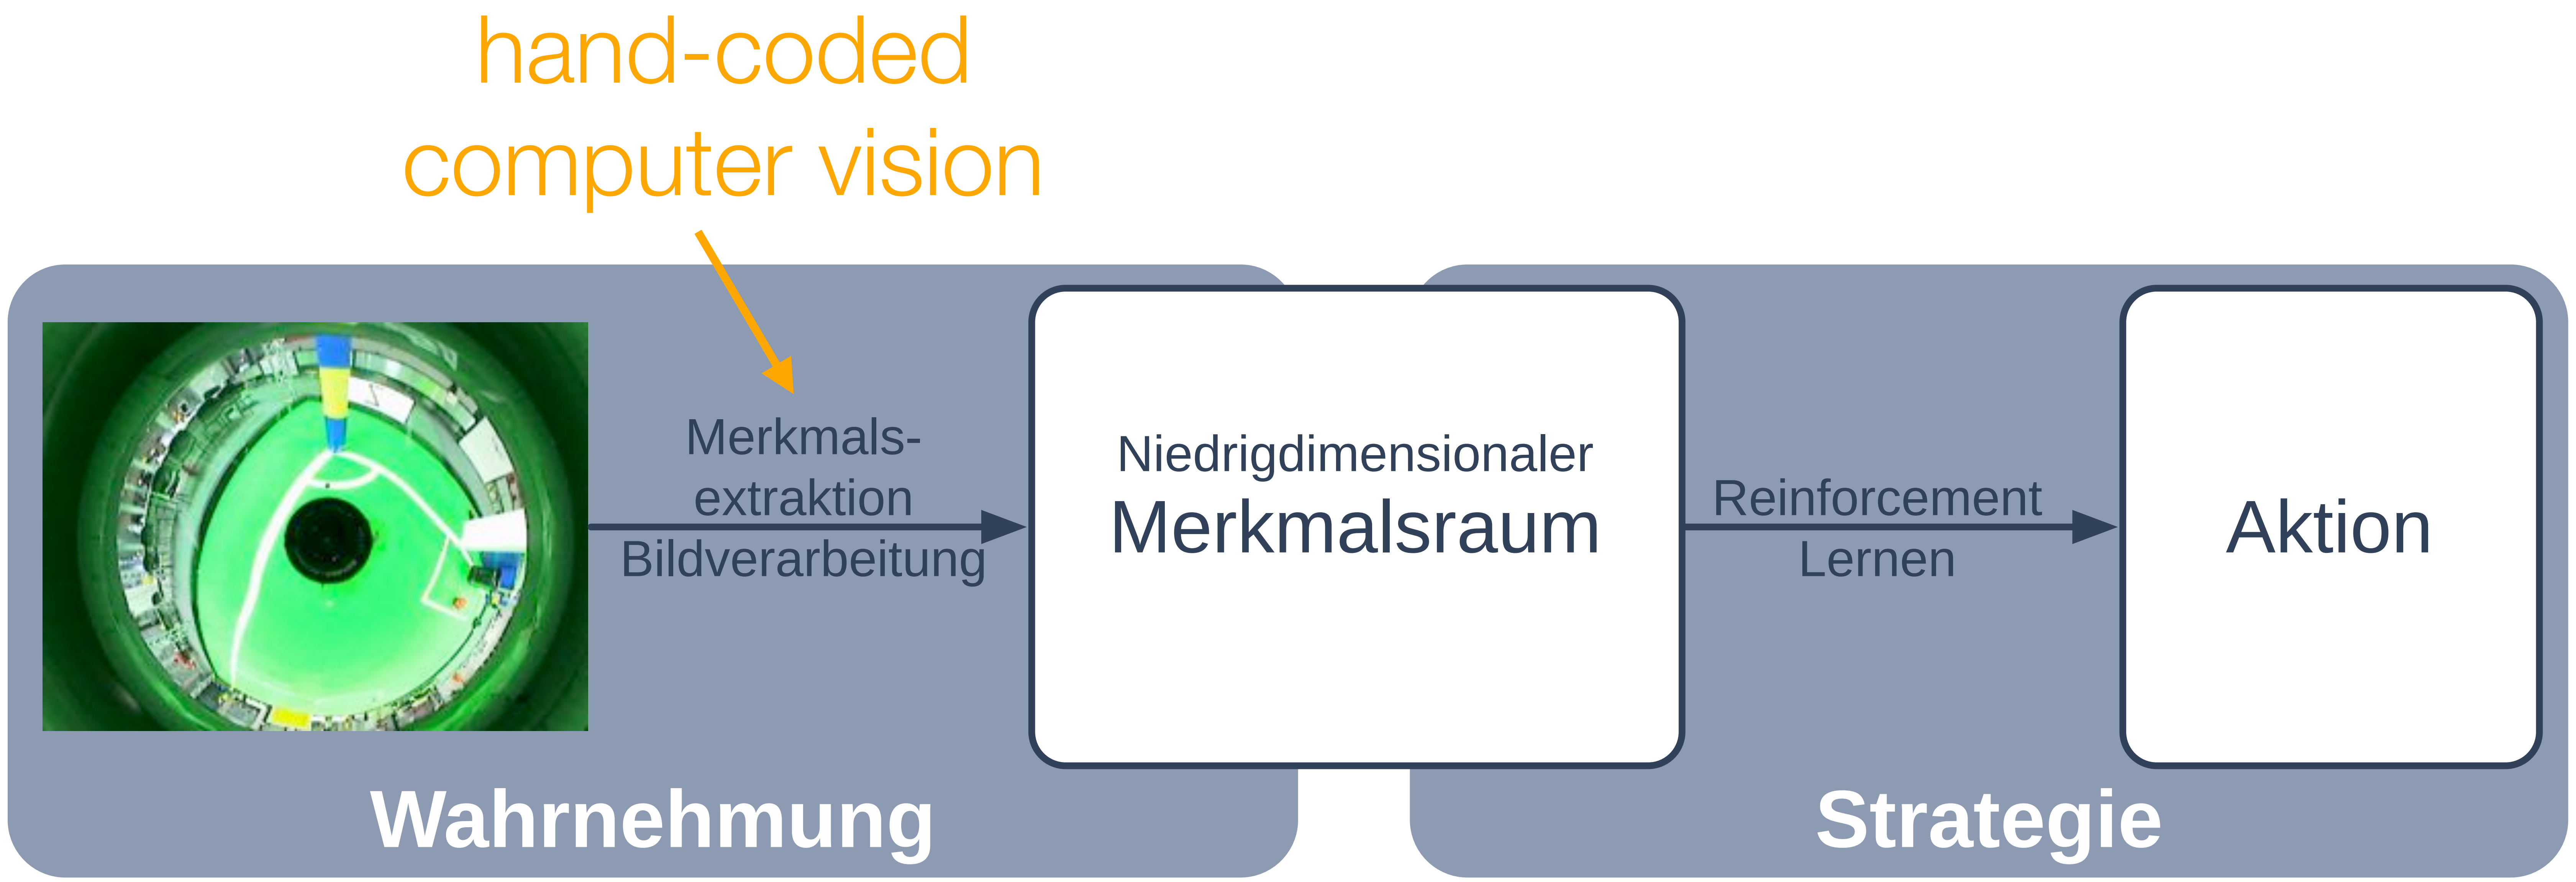

Classical Solution: Two Separate Processing Steps

- Relevant information must be extracted from the images and "properly" encoded in a first processing step

Training of Deep Auto Encoders

Idea: Integration of Deep Learning into the Reinforcement Learning procedure

- Deep Auto Encoder: learn low-dimensional feature spaces
- Reinforcement Learning: learn strategies

My Algorithm: Deep Fitted Q-Iteration (DFQ, 2009)

First Experiment: Synthetic Grid World

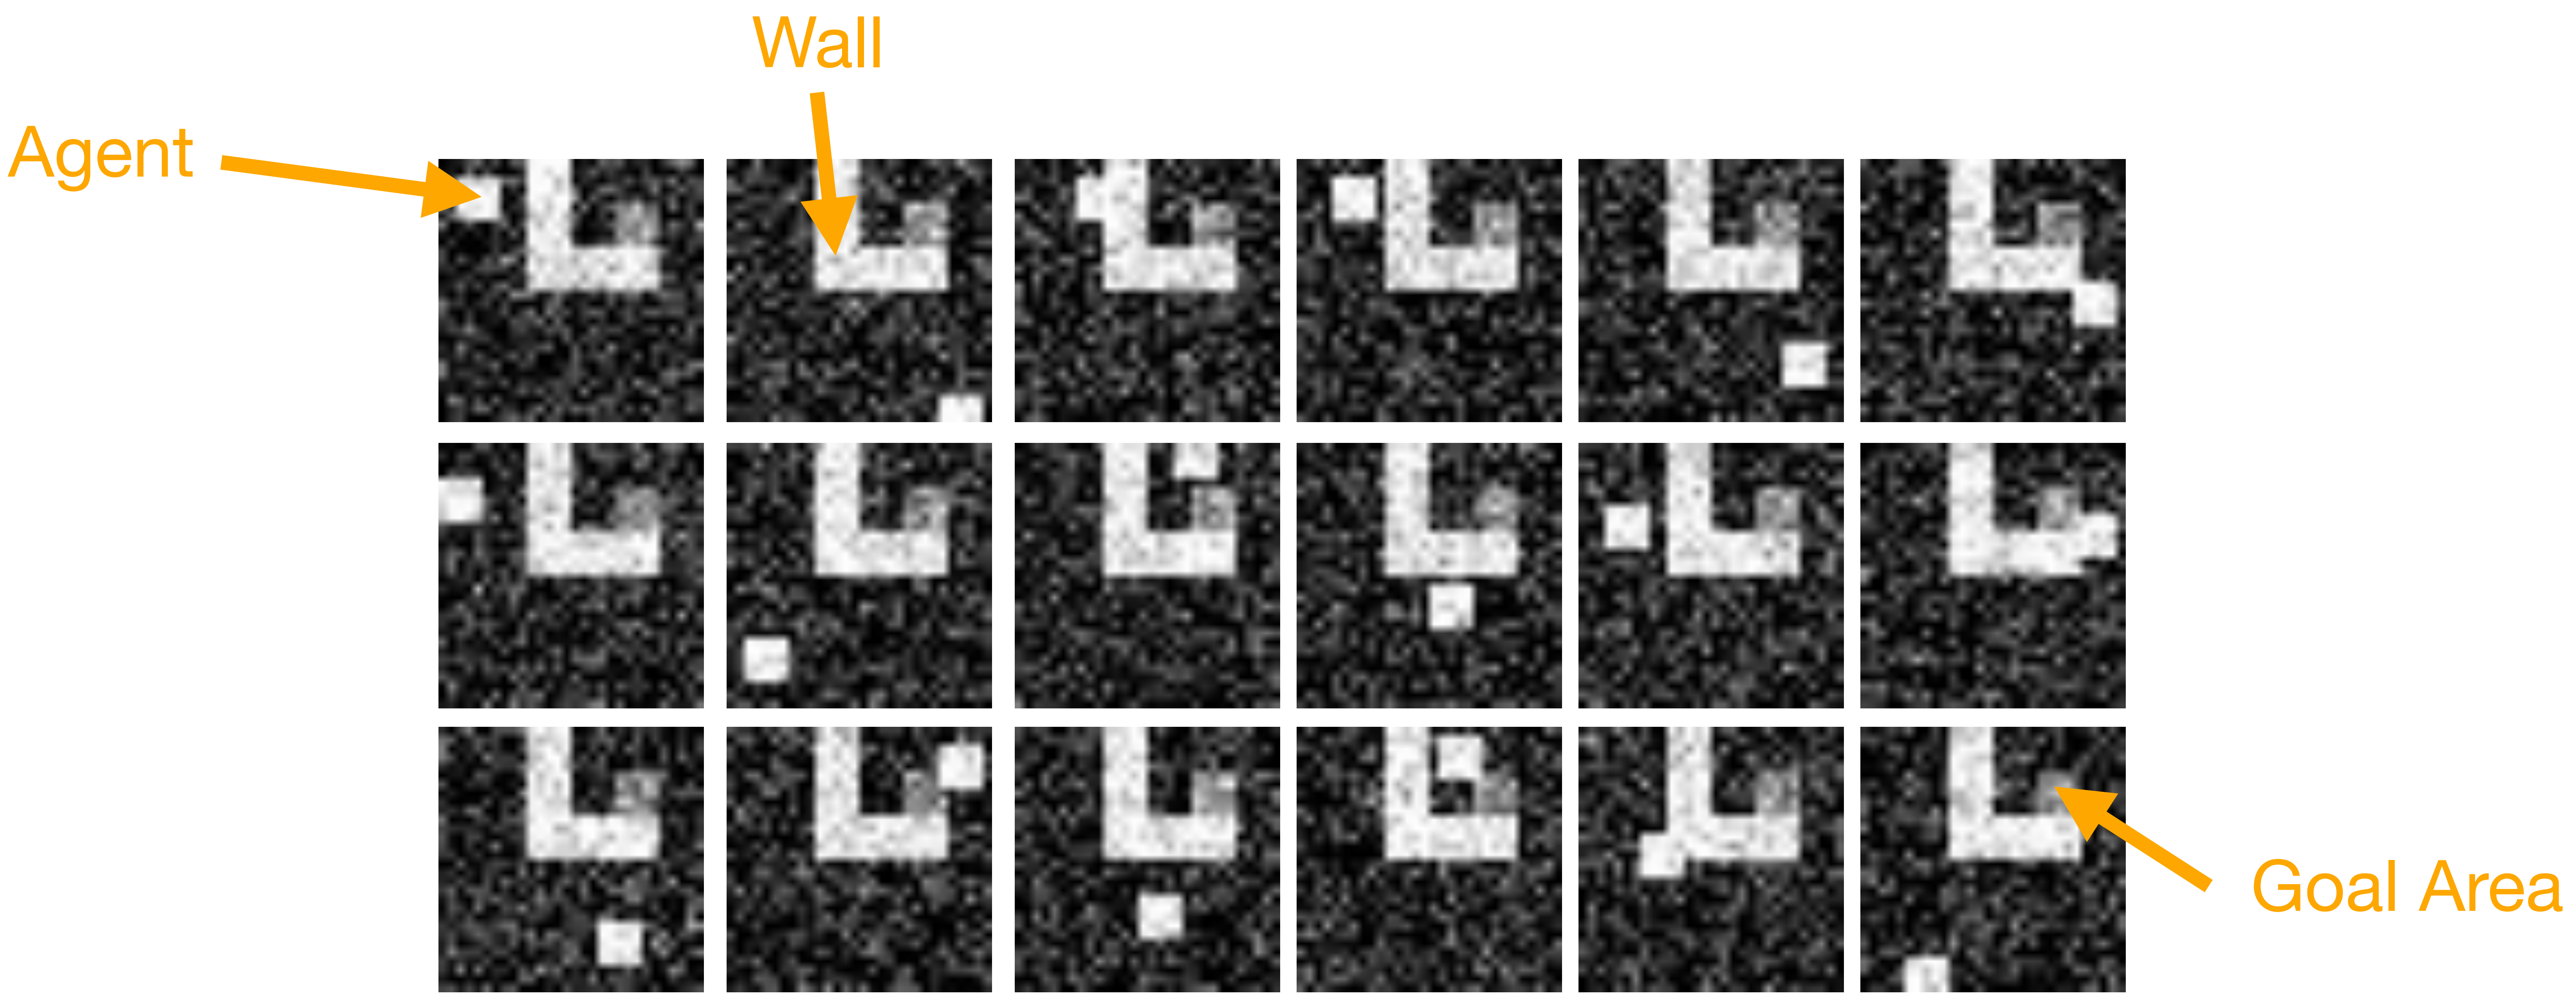

30x30 Pixel, independent white noise on pixels with $\sigma=0.1$

Computer should learn to walk the agent into the goal area as fast as possible. On the images!

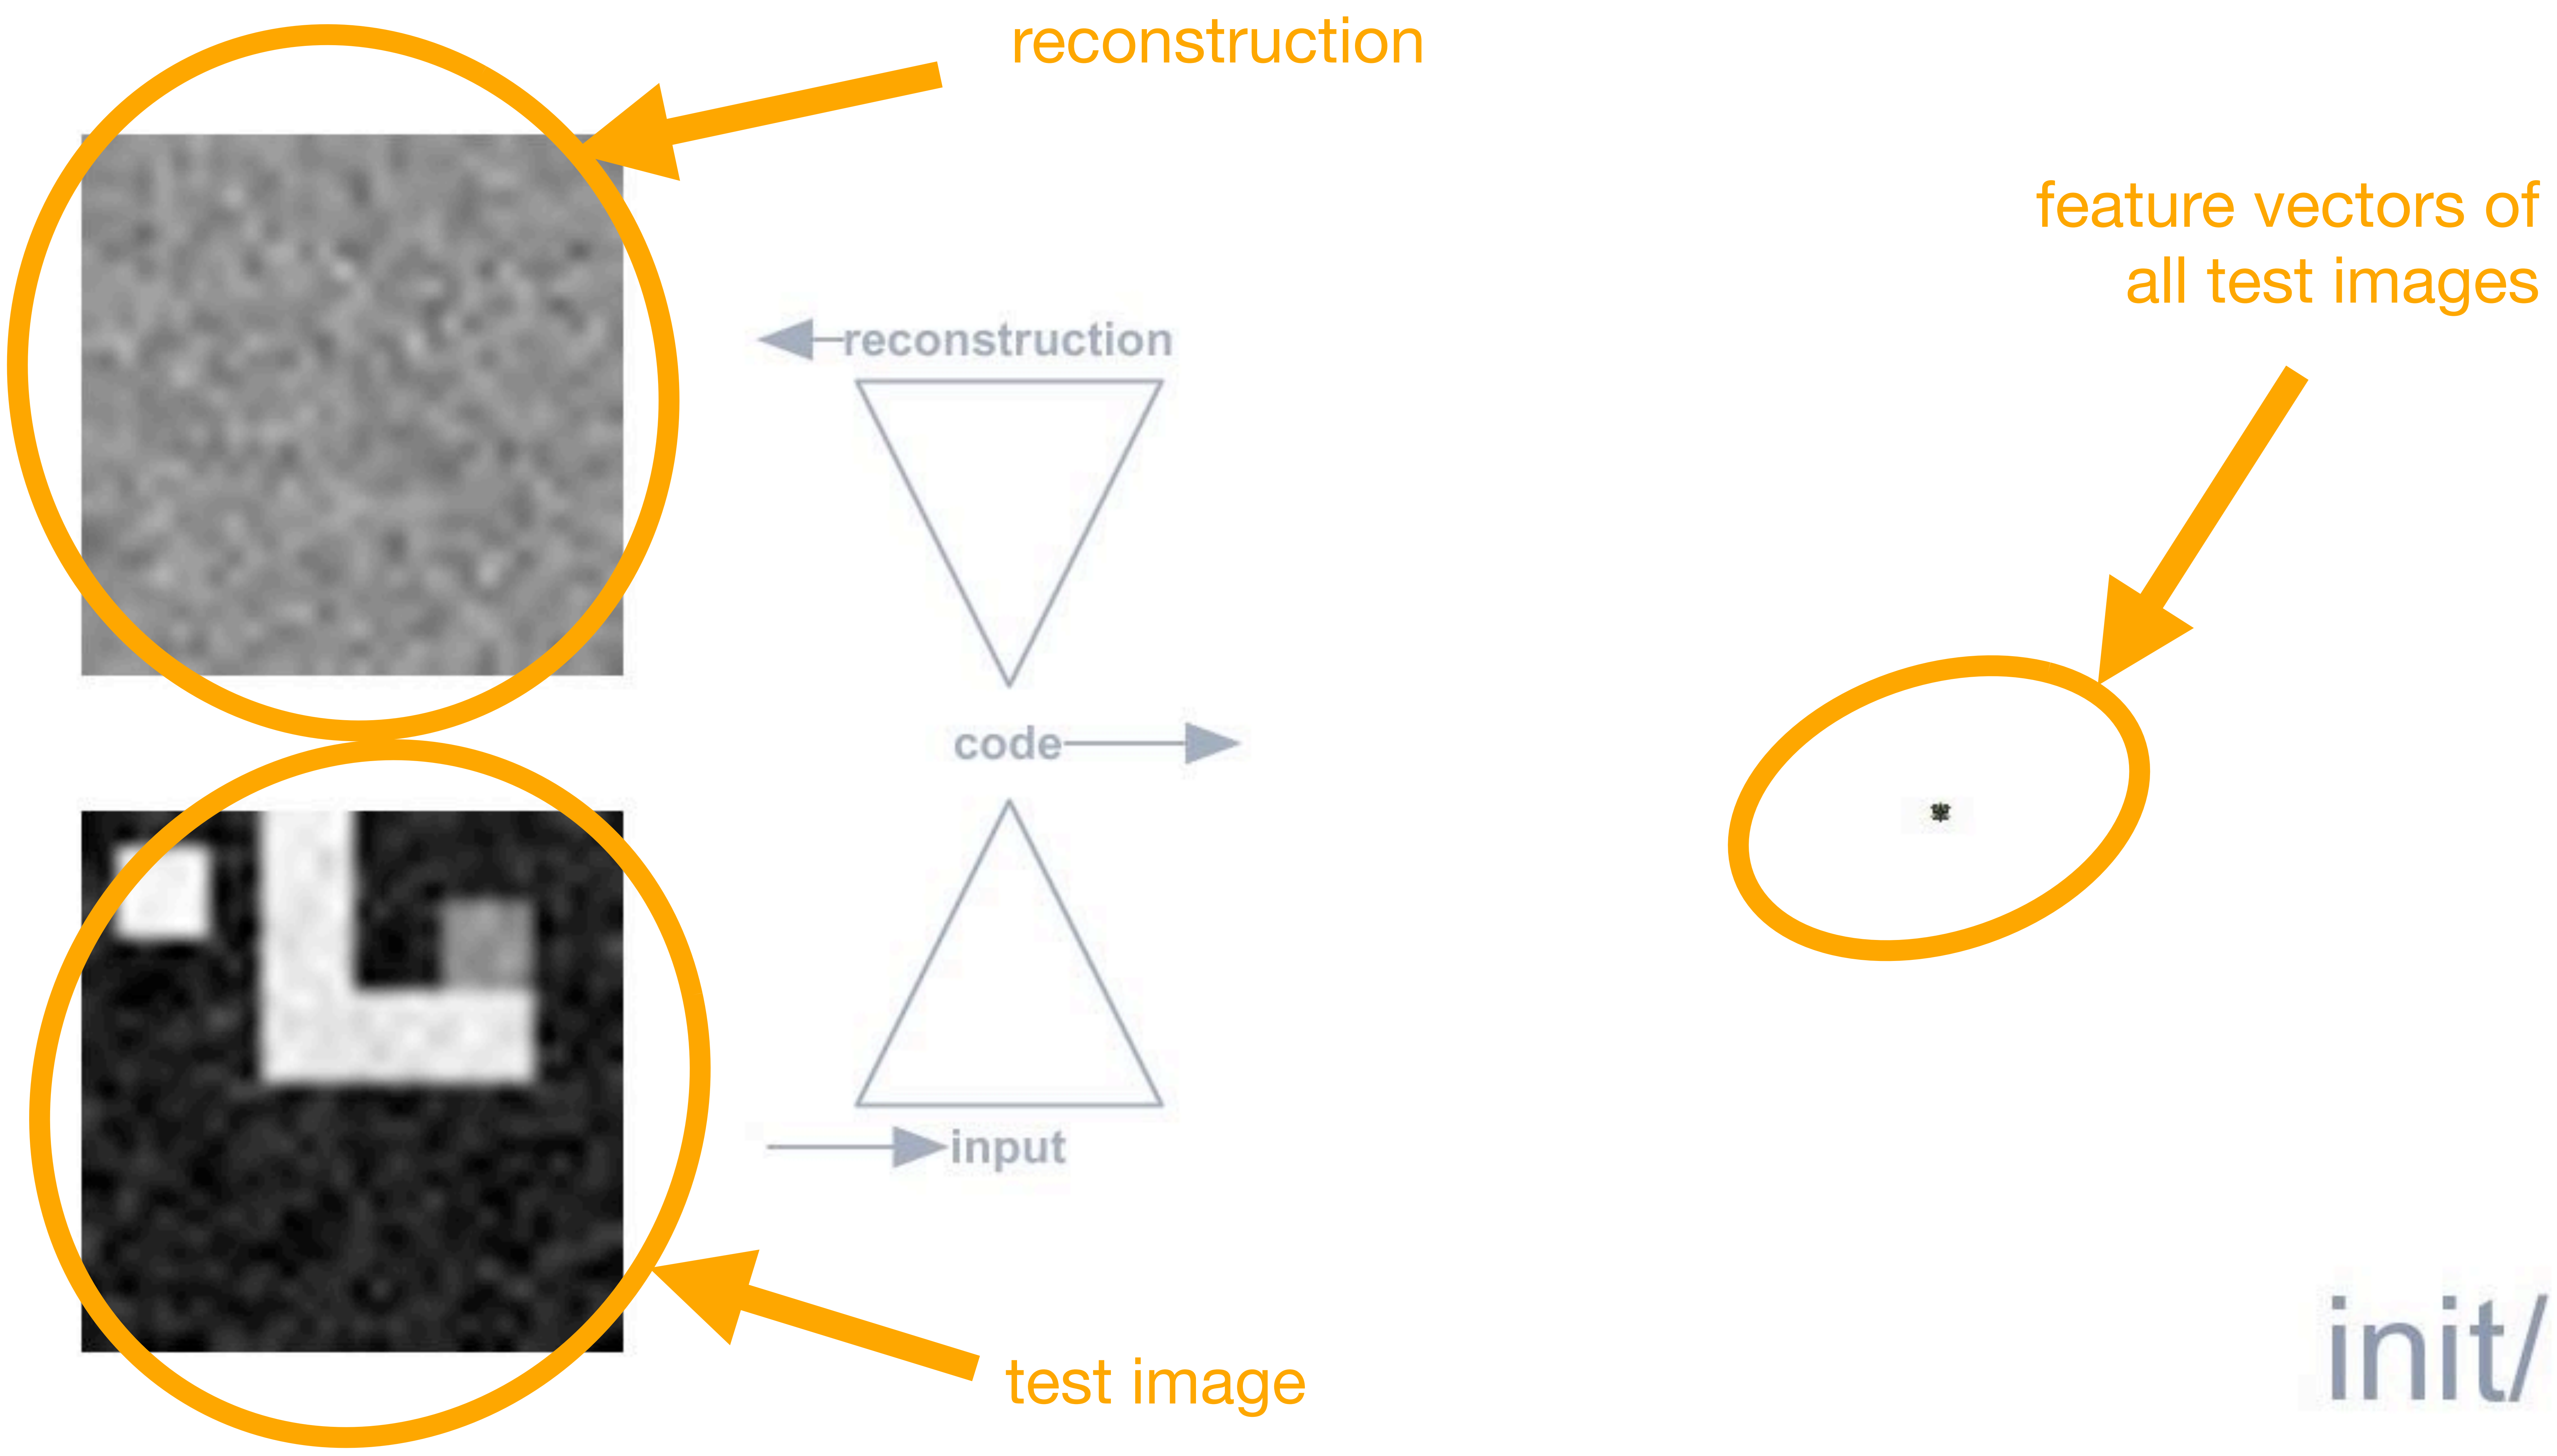
reconstruction
feature vectors of
all test images
reconstruction
code
input
test image
init/

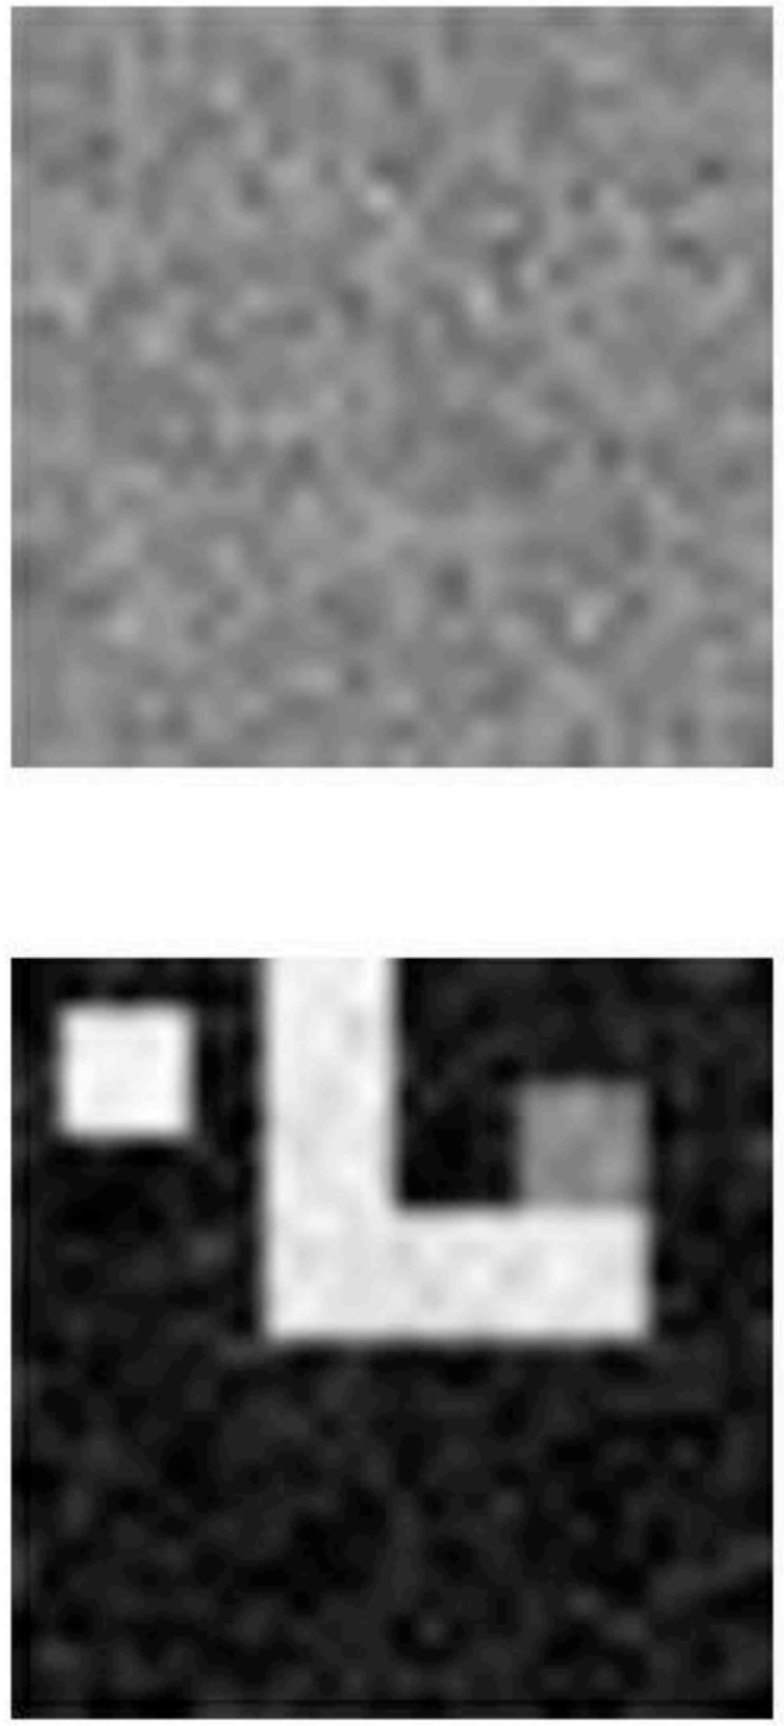

Loop over manually selected
test images

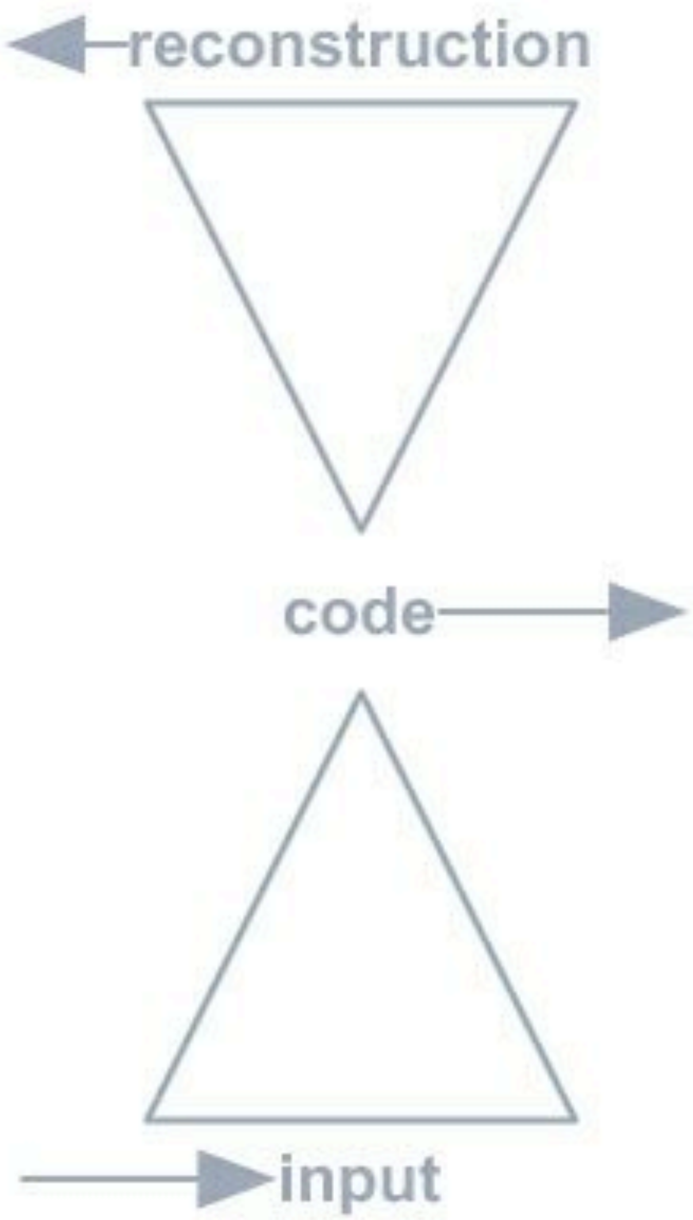

init/

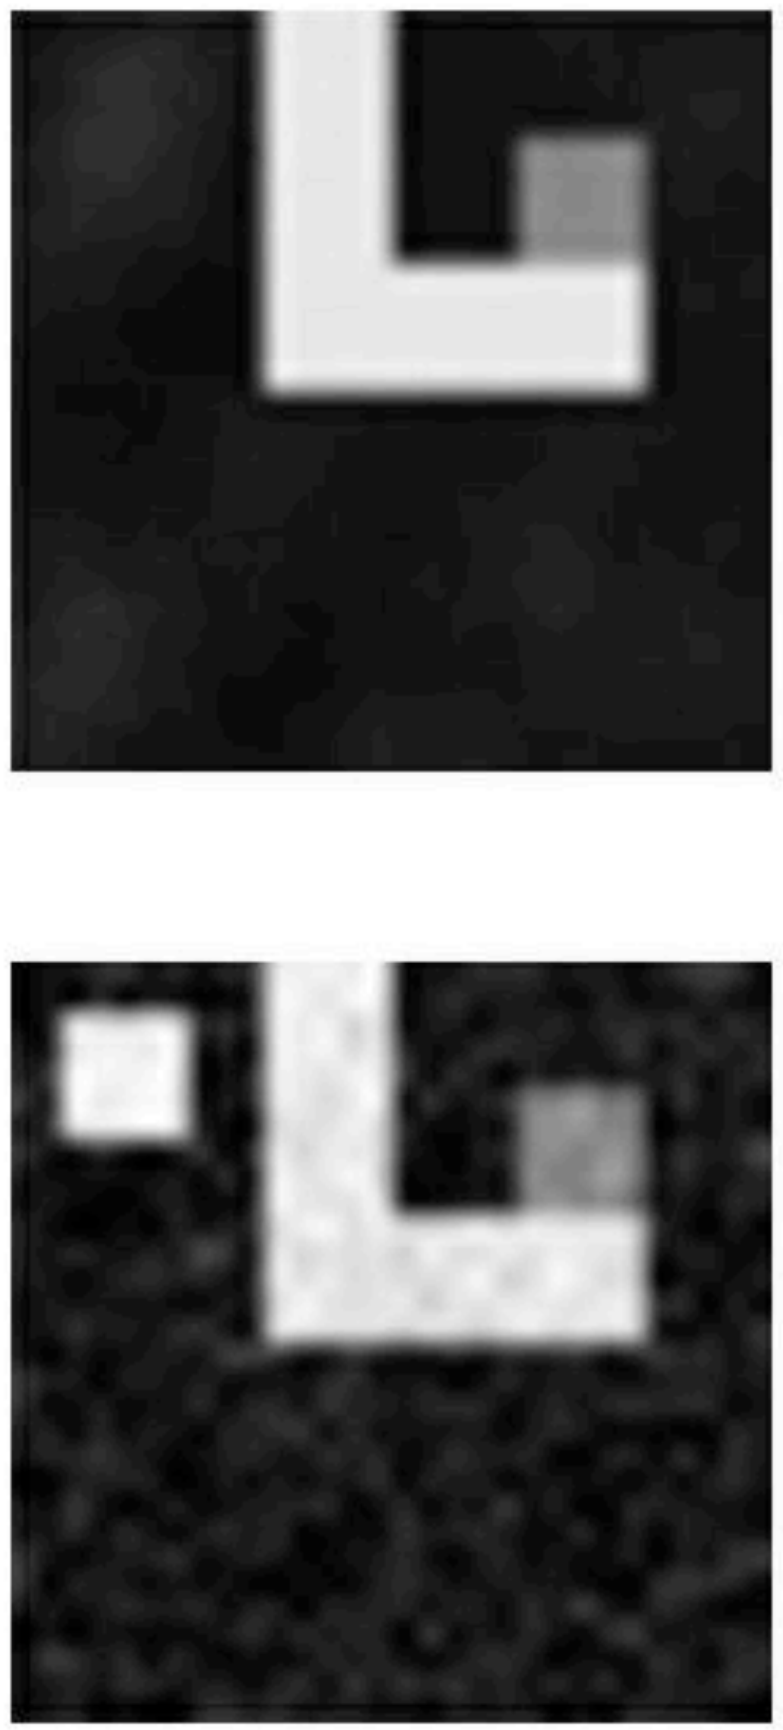

Loop over manually selected test images

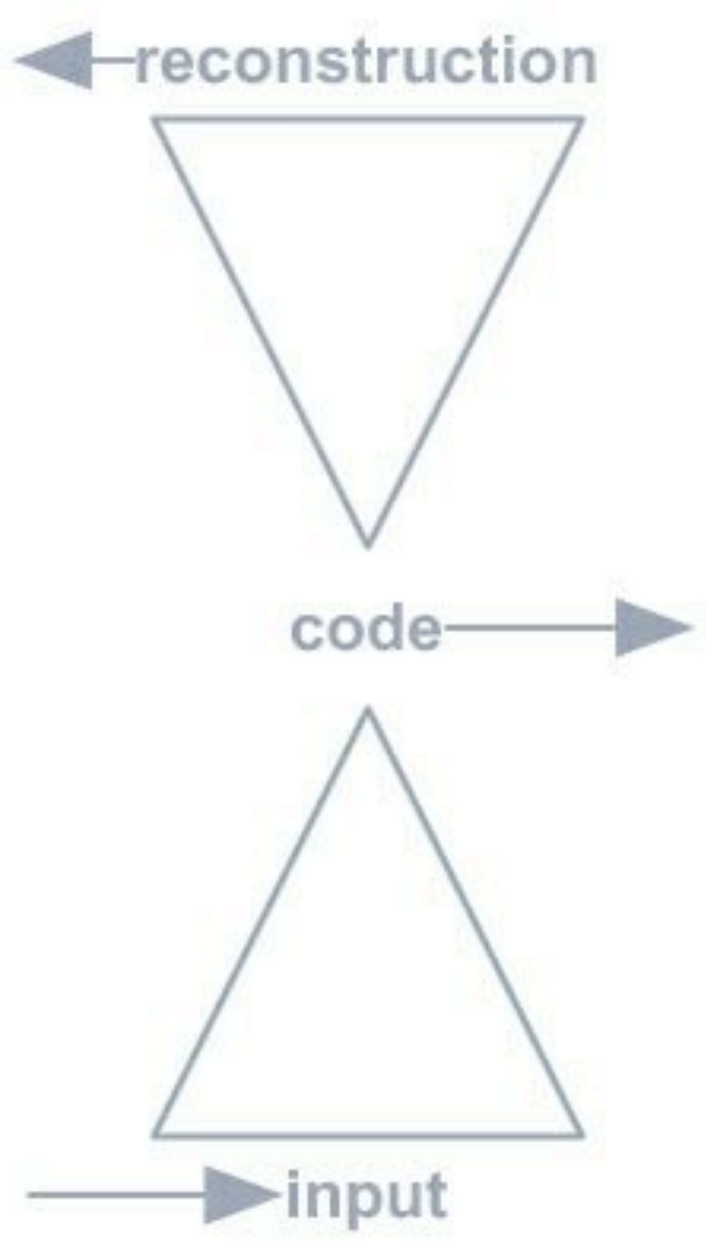

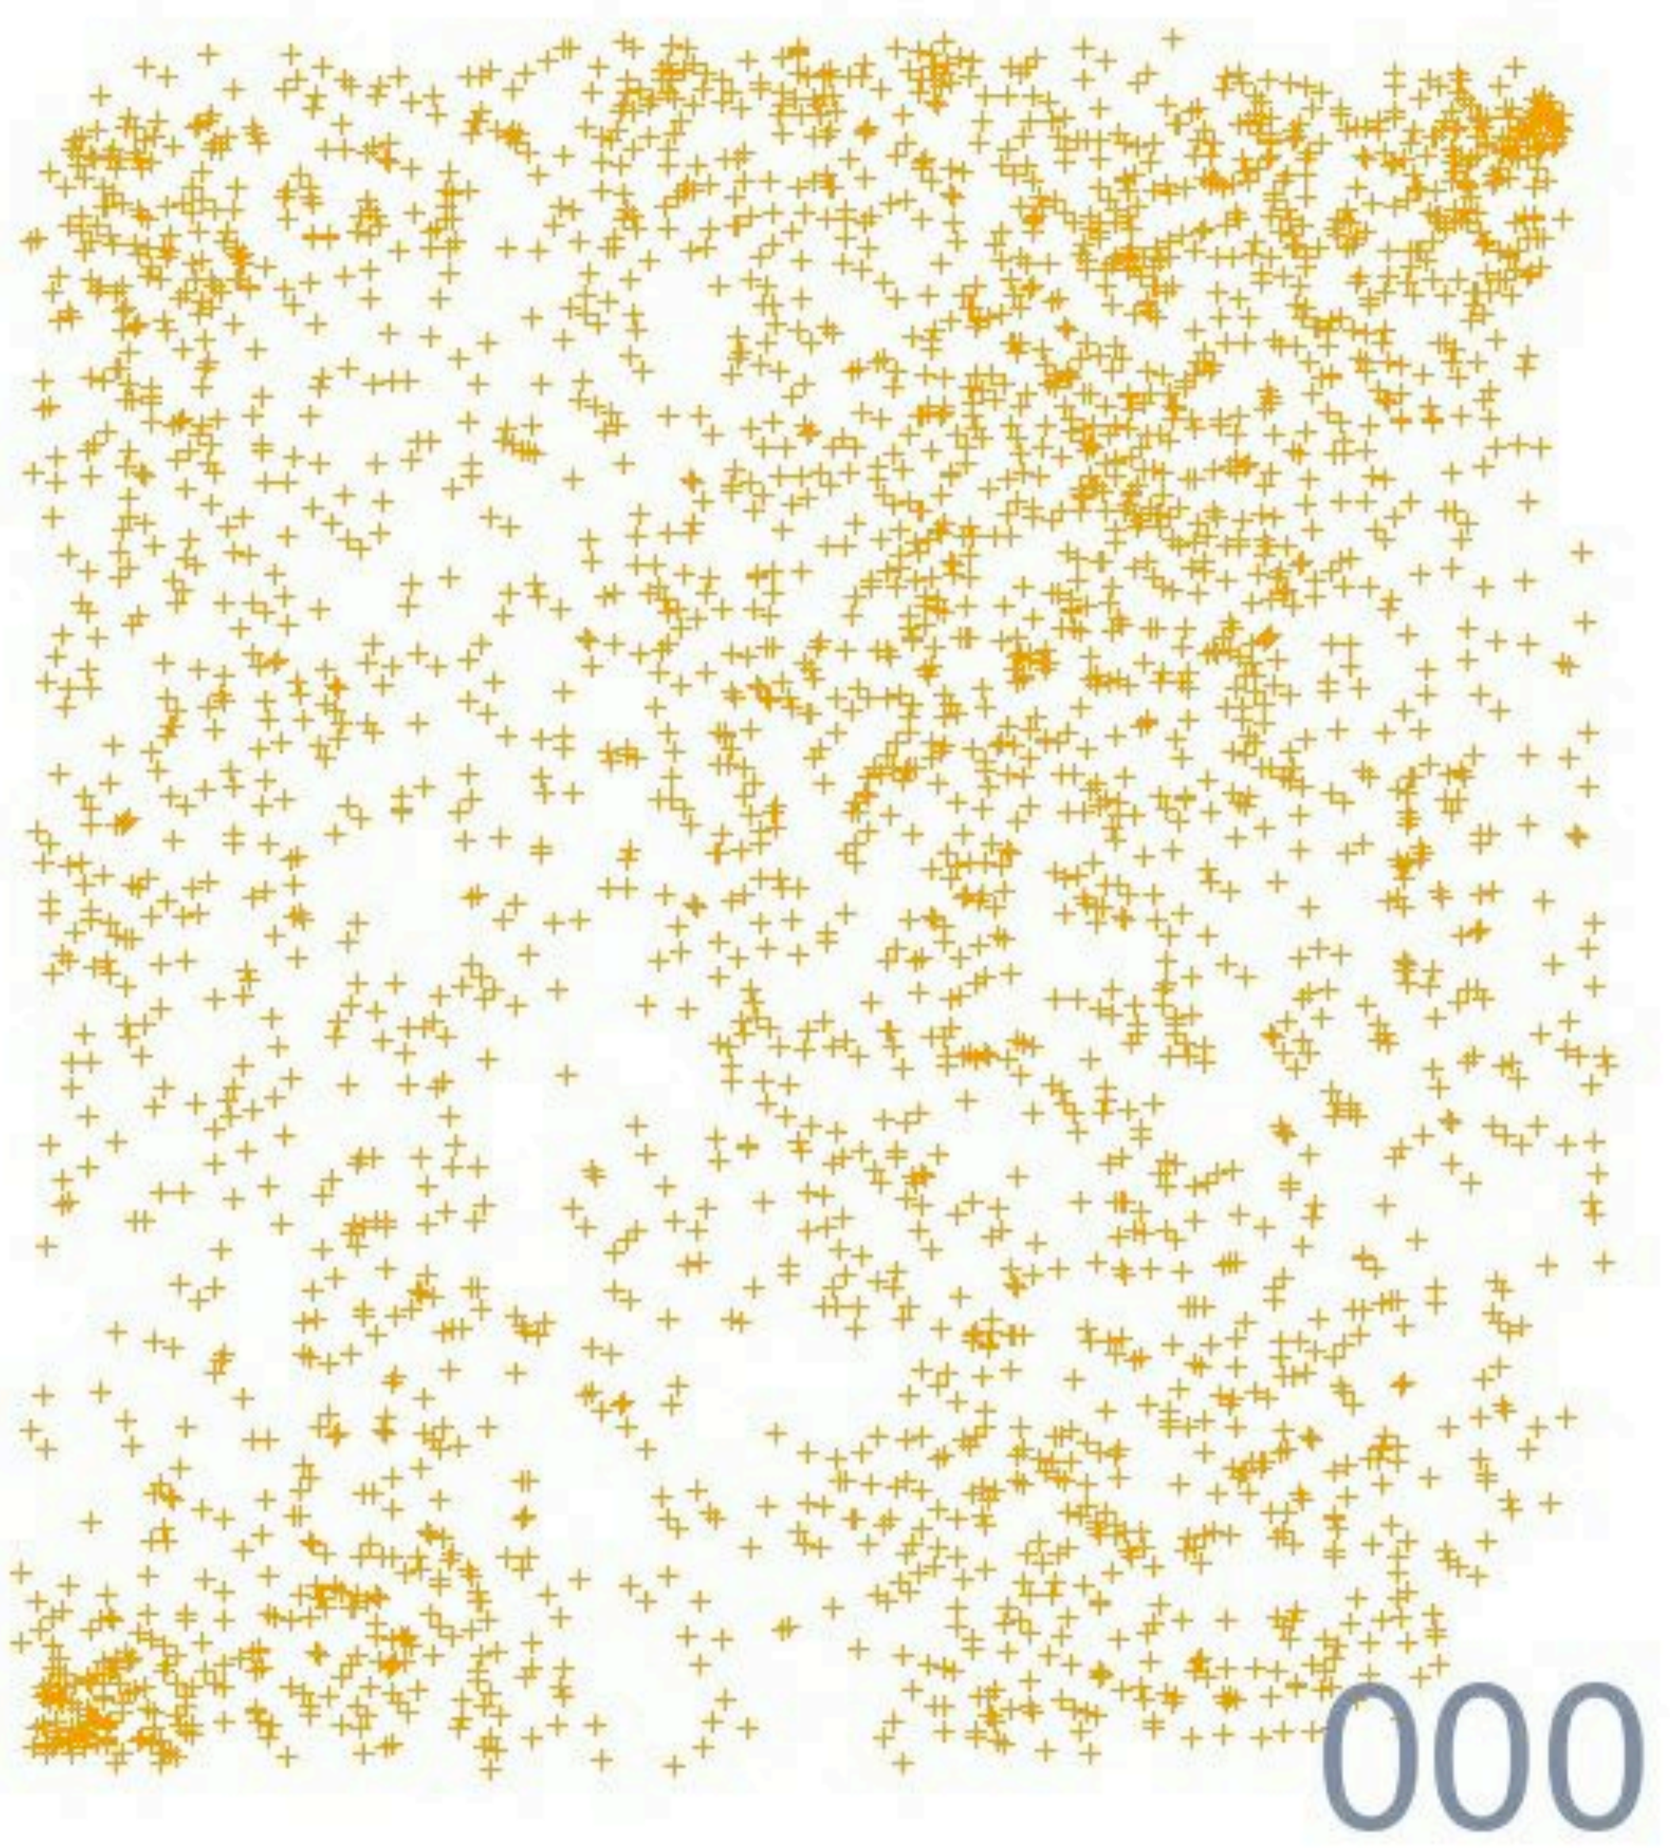

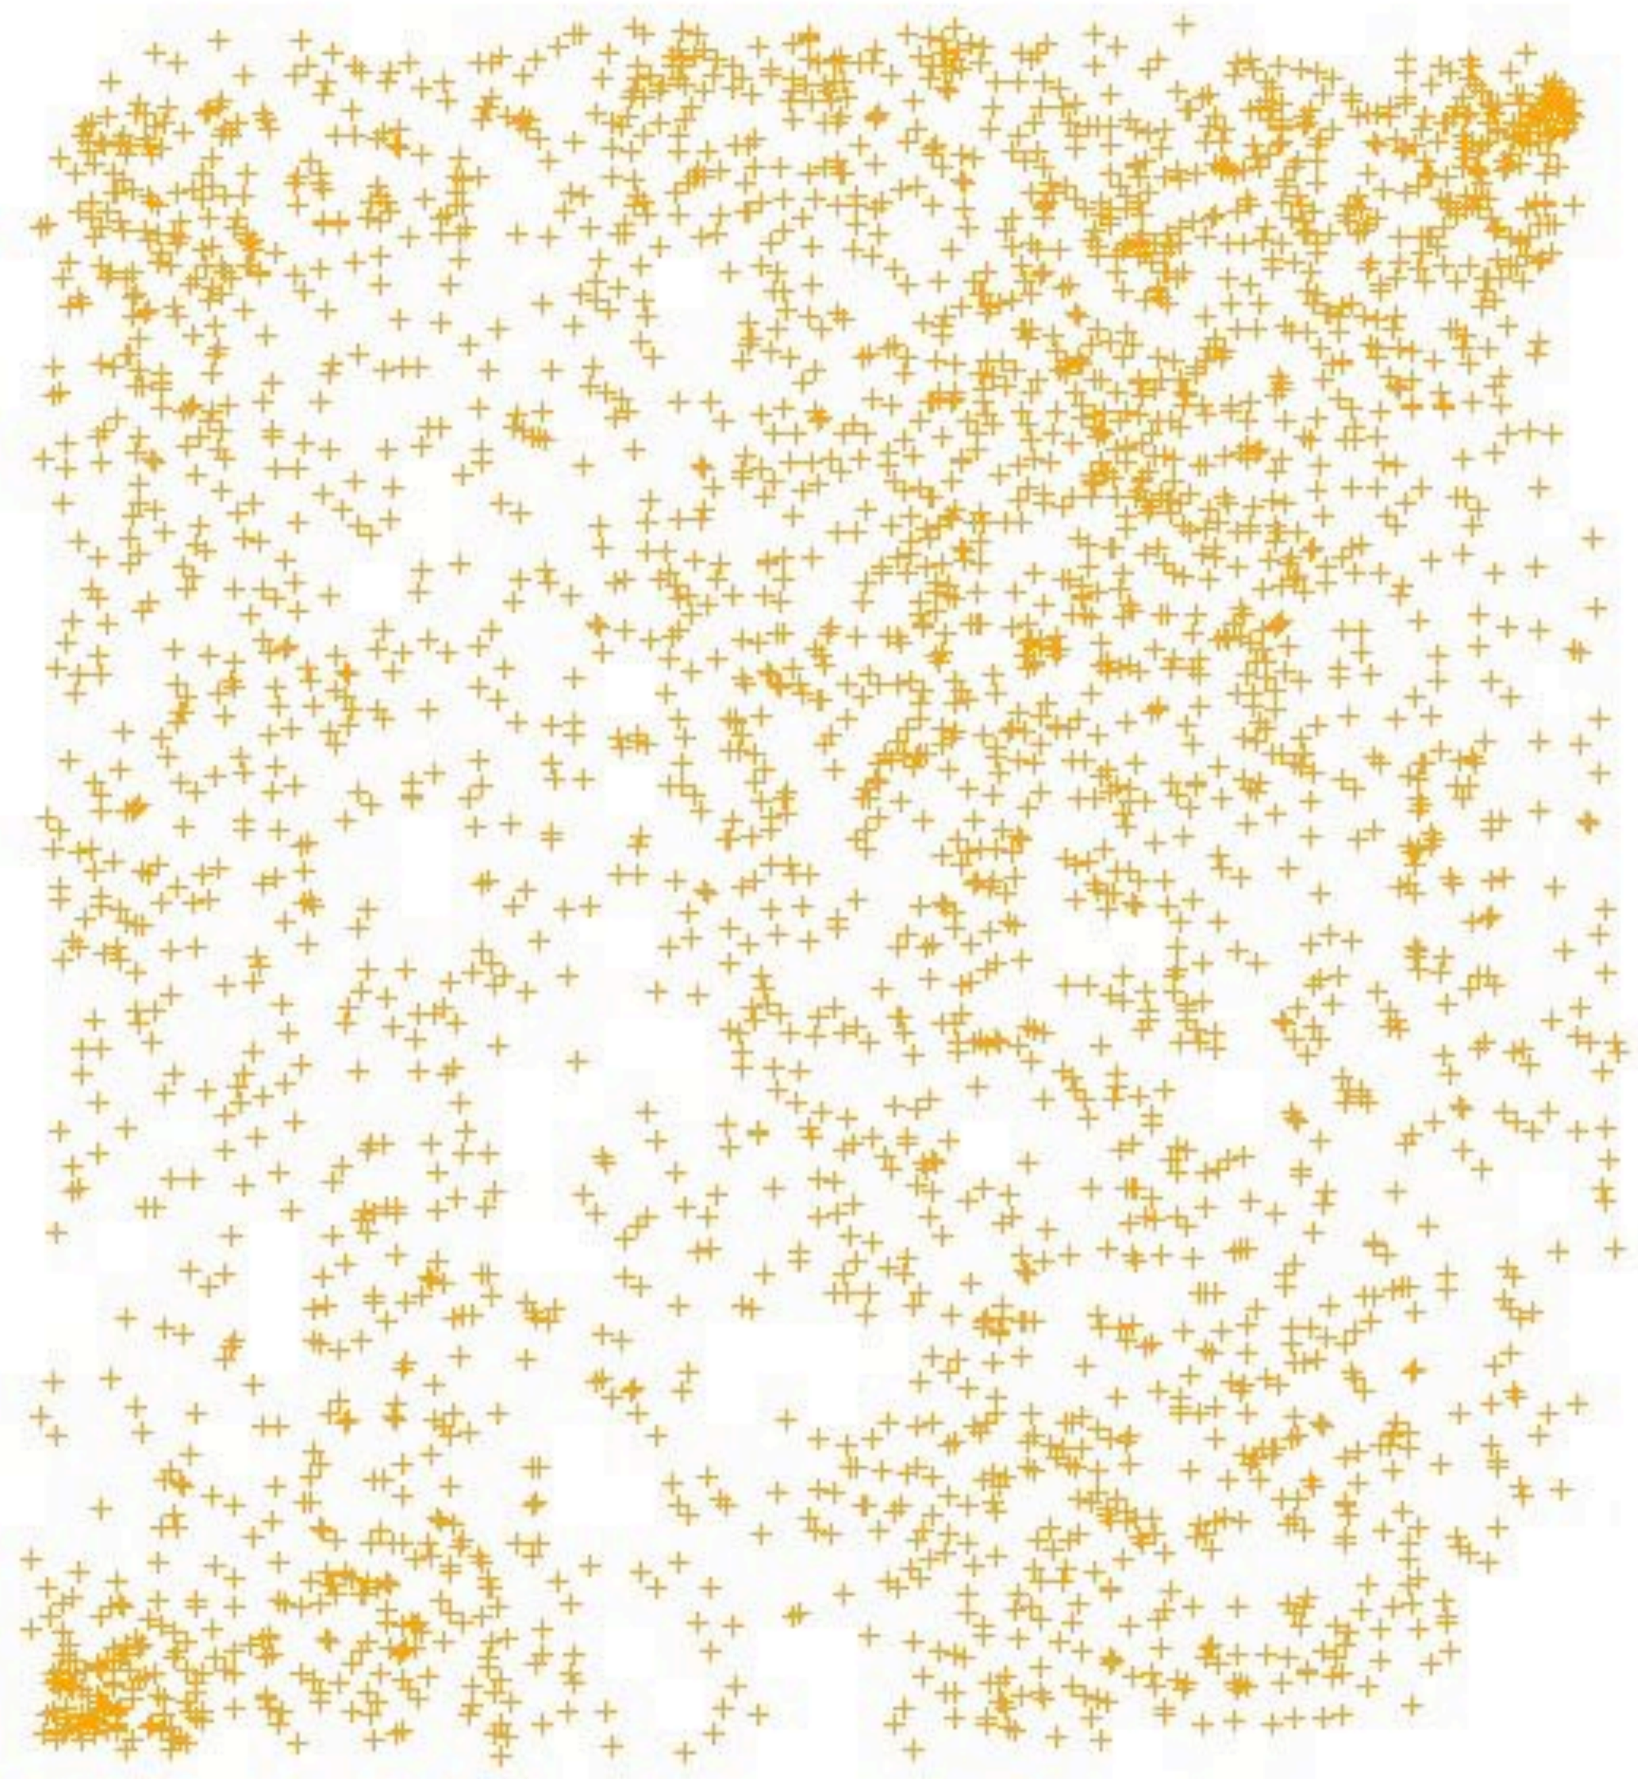

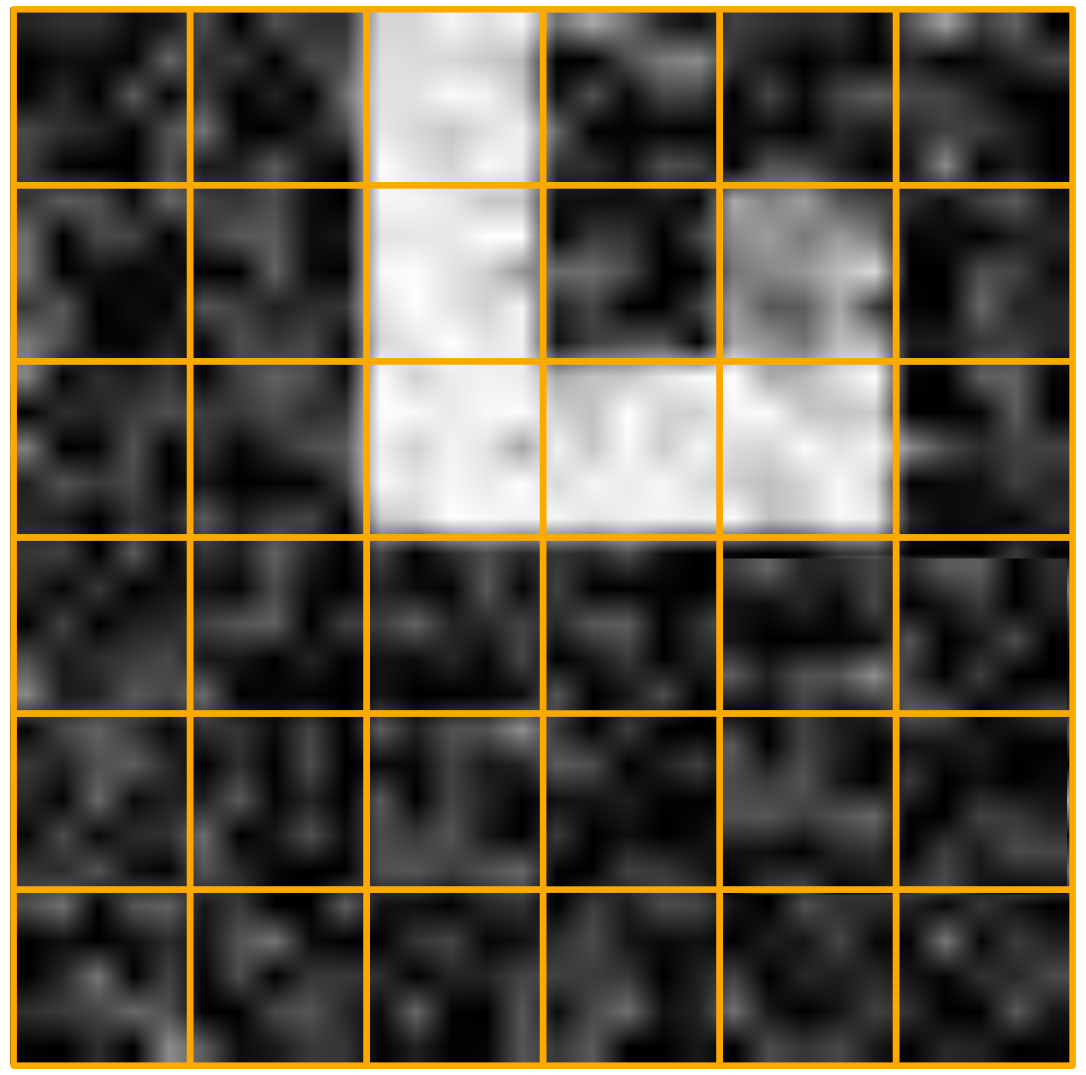

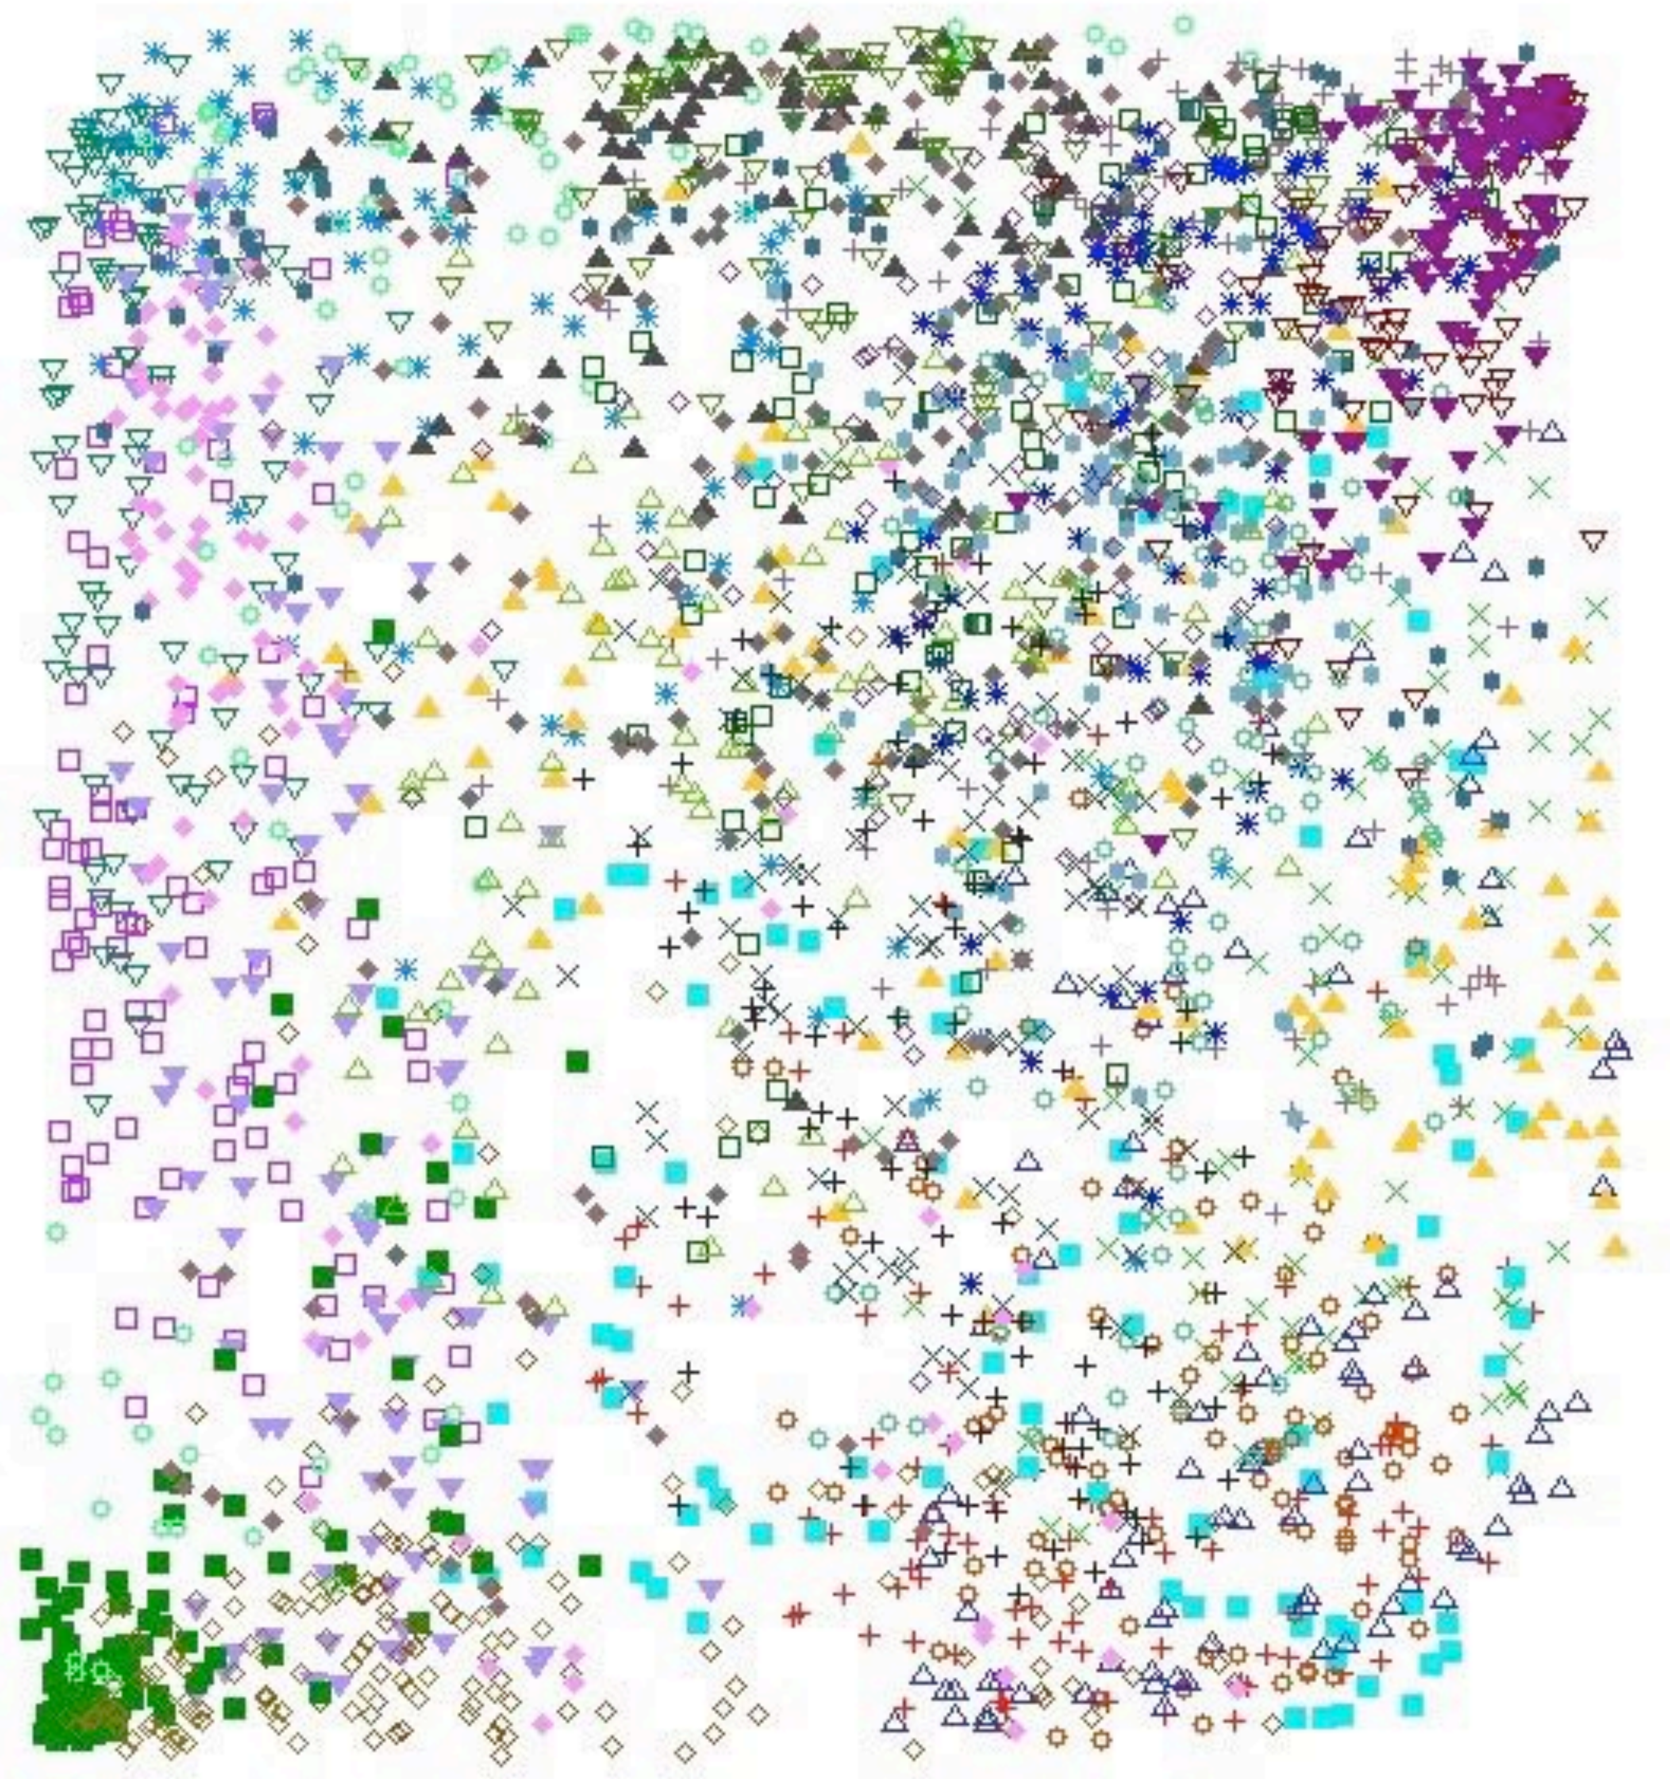

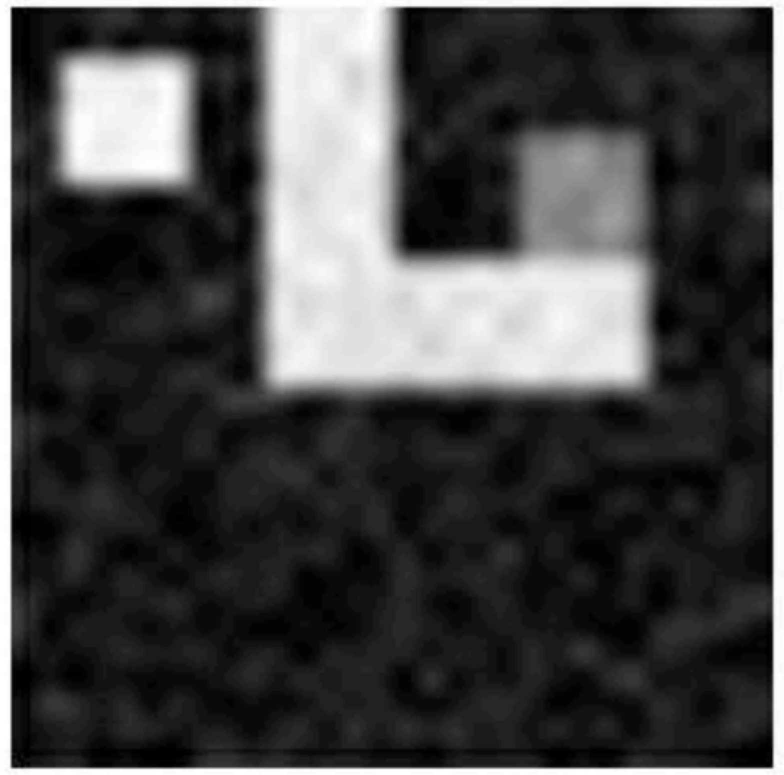

Loop over manually selected test images

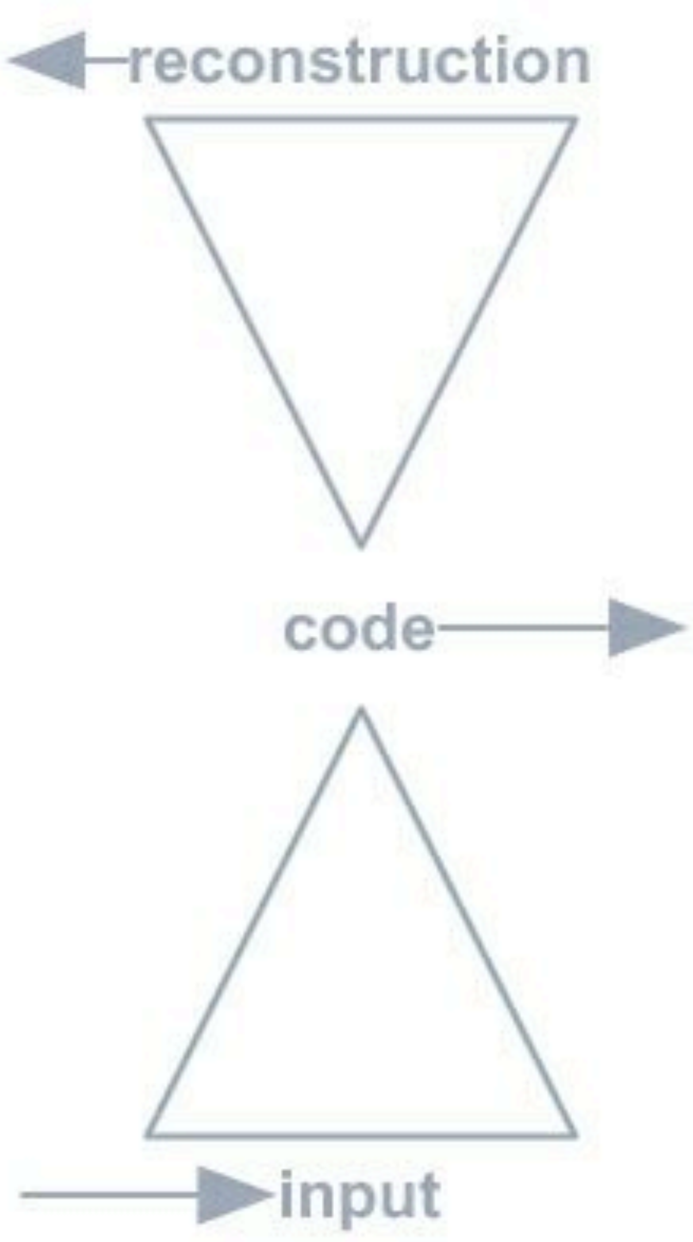

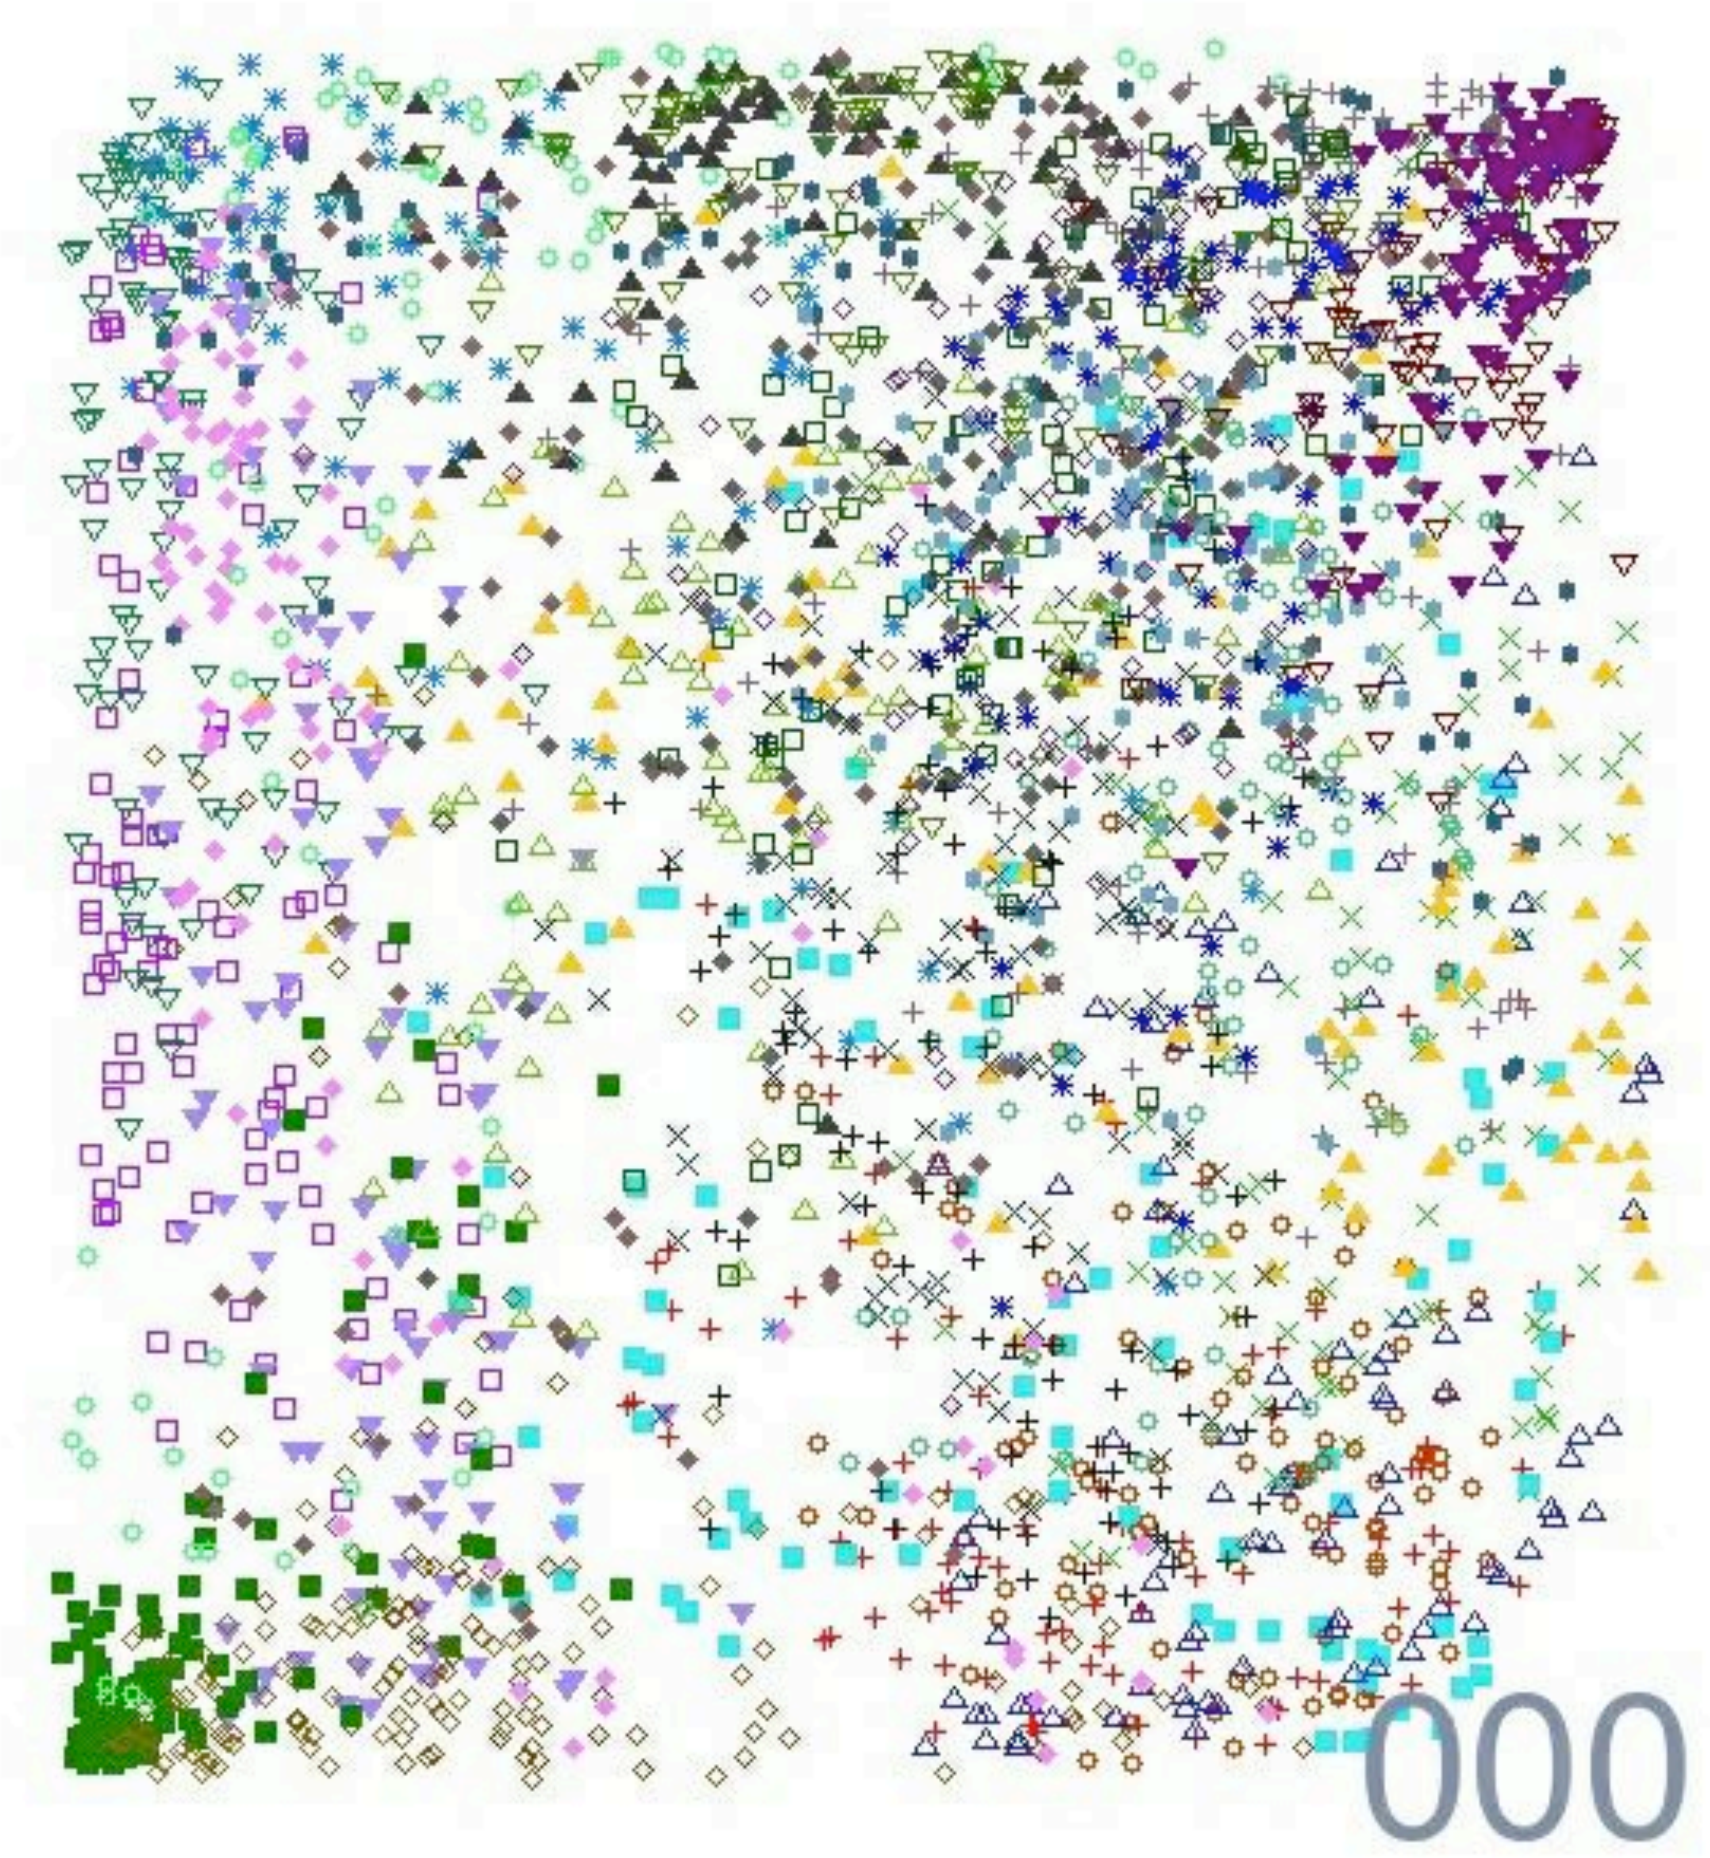

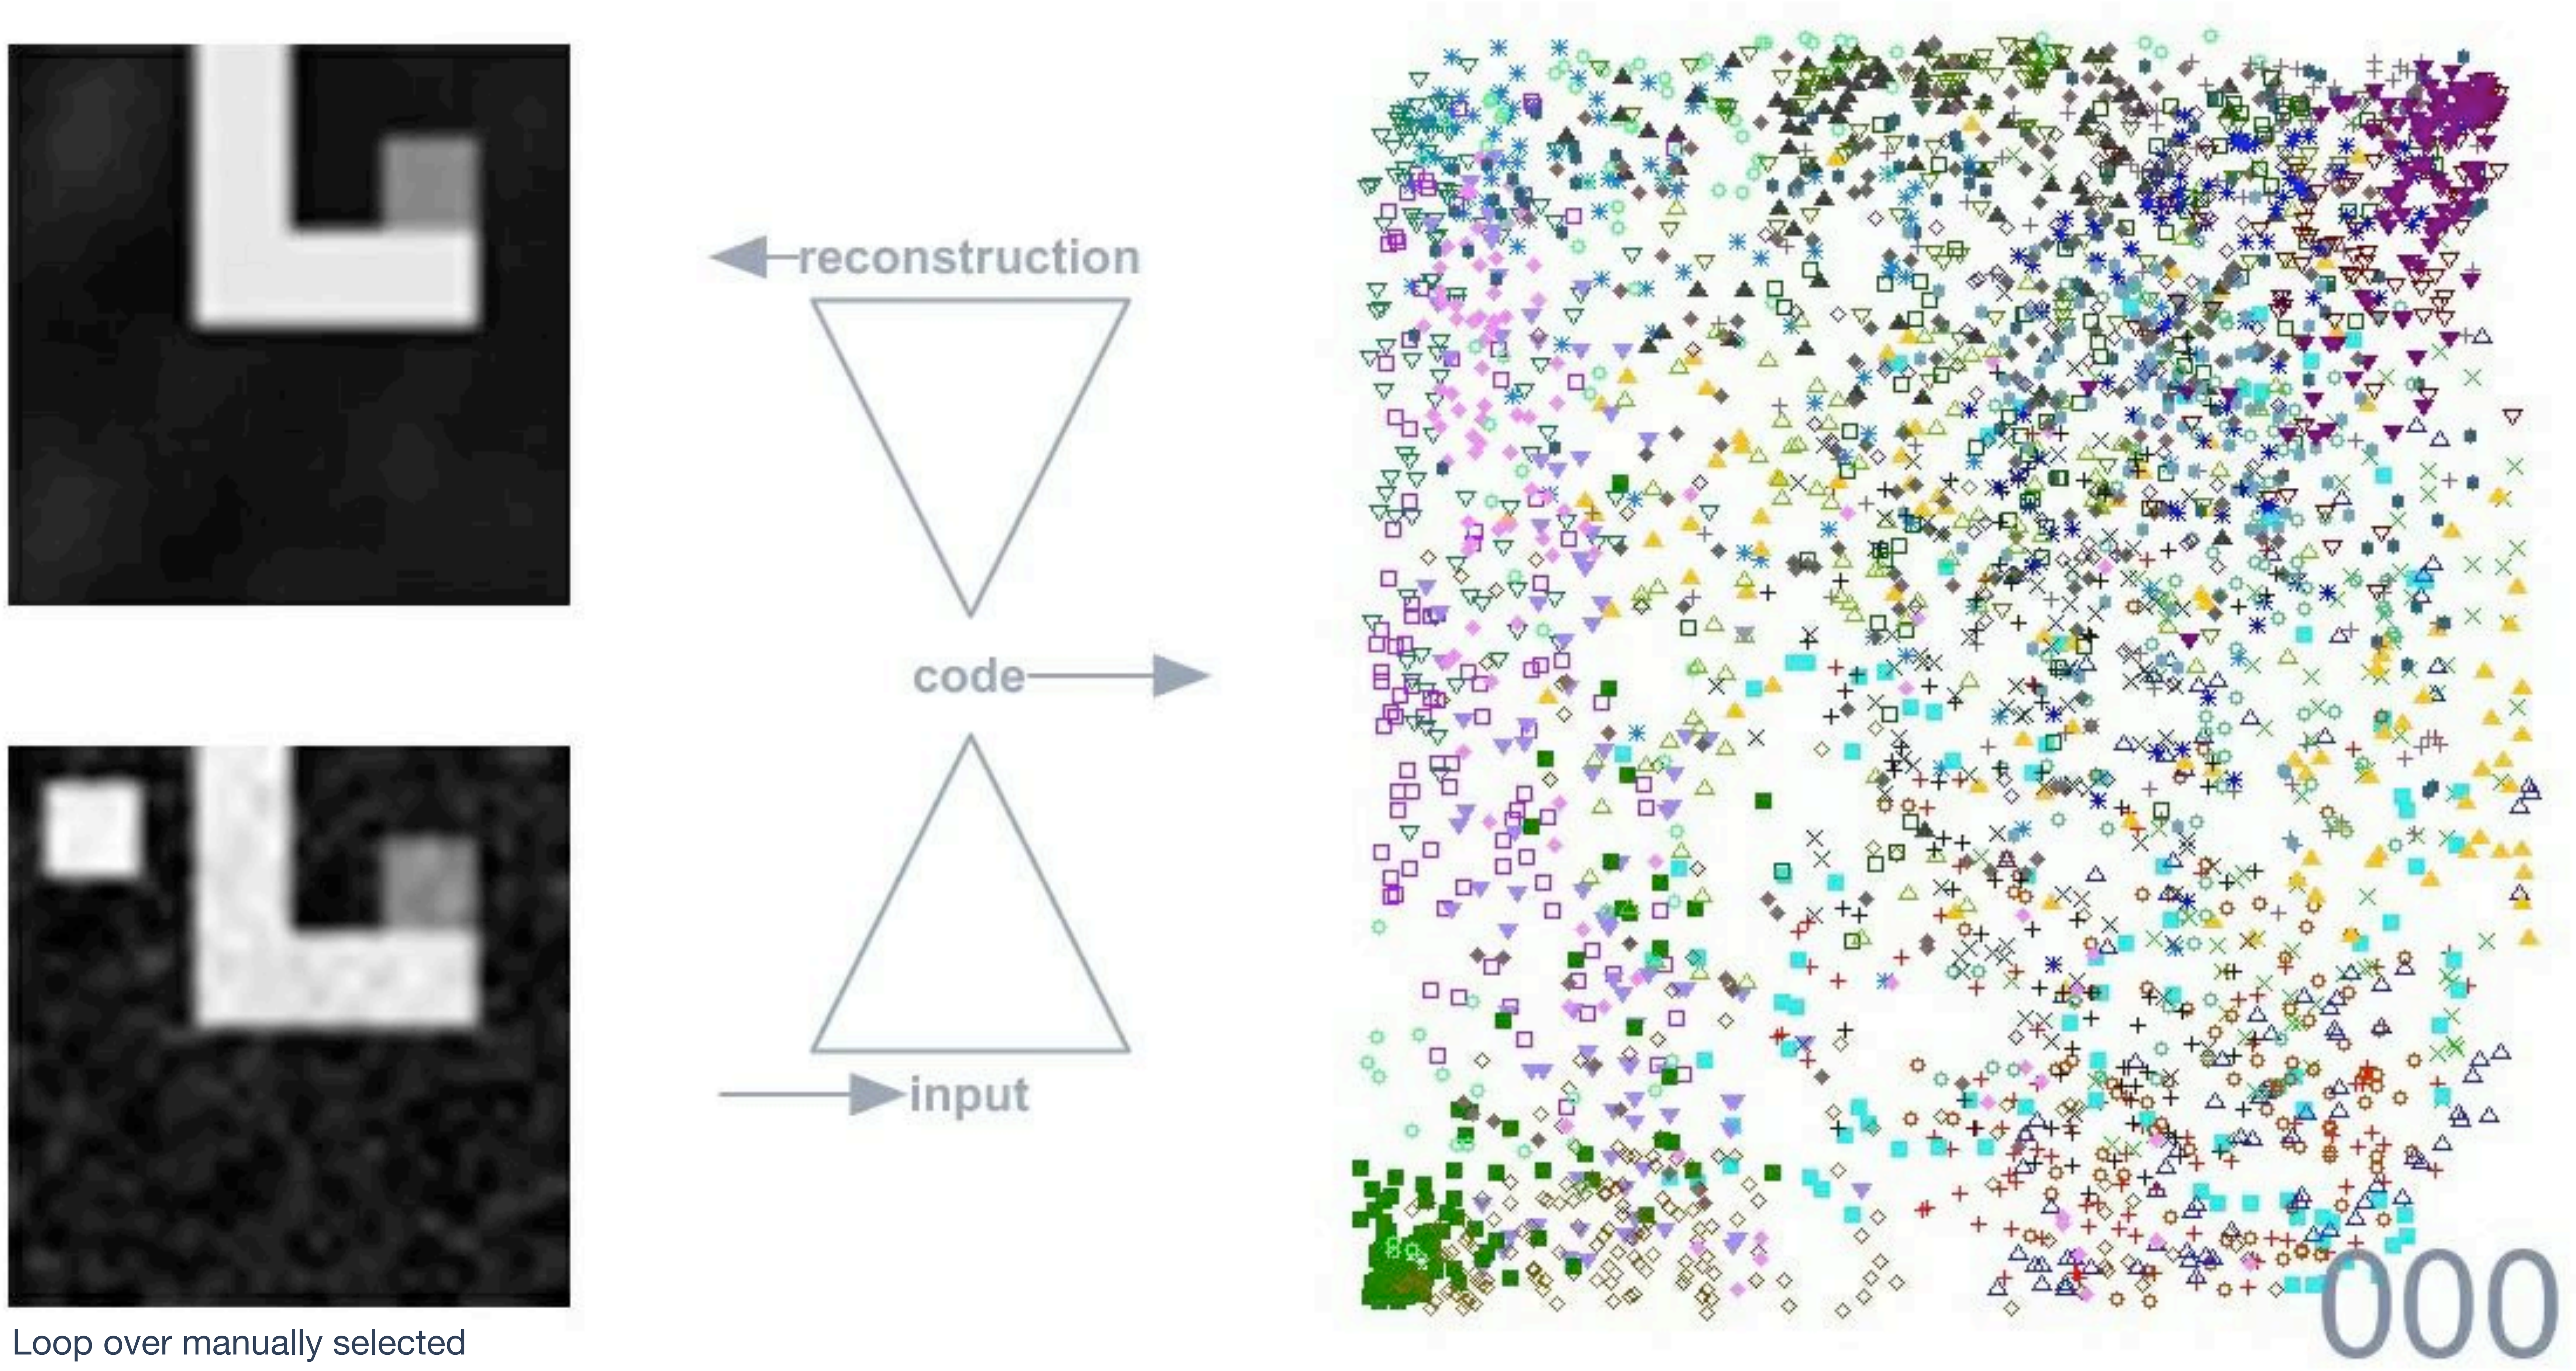

Loop over manually selected
test images

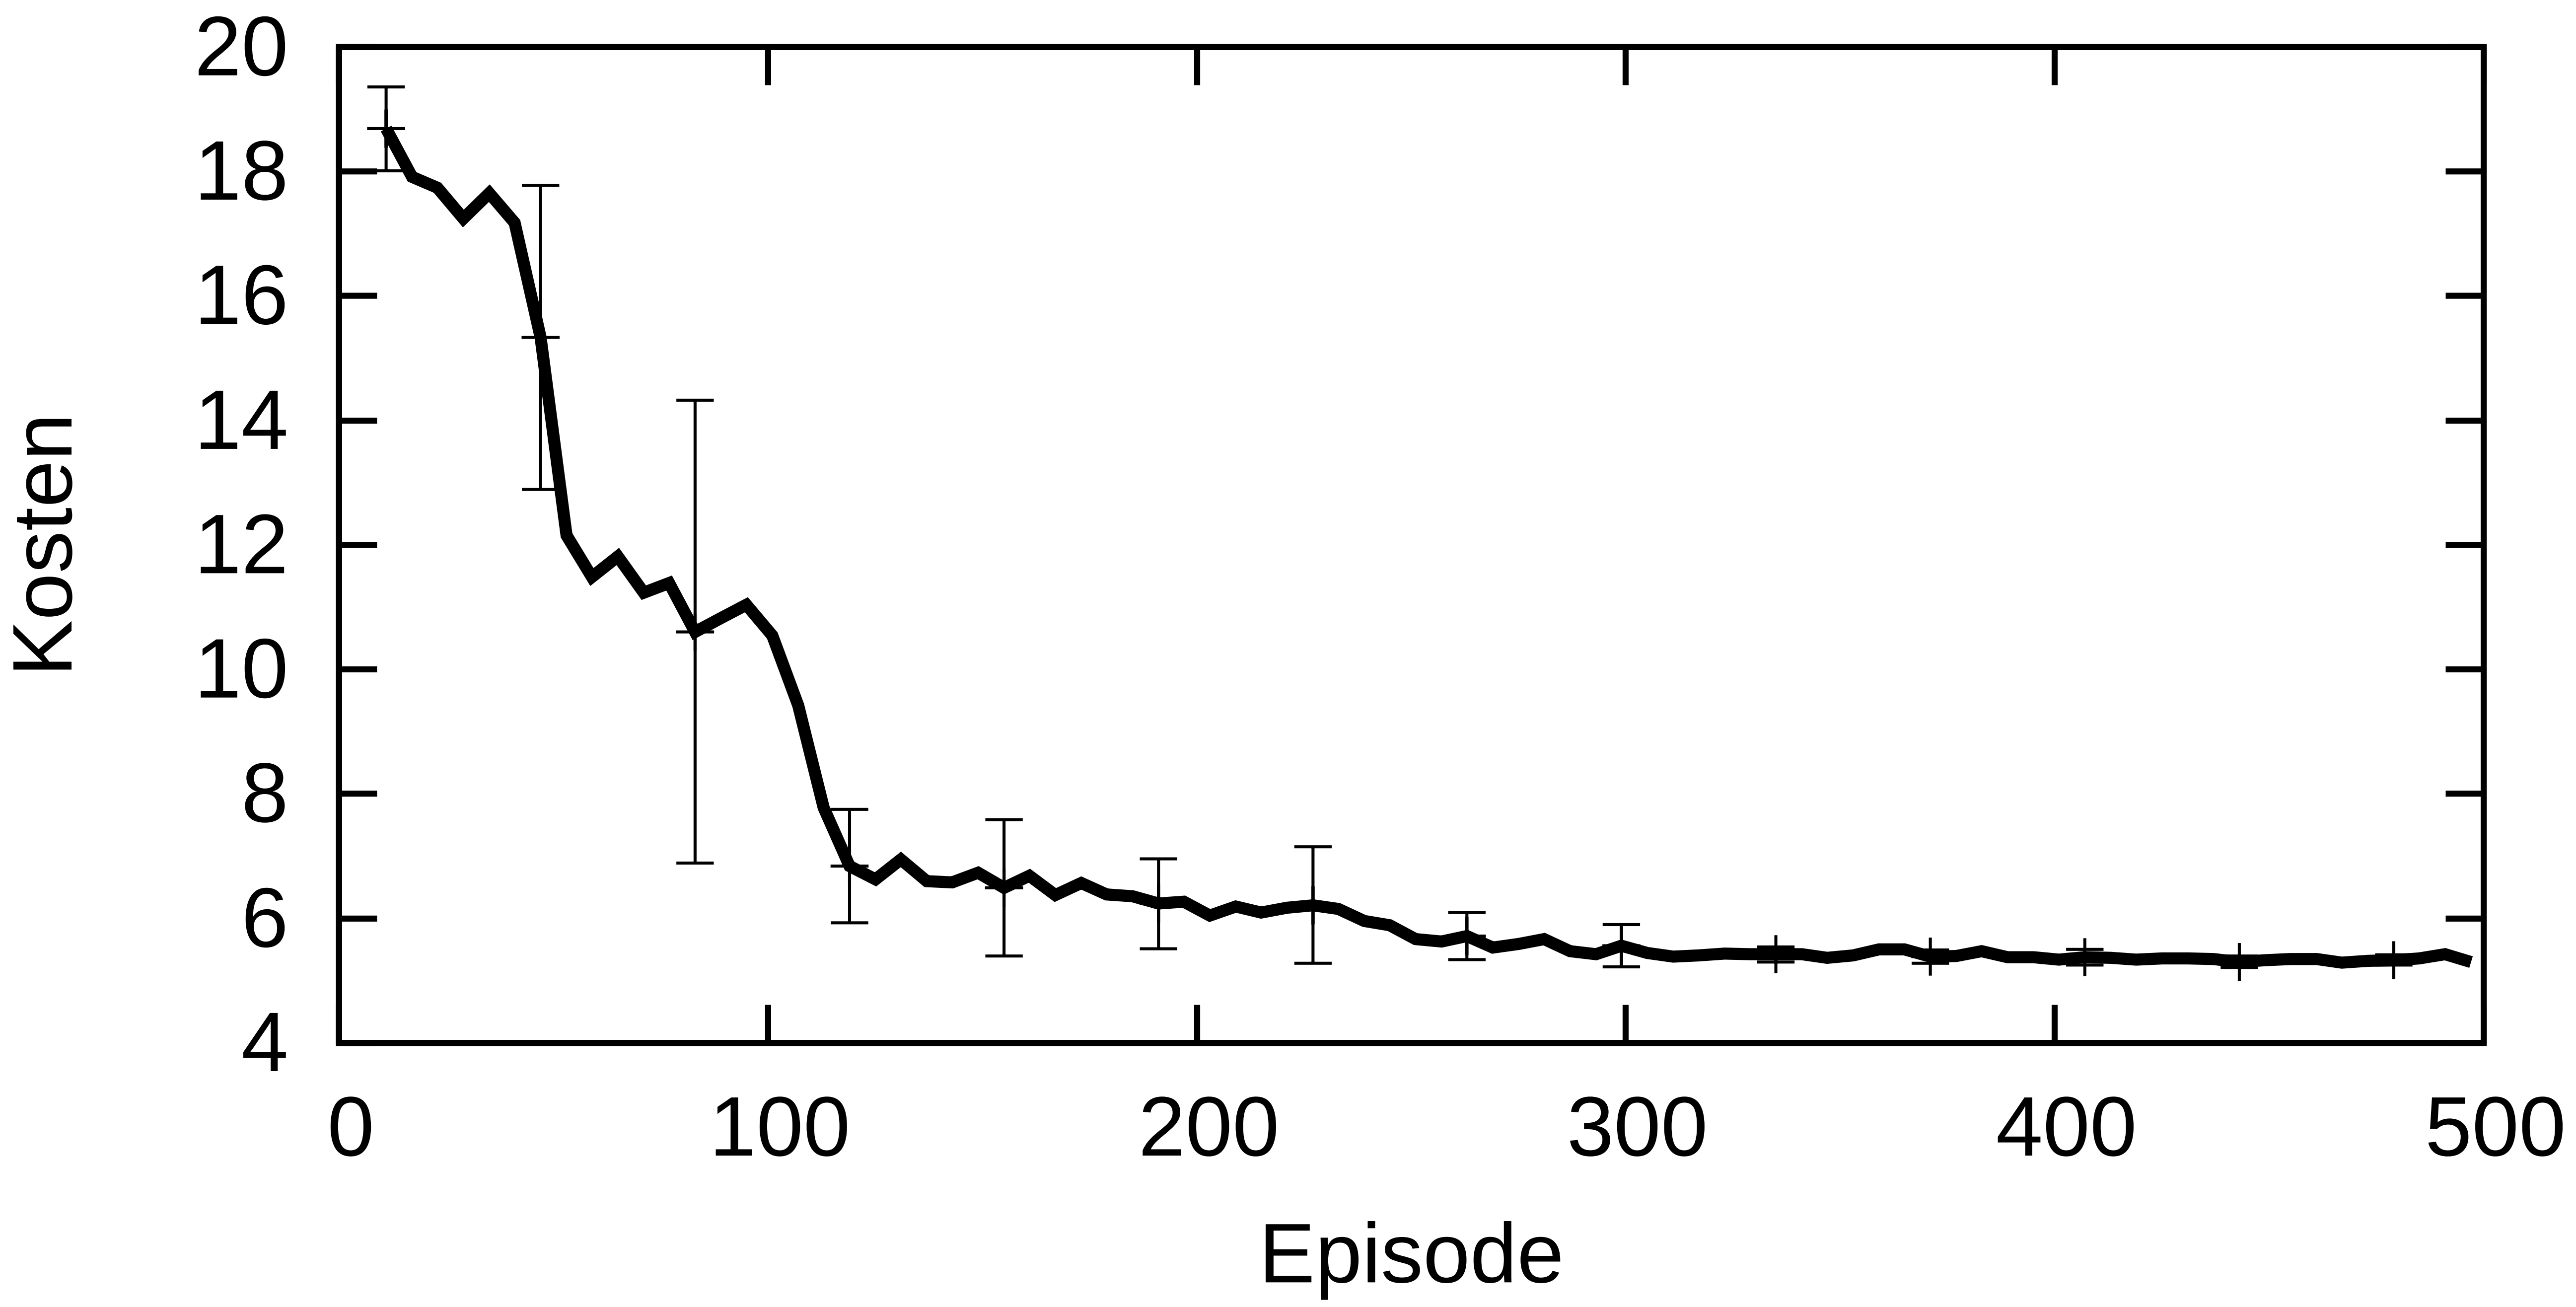

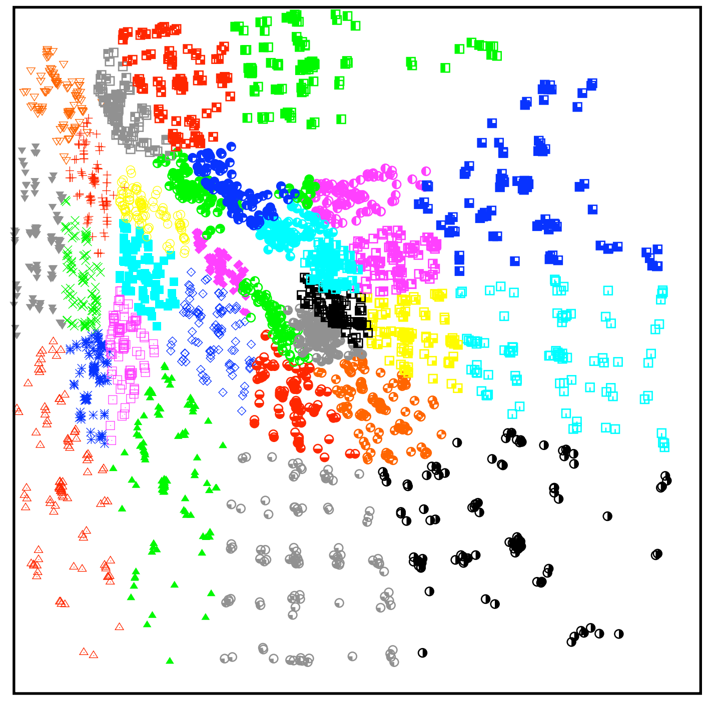

Learned feature space

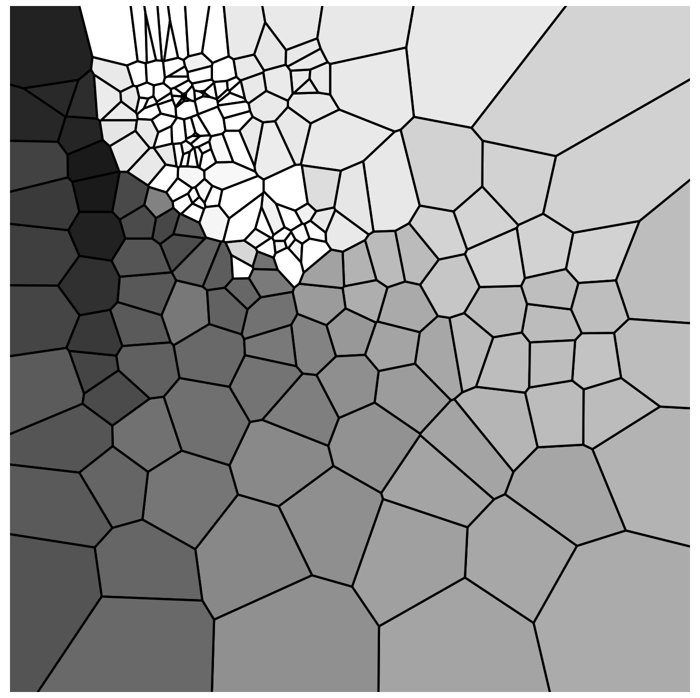

Learned value function in feature space

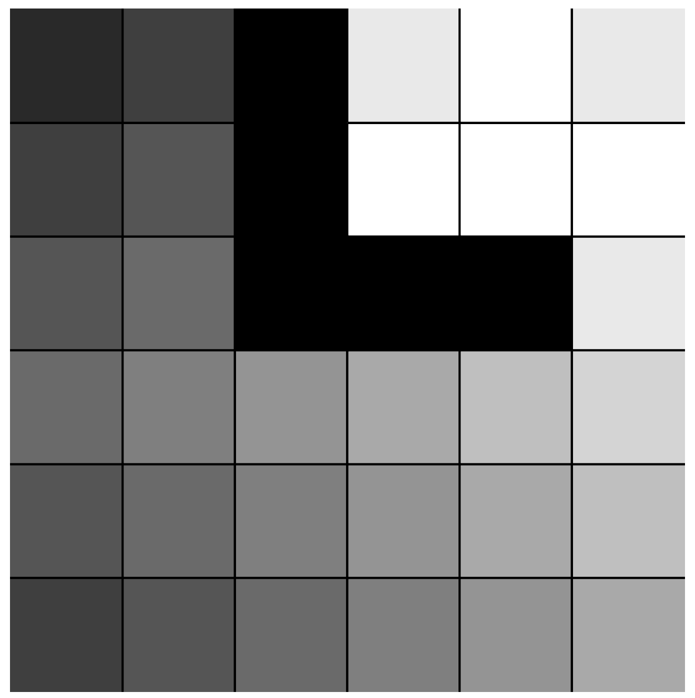

Learned value function in state space

- On average, loses less than 1/10 of a step in comparison with the optimal strategy.


High speed curve
danger
danger

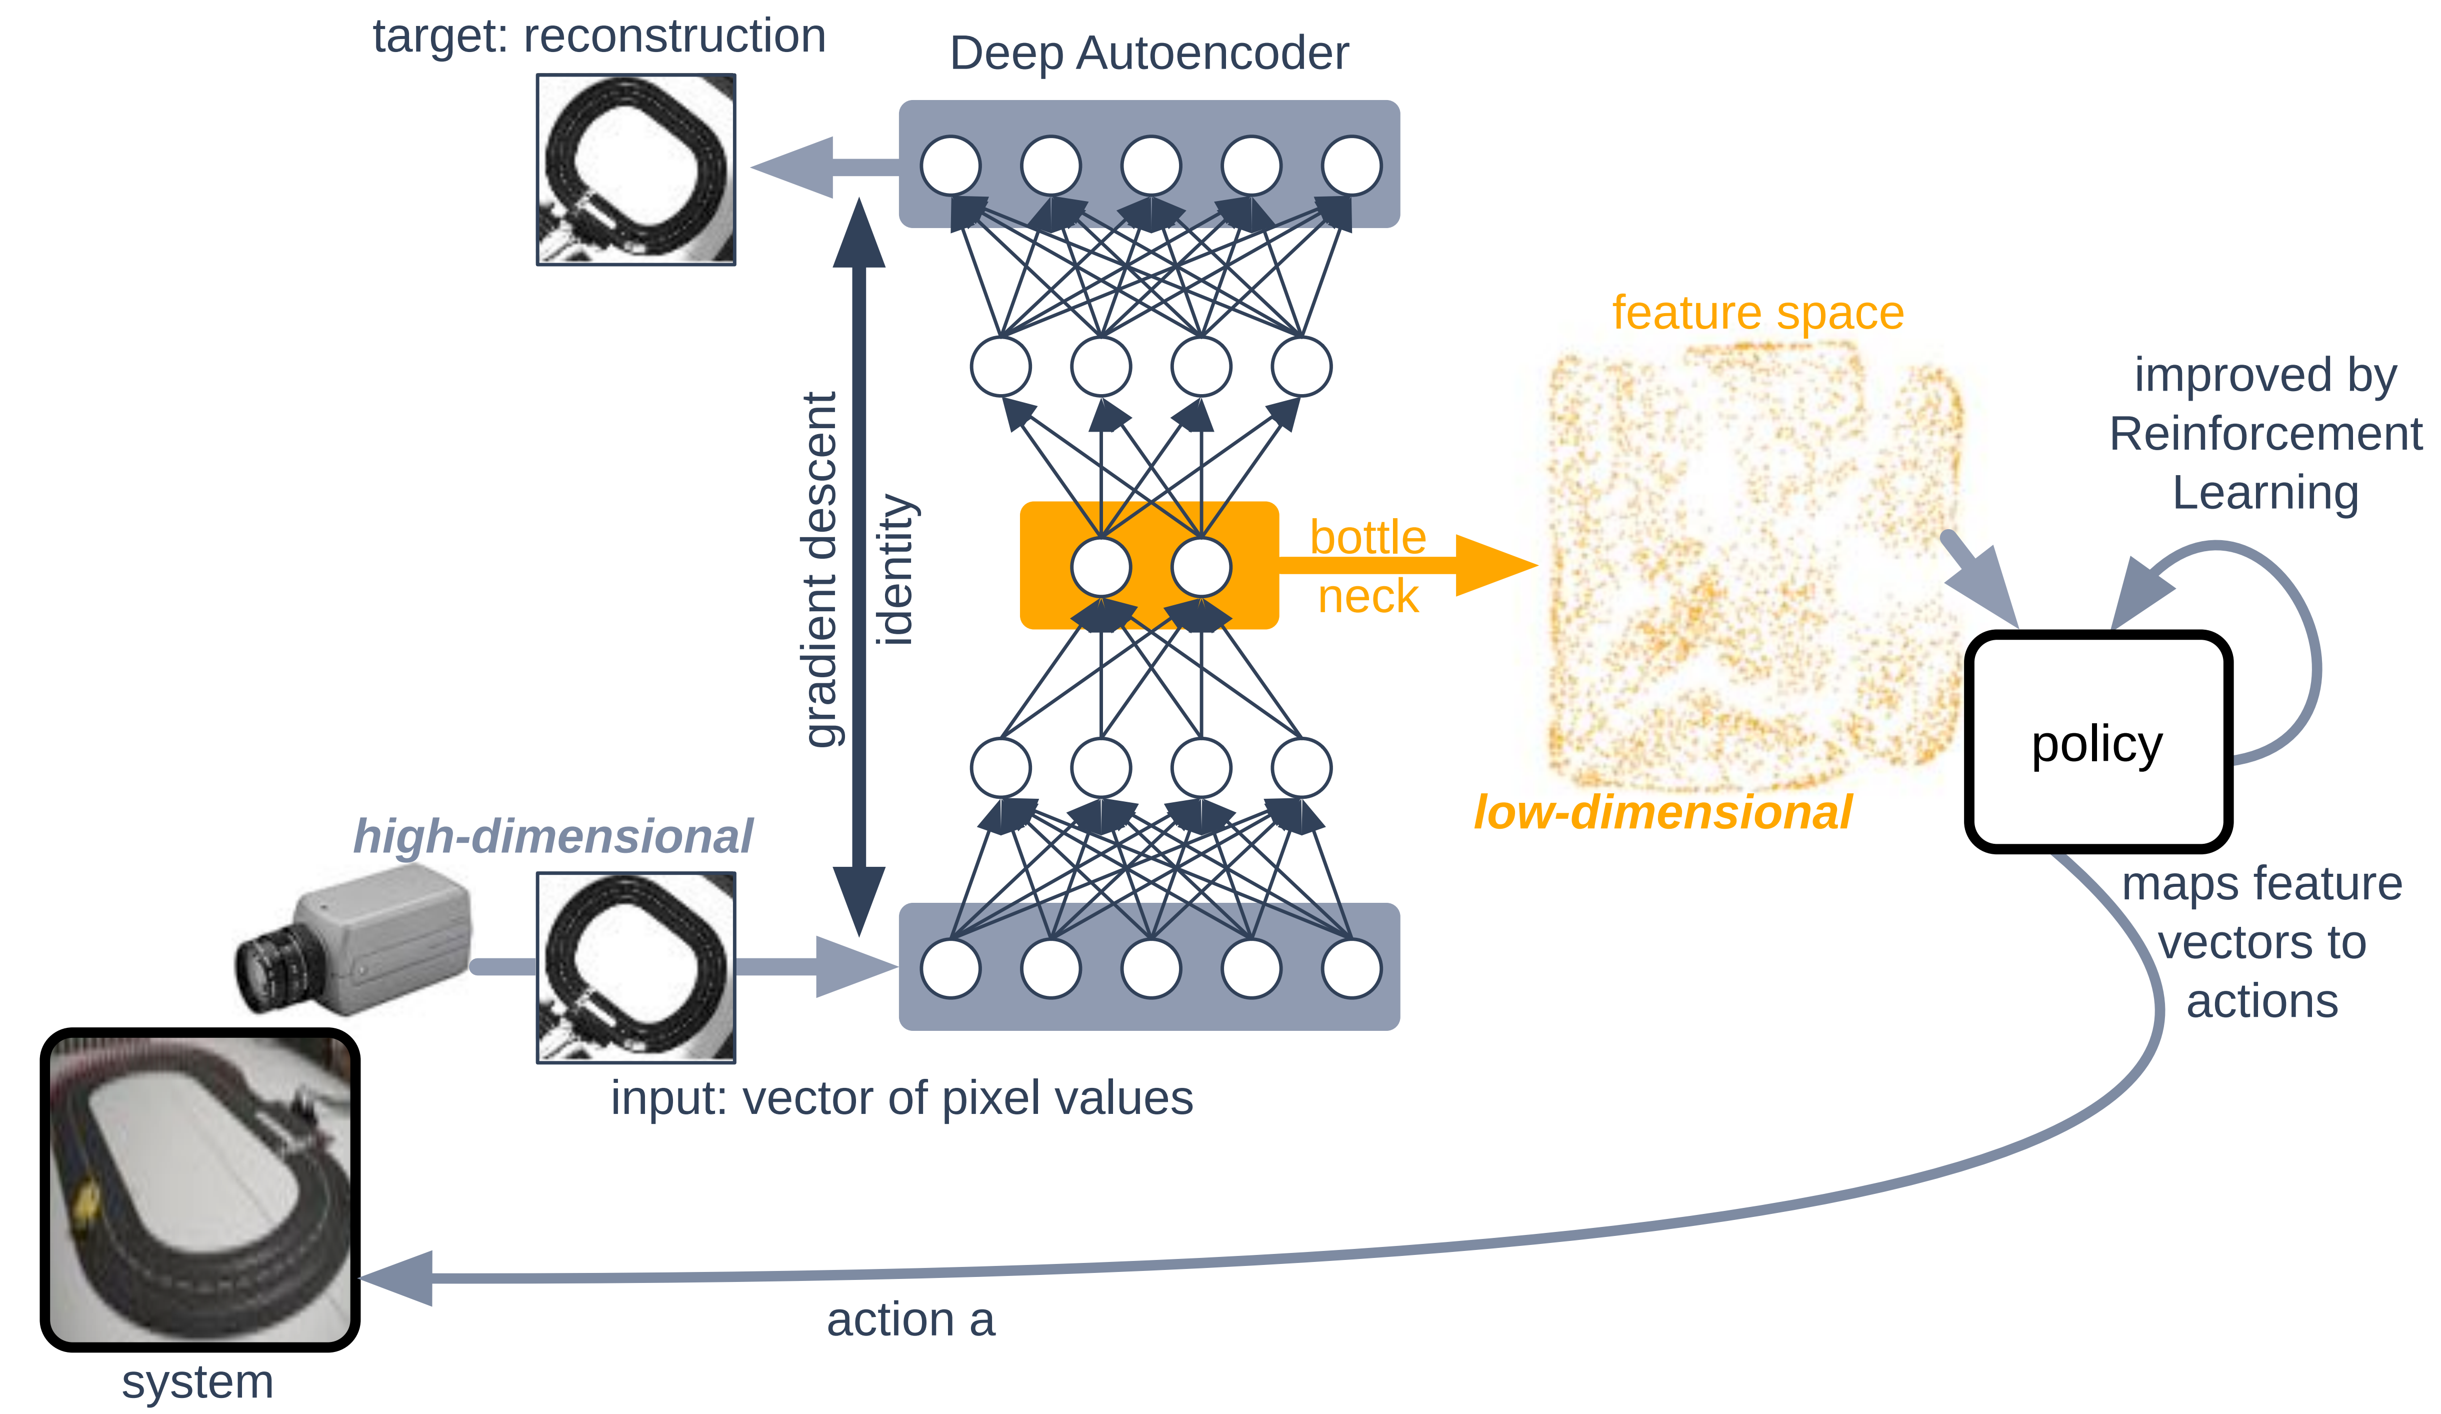
target: reconstruction
Deep Autoencoder
gradient descent
identity
bottle
neck
feature space
low-dimensional
improved by
Reinforcement
Learning
policy
maps feature
vectors to
actions
high-dimensional
input: vector of pixel values
system
action a

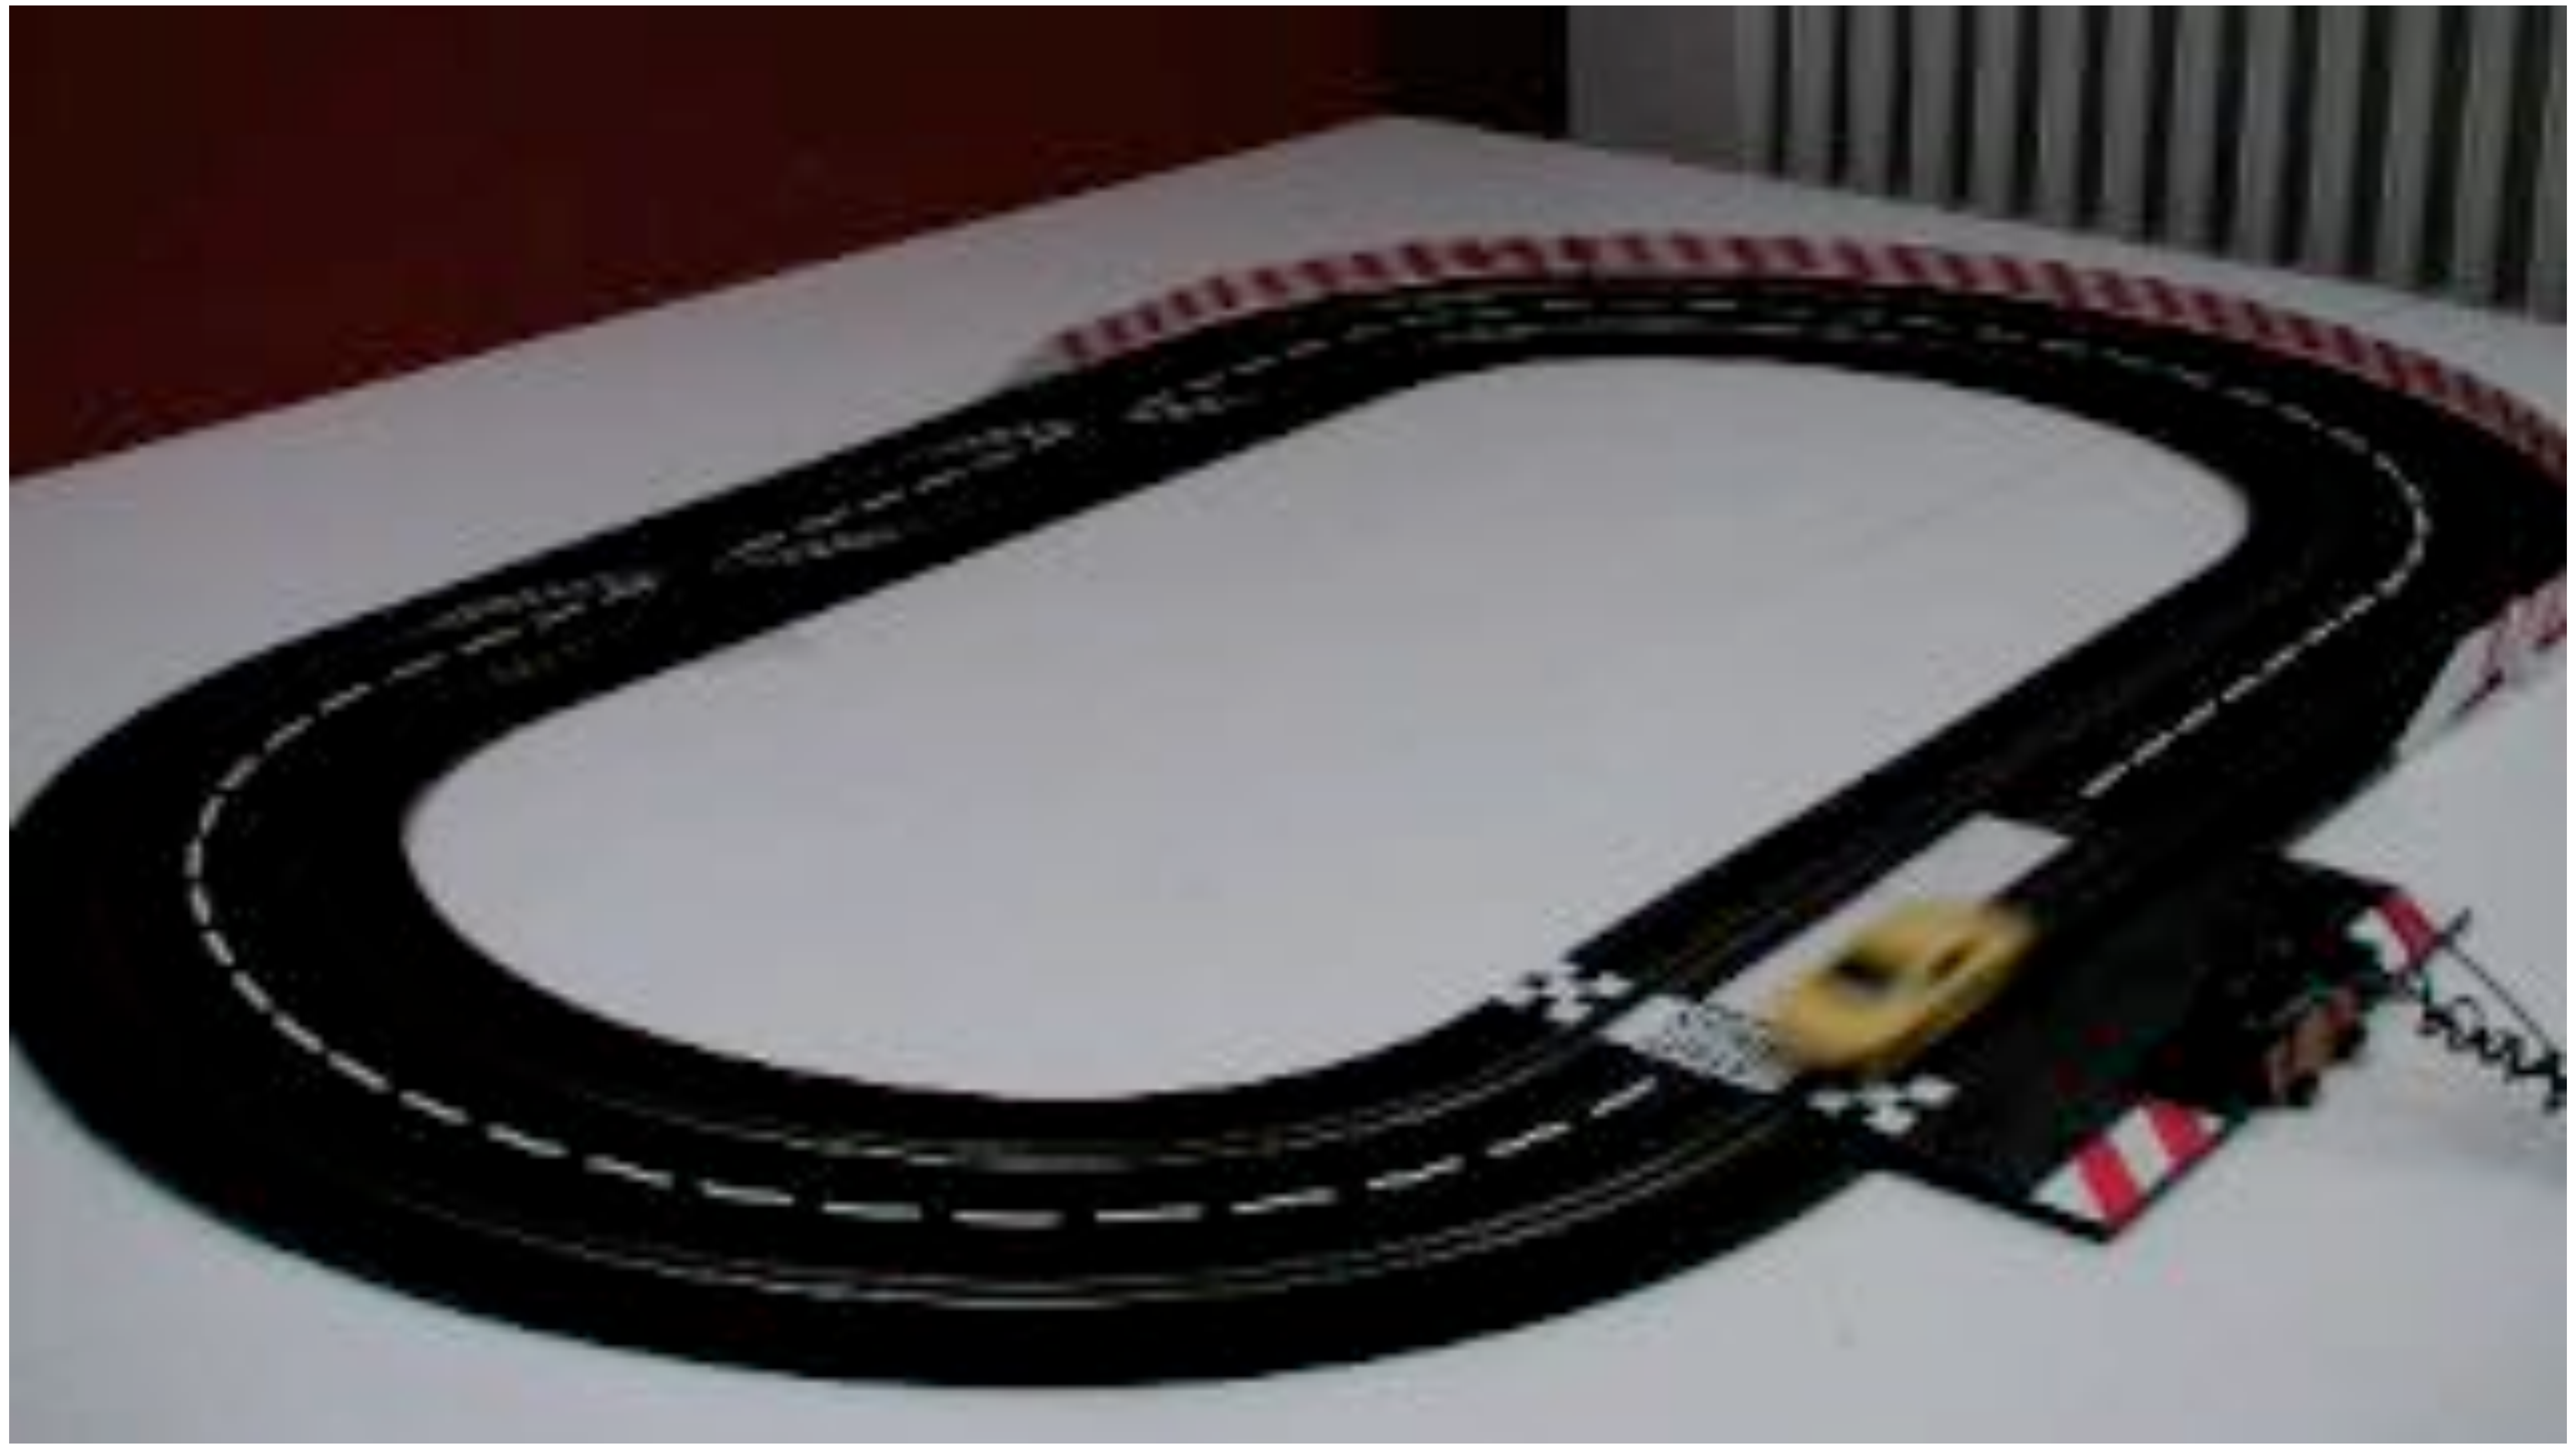

Slot Car in Latent Feature Space

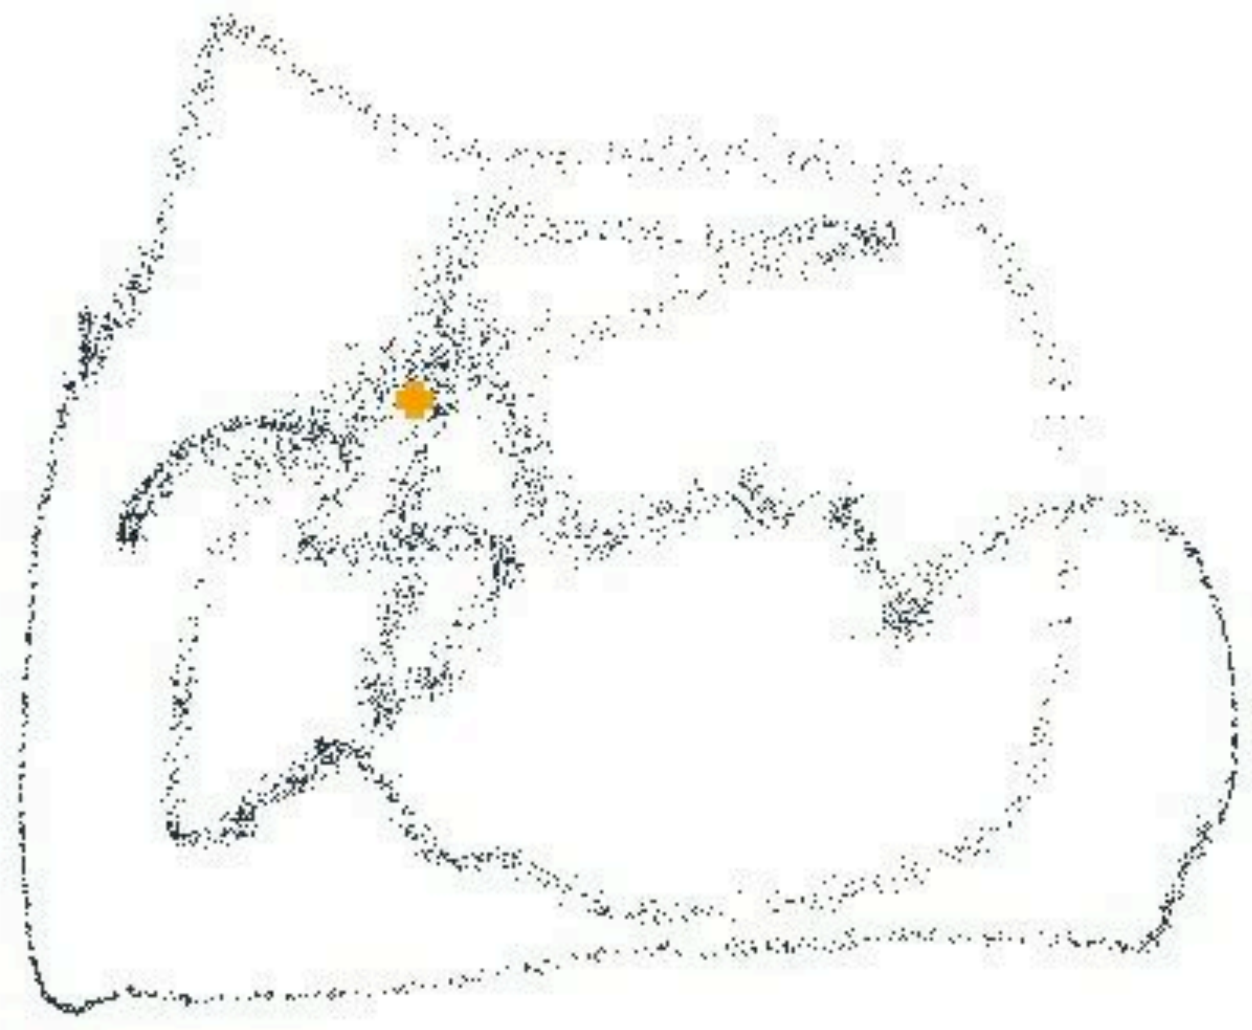

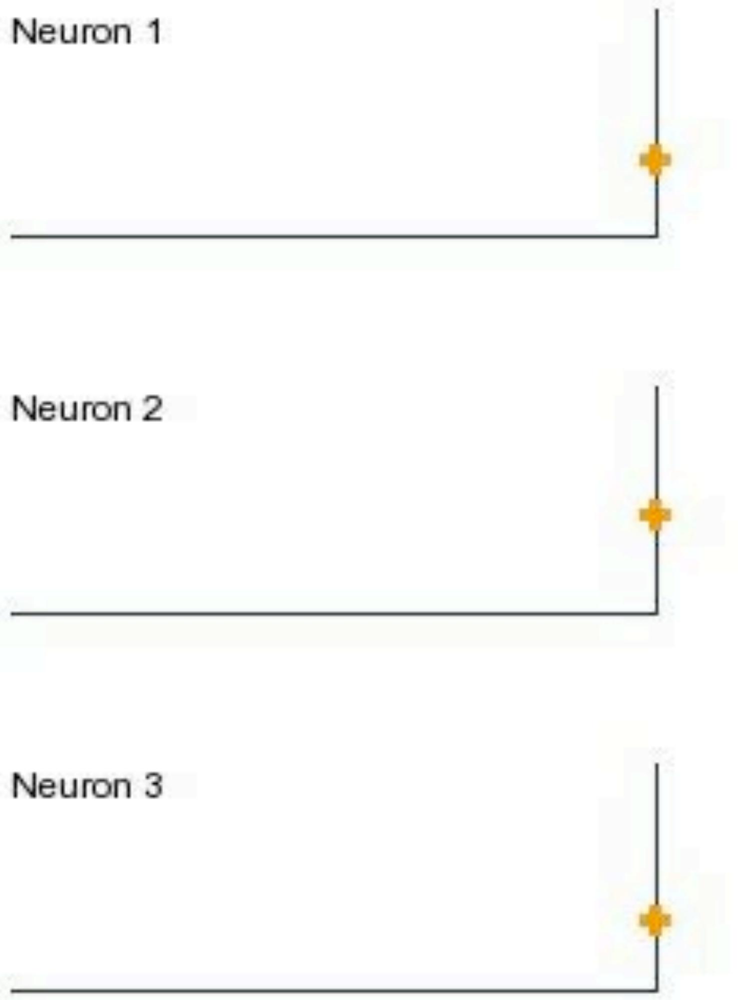

Slow motion with 1/4 of real velocity

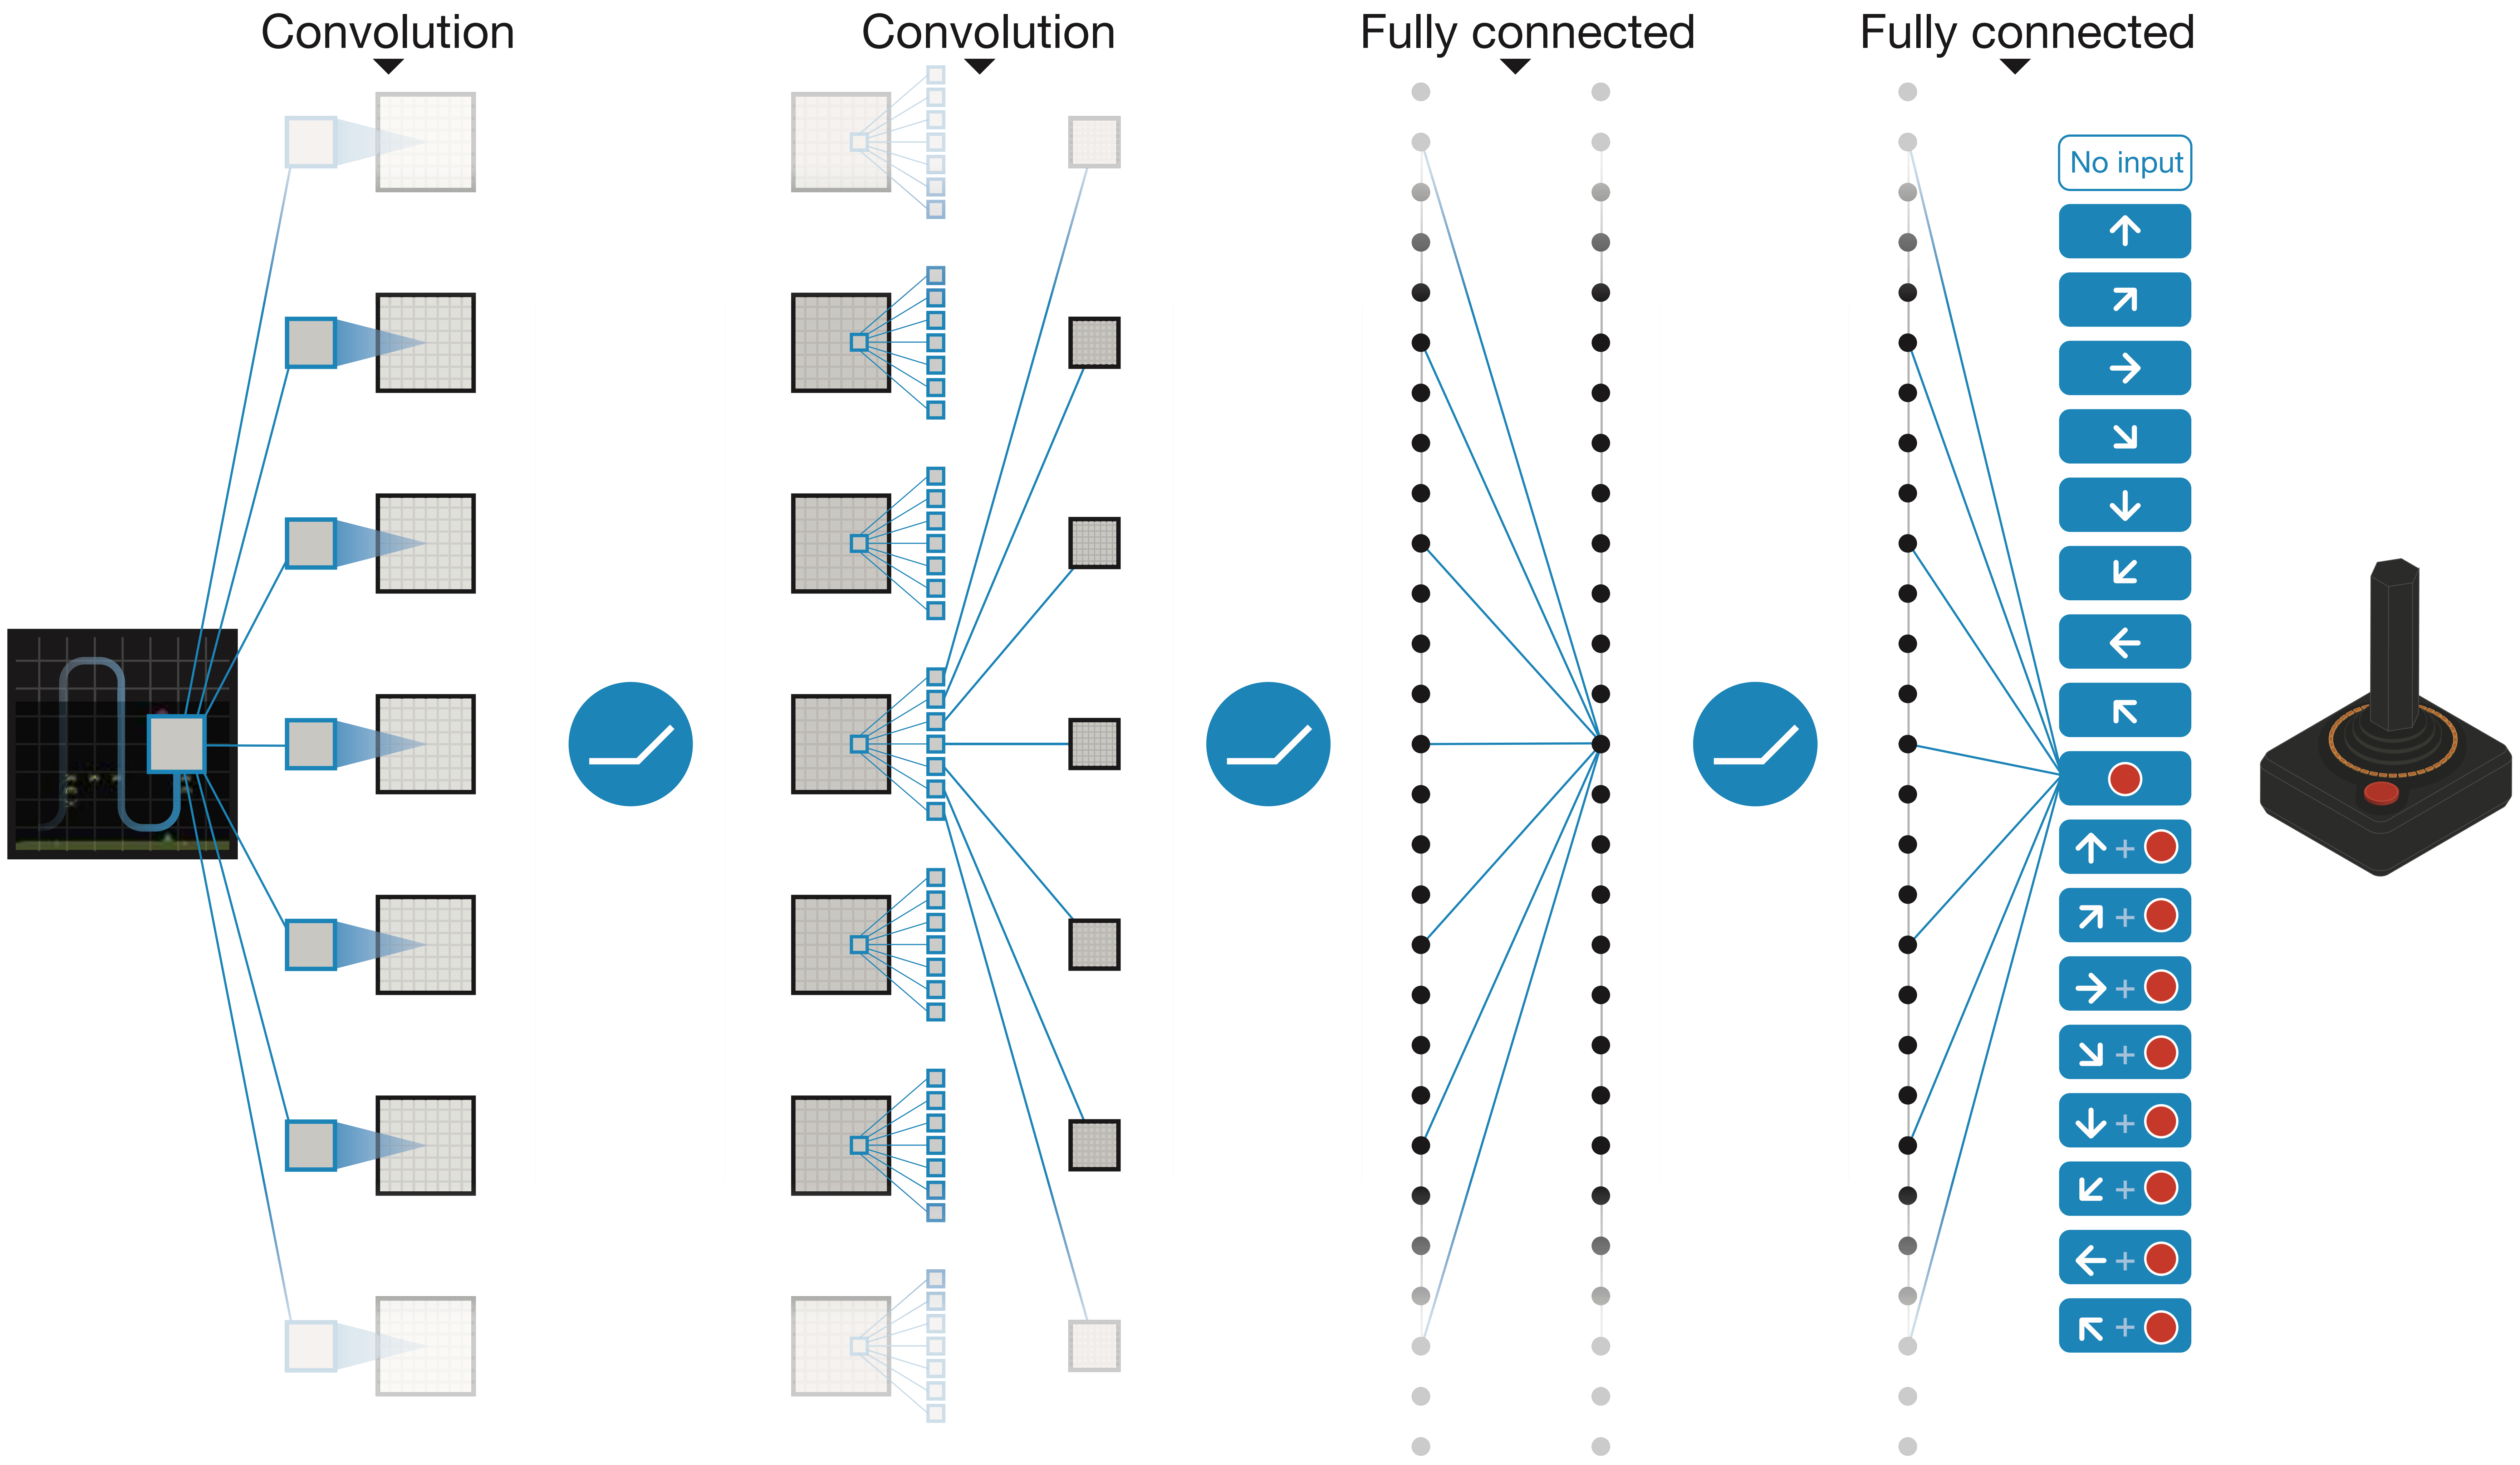

LETTER

doi:10.1038/nature14236

Human-level control through deep reinforcement learning

Volodymyr Mnih[1]*, Koray Kavukcuoglu[1]*, David Silver[1]*, Andrei A. Rusu[1], Joel Veness[1], Marc G. Bellemare[1], Alex Graves[1], Martin Riedmiller[1], Andreas K. Fidjeland[1], Georg Ostrovski[1], Stig Petersen[1], Charles Beattie[1], Amir Sadik[1], Ioannis Antonoglou[1], Helen King[1], Dharshan Kumaran[1], Daan Wierstra[1], Shane Legg[1] & Demis Hassabis[1]

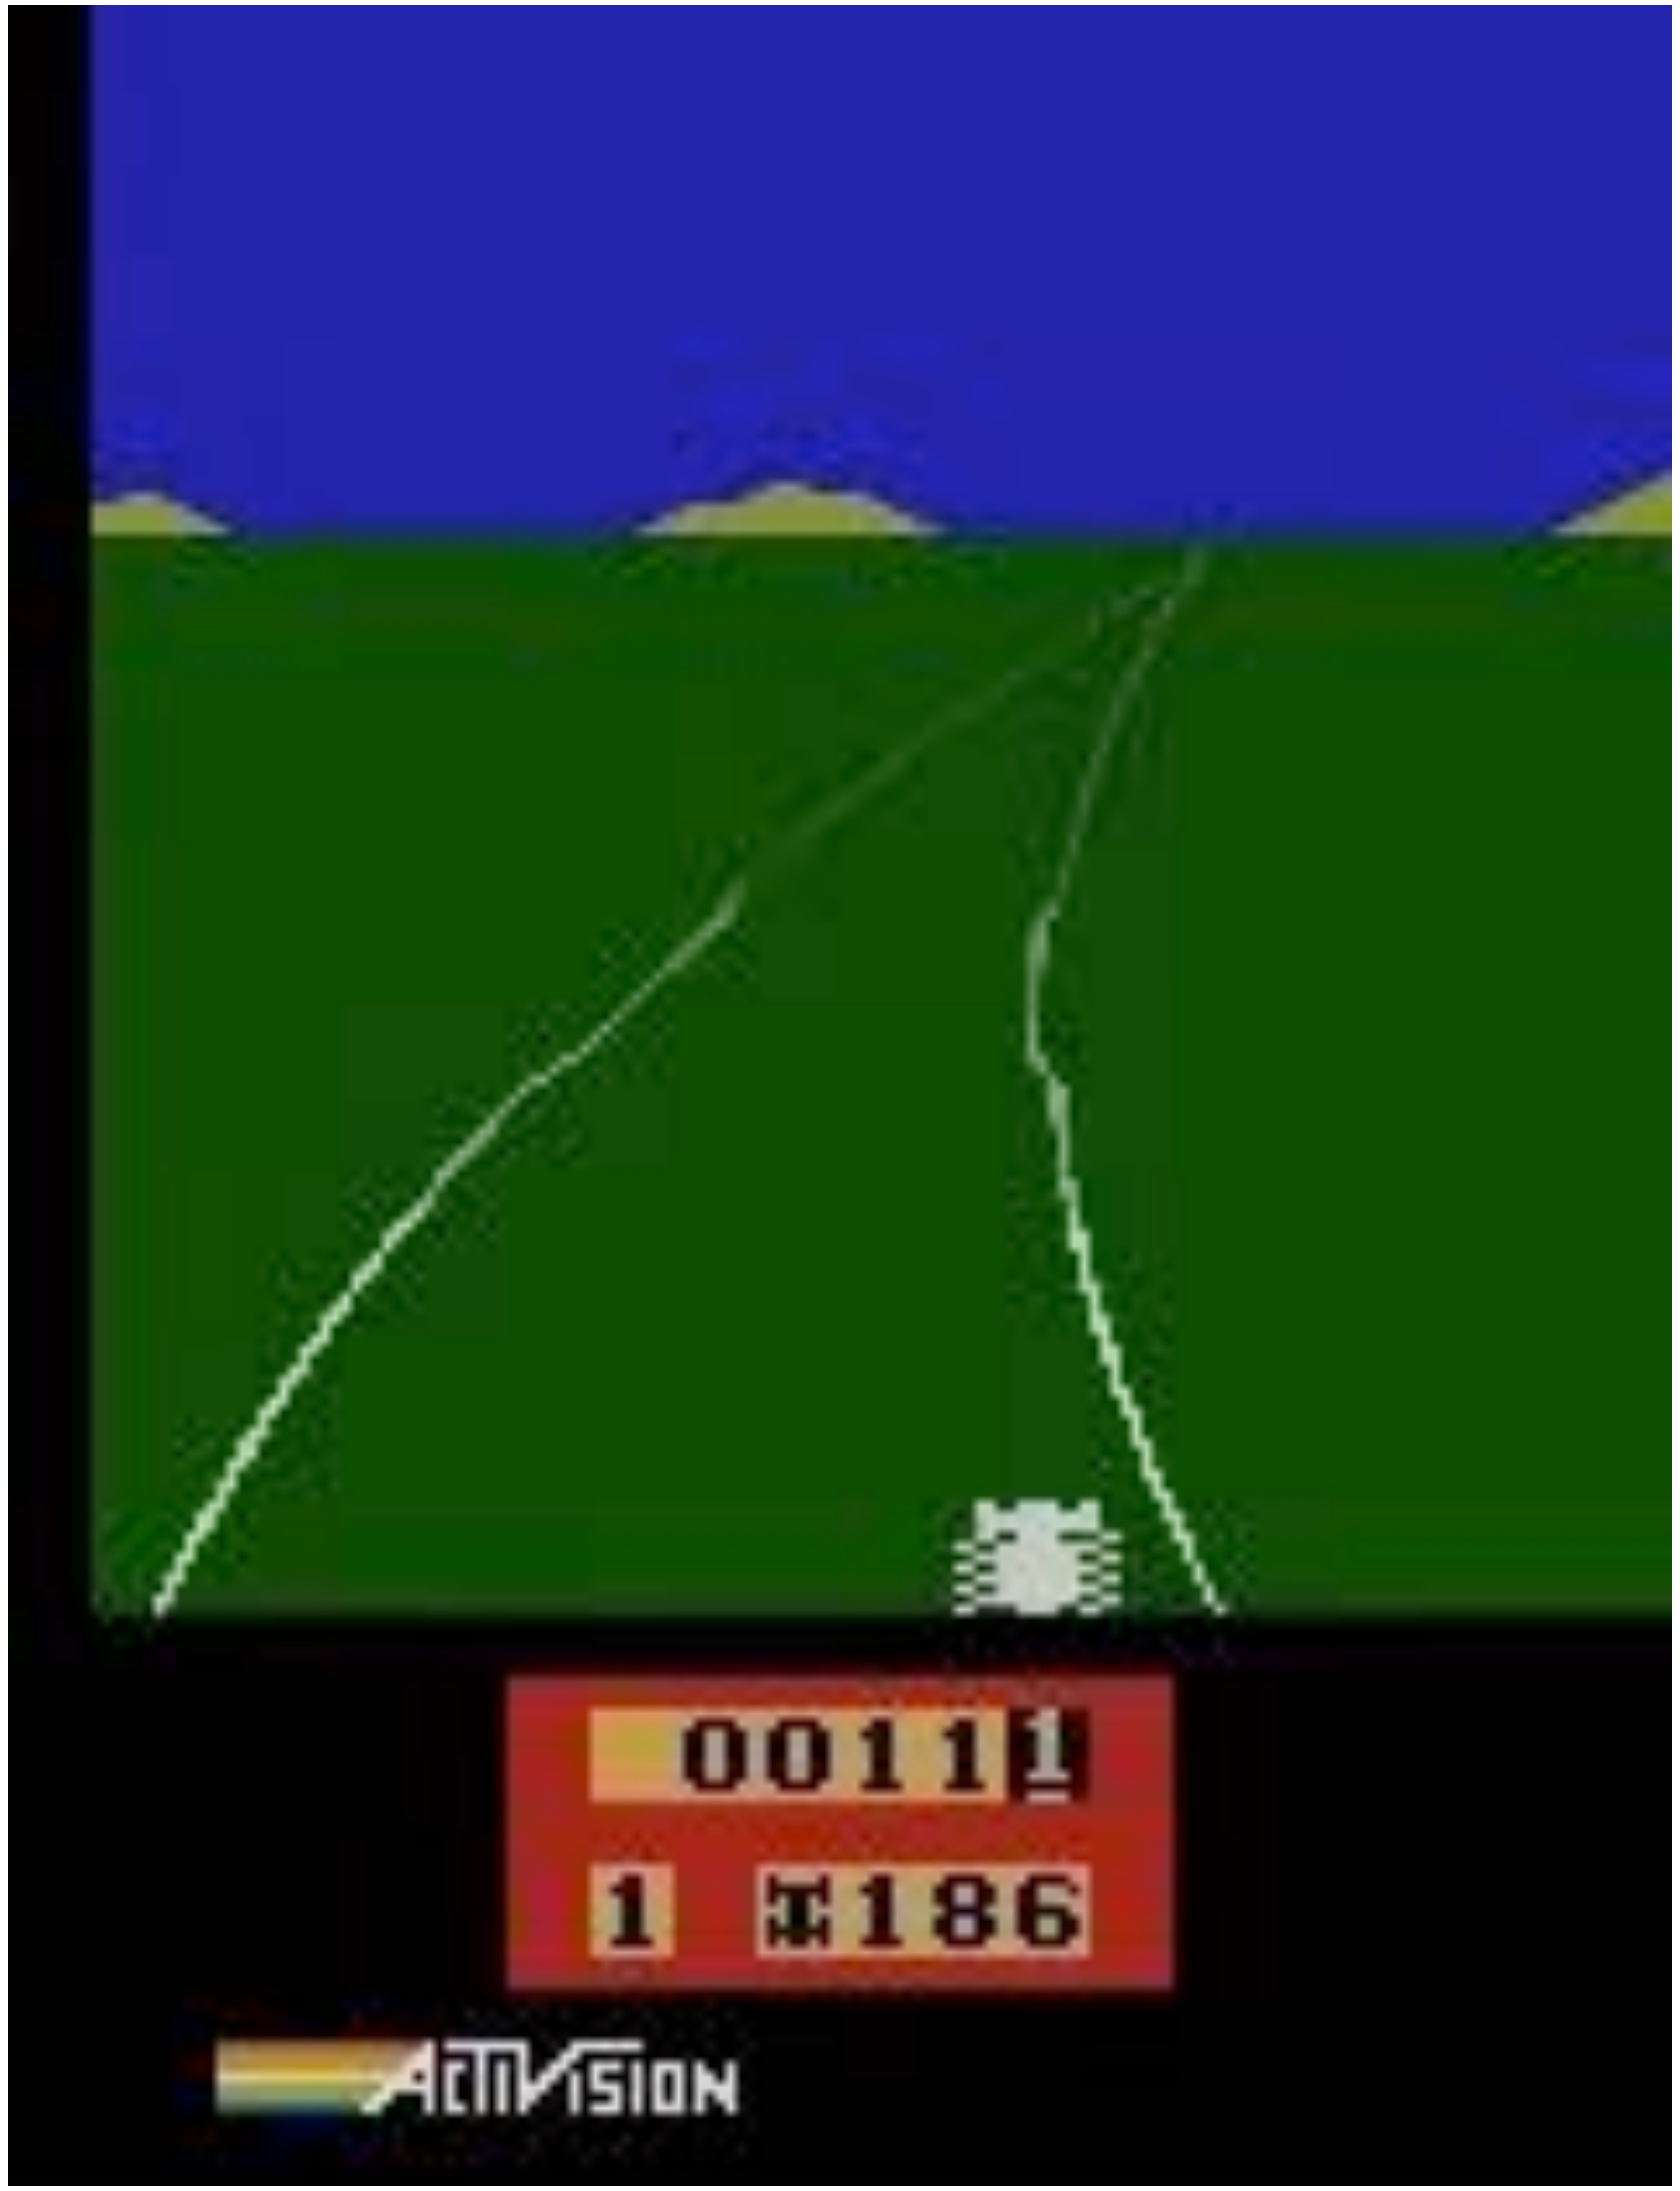

00111
1 186
ACTIVISION

0

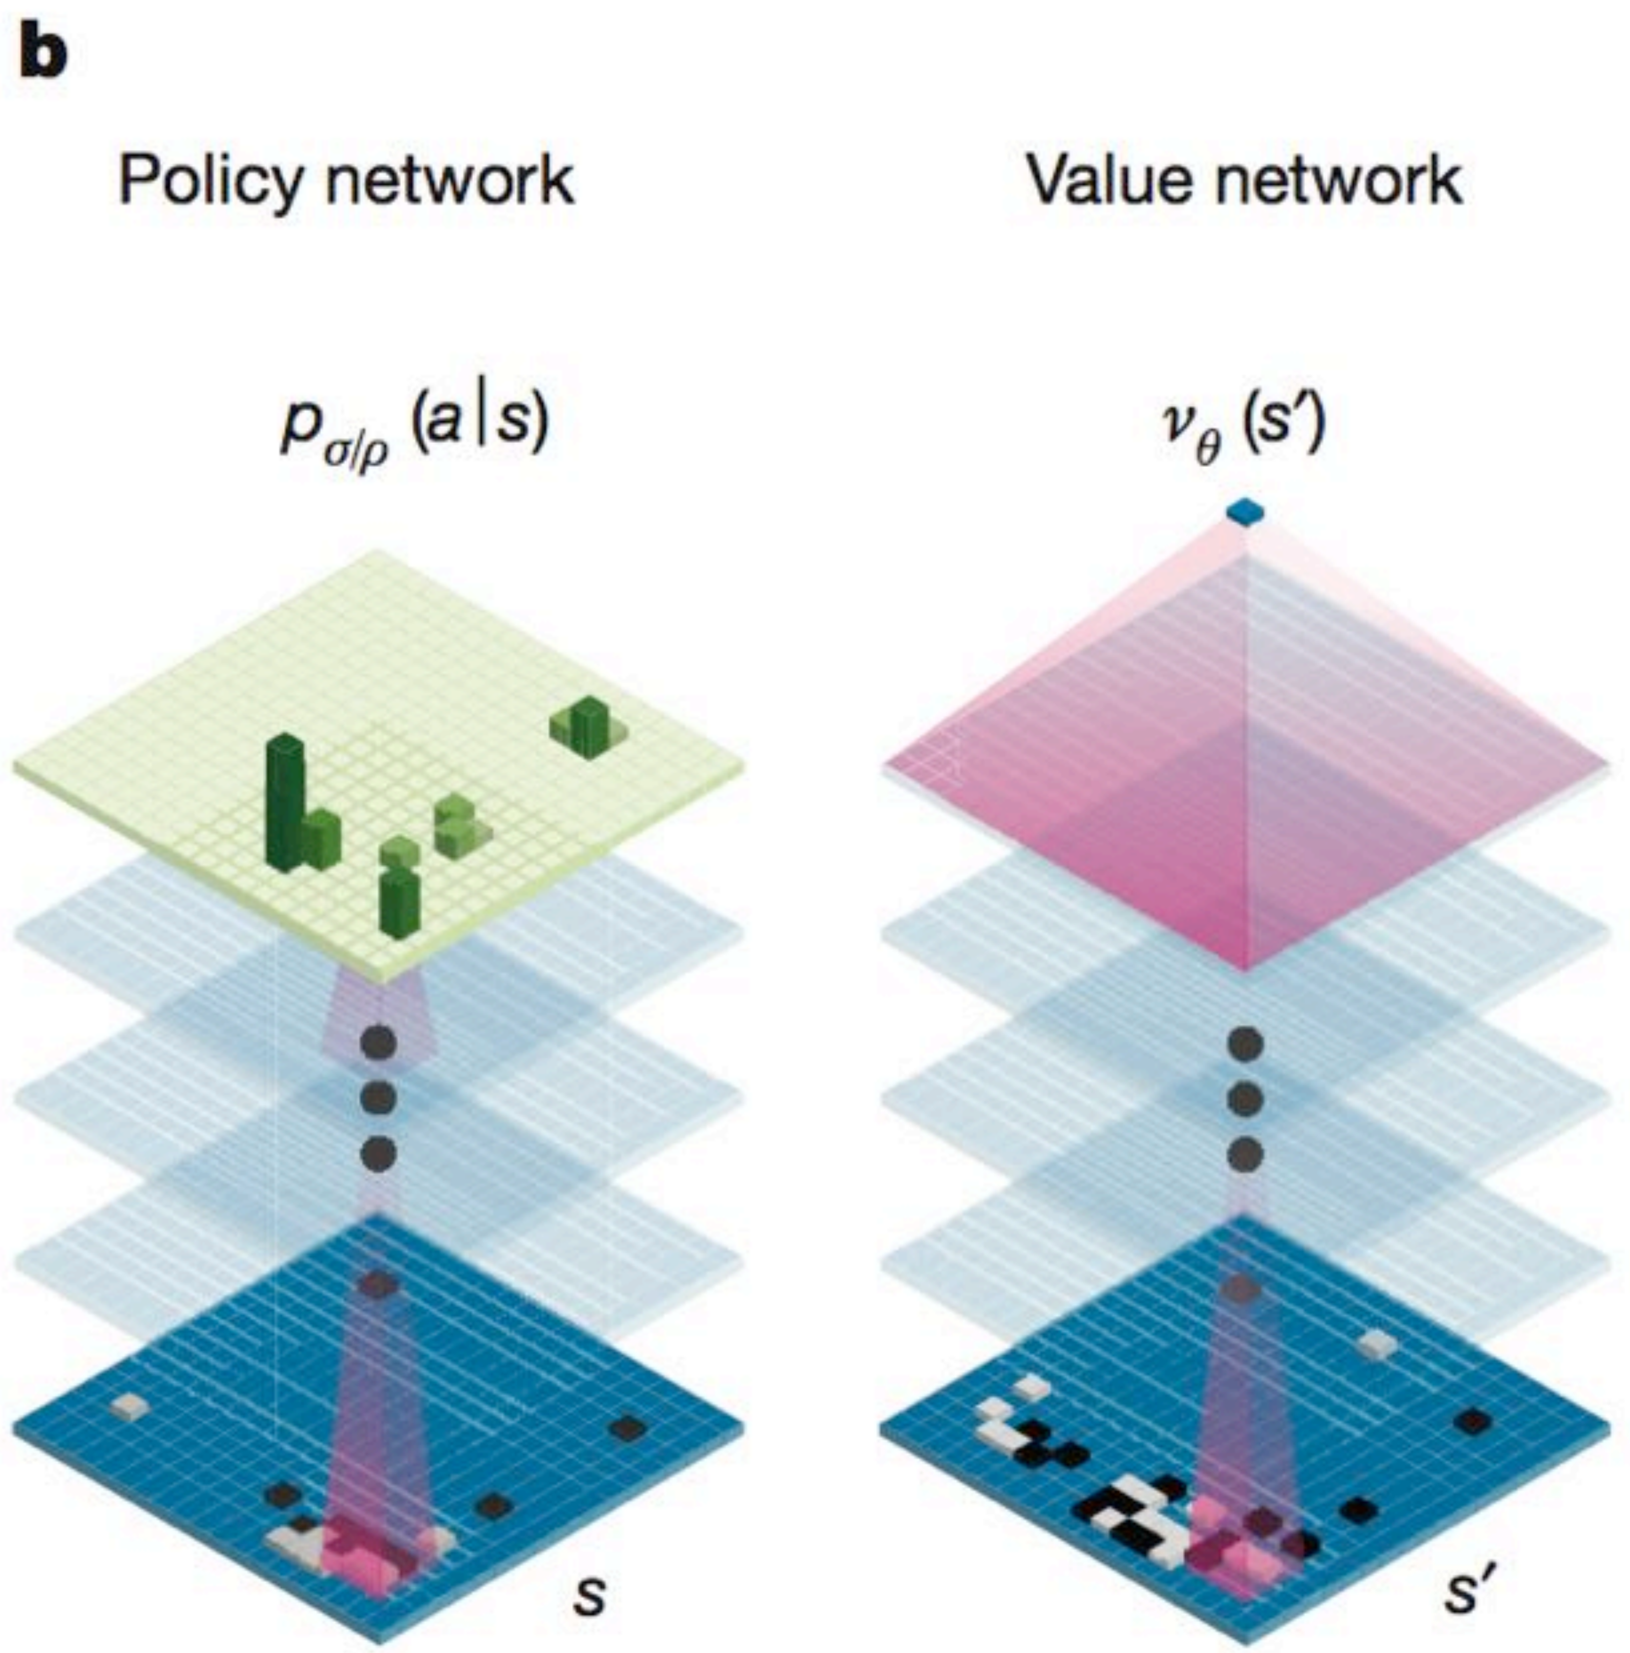

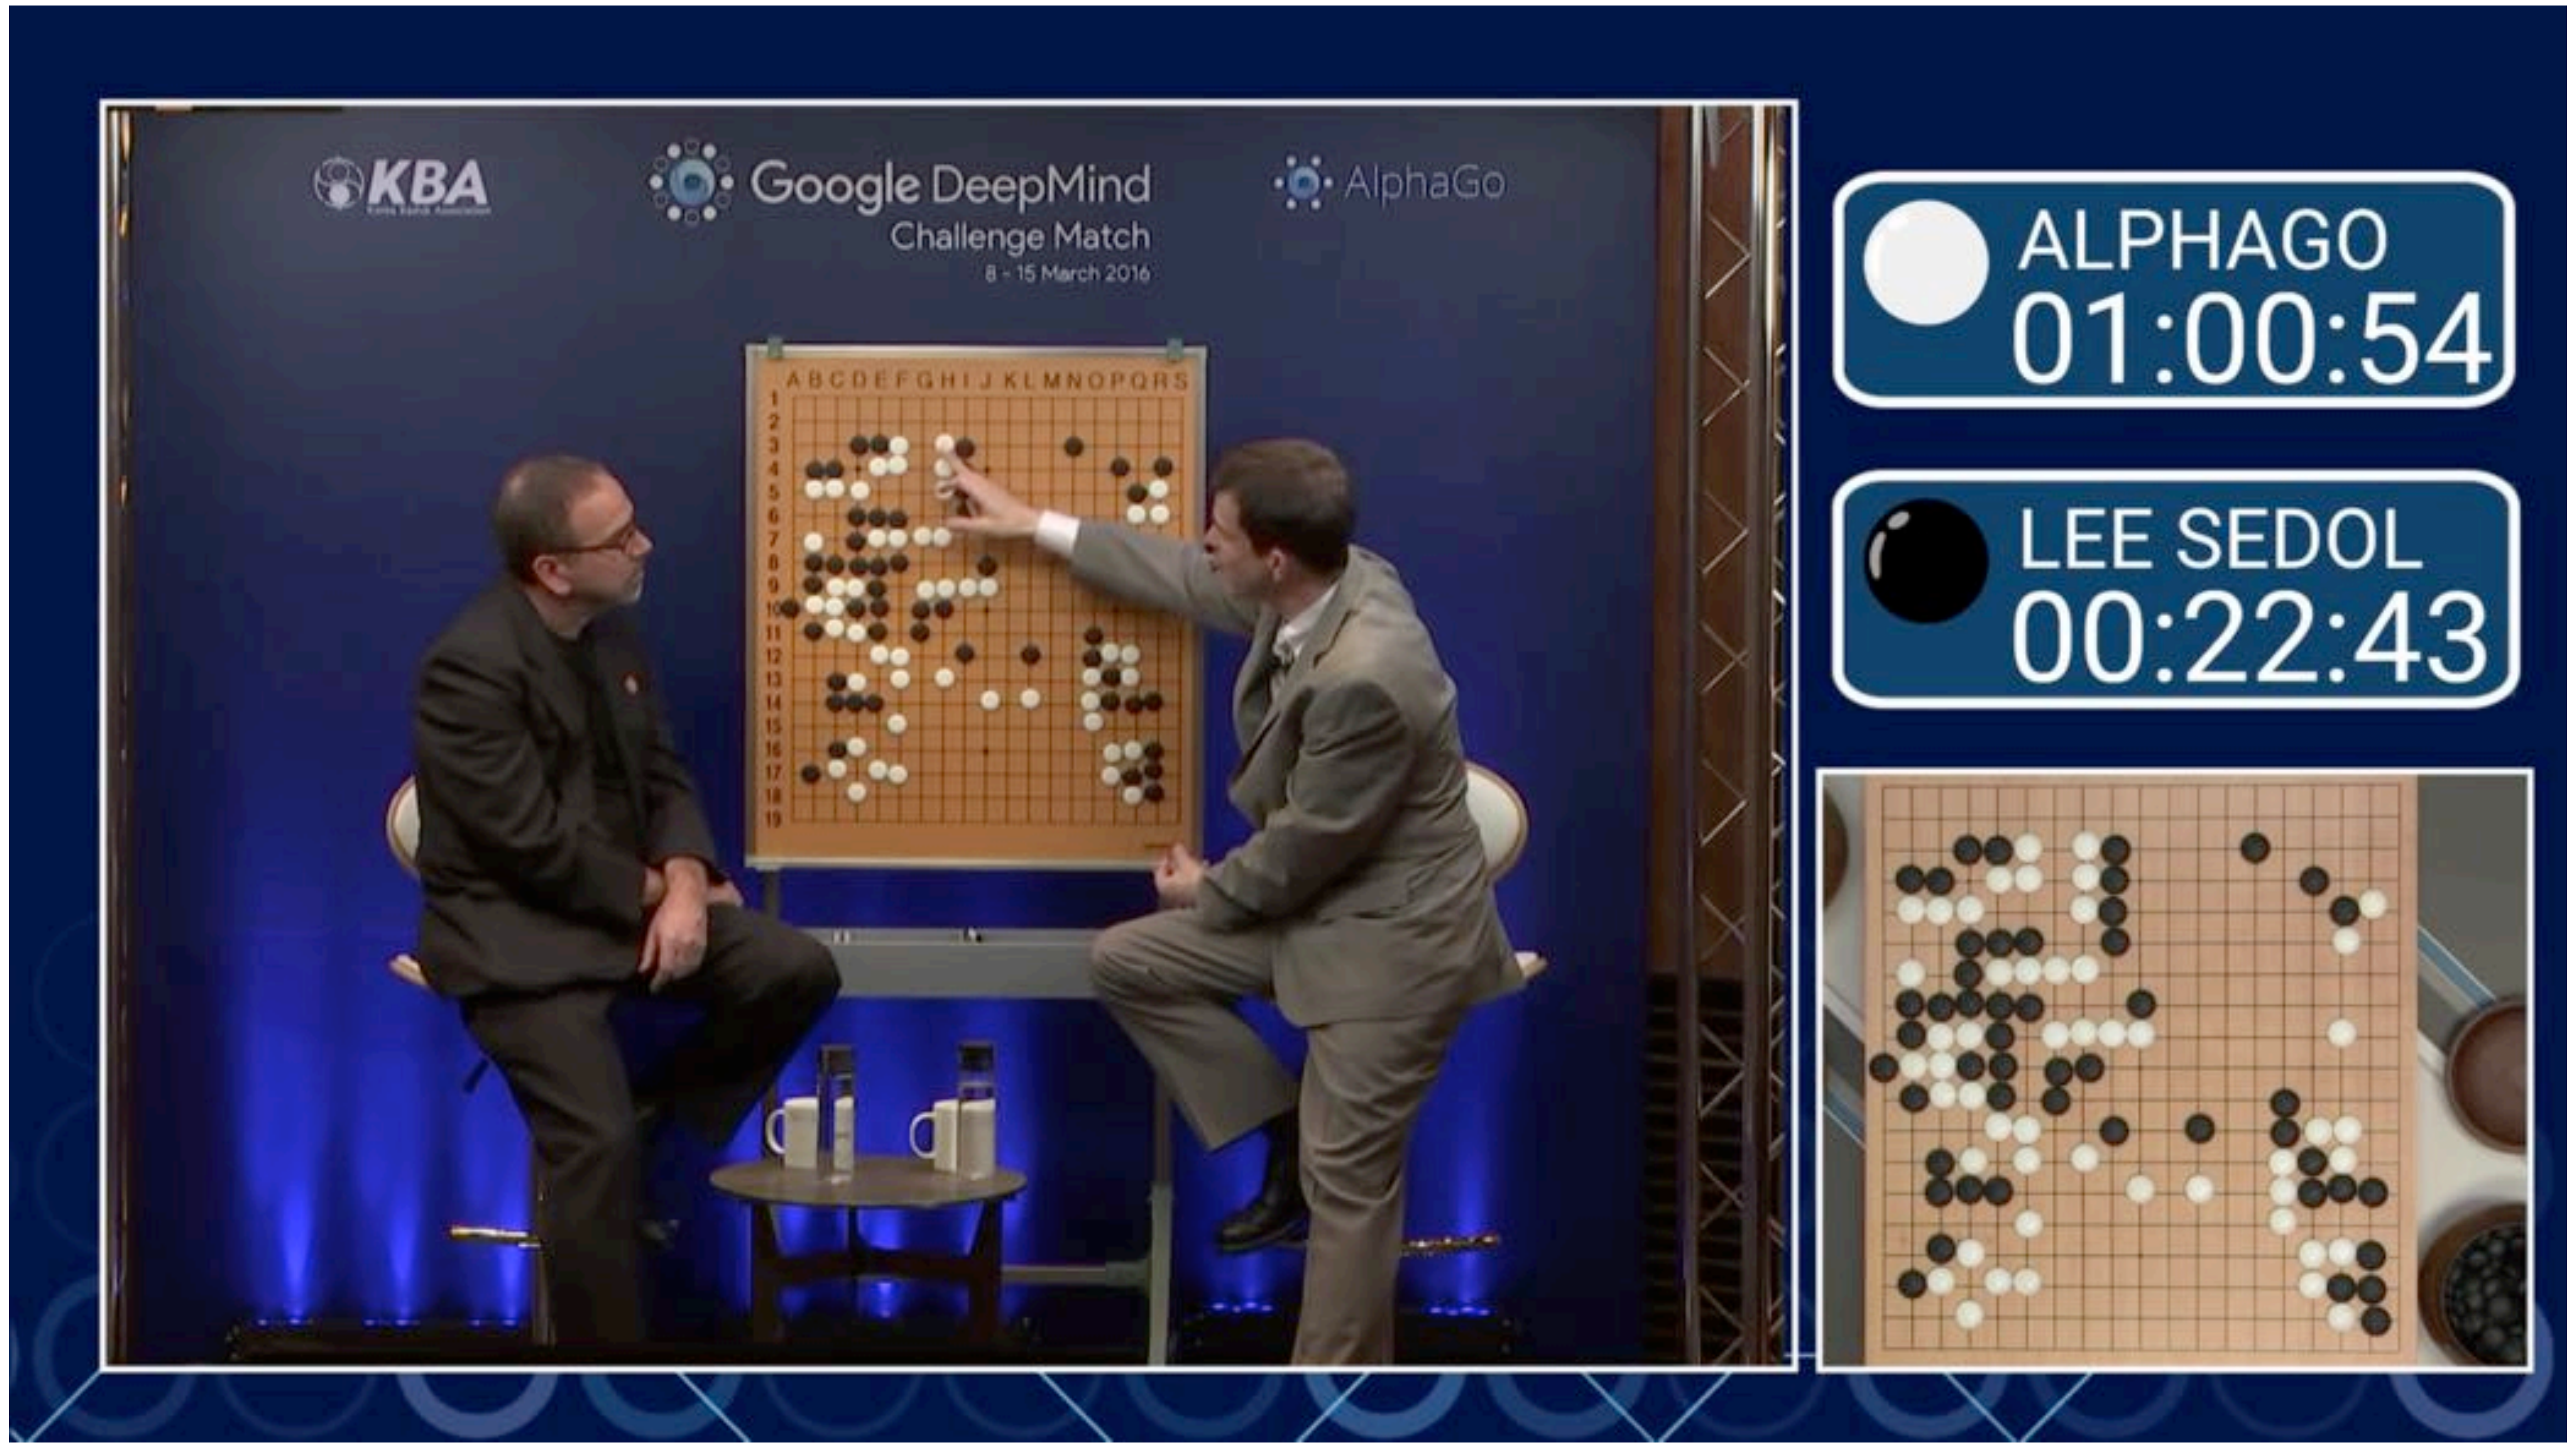

Mastering the game of Go with deep neural networks and tree search

David Silver[1]*, Aja Huang[1]*, Chris J. Maddison[1], Arthur Guez[1], Laurent Sifre[1], George van den Driessche[1], Julian Schrittwieser[1], Ioannis Antonoglou[1], Veda Panneershelvam[1], Marc Lanctot[1], Sander Dieleman[1], Dominik Grewe[1], John Nham[2], Nal Kalchbrenner[1], Ilya Sutskever[2], Timothy Lillicrap[1], Madeleine Leach[1], Koray Kavukcuoglu[1], Thore Graepel[1] & Demis Hassabis[1]

Source: Science

AI beats Human whenever it enters the competition

Applications

Unsupervised Learning

- There is just the data, no target, no feedback.
  Example: Probability Density Estimation, Clustering, Mannifold Learning, Feature Learning.

Visualization & Representation

PCA - t-SNE - Deep Autoencoder

Visualization of Gene Expression Data of Cancer Patients

3000 cancer patients,
20000 expressed genes per patient

1007_s_at	11.113	9.0016	9.73082	10.0858	9.02319	10.8084	8.53953	9.62217	9.46074	...	7.08799	7.34367	7.14374	6.98799
1053_at	8.63353	7.13359	8.75026	8.03603	8.25018	7.44034	7.33812	9.09907	8.76266	...	7.79403	6.75794	6.16958	6.16657
117_at	5.70633	5.60056	5.53539	5.58342	5.63654	5.57686	5.5821	5.85586	5.6893	...	6.3801	10.4054	9.55094	8.52637
121_at	8.43286	11.214	8.75827	8.715	8.50598	8.96119	8.62649	8.57374	10.1316	...	8.0948	7.91892	8.01743	7.78107
1255_g_at	3.6126	3.50758	3.49087	3.50939	3.48815	3.48791	3.78254	3.46632	3.4939	...	3.51056	3.54645	3.51243	3.55127
1294_at	6.56716	6.27671	6.48012	6.52817	6.07641	6.41791	6.99648	6.11197	5.89823	...	8.58388	7.64425	7.55718	7.35374
1316_at	5.32575	5.26078	5.5346	5.30245	5.41222	5.85616	5.35442	5.30053	5.24834	...	5.26625	5.65879	5.49398	5.96872
1320_at	5.0257	4.32289	4.14413	4.43518	5.70804	5.02017	5.47859	4.38358	4.6995	...	3.91704	4.37035	3.98443	3.87286
1405_i_at	4.33825	4.05247	3.8998	3.89405	3.80096	3.72654	3.89077	5.18569	3.679	...	5.66482	9.92337	9.64229	8.72453
1431_at	3.52288	3.57089	3.59646	3.54154	4.93521	3.88633	3.66972	4.05415	3.68796	...	3.54359	3.6156	3.75262	3.58121
1438_at	6.2833	6.25143	8.56186	6.31753	6.06036	6.53778	5.98391	5.73928	5.63754	...	4.67545	5.4586	5.58399	5.162
1487_at	7.57014	7.87256	8.41749	8.98784	7.46501									
1494_f_at	5.65639	5.70568	5.57135	5.59633	5.7263									
1552256_a_at	8.71516	8.60713	10.6836	9.31663	8.71652									

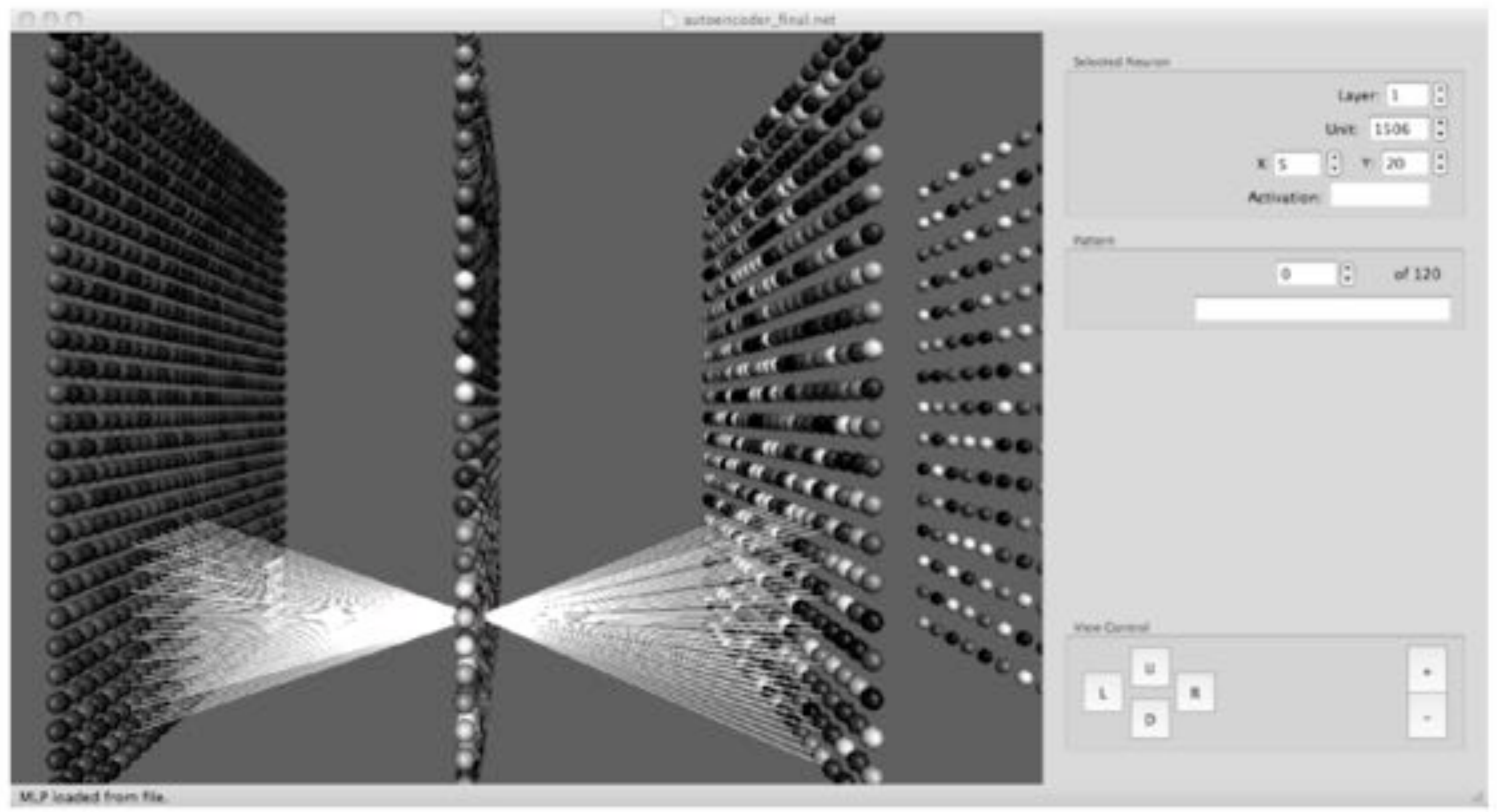

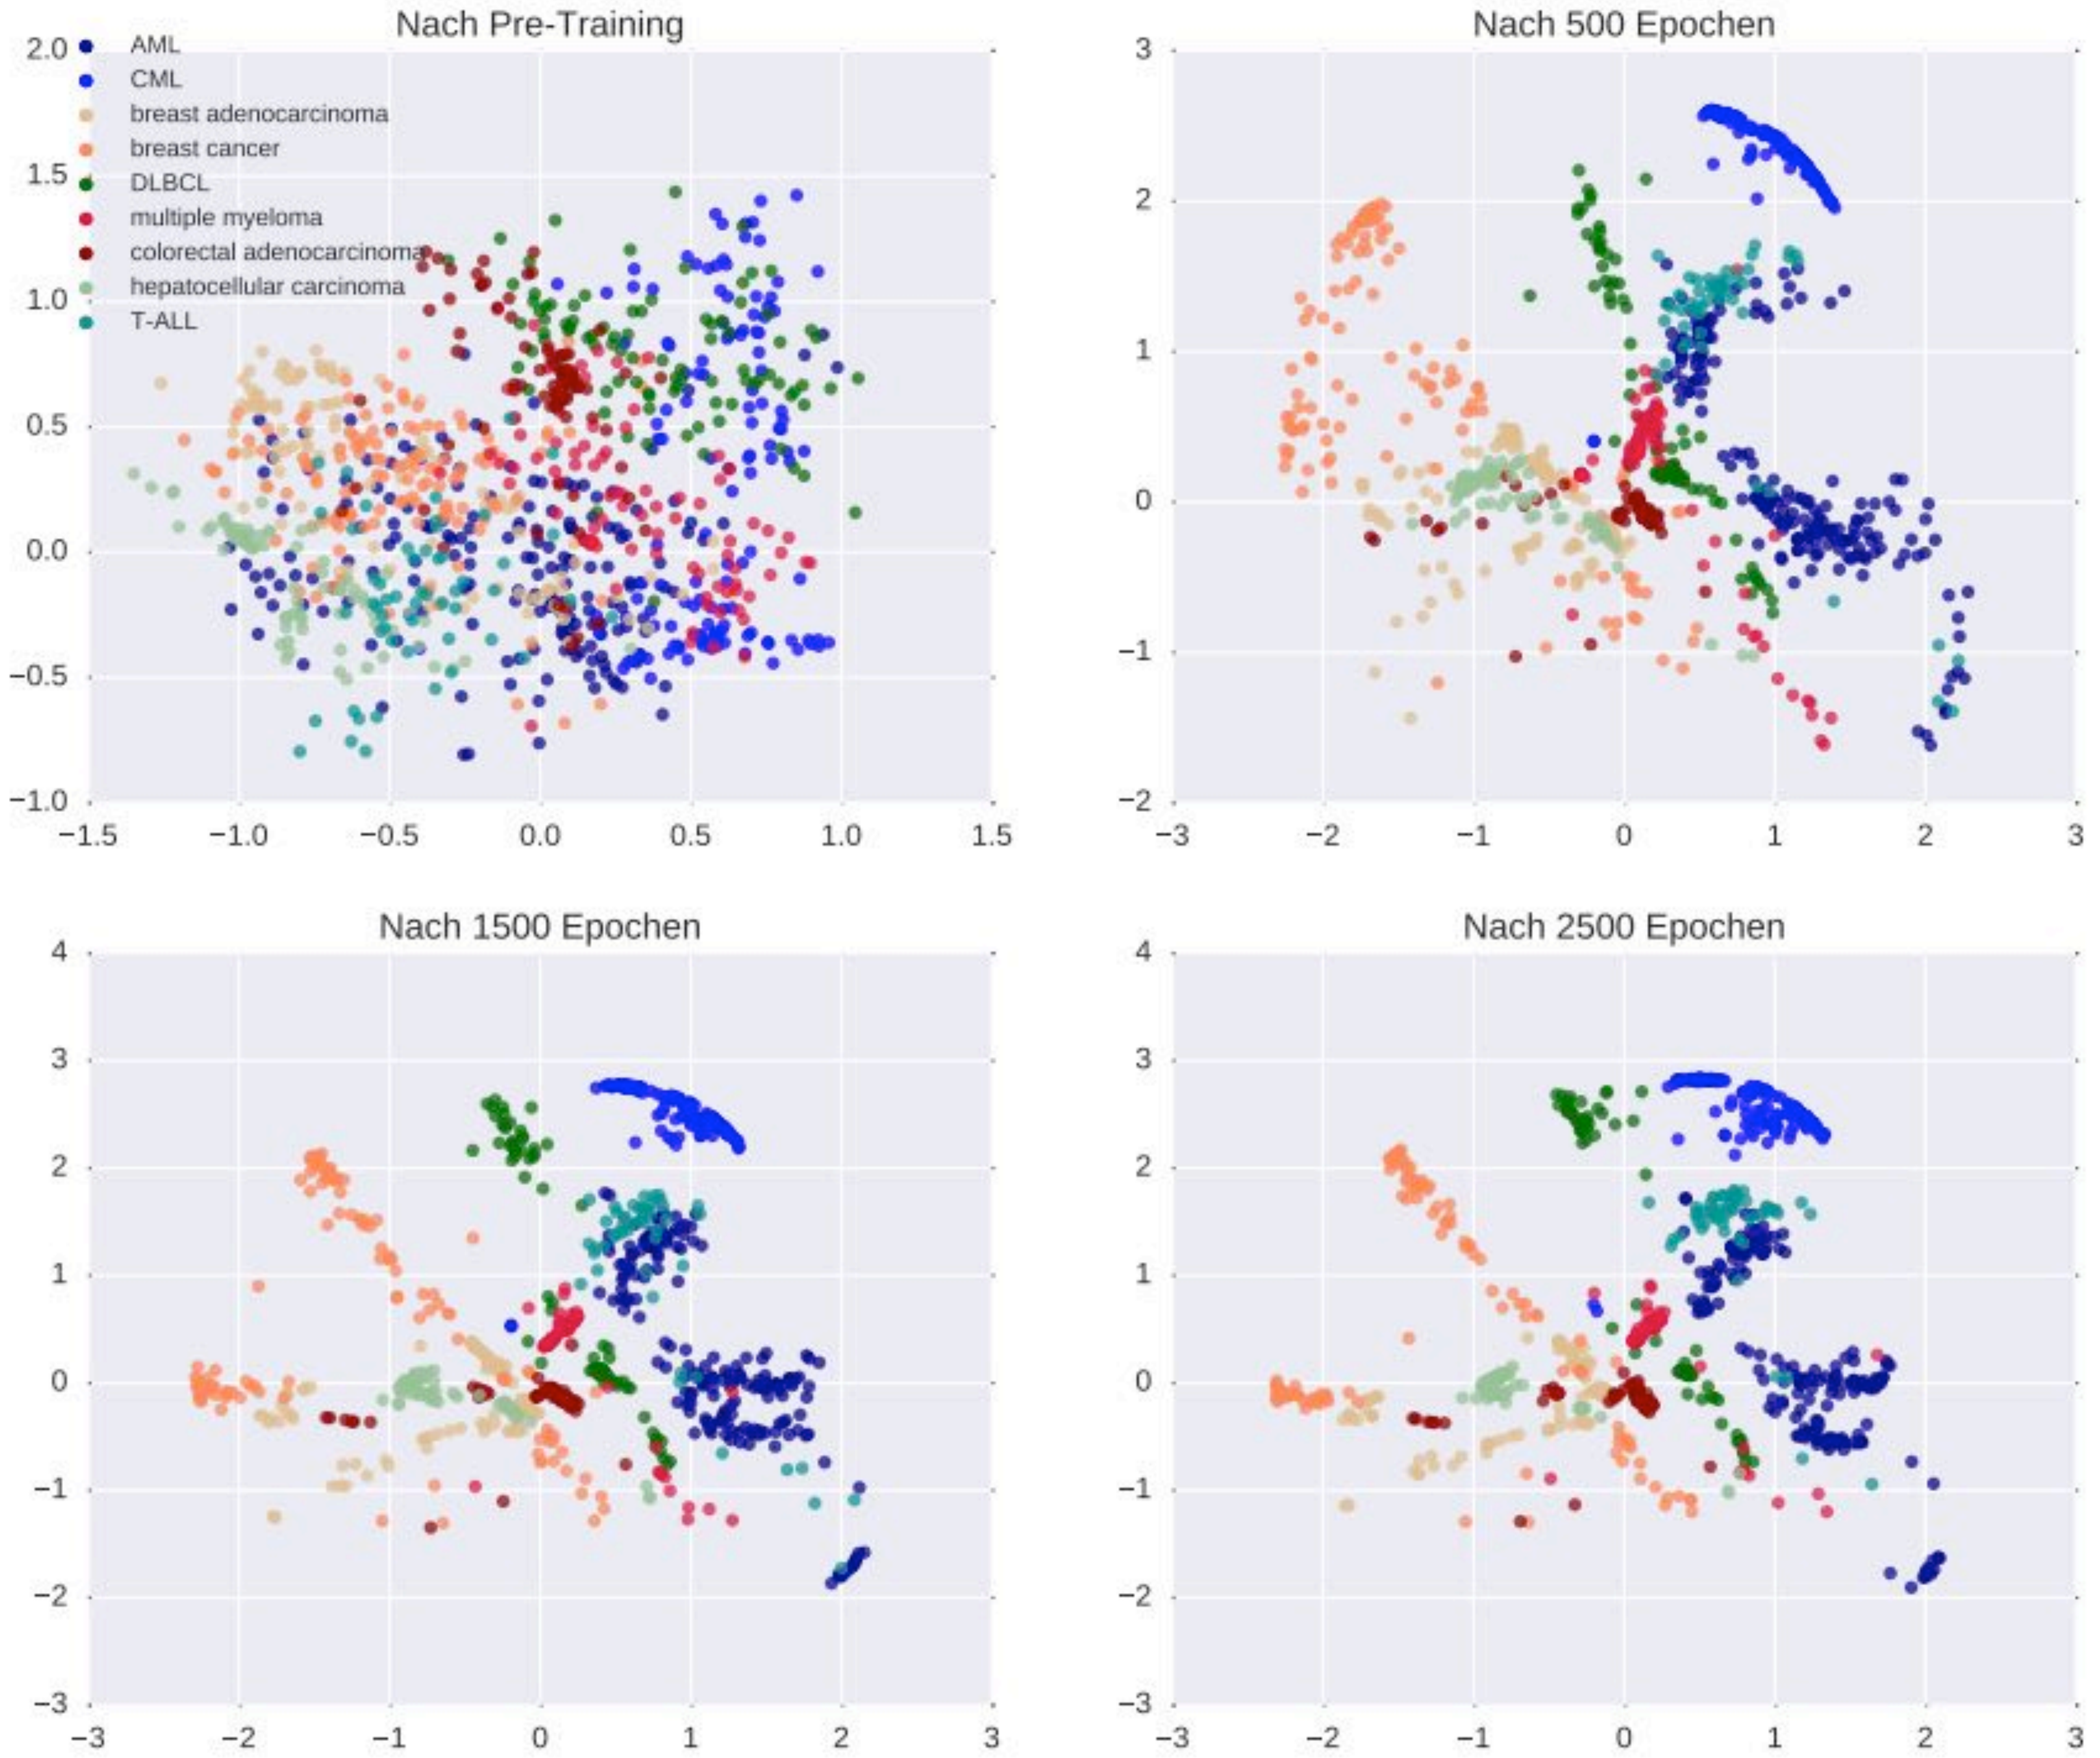

Klinikum Augsburg


Windows aktivieren
Wechseln Sie zu den Einstellungen, um

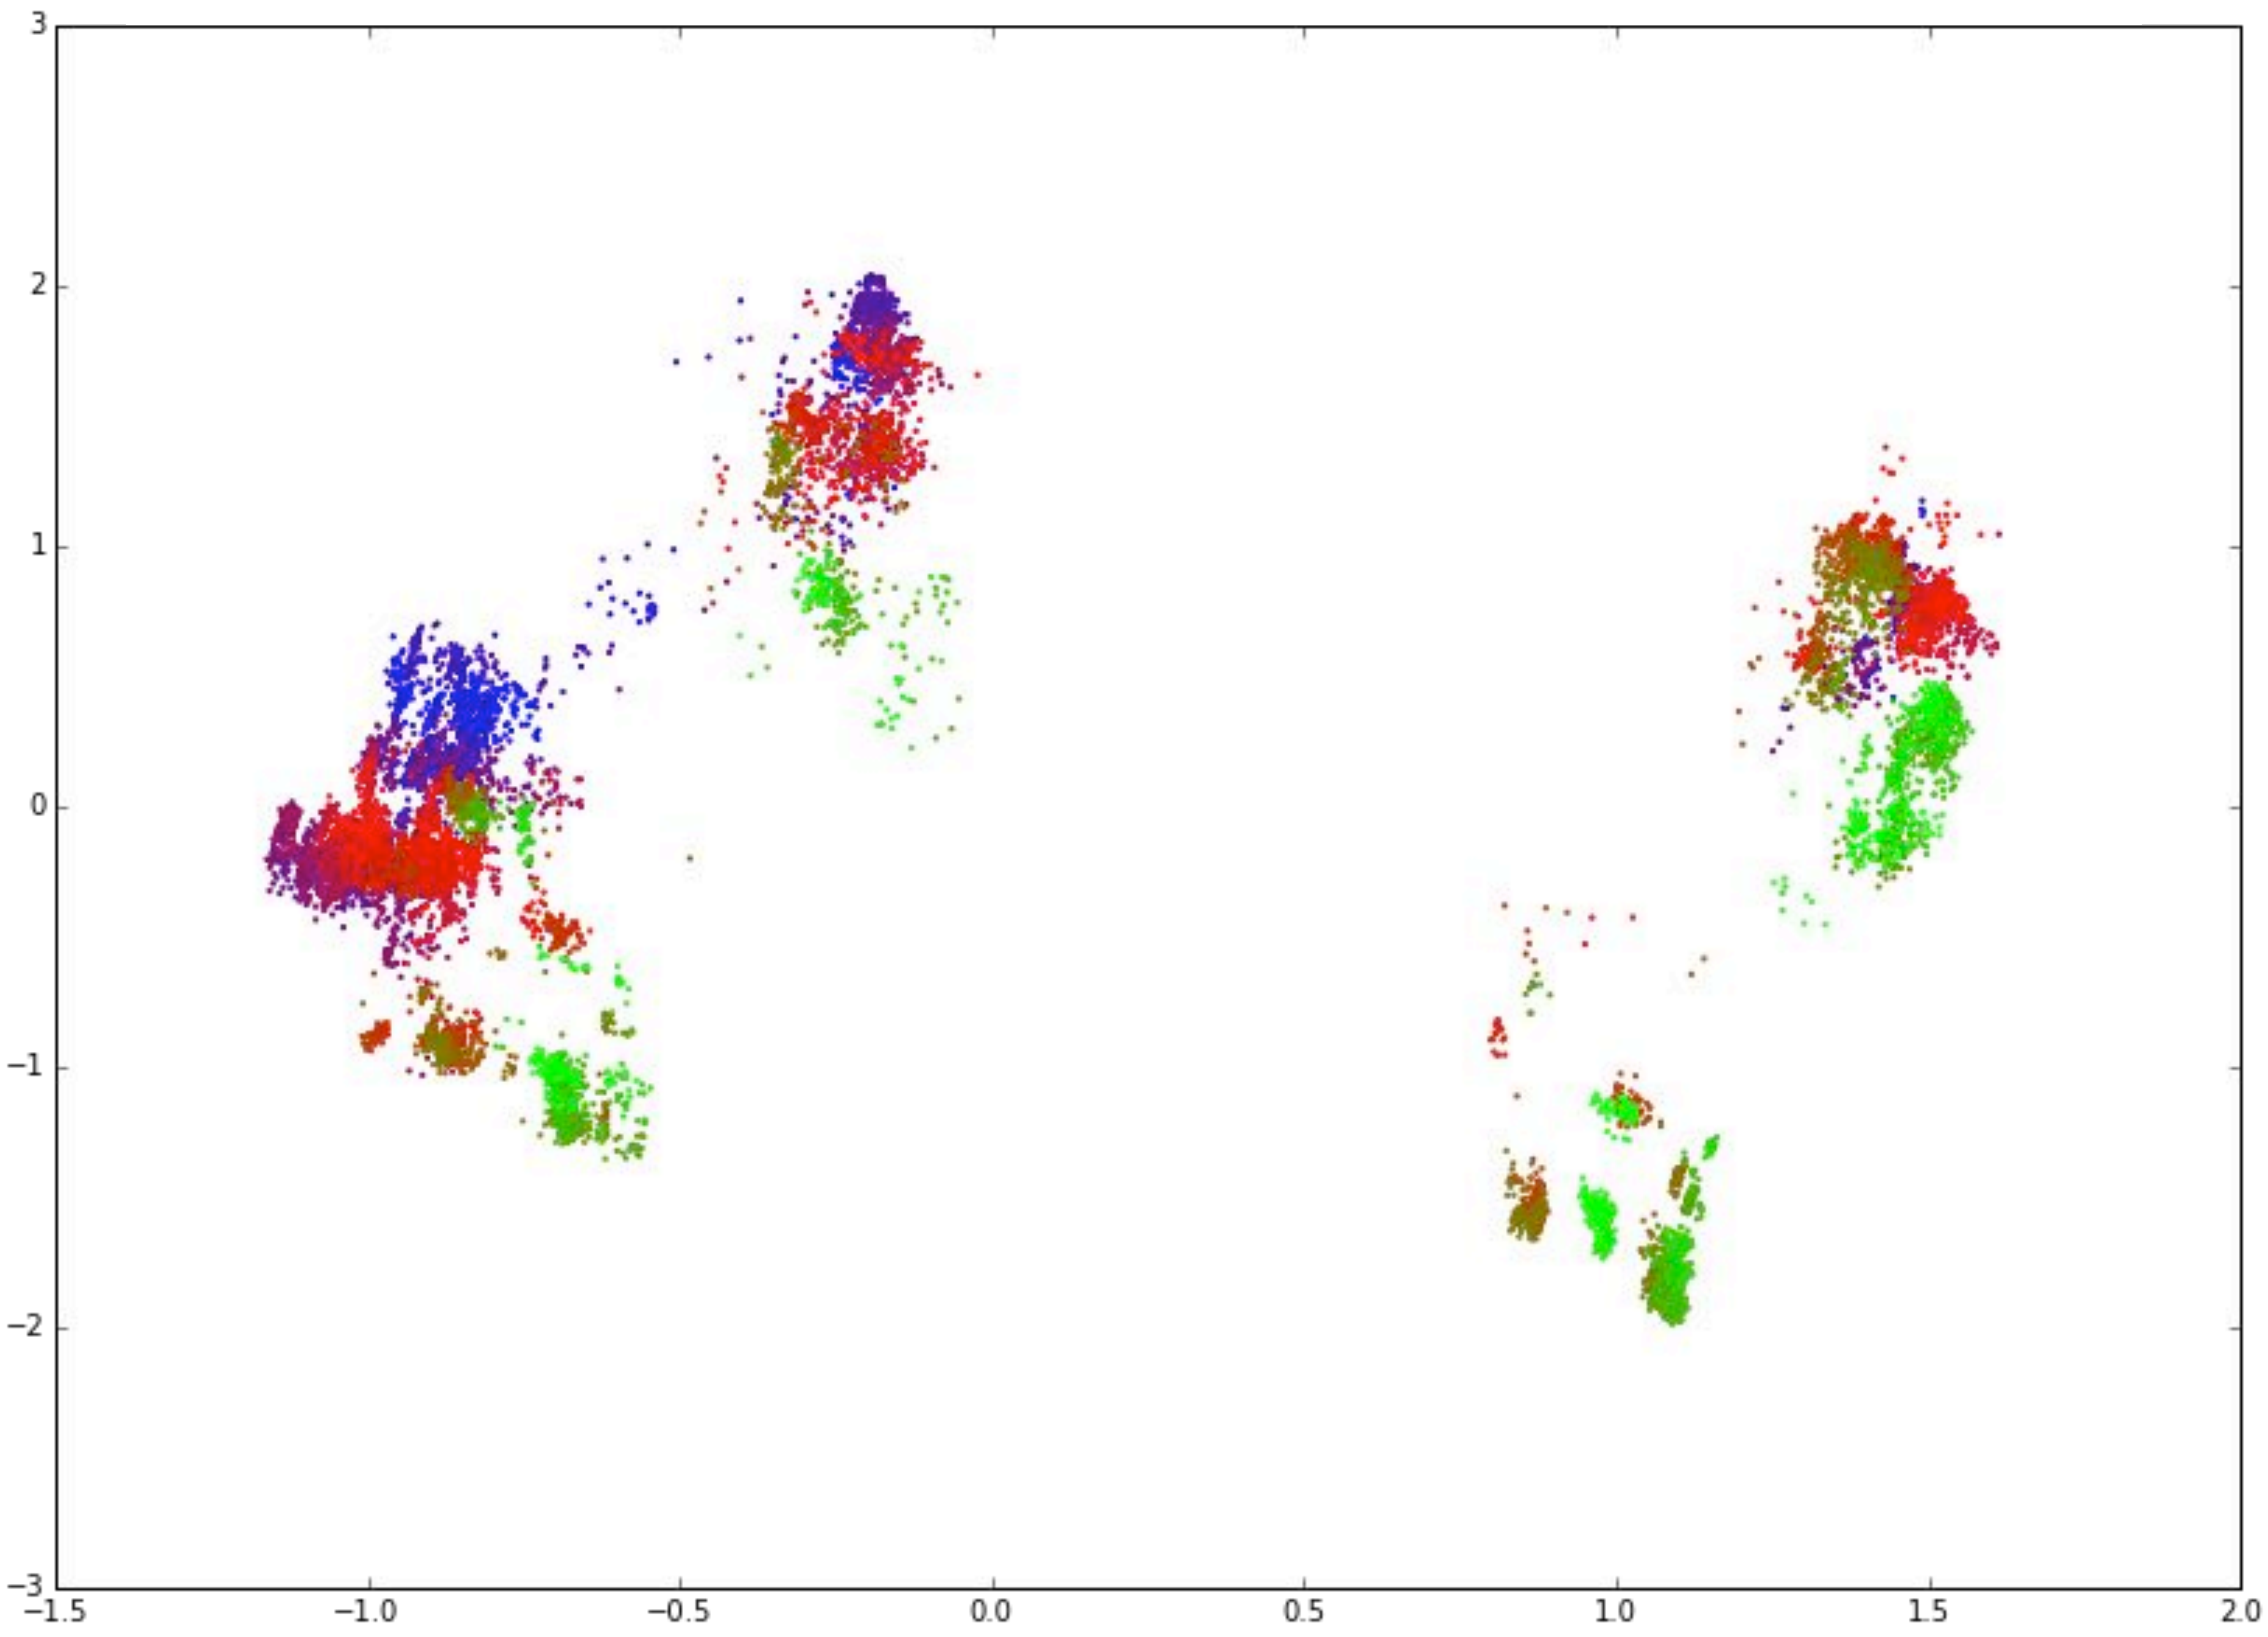
3
2
1
0
−1
−2
−3
−1.5
−1.0
−0.5
0.0
0.5
1.0
1.5
2.0

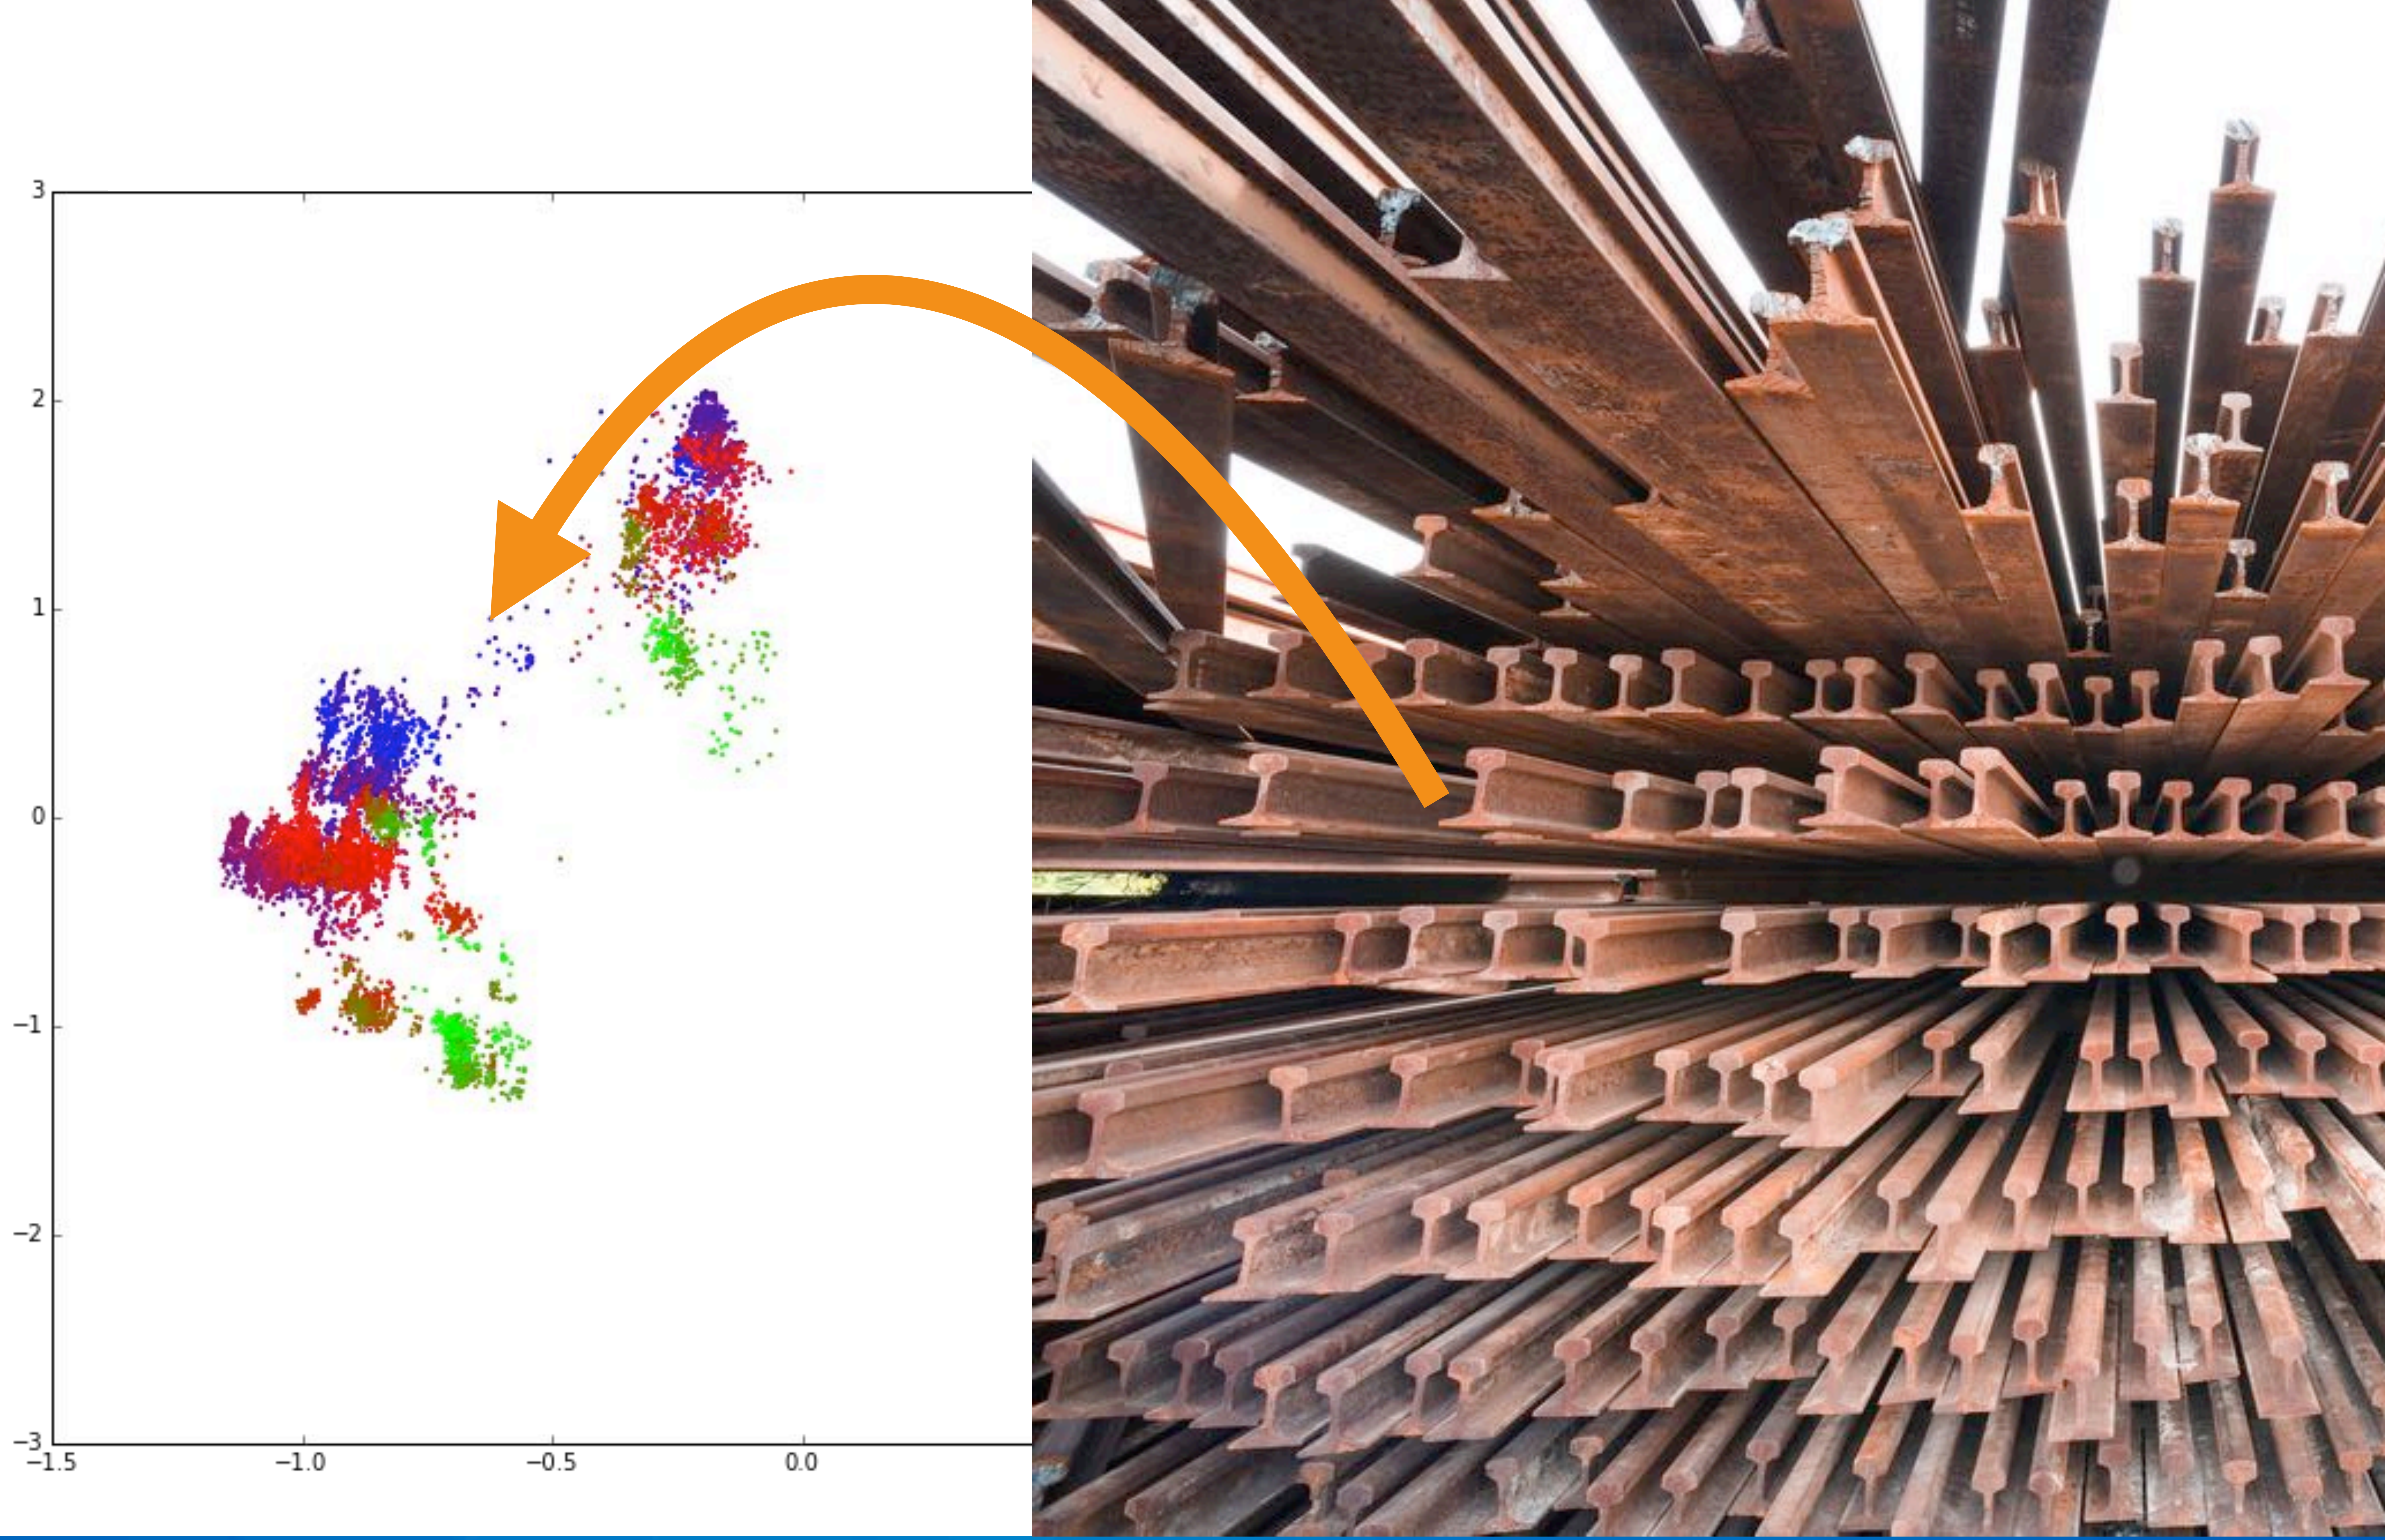
3
2
1
0
−1
−2
−3
−1.5
−1.0
−0.5
0.0

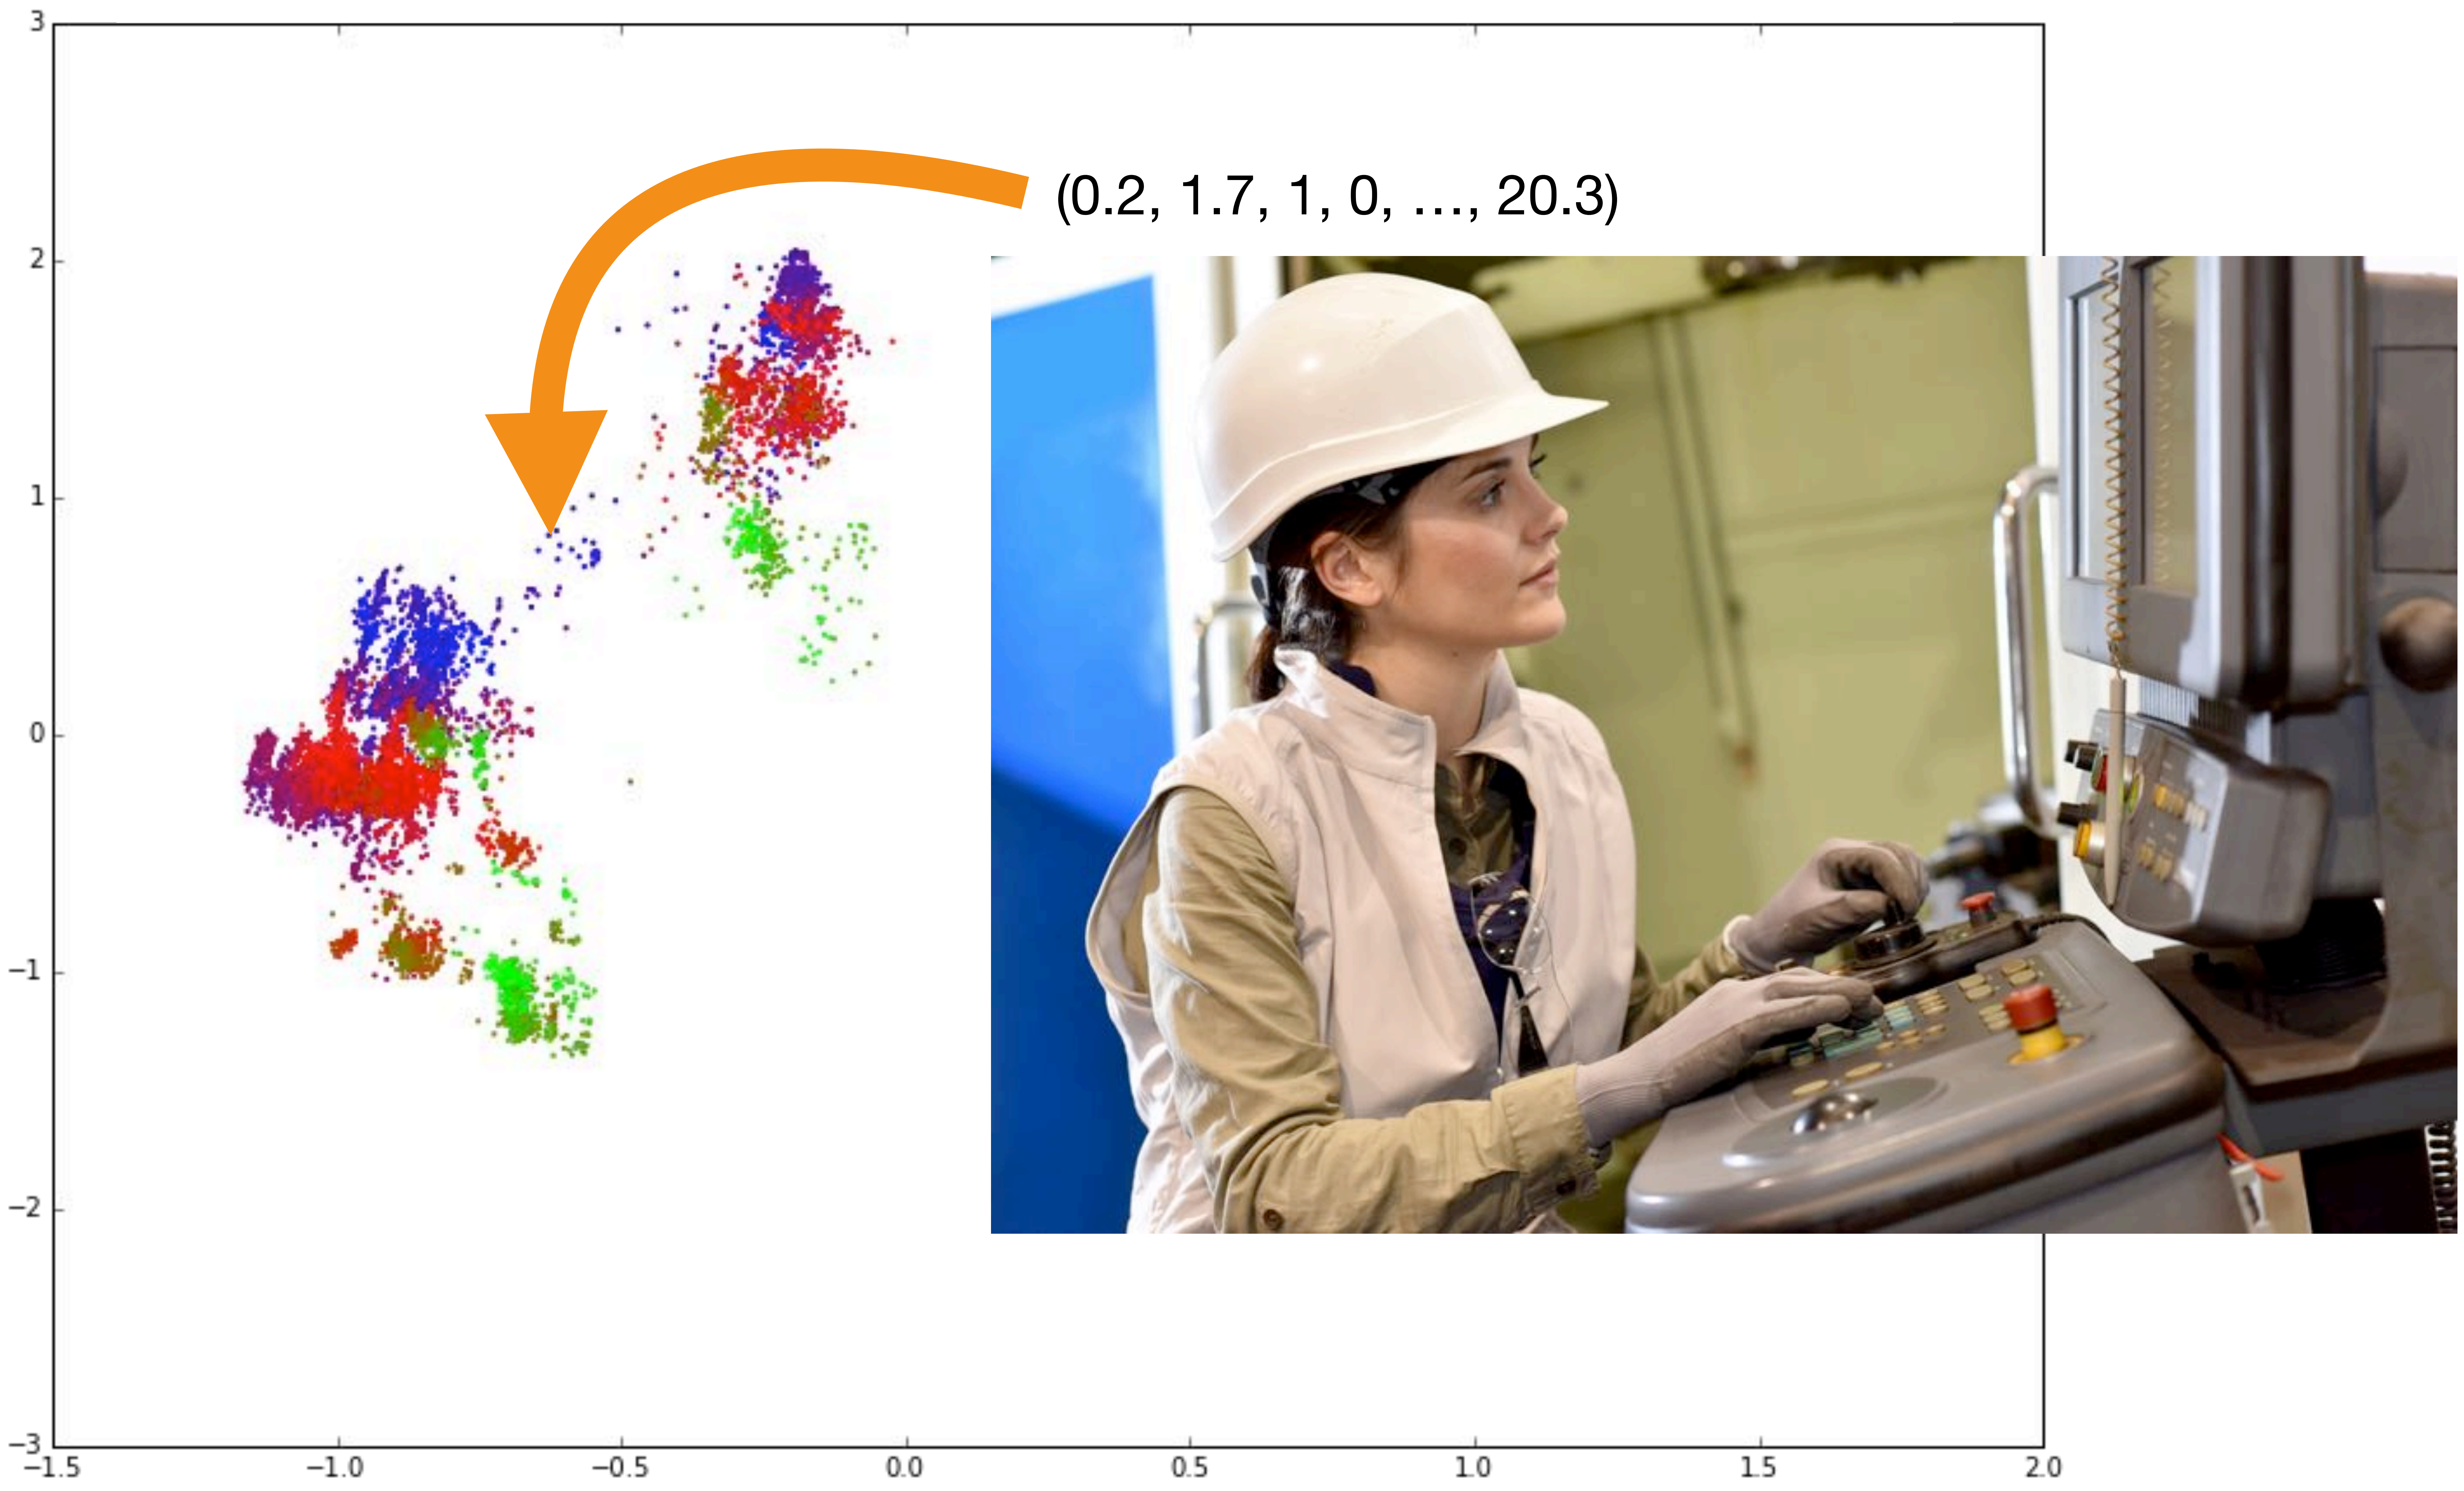
(0.2, 1.7, 1, 0, …, 20.3)
3
2
1
0
−1
−2
−3
−1.5
−1.0
−0.5
0.0
0.5
1.0
1.5
2.0

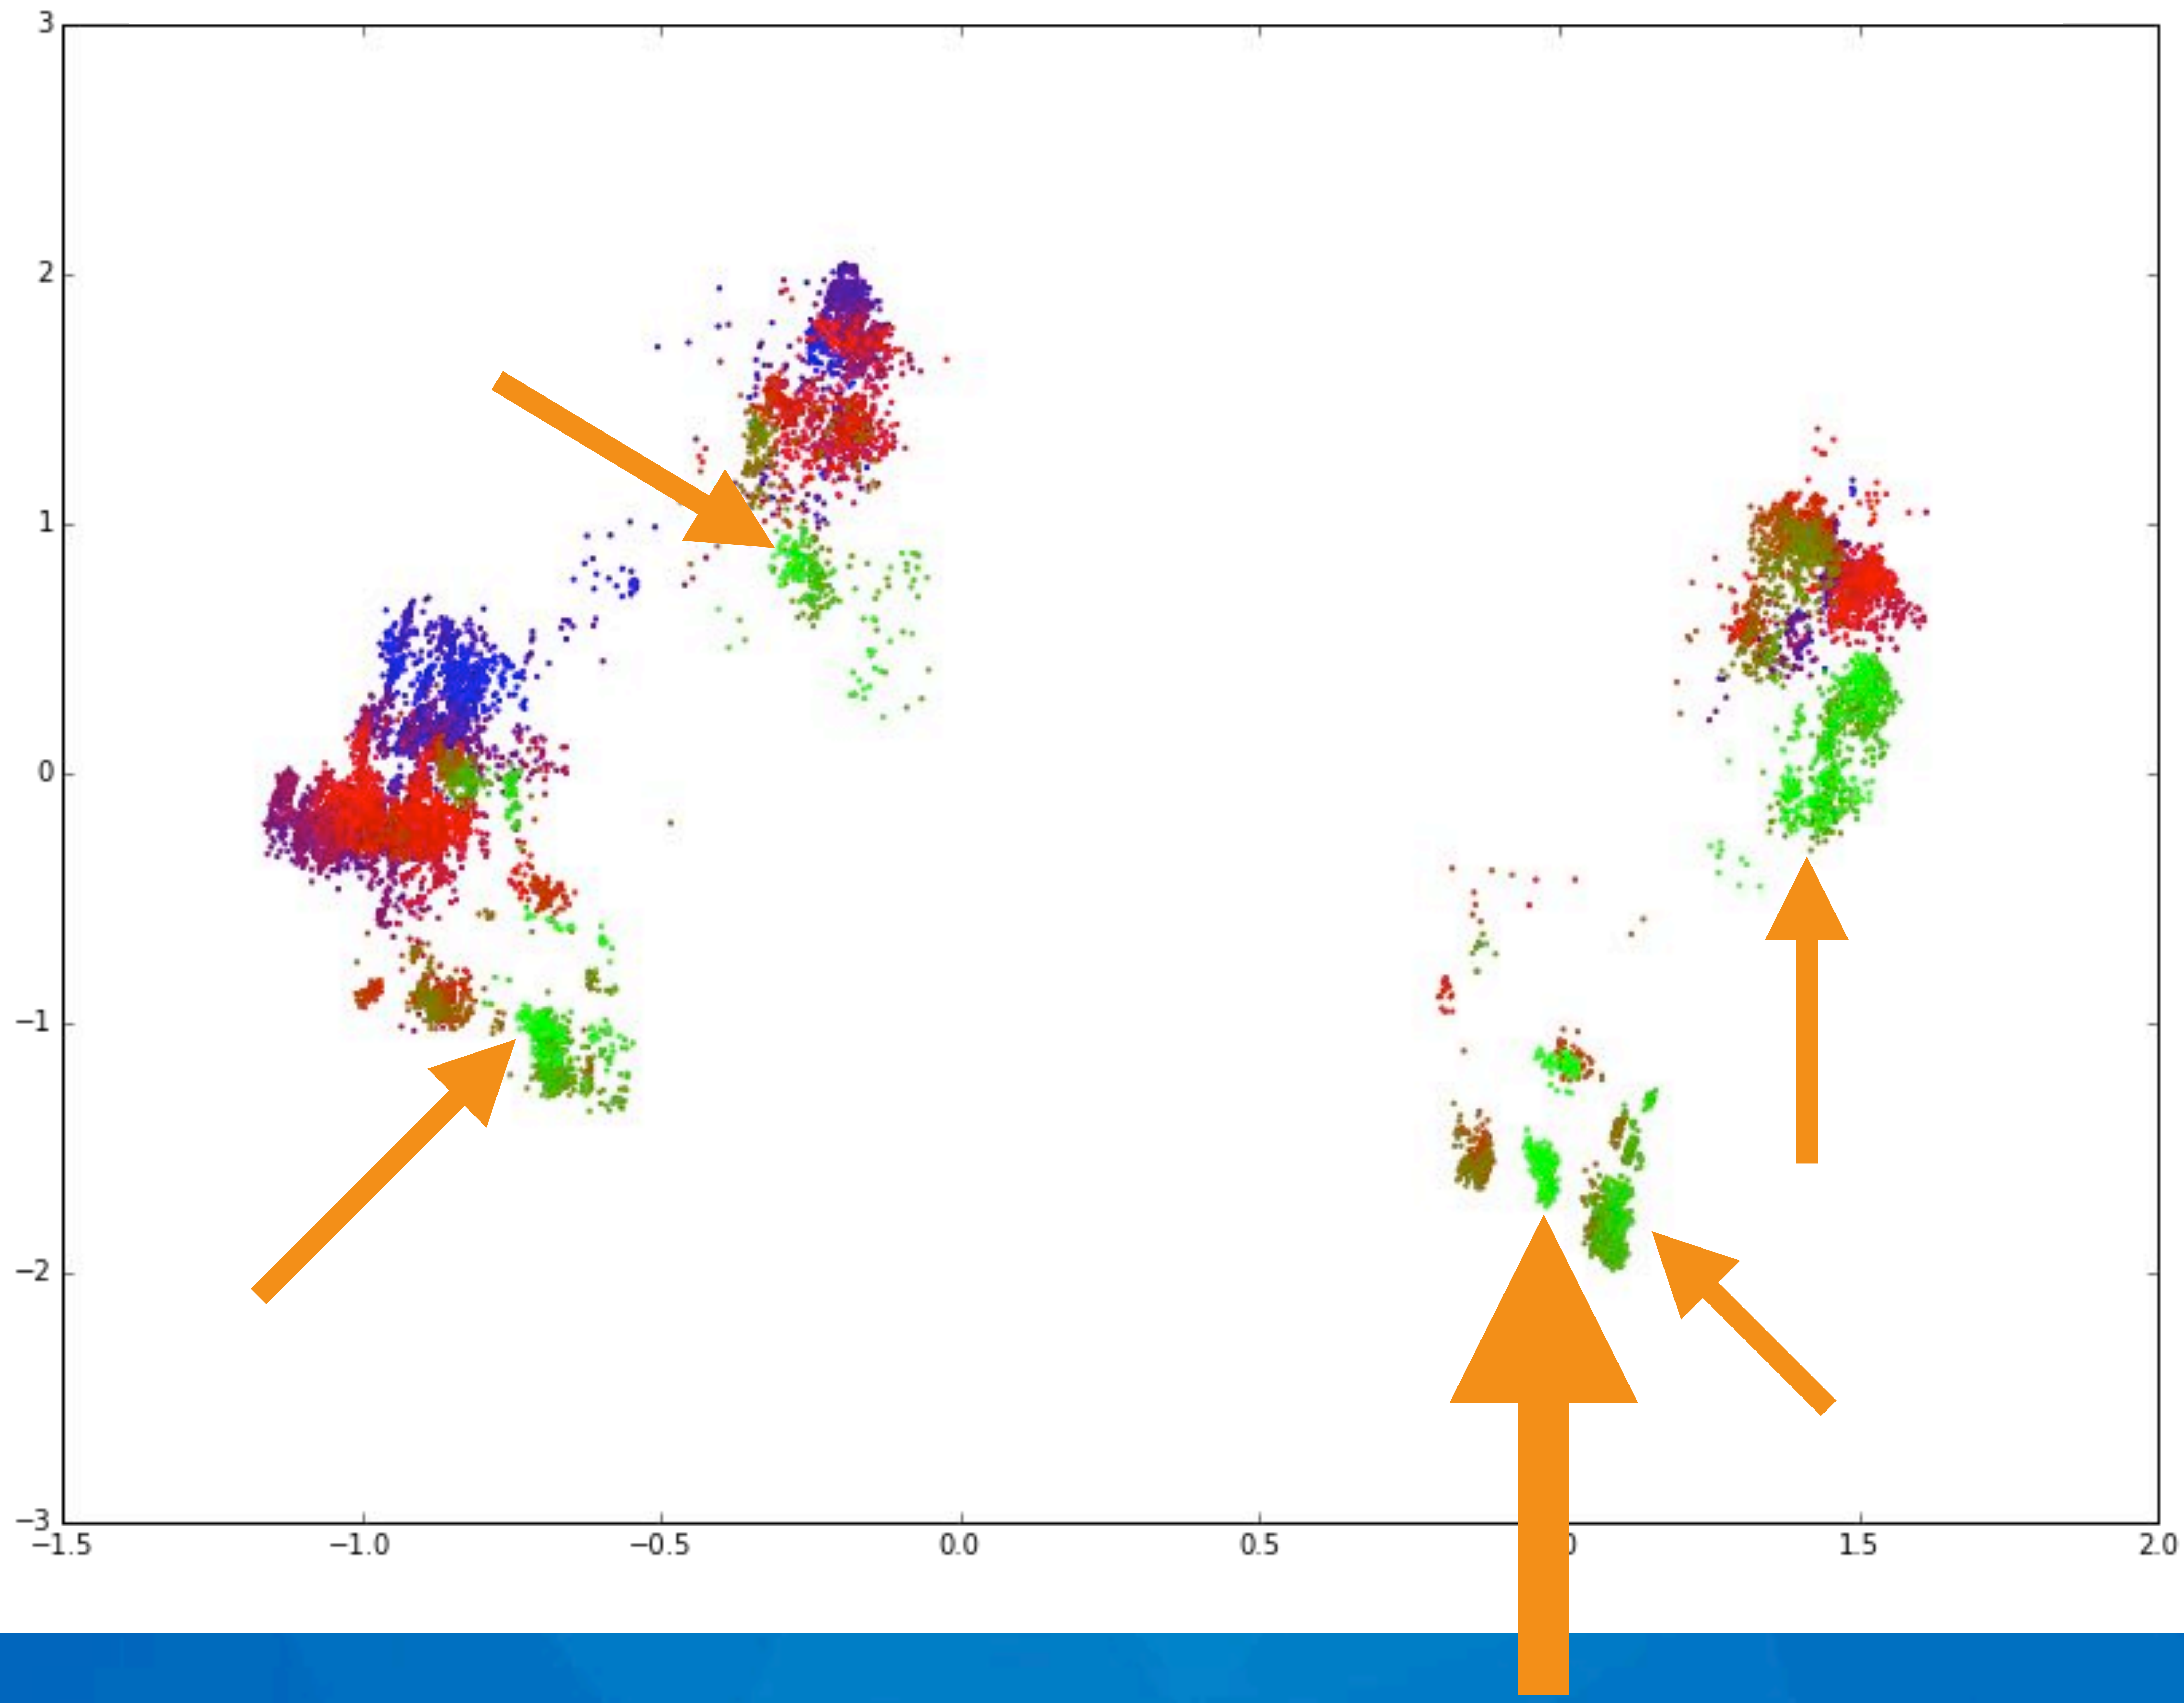
3
2
1
0
−1
−2
−3
−1.5
−1.0
−0.5
0.0
0.5
1.5
2.0

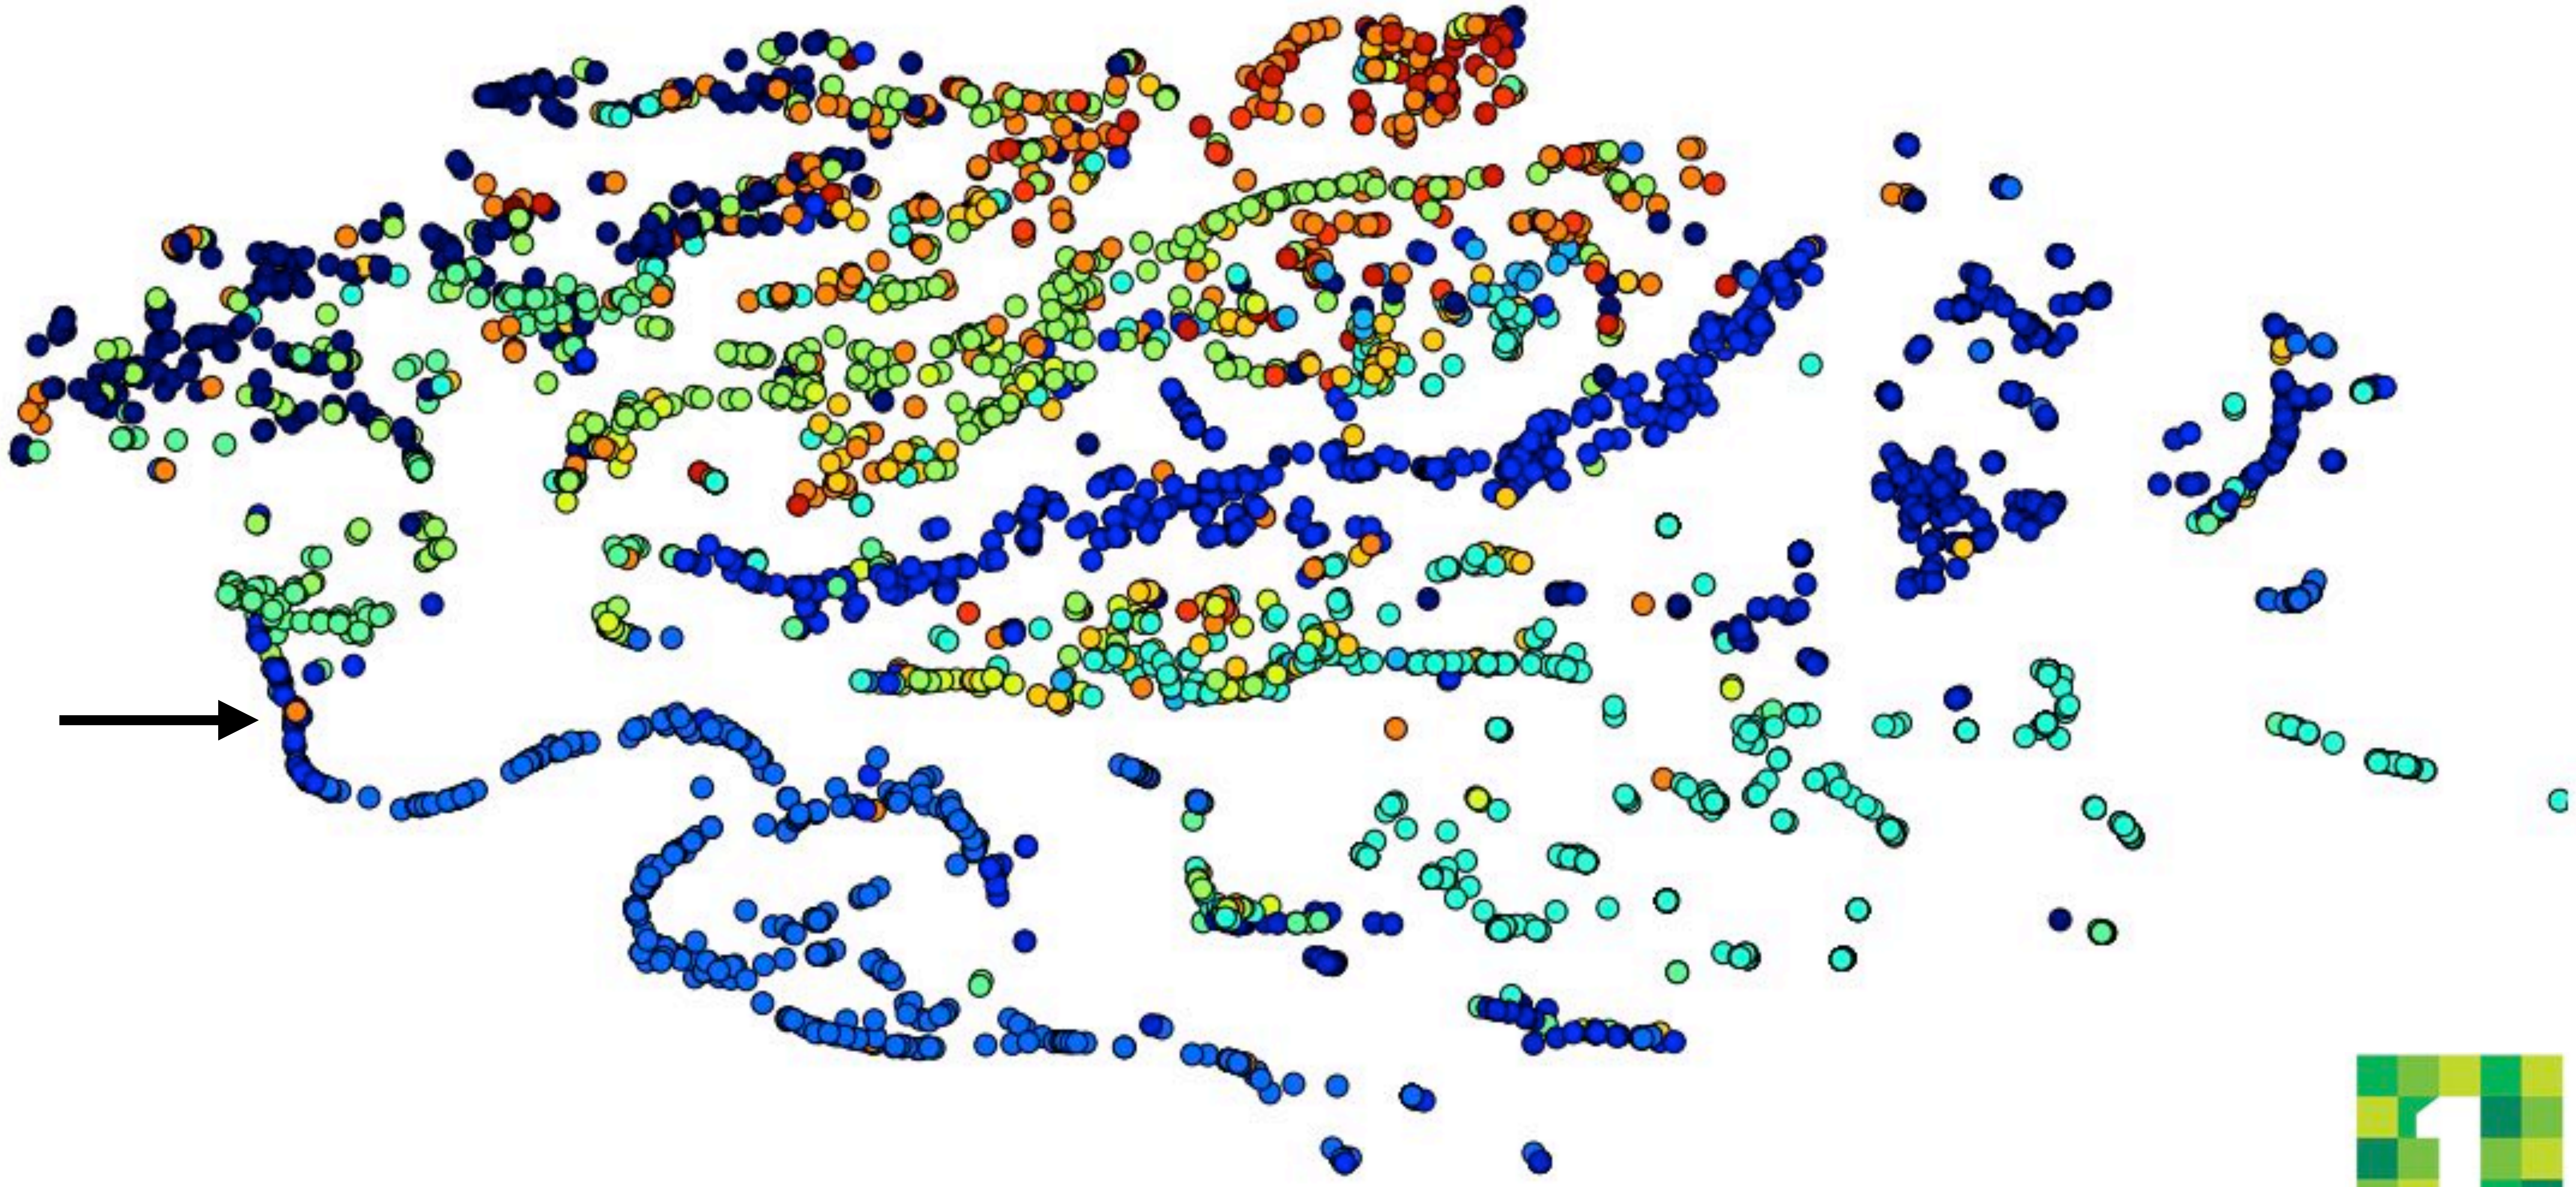

Deep factor model for financial time series

- **input**
  - time series of market data, index returns, macroeconomic data, fundamentals, …
- **feature learning**
  - the network autonomously learns predictive features from the input variables, simultaneously modeling covariate information and outcome
- **factor loadings**
  - latent factors are extracted using the learned features from the preceding layers
- **asset returns**
  - returns for a universe of financial instruments are obtained as a linear combination of the factor loadings

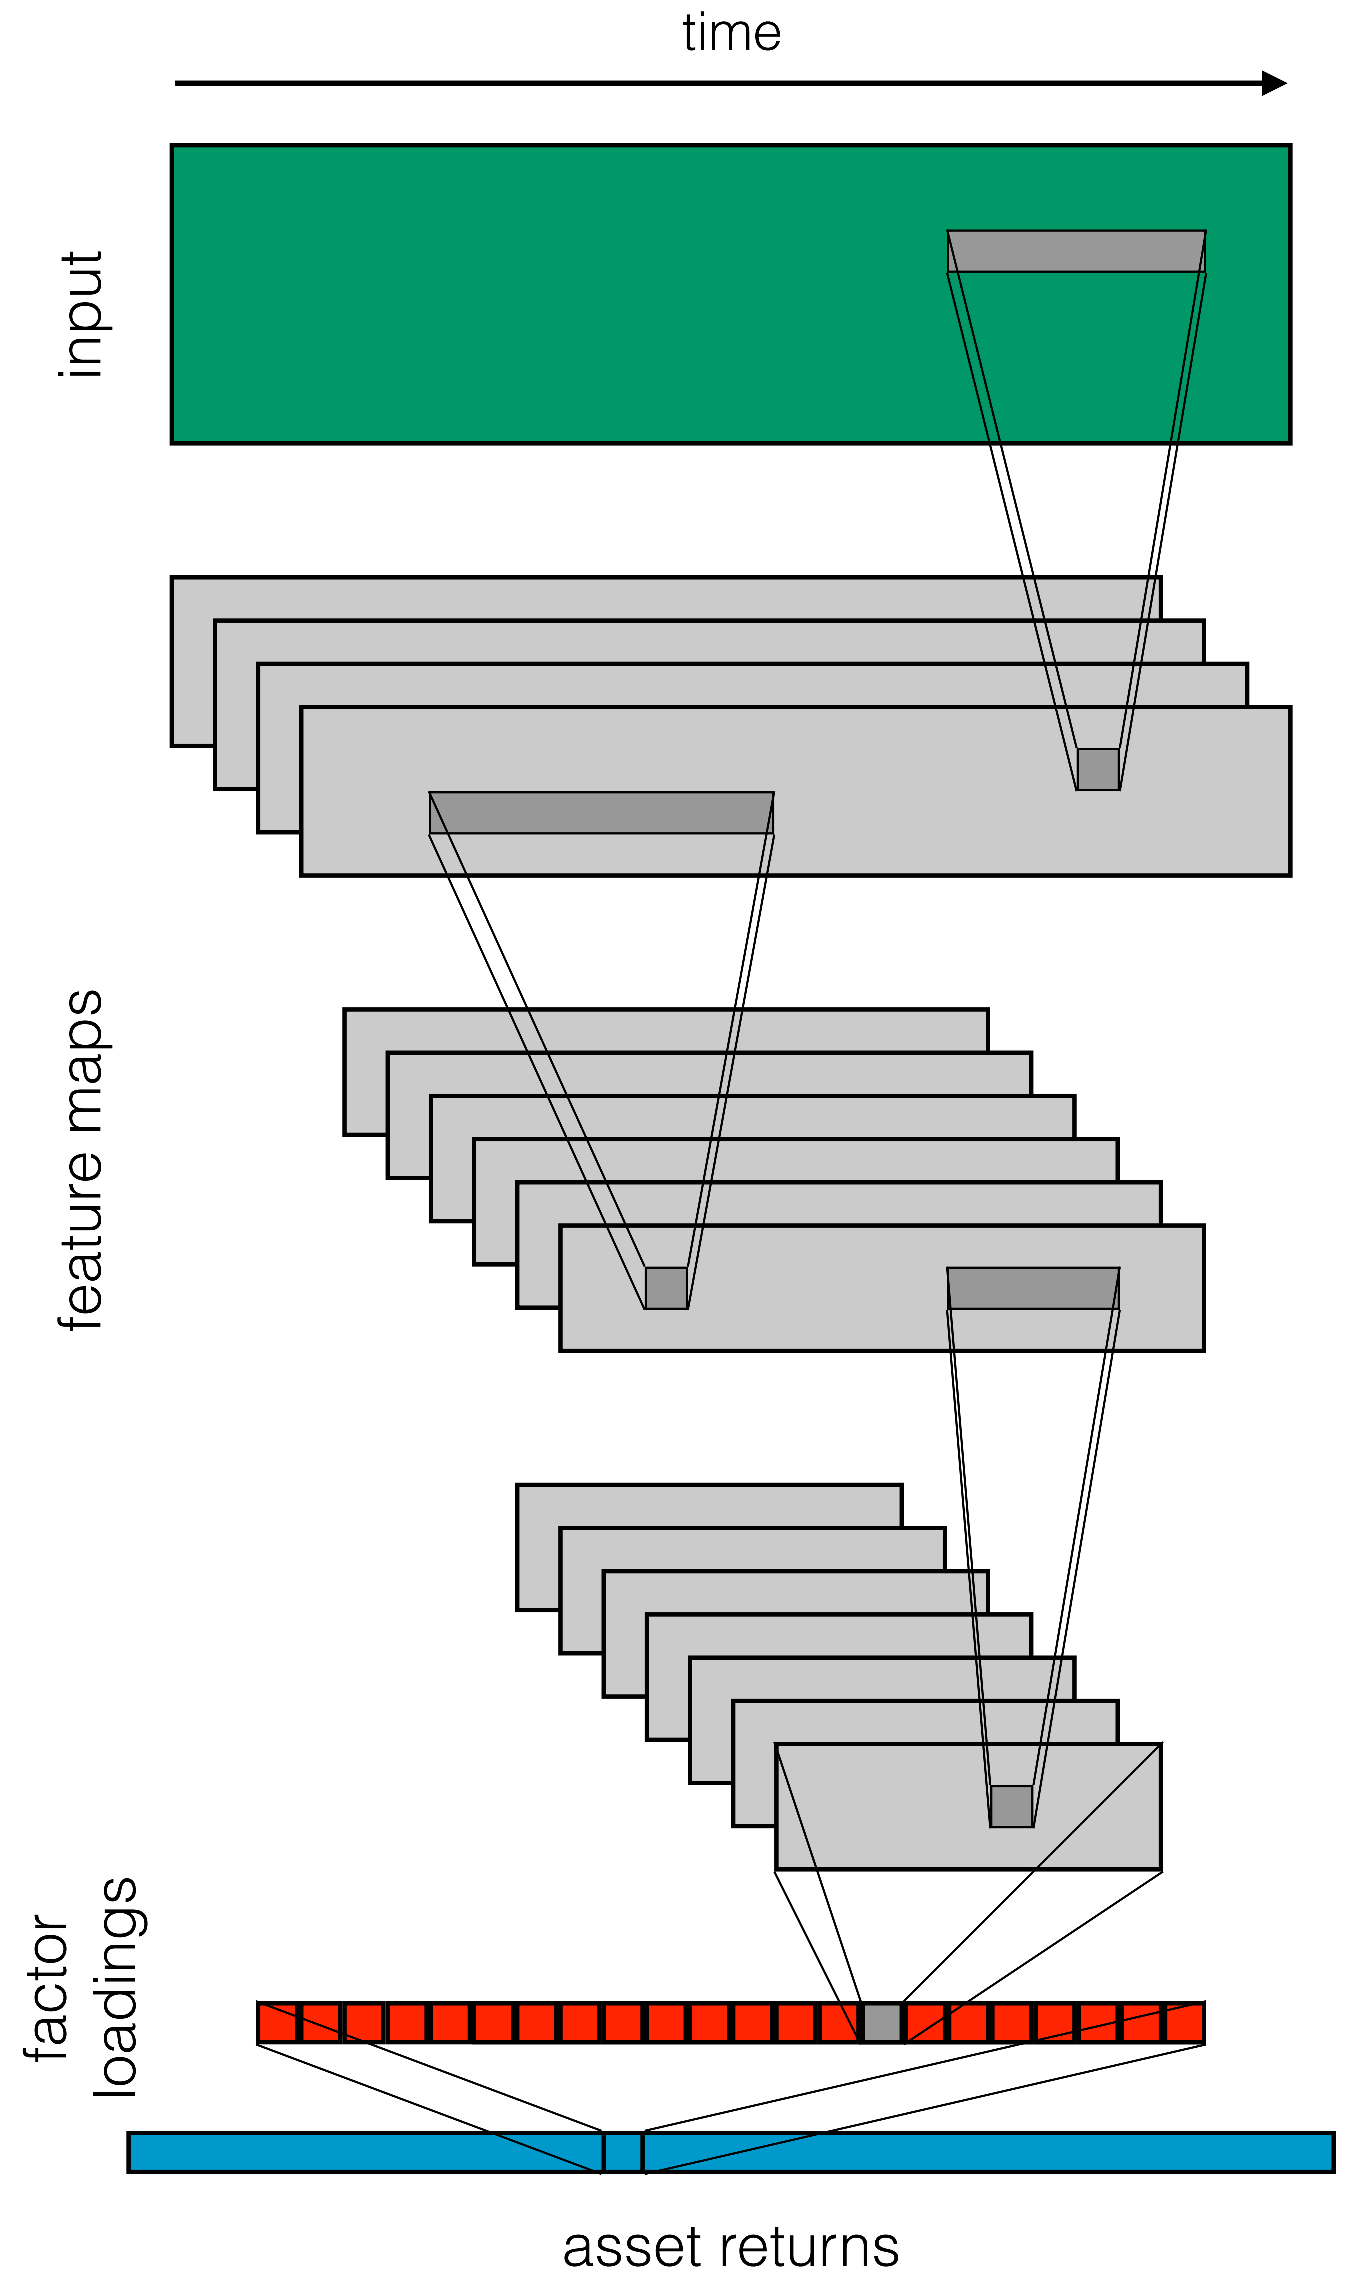

PREVEDEX

Supervised Learning

- External teacher provides "samples" of input and desired output.
  Example: Data from the past, existing users.

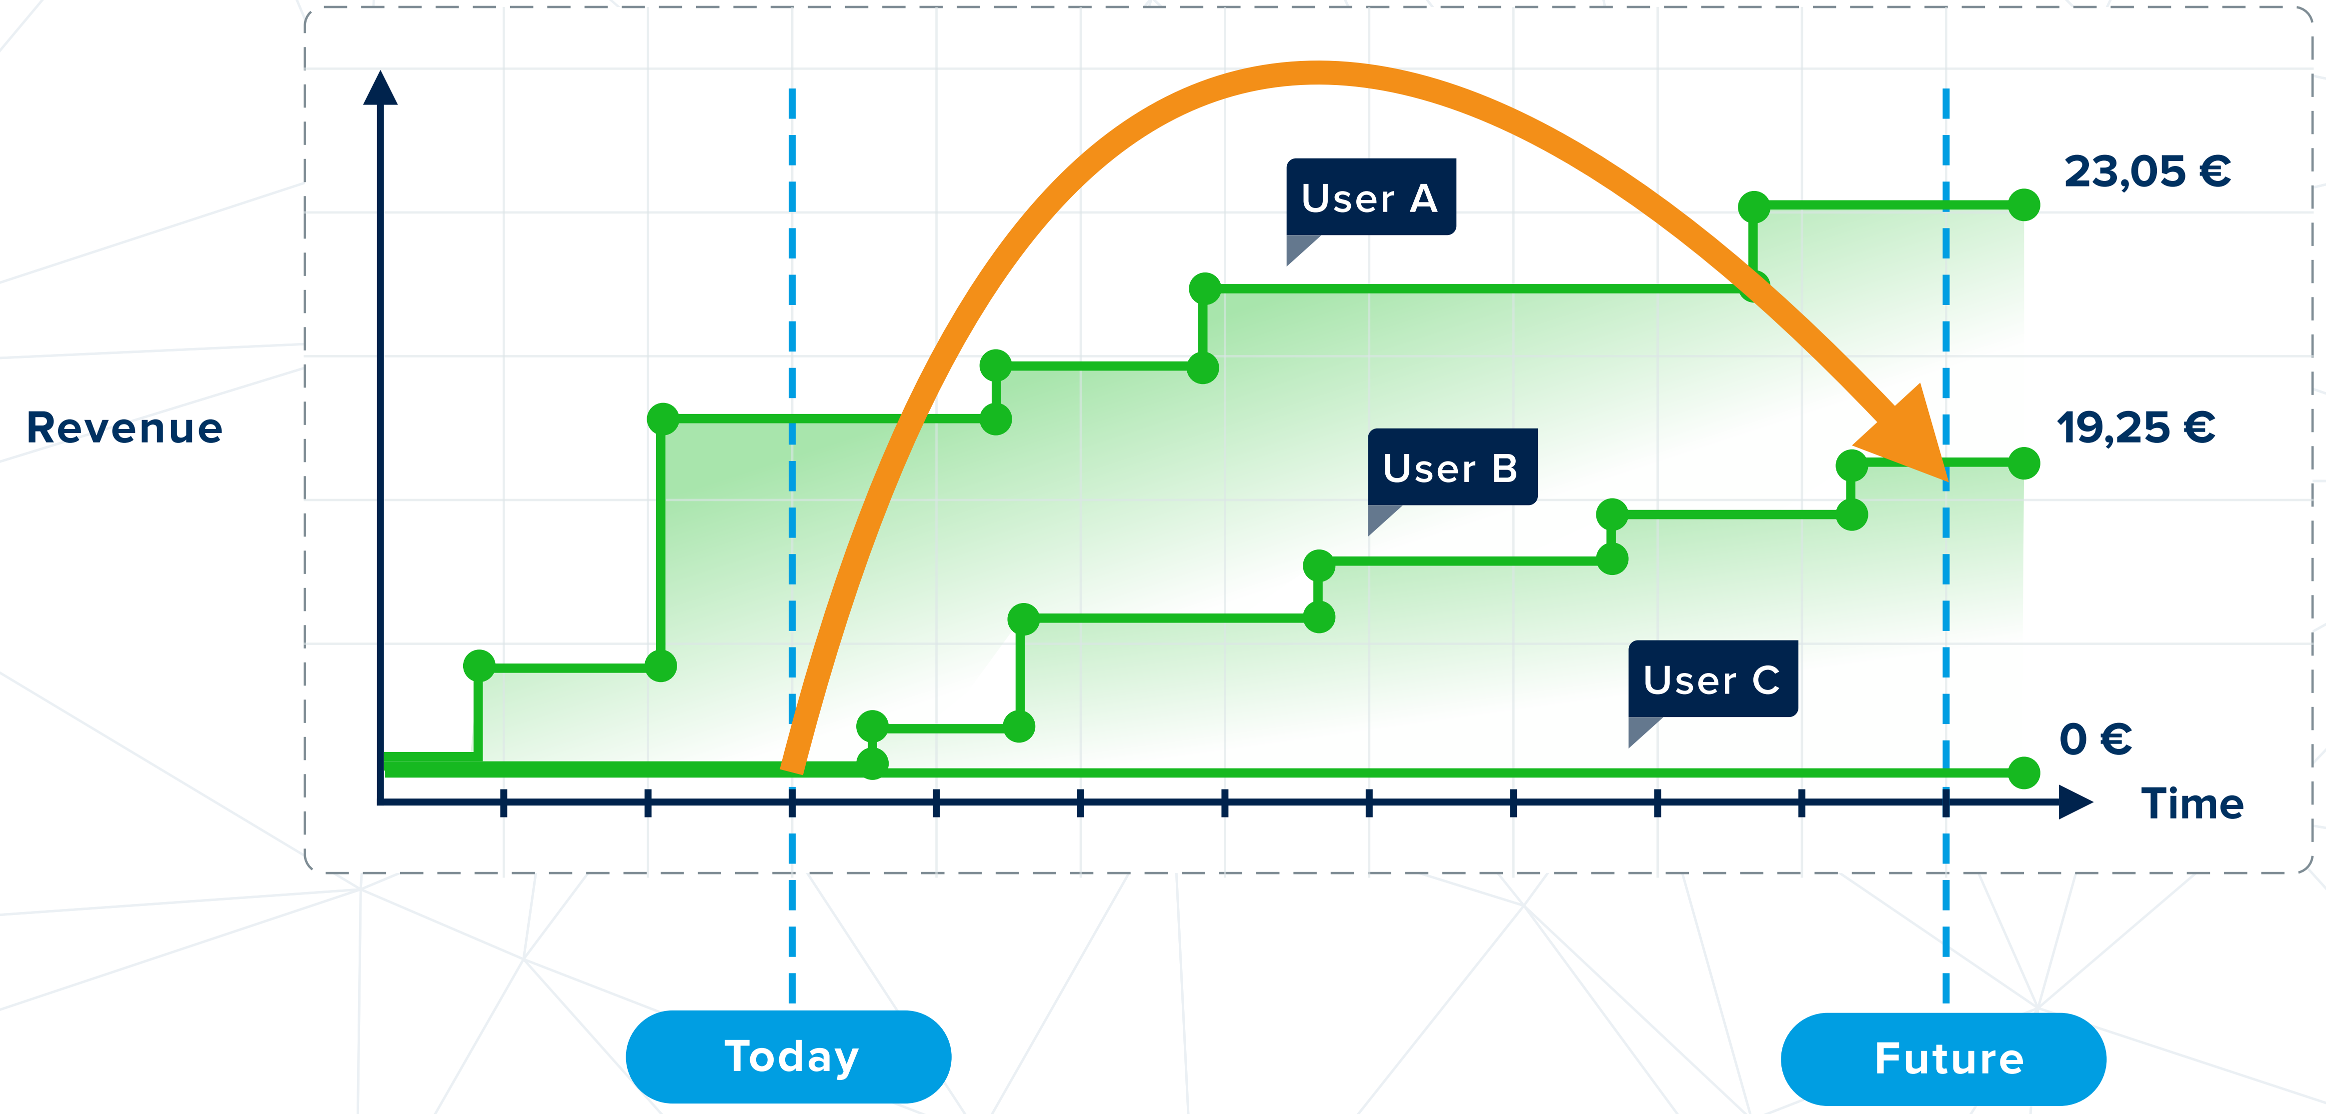
Revenue
User A
User B
User C
23,05 €
19,25 €
0 €
Time
Today
Future

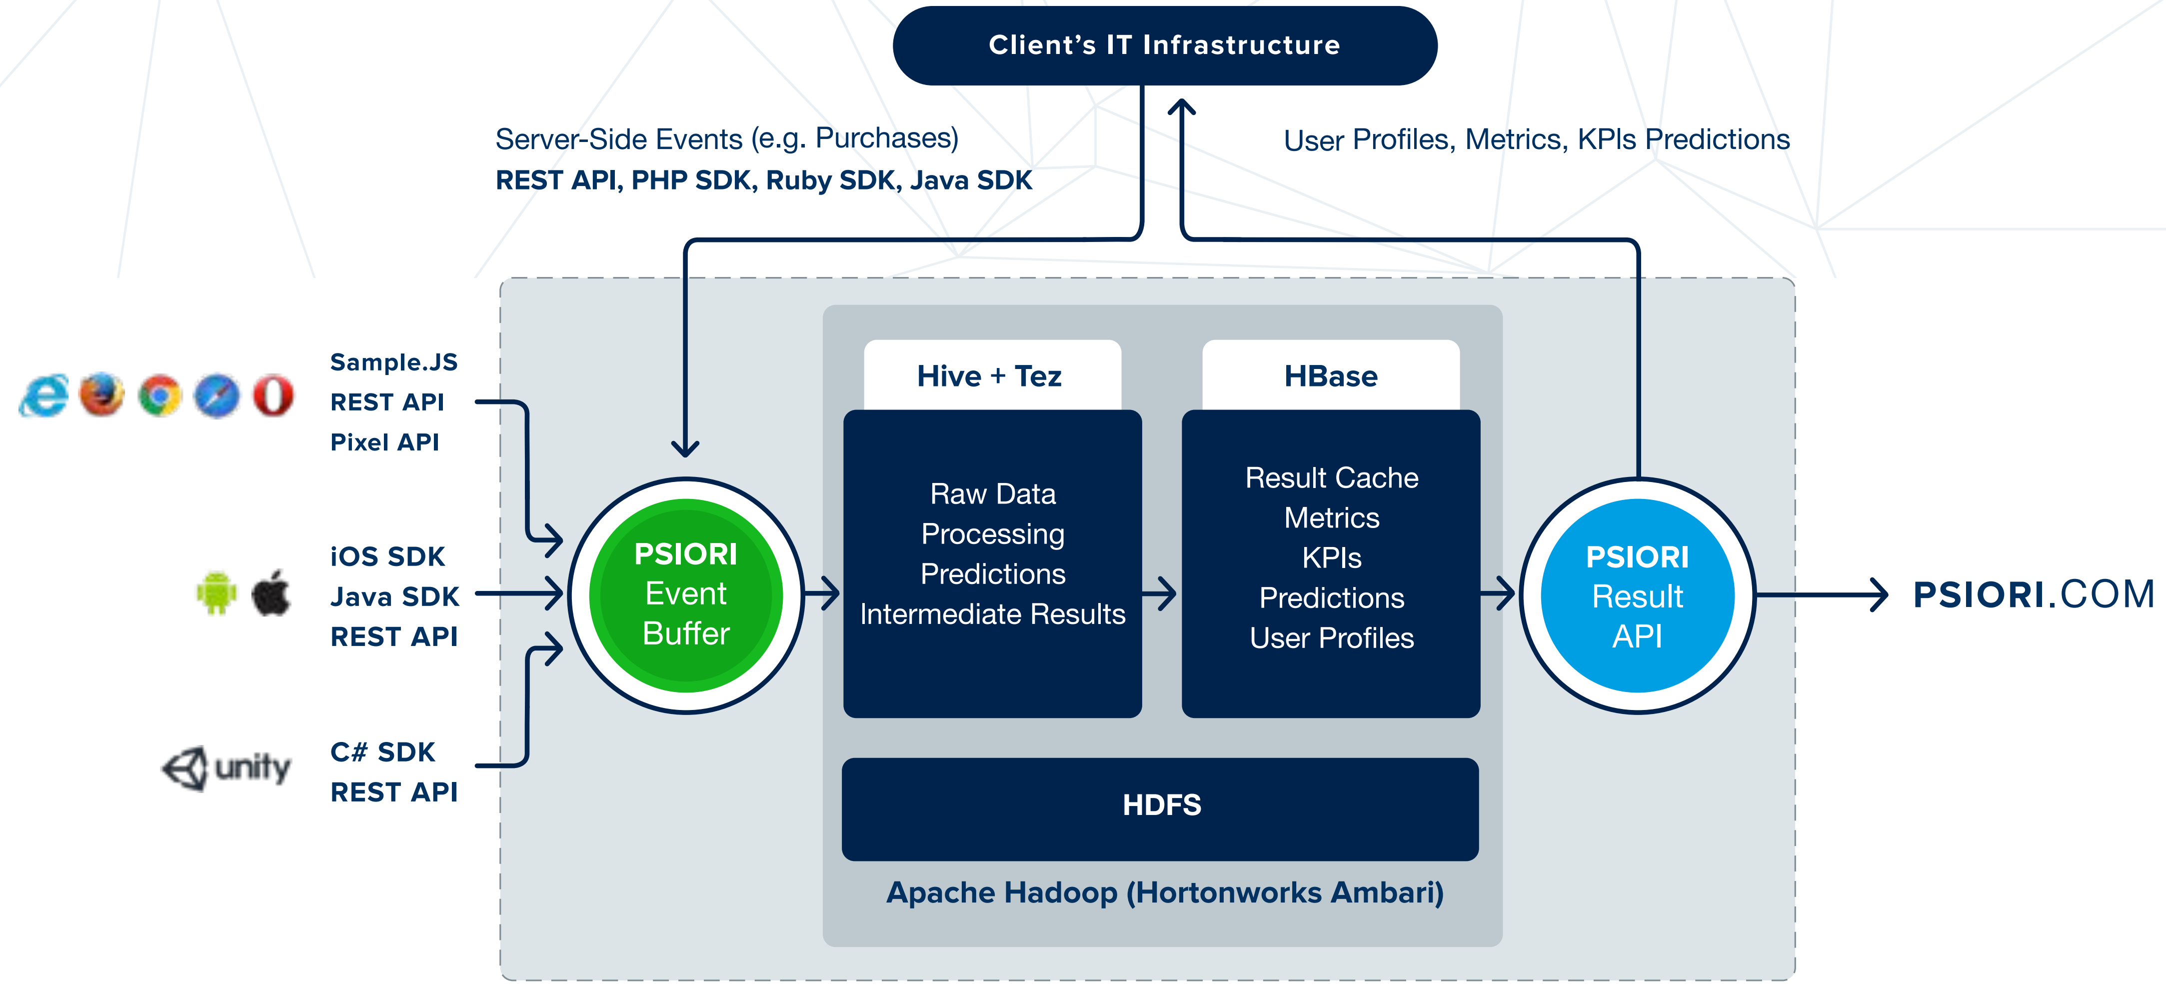
Client's IT Infrastructure
Server-Side Events (e.g. Purchases)
REST API, PHP SDK, Ruby SDK, Java SDK
User Profiles, Metrics, KPIs Predictions
Sample.JS
REST API
Pixel API
iOS SDK
Java SDK
REST API
unity
C# SDK
REST API
PSIORI Event Buffer
Hive + Tez
Raw Data
Processing
Predictions
Intermediate Results
HBase
Result Cache
Metrics
KPIs
Predictions
User Profiles
HDFS
Apache Hadoop (Hortonworks Ambari)
PSIORI Result API
PSIORI.COM


www.psiori.com
PSIORI
Insights
Blog
Help
Sascha Lange
Dashboard
Channels
Comparison
Cumulative
Overview
Info
Rights
1
10
100
1000
linear
log
Organic Customer Value 180
1,12 EUR
3,38 EUR
Non-Organic CV 180
0,93 EUR
3,47 EUR
facebook campaign #13
User 2013
CPC 0,76 €
Profit 1,27 €
increase budget
Referrer / Group
Organic
Registrations
9982 1.062
Profit
5636.75 € 8187.00 €
Referrer / Group
0013-00
demo-app
Budget
4.667,00 € 8.000,00 €
Registrations
240 17.062
Profit
2.18 € 8187.00 €

Machine Learning for Product Categorization

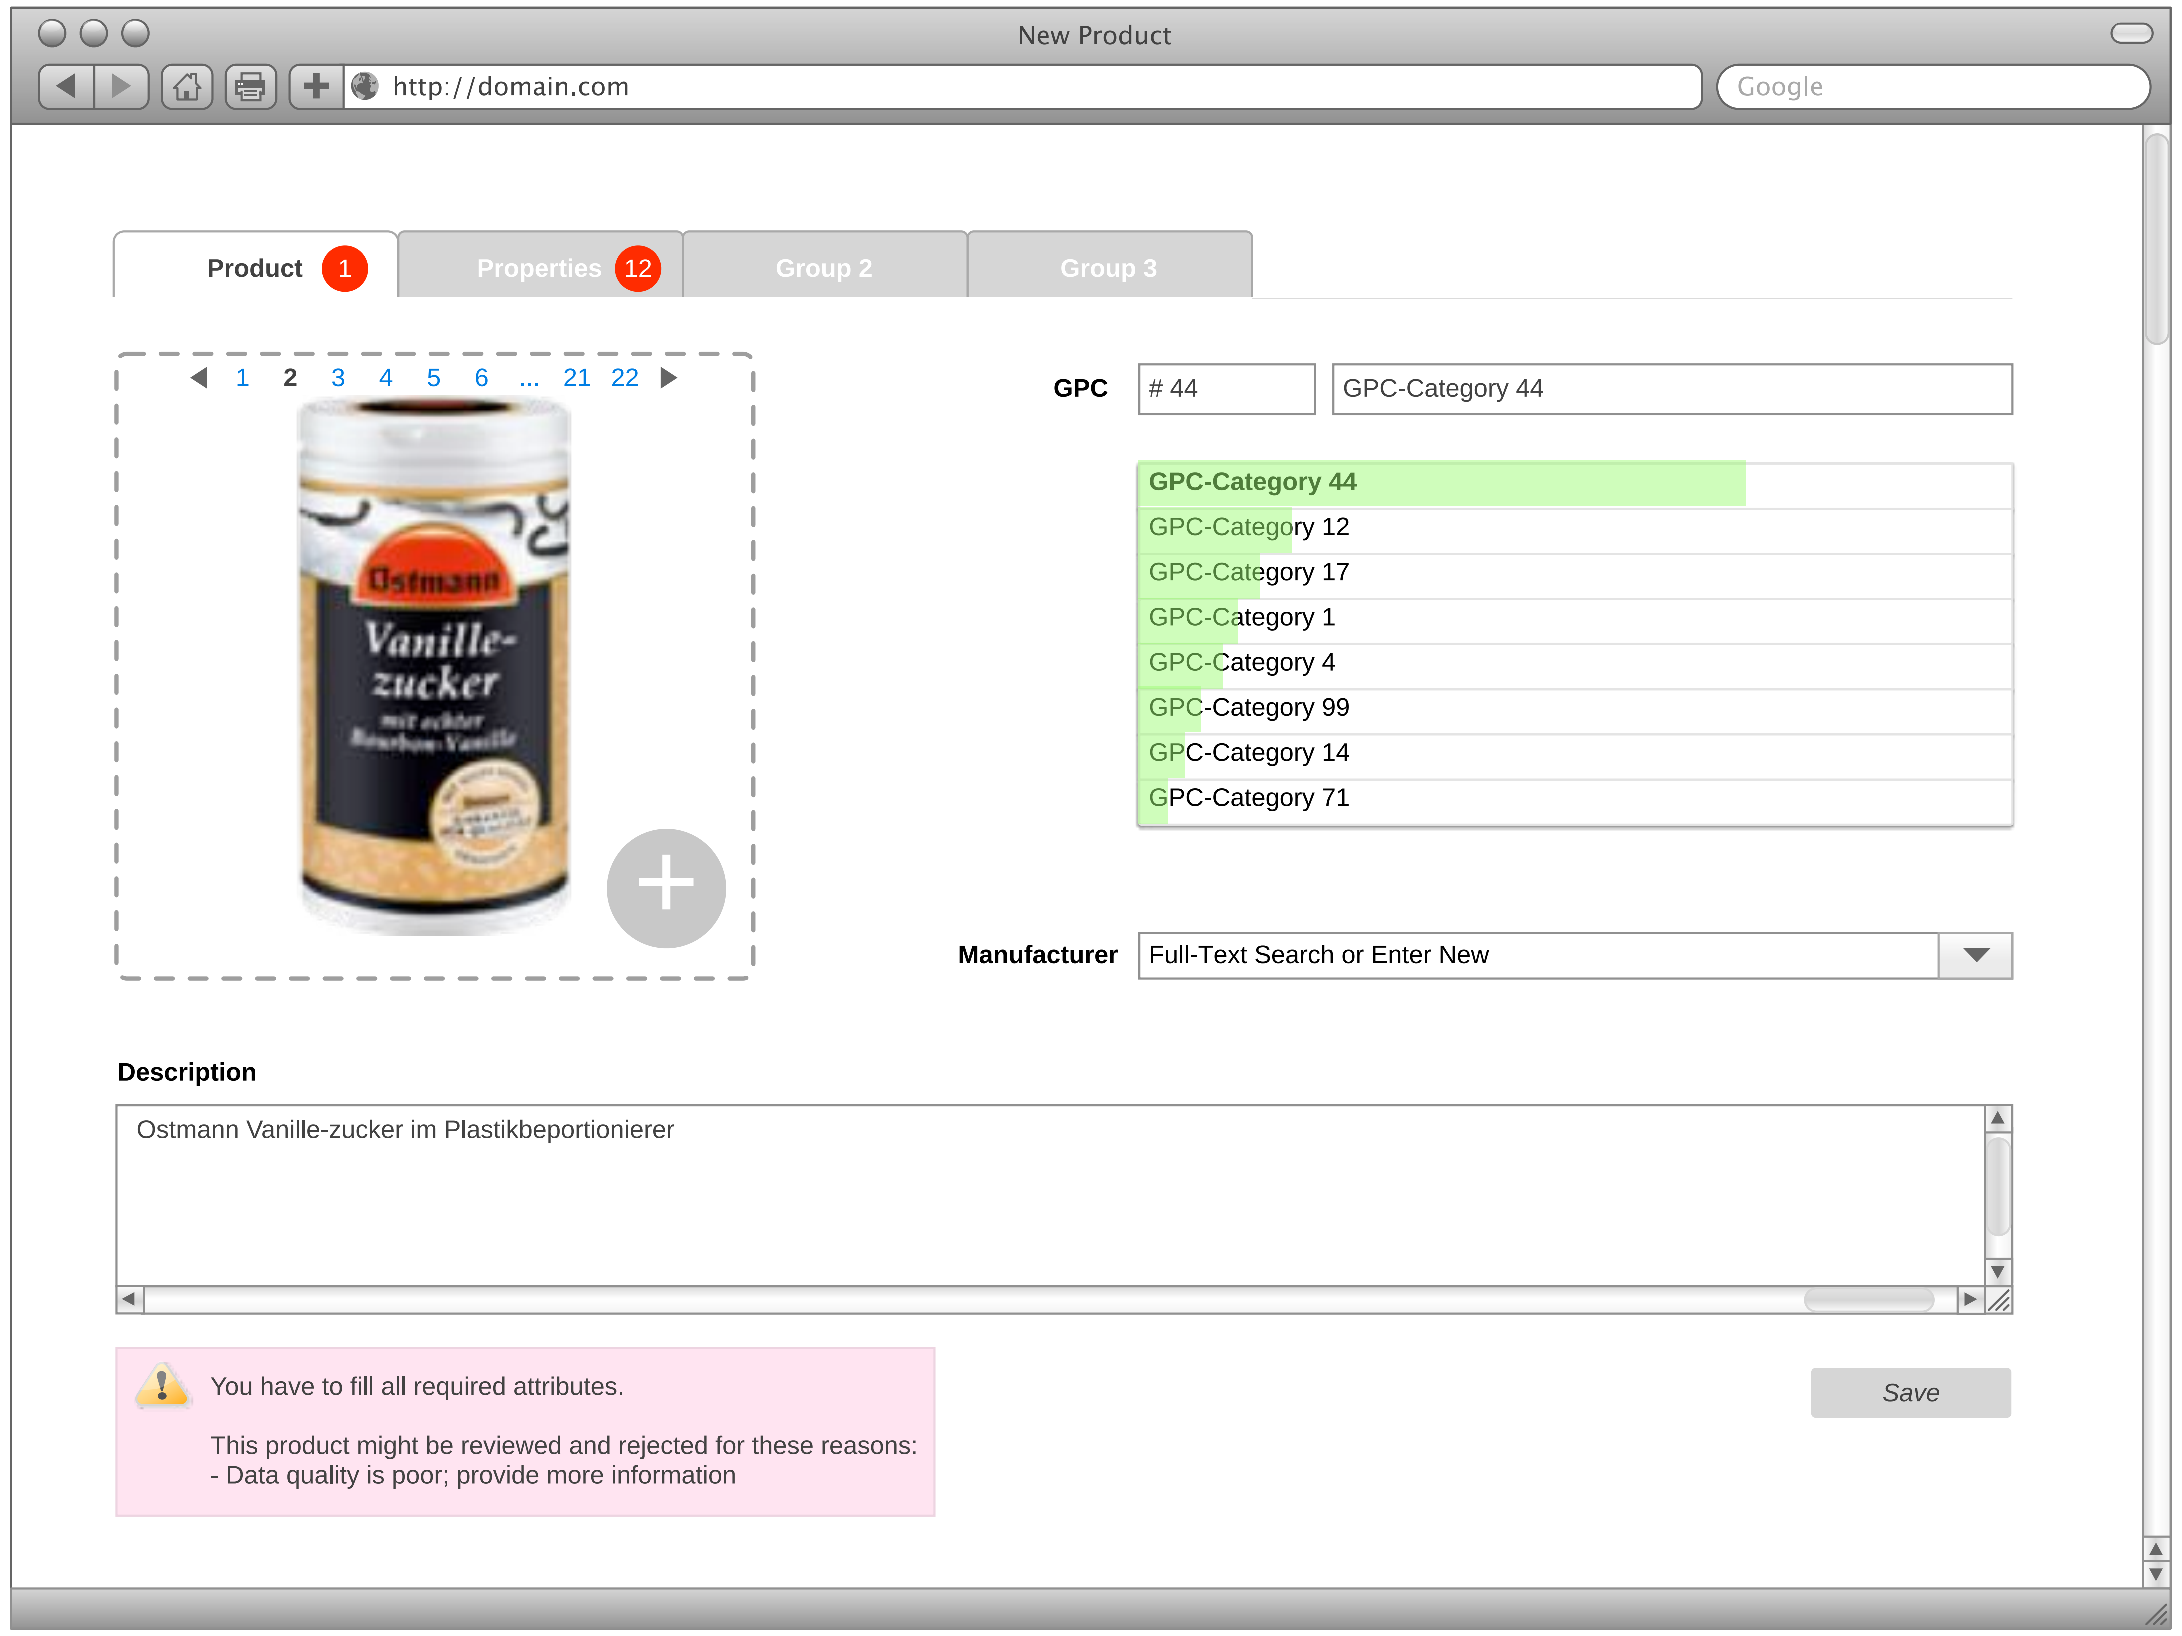

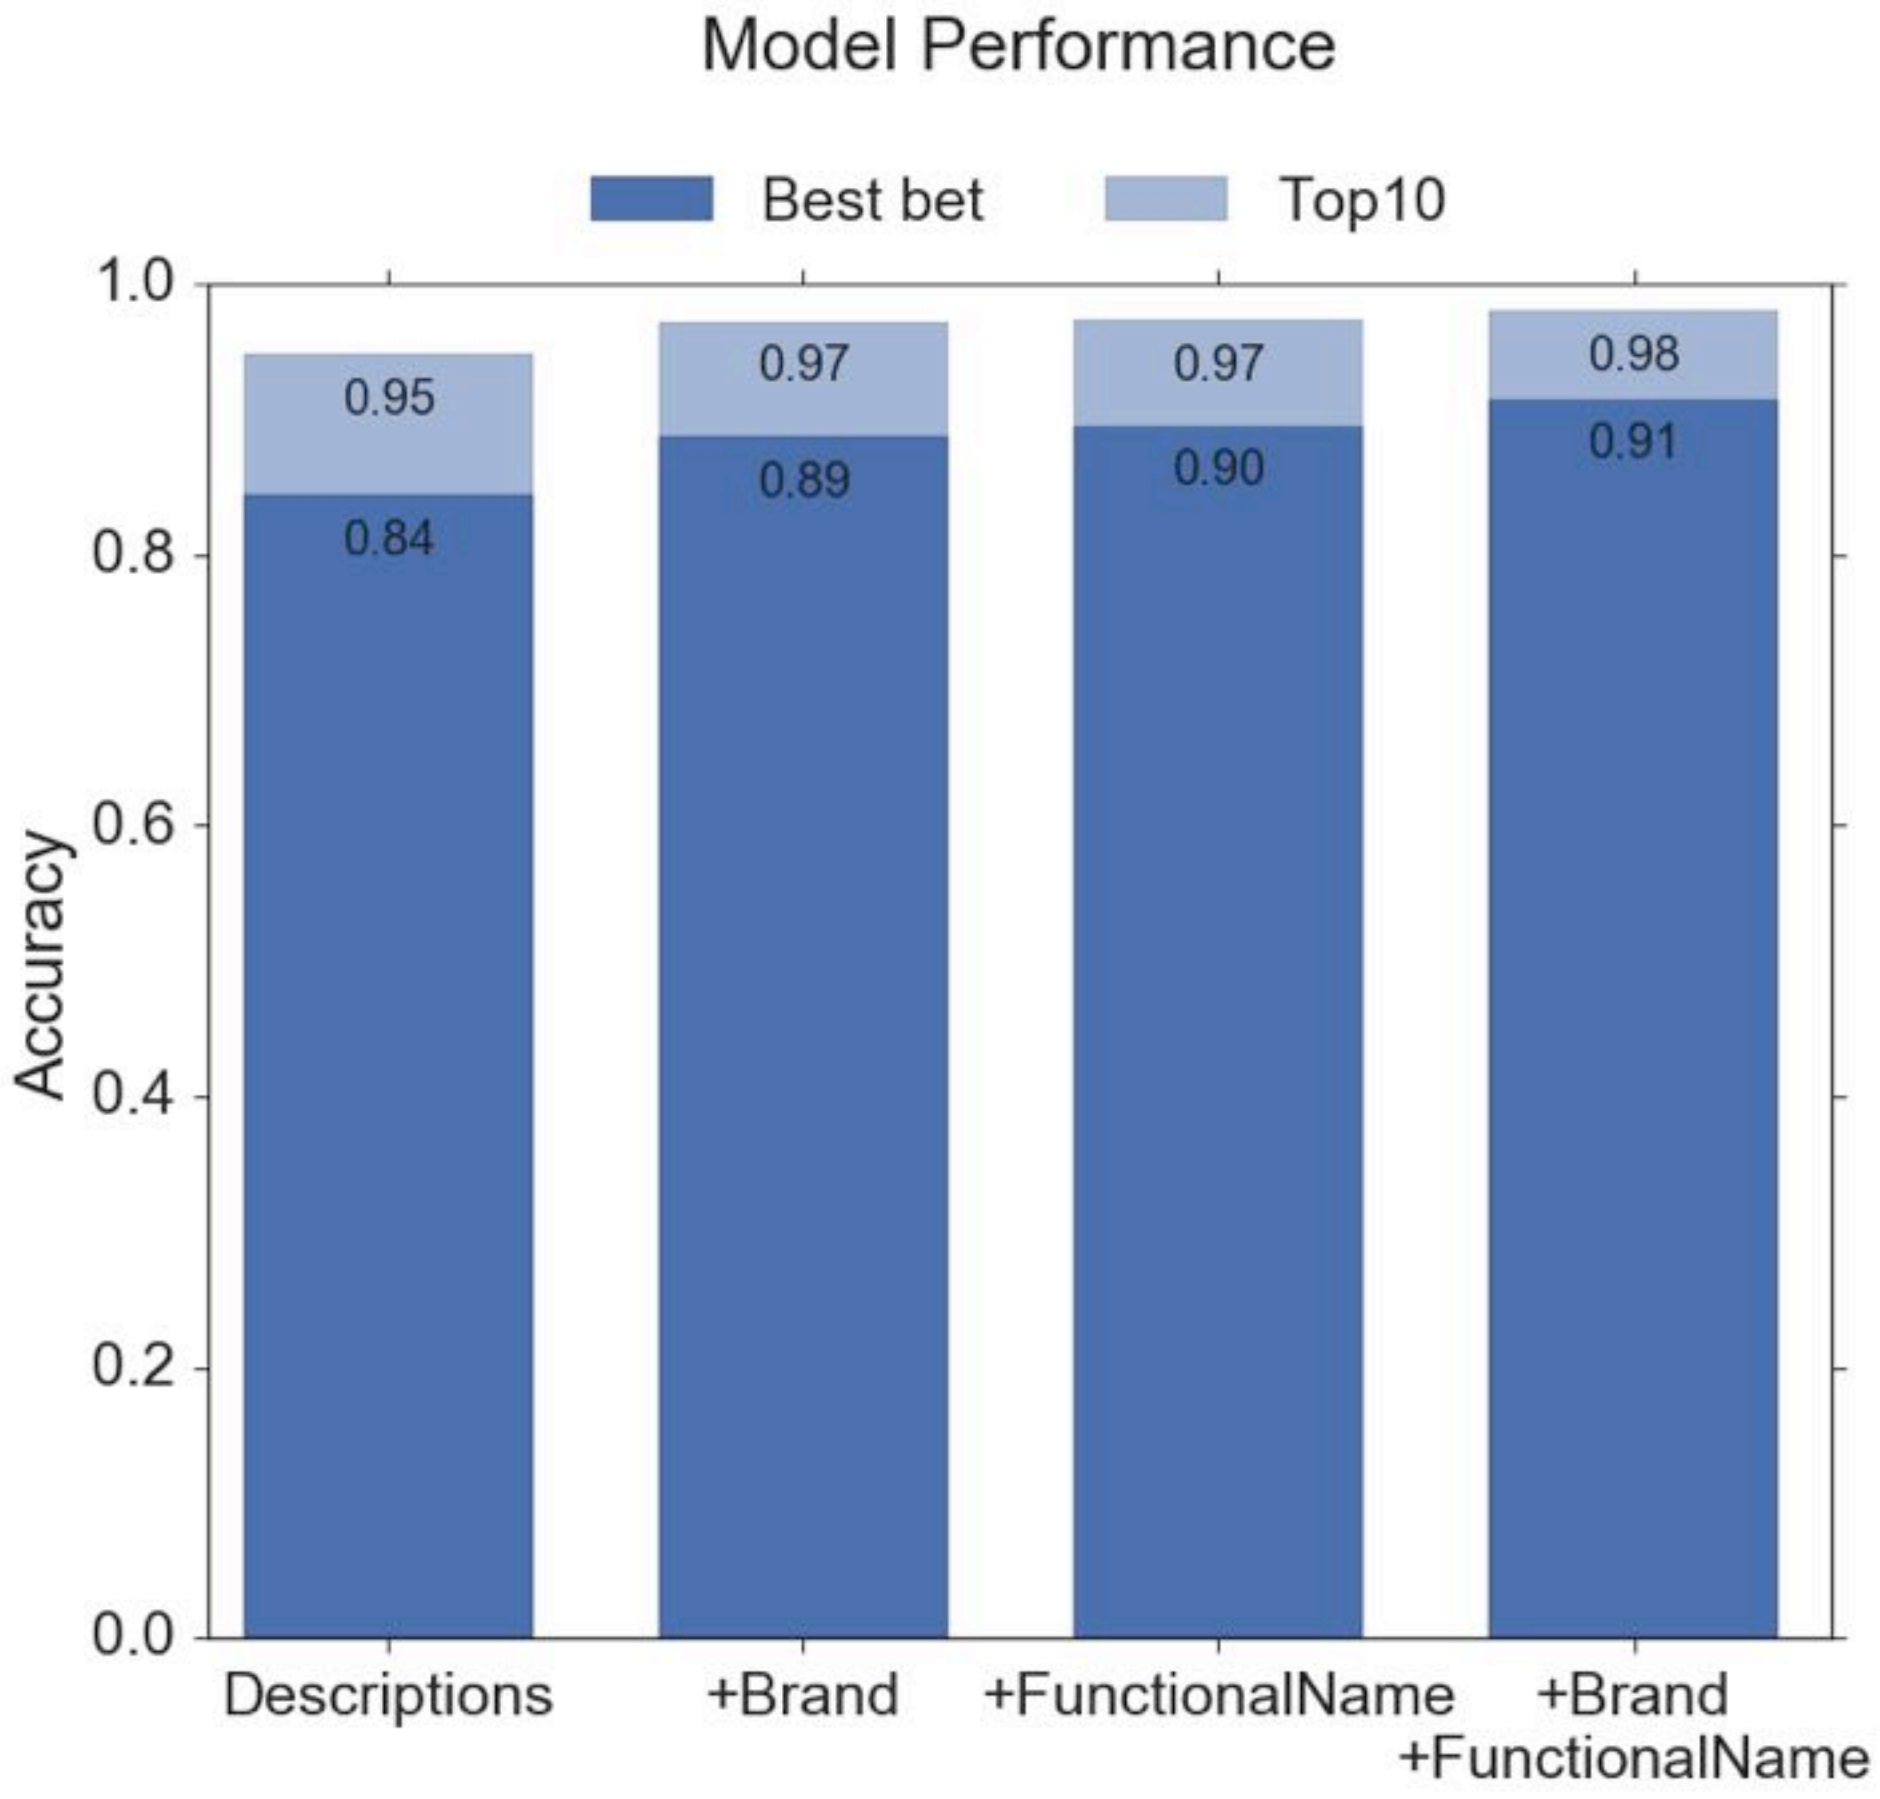

WORLDSYNC

SMARTDATAONE®
HIGH QUALITY PRODUCT INFORMATION

ThyssenKrupp Presta: Faulty Parts Prediction

- Predictive Model that selects flawed components very early, before delivery to customer and final assembly, saves ~ $25k/week
- Input: raw sensor data from robots, Target: quality of manufactured part

- Accepting 2% false positives → detect 25% of bad parts
- Accepting 10% false positives → detect 67% of bad parts

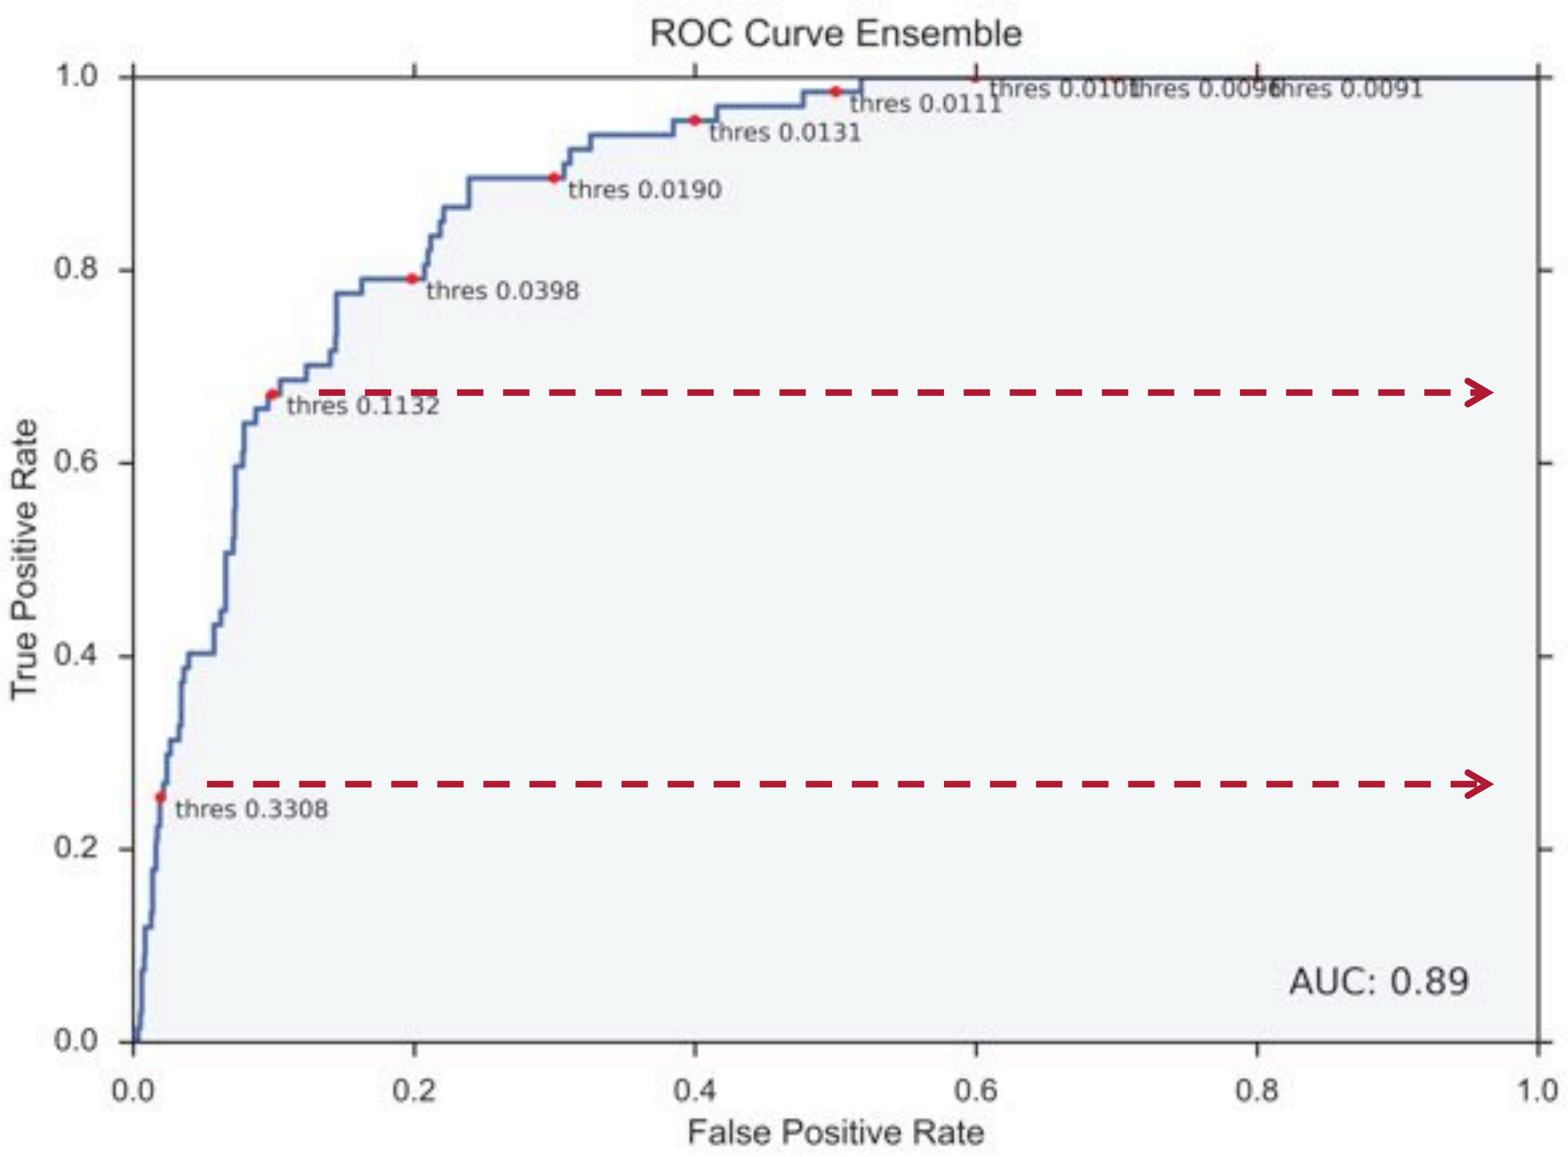

	Pred Neg	Pred Pos
Cond Neg	1146	22
Cond Pos	50	17

With thres 0.3308

	Pred Neg	Pred Pos
Cond Neg	1052	116
Cond Pos	22	45

With thres 0.1132

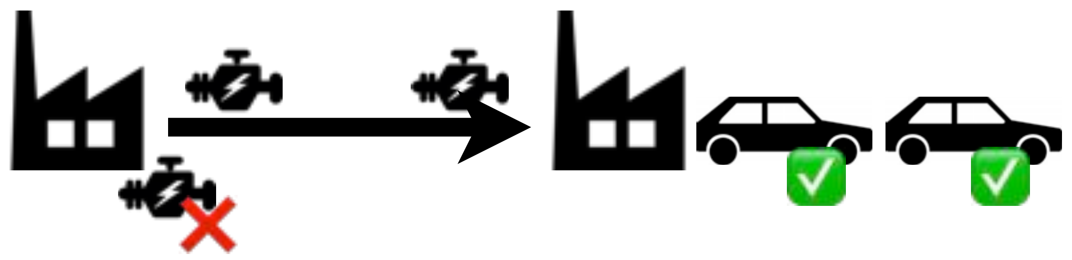

Predictive Maintenance & Asset Health Index

Goal: anticipate end-of-life of critical machines

- Quantify the health status based on sensor data or inspection protocols
- Predict remaining lifespan to prepare replacement
- Detect abuse and measure effect of maintenance

Reference: Wood processing tools, expensive and delicate

- Customers report breakdowns
- Manufacturer analyzes usage and predicts asset health, interested in improper use

HOMAG GROUP

pmOne

LipSync Face Tracking & Transfer

- Marker-less motion capturing
- Track head pose and facial features
- Improves on open source algorithms

metricminds

Audi

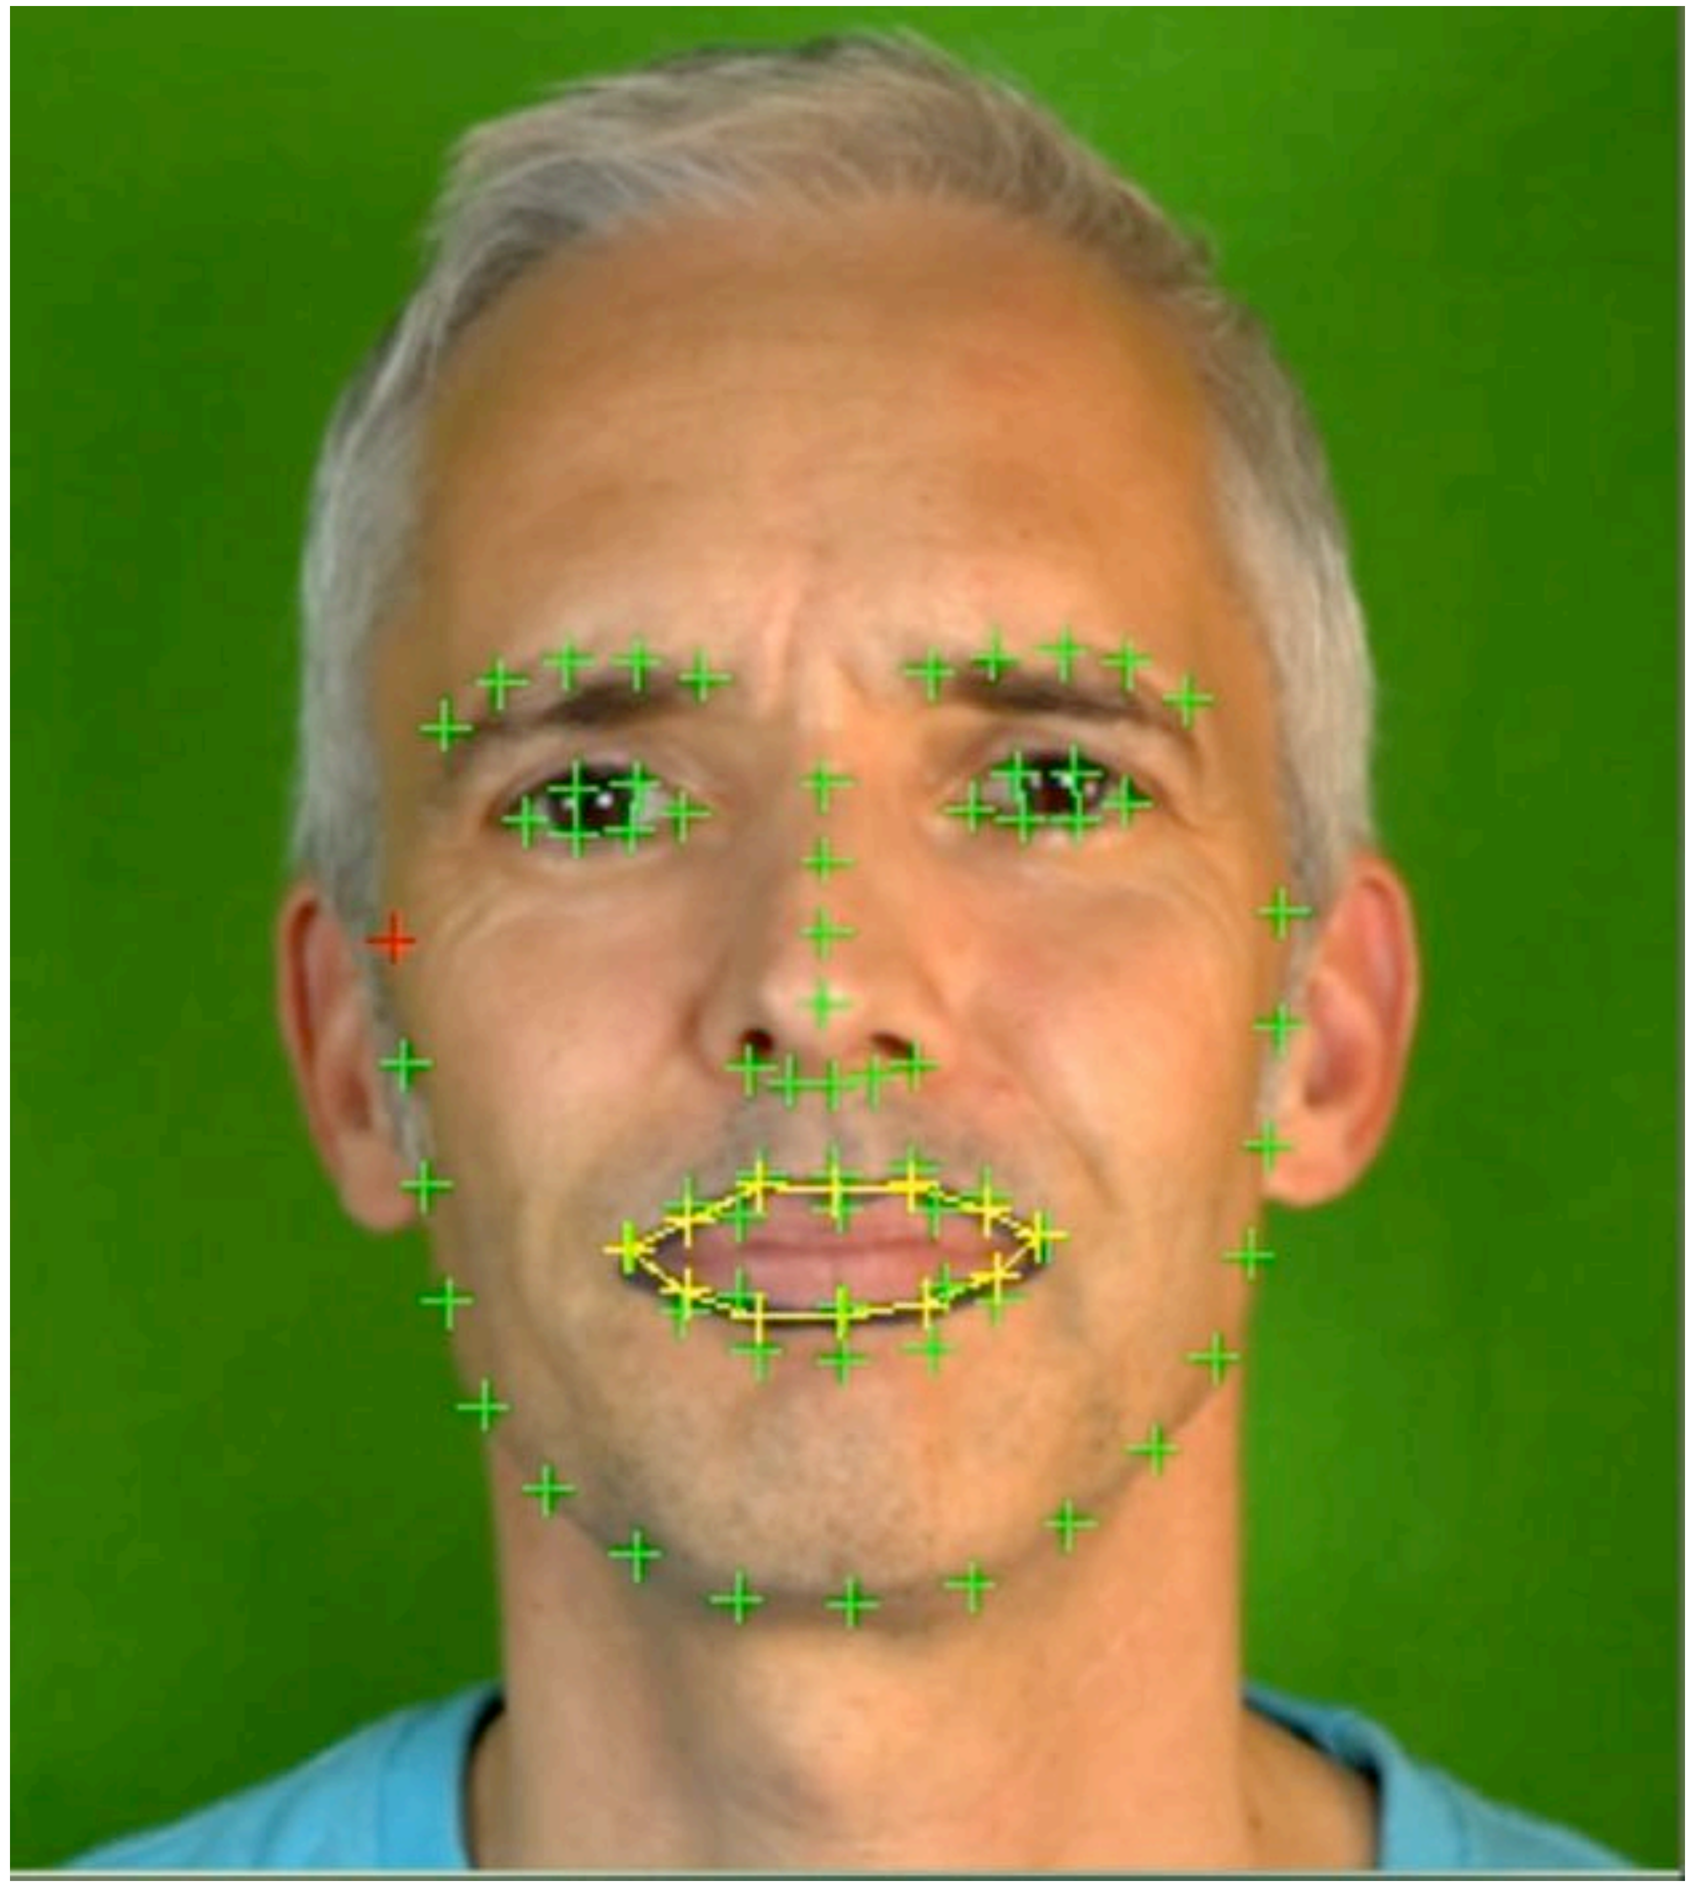

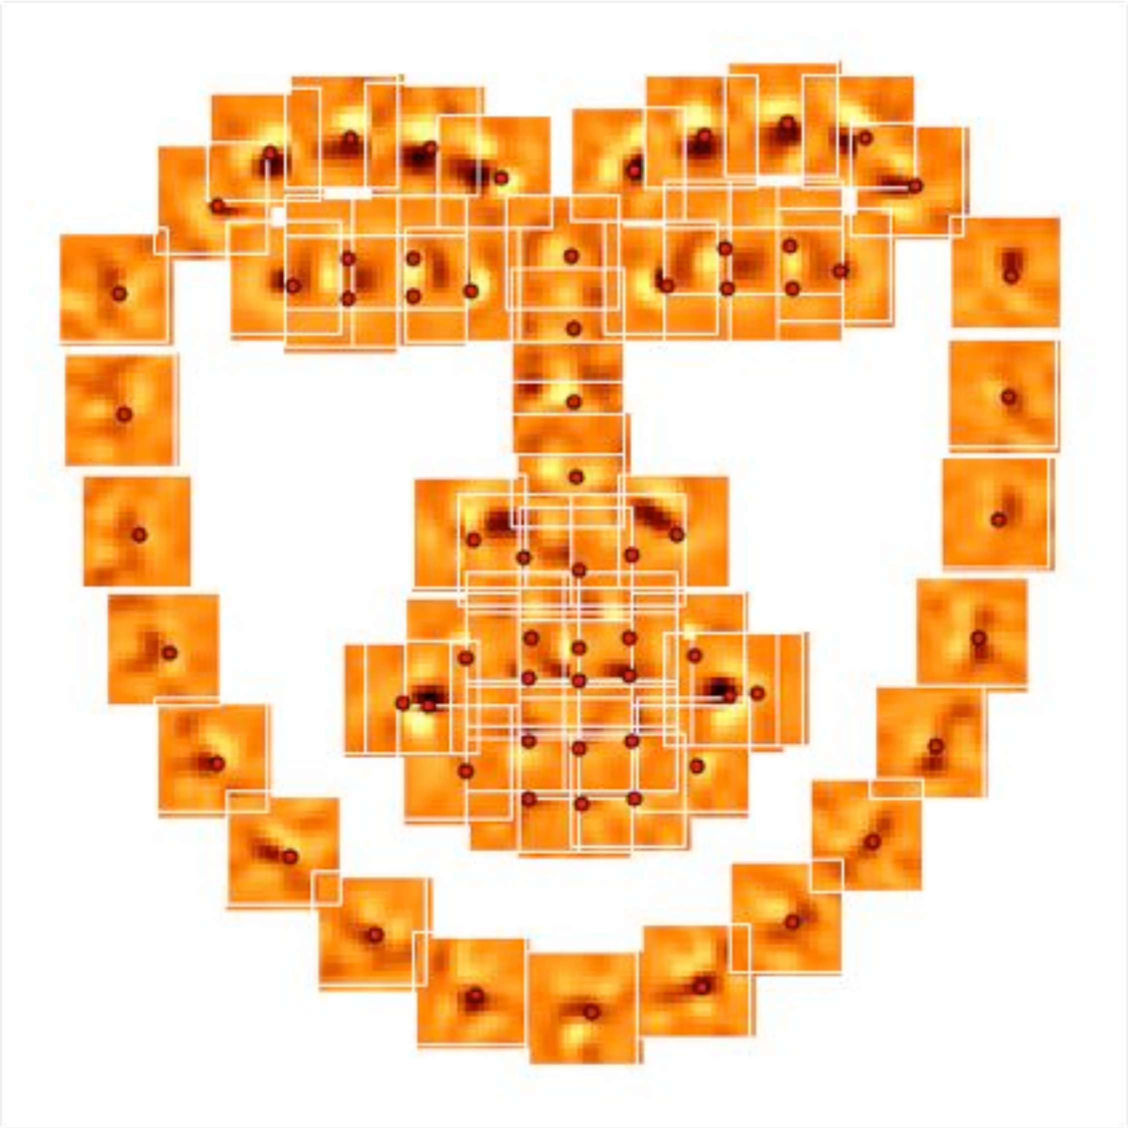

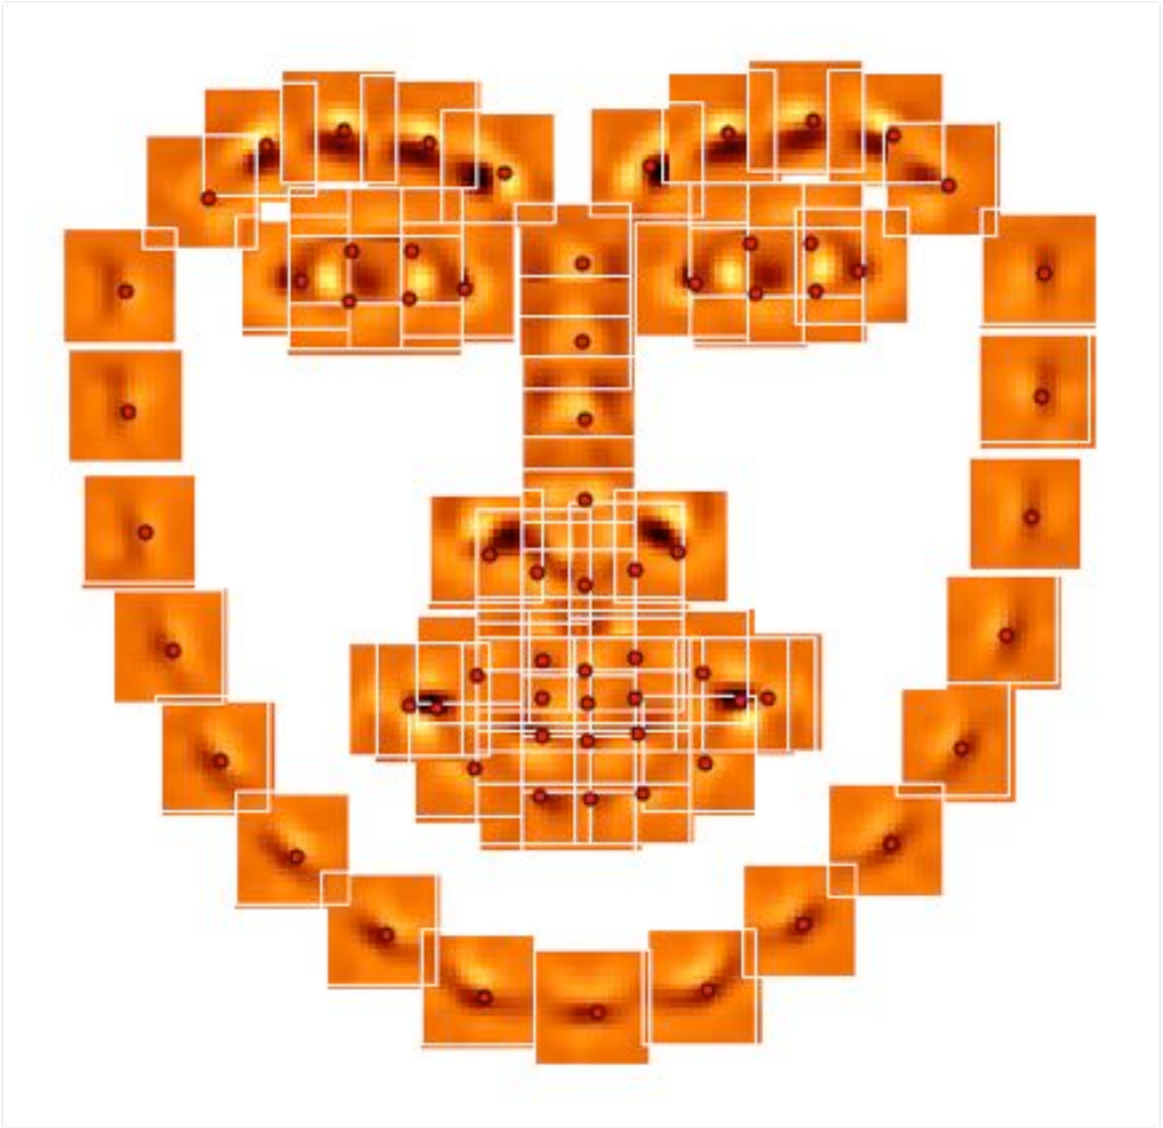

better → Deep Conv Nets

(best open source) (our current state)


Audi RS7 Sportback

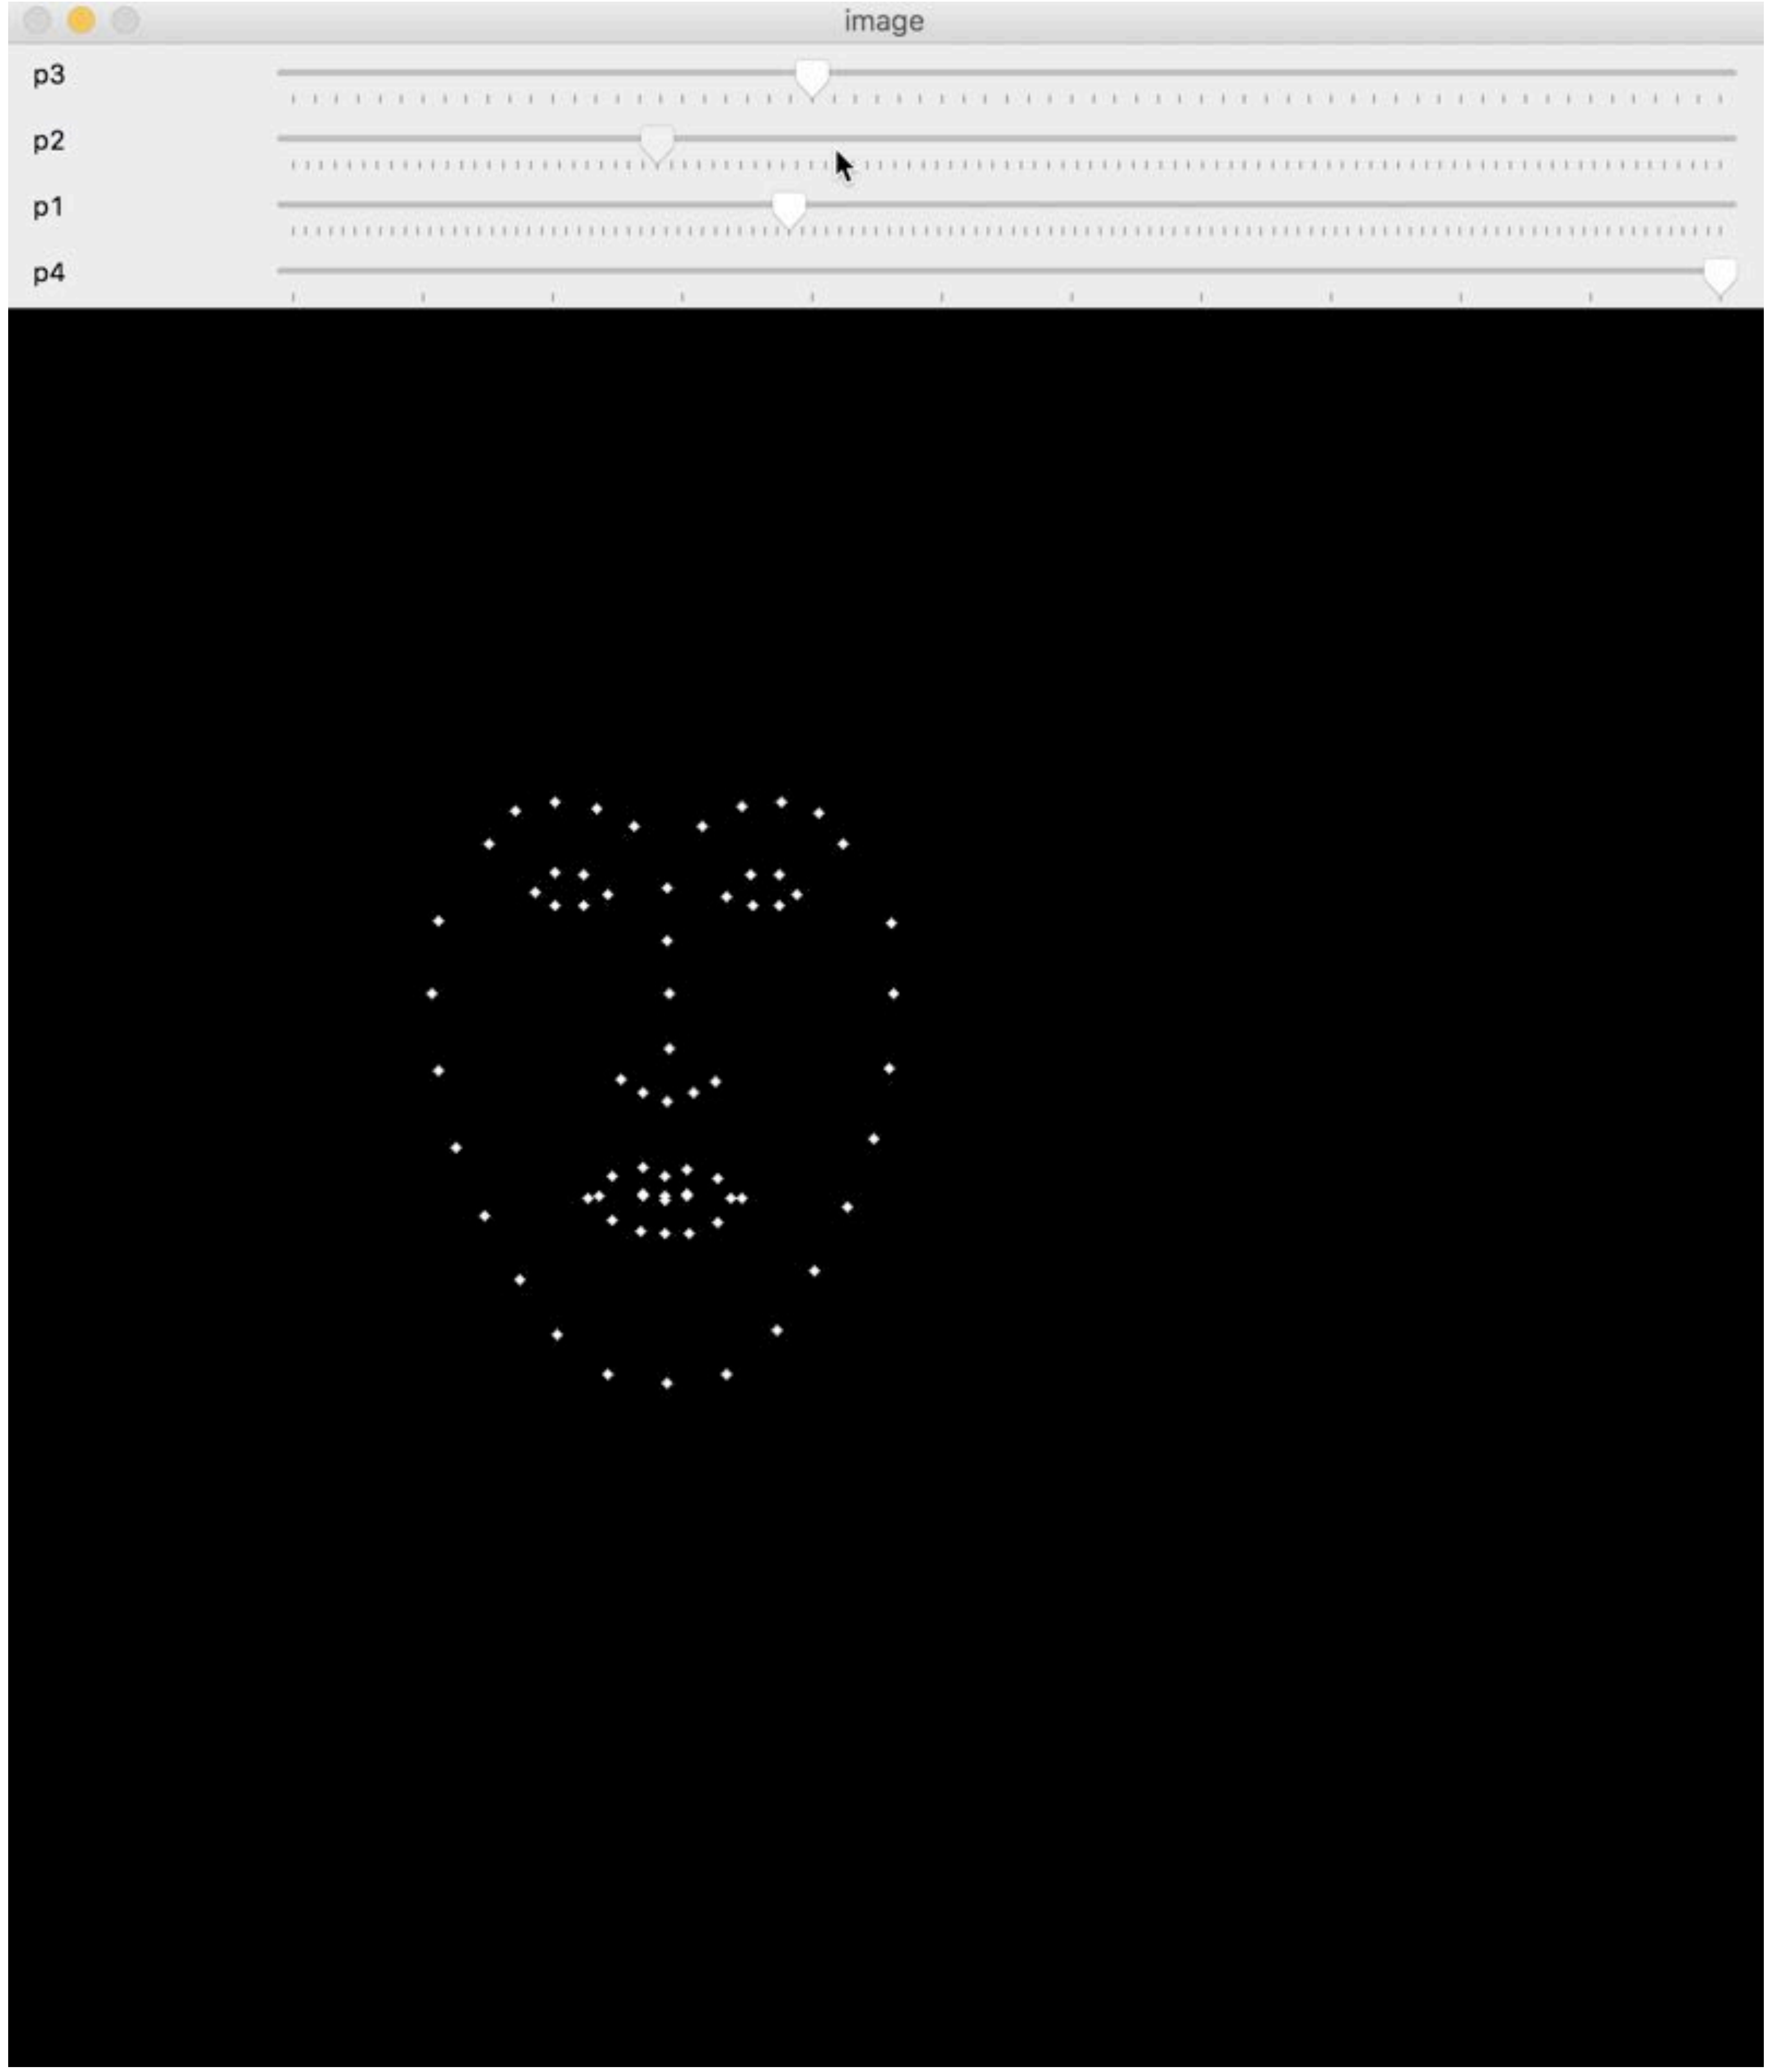
image
p3
p2
p1
p4

Deep Learning for Multivariate Time Series Prediction

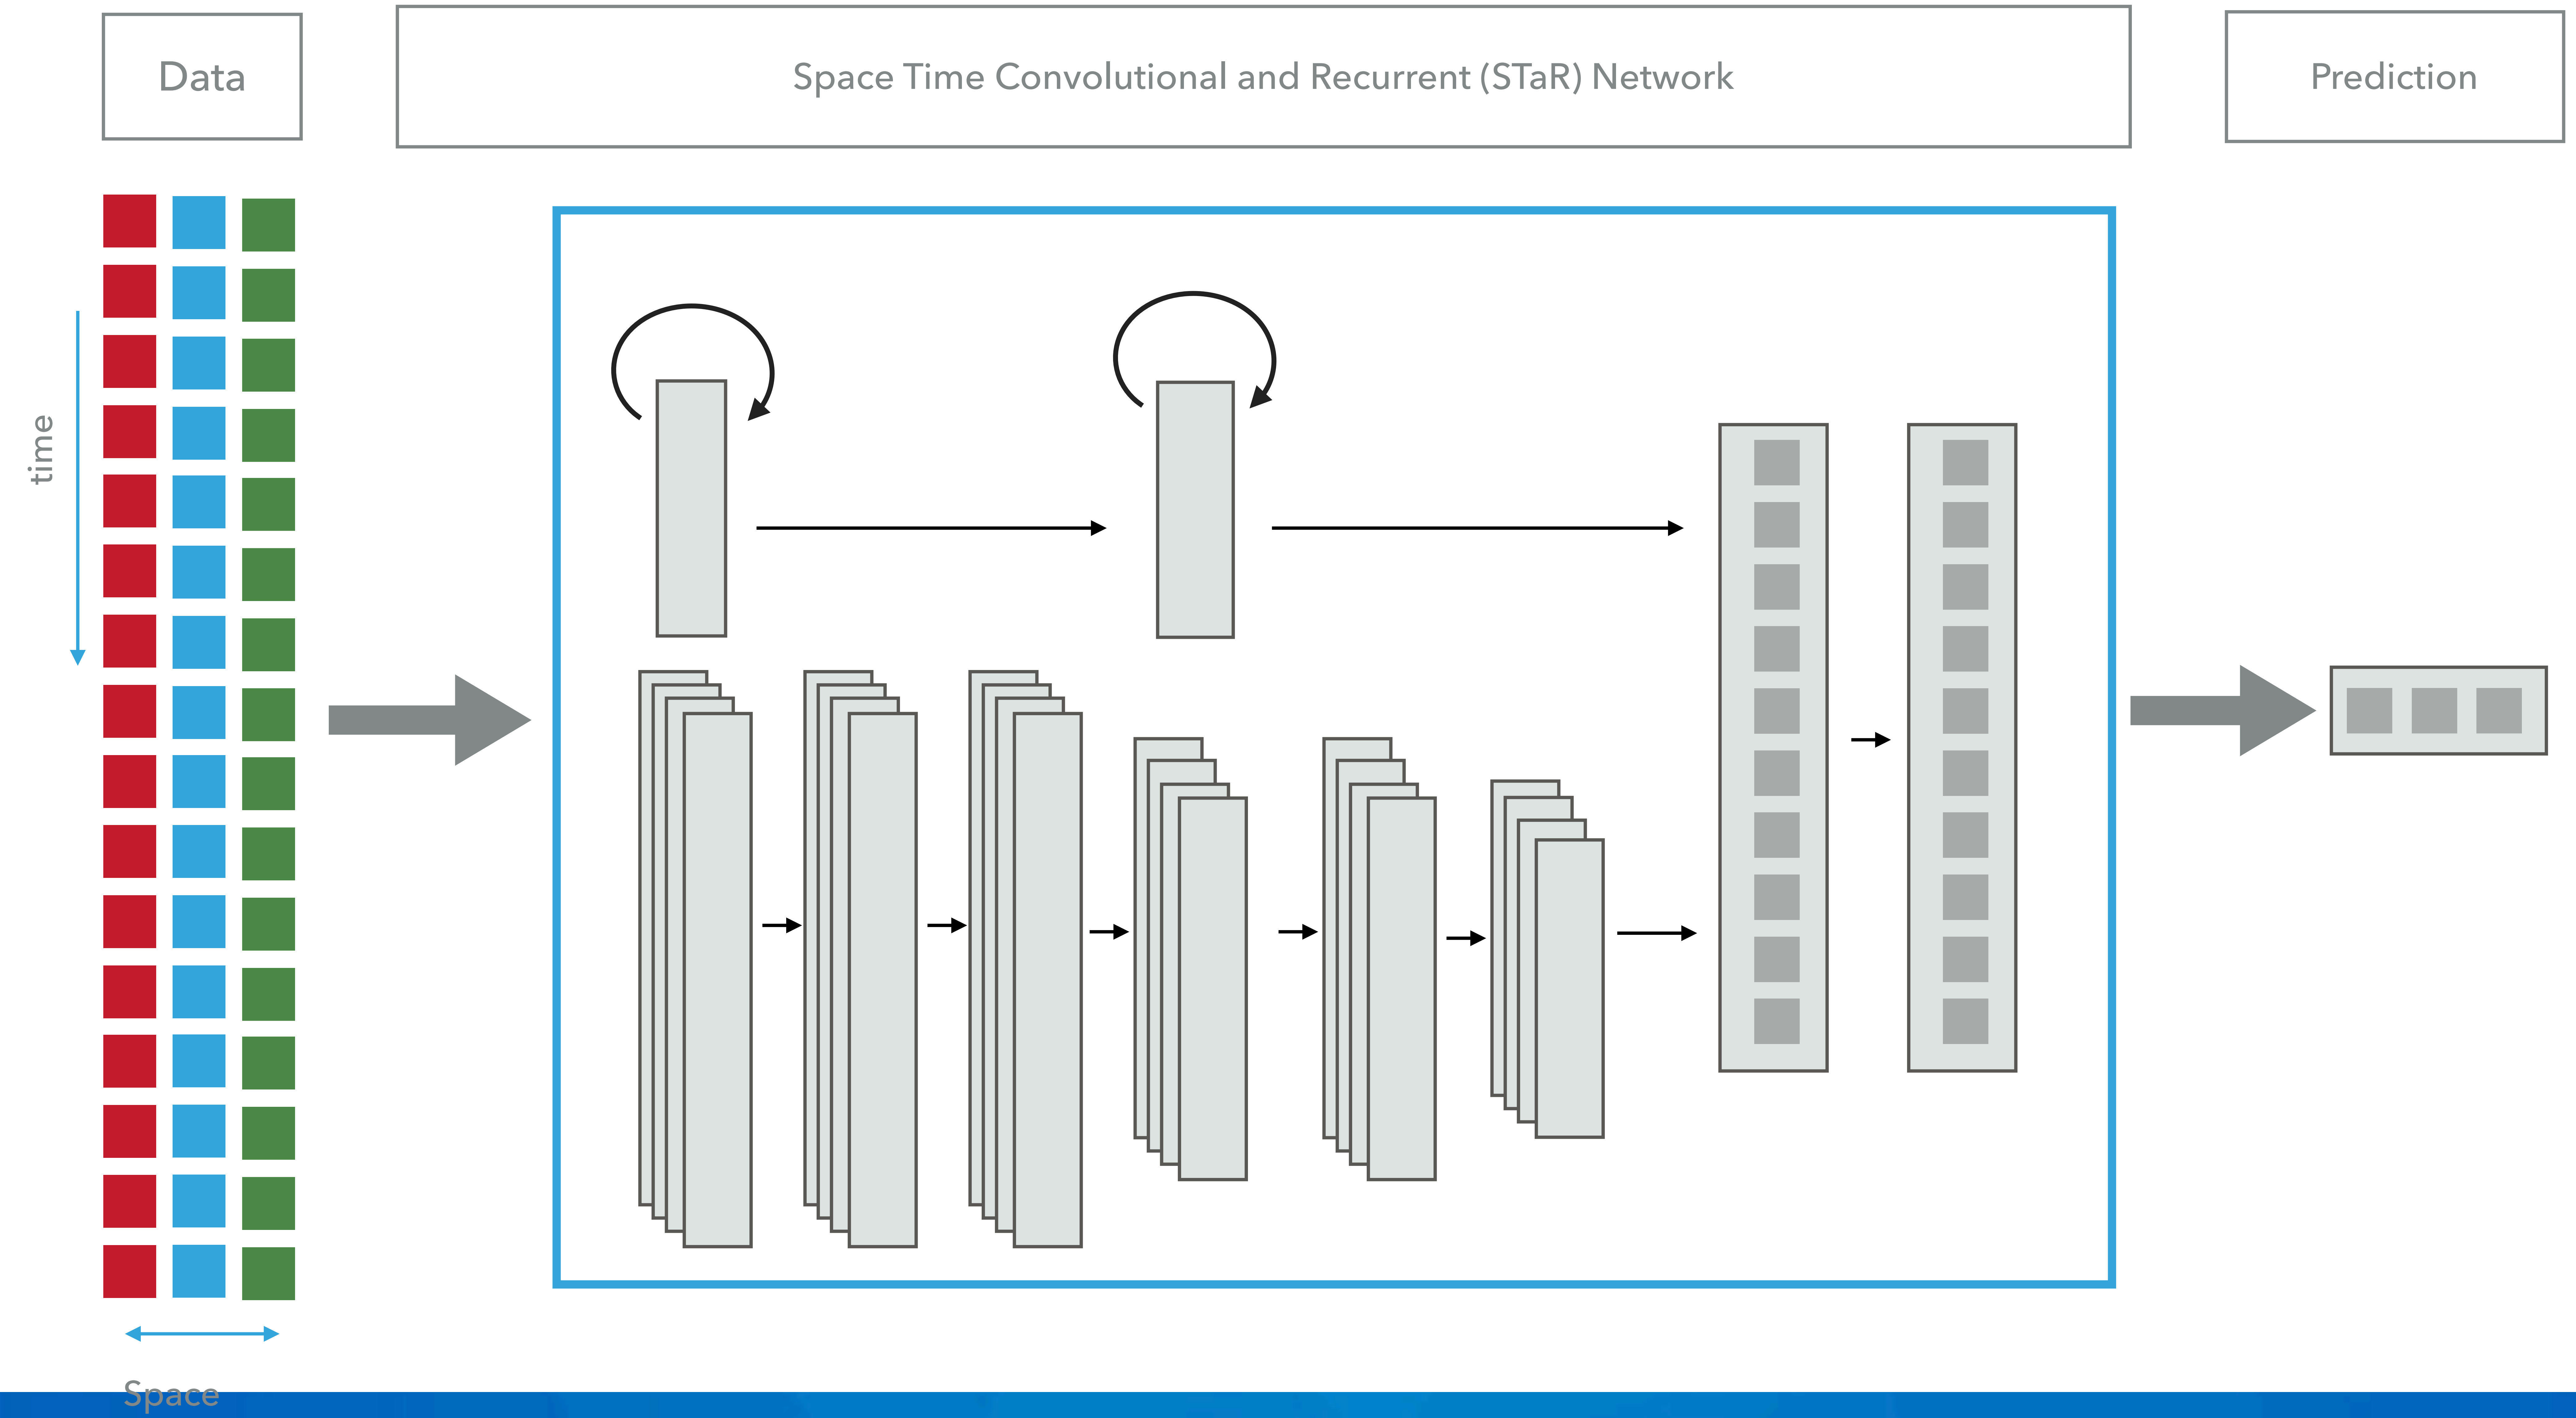

EPEX - European Power EXchange

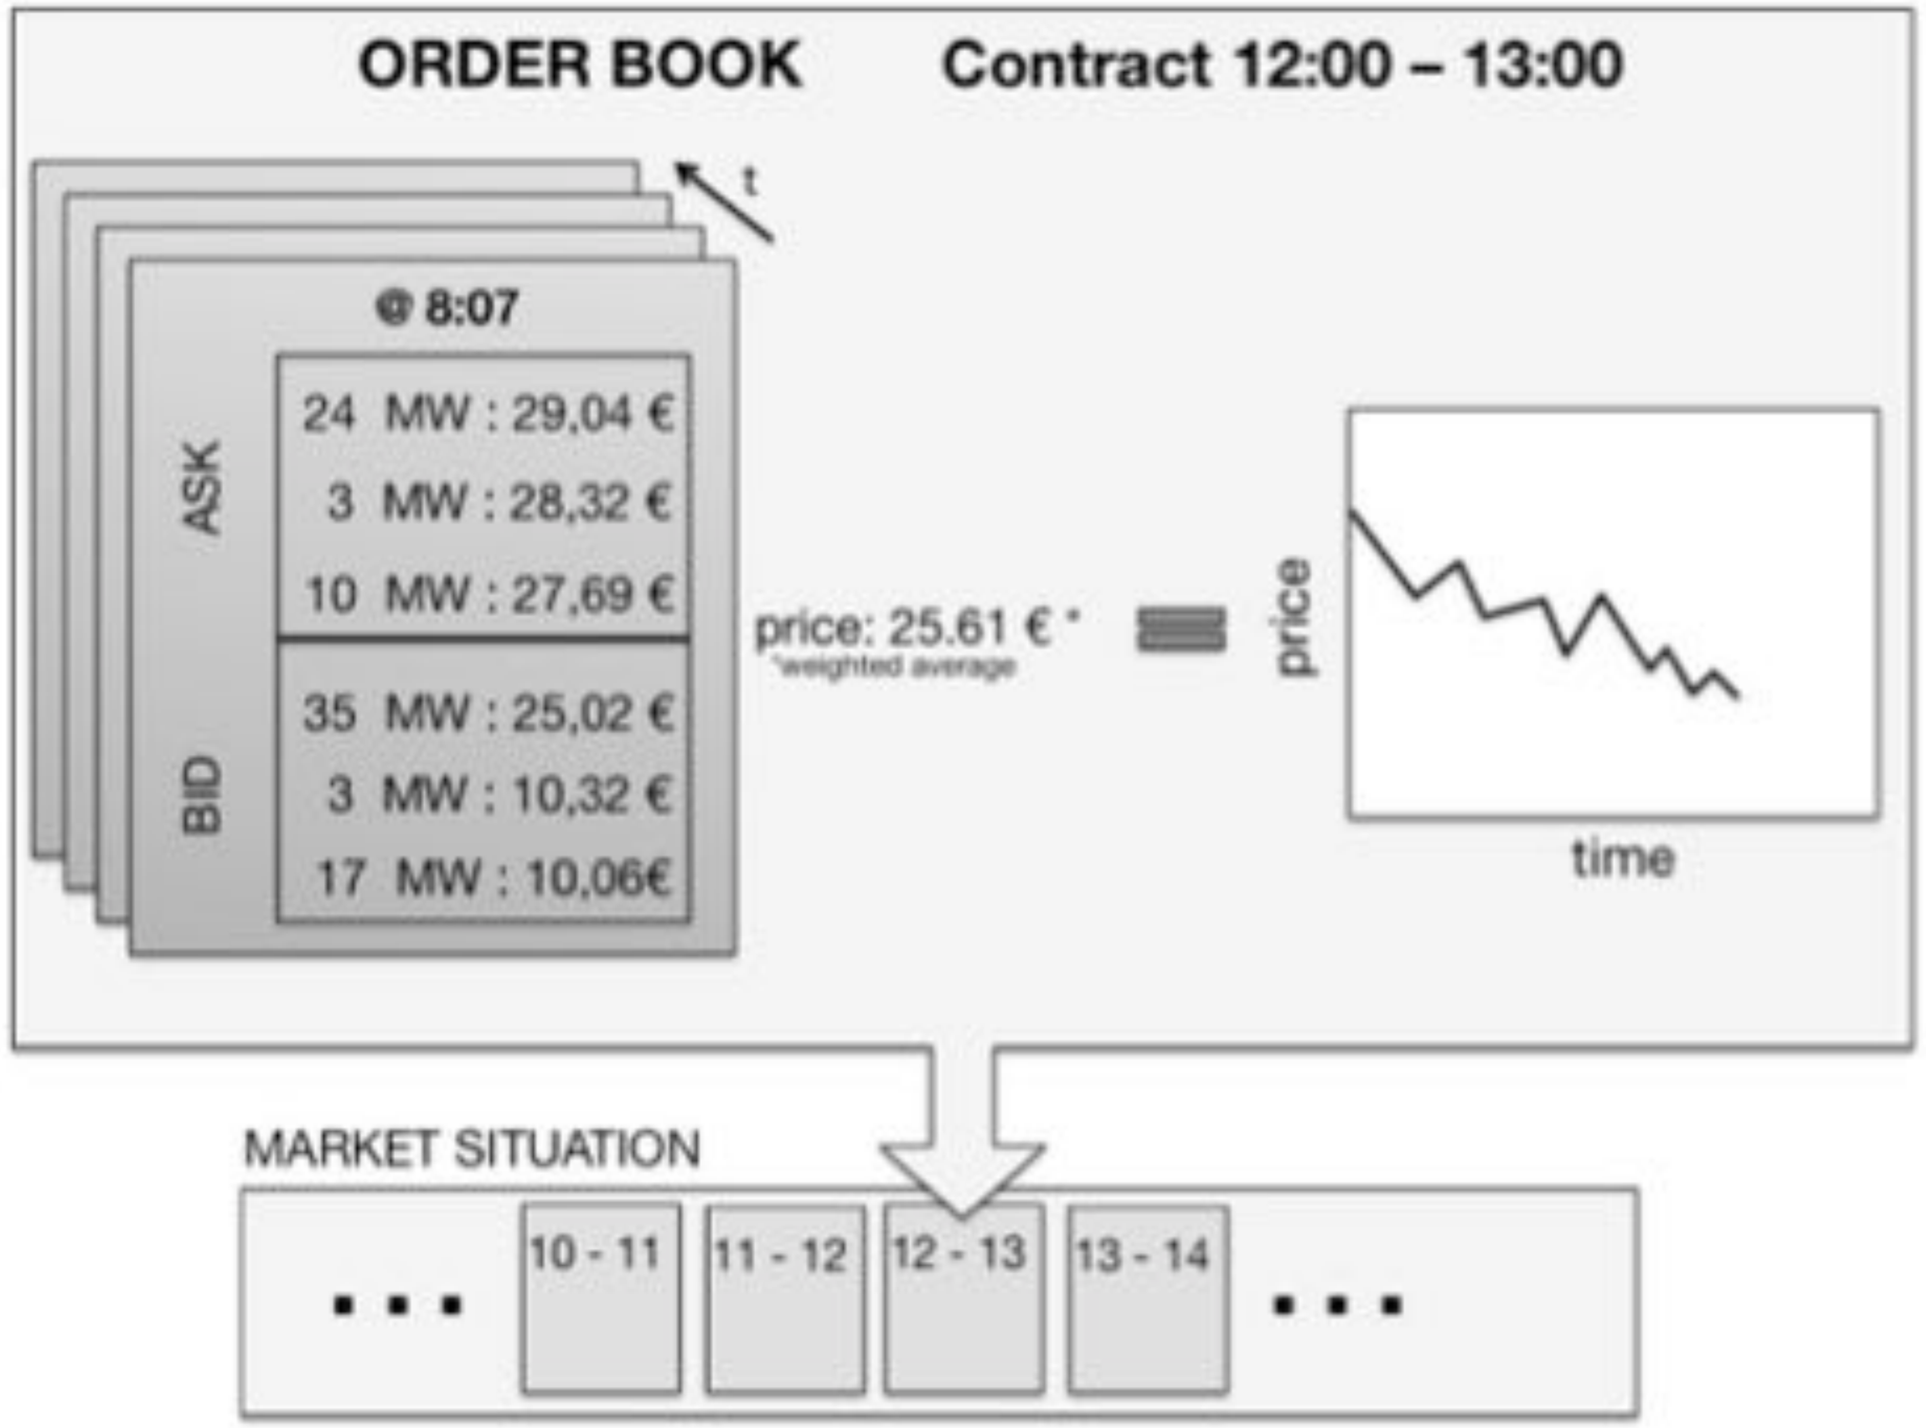

EPEXSPOT

2016 Backtest

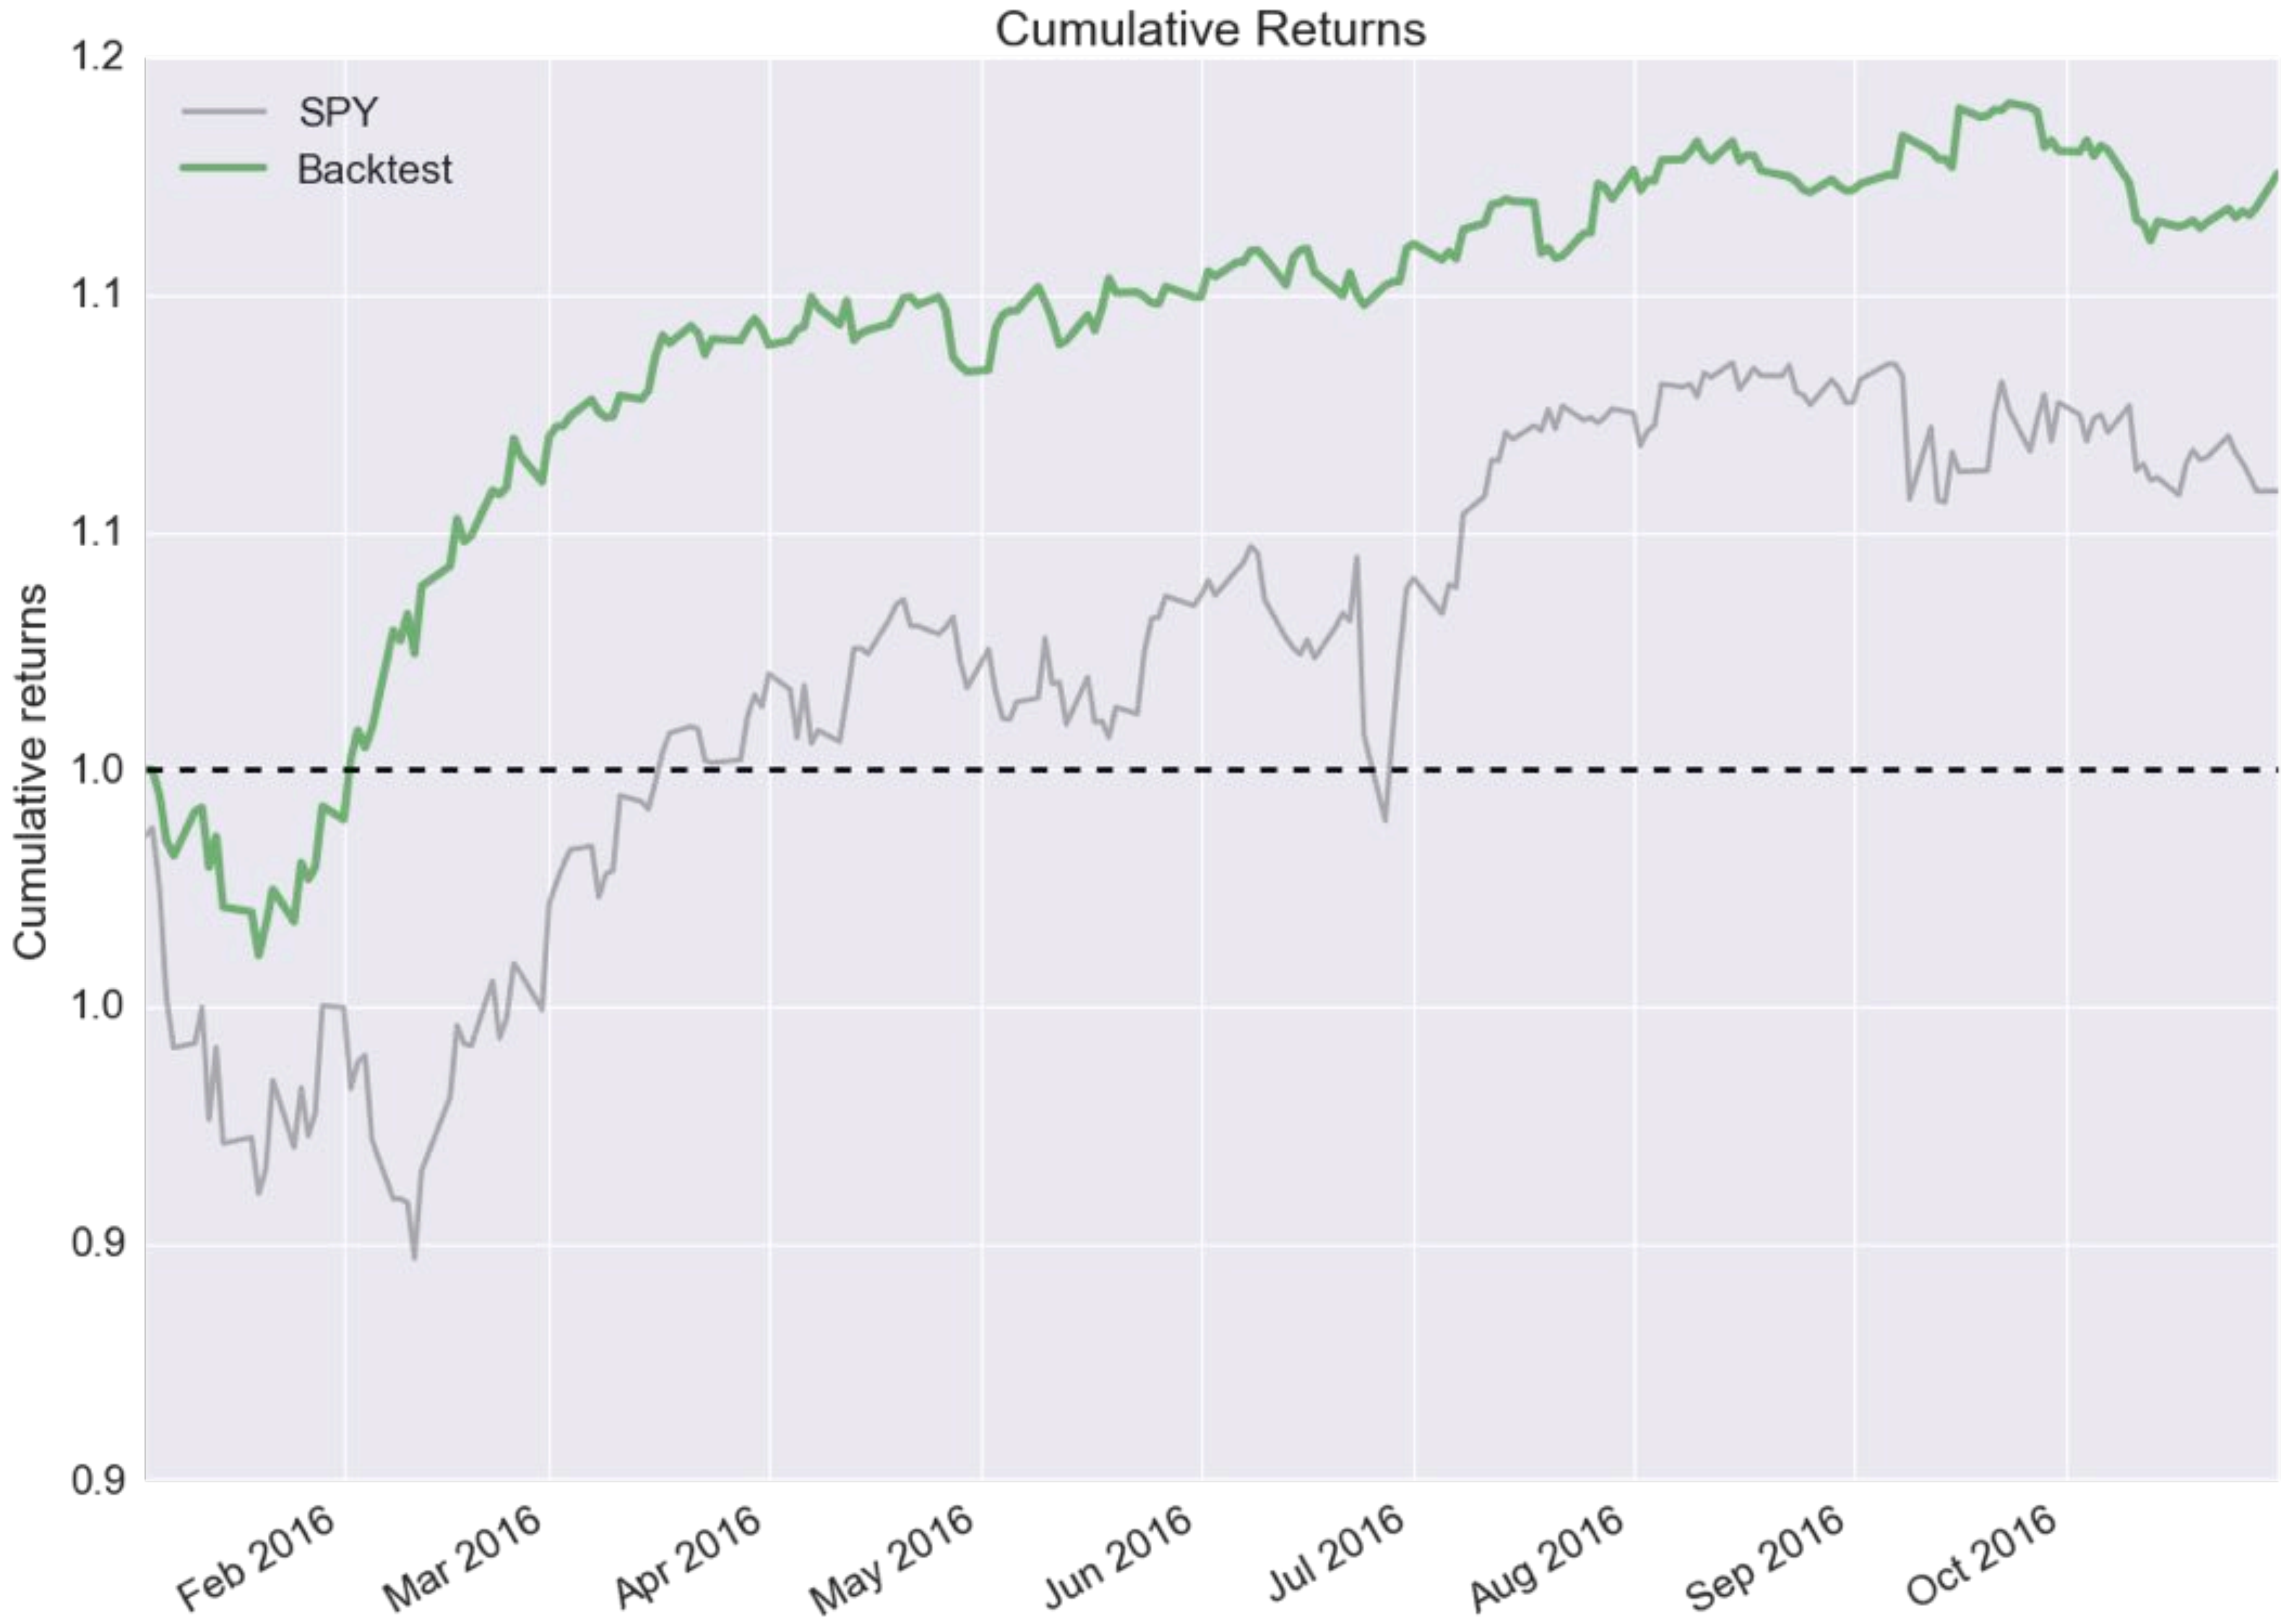

PSIORI “momentum”

Index

Return: 15%
Max Drawdown: 4%
Sharpe Ratio: 2.0

PREVEDEX

Seizure detection in EEG using Deep Neural Networks

- **first successful approach** to seizure detection on **raw intracranial EEG** data using deep convolutional neural networks
- the project aims towards an **implantable device**, detecting abnormal neuronal activity early on and **suppressing seizures** using deep brain stimulation
- the model is able to effectively extract features from extremely noisy multi-variate time series
- training utilizes gigabytes of EEG recordings from 27 different patients
- seizure detection improves upon previous methods in terms of sensitivity, specificity and detection delay by a large margin

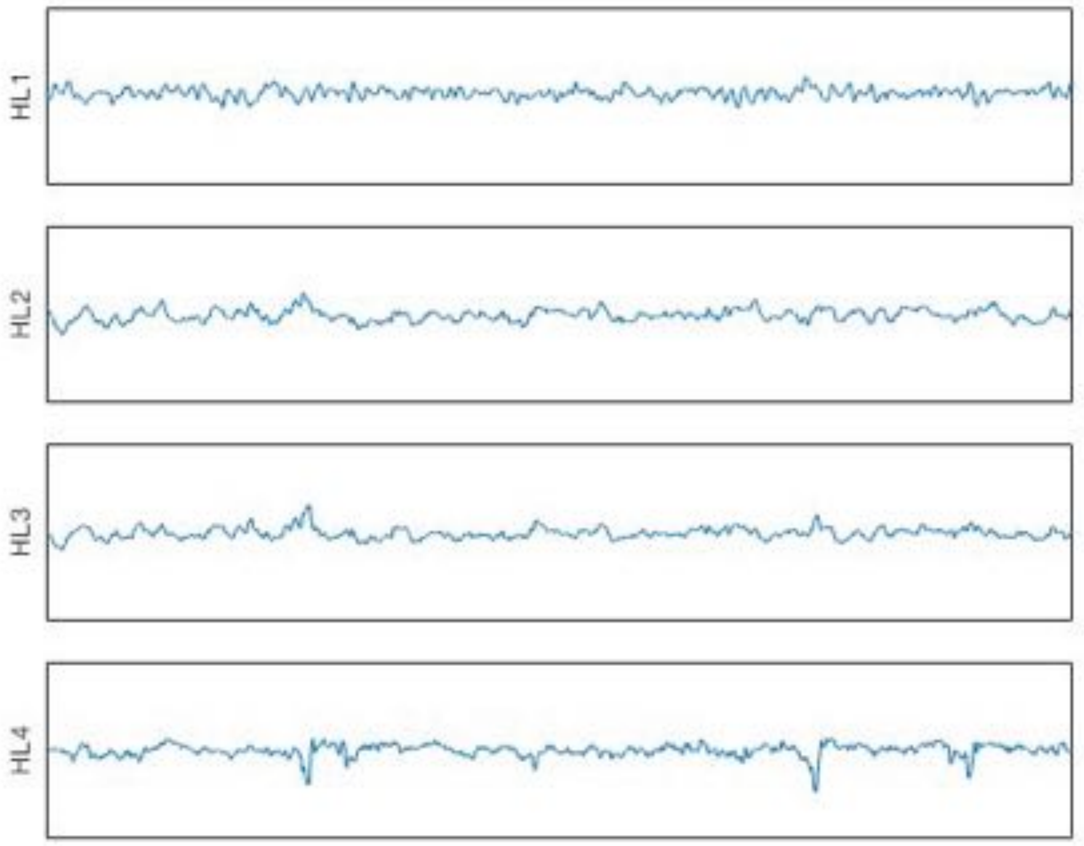

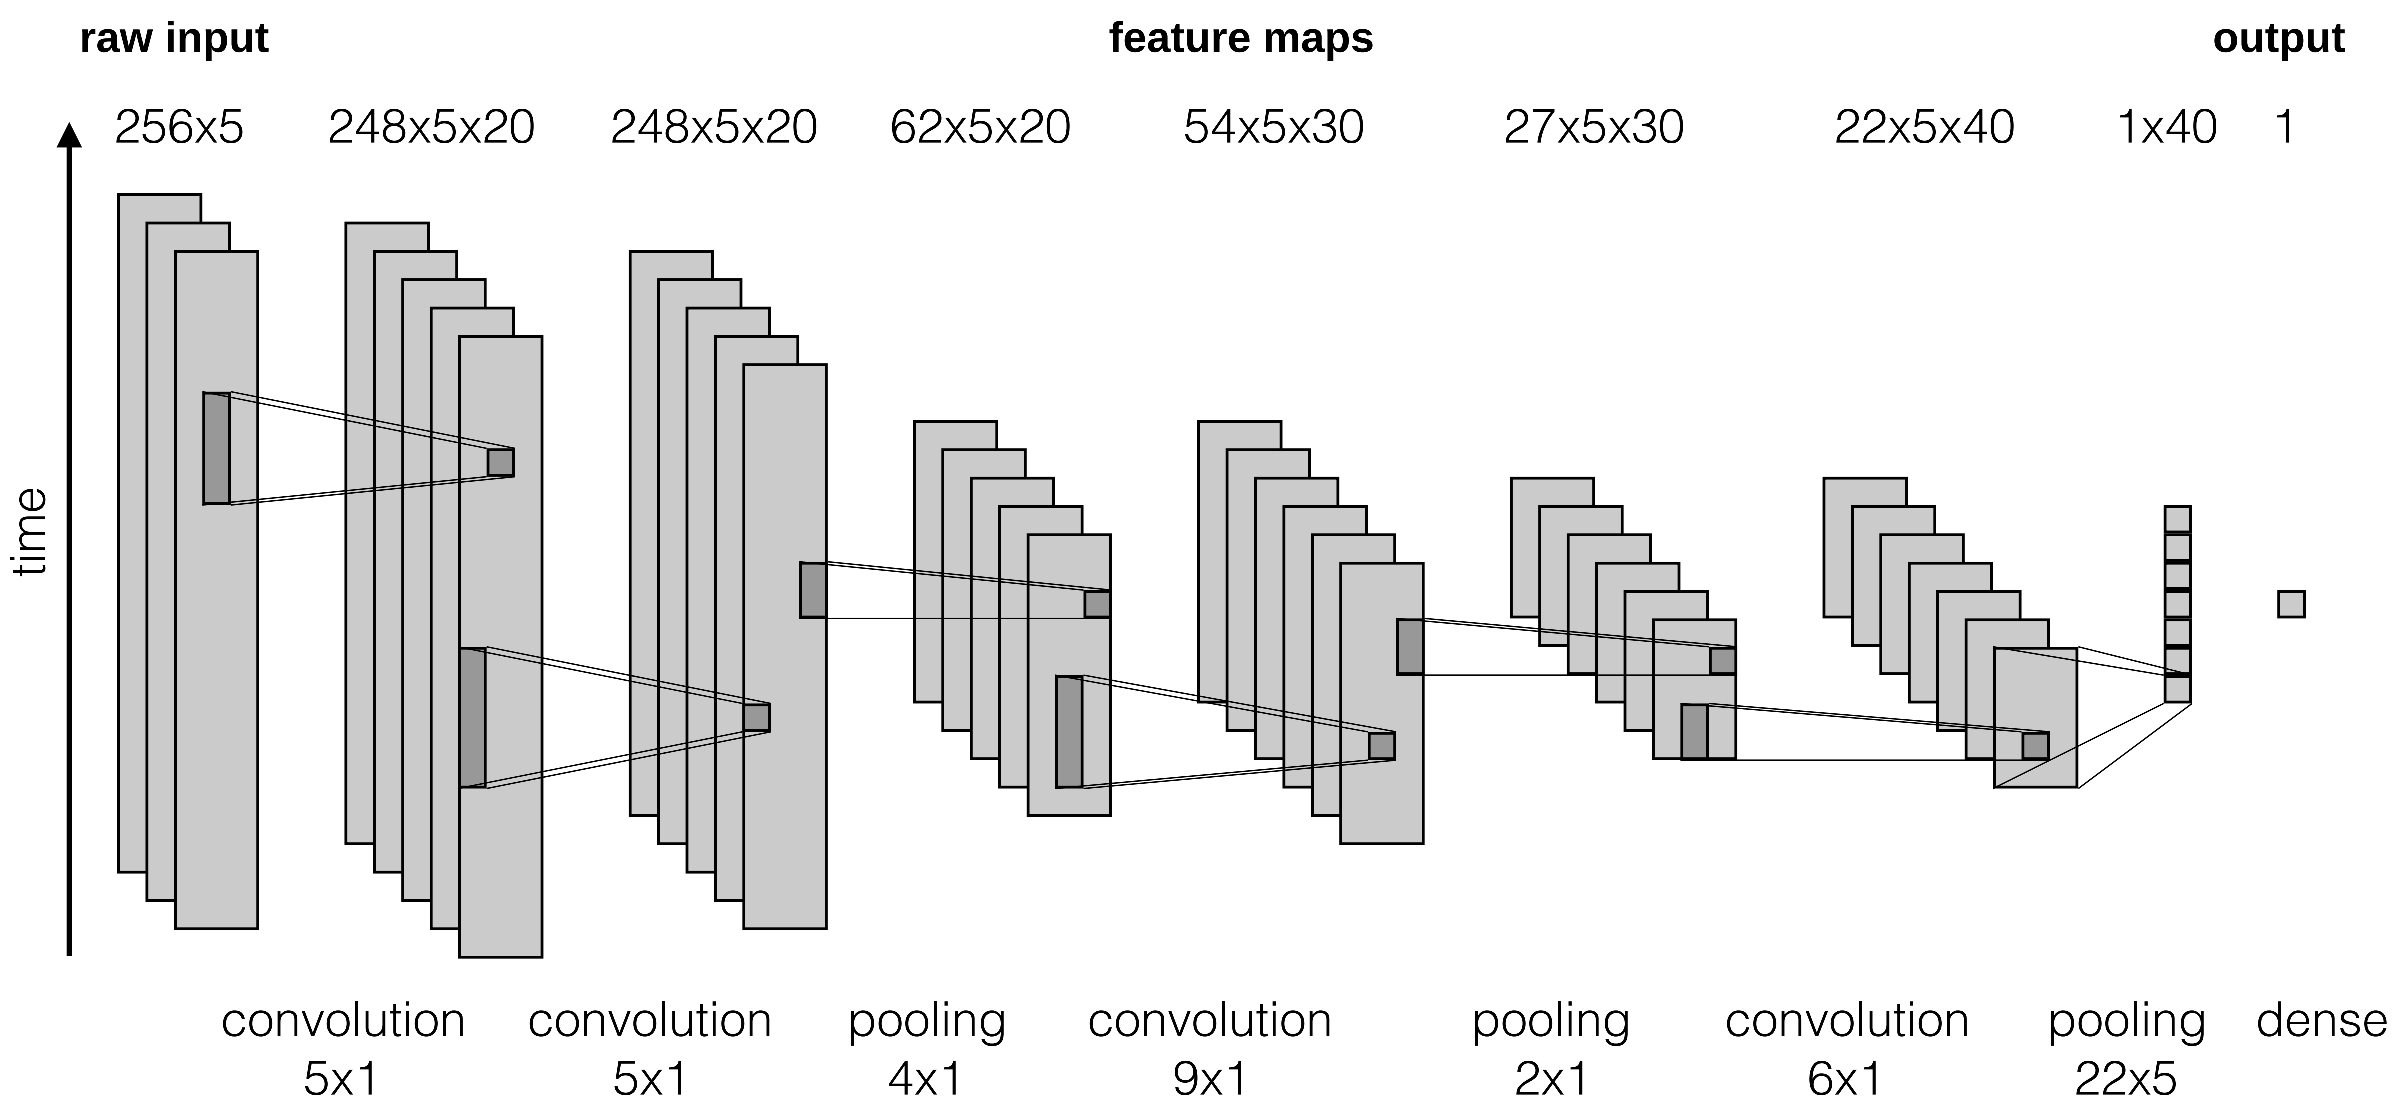

BrainLinks BrainTools
acting.thoughts

Reinforcement Learning

- Learning through interaction, by means of trial and error.
  Examples: Closed Loop Control, Gas Turbine, Engine .

Three Projects in Pipeline

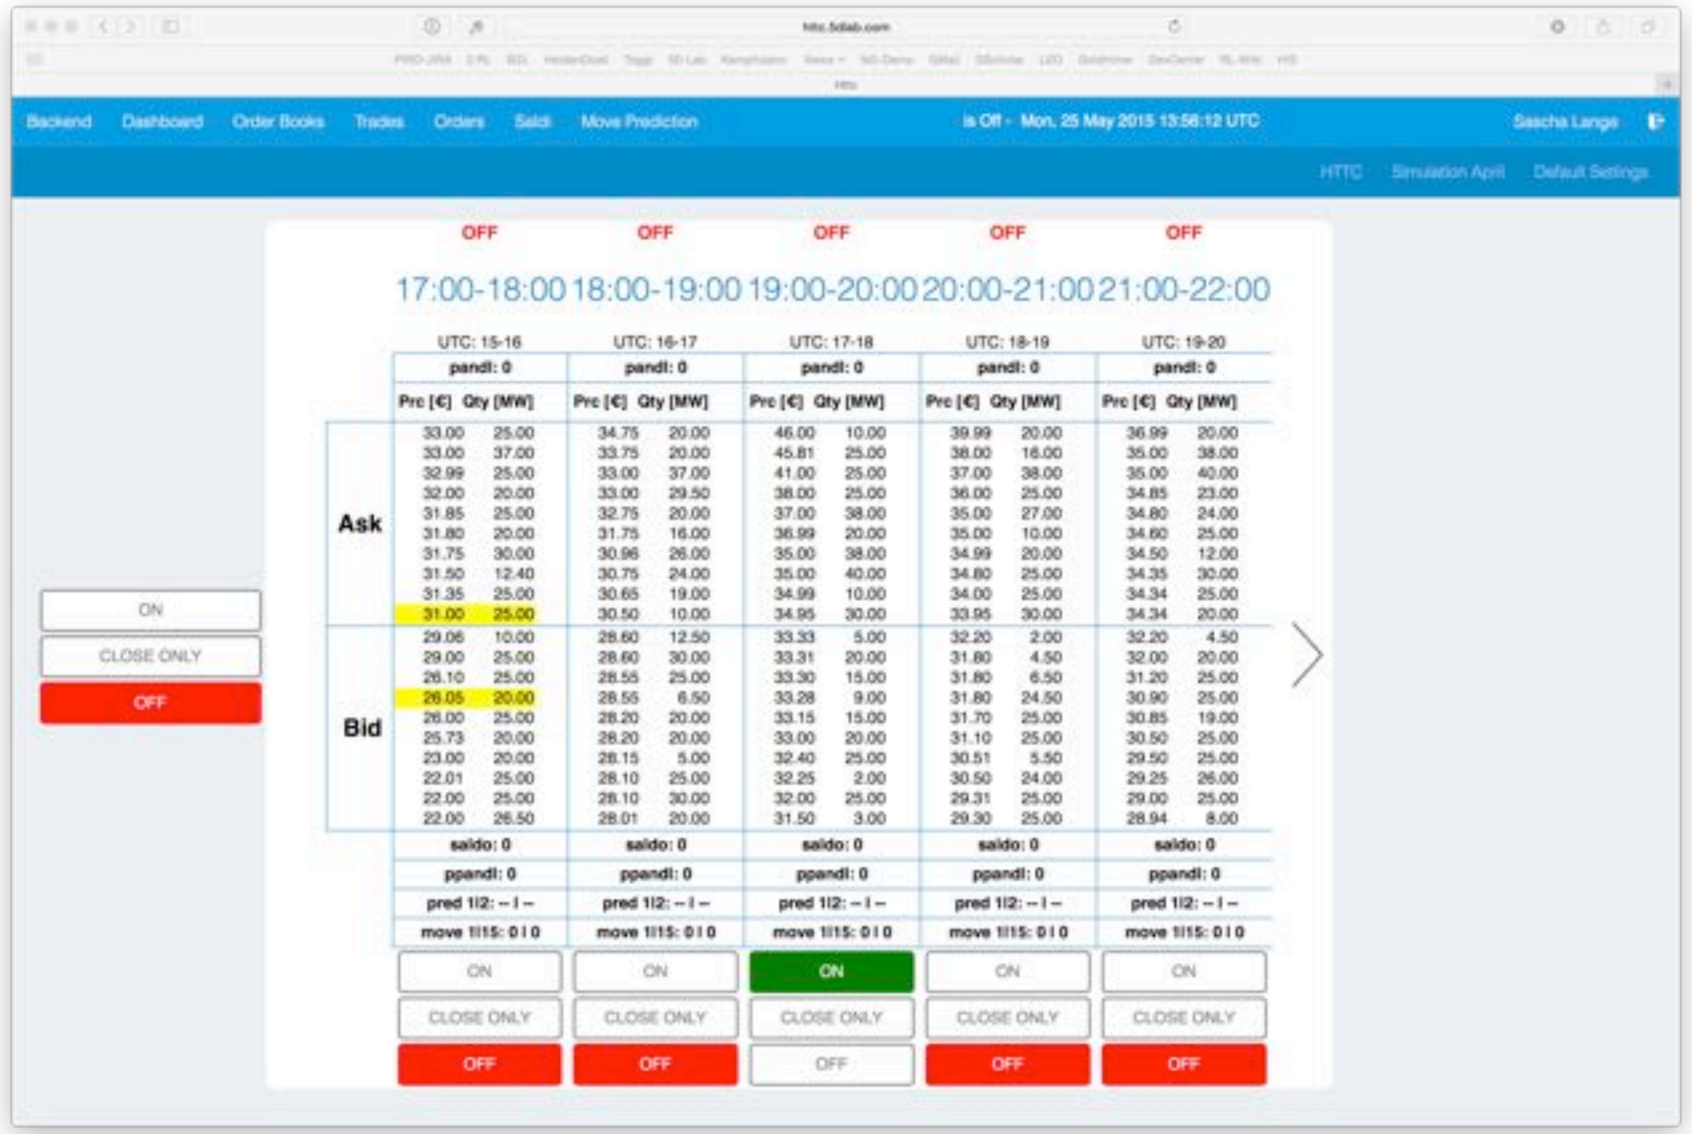

- Prevedex Deep RL Trading Bot — Planned Deployment: end of 2018
- Deep RL for Optimal Engine Control
- Batch Reinforcement Learning for Optimal Budget Allocation to Marketing Channels

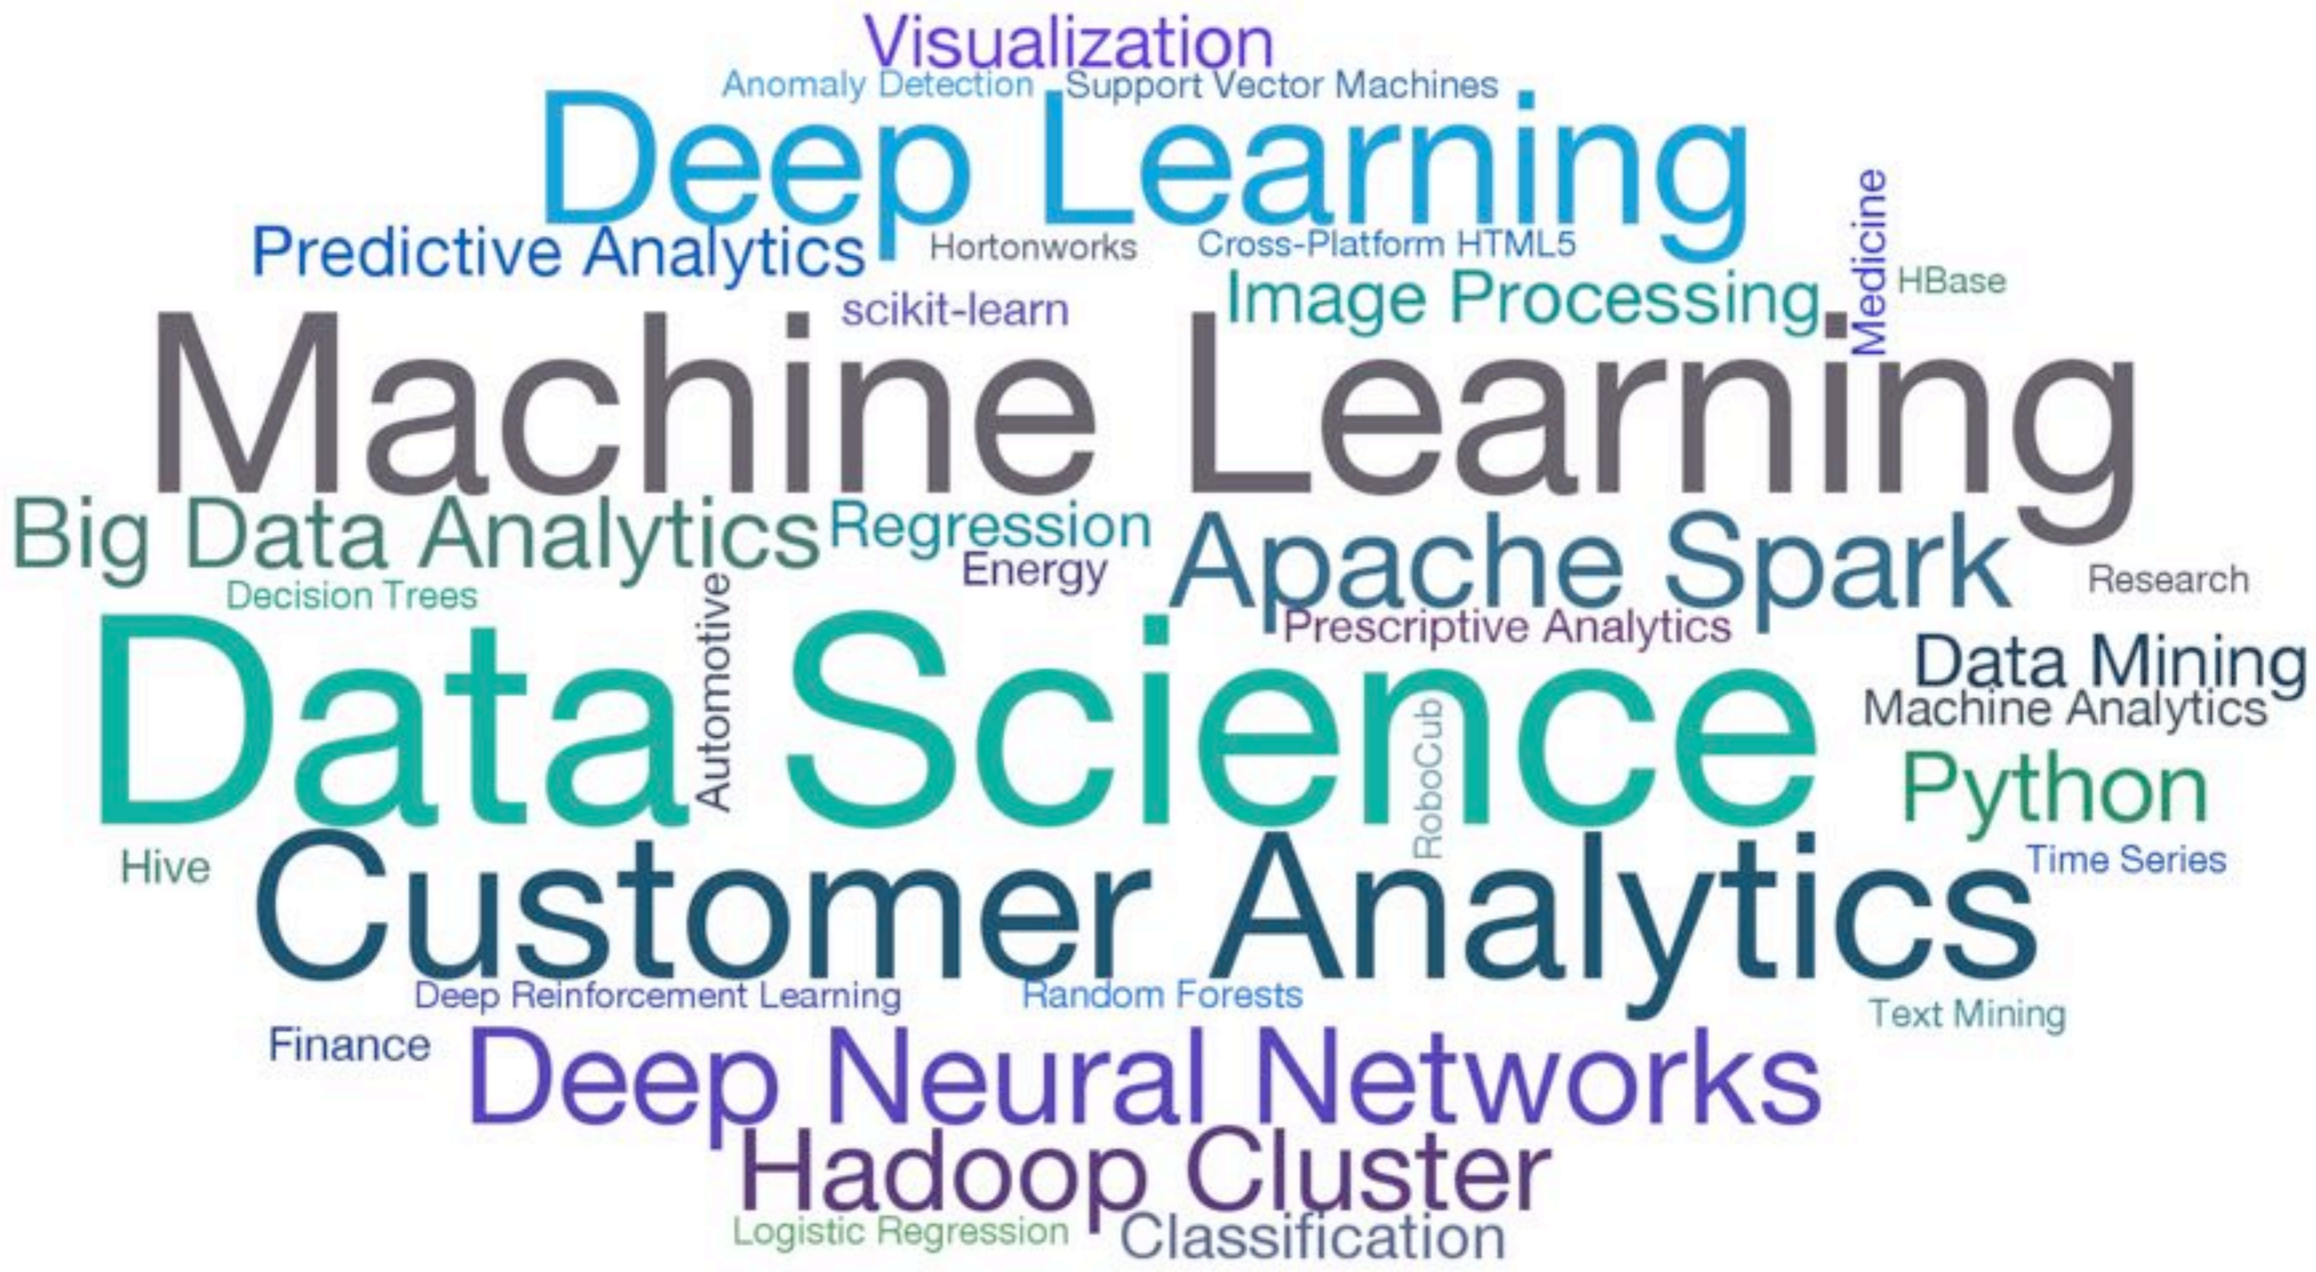

Recap

- Part I: The Research Perspective:

  Why is Artificial Intelligence hyped **right now**!
  What is Machine Learning!
  What is Deep Learning!
  What is Deep Reinforcement Learning!

- Part 2: The Business Perspective:

  Applications:
  Visualizing
  Predicting
  (Acting)