Statistical learning in the nervous system/ Intelligent systems Introduction

Gergő Orbán
golab.wigner.mta.hu

Anna Székely

Quiz

Quiz

- $\underline{a} \cdot \underline{b}$

Quiz

- $\underline{a} \cdot \underline{b}$
- $\text{Normal}(x, \mu, \Sigma)$

Quiz

- $\underline{a} \cdot \underline{b}$
- $\mathrm{Normal}(x, \mu, \Sigma)$
- Gauss function vs normal distribution

Quiz

- $\underline{a} \cdot \underline{b}$
- $\text{Normal}(x, \mu, \Sigma)$
- Gauss function vs normal distribution
- $\displaystyle\int f(x)\, dx$

Quiz

- $\underline{a} \cdot \underline{b}$
- $\mathrm{Normal}(x, \mu, \Sigma)$
- Gauss function vs normal distribution
- $\int f(x)\, dx$
- $p(a \,|\, b)$

How does the brain work?

Modern neuroscience

Modern neuroscience

a wired.com szerint

CHELSEA LEU SCIENCE 12.12.16 3:00 PM

WATCH A RESTING BRAIN LIGHT UP WITH ACTIVITY

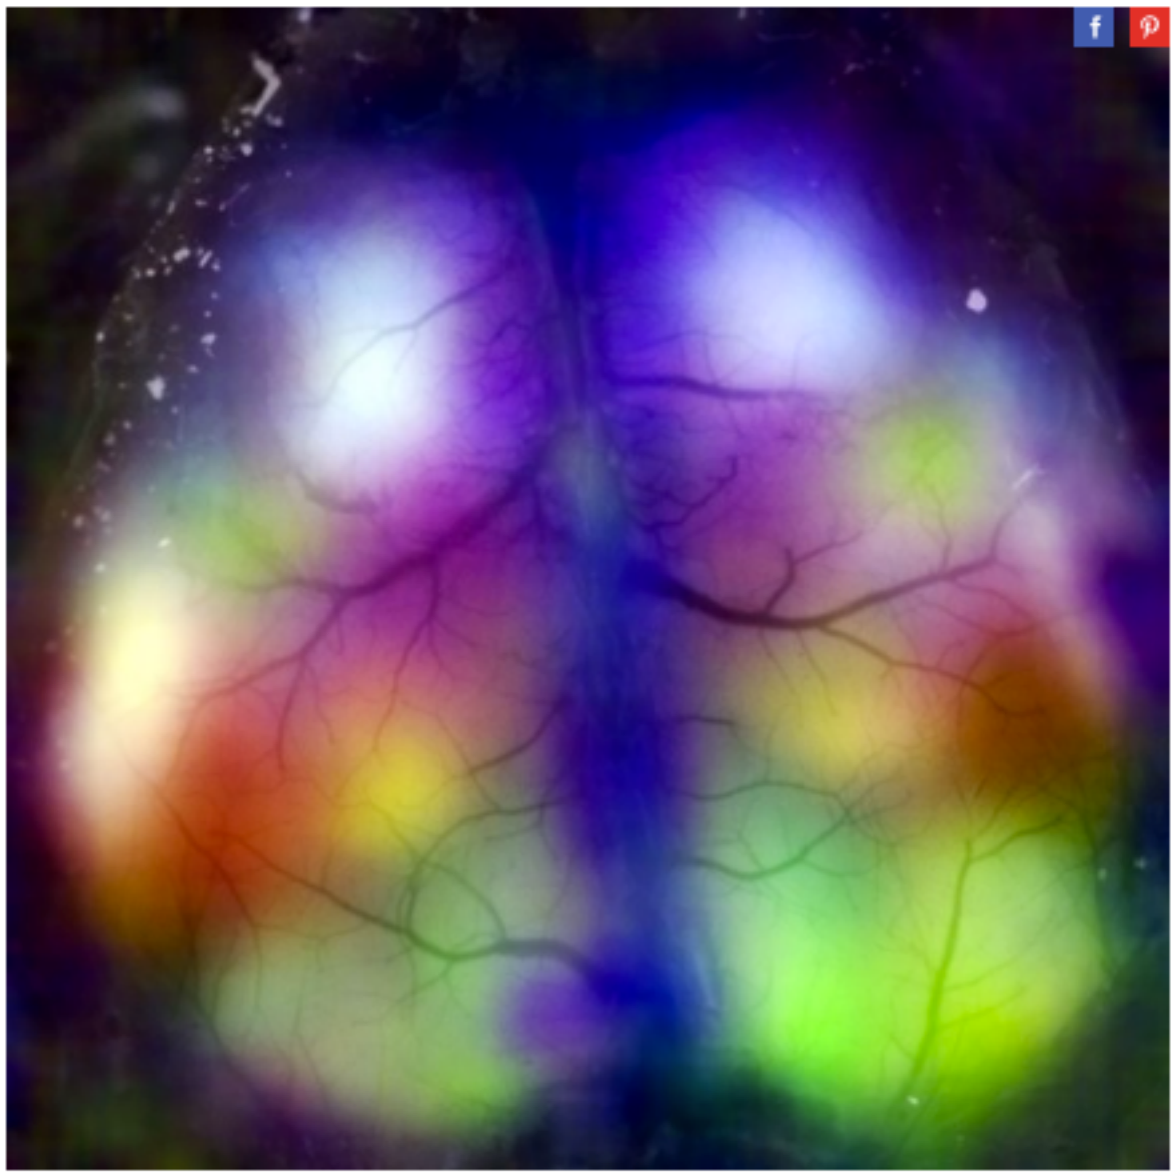

YING MA AND ELIZABETH HILLMAN/COLUMBIA'S ZUCKERMAN INSTITUTE

Modern neuroscience

a wired.com szerint

CHELSEA LEU SCIENCE 12.12.16 3:00 PM

WATCH A RESTING BRAIN LIGHT UP WITH ACTIVITY

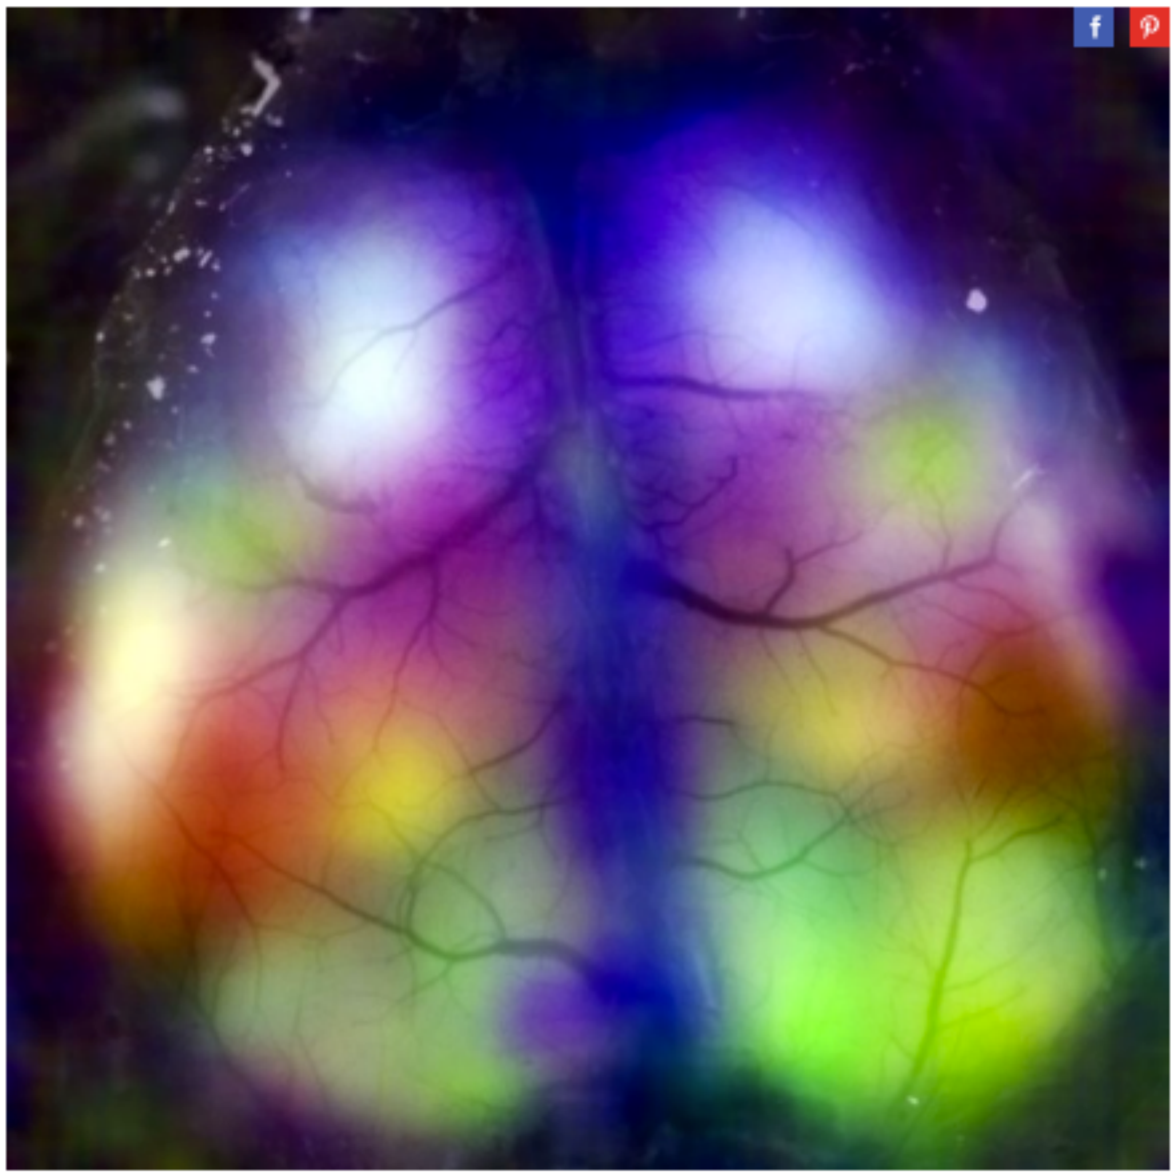

YING MA AND ELIZABETH HILLMAN/COLUMBIA'S ZUCKERMAN INSTITUTE

CHELSEA LEU SCIENCE 07.20.16 1:00 PM

A NEW MAP OF THE BRAIN REDRAWS THE BOUNDARIES OF NEUROSCIENCE

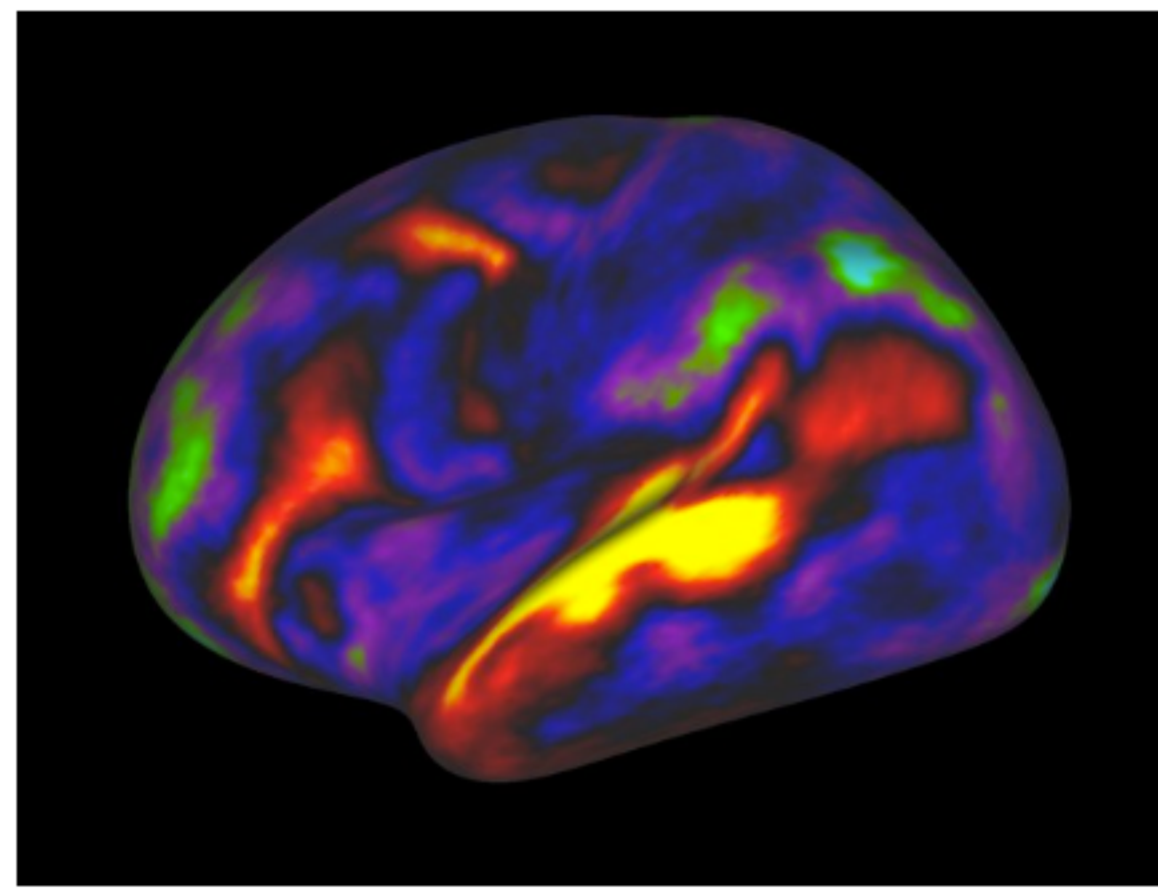

The image shows the pattern of brain activation (red, yellow) and deactivation (blue, green) in the left hemisphere when listening to stories in the MRI scanner. MATTHEW F. GLASSER/DAVID C. VAN ESSEN

Modern neuroscience

a wired.com szerint

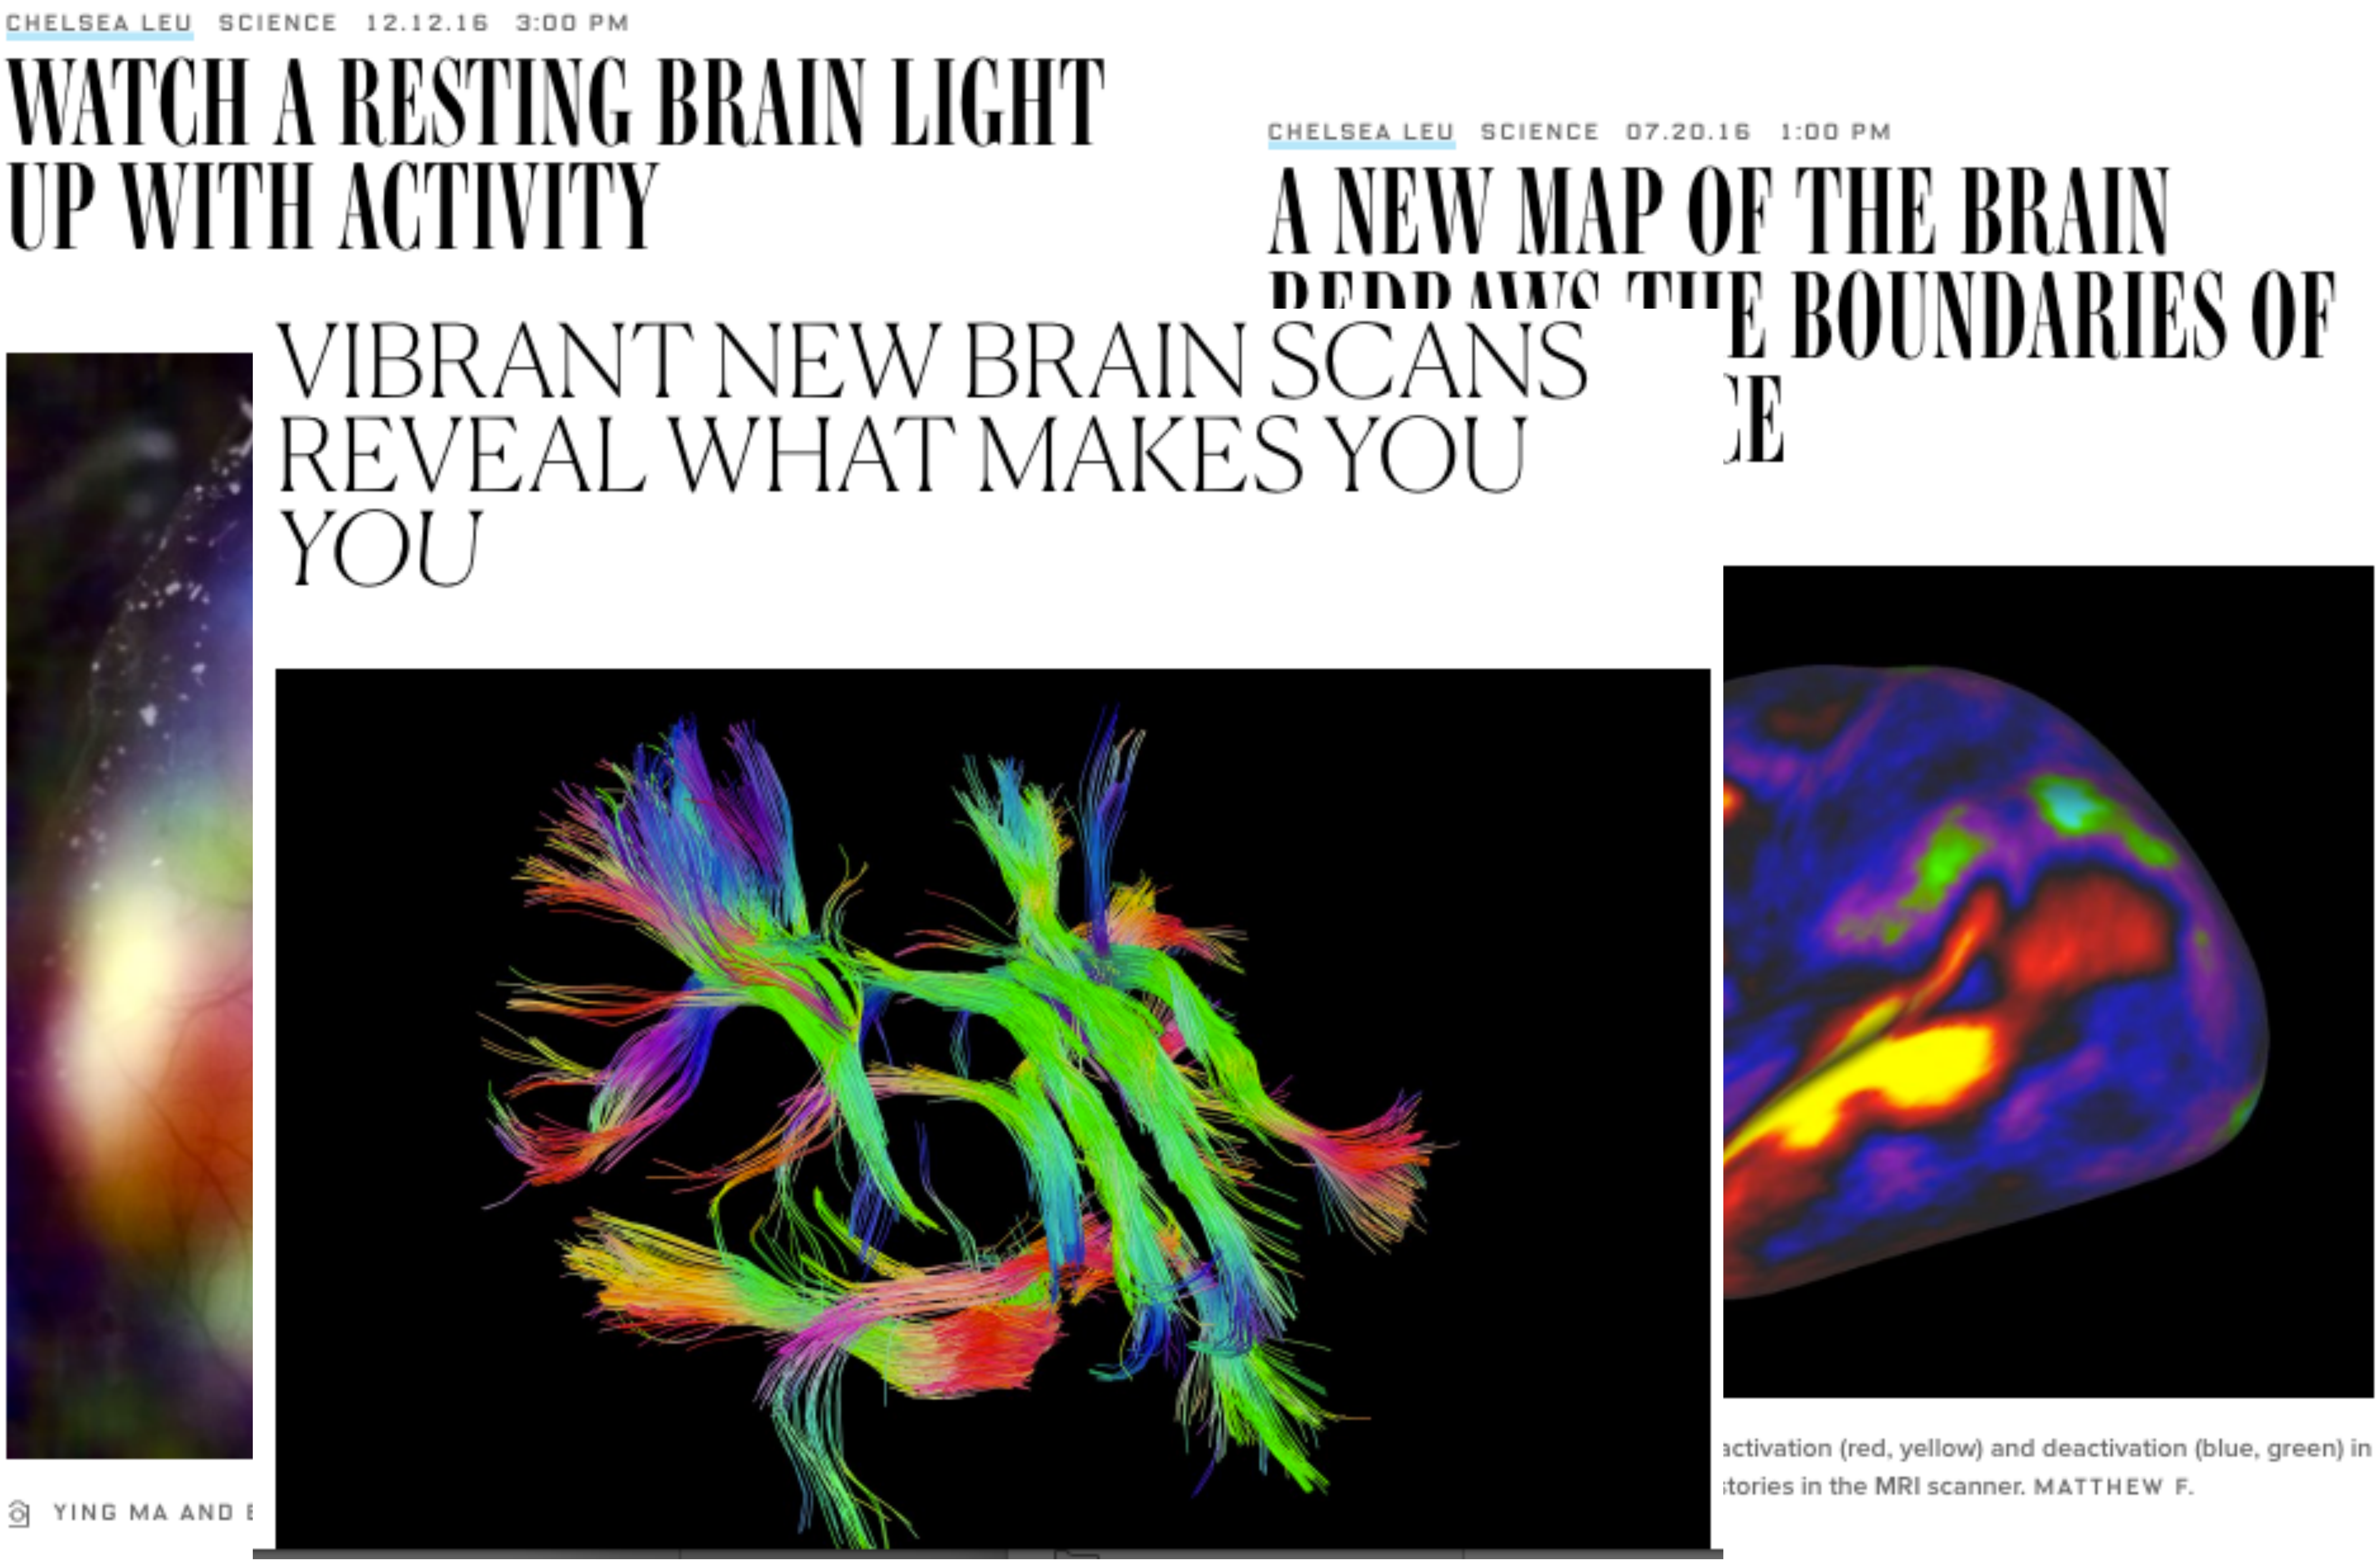

Modern neuroscience

a wired.com szerint

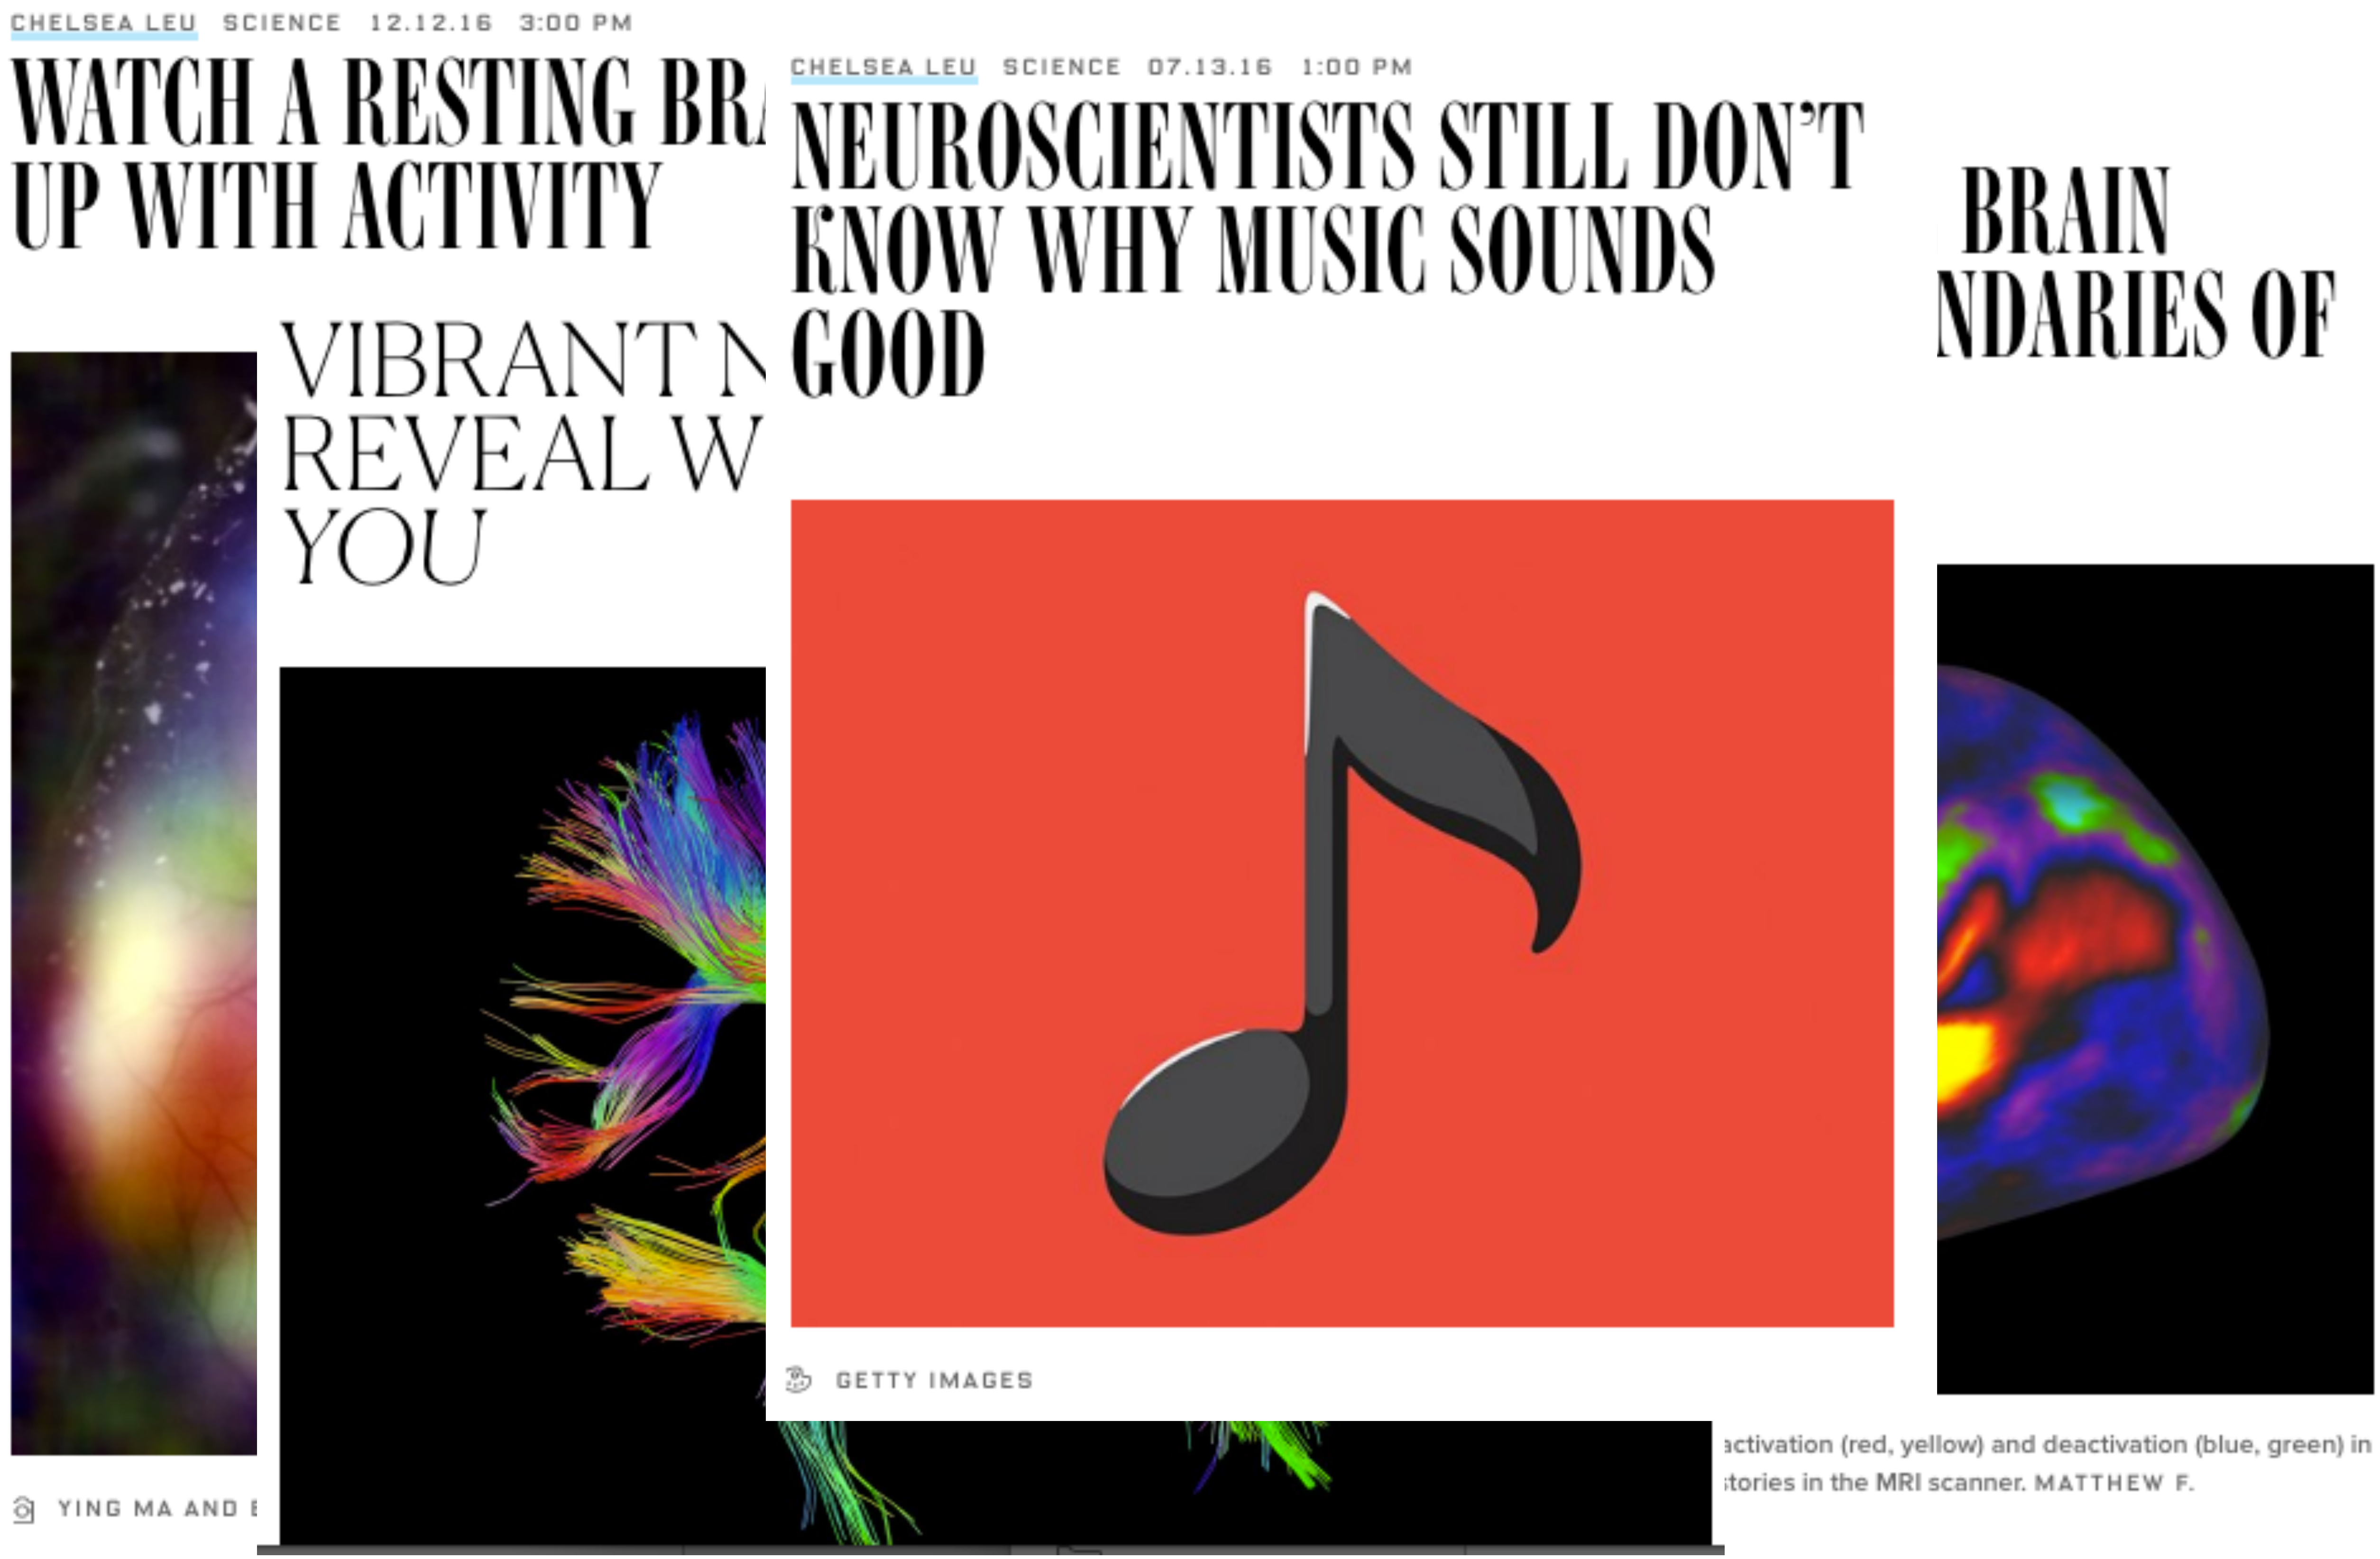

How does the brain work?

- What does this question actually mean? What are the answers that we would accept?
- Do we have an answer to this question?

How does an intelligent system look like?

Artificial intelligence according to wired

Artificial intelligence according to wired

a wired.com szerint

CADE METZ BUSINESS 01.30.17 7:00 AM

GOOGLE'S GO-PLAYING MACHINE OPENS THE DOOR TO ROBOTS THAT LEARN

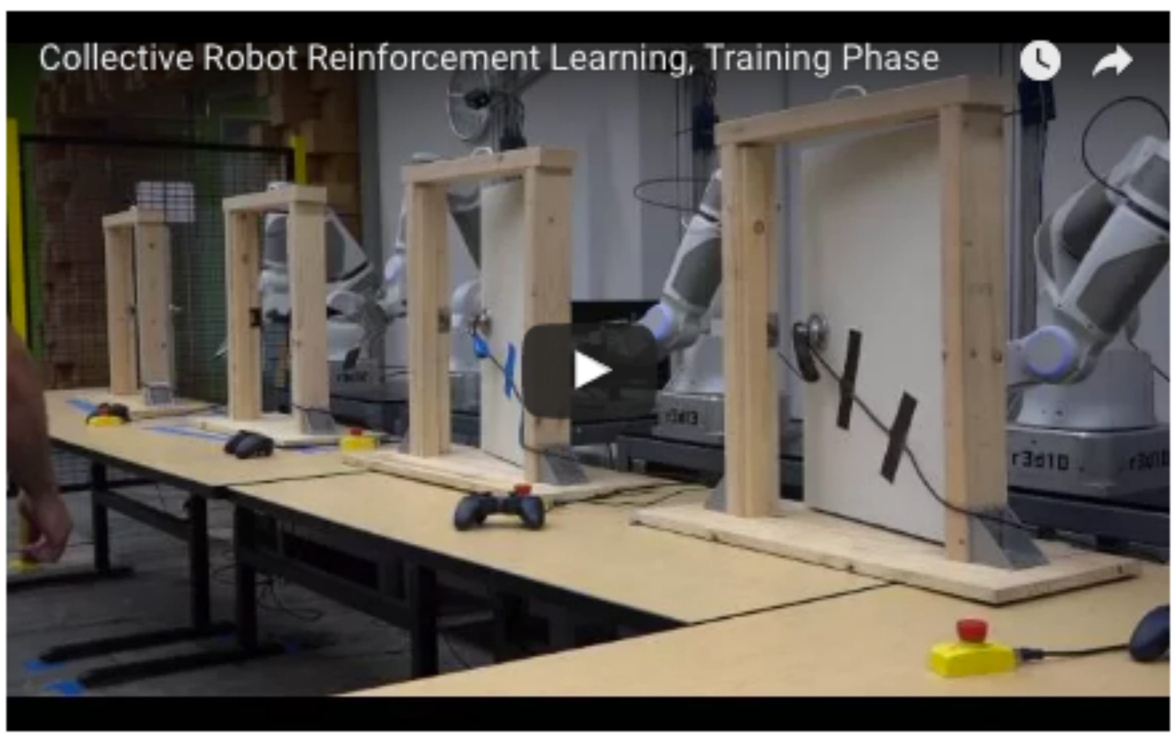

Artificial intelligence according to wired

a wired.com szerint

CADE METZ BUSINESS 01.30.17 7:00 AM

GOOGLE'S GO-PLAYING MACHINE OPENS THE DOOR TO ROBOTS THAT LEARN

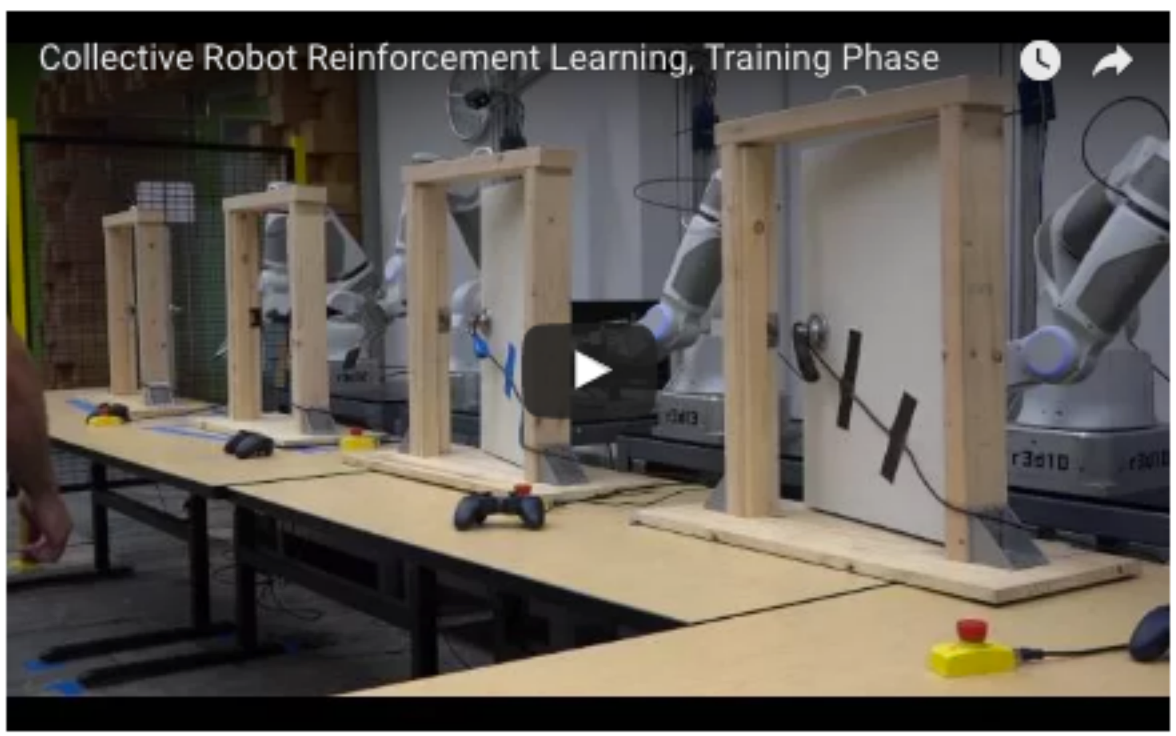

CADE METZ BUSINESS 01.31.17 7:00 AM

A MYSTERY AI JUST CRUSHED THE BEST HUMAN PLAYERS AT POKER

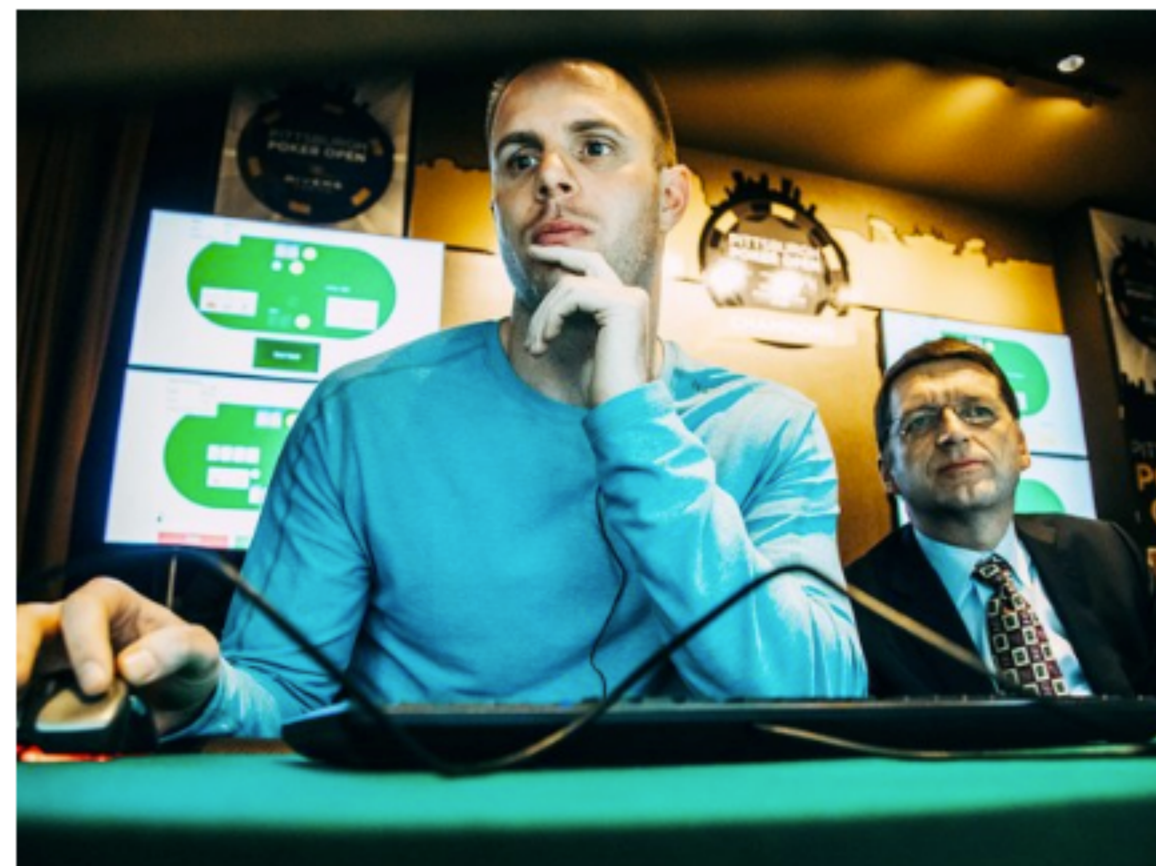

Professional poker player Jason Les plays against "Libratus," at Rivers Casino in Pittsburgh, on January 11, 2017. ANDREW RUSH/PITTSBURGH POST-GAZETTE/AP

Artificial intelligence according to wired

a wired.com szerint

CADE METZ B

GOOG
OPENS
THAT

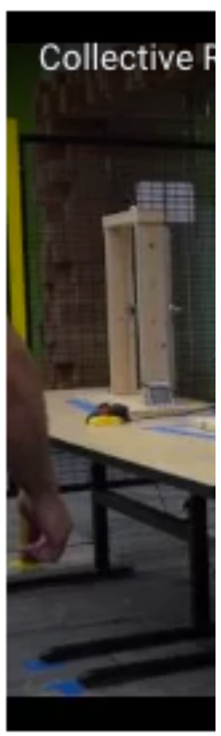

MEGAN MOLTENI SCIENCE 01.25.17 1:00 PM

IF YOU LOOK AT X-RAYS OR MOLES FOR A LIVING, AI IS COMING FOR YOUR JOB

MATT YOUNG

CADE METZ BUSINESS 01.31.17 7:00 AM

A MYSTERY AI JUST CRUSHED THE BEST HUMAN PLAYERS AT POKER

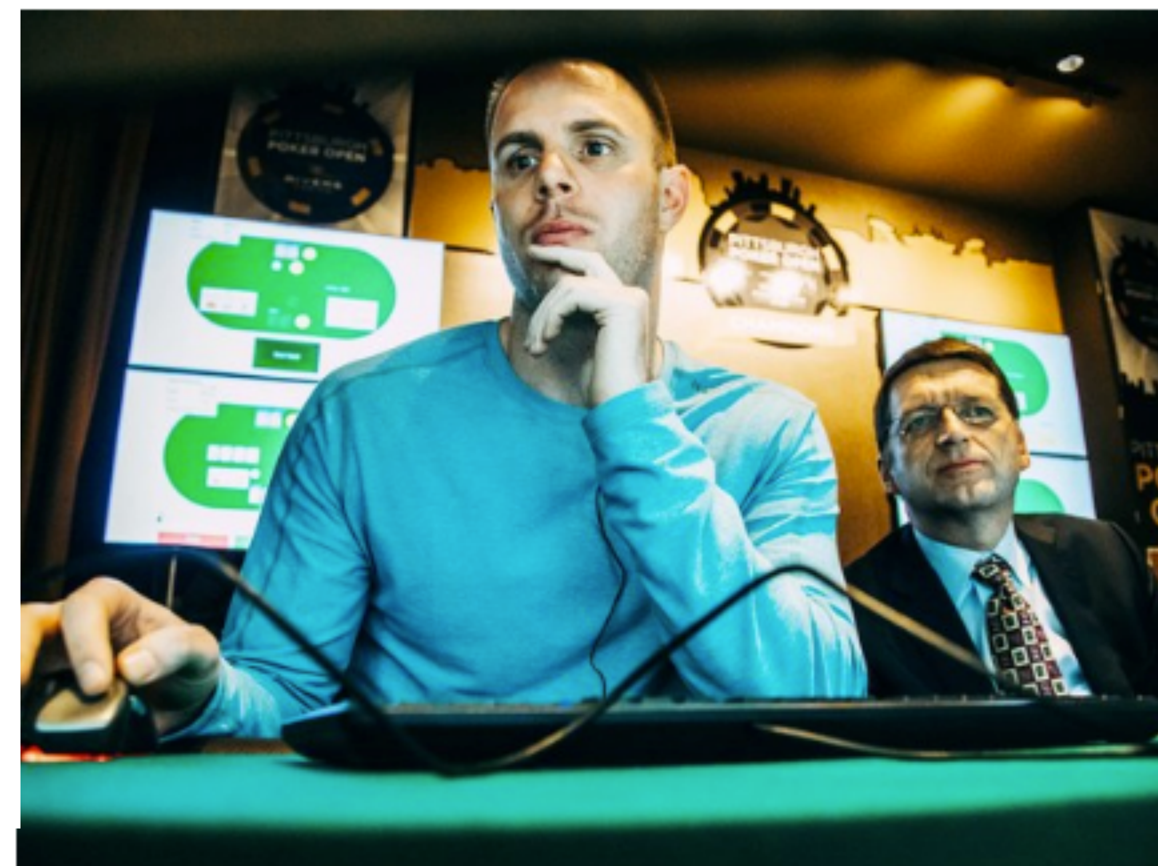

Professional poker player Jason Les plays against "Libratus," at Rivers Casino in Pittsburgh, on January 11, 2017. ANDREW RUSH/PITTSBURGH POST-GAZETTE/AP

Artificial intelligence according to wired

a wired.com szerint

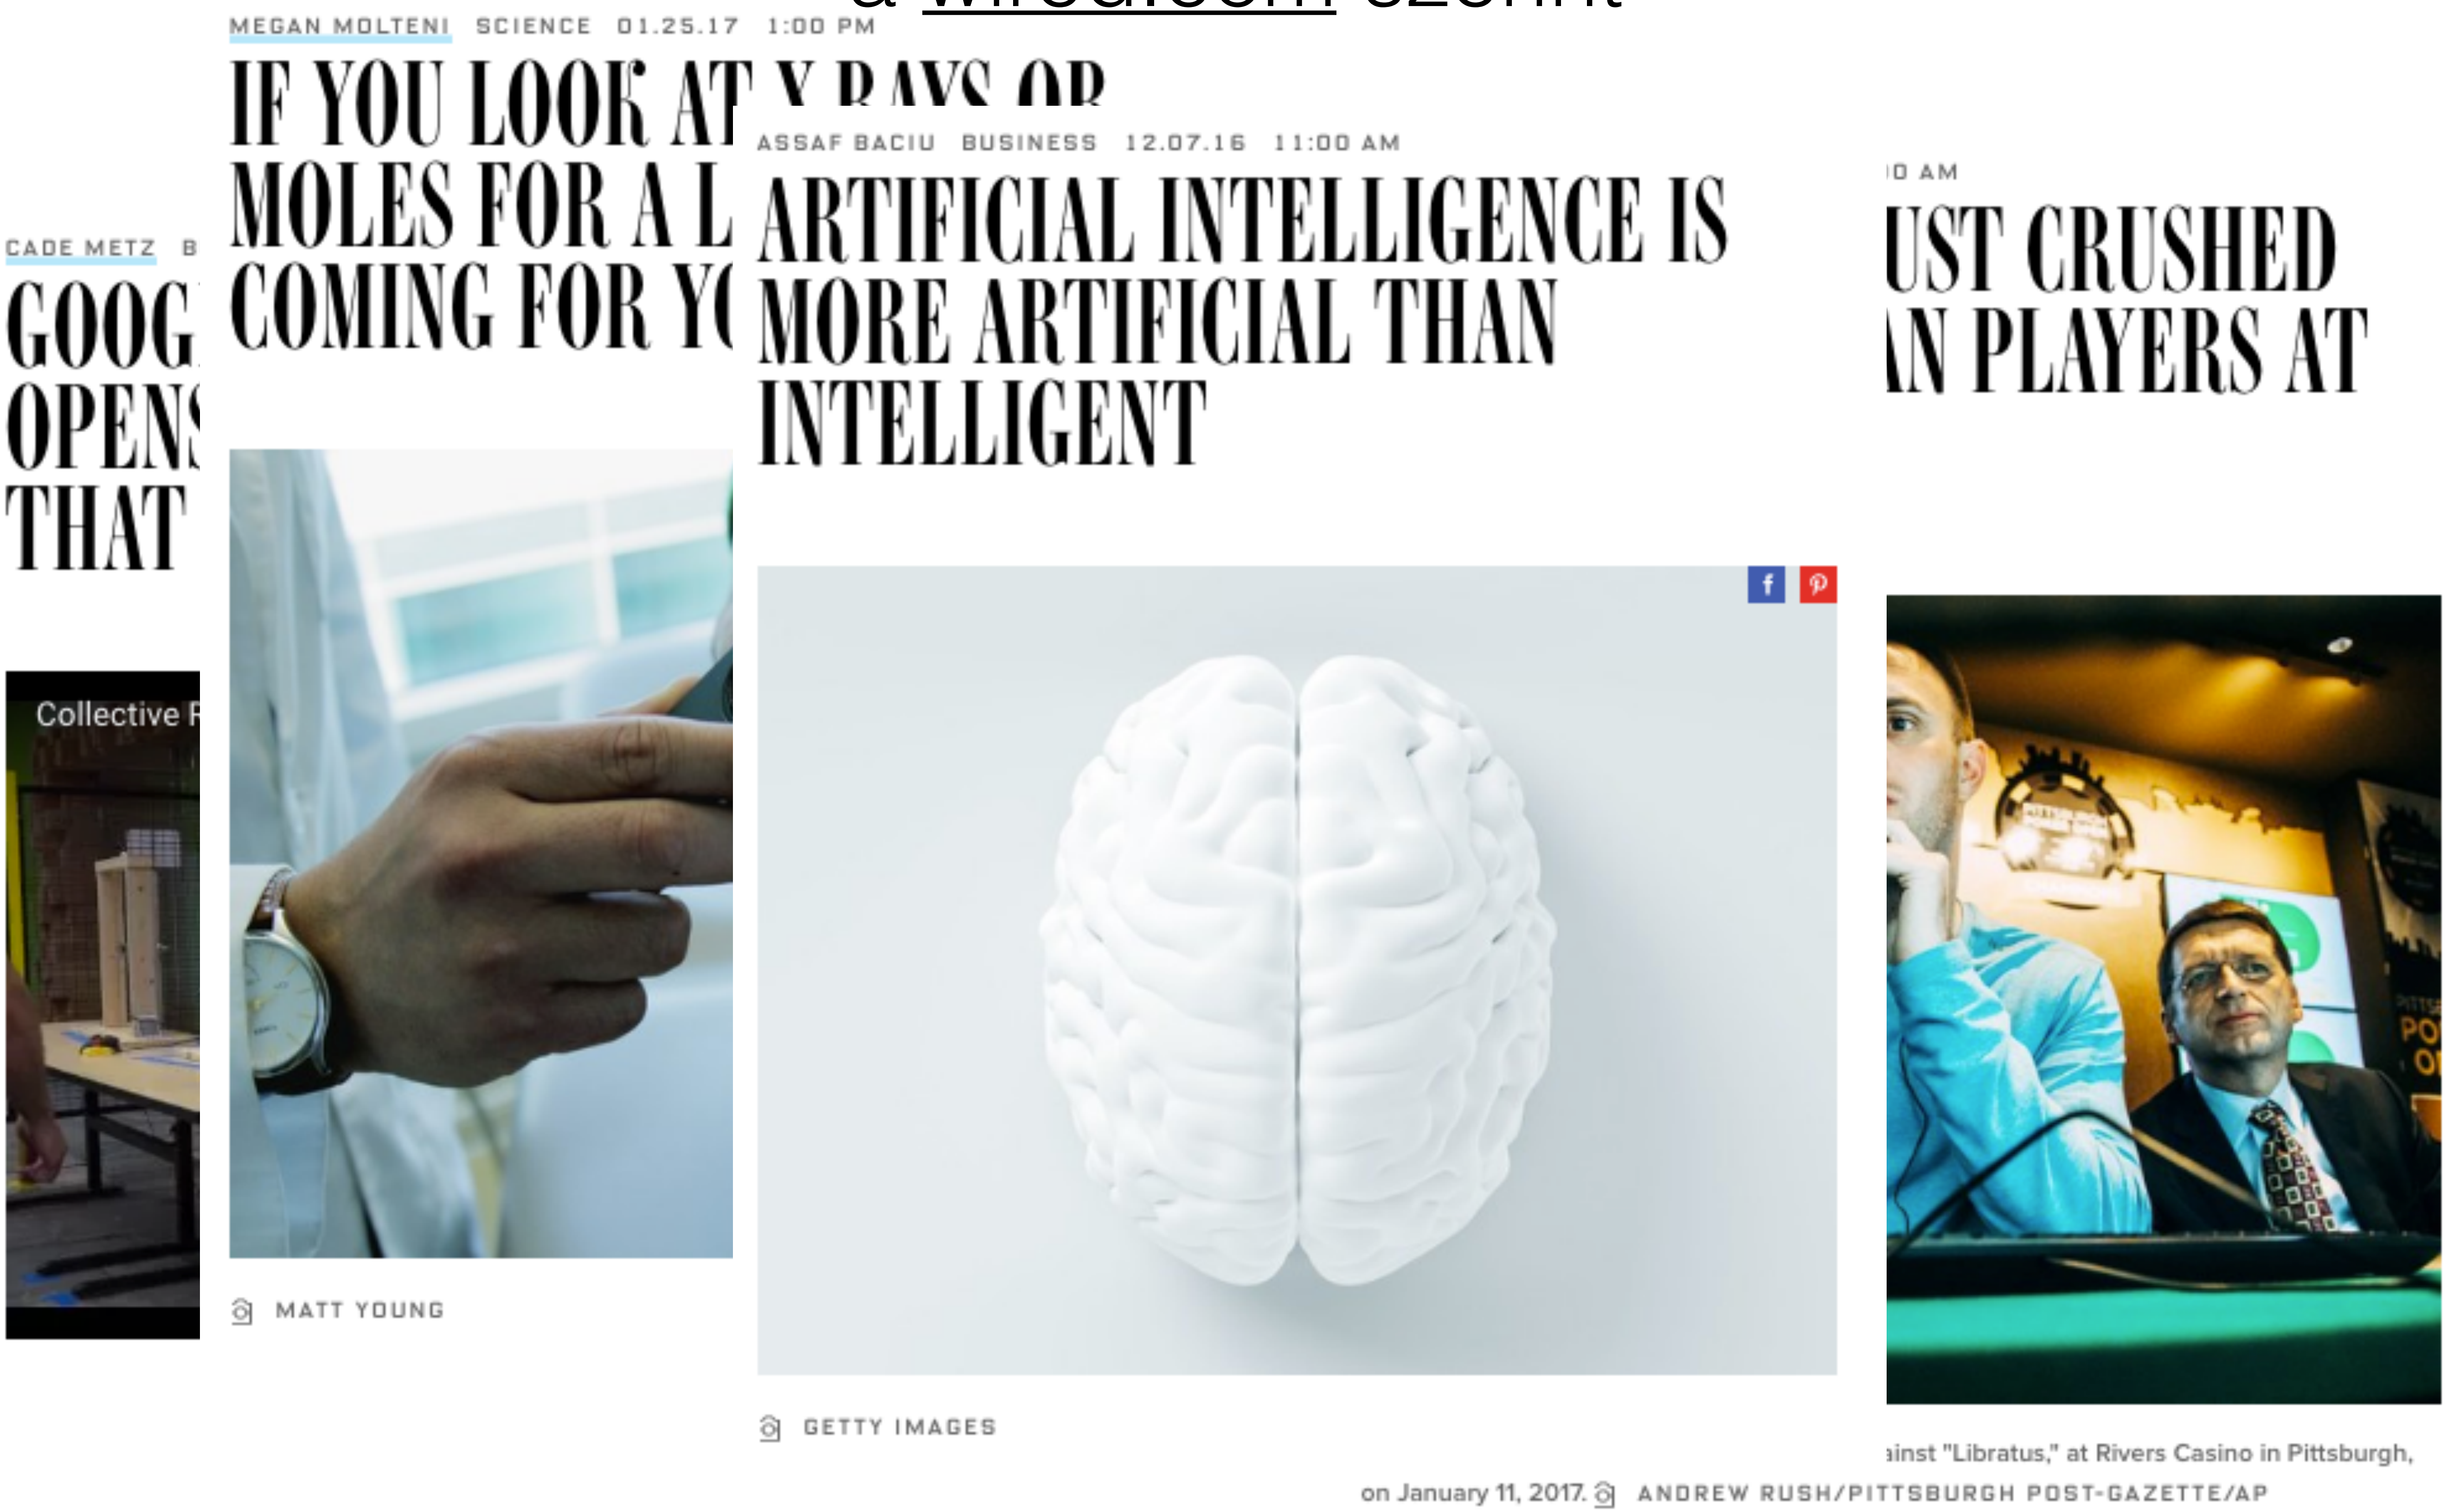

How does an intelligent system look like?

- What does this question actually mean? What are the answers that we would accept?
- Do we have an answer to this question?

Theoretical neuroscience

Theoretical neuroscience

- We approach the brain from a science point of view. We want theories that are

Theoretical neuroscience

- We approach the brain from a science point of view. We want theories that are
    - predictive

Theoretical neuroscience

- We approach the brain from a science point of view. We want theories that are
  - predictive
  - is normative

Theoretical neuroscience

- We approach the brain from a science point of view. We want theories that are
  - predictive
  - is normative
  - synthesizing

Theoretical neuroscience

- We approach the brain from a science point of view. We want theories that are
  - predictive
  - is normative
  - synthesizing
  - provides biological insights

Theoretical neuroscience

- We approach the brain from a science point of view. We want theories that are

  - predictive

  - is normative

  - synthesizing

  - provides biological insights

- We seek the mathematical models of brain functions

Theoretical neuroscience

- V f view.
  V

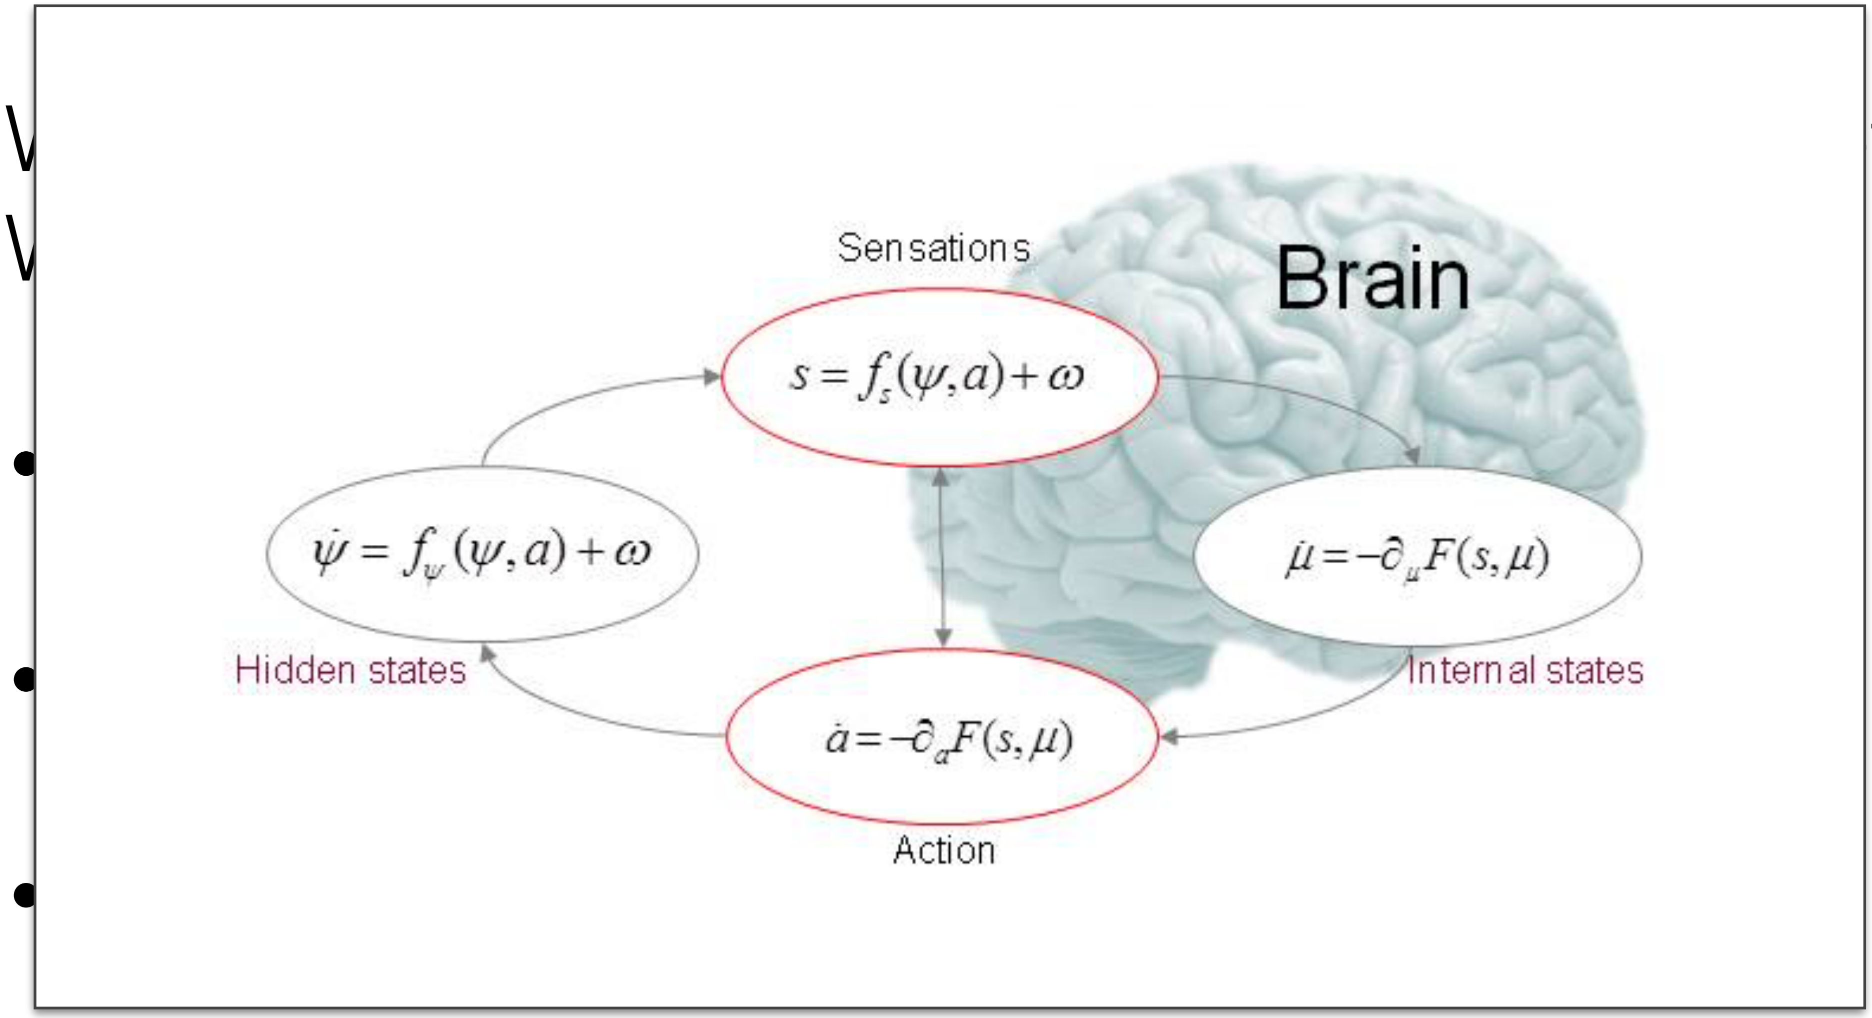

  - provides biological insights

- We seek the mathematical models of brain functions

Theoretical neuroscience

- We approach the brain from a science point of view. We want theories that are
  - predictive
  - is normative
  - synthesizing
  - provides biological insights
- We seek the mathematical models of brain functions

Theoretical neuroscience

- We approach the brain from a science point of view. We want theories that are
  - predictive
  - is normative
  - synthesizing
  - provides biological insights
- We seek the mathematical models of brain functions

Theoretical neuroscience

- We approach the brain from a science point of view. We want theories that are
  - predictive
  - is normative
  - synthesizing
  - provides biological insights
- We seek the mathematical models of brain functions

Predictive models of brain

- we build models (theories) of observations and we intend to predict phenomena in novel situations
- Computational neuroscience
  - predicting biophysical quantities based on **physiological** measurements
- Computational cognitive science
  - predicting behavioral quantities based on **psychophysics** experiments

What do we want to predict?

- I.e. what are the properties of the nervous system that are relevant and which are those that we deem irrelevant?
- When would we be more satisfied:
  - A model that describes neuronal responses well but provides little insight about its functioning?
  - A model that reproduces the behavior well but has little resemblance to the actual brain?

Levels of abstraction

- **Computation** - determining the challenge:
  through the goals of the biological agent phrasing a mathematical model that is capable of addressing the challenge

- **Algorithm** - solving the challenge can be achieved in many different ways that can be weighed based on different factors

- **Implementation** - physical realization of the algorithm, in the case of neuroscience, based on neurons, spikes, etc

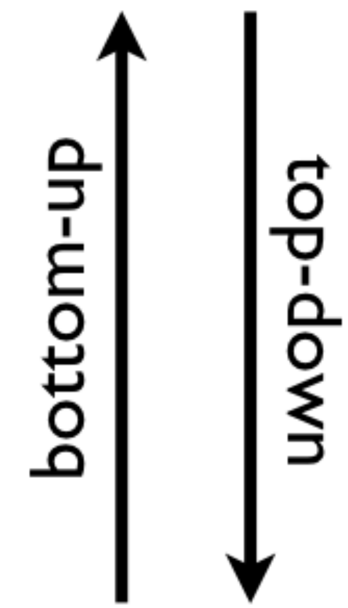

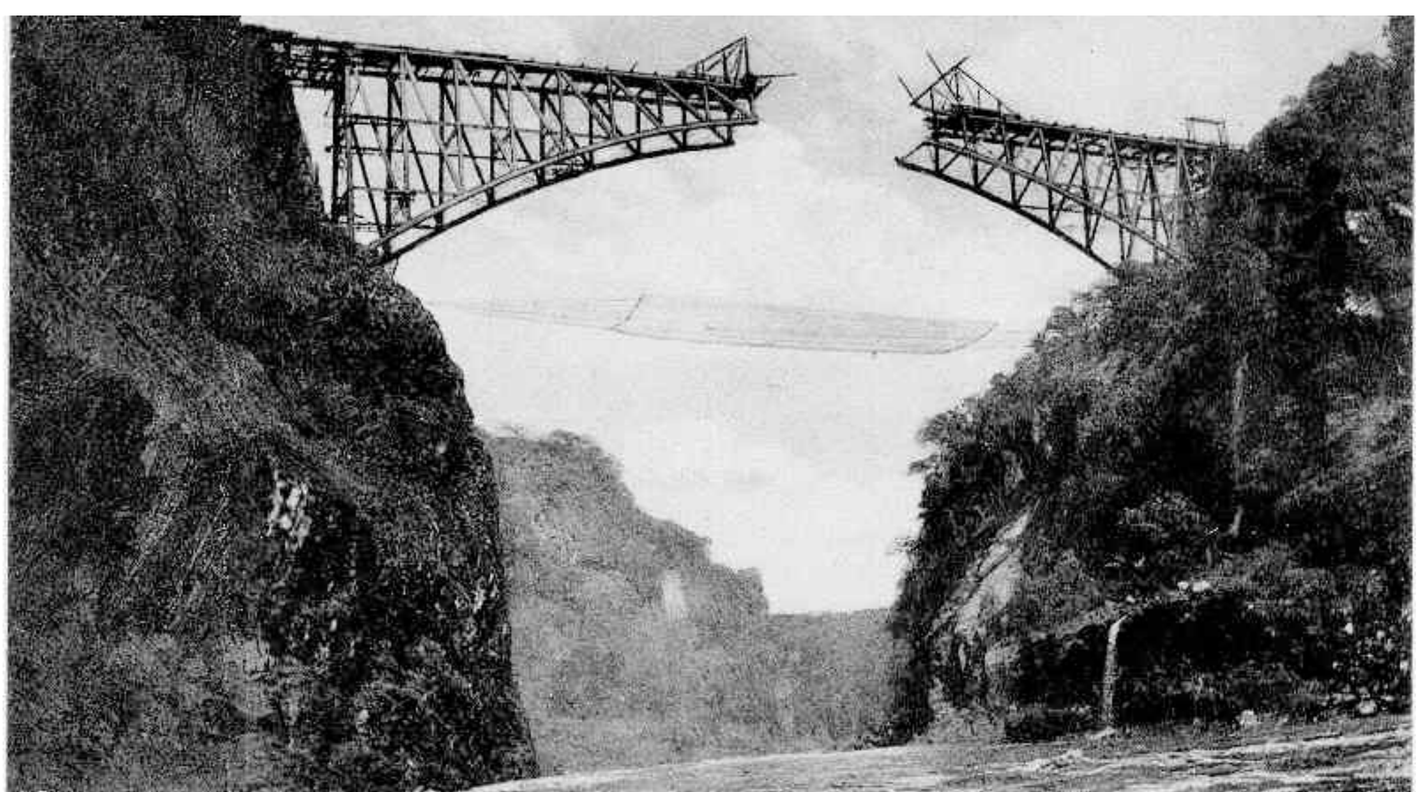

David Marr, 1976

Normative approach

- First, we want to reproduce the high-level properties of the system
  - next, we want to identify the structural similarity
  - is called top-down modelling
- In contrast, in bottom-up modelling
  - attempts to best describe the structural and dynamical properties
  - the function (behavior) is an emergent property of the system

Normative approach

- First, we want to reproduce the high-level properties of the system
  - next, we want to identify the structural similarity
  - is called top-down modelling
- In contrast, in bottom-up modelling
  - attempts to best describe the structural and dynamical properties
  - the function (behavior) is an emergent property of the system

Normative approach

•

“A wing would be the most mystifying structure if one did not know that birds flew. One might observe that it could be extended a considerable distance, that it had smooth covering of feathers with conspicuous markings, that it was operated by powerful muscles, and that strength and lightness were prominent features of its constructions. These are important facts, but in themselves they do not tell us that birds fly. Yet, without knowing this and without understanding something of the principles of flight, a more detailed examination of the wing itself would probably be unrewarding.”

Horace Barlow

•

properties

- the function (behavior) is an emergent property of the system

predictive power

predictive power

Theoretical neuroscience

- We approach the brain from a science point of view. We want theories that are
  - predictive
  - is normative
  - synthesizing
  - provides biological insights
- We seek the mathematical models of brain functions

Synthesizing theories

- The main goal of neuroscience to find theories that explains main aspects of neural processes and behavior using unufying principles
- Synthesizing theories from other disciplines
  - Laws of Newton in physics
  - Principle of evolution in biology
  - Computational complexity & Turing completeness in computer science

The role of theories in neuroscience

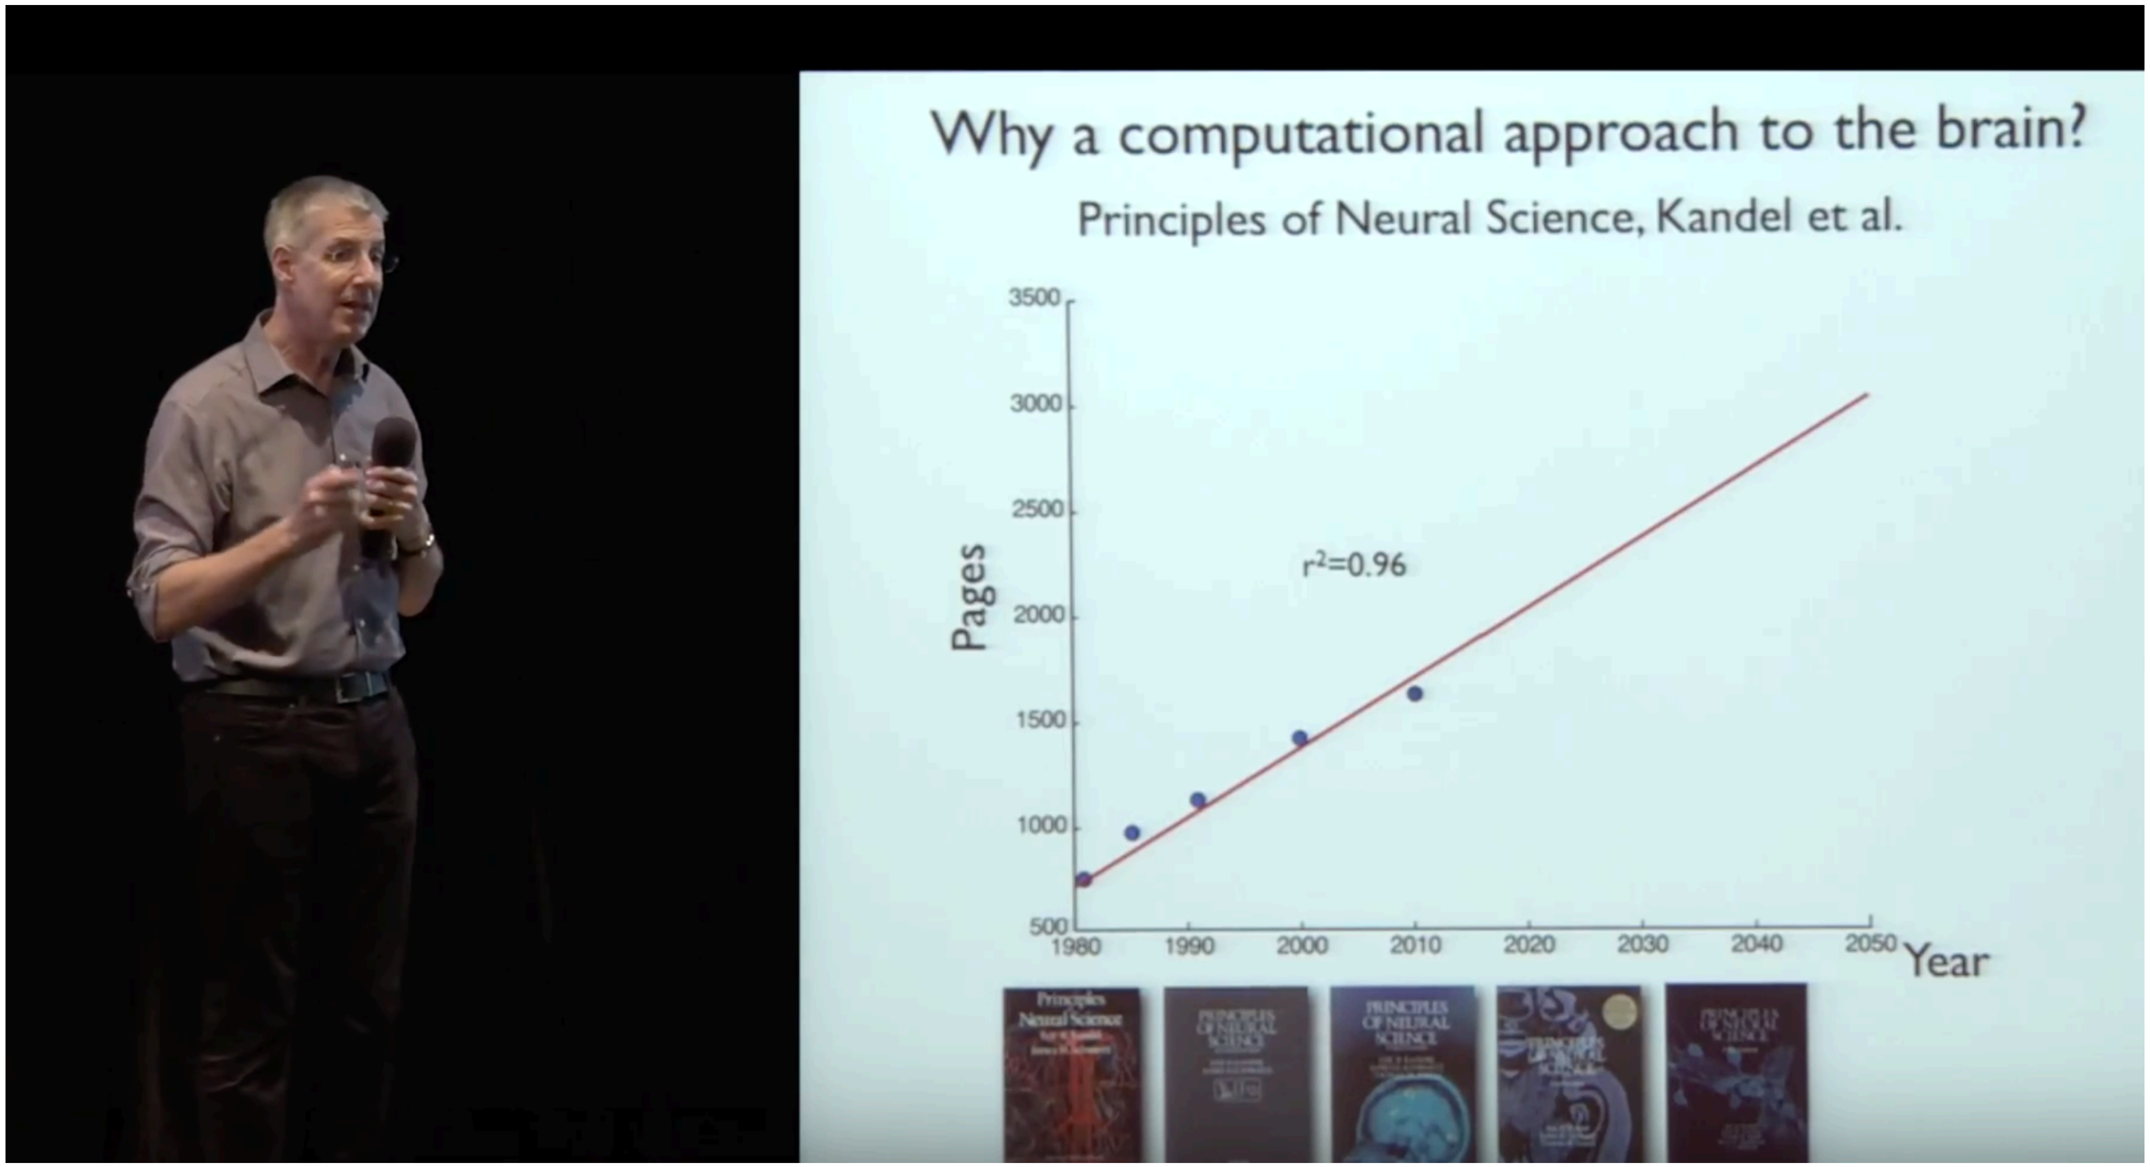

Daniel Wolpert

https://www.youtube.com/watch?v=wTYHF4LAKQI

The role of theories in neuroscience

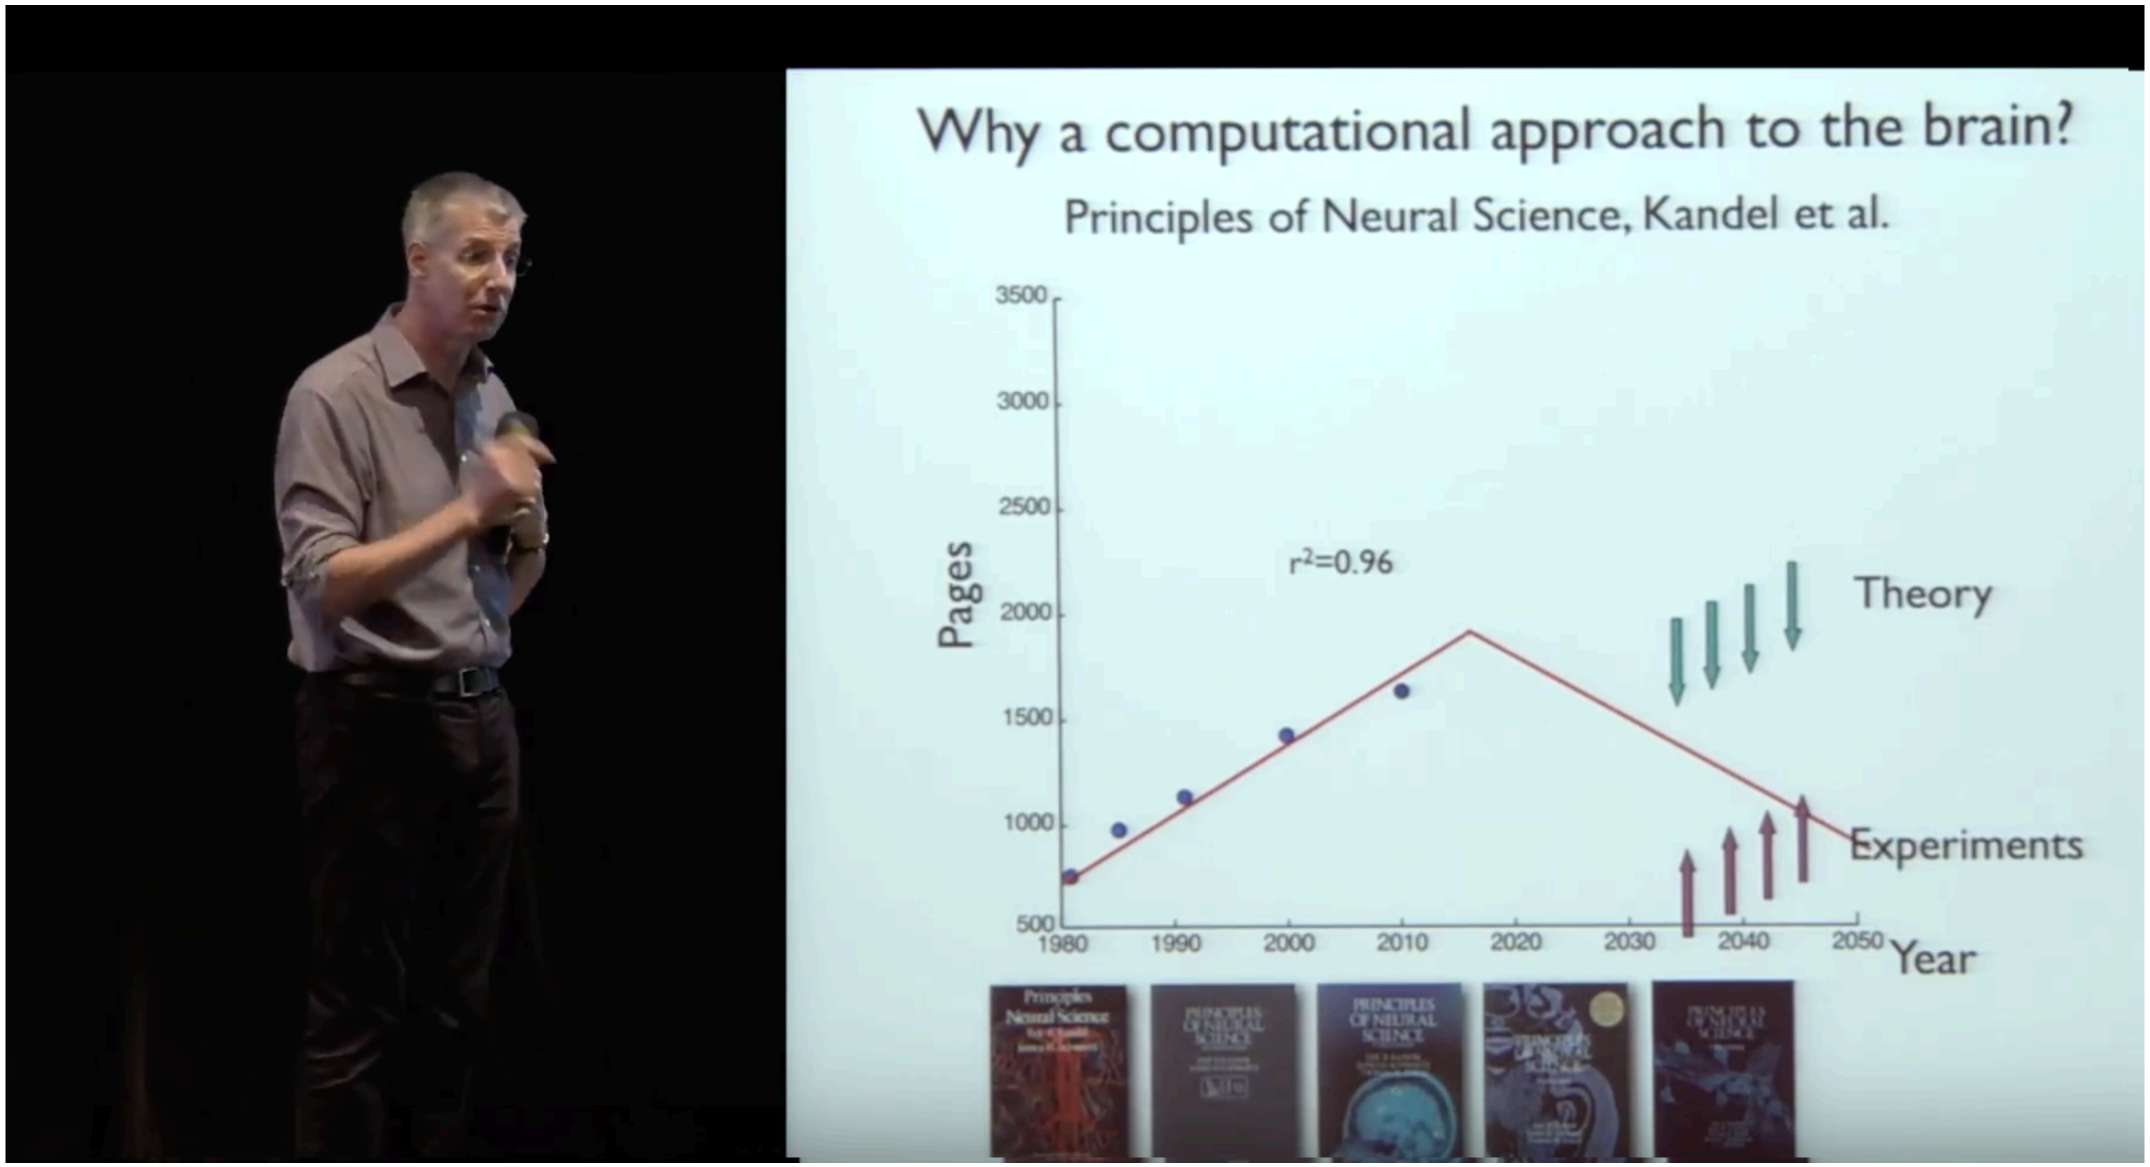

Daniel Wolpert

https://www.youtube.com/watch?v=wTYHF4LAKQI

Principles vs. experiements

Principles vs. experiements

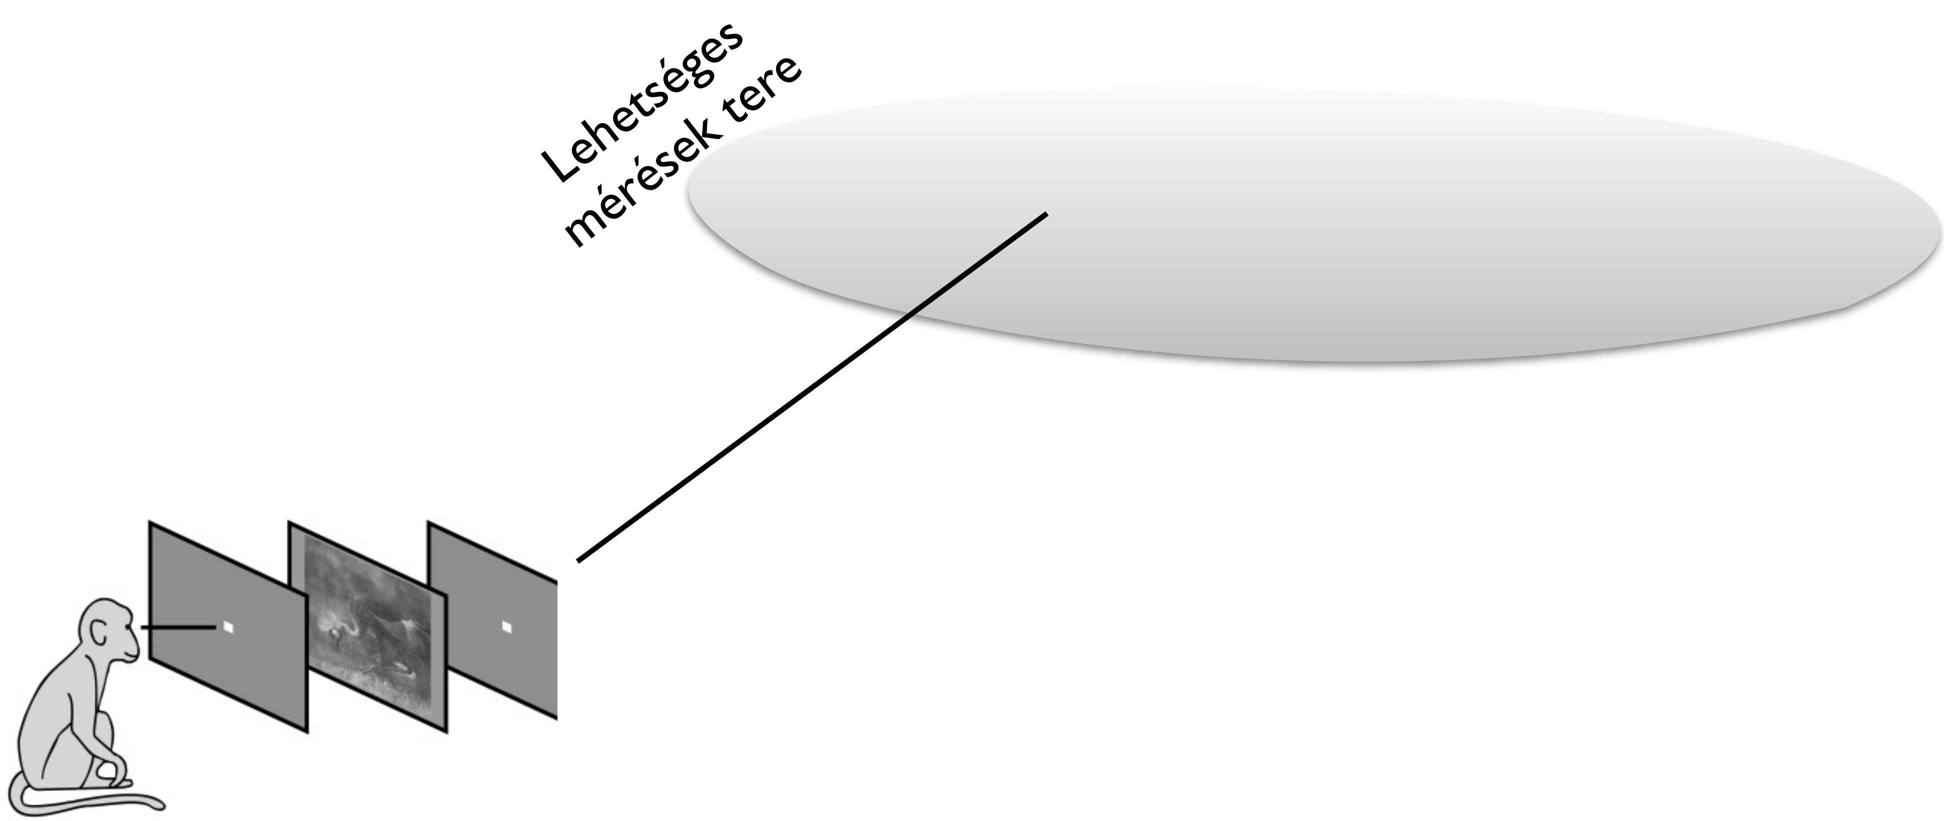

Principles vs. experiements

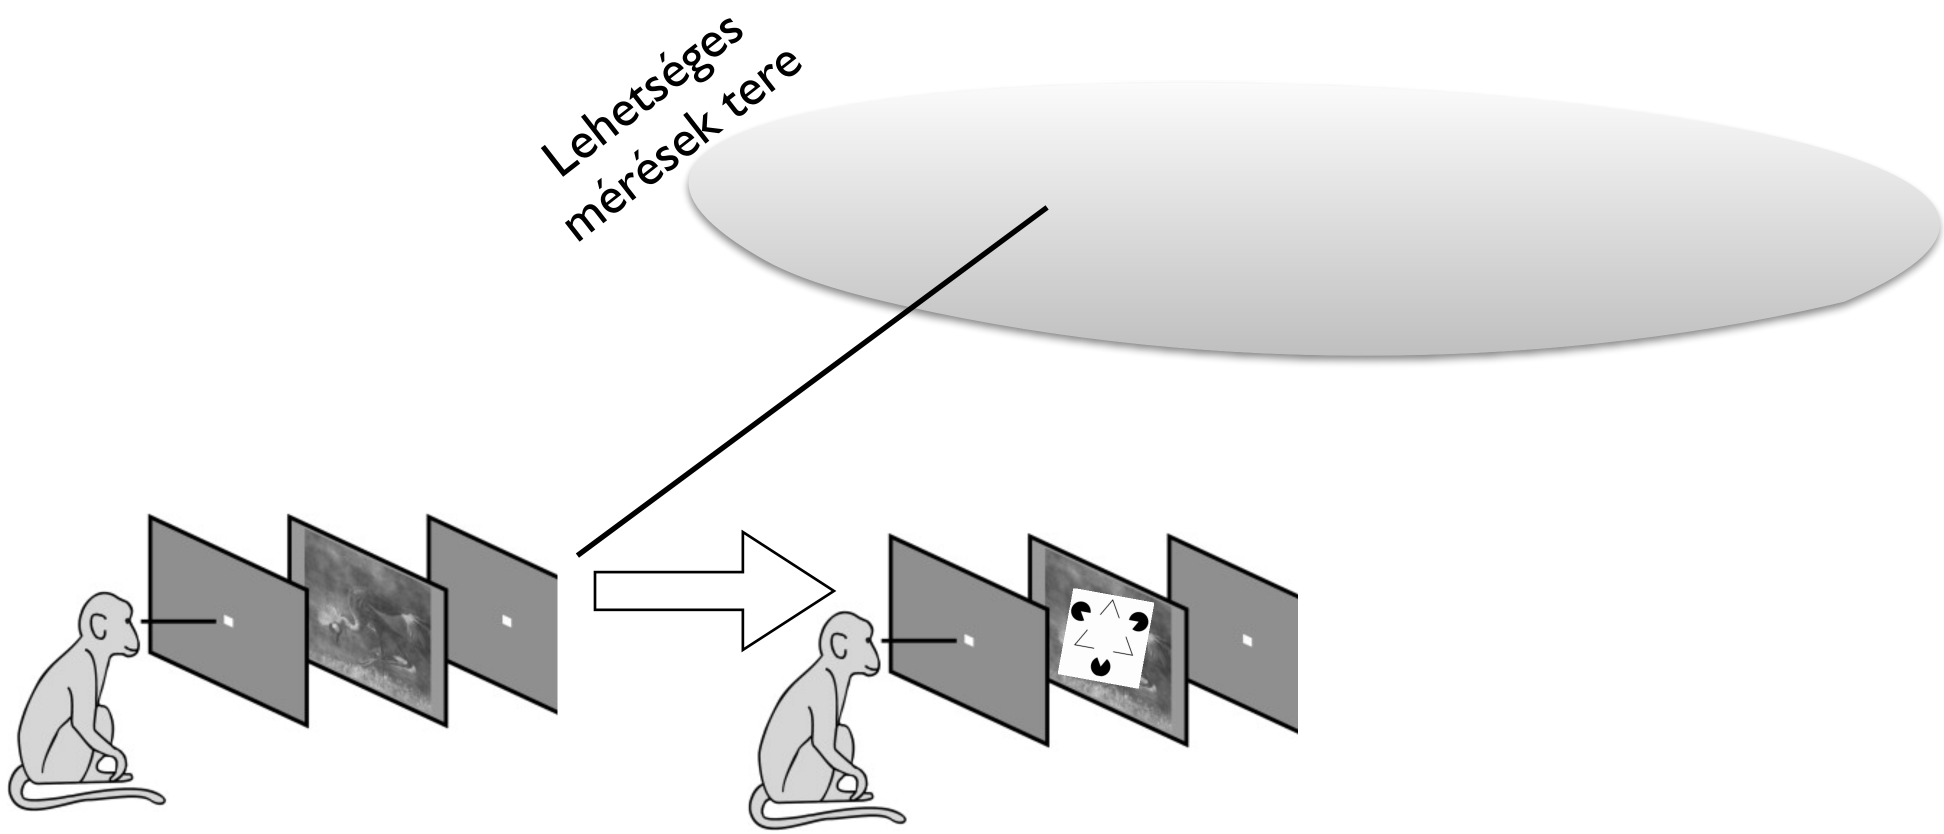

Principles vs. experiements

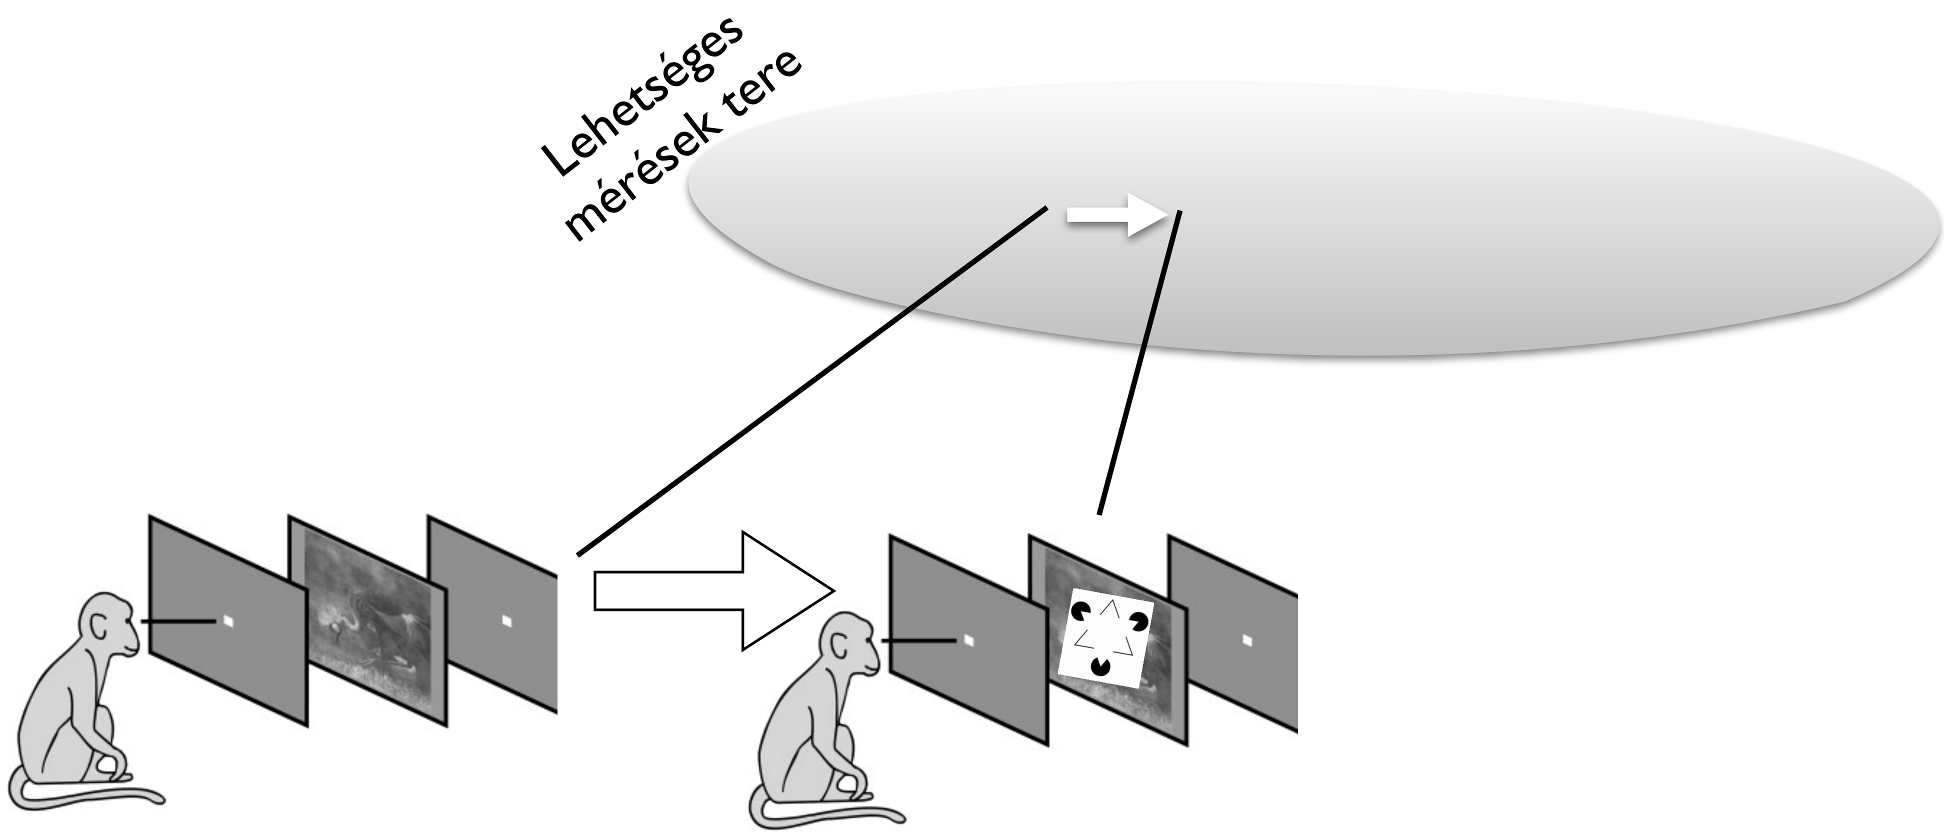

Principles vs. experiements

It may be that brain..

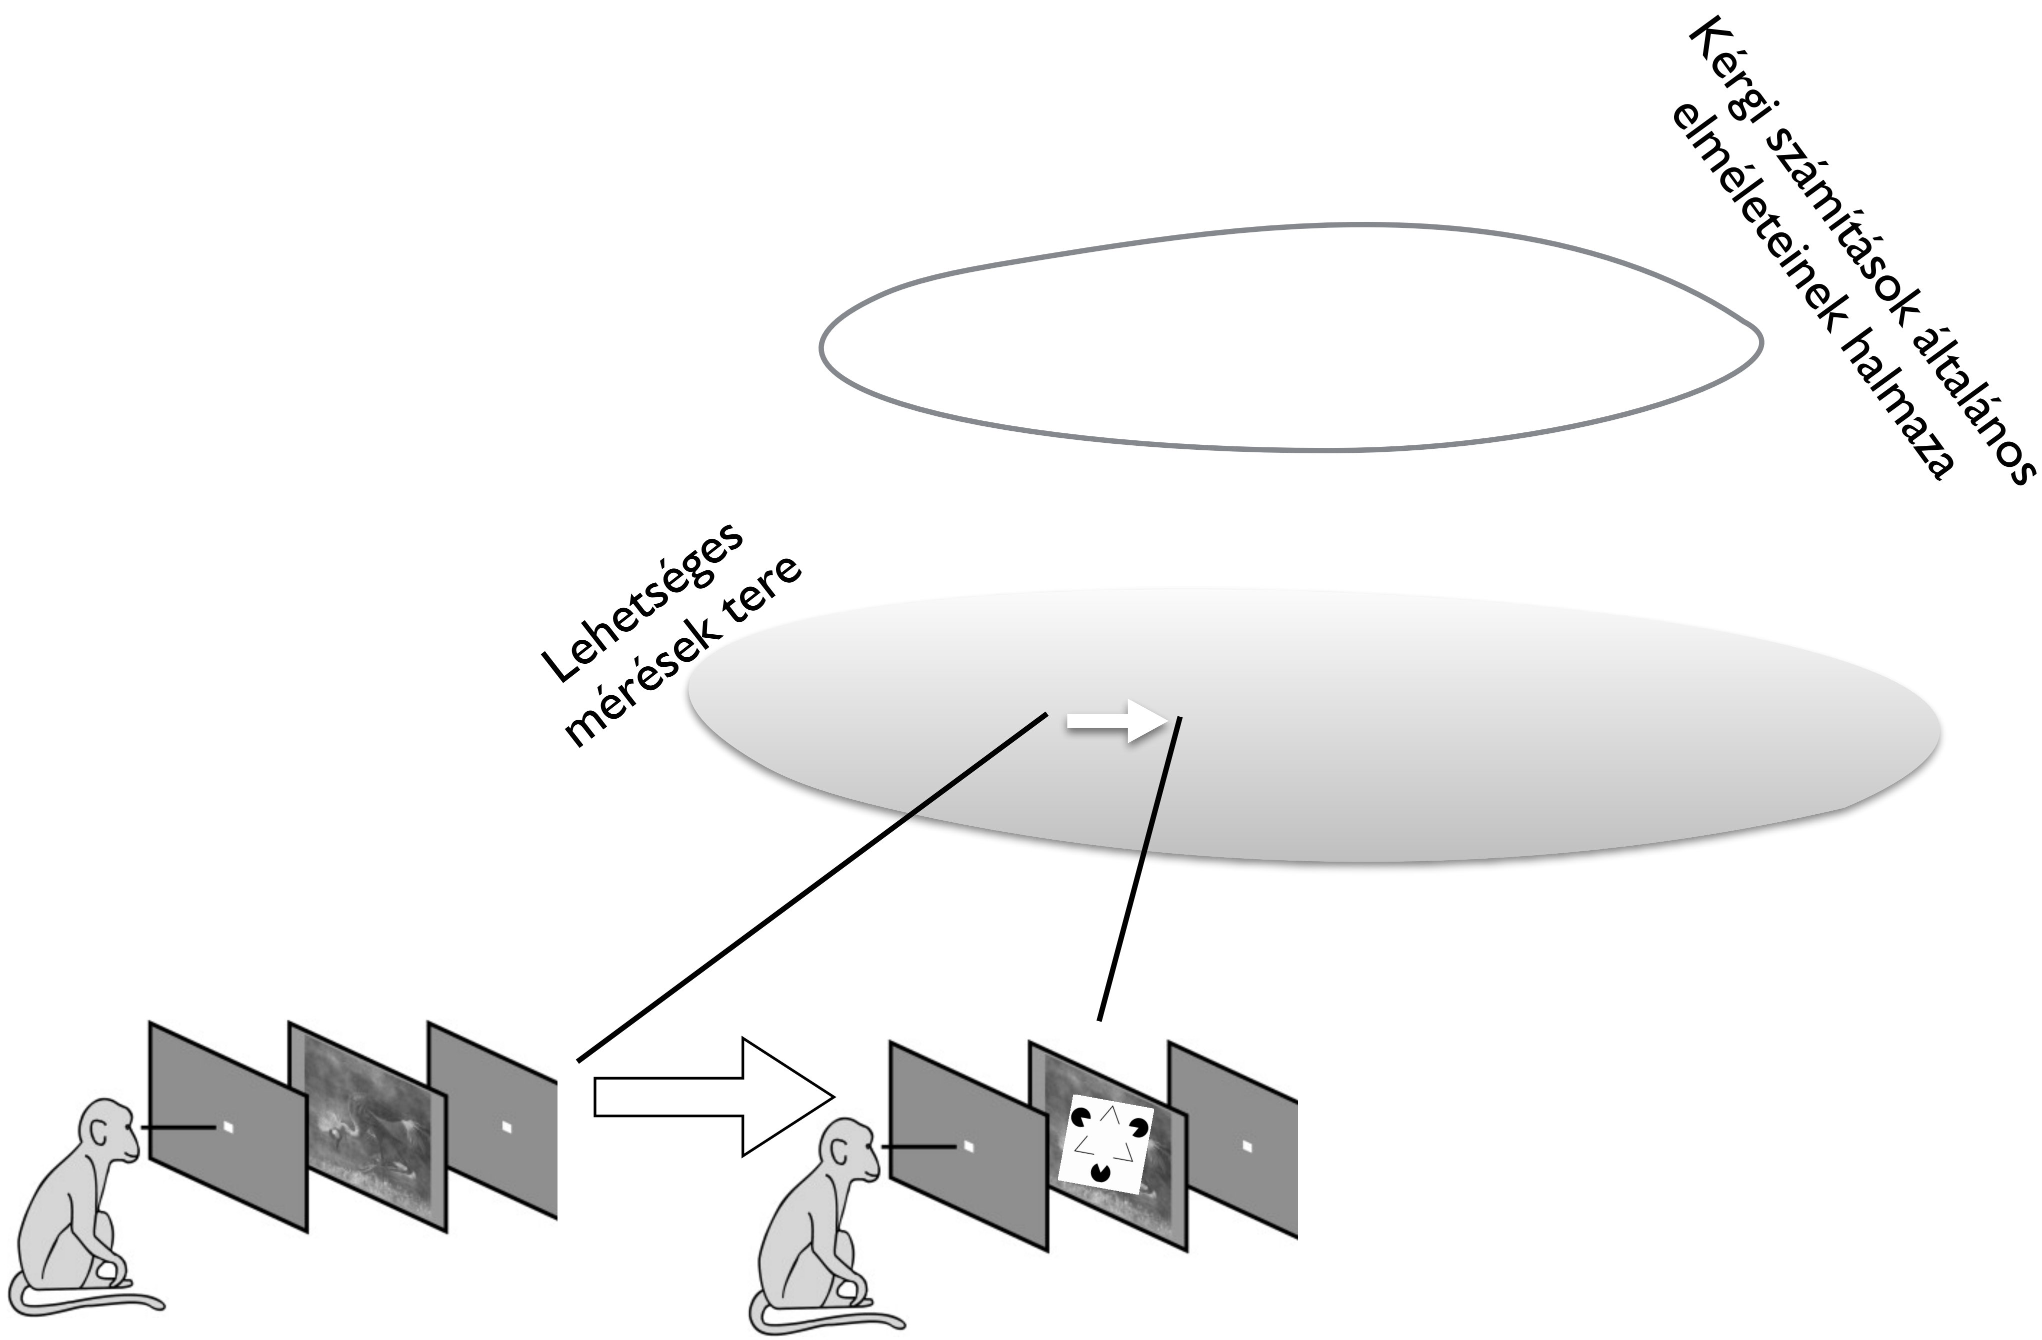

Principles vs. experiements

It may be that brain..

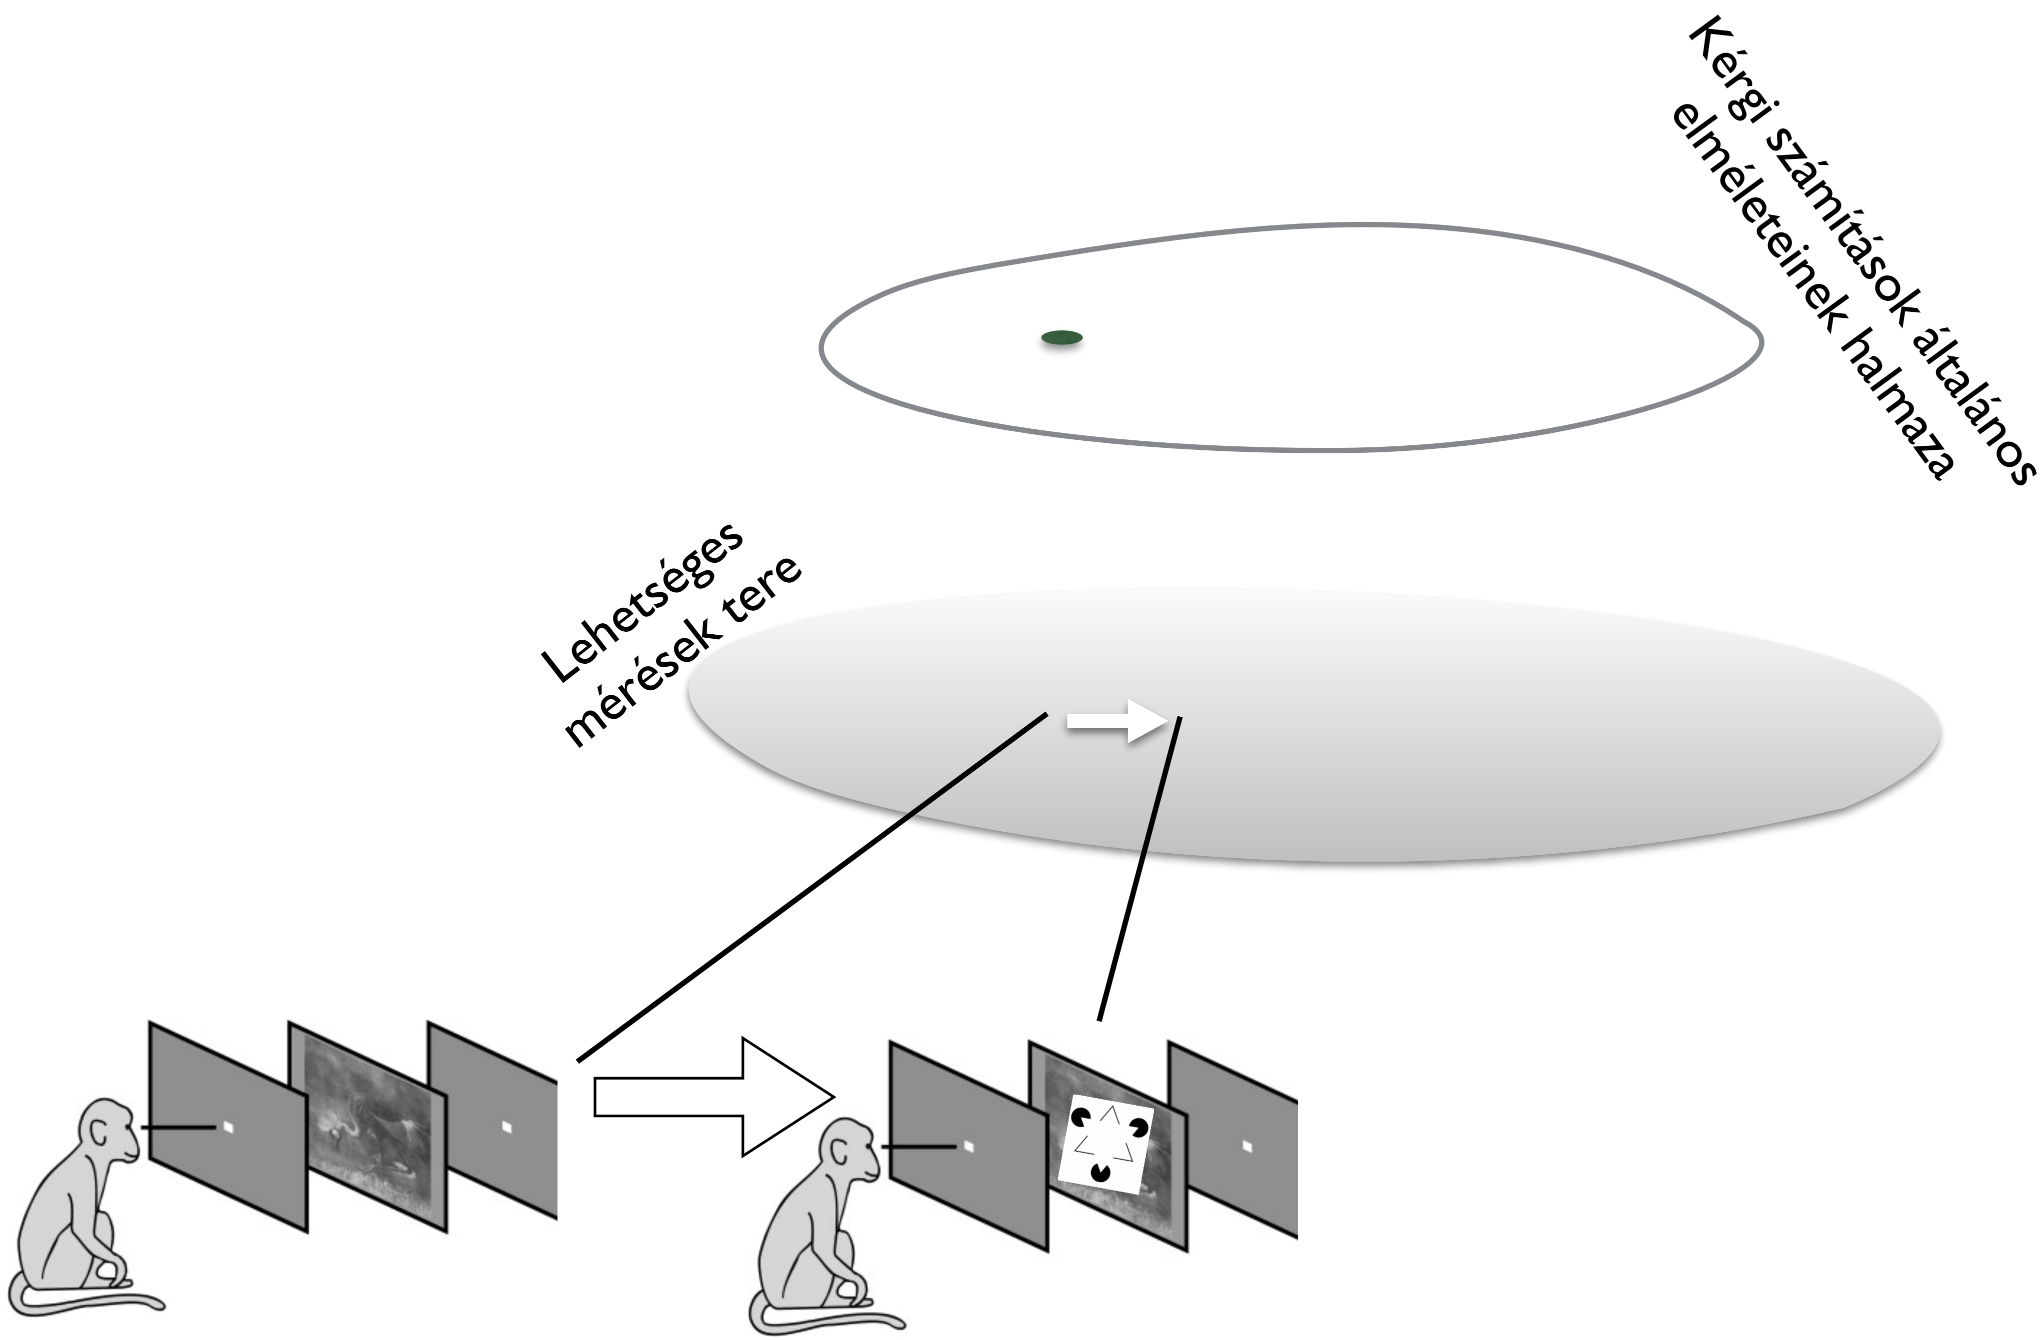

Principles vs. experiements

It may be that brain..

…prepares the inputs
for classification

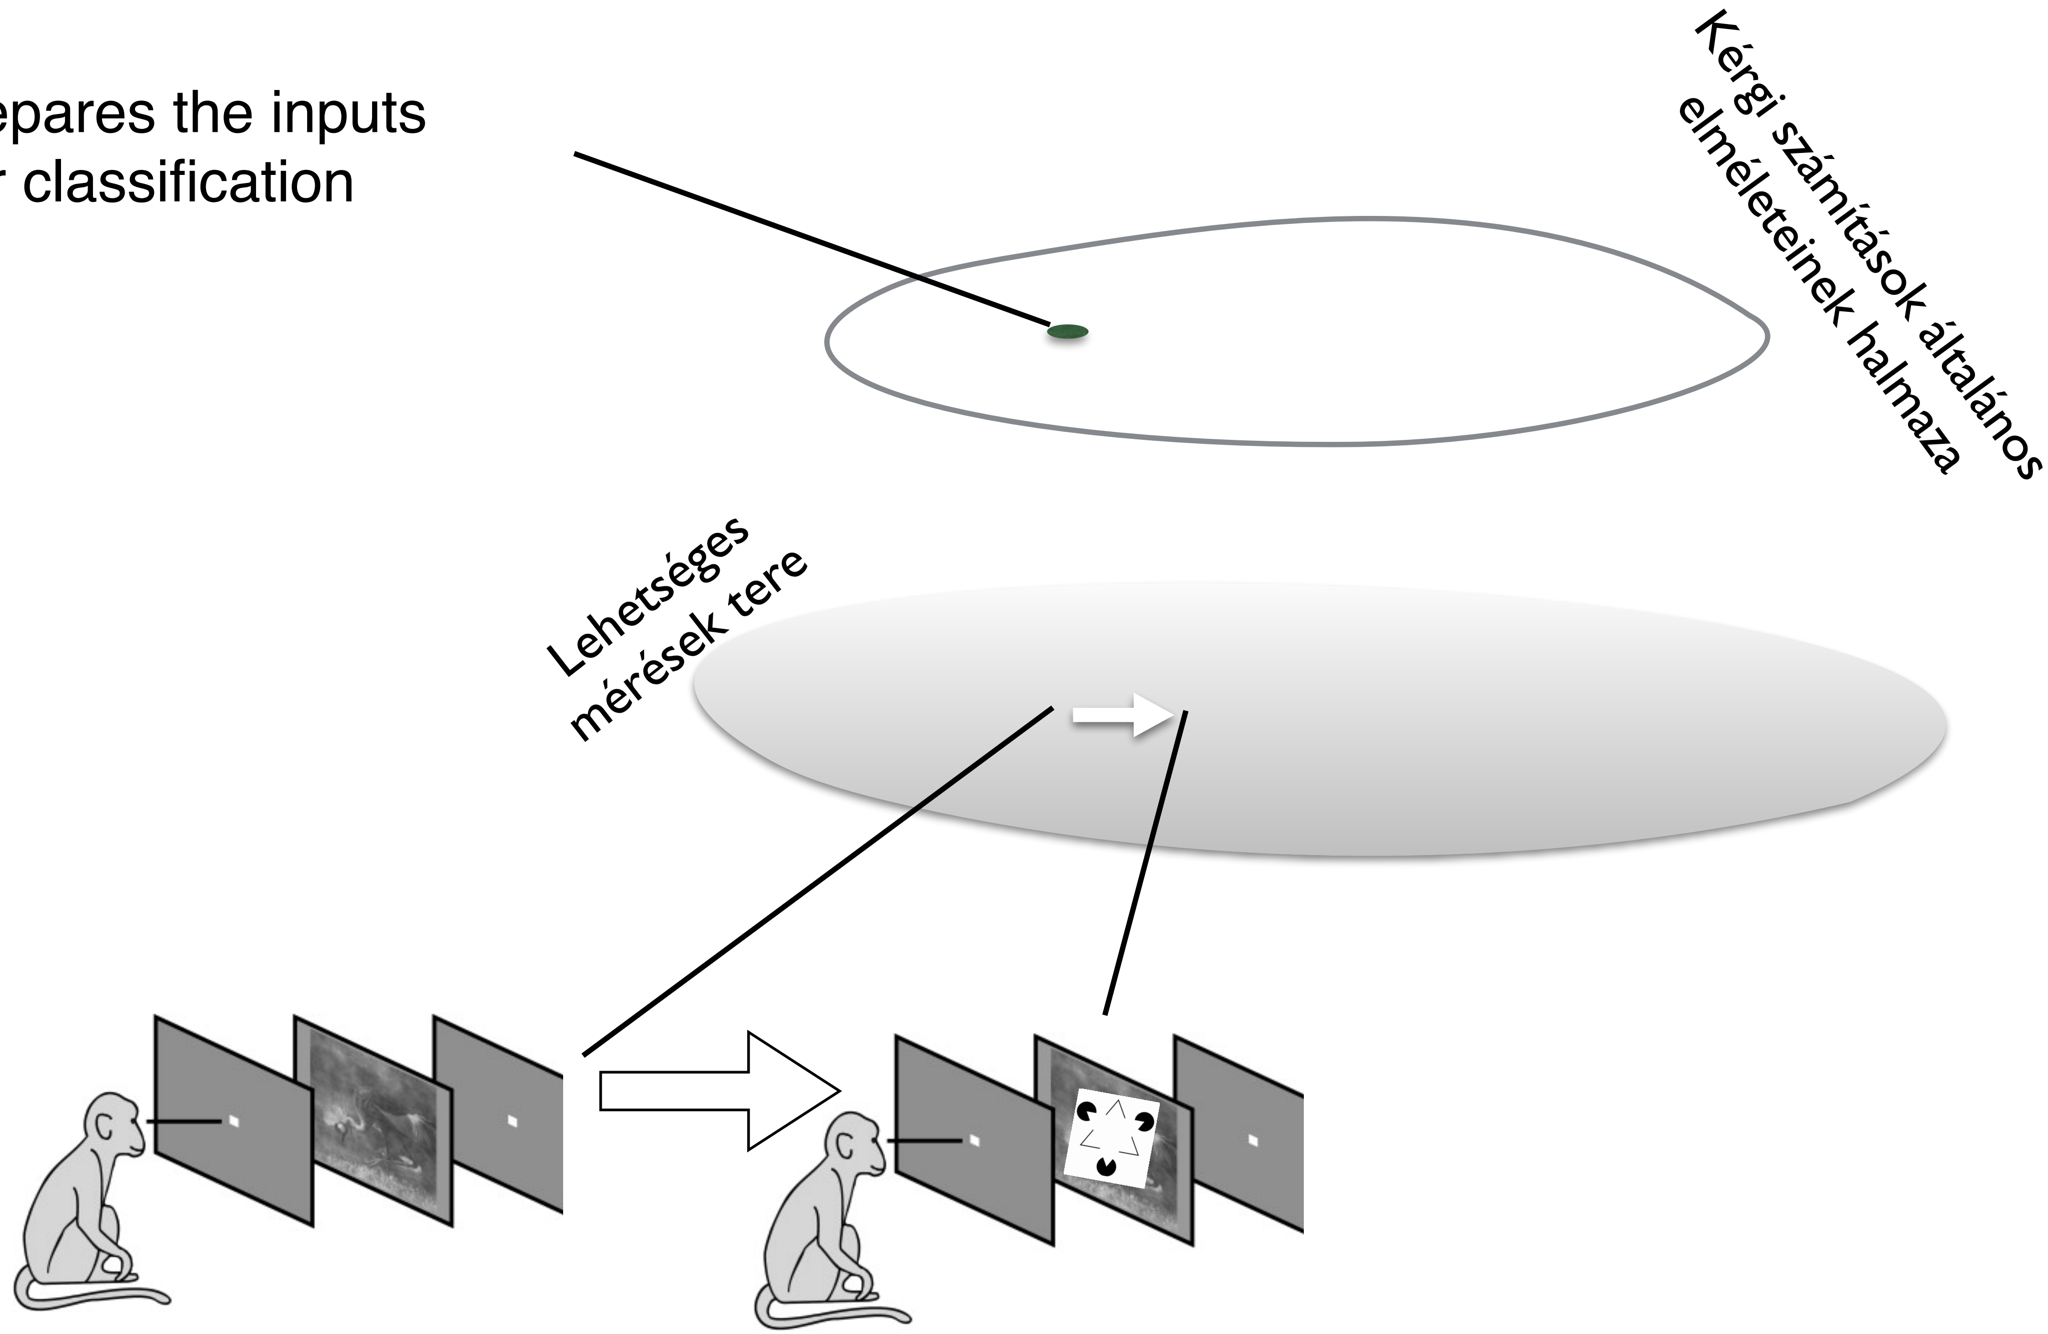

Principles vs. experiements

It may be that brain..

…prepares the inputs for classification

Hierarchical models of object recognition in cortex

Maximilian Riesenhuber & Tomaso Poggio

Nature Neuroscience **2**, 1019–1025 (1999) | Download Citation

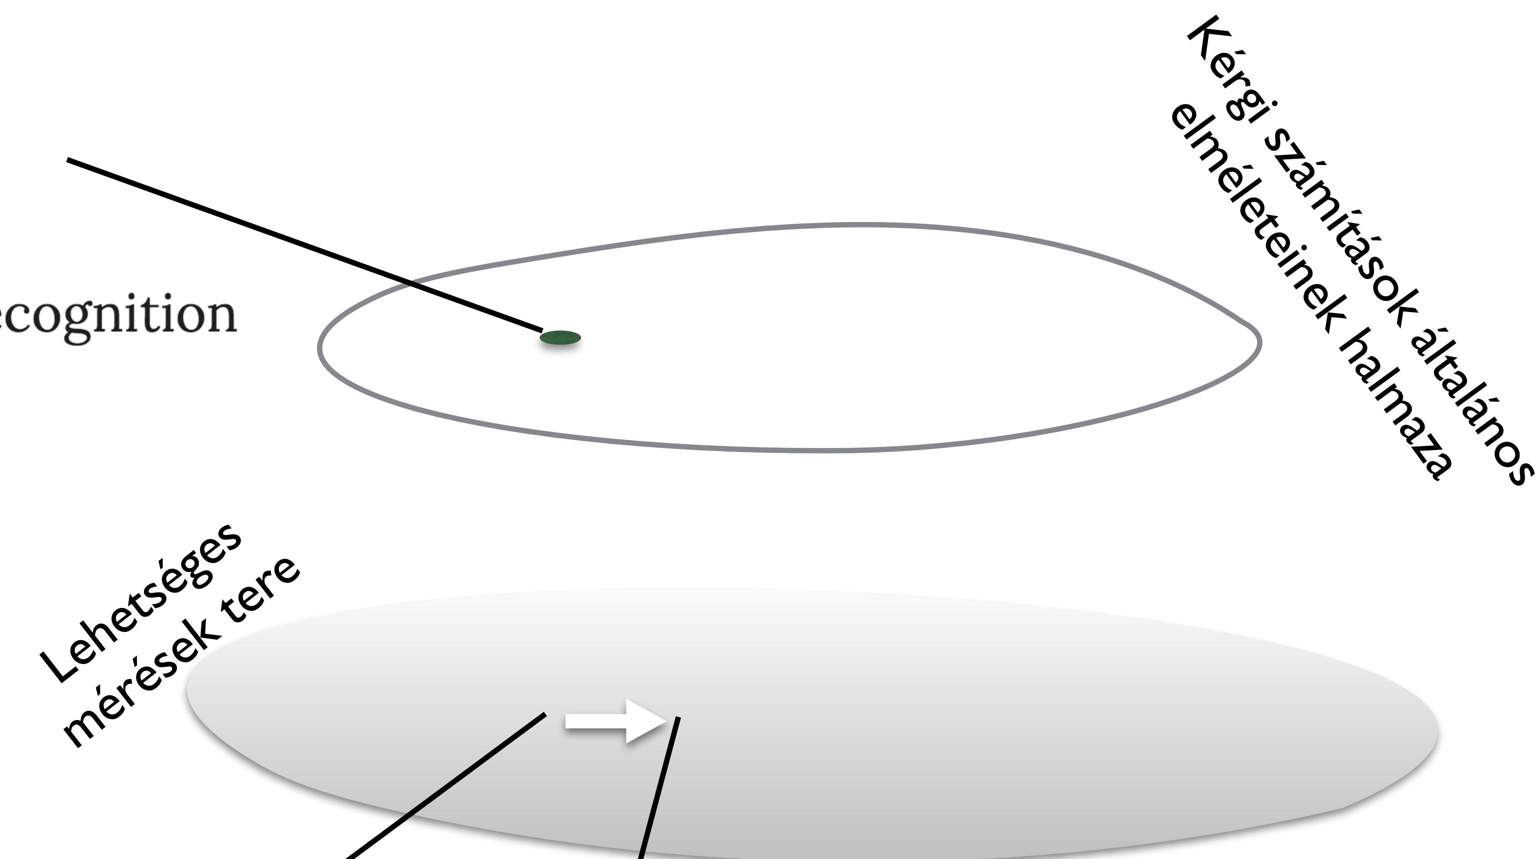

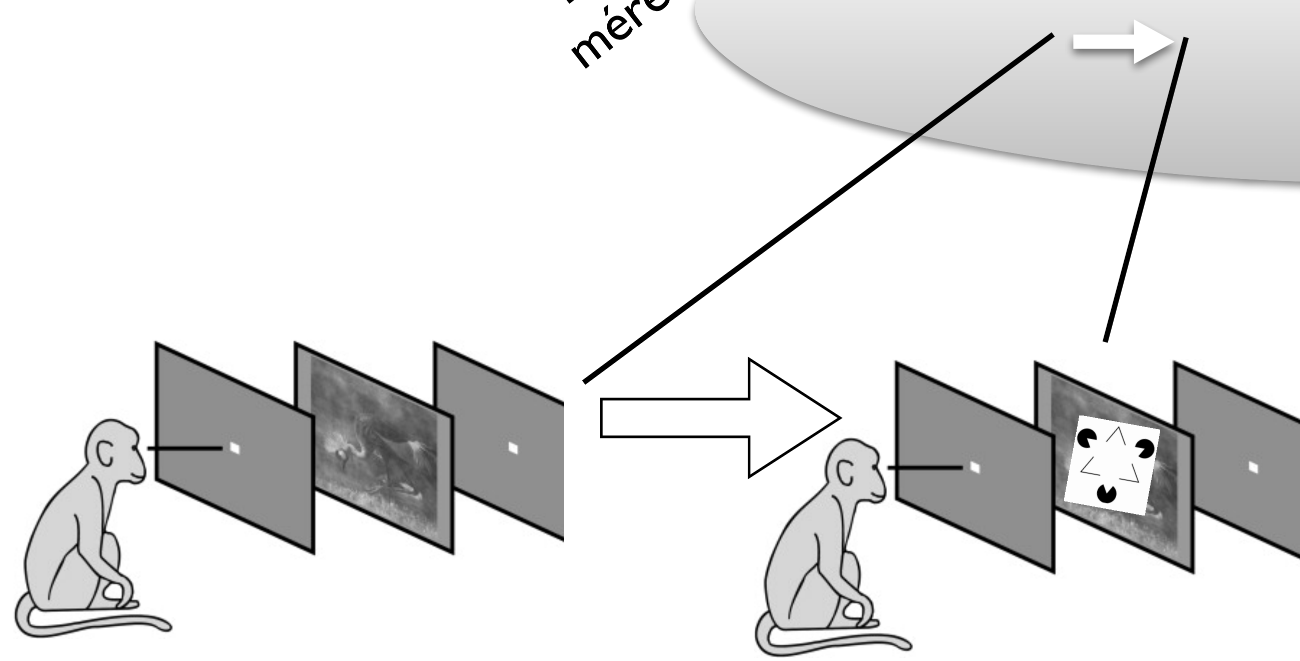

Principles vs. experiements

It may be that brain..

...prepares the inputs for classification

Hierarchical models of object recognition in cortex

Maximilian Riesenhuber & Tomaso Poggio

Nature Neuroscience **2**, 1019–1025 (1999) | Dov

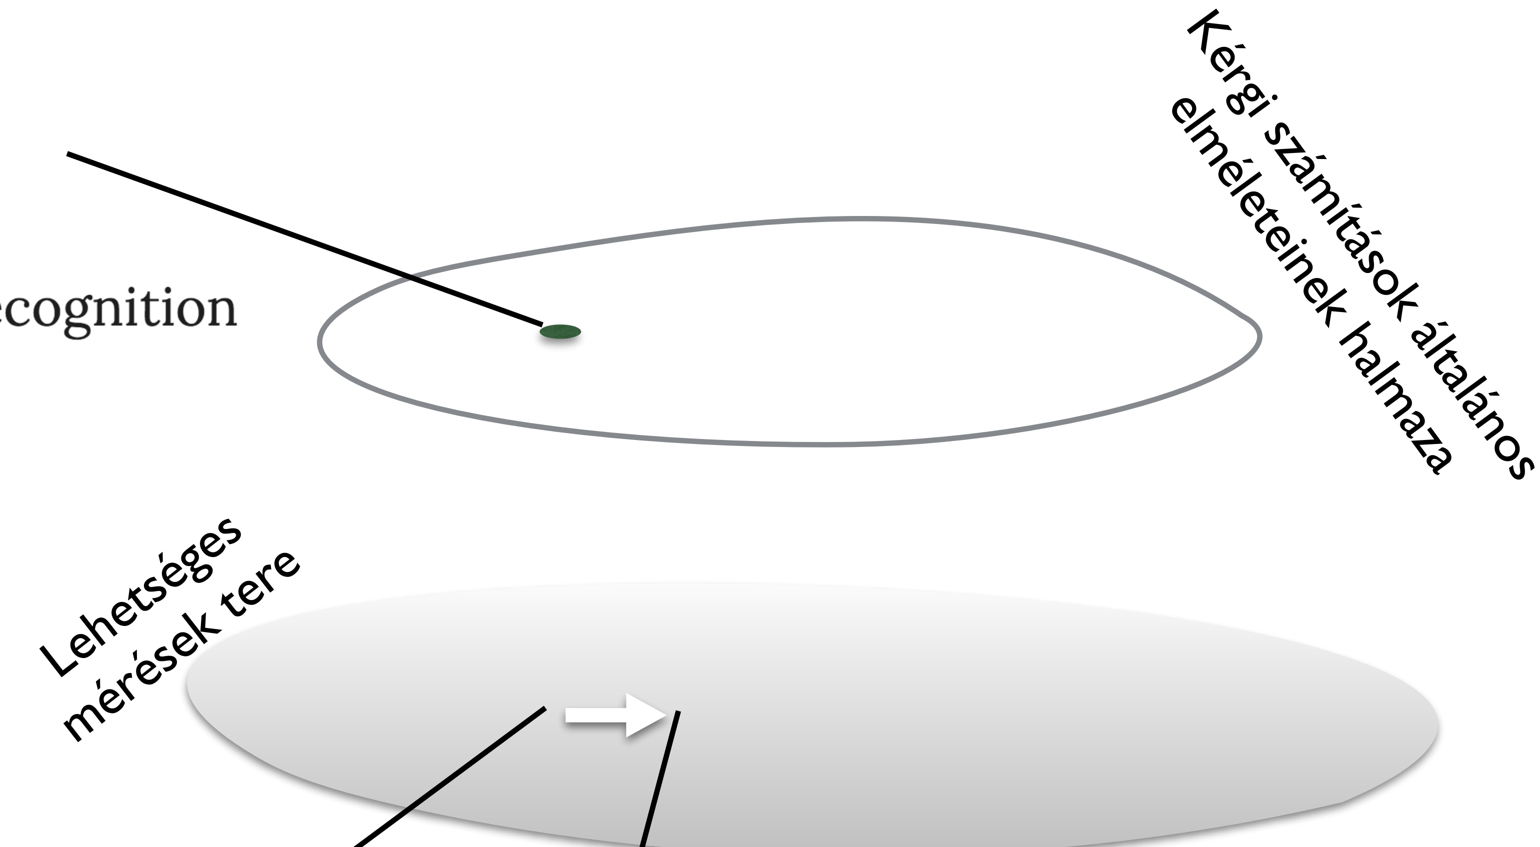

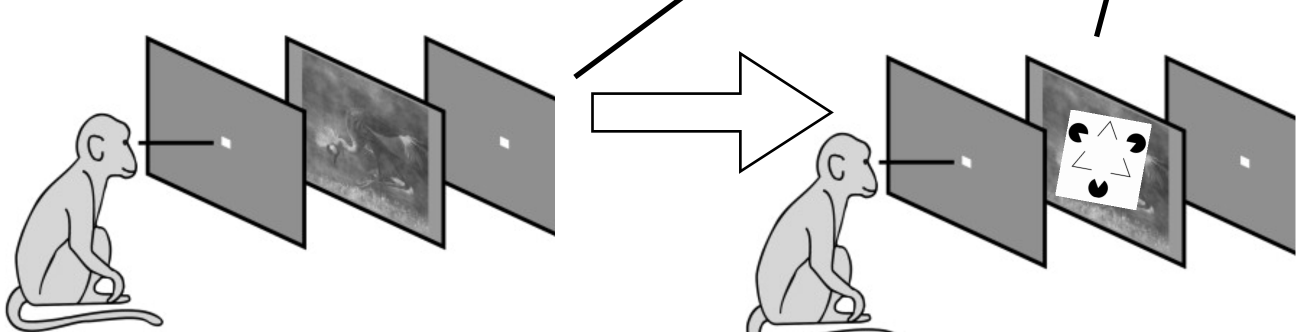

Principles vs. experiements

It may be that brain..

...prepares the inputs
for classification

Hierarchical models of object recognition in cortex

Maximilian Riesenhuber & Tomaso Poggio

Nature Neuroscience **2**, 1019–1025 (1999) | Dov

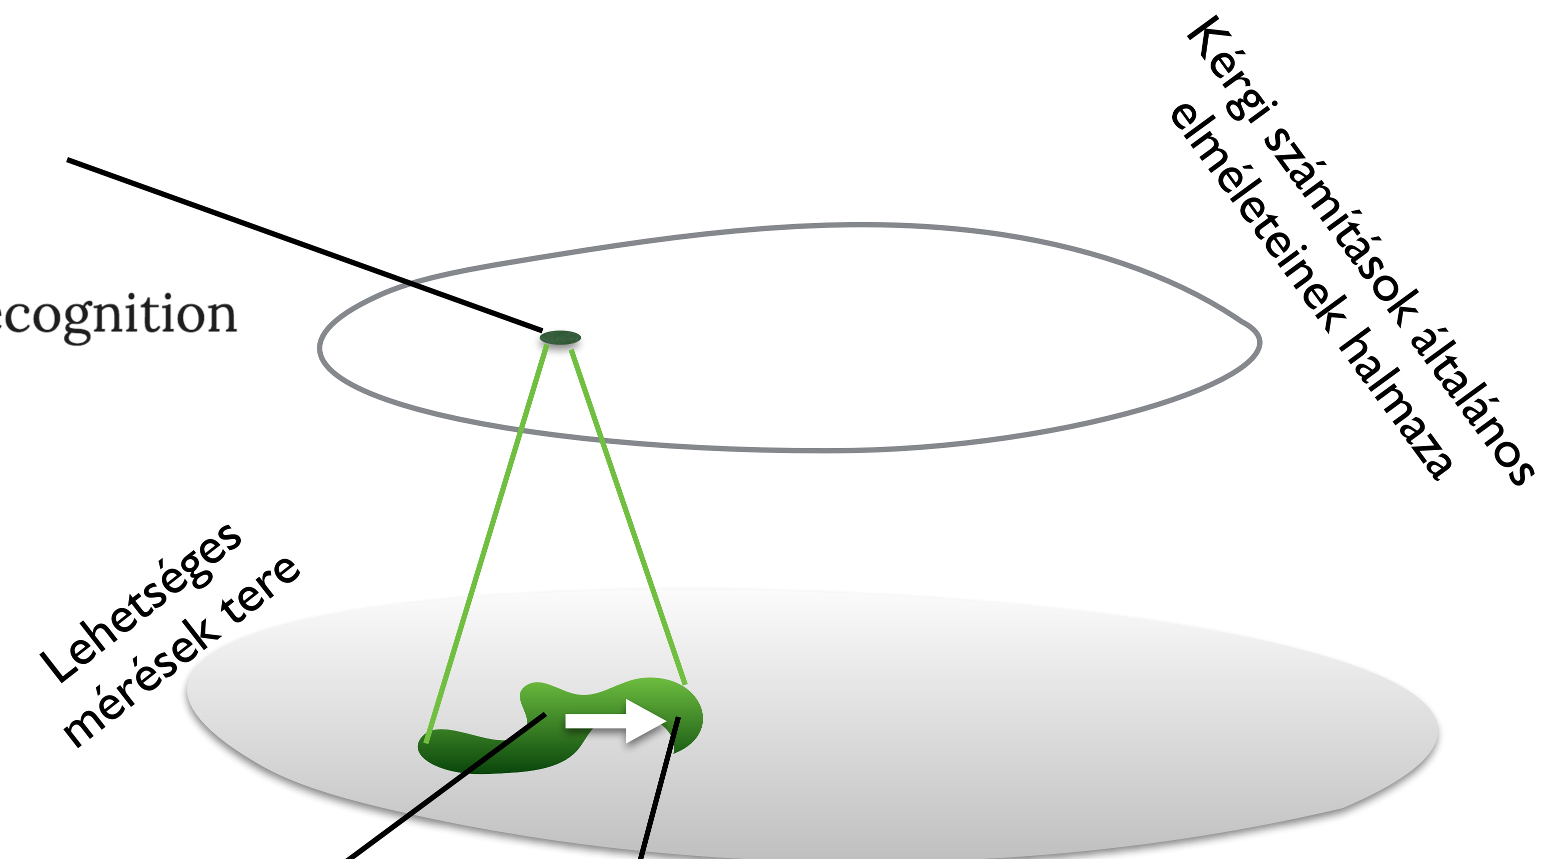

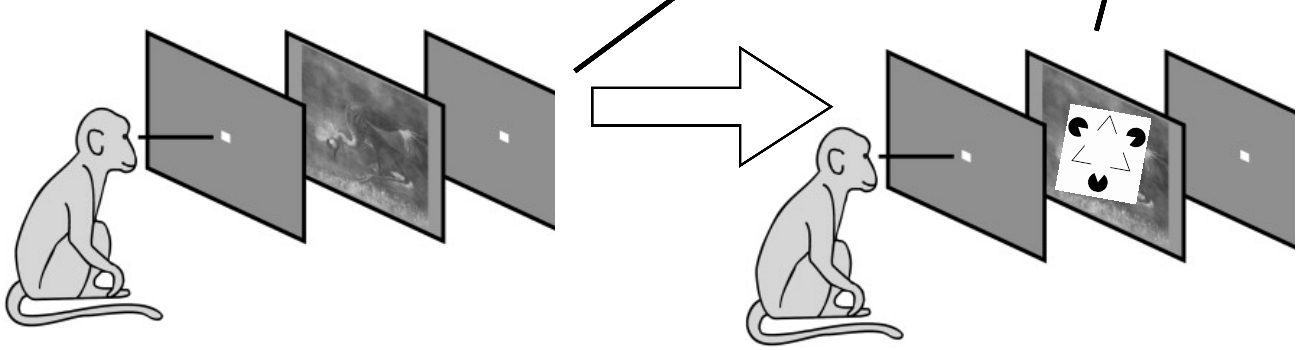

Principles vs. experiements

… implements probabilistic inference

It may be that brain..

…prepares the inputs for classification

Hierarchical models of object recognition in cortex

Maximilian Riesenhuber & Tomaso Poggio

Nature Neuroscience **2**, 1019–1025 (1999) | Dov

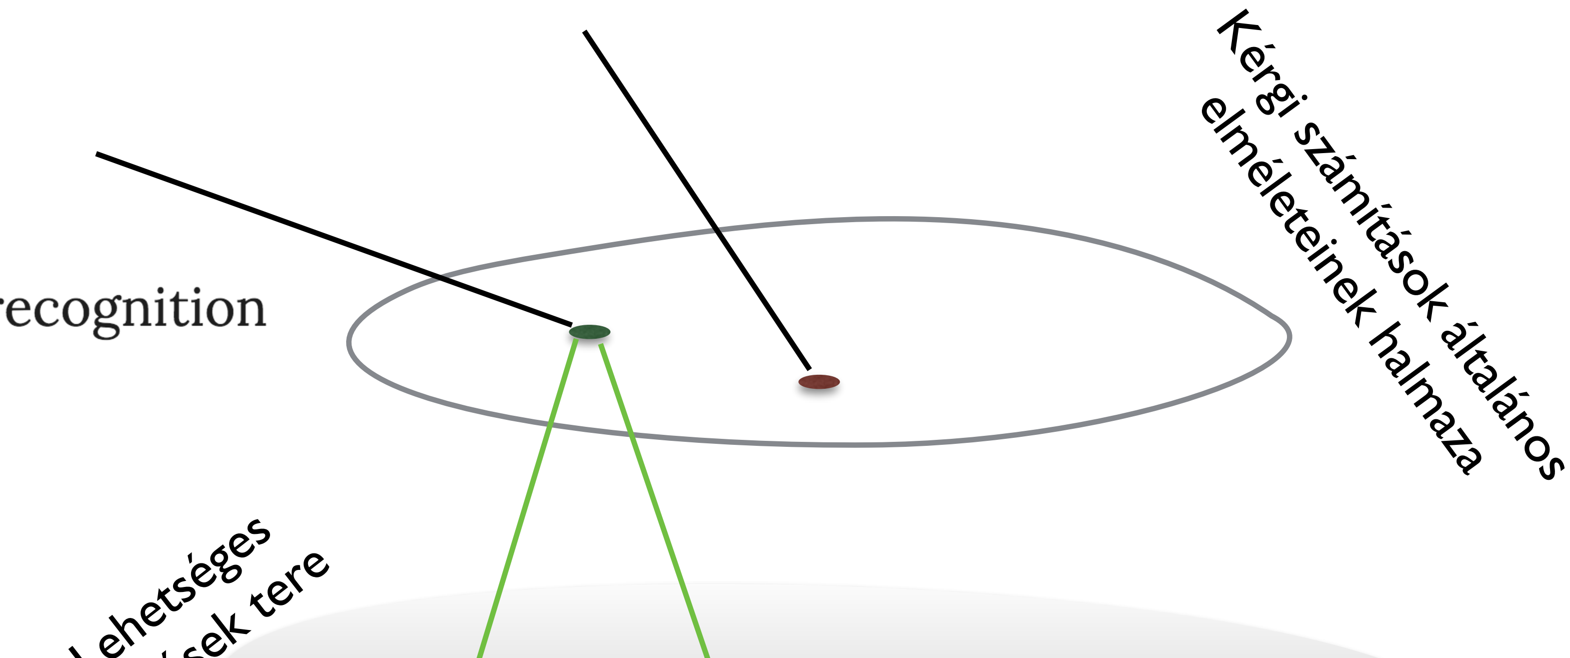
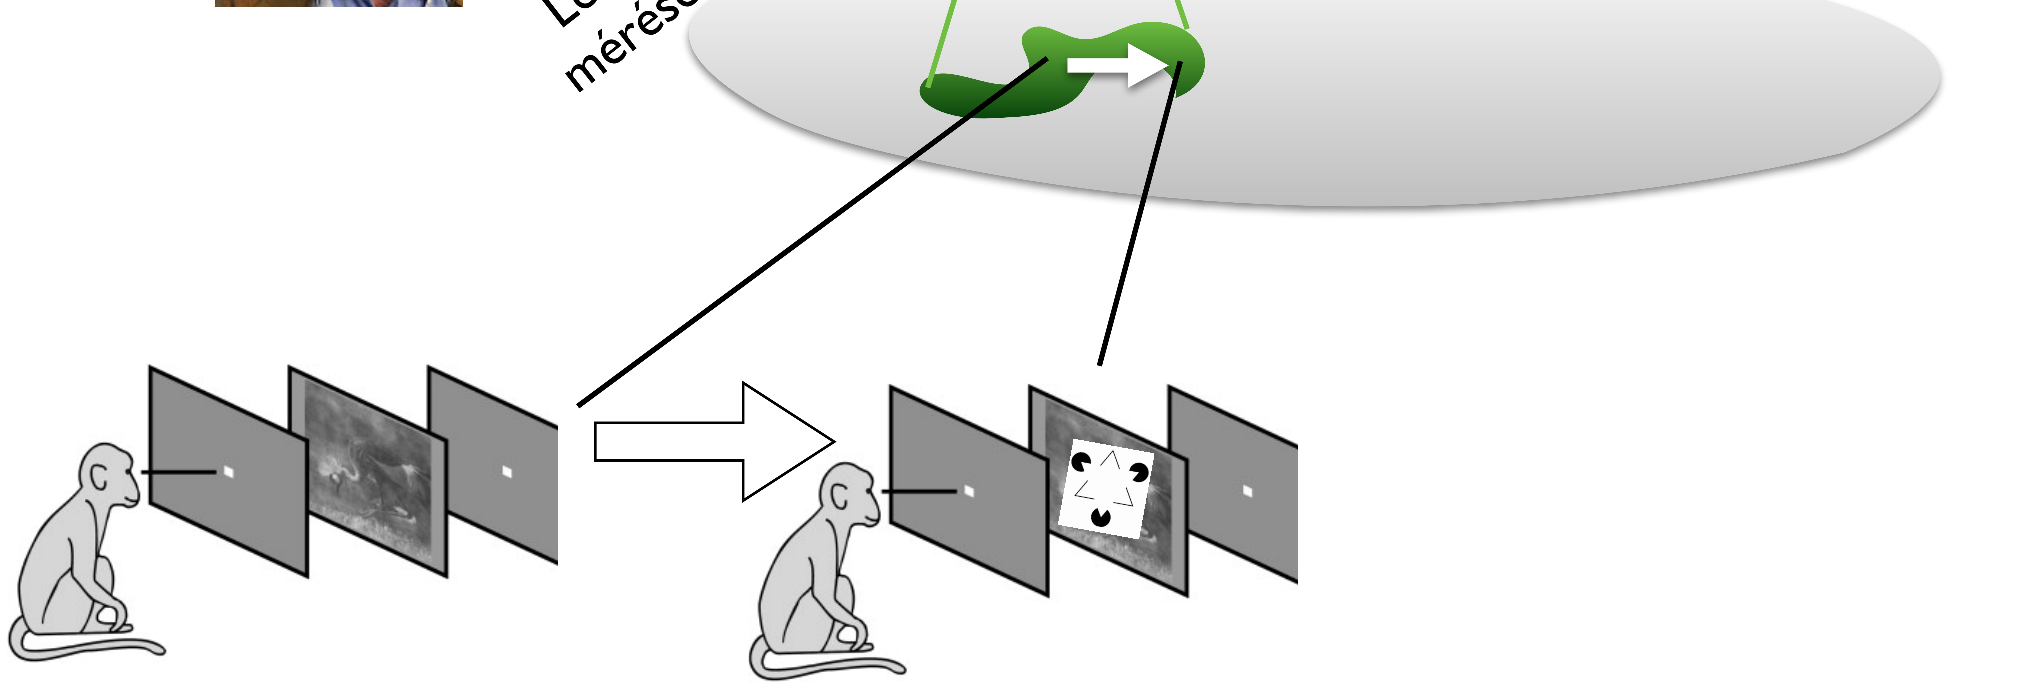

Principles vs. experiements

… implements probabilistic inference

Hierarchical Bayesian inference in the visual cortex

Tai Sing Lee and David Mumford

Journal of the Optical Society of America A

Vol. 20, Issue 7, pp. 1434-1448 (2003)

It may be that brain..

…prepares the inputs for classification

Hierarchical models of object recognition in cortex

Maximilian Riesenhuber & Tomaso Poggio

Nature Neuroscience **2**, 1019–1025 (1999) | Do

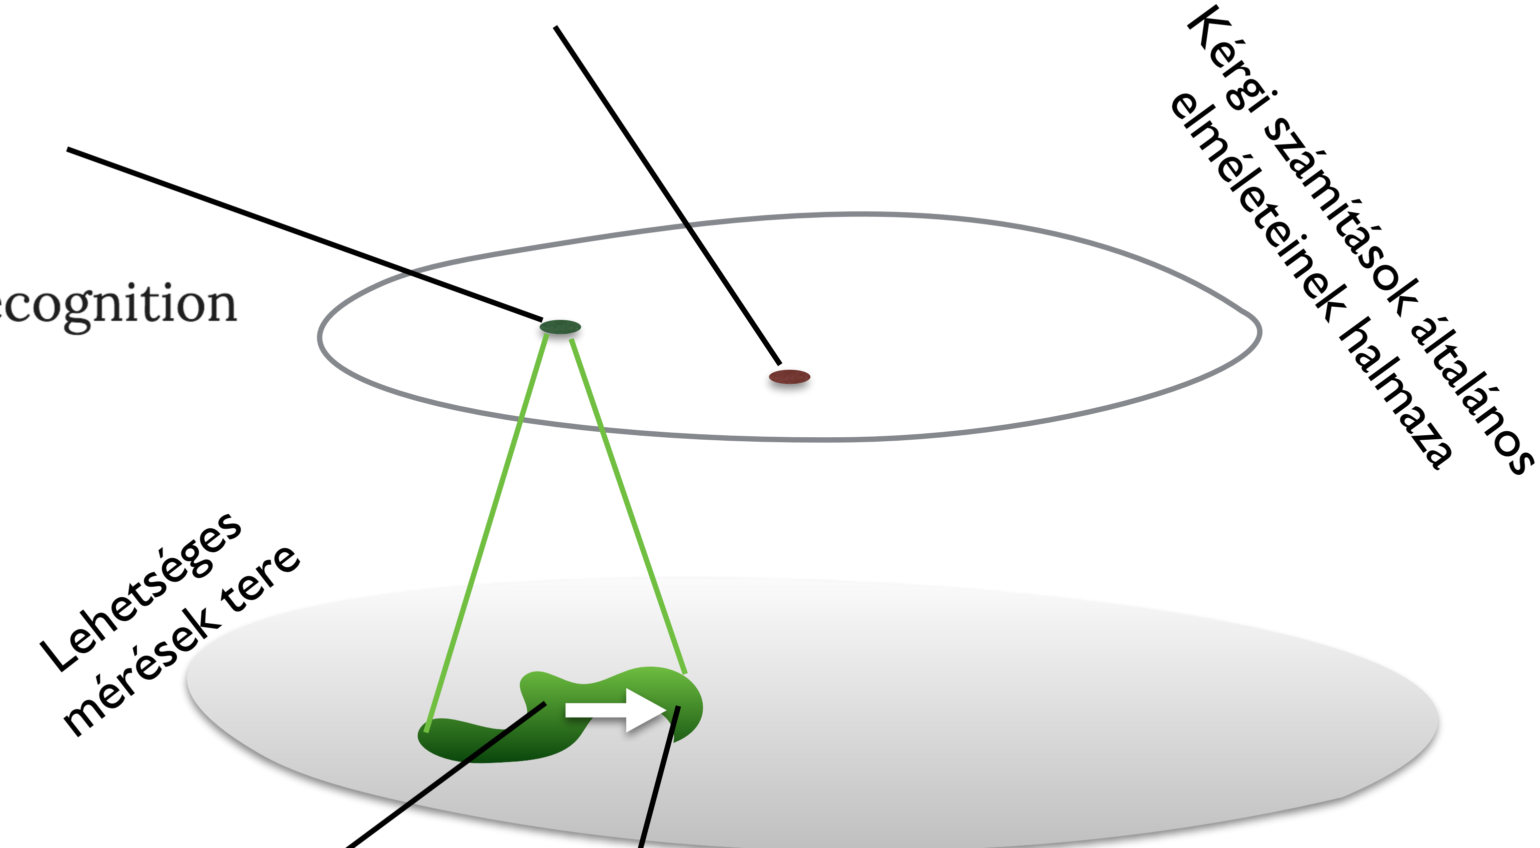

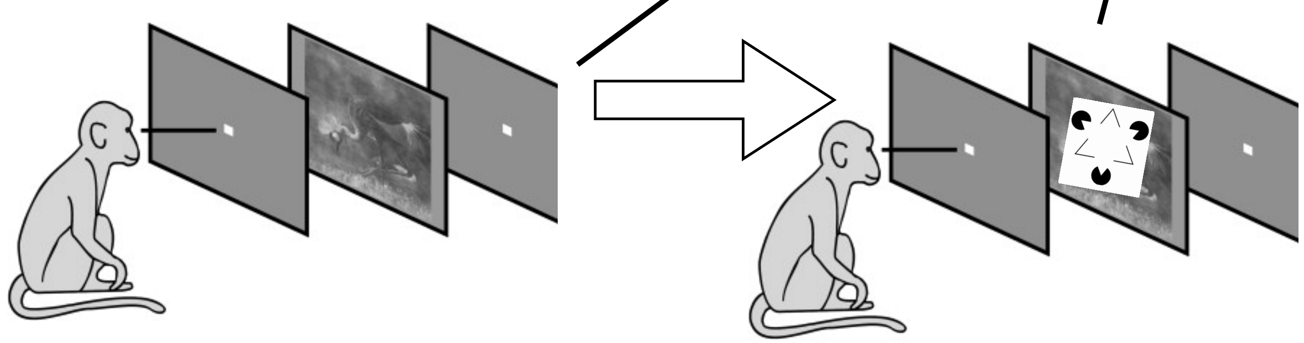

Principles vs. experiements

… implements probabilistic inference

Hierarchical Bayesian inference in the visual cortex

Tai Sing Lee and David Mumford

Journal of the Optical Society of America A

Vol. 20, Issue 7, pp. 1434-1448 (2003)

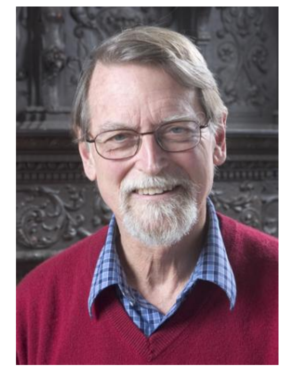

It may be that brain..

…prepares the inputs for classification

Hierarchical models of object recognition in cortex

Maximilian Riesenhuber & Tomaso Poggio

Nature Neuroscience **2**, 1019–1025 (1999) | Do

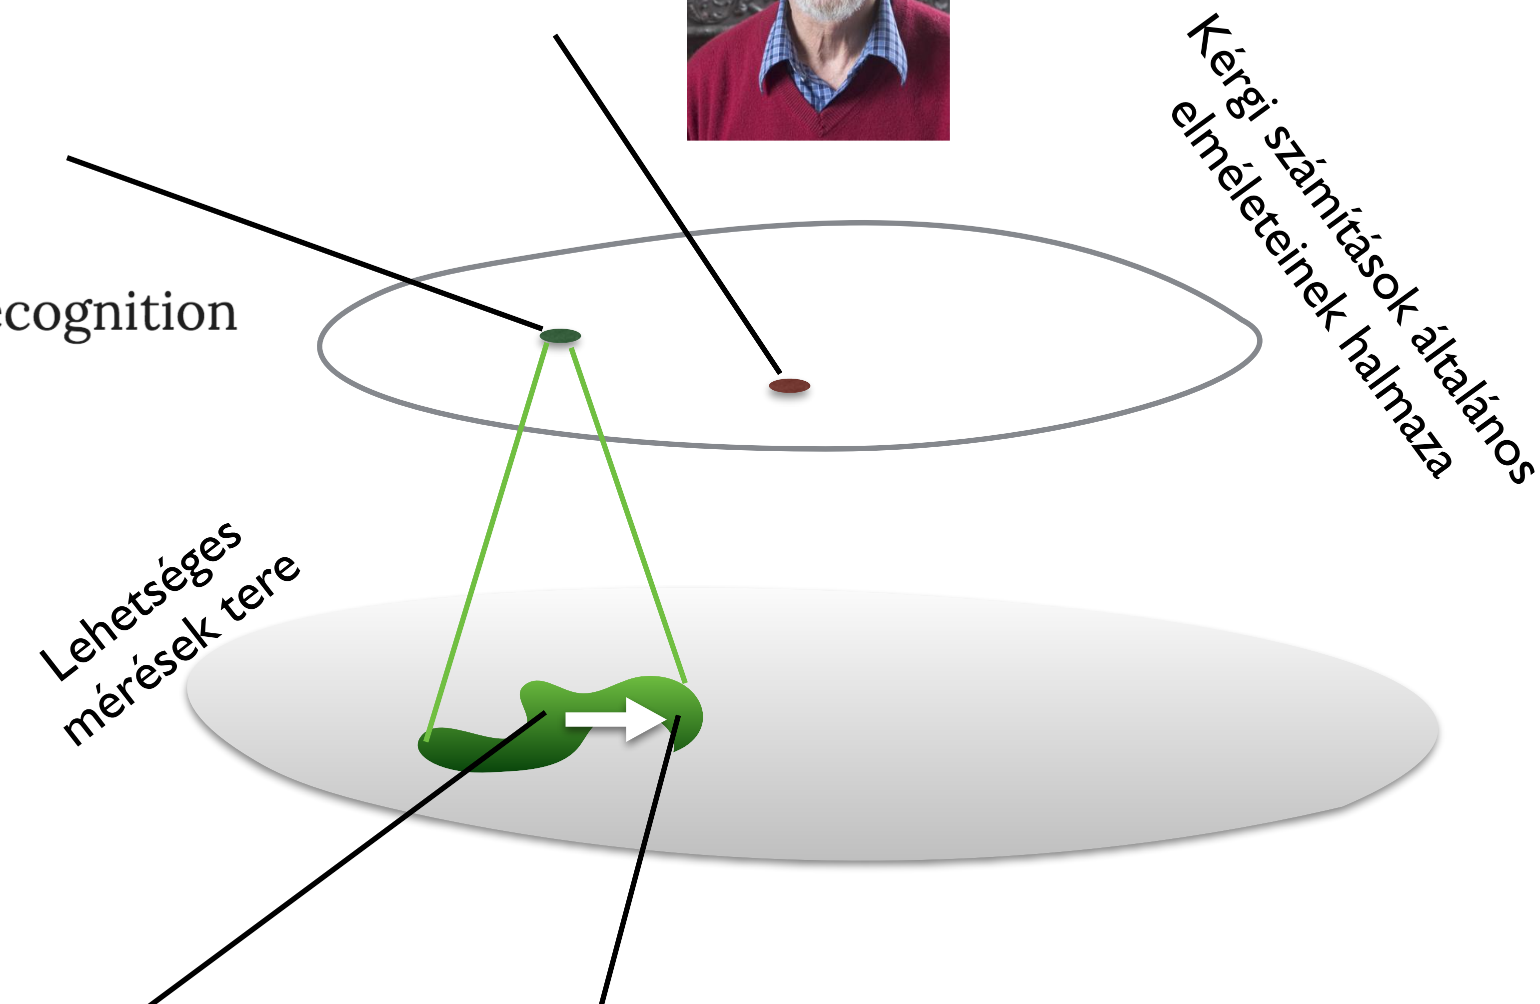

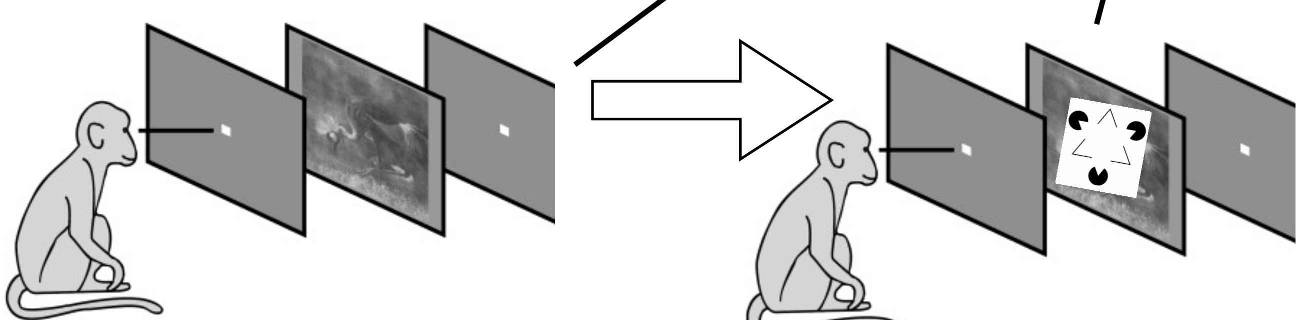

Principles vs. experiements

… implements probabilistic inference

Hierarchical Bayesian inference in the visual cortex

Tai Sing Lee and David Mumford

Journal of the Optical Society of America A

Vol. 20, Issue 7, pp. 1434-1448 (2003)

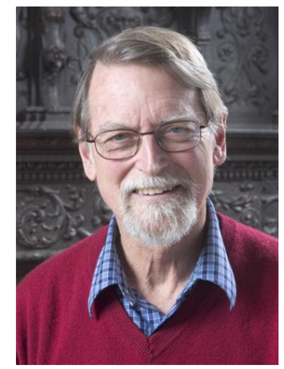

It may be that brain..

…prepares the inputs for classification

Hierarchical models of object recognition in cortex

Maximilian Riesenhuber & Tomaso Poggio

Nature Neuroscience **2**, 1019–1025 (1999) | Do

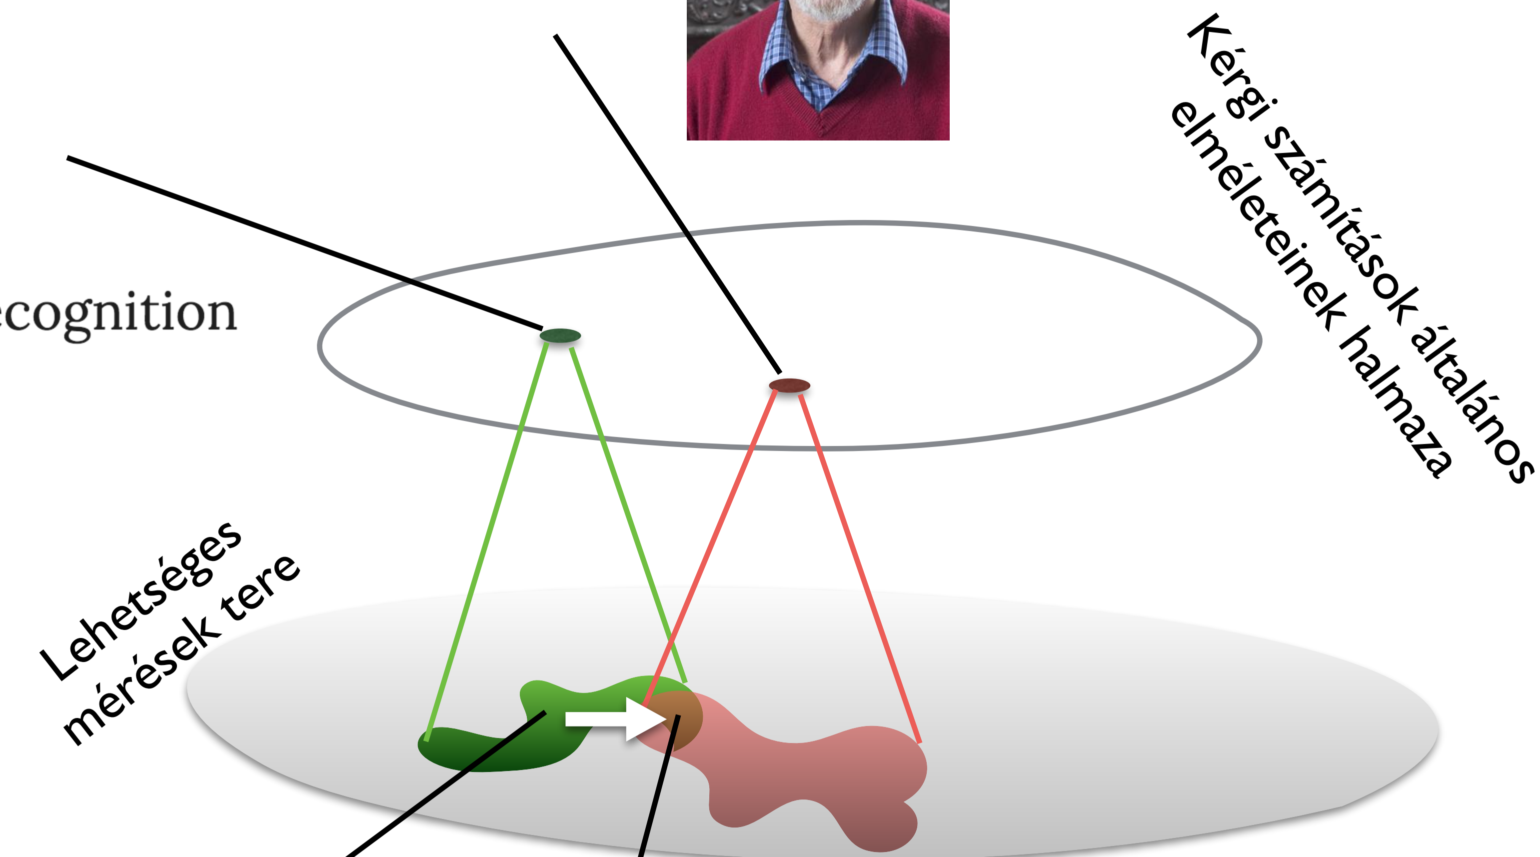

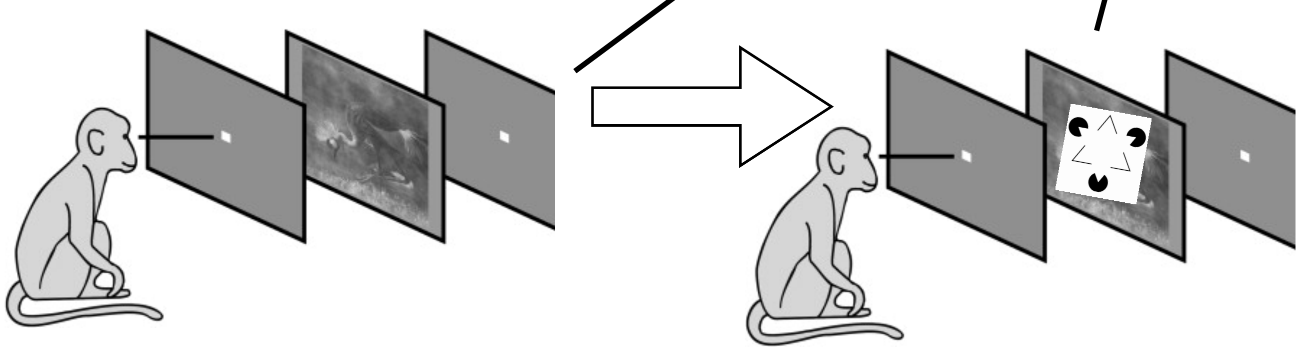

Principles vs. experiements

…maximizes the mutual information with future stimuli

… implements probabilistic inference

Hierarchical Bayesian inference in the visual cortex

Tai Sing Lee and David Mumford

Journal of the Optical Society of America A

Vol. 20, Issue 7, pp. 1434-1448 (2003)

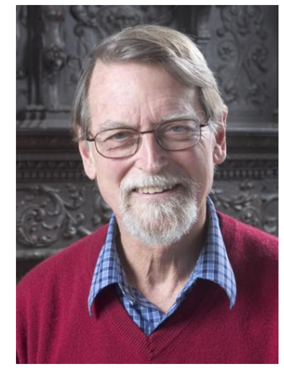

It may be that brain..

…prepares the inputs for classification

Hierarchical models of object recognition in cortex

Maximilian Riesenhuber & Tomaso Poggio

Nature Neuroscience **2**, 1019–1025 (1999) | Do

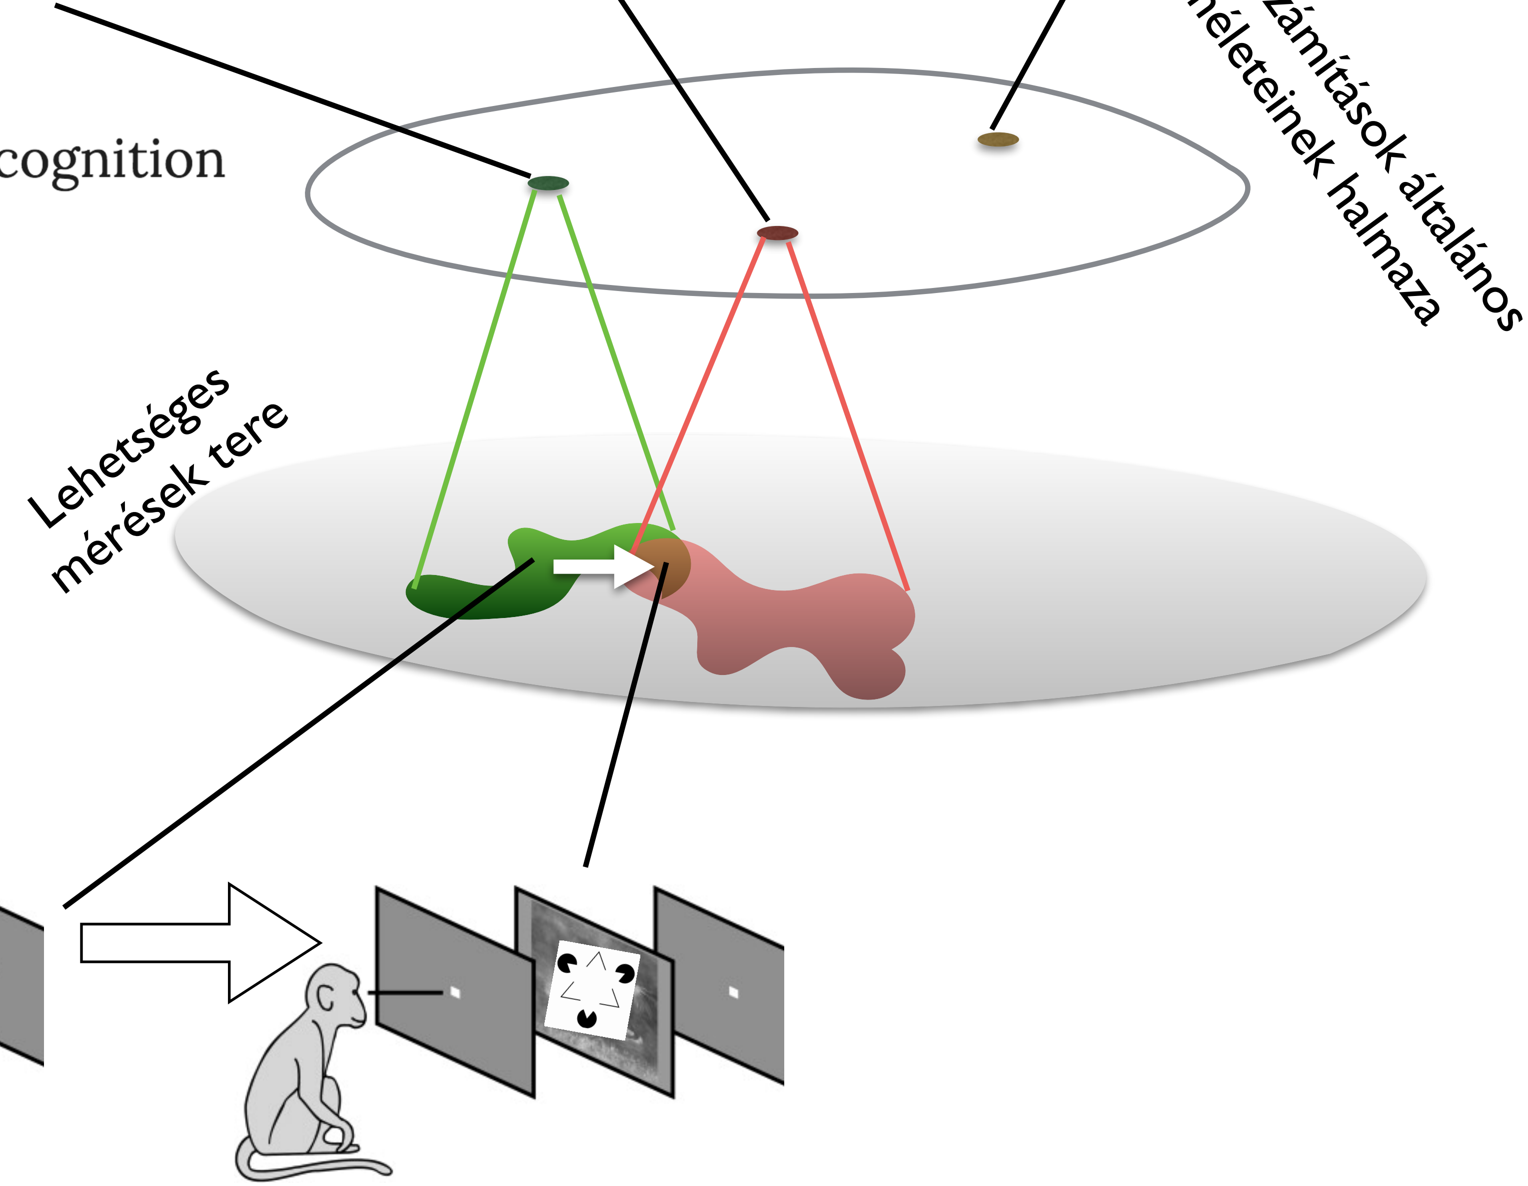

Principles vs. experiements

… implements probabilistic inference

Hierarchical Bayesian inference in the visual cortex

Tai Sing Lee and David Mumford

Journal of the Optical Society of America A

Vol. 20, Issue 7, pp. 1434-1448 (2003)

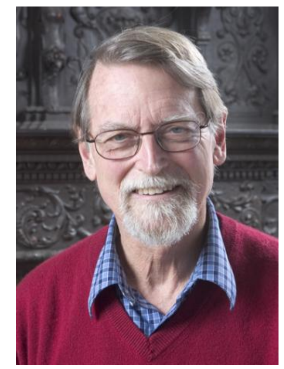

…maximizes the mutual information with future stimuli

Predictability, Complexity, and Learning

William Bialek, Ilya Nemenman and Naftali Tishby

Posted Online March 13, 2006

https://doi.org/10.1162/089976601753195969

© 2001 Massachusetts Institute of Technology

Neural Computation

Volume 13 | Issue 11 | November 2001

p.2409-2463

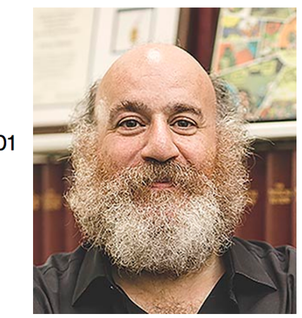

It may be that brain..

…prepares the inputs for classification

Hierarchical models of object recognition in cortex

Maximilian Riesenhuber & Tomaso Poggio

Nature Neuroscience **2**, 1019–1025 (1999) | Do

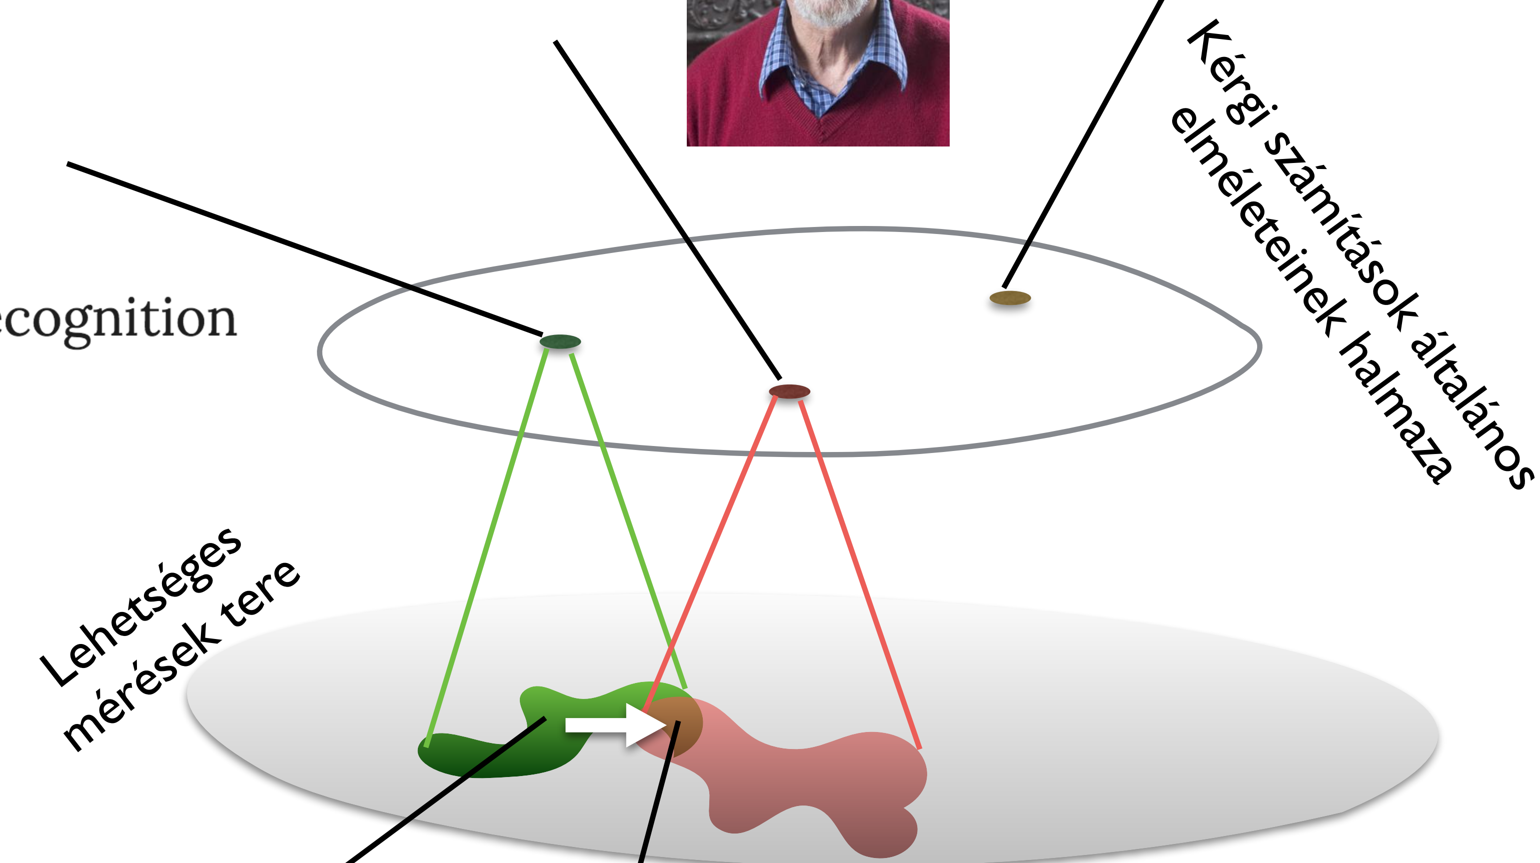

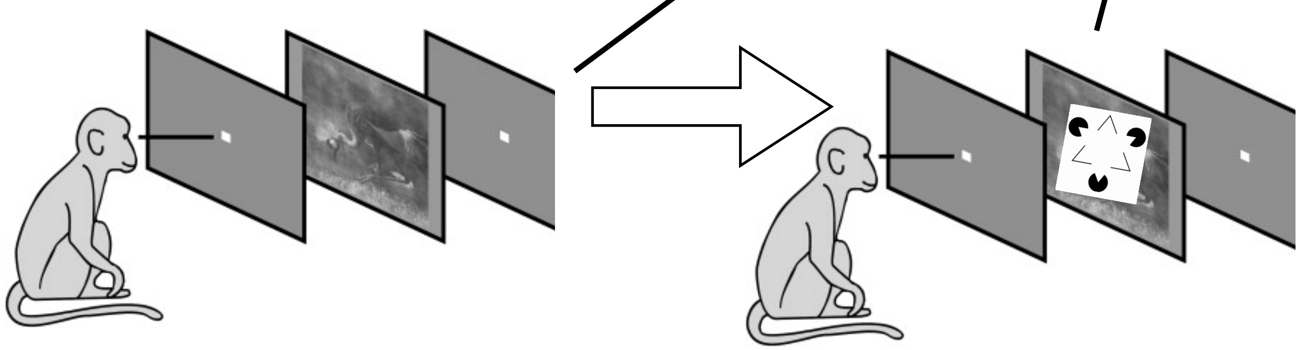

Principles vs. experiements

…maximizes the mutual information with future stimuli

… implements probabilistic inference

Predictability, Complexity, and Learning

William Bialek, Ilya Nemenman and Naftali Tishby

Posted Online March 13, 2006
https://doi.org/10.1162/089976601753195969

Neural Computation
Volume 13 | Issue 11 | November 2001
p.2409-2463

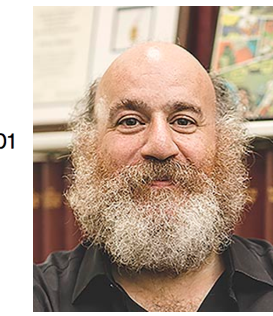

Hierarchical Bayesian inference in the visual cortex

Tai Sing Lee and David Mumford

Journal of the Optical Society of America A

Vol. 20, Issue 7, pp. 1434-1448 (2003)

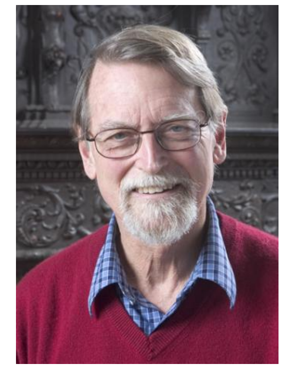

It may be that brain..

…prepares the inputs for classification

Hierarchical models of object recognition in cortex

Maximilian Riesenhuber & Tomaso Poggio

Nature Neuroscience **2**, 1019–1025 (1999) | Do

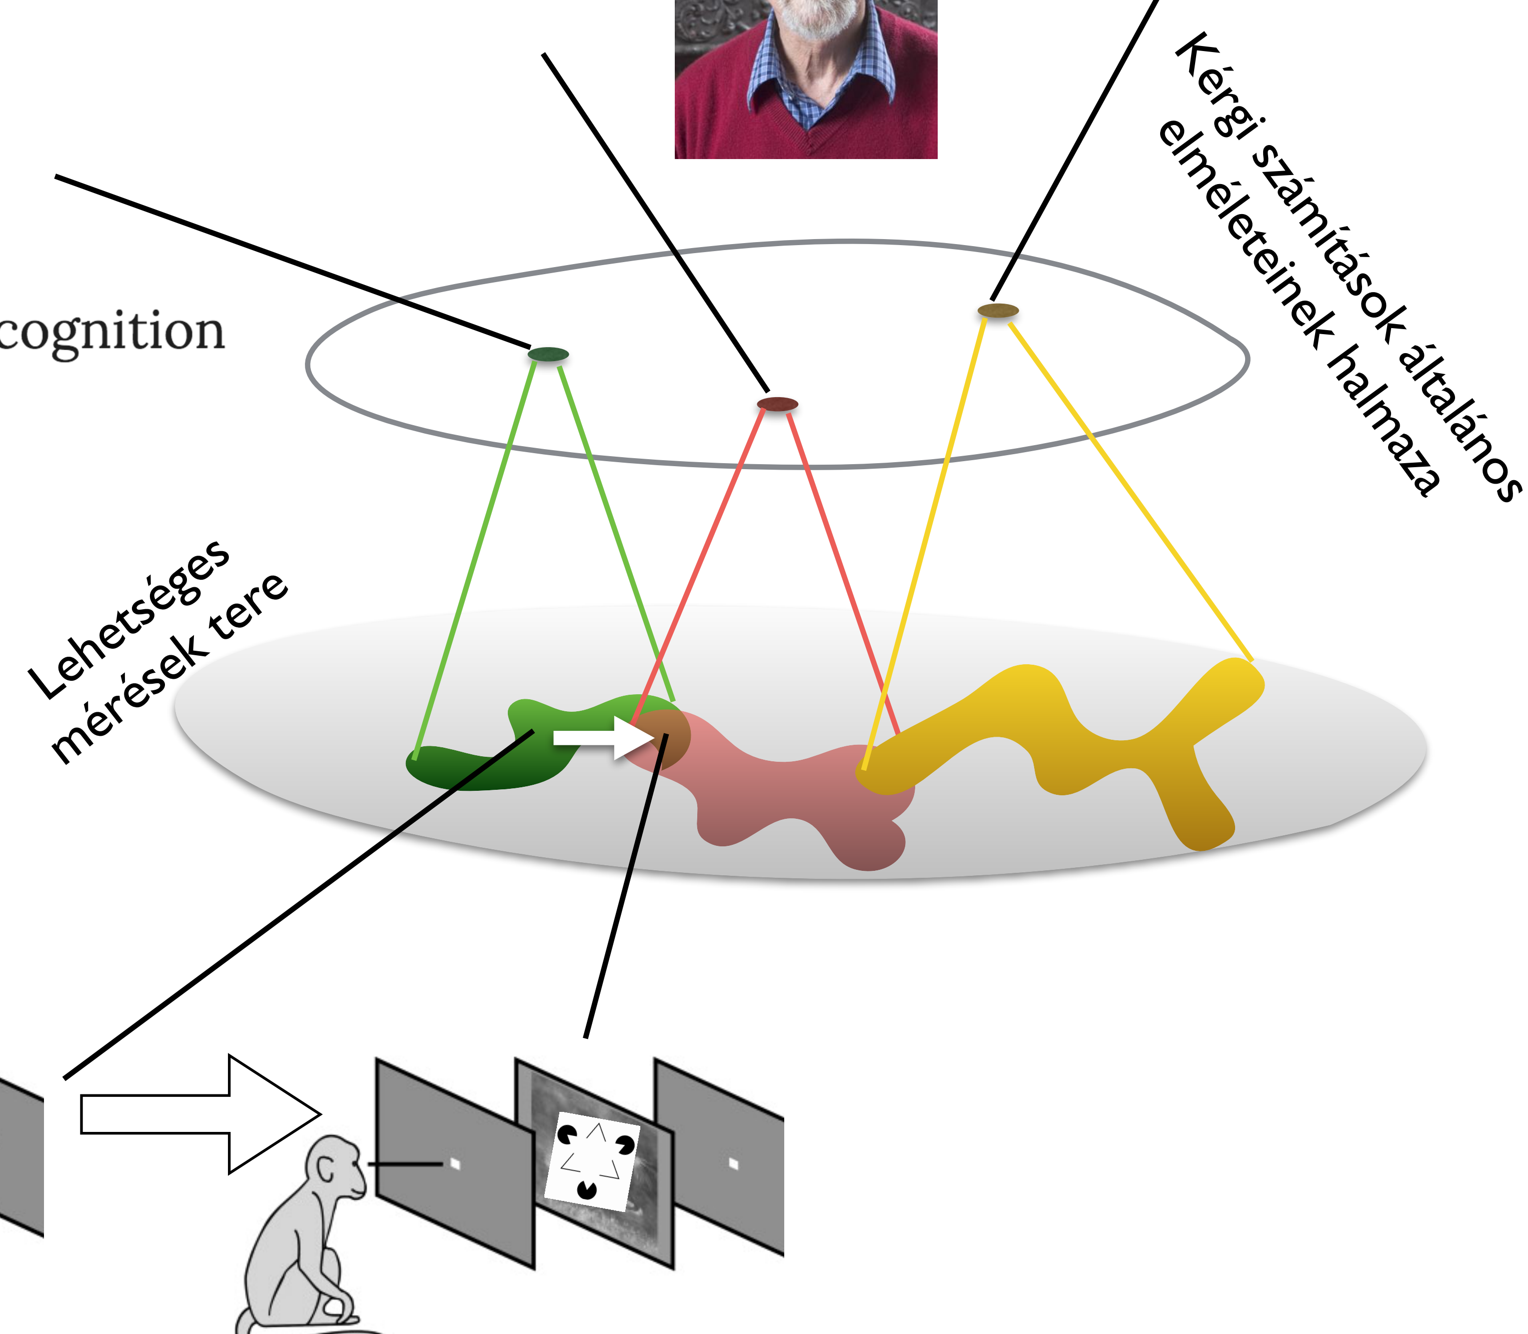

Principles vs. experiements

… implements probabilistic inference

Hierarchical Bayesian inference in the visual cortex

Tai Sing Lee and David Mumford

Journal of the Optical Society of America A

Vol. 20, Issue 7, pp. 1434-1448 (2003)

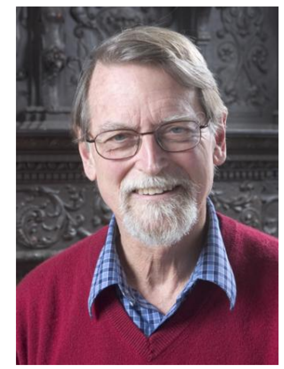

…maximizes the mutual information with future stimuli

Predictability, Complexity, and Learning

William Bialek, Ilya Nemenman and Naftali Tishby

Posted Online March 13, 2006

https://doi.org/10.1162/089976601753195969

© 2001 Massachusetts Institute of Technology

Neural Computation
Volume 13 | Issue 11 | November 2001
p.2409-2463

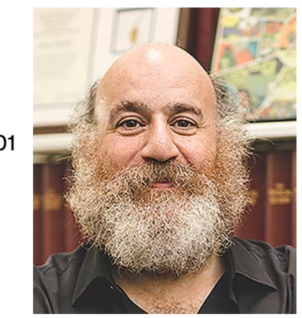

It may be that brain..

…prepares the inputs for classification

Hierarchical models of object recognition in cortex

Maximilian Riesenhuber & Tomaso Poggio

Nature Neuroscience **2**, 1019–1025 (1999) | Do

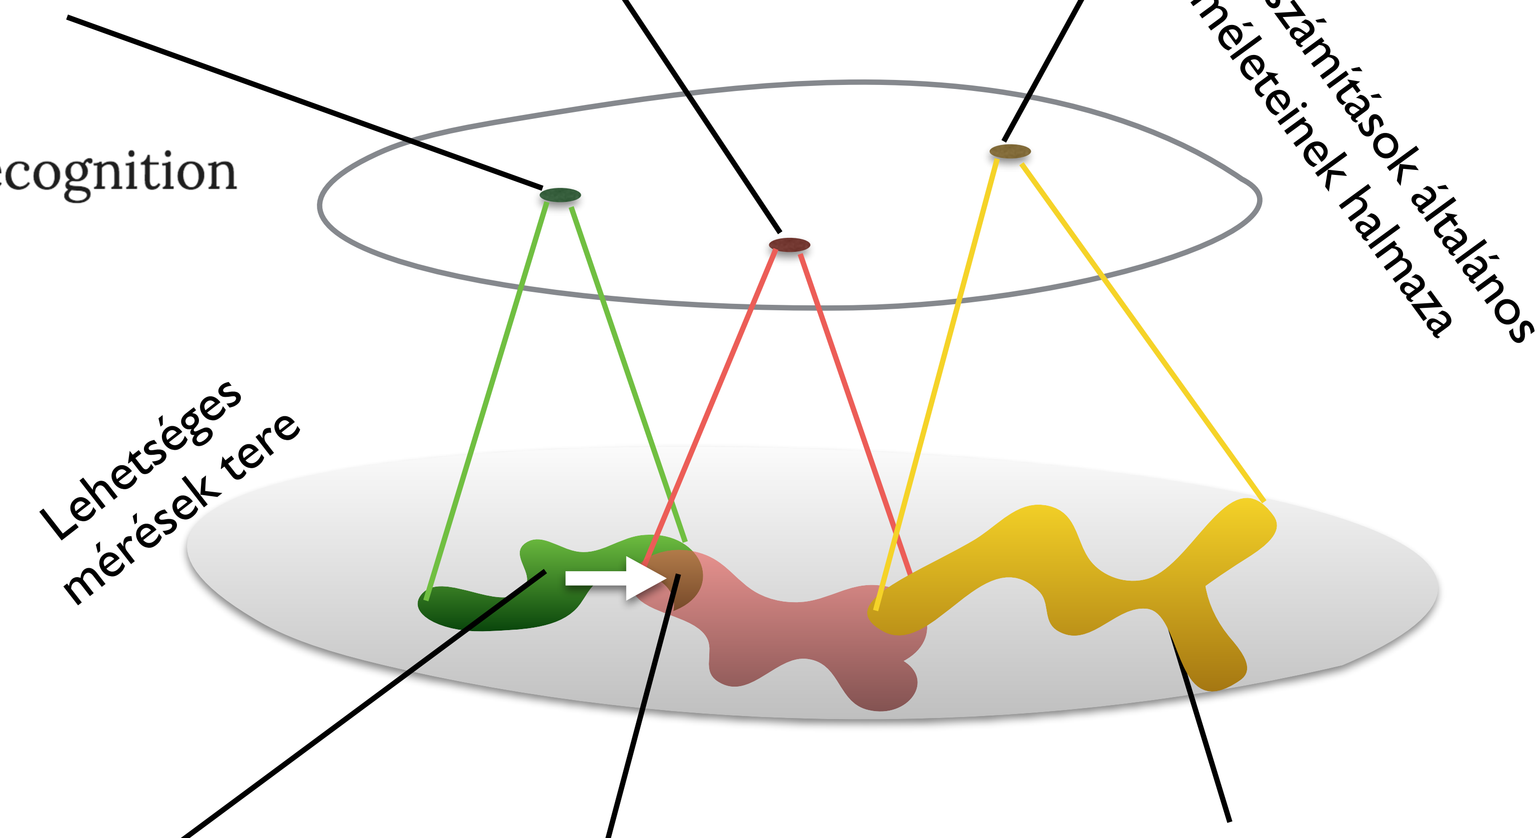

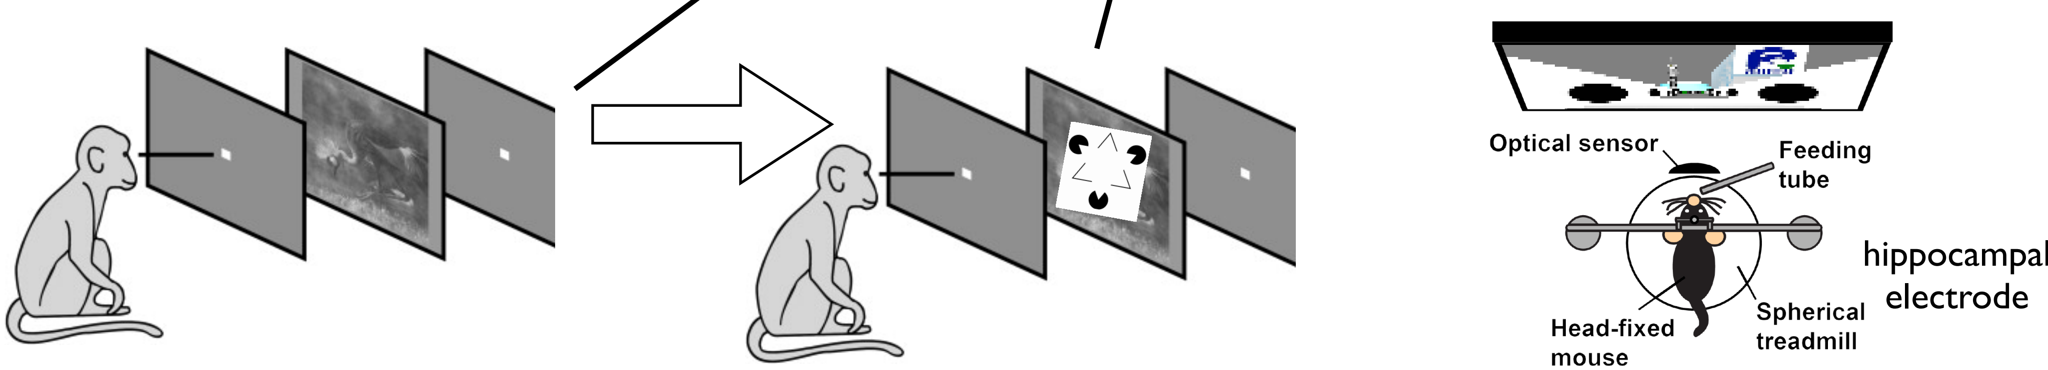

Principles vs. experiements

…maximizes the mutual information with future stimuli

… implements probabilistic inference

Predictability, Complexity, and Learning

William Bialek, Ilya Nemenman and Naftali Tishby

Posted Online March 13, 2006
https://doi.org/10.1162/089976601753195969

Neural Computation
Volume 13 | Issue 11 | November 2001
p.2409-2463

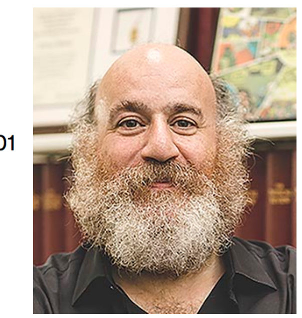

Hierarchical Bayesian inference in the visual cortex

Tai Sing Lee and David Mumford

Journal of the Optical Society of America A

Vol. 20, Issue 7, pp. 1434-1448 (2003)

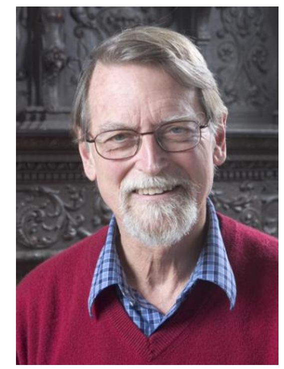

It may be that brain..

…prepares the inputs for classification

Hierarchical models of object recognition in cortex

Maximilian Riesenhuber & Tomaso Poggio

Nature Neuroscience 2, 1019–1025 (1999) | Do

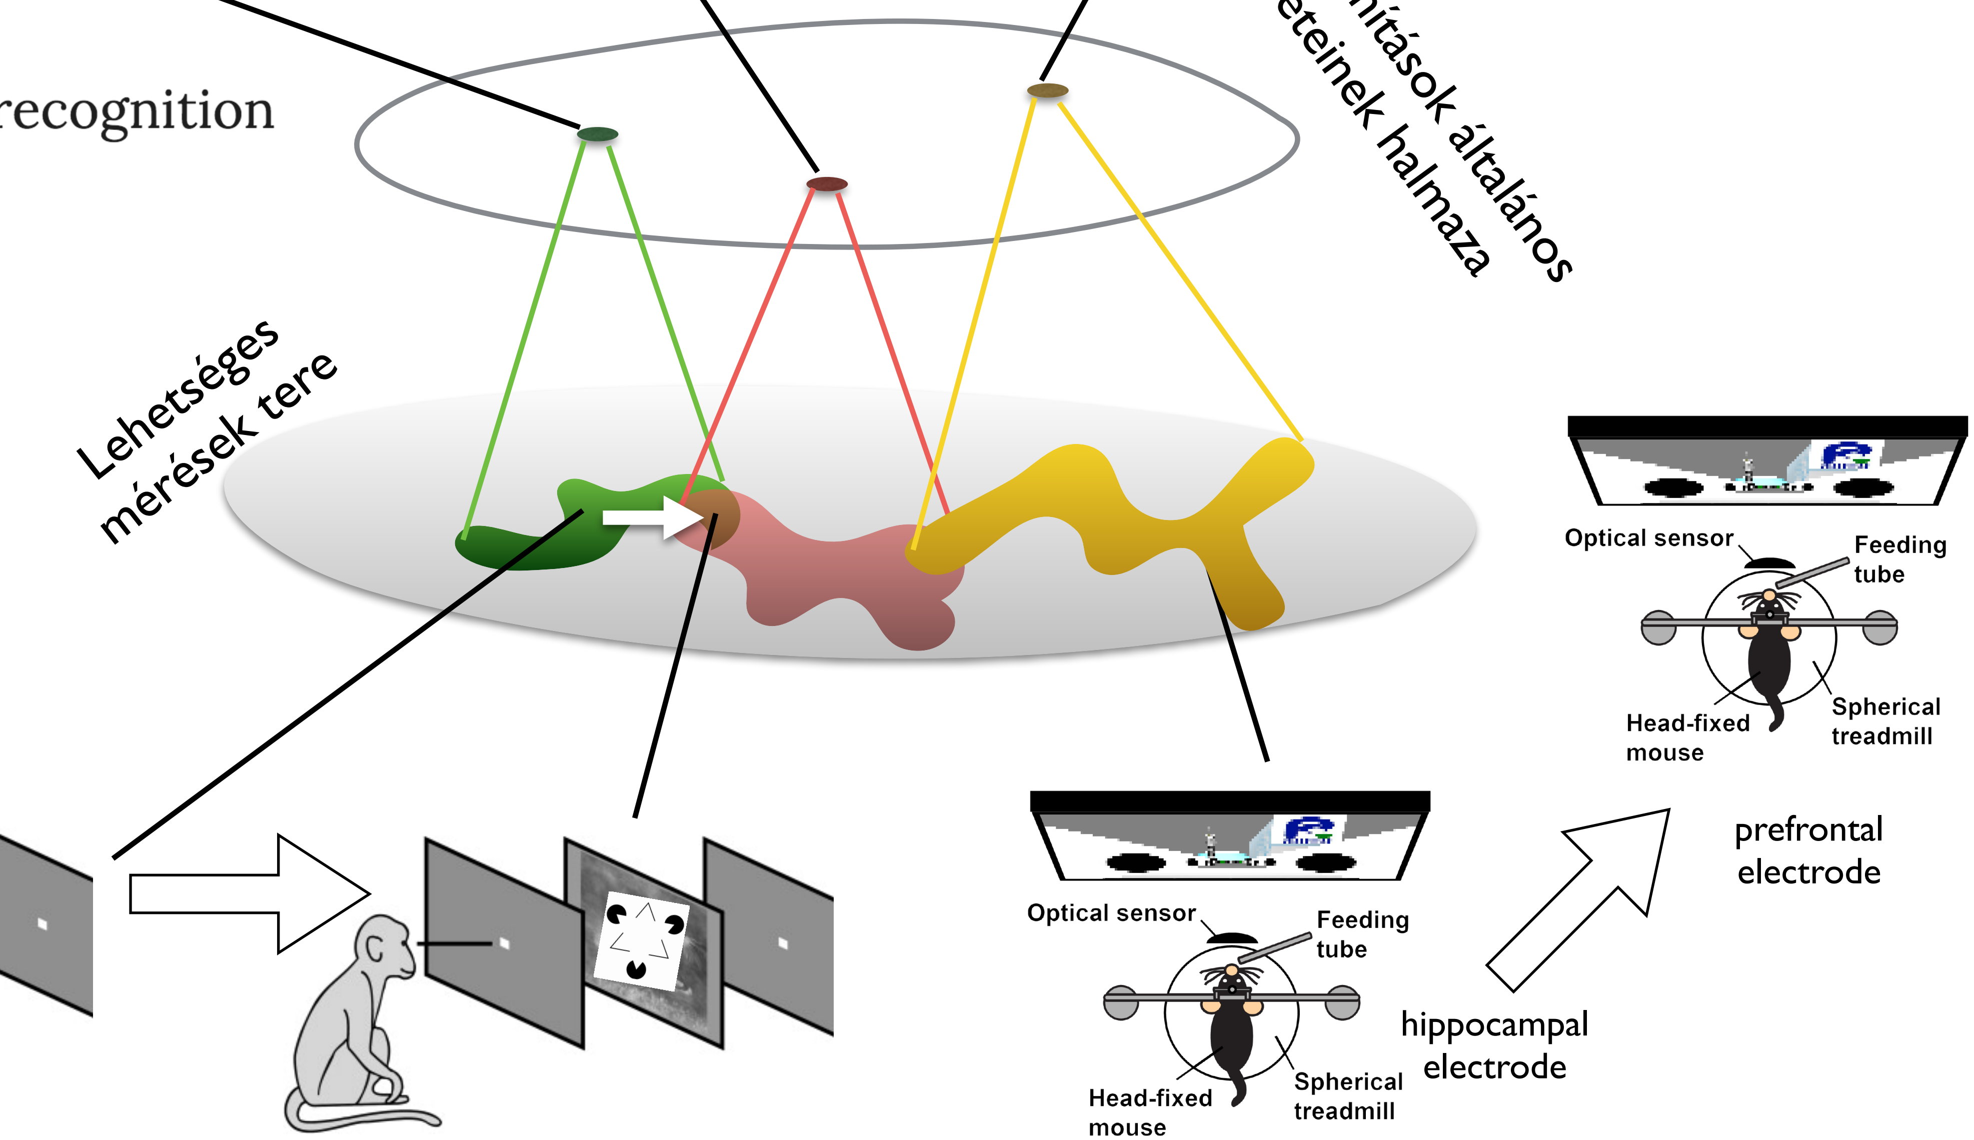

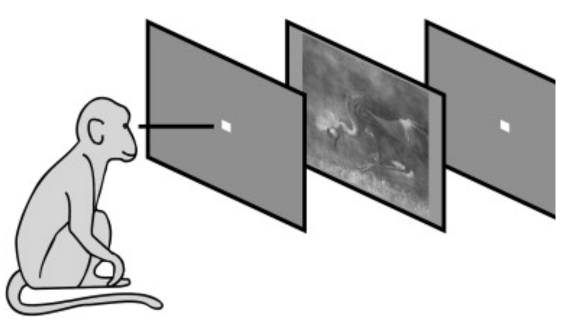

Principles vs. experiements

…maximizes the mutual information with future stimuli

… implements probabilistic inference

Predictability, Complexity, and Learning

William Bialek, Ilya Nemenman and Naftali Tishby

Posted Online March 13, 2006
https://doi.org/10.1162/089976601753195969

Neural Computation
Volume 13 | Issue 11 | November 2001
p.2409-2463

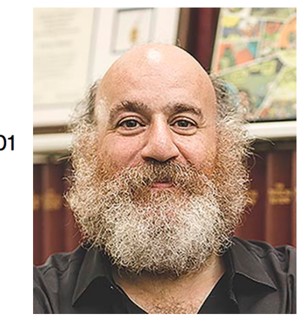

Hierarchical Bayesian inference in the visual cortex

Tai Sing Lee and David Mumford

Journal of the Optical Society of America A

Vol. 20, Issue 7, pp. 1434-1448 (2003)

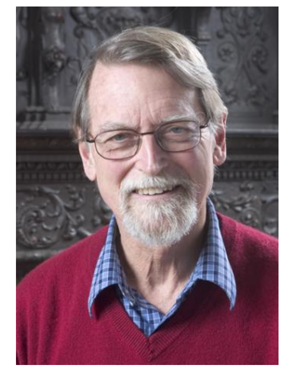

It may be that brain..

…prepares the inputs for classification

Hierarchical models of object recognition in cortex

Maximilian Riesenhuber & Tomaso Poggio

Nature Neuroscience **2**, 1019–1025 (1999) | Do

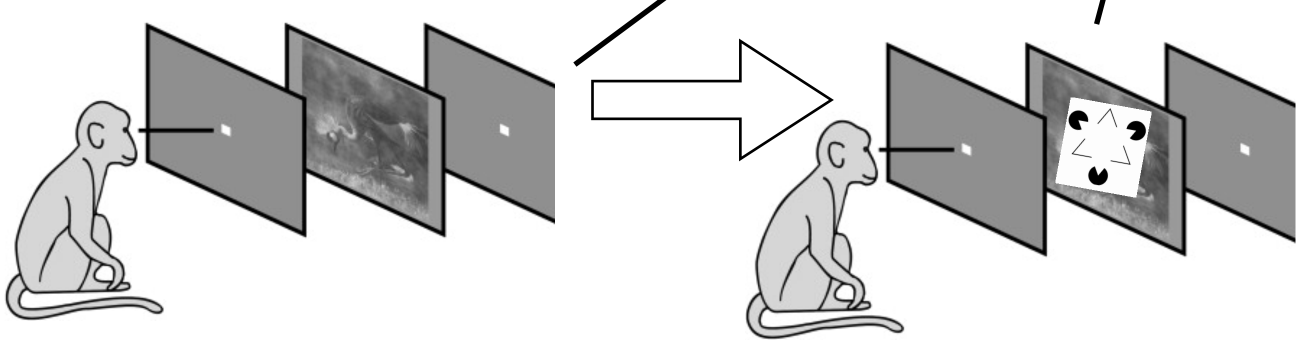

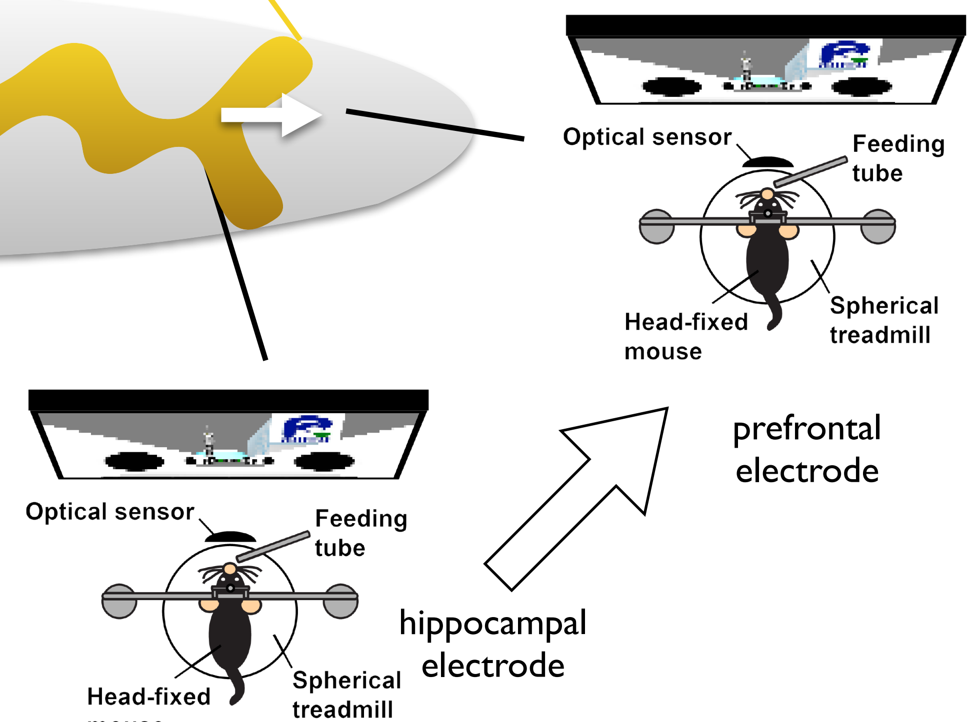

Levels of abstraction

- **Computation** - determining the challenge:
  through the goals of the biological agent phrasing a mathematical model that is capable of addressing the challenge

- **Algorithm** - solving the challenge can be achieved in many different ways that can be weighed based on different factors

- **Implementation** - physical realization of the algorithm, in the case of neuroscience, based on neurons, spikes, etc

Levels of abstraction

- **Computation** - determining the challenge:
  through the goals of the biological agent phrasing a mathematical model that is capable of addressing the challenge

- **Algorithm** - solving the challenge can be achieved in many different ways that can be weighed based on different factors

- **Implementation** - physical realization of the algorithm, in the case of neuroscience, based on neurons, spikes, etc

What is the goal of the brain?

What is the goal of the brain?

- Muscle control
  - for this, we need to make desicions
  - this is simple if we are satisfied with reflexes
  - otherwise, we need to remember things — we need to construct the representation of the knowledge we acquired in the neural tissue
  - based on this knowledge we interpret our sensnory data — **inference**
  - we need to constantly update the knowledge base based on novel information — **learning**

What is the goal of the brain?

zsákállatok

- Muscle control
  - for this, we need to make desicions
  - this is simple if we are satisfied with reflexes
  - otherwise, we need to remember things — we need to construct the representation of the knowledge we acquired in the neural tissue
  - based on this knowledge we interpret our sensnory data — **inference**
  - we need to constantly update the knowledge base based on novel information — **learning**

What is the goal of the brain?

zsákállatok

- Muscle control
  - for this, we need to make desicions
  - this is simple if we are satisfied with reflexes
  - otherwise, we need to remember things — we need to construct the representation of the knowledge we acquired in the neural tissue
  - based on this knowledge we interpret our sensnory data — **inference**
  - we need to constantly update the knowledge base based on novel information — **learning**
- The brain uses the acquired knowledge to predict the outcome of decisions — building a **model** of the environment - (well, we are modelling processes that model some other processes)
  - in neuroscience we can use models for analyzing data, or we can speak of models that the brain builds — not to be confused

What is the goal of the brain?


zsákállatok

- Muscle control
  - for this, we need to make desicions
  - this is simple if we are satisfied with reflexes
  - otherwise, we need to remember things — we need to construct the representation of the knowledge we acquired in the neural tissue
  - based on this knowledge we interpret our sensnory data — **inference**
  - we need to constantly update the knowledge base based on novel information — **learning**
- The brain uses the acquired knowledge to predict the outcome of decisions — building a **model** of the environment - (well, we are modelling processes that model some other processes)
  - in neuroscience we can use models for analyzing data, or we can speak of models that the brain builds — not to be confused

What is the goal of the visual system?

- To gain relevant information (objects) from a huge data set (wavelength of incoming light)

What is the goal of the visual system?

- To gain relevant information (objects) from a huge data set (wavelength of incoming light)

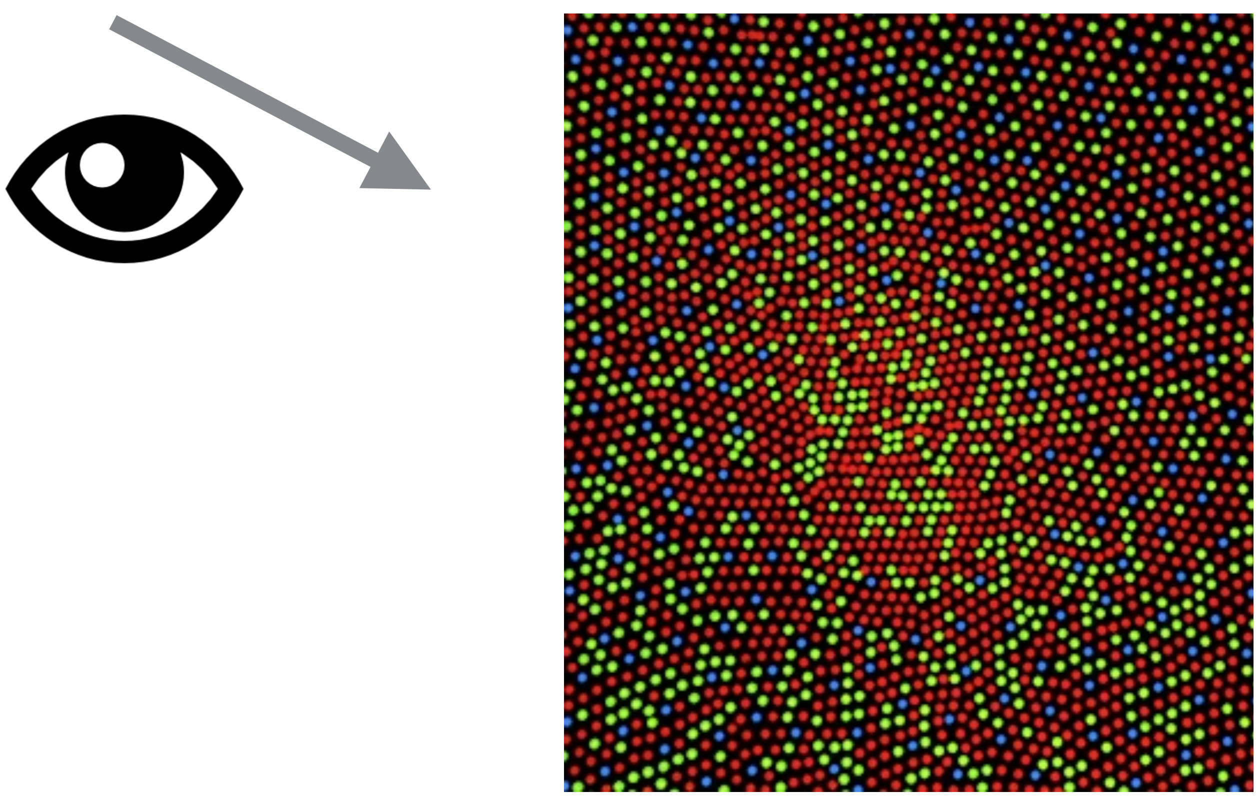

What is the goal of the visual system?

- To gain relevant information (objects) from a huge data set (wavelength of incoming light)

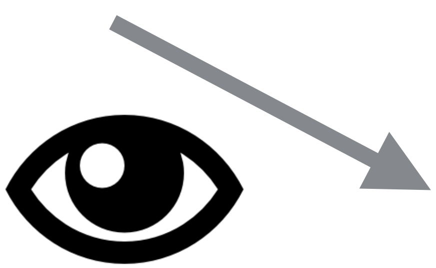

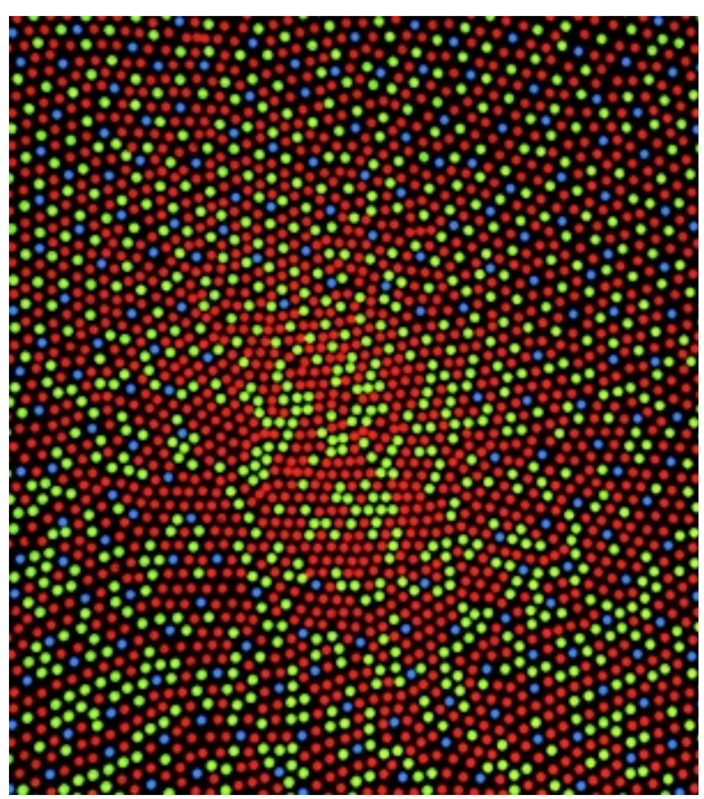

What is the goal of the visual system?

- To gain relevant information (objects) from a huge data set (wavelength of incoming light)
- But what are the useful objects and how do we know how to recognize them?

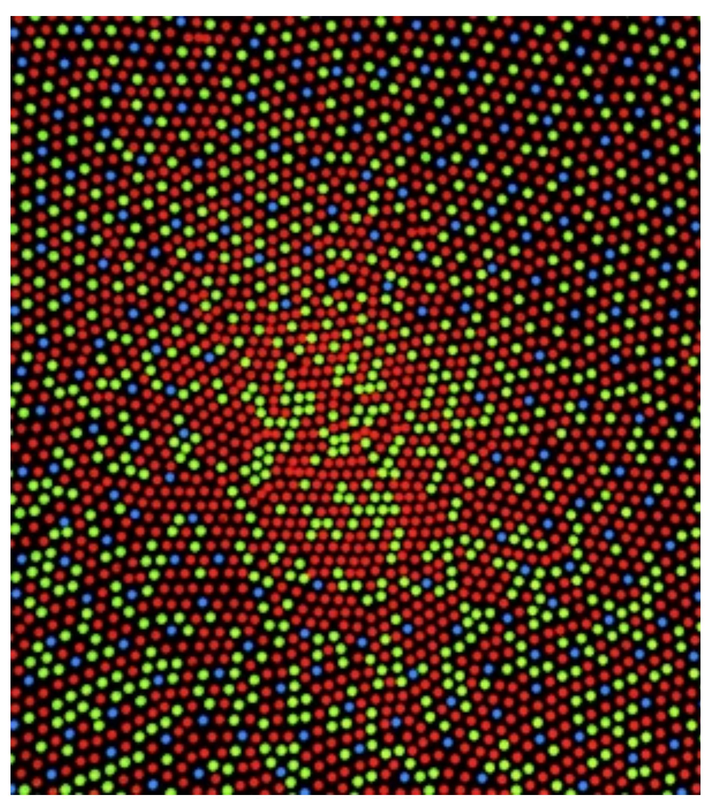

Ambiguity of observations

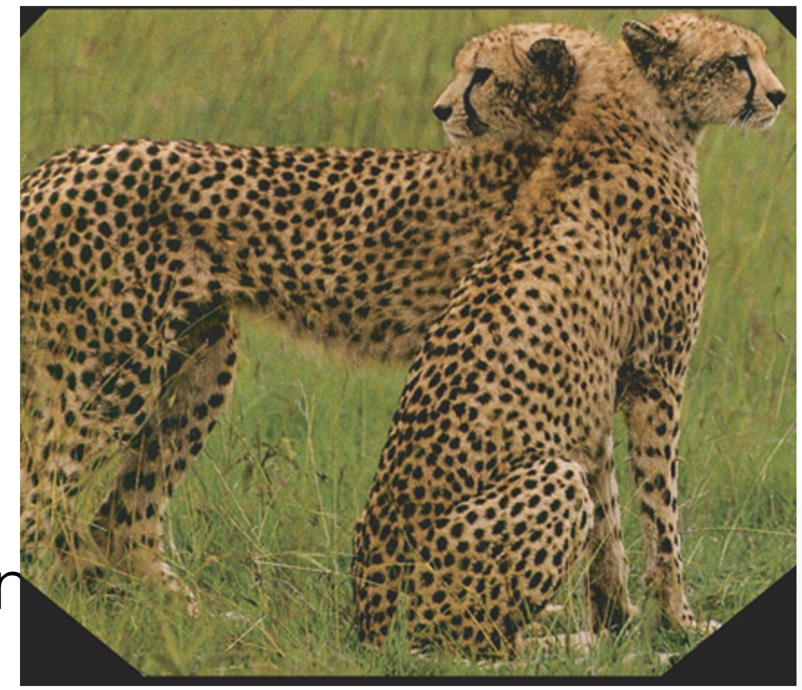

- An inherent property of the the environment is that multiple interpretations are compatible with observations
    - this is the rule, not the exception
- Sensors are noisy (e.g. cones) but the main source of uncertainty is limited data (insufficient observations)
- As a consequence, perception is essentially a challenge of inference: based on observations we reconstruct the state of the environment

Ernst & Bülthoff, Trends Cogn Sci (2004)

Ambiguity of observations

- An inherent property of the the environment is that multiple interpretations are compatible with observations
  - this is the rule, not the exception

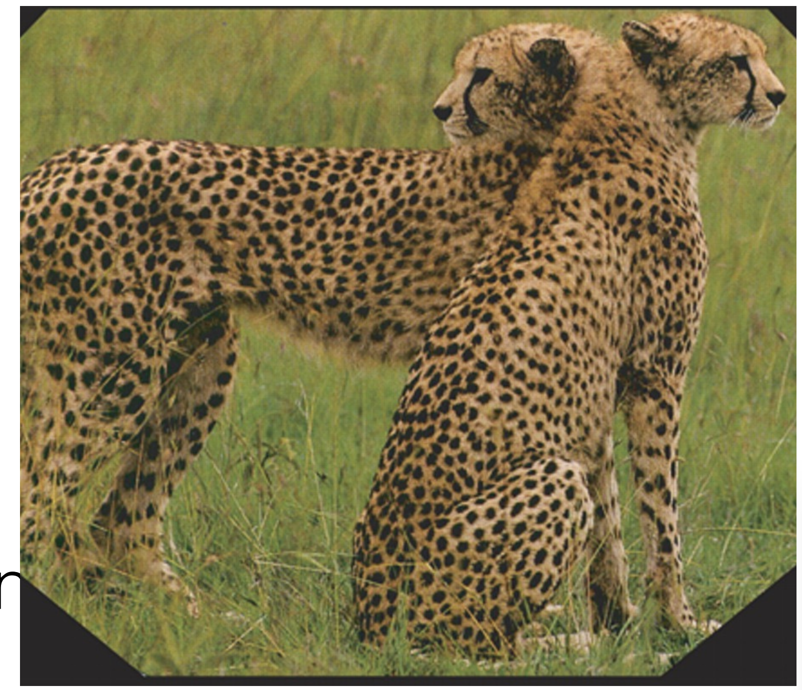

- Sensors are noisy (e.g. cones) but the main source of uncertainty is limited data (insufficient observations)
- As a consequence, perception is essentially a challenge of inference: based on observations we reconstruct the state of the environment

Ernst & Bülthoff, Trends Cogn Sci (2004)

Ambiguity of observations

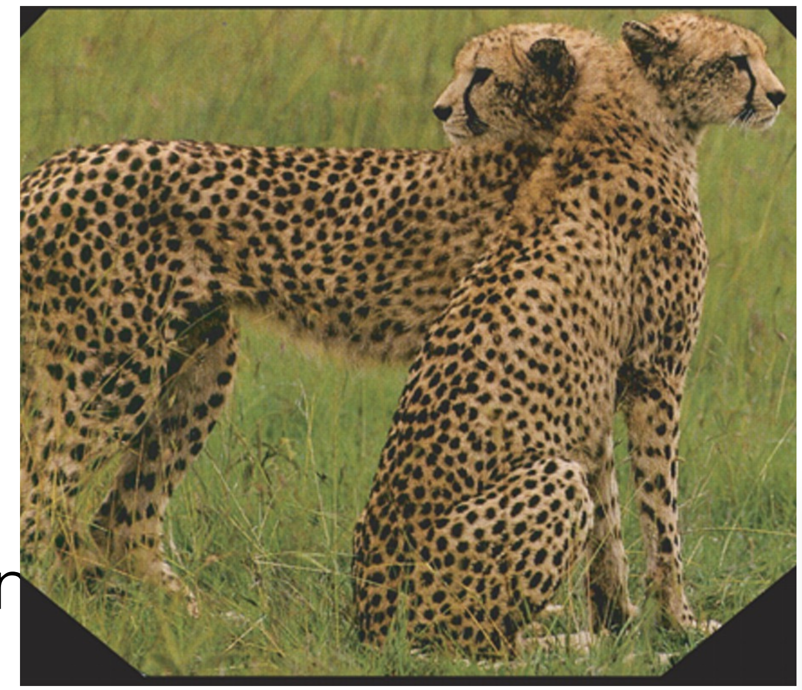

- An inherent property of the the environment is that multiple interpretations are compatible with observations
    - this is the rule, not the exception
- Sensors are noisy (e.g. cones) but the main source of uncertainty is limited data (insufficient observations)
- As a consequence, perception is essentially a challenge of inference: based on observations we reconstruct the state of the environment

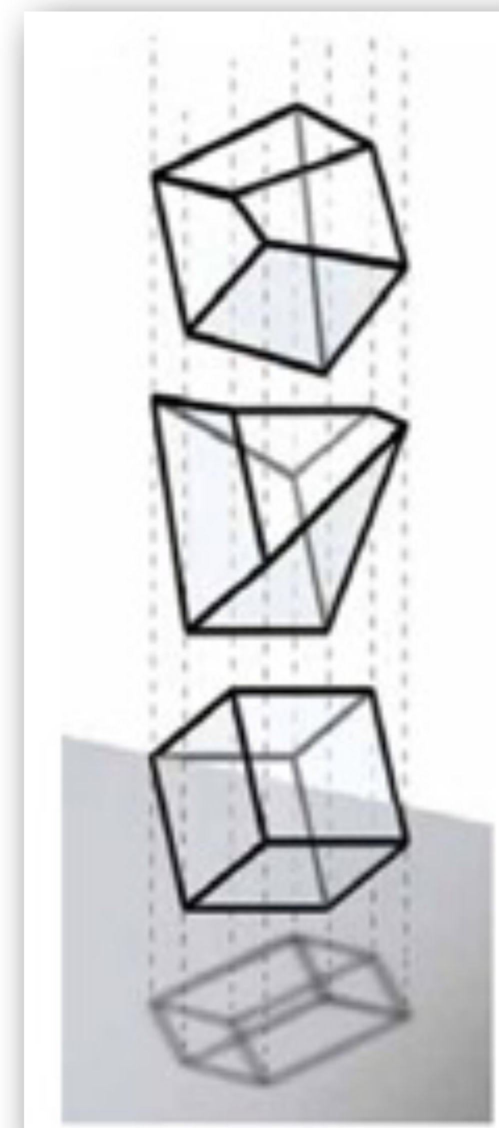

Ernst & Bülthoff, Trends Cogn Sci (2004)

What is the goal of the brain?

- Muscle control
  - for this, we need to make desicions
  - this is simple if we are satisfied with reflexes
  - otherwise, we need to remember things — we need to construct the representation of the knowledge we acquired in the neural tissue
  - based on this knowledge we interpret our sensory data — **inference**
  - we need to constantly update the knowledge base based on novel information — **learning**
- The brain uses the acquired knowledge to predict the outcome of decisions — building a **model** of the environment (well, we are modelling processes that model some other processes)
  - in neuroscience we can use models for analyzing data, or we can speak of models that the brain builds — not to be confused

What is the goal of the brain?

- Muscle control
    - for this, we need to make desicions
    - this is simple if we are satisfied with reflexes
    - otherwise, we need to remember things — we need to construct the representation of the knowledge we acquired in the neural tissue
    - based on this knowledge we interpret our sensory data — **inference**
    - we need to constantly update the knowledge base based on novel information — **learning**
- The brain uses the acquired knowledge to predict the outcome of decisions — building a **model** of the environment (well, we are modelling processes that model some other processes)
    - in neuroscience we can use models for analyzing data, or we can speak of models that the brain builds — not to be confused

Perception as inference

- what are the directly observable quantities for the brain?

Perception as inference

- what are the directly observable quantities for the brain?
    - activations of sensory neurons (cones and rodes, etc, …)

Perception as inference

- what are the directly observable quantities for the brain?
    - activations of sensory neurons (cones and rodes, etc, …)
- what are the relevant quantities?

Perception as inference

- what are the directly observable quantities for the brain?
    - activations of sensory neurons (cones and rodes, etc, …)
- what are the relevant quantities?
    - combination of complex features, such as if there is a lion in the vicinity?

Perception as inference

- what are the directly observable quantities for the brain?
  - activations of sensory neurons (cones and rodes, etc, …)
- what are the relevant quantities?
  - combination of complex features, such as if there is a lion in the vicinity?
- generative models describe the relationship between the unobserved (latent) quantities and observed quantities

Perception as inference

- what are the directly observable quantities for the brain?
  - activations of sensory neurons (cones and rodes, etc, …)
- what are the relevant quantities?
  - combination of complex features, such as if there is a lion in the vicinity?
- generative models describe the relationship between the unobserved (latent) quantities and observed quantities
  - In this model we need to make inference on the latents based on observations (based on pixel observations tell whether we see a lion)

Perception as inference

- what are the directly observable quantities for the brain?
    - activations of sensory neurons (cones and rodes, etc, …)
- what are the relevant quantities?
    - combination of complex features, such as if there is a lion in the vicinity?
- generative models describe the relationship between the unobserved (latent) quantities and observed quantities
    - In this model we need to make inference on the latents based on observations (based on pixel observations tell whether we see a lion)
    - If we know the model we actually know how to produce observations (e.g. how to draw a lion, dream a lion)

making high-level inferences

- How many whees does a car have?

making high-level inferences

- How many whees does a car have?
    - usually we only see three of them, bu surely, we would say four

making high-level inferences

- How many whees does a car have?
  - usually we only see three of them, bu surely, we would say four
  - this is learned information: once we g around the car we see all of them

making high-level inferences

- How many whees does a car have?
  - usually we only see three of them, bu surely, we would say four
  - this is learned information: once we g around the car we see all of them
    - however, this requires recognizing object constancy

making high-level inferences

- How many whees does a car have?
  - usually we only see three of them, bu surely, we would say four
  - this is learned information: once we g around the car we see all of them
    - however, this requires recognizing object constancy
    - sounds trivial. For you. Not for an infant.

making high-level inferences

- How many whees does a car have?
  - usually we only see three of them, bu surely, we would say four
  - this is learned information: once we g around the car we see all of them
    - however, this requires recognizing object constancy
    - sounds trivial. For you. Not for an infant.
  - we could say that if the car had only three wheels, it would be unstable

making high-level inferences

- How many whees does a car have?
  - usually we only see three of them, bu surely, we would say four
  - this is learned information: once we g around the car we see all of them
    - however, this requires recognizing object constancy
    - sounds trivial. For you. Not for an infant.
  - we could say that if the car had only three wheels, it would be unstable
    - this requires even more knowledge, unfortunately (about weight distributions, and about gravitation)

What is the goal of the brain?

- Muscle control
  - for this, we need to make desicions
  - this is simple if we are satisfied with reflexes
  - otherwise, we need to remember things — we need to construct the representation of the knowledge we acquired in the neural tissue
  - based on this knowledge we interpret our sensory data — **inference**
  - we need to constantly update the knowledge base based on novel information — **learning**
- The brain uses the acquired knowledge to predict the outcome of decisions — biulding a **model** of the environment - **modellt épít** a világról (well, we are modelling processes that model some other processes)
  - in neuroscience we can use models for analyzing data, or we can speak of models that the brain builds — not to be confused

What is the goal of the brain?

- Muscle control
  - for this, we need to make desicions
  - this is simple if we are satisfied with reflexes
  - otherwise, we need to remember things — we need to construct the representation of the knowledge we acquired in the neural tissue
  - based on this knowledge we interpret our sensory data — **inference**
  - we need to constantly update the knowledge base based on novel information — **learning**
- The brain uses the acquired knowledge to predict the outcome of decisions — biulding a **model** of the environment - **modellt épít** a világról (well, we are modelling processes that model some other processes)
  - in neuroscience we can use models for analyzing data, or we can speak of models that the brain builds — not to be confused

What is learning?

What is learning?

- Building the generative model of the environment based on observations

What is learning?

- Building the generative model of the environment based on observations
- We need to find the regularities characterizing the observations, yielding a model that provides a concise description of observations

What is learning?

- Building the generative model of the environment based on observations

- We need to find the regularities characterizing the observations, yielding a model that provides a concise description of observations

    - e.g. the identity of objects is not affected by a number of properties, such as pose, viewing angle, lighting, …

A tanult információ és a megfigyelések integrációja

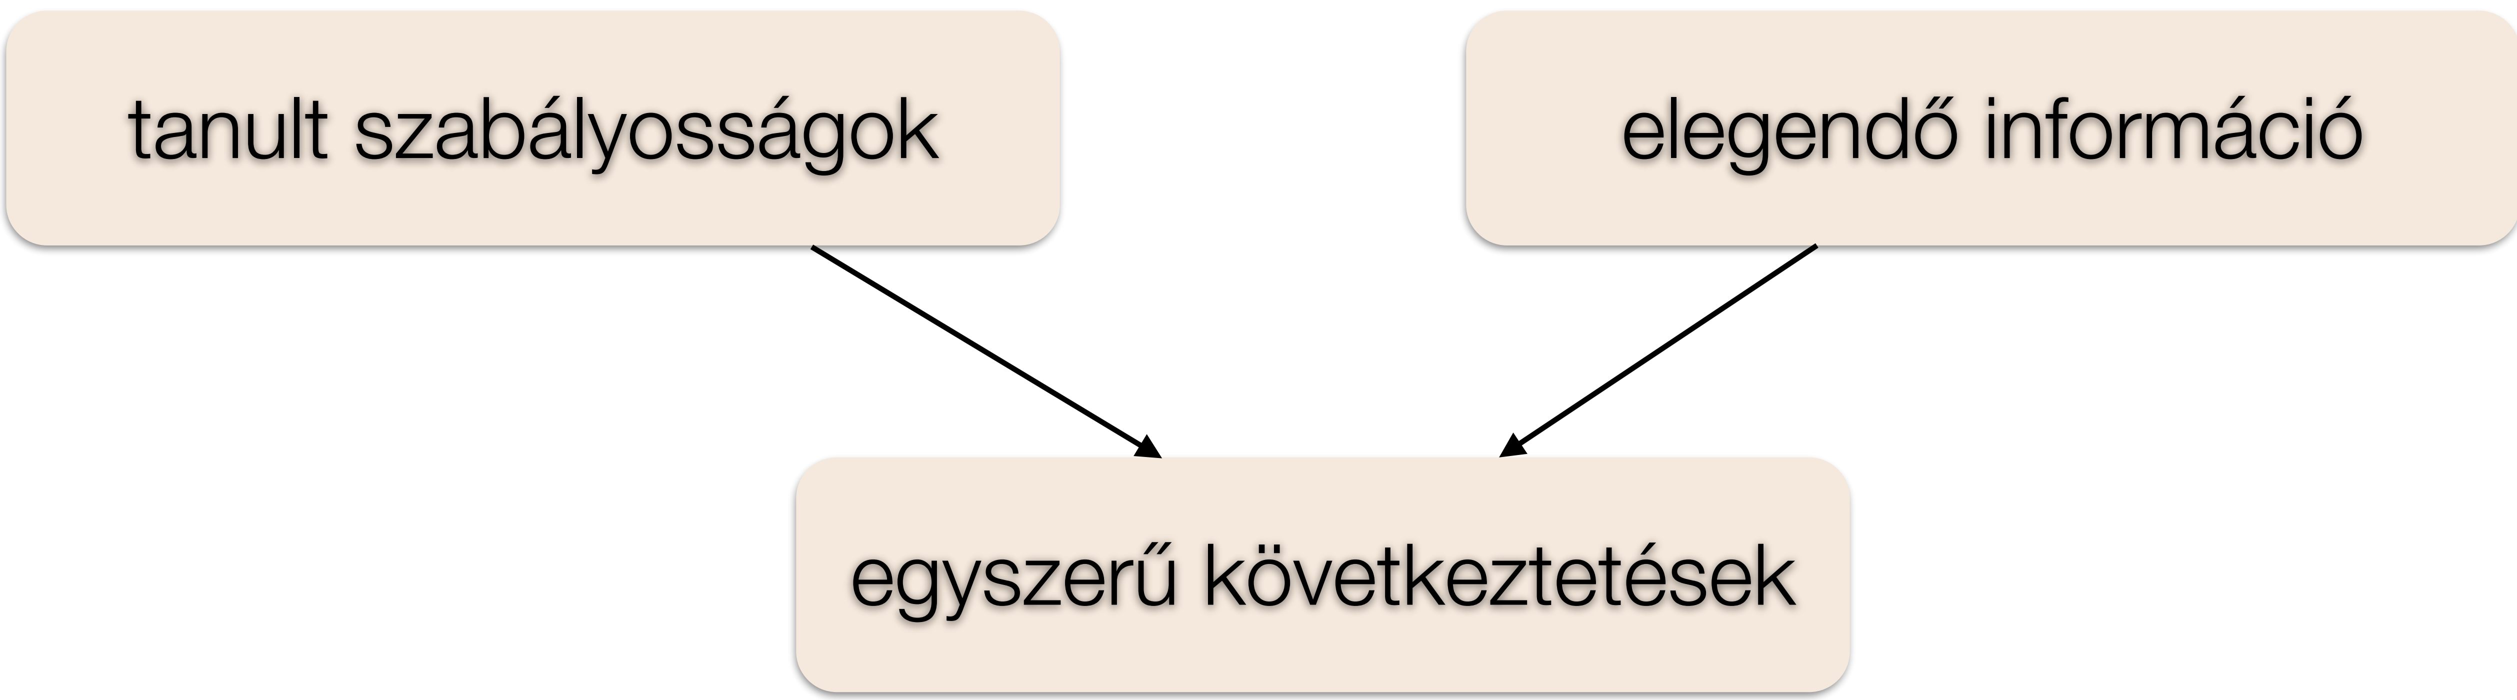

A tanult információ és a megfigyelések integrációja

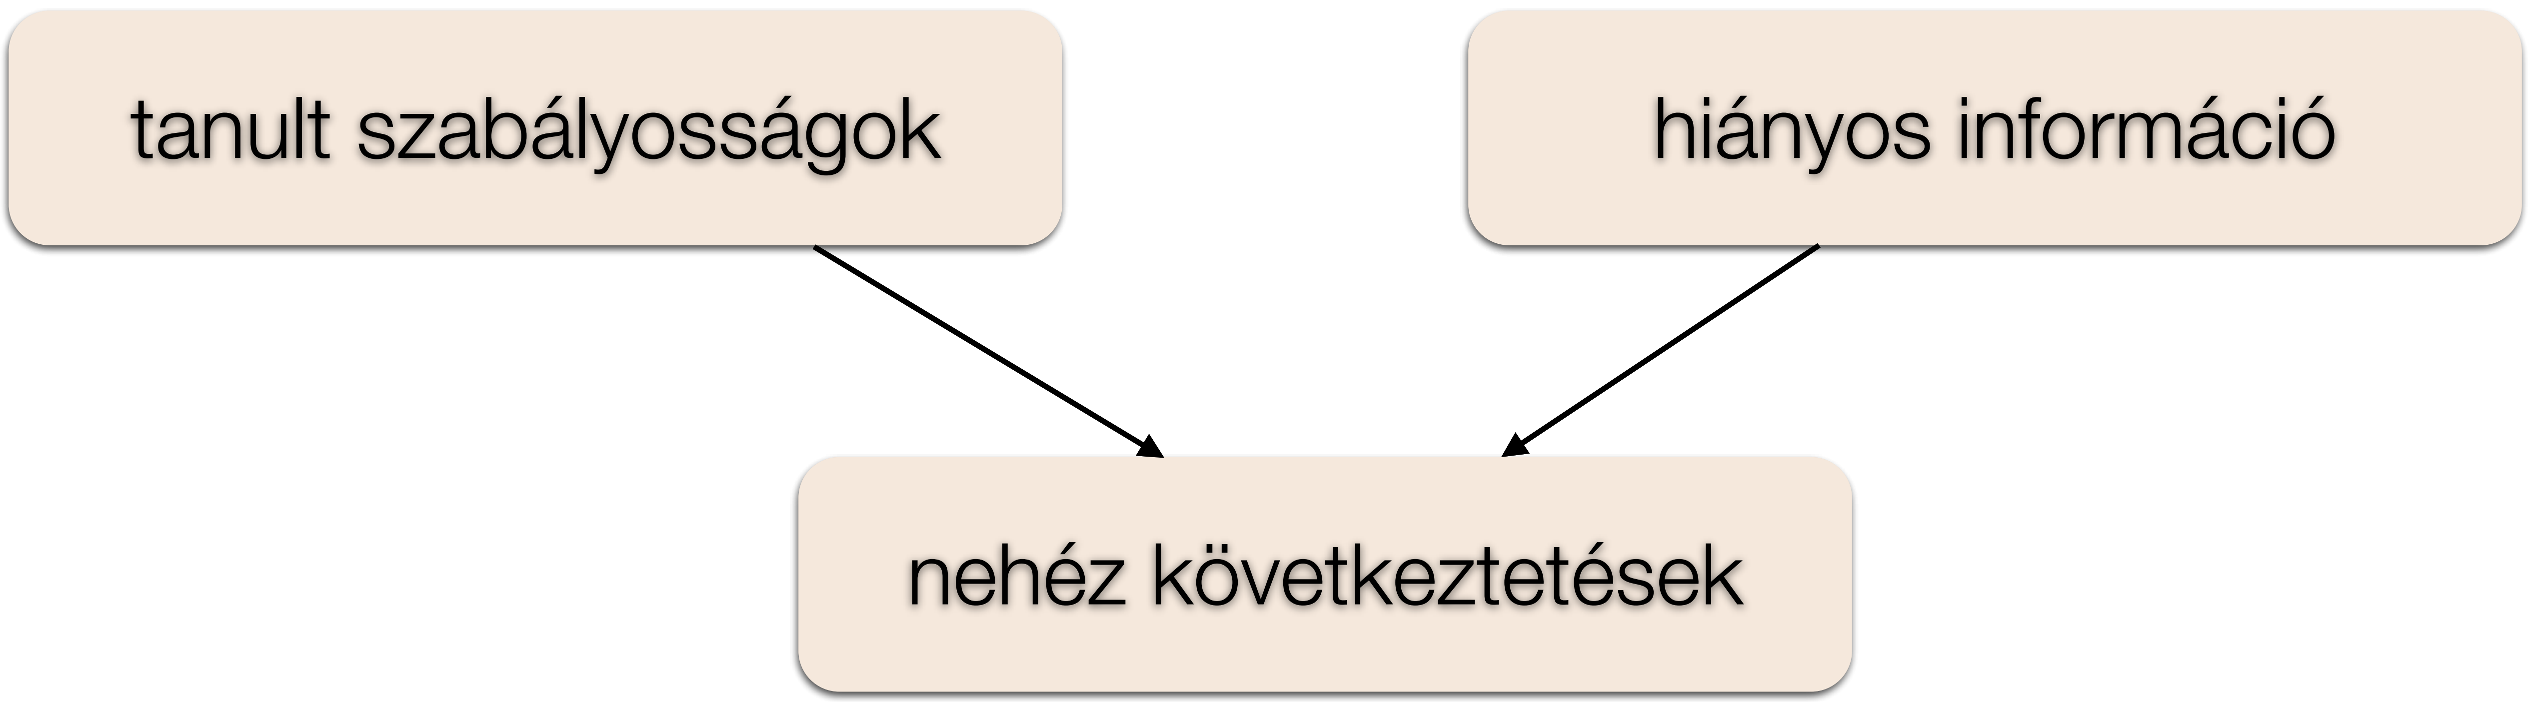

A tanult információ és a megfigyelések integrációja

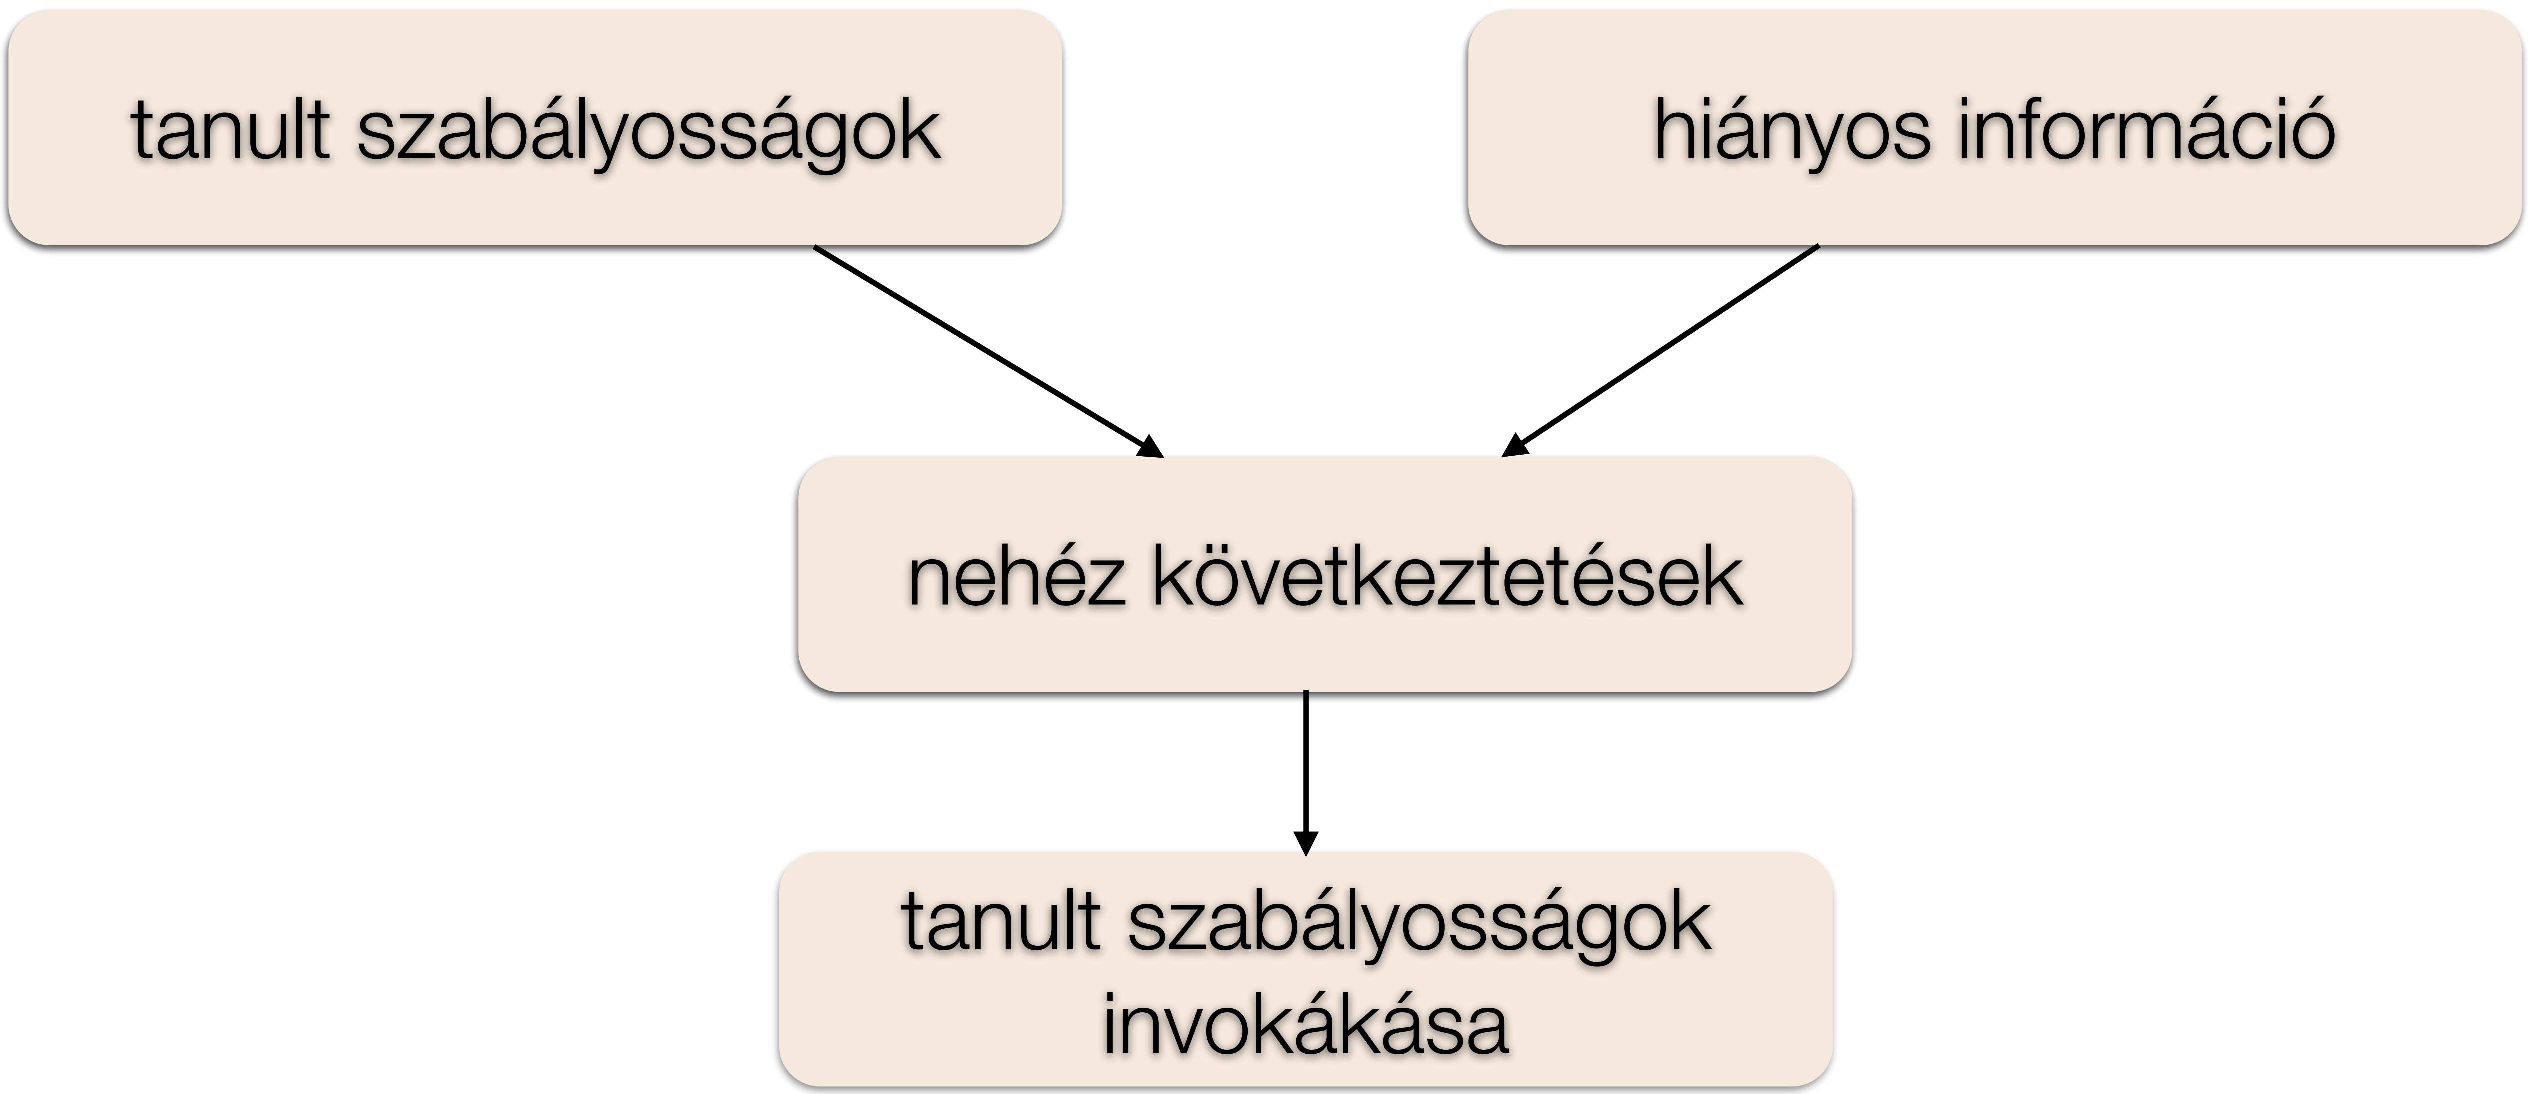

A tanult információ és a megfigyelések integrációja

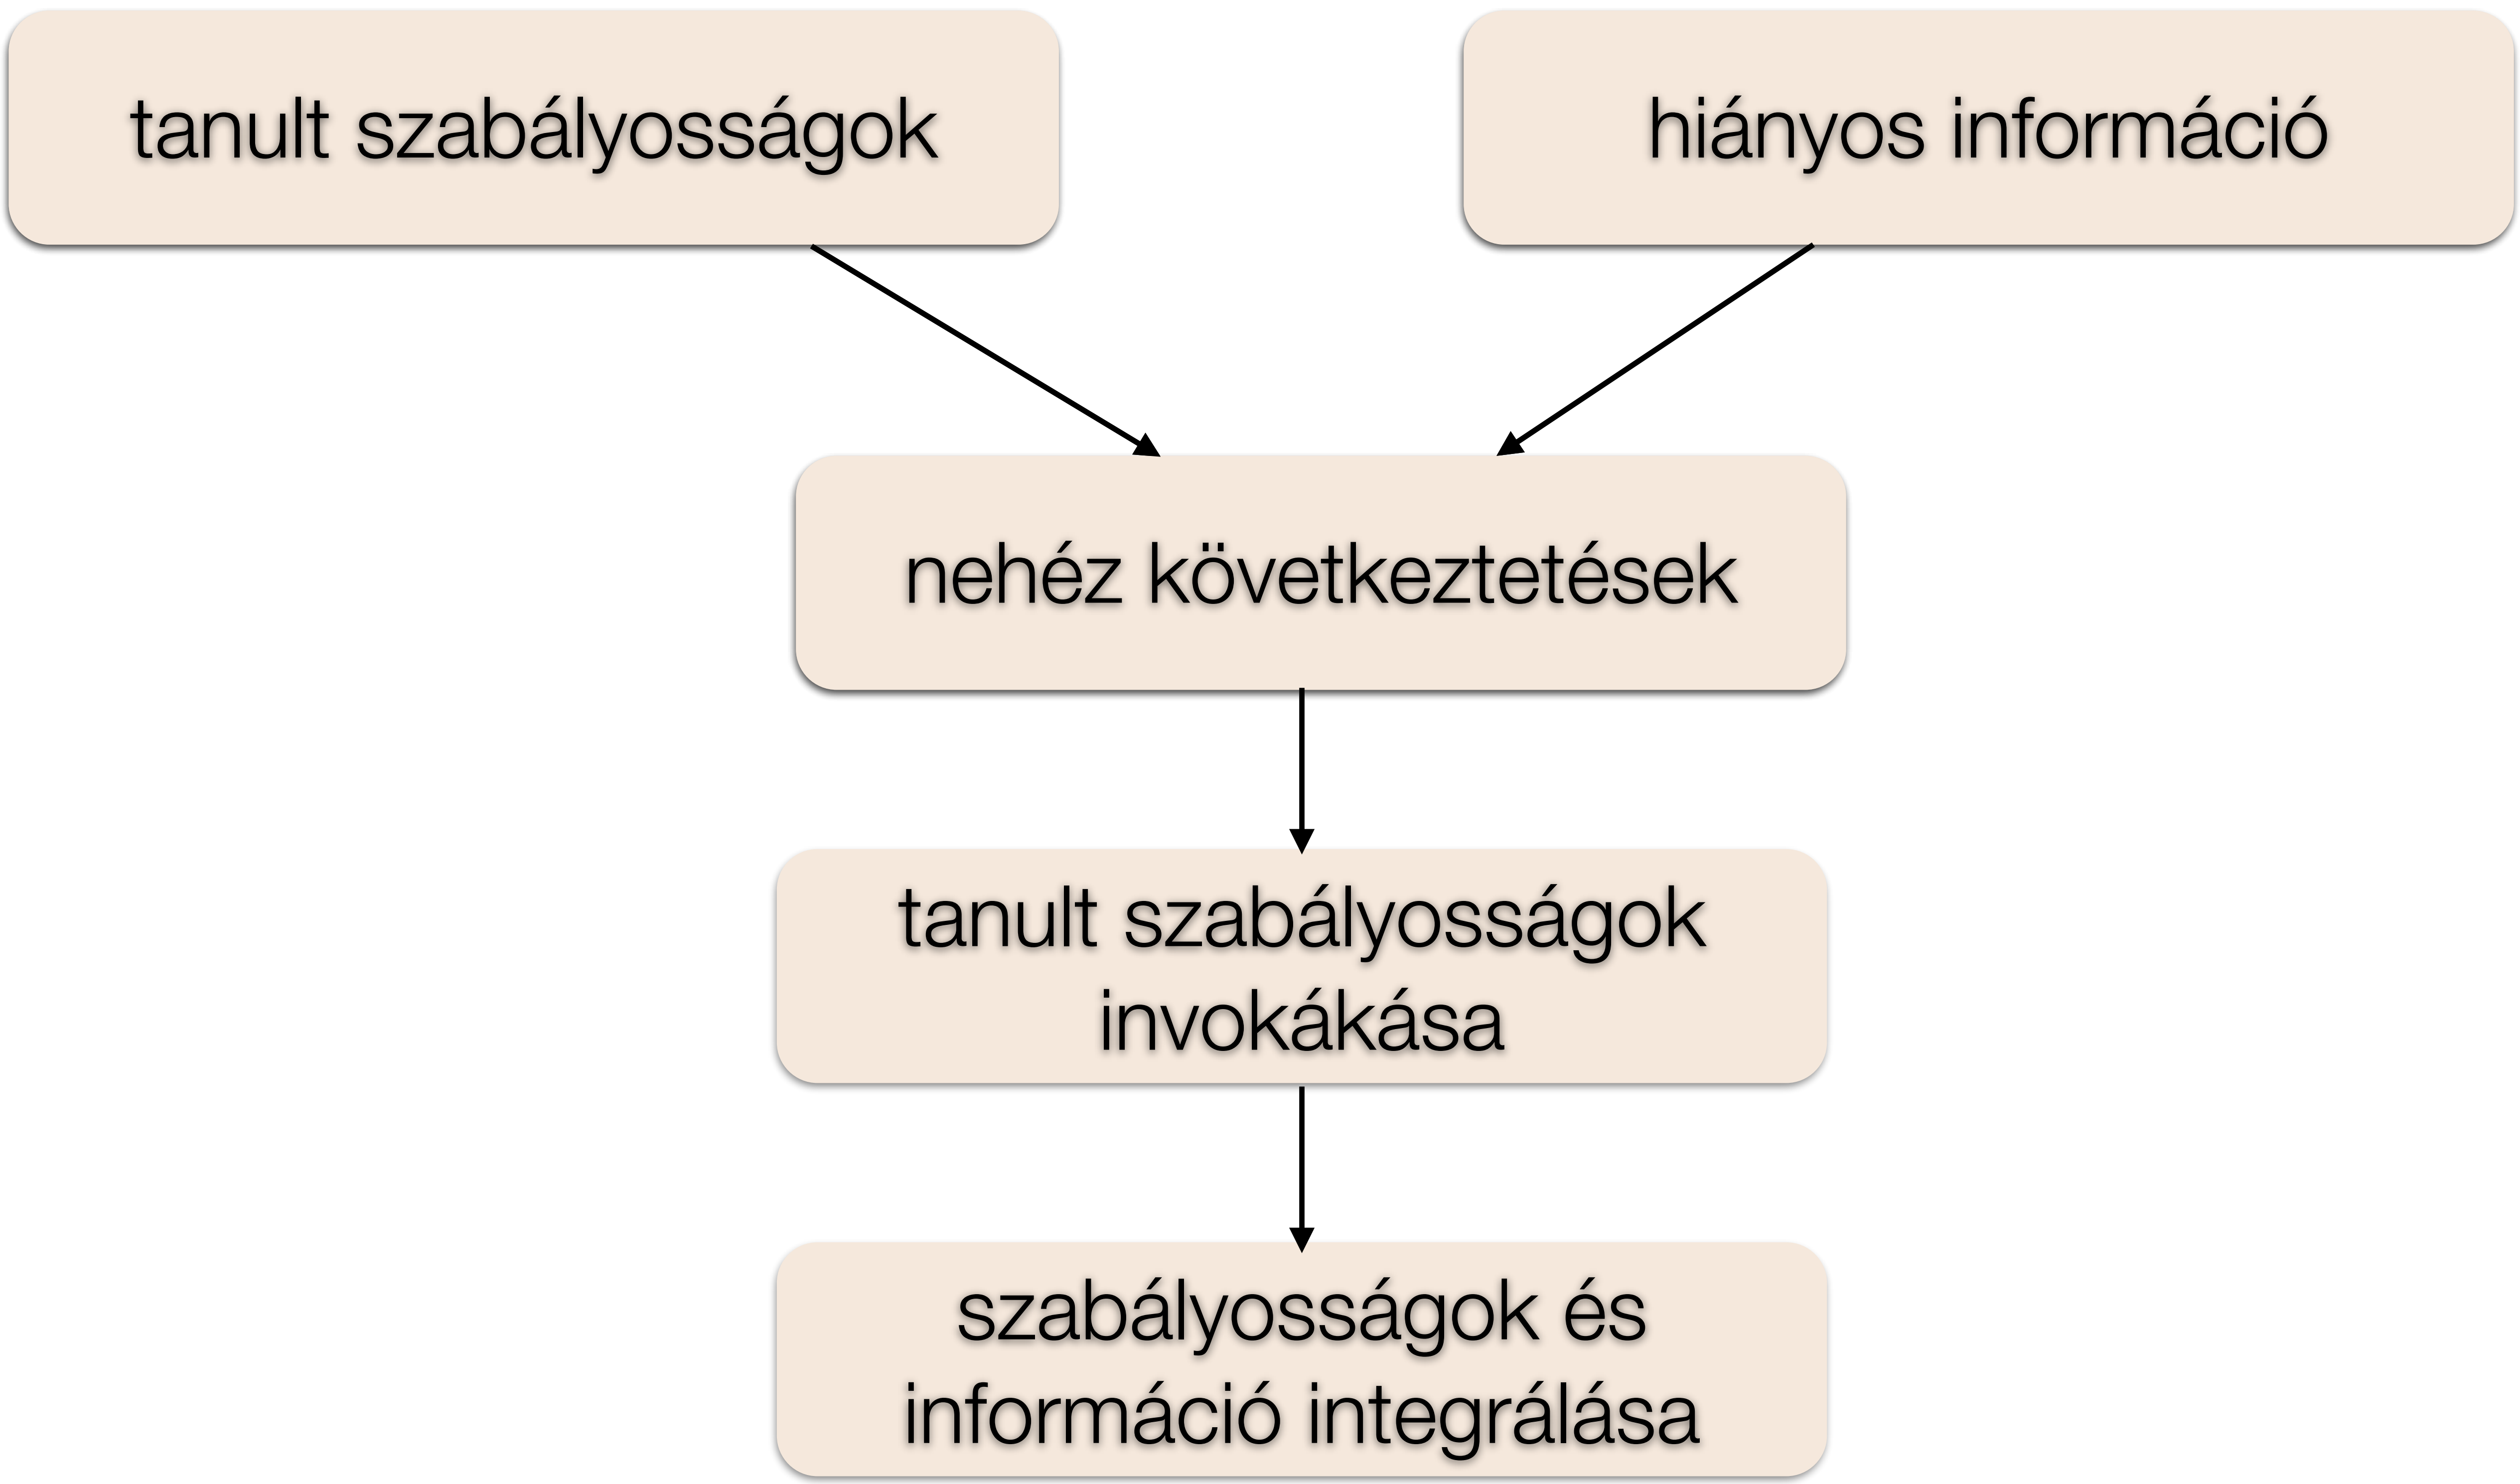

Rule of thumbs for identifying objects: Gestalt rules

- Proximity

Rule of thumbs for identifying objects: Gestalt rules

- Proximity

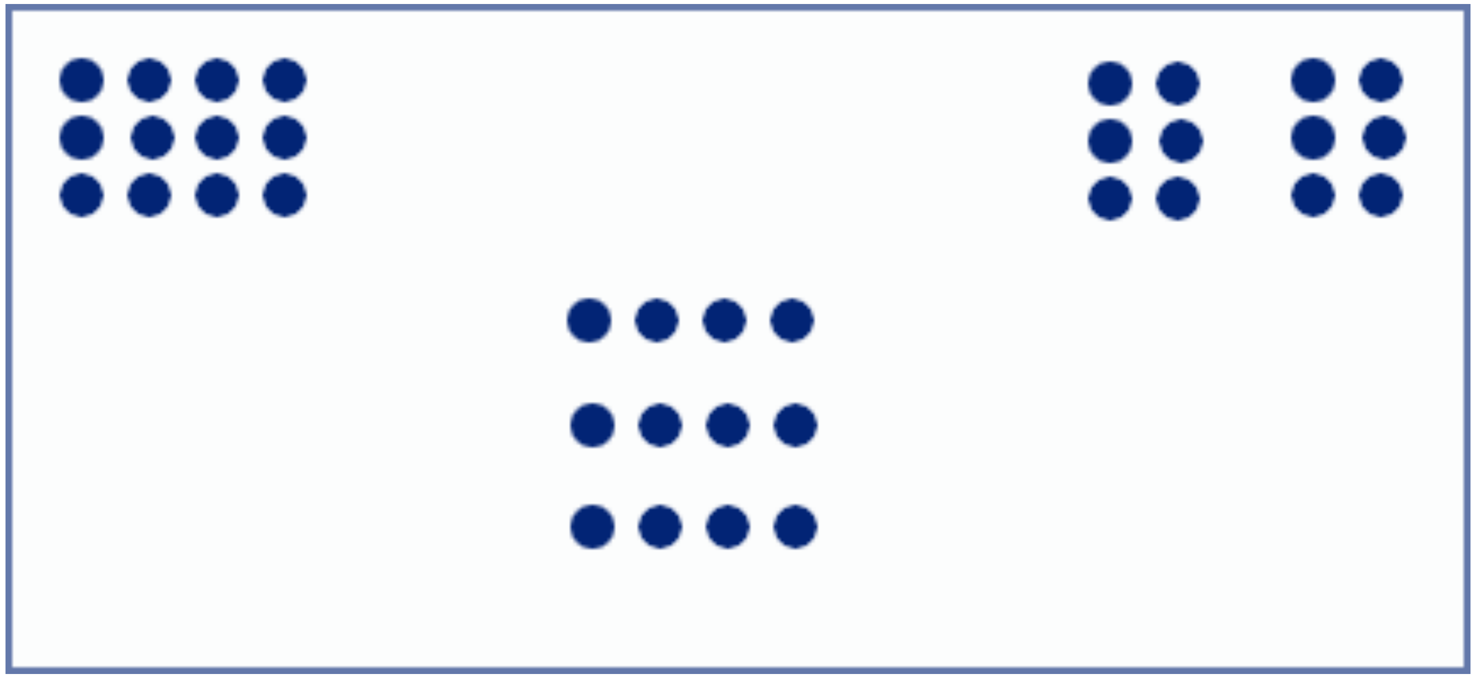

Rule of thumbs for identifying objects: Gestalt rules

- Proximity

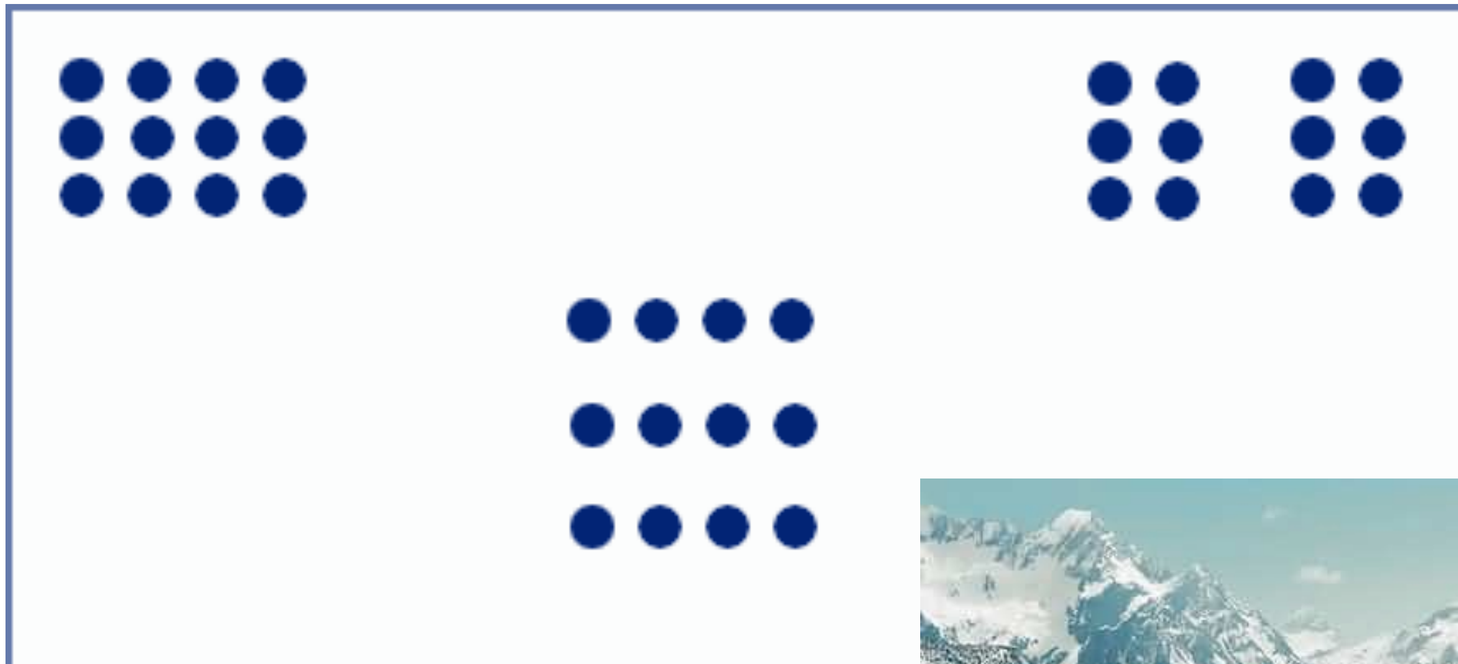

Rule of thumbs for identifying objects: Gestalt rules

- Similarity

Rule of thumbs for identifying objects: Gestalt rules

- Similarity

Rule of thumbs for identifying objects: Gestalt rules

- Similarity

Rule of thumbs for identifying objects: Gestalt rules

- Good continuity

Rule of thumbs for identifying objects: Gestalt rules

- Good continuity

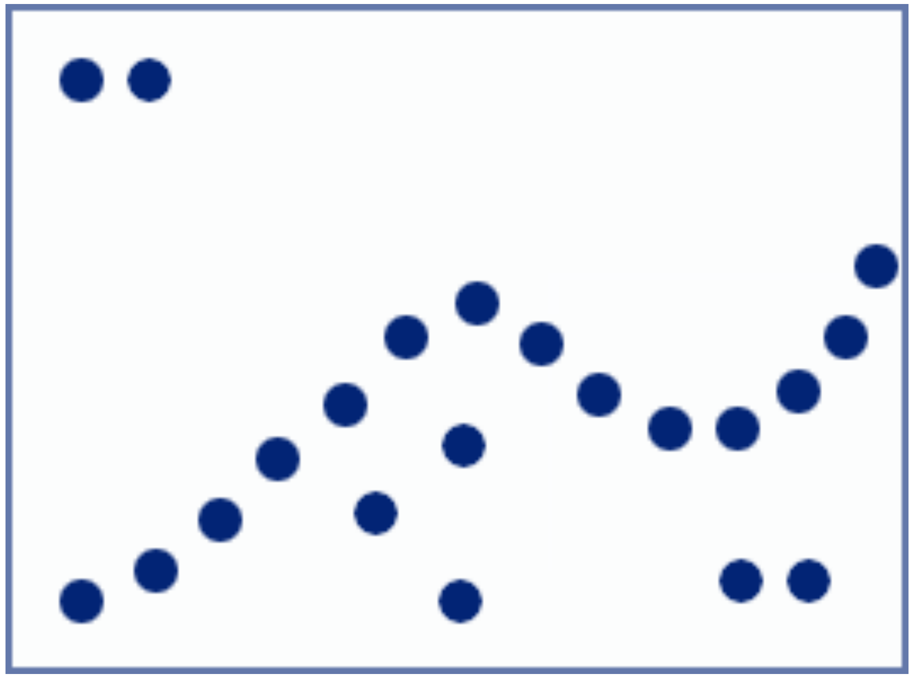

Rule of thumbs for identifying objects: Gestalt rules

- Good continuity

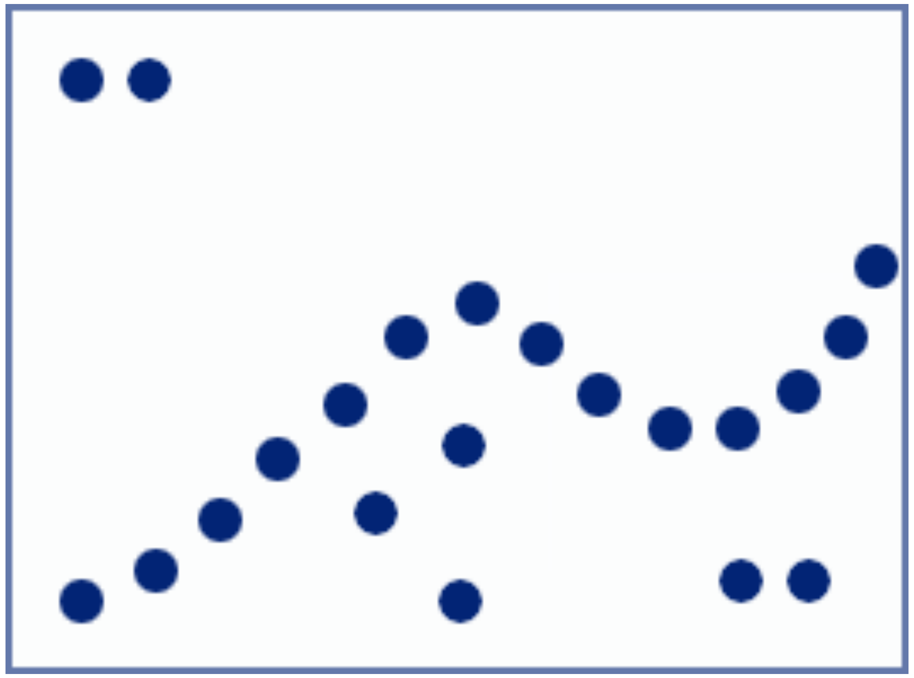

Rule of thumbs for identifying objects: Gestalt rules

- Common fate

Rule of thumbs for identifying objects: Gestalt rules

- Common fate

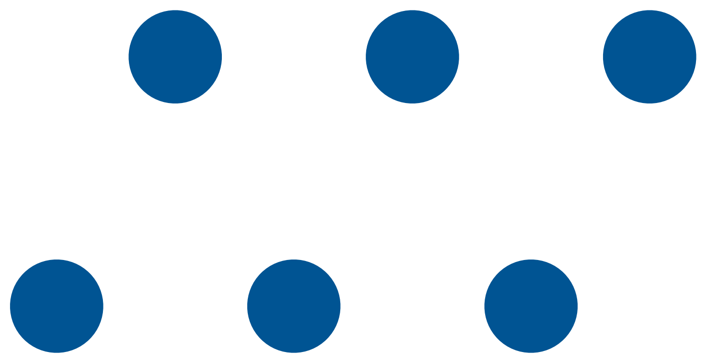

Rule of thumbs for identifying objects: Gestalt rules

- Common fate

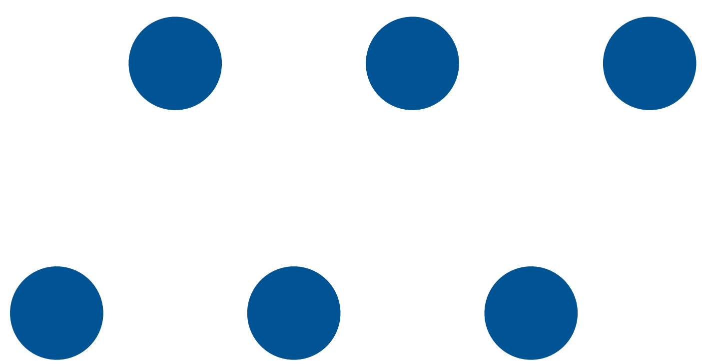

Integration of learned regularities and observations

Where is the sun?

Jennifer Sun[1] and Pietro Perona[1,2]

[1] *California Institute of Technology 136-93, Pasadena, California 91125, USA*
[2] *Universita di Padova, Via Ognissanti 72, 35131 Padova, Italy*
Correspondence should be addressed to P.P. (perona@vision.caltech.edu)

nature *neuroscience* • volume 1 no 3 • july 1998

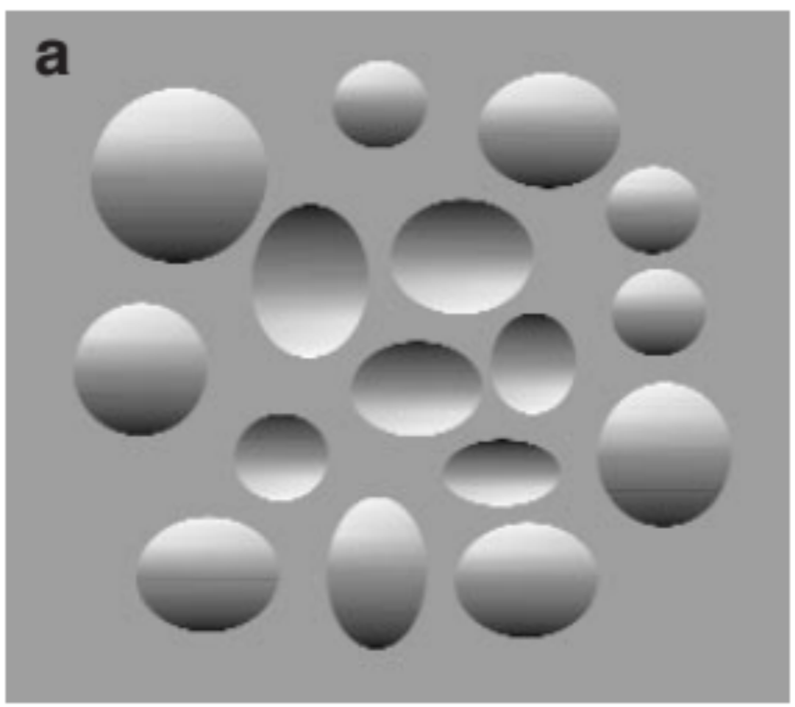

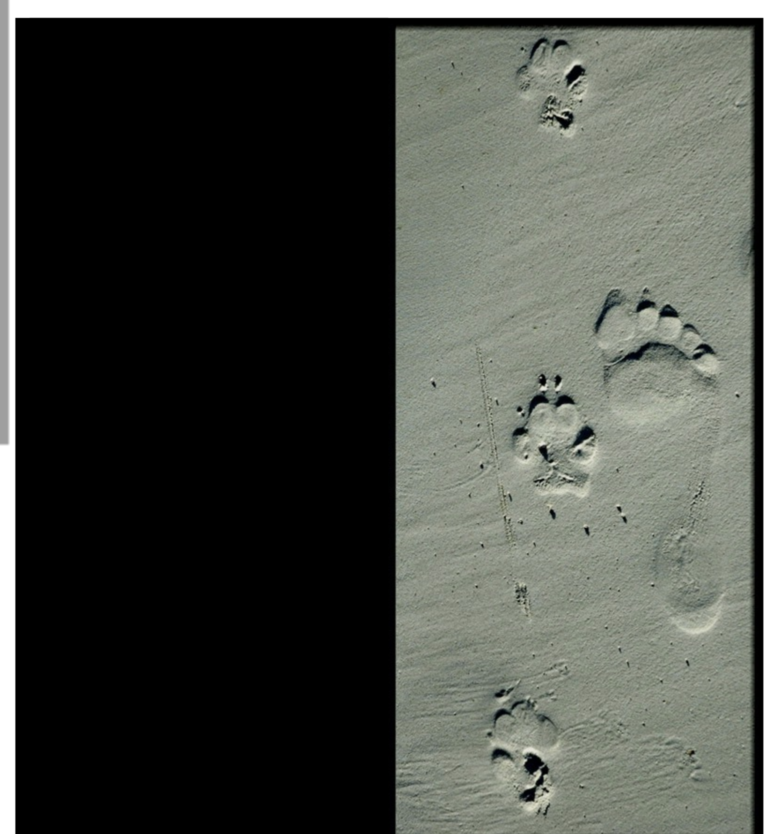

Integration of learned regularities and observations

Where is the sun?

Jennifer Sun[1] and Pietro Perona[1,2]

[1] *California Institute of Technology 136-93, Pasadena, California 91125, USA*
[2] *Universita di Padova, Via Ognissanti 72, 35131 Padova, Italy*
Correspondence should be addressed to P.P. (perona@vision.caltech.edu)

nature *neuroscience* • volume 1 no 3 • july 1998

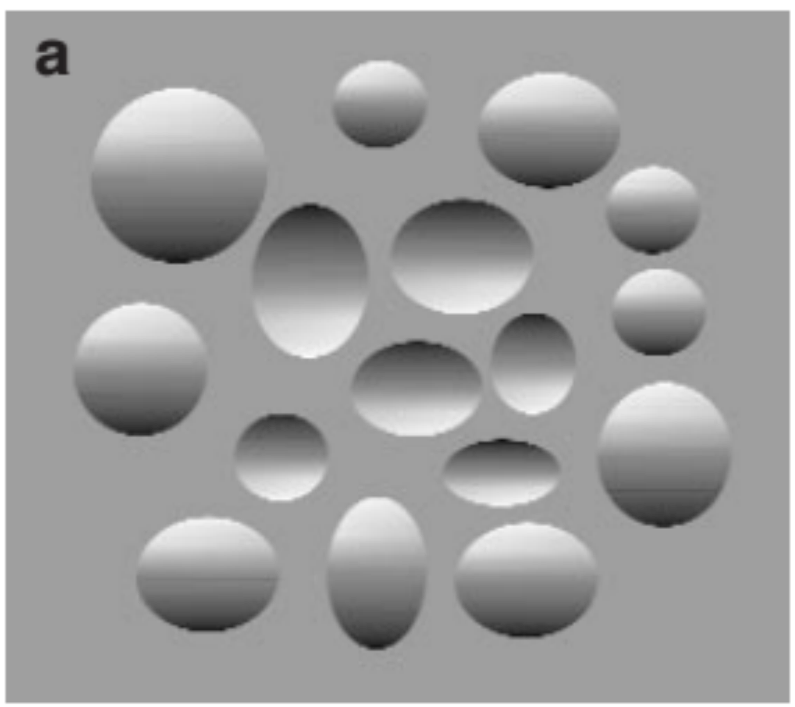

Integration of learned regularities and observations

Where is the sun?

Jennifer Sun[1] and Pietro Perona[1,2]

[1] California Institute of Technology 136-93, Pasadena, California 91125, USA
[2] Universita di Padova, Via Ognissanti 72, 35131 Padova, Italy
Correspondence should be addressed to P.P. (perona@vision.caltech.edu)

nature neuroscience • volume 1 no 3 • july 1998

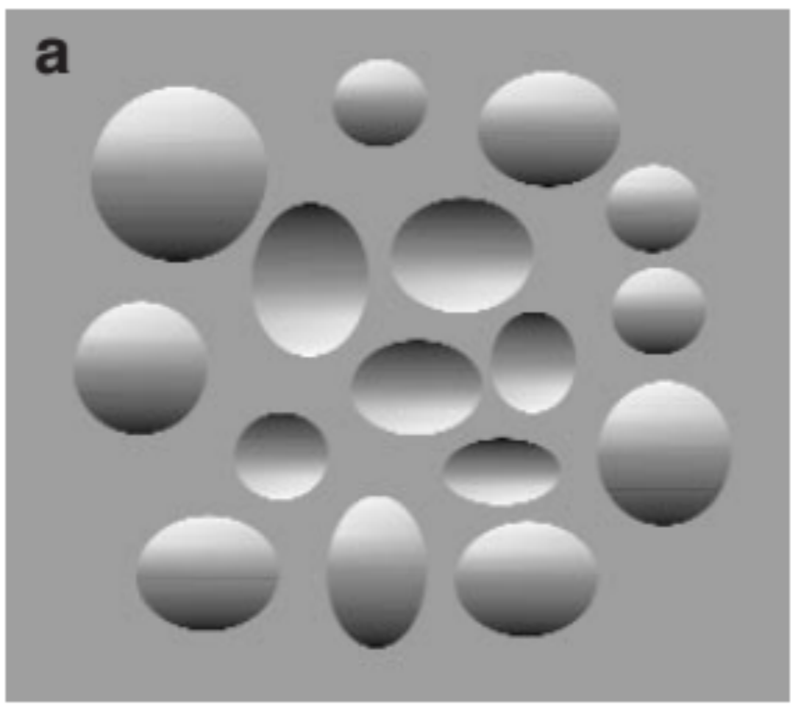

Integration of learned regularities and observations

Where is the sun?

Jennifer Sun[1] and Pietro Perona[1,2]

[1] *California Institute of Technology 136-93, Pasadena, California 91125, USA*
[2] *Universita di Padova, Via Ognissanti 72, 35131 Padova, Italy*
Correspondence should be addressed to P.P. (perona@vision.caltech.edu)

nature *neuroscience* • volume 1 no 3 • july 1998

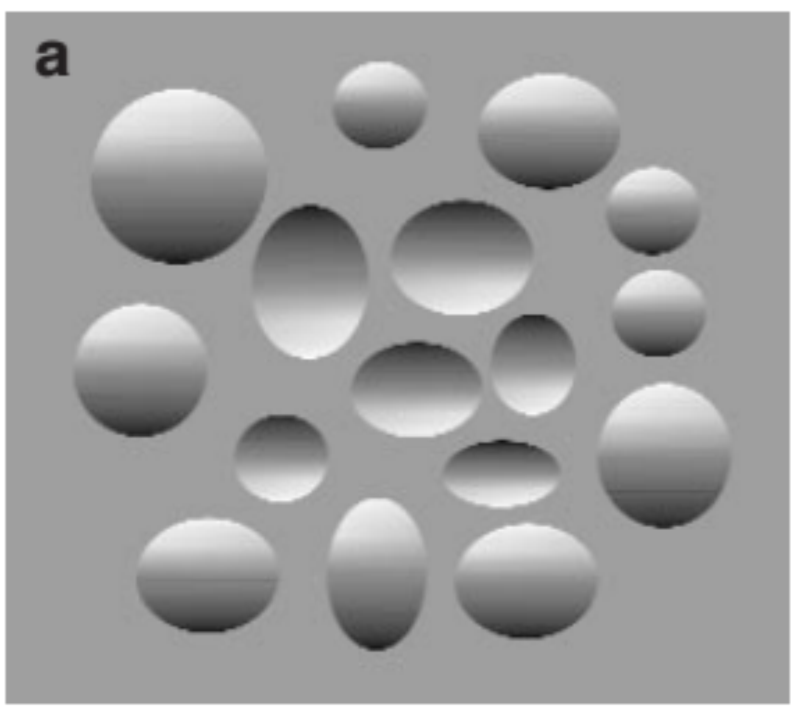

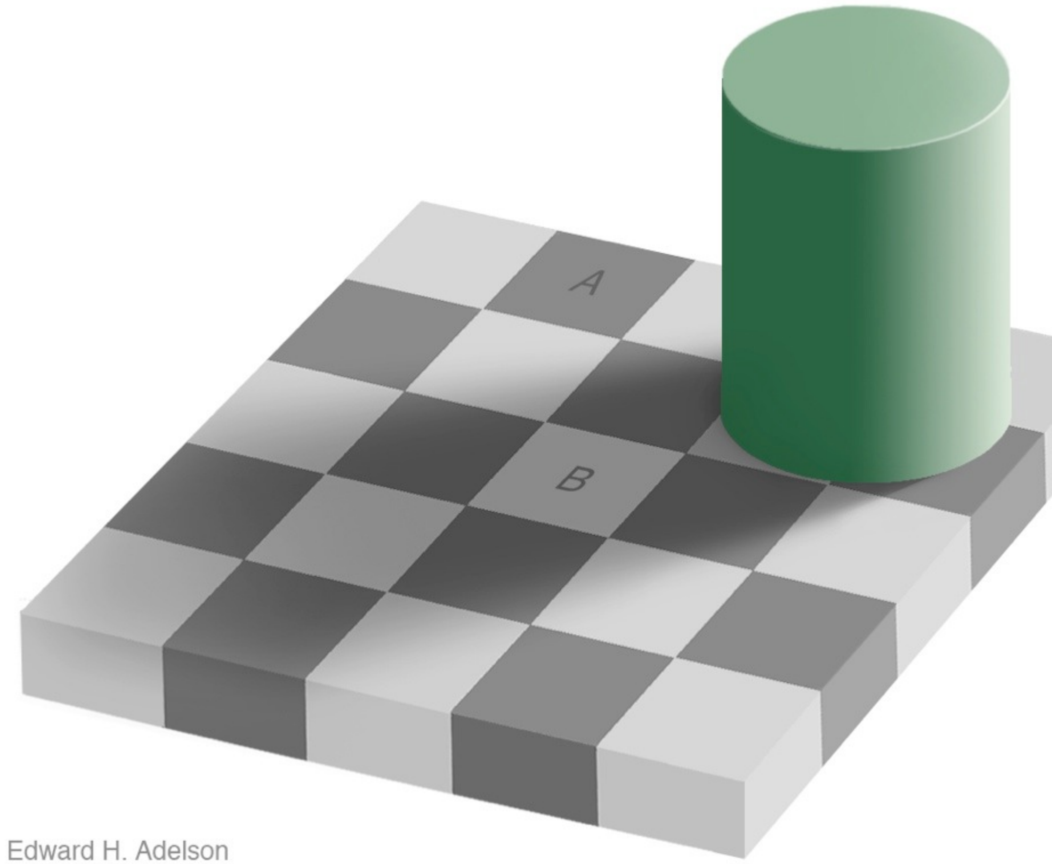

Edward H. Adelson

Integration of learned regularities and observations

Where is the sun?

Jennifer Sun[1] and Pietro Perona[1,2]

[1] California Institute of Technology 136-93, Pasadena, California 91125, USA
[2] Universita di Padova, Via Ognissanti 72, 35131 Padova, Italy
Correspondence should be addressed to P.P. (perona@vision.caltech.edu)

nature neuroscience • volume 1 no 3 • july 1998

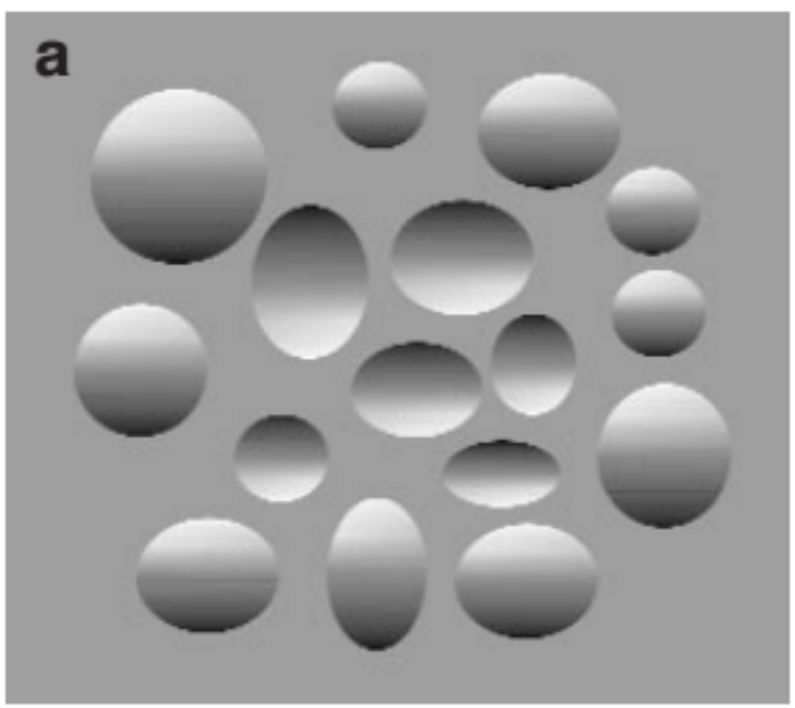

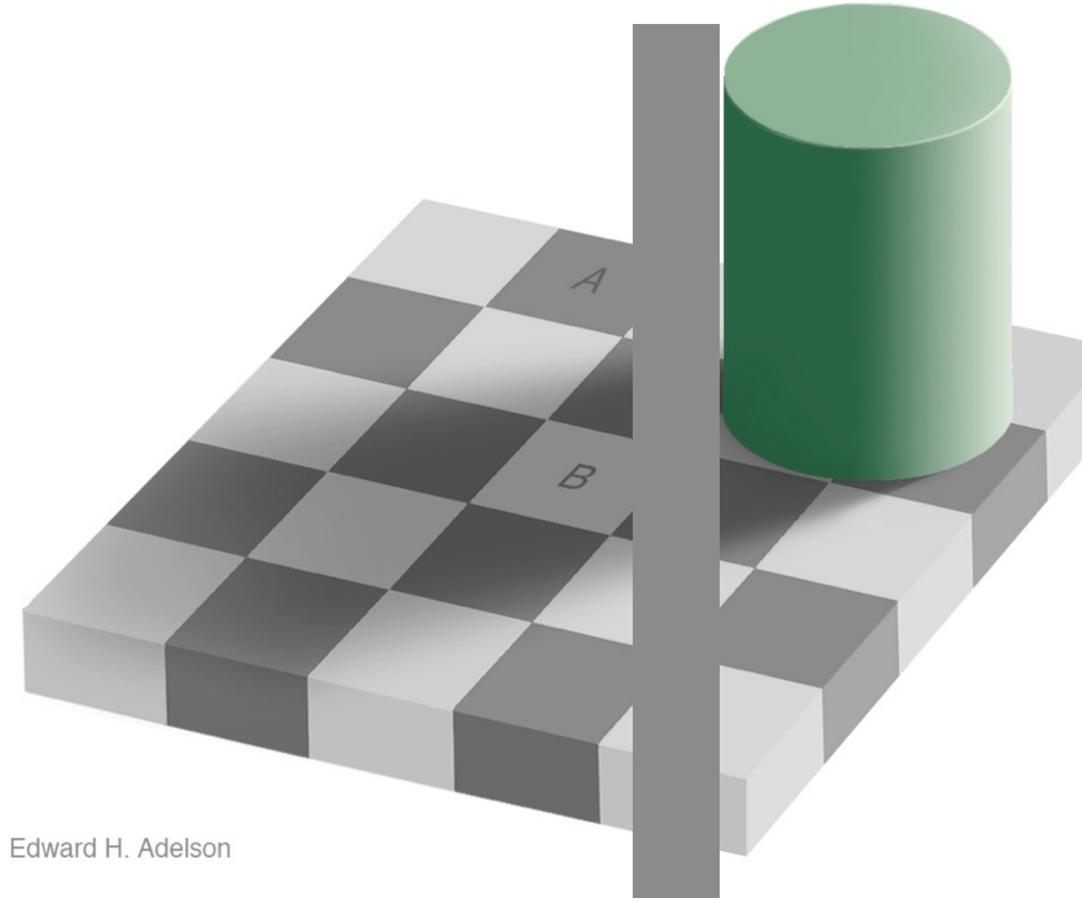

Edward H. Adelson

What is the goal of the brain?

- Muscle control
  - for this, we need to make desicions
  - this is simple if we are satisfied with reflexes
  - otherwise, we need to remember things — we need to construct the representation of the knowledge we acquired in the neural tissue
  - based on this knowledge we interpret our sesnory data — **inference**
  - we need to constantly update the knowledge base based on novel information — **learning**
- The brain uses the acquired knowledge to predict the outcome of decisions — biulding a **model** of the environment - **modellt épít** a világról (well, we are modelling processes that model some other processes)
  - in neuroscience we can use models for analyzing data, or we can speak of models that the brain builds — not to be confused

What is the goal of the brain?

- Muscle control
  - for this, we need to make desicions
  - this is simple if we are satisfied with reflexes
  - otherwise, we need to remember things — we need to construct the representation of the knowledge we acquired in the neural tissue
  - based on this knowledge we interpret our sesnory data — **inference**
  - we need to constantly update the knowledge base based on novel information — **learning**
- The brain uses the acquired knowledge to predict the outcome of decisions — biulding a **model** of the environment - **modellt épít** a világról (well, we are modelling processes that model some other processes)
  - in neuroscience we can use models for analyzing data, or we can speak of models that the brain builds — not to be confused

A reprezentáció kérdései

- Mik azok a környezetet leíró mennyiségek, amelyekre az embereknek (és állatoknak) szükségük van a döntéseik meghozatalához?
    - ezek nem feltétlenül ugyanazok, mint amiket ténylegesen kiválasztanak, de remélhetőleg van kapcsolat
- Mik azok a köztes mennyiségek ezen mennyiségek és a megfigyelések között amelyek hasznosak a számítás során?
- Ha erre a két kérdésre tudjuk a választ, az megadja a mentális modell struktúráját - általában egyszerre ennek csak egy kis részét vizsgáljuk
- Milyen algoritmusok számítják ki ezeket a mennyiségeket?
- Létezik feladatfüggetlen mentális világmodell, vagy az agy különböző feladatokhoz különböző modelleket használ, amiket a kontextus alapján váltogat?
- Hogyan tudunk történeteket kitalálni - azaz hogyan használjuk a mentális modellt arra, hogy akár a tudásunkkal ellentmondó hipotéziseket fogalmazzunk meg és vizsgáljuk (mi lenne, ha…)?

Mi a jó reprezentáció?

- A környezetet leíró változók leképezése az agy mint formális rendszer által megvalósított világmodell változóira
- A jó reprezentáció tömöríti a megfigyeléseket
    - nem emlékezhetünk minden szituáció minden részletére egész életünkben - egyszer megharapott egy fekete kutya, egyszer meg egy barna
    - túl nagy mennyiségű adat lenne
    - nem tudnánk általánosítani az ismereteinket új megfigyelésekre - nem tudnám, mi lesz, ha jön egy fehér kutya

A neurális kód

- Hogyan kapcsolható a mentális modell reprezentációja biofizikai mennyiségekhez?
- Mindenképpen szükségünk van erre?
    - Ha rendelkeznénk a mentális modell és a hozzátartozó algoritmusok teljes leírásával, és kizárólag viselkedéstakarunk prediktálni, akkor nem
    - a gyakorlatban ezzel nem rendelkezünk, és segít a körülhatárolásában, ha biológiai szempontú megszorításokat vehetünk figyelembe
    - természetesen orvosi alkalmazásokhoz szükségesek a neurális szintű leírások

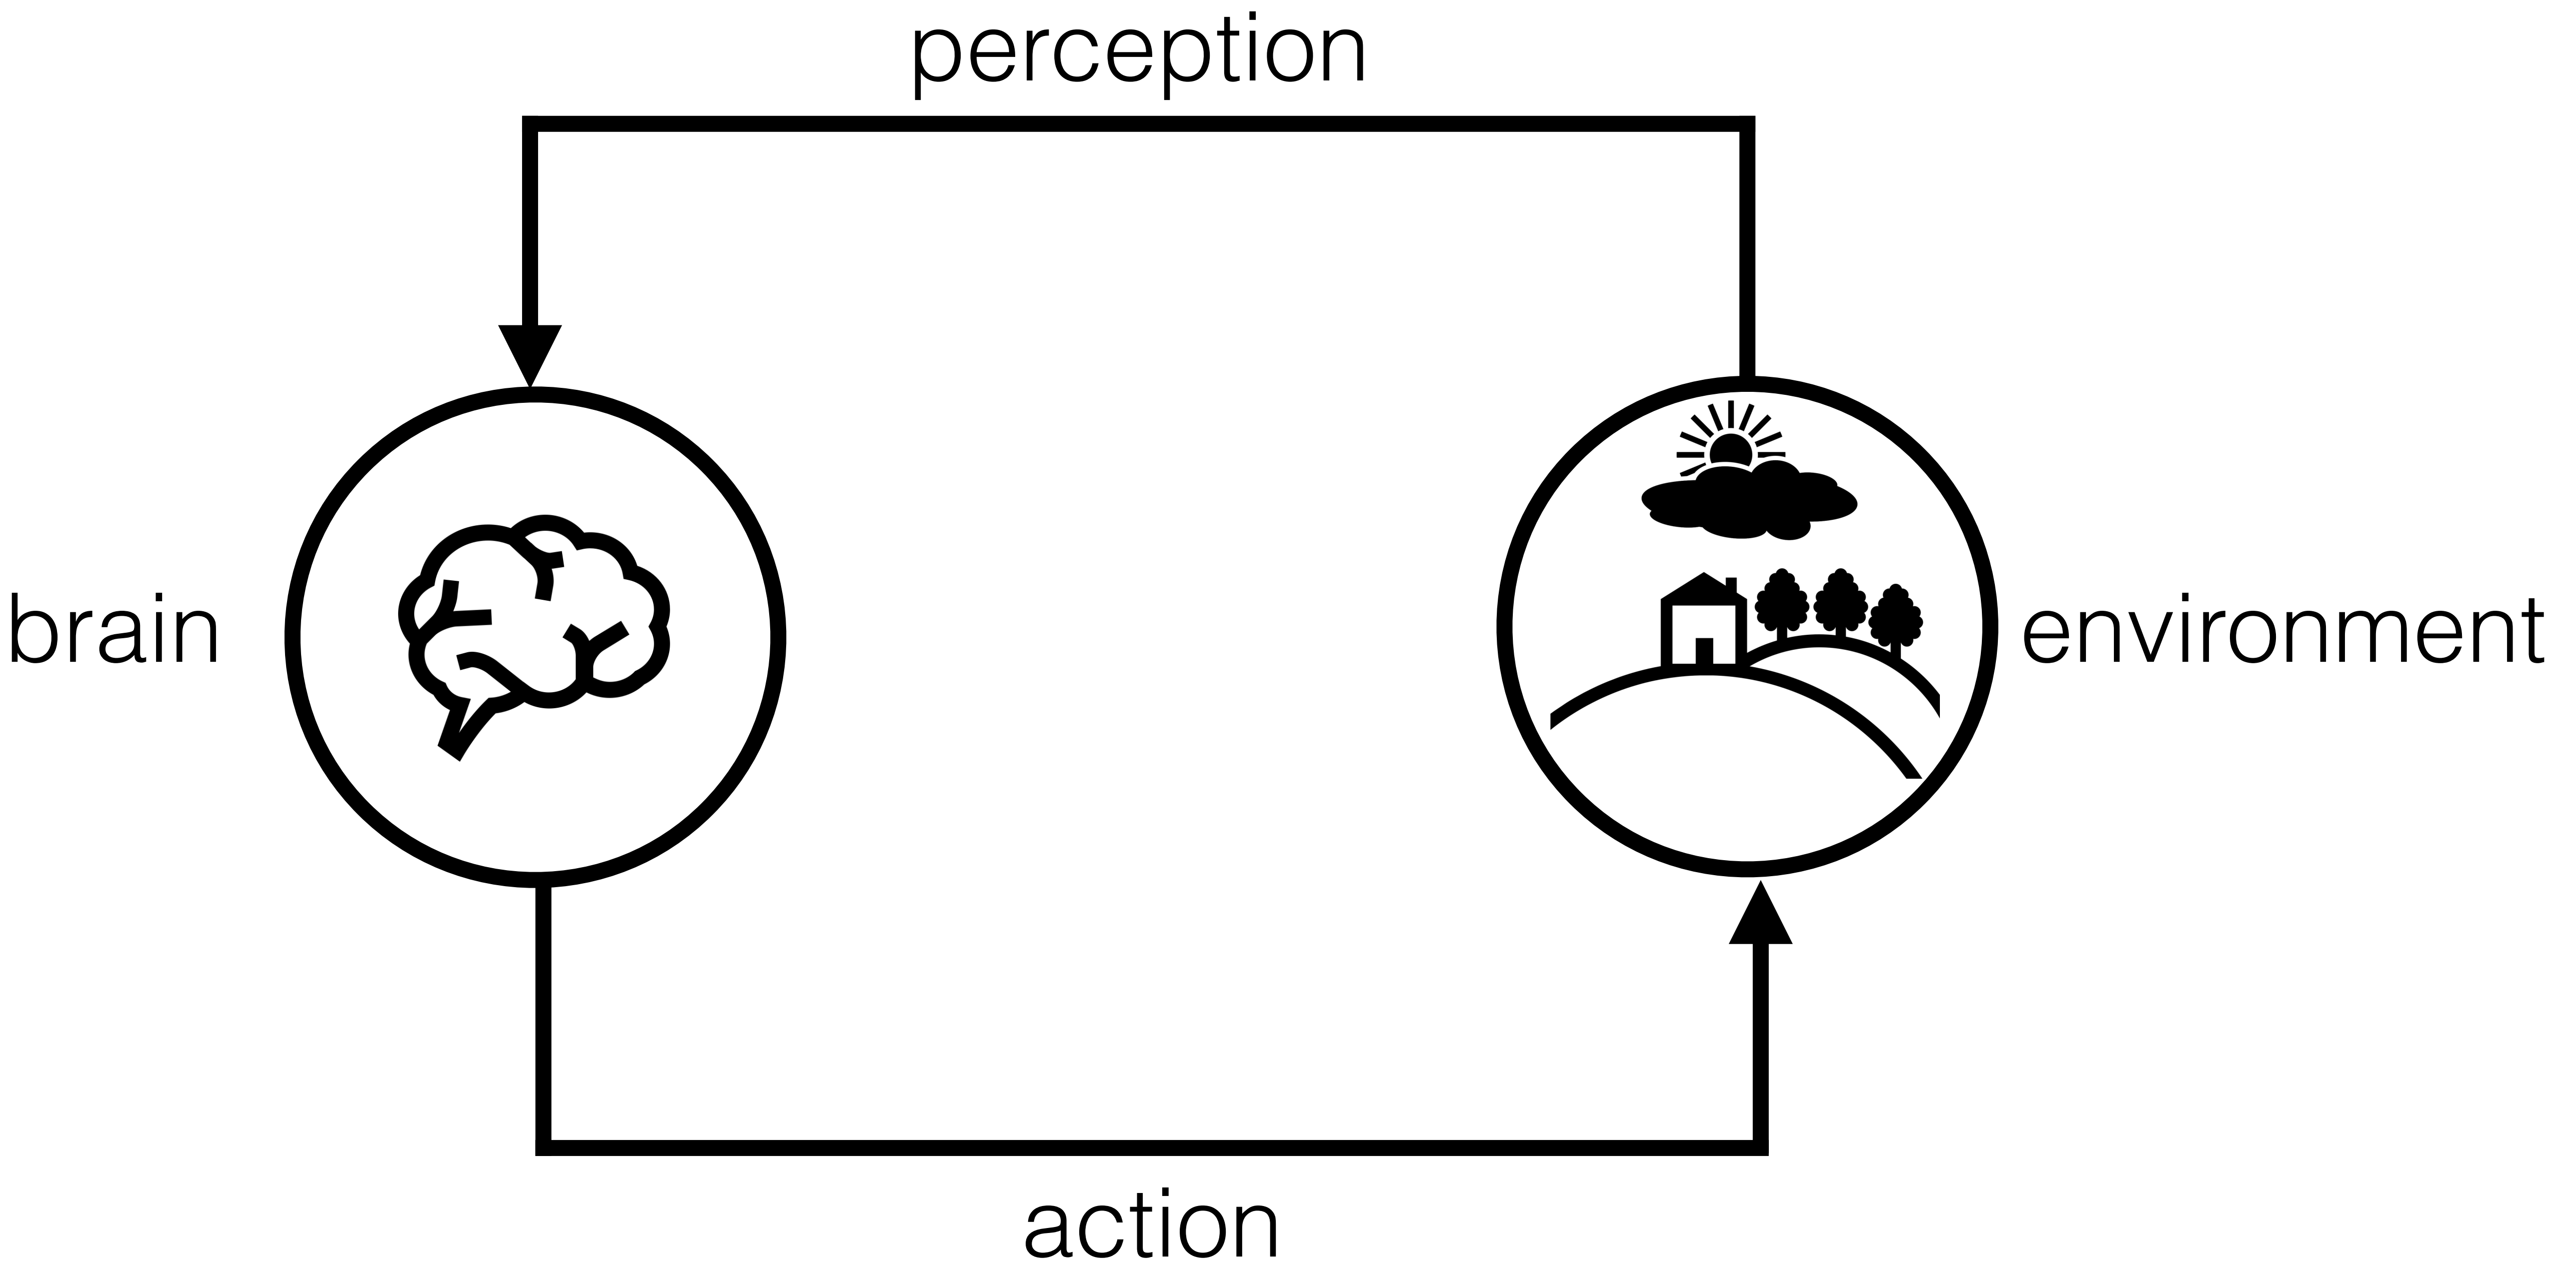
perception
brain
environment
action

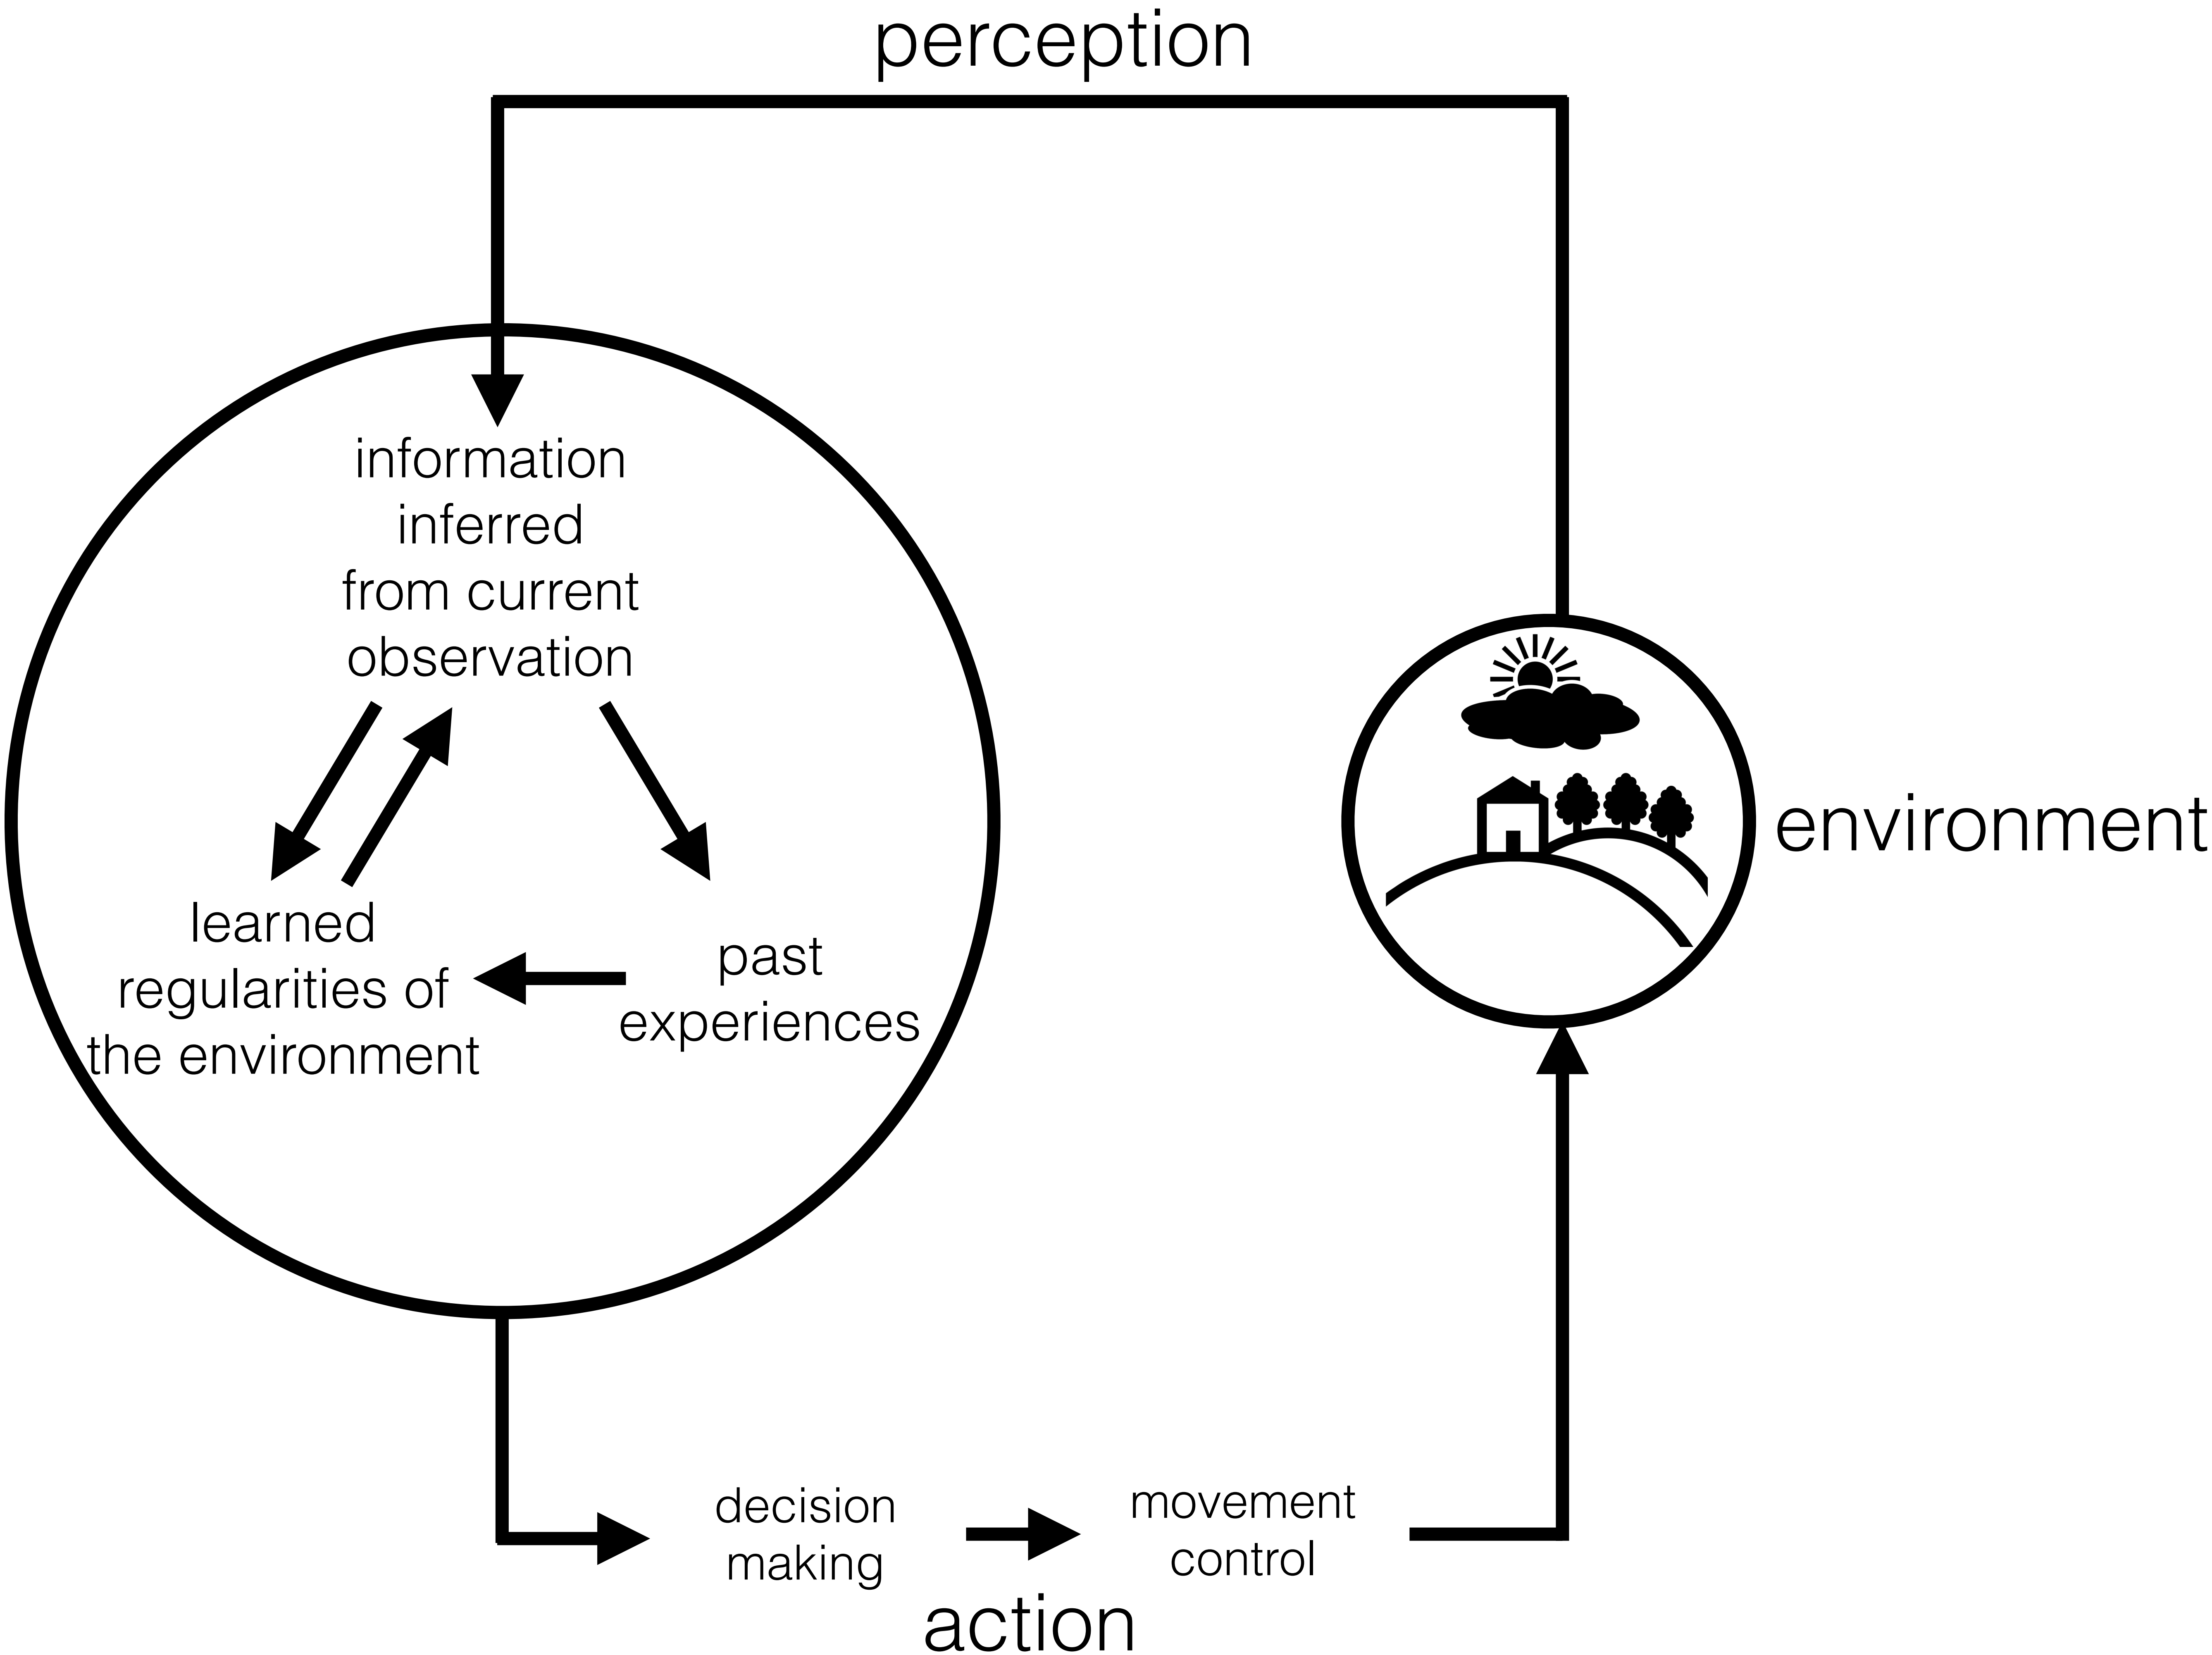
perception
information
inferred
from current
observation
learned
regularities of
the environment
past
experiences
environment
decision
making
movement
control
action

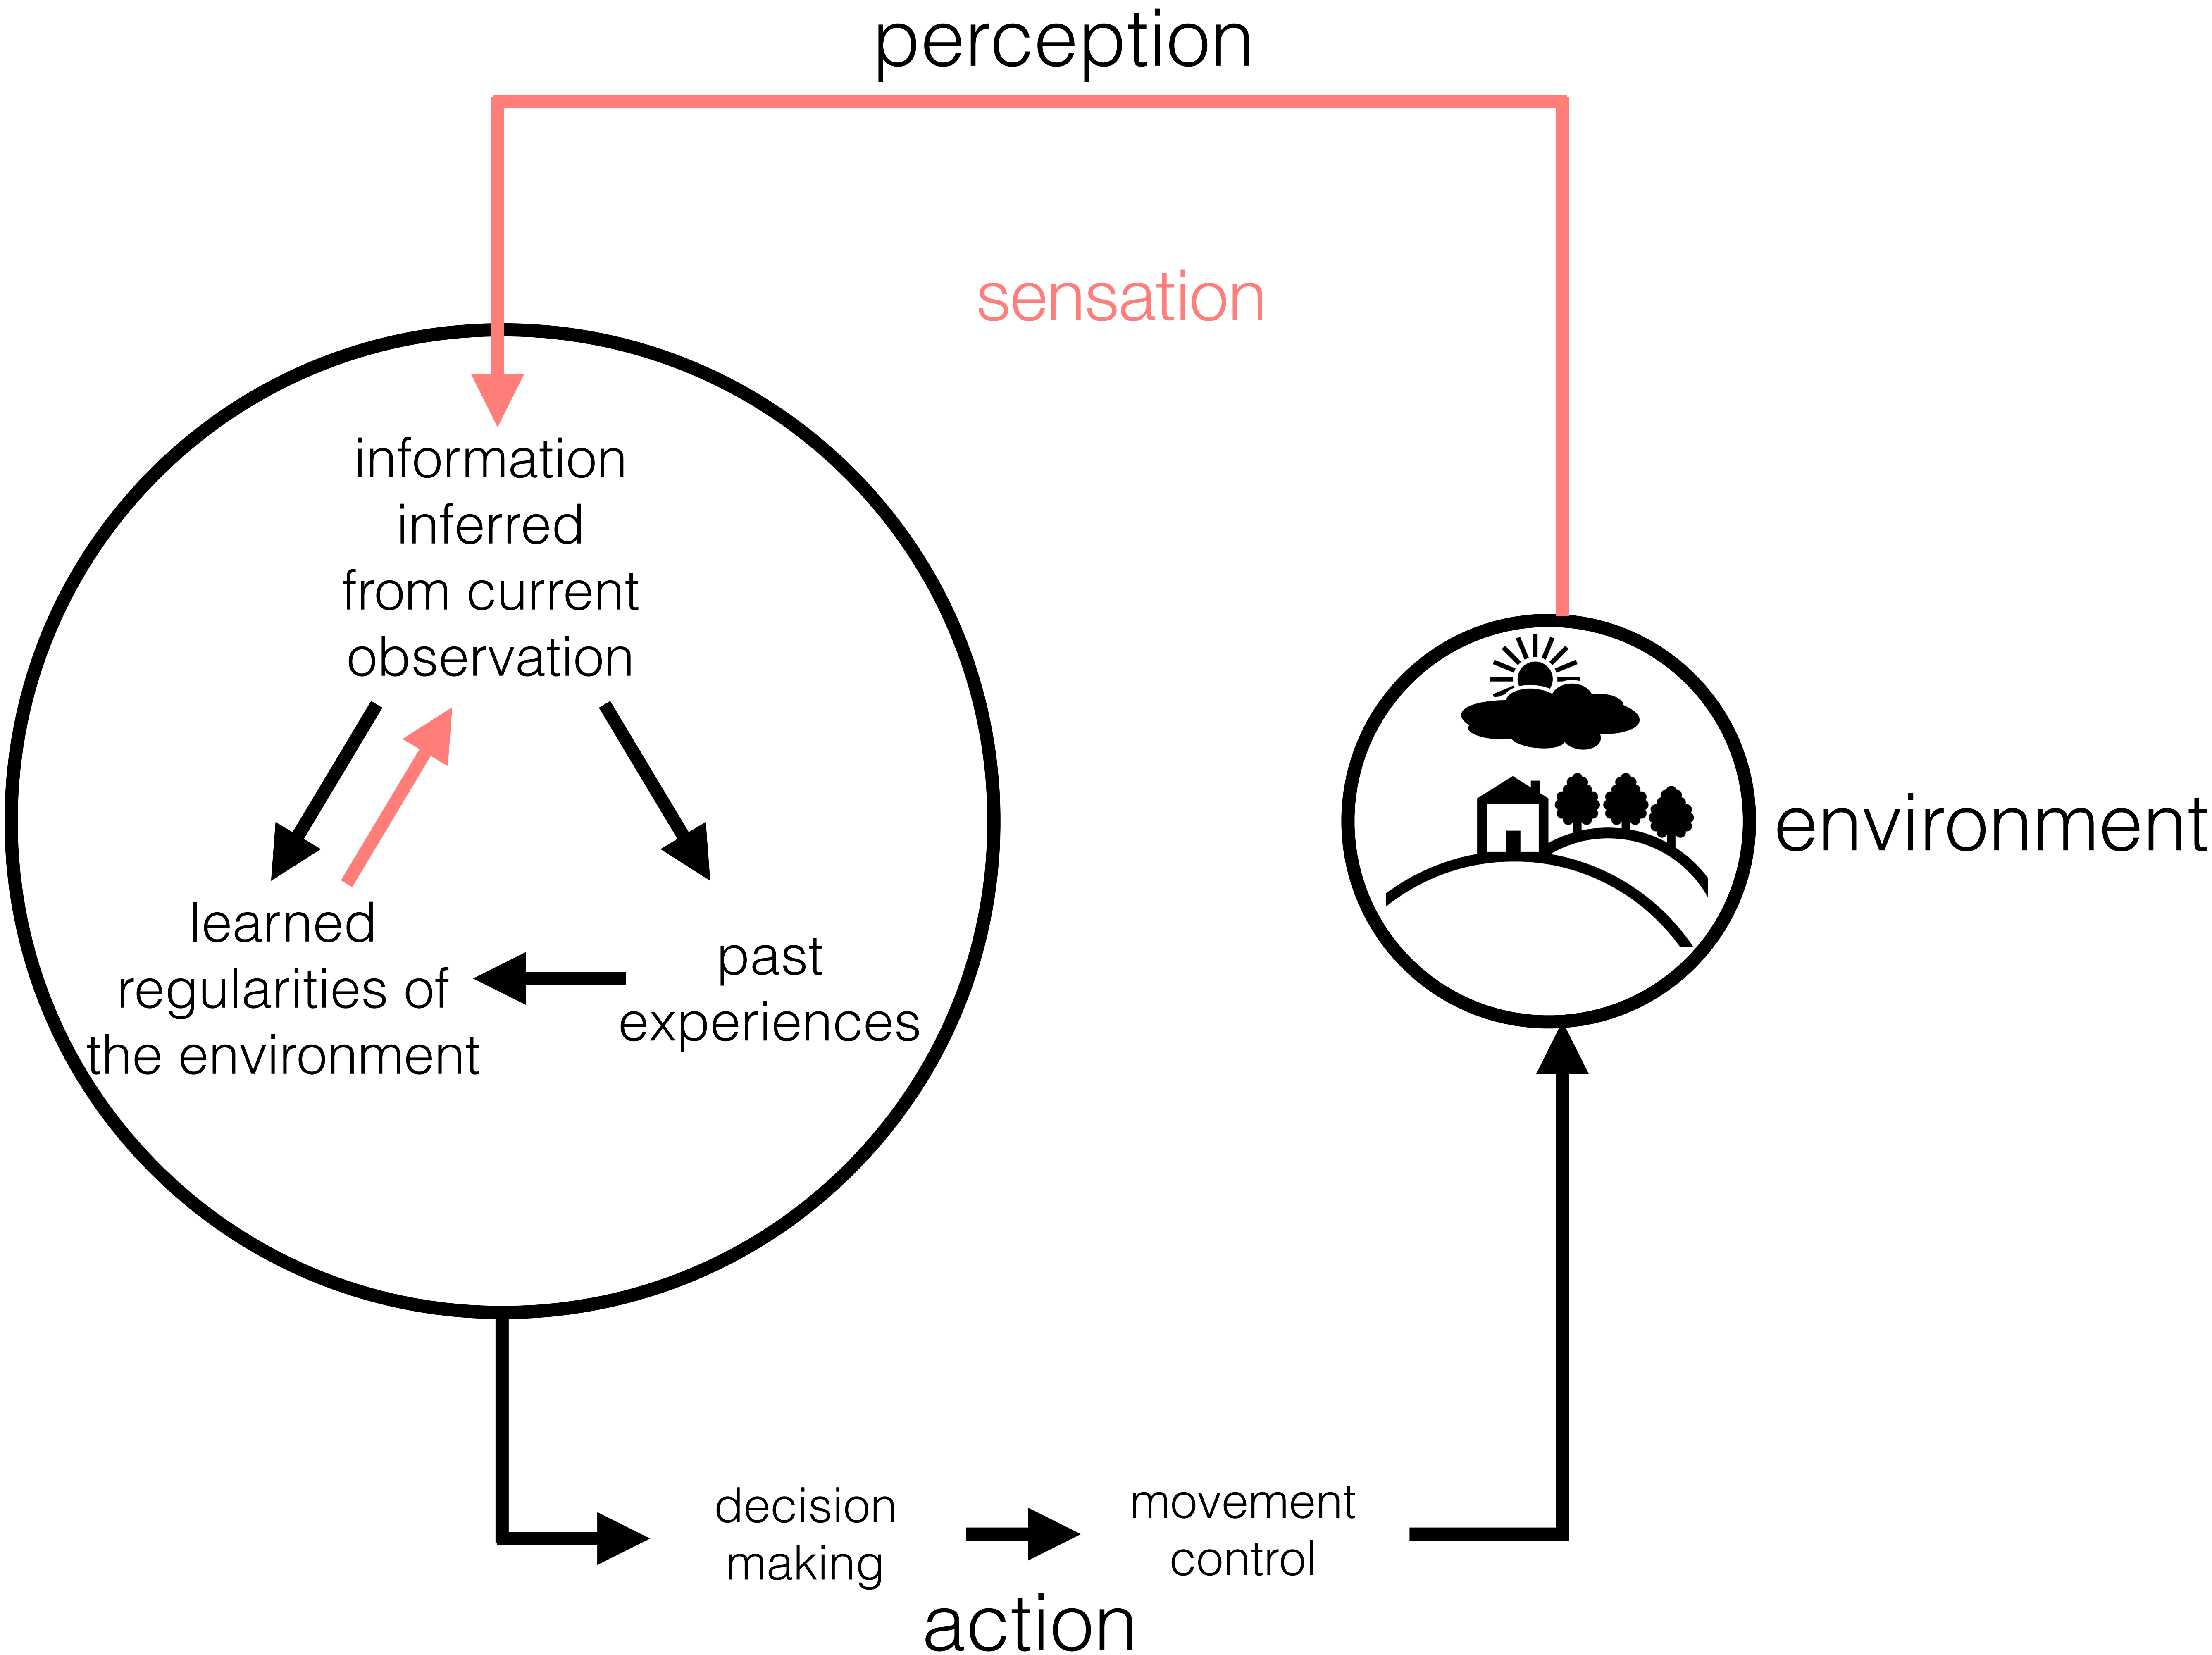
perception
sensation
information
inferred
from current
observation
learned
regularities of
the environment
past
experiences
environment
decision
making
movement
control
action

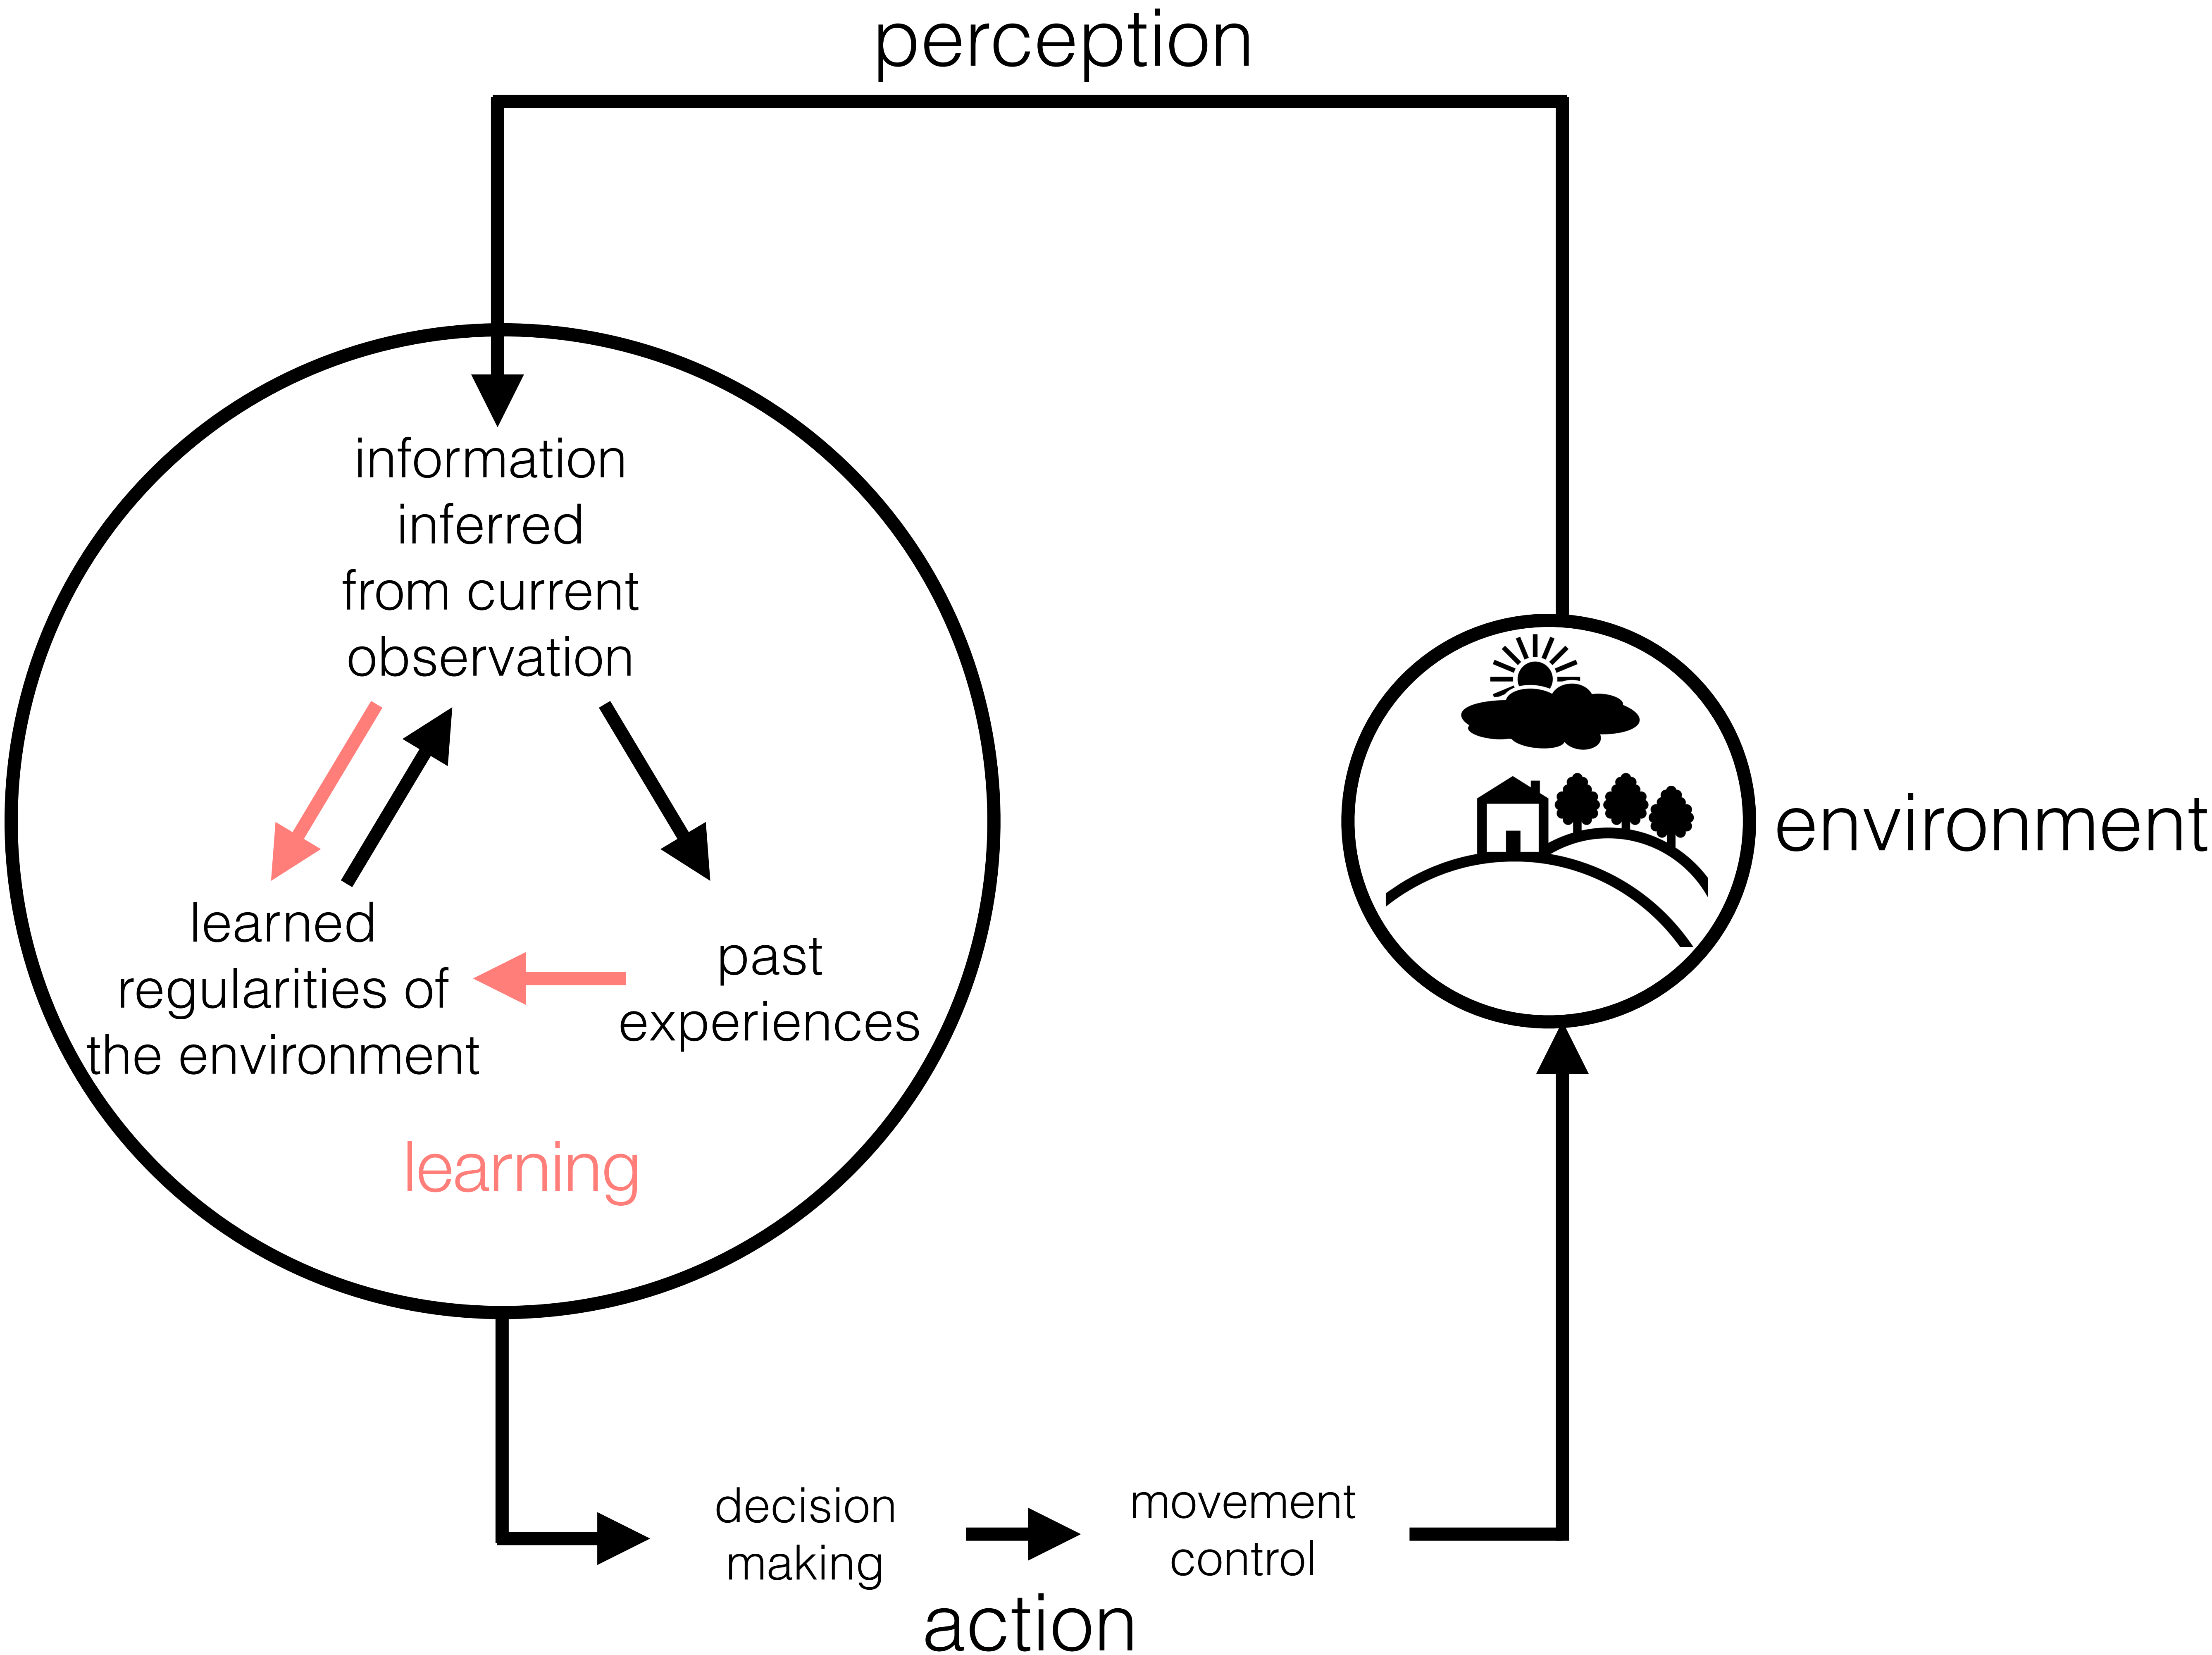
perception
information
inferred
from current
observation
learned
regularities of
the environment
past
experiences
learning
environment
decision
making
movement
control
action

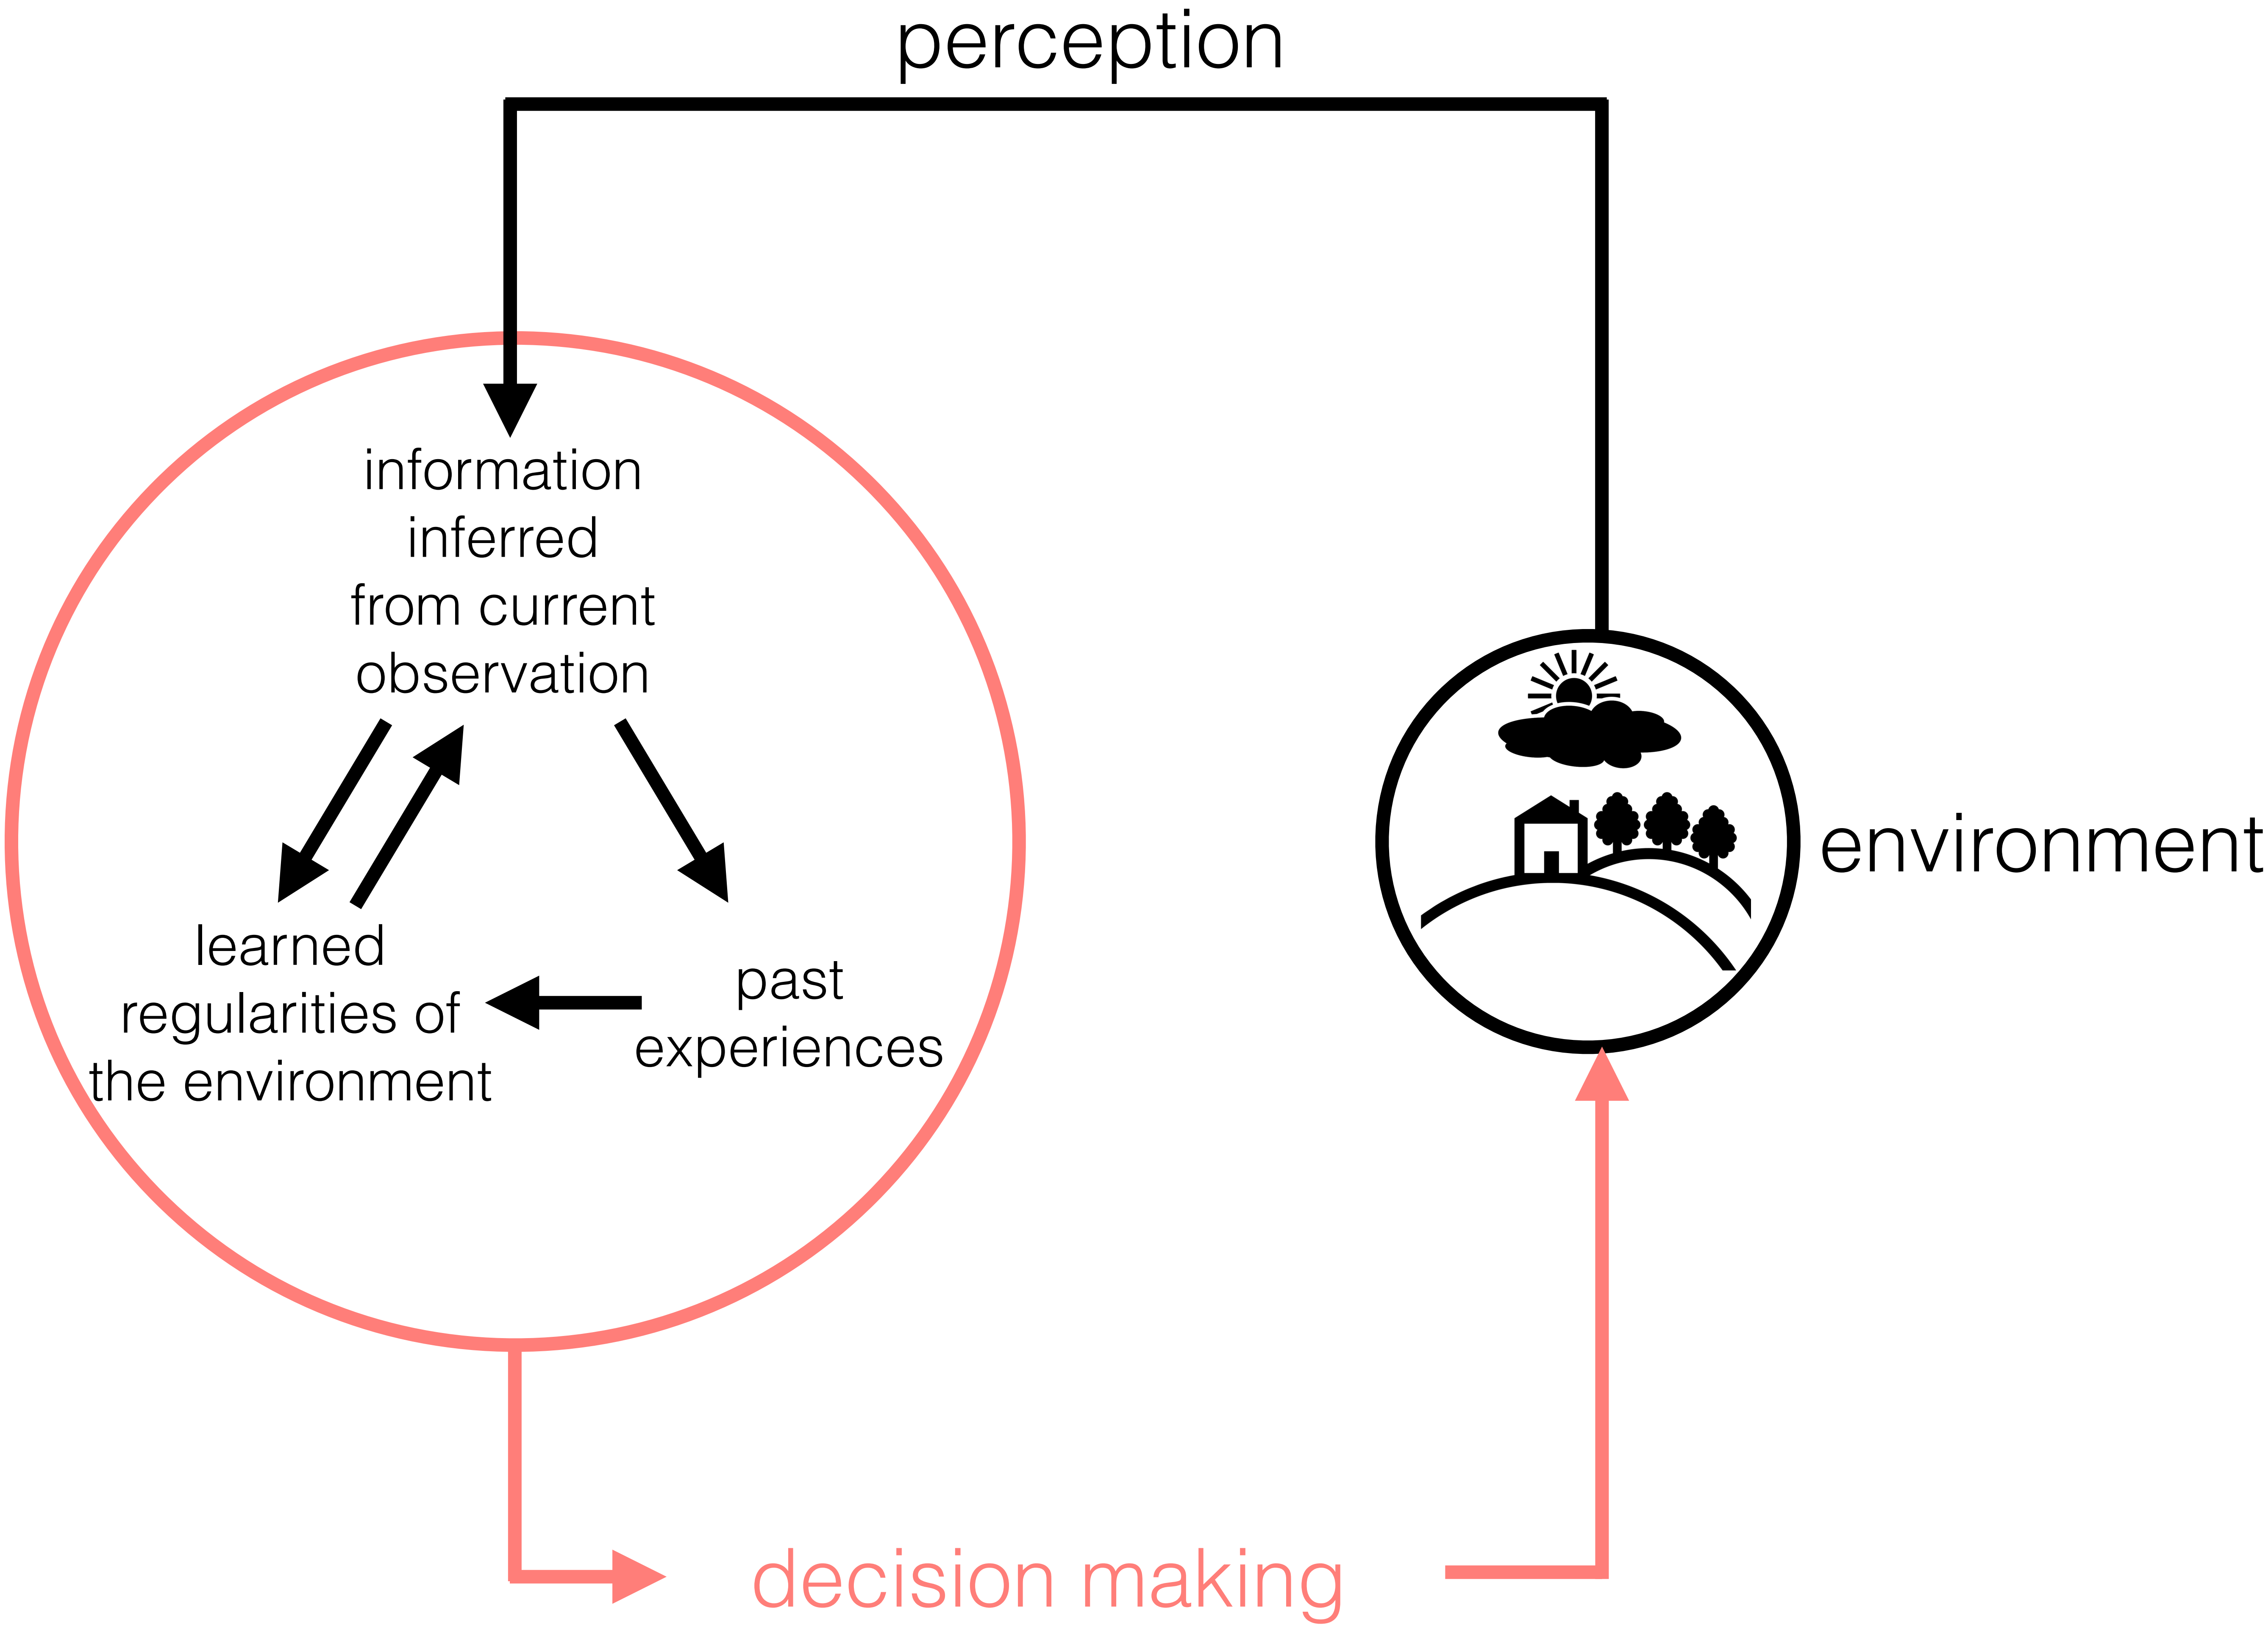
perception
information
inferred
from current
observation
learned
regularities of
the environment
past
experiences
environment
decision making

Artificial and biological intelligence

Brain computer analogy

- We can attempt to think of the brain as a digital computer

Artificial and biological intelligence

Brain computer analogy

- We can attempt to think of the brain as a digital computer
  - the hardware is the tissue, the software is its behavior

Artificial and biological intelligence

Brain computer analogy

- We can attempt to think of the brain as a digital computer
  - the hardware is the tissue, the software is its behavior
- These days software solve tasks similar to those of the human brain

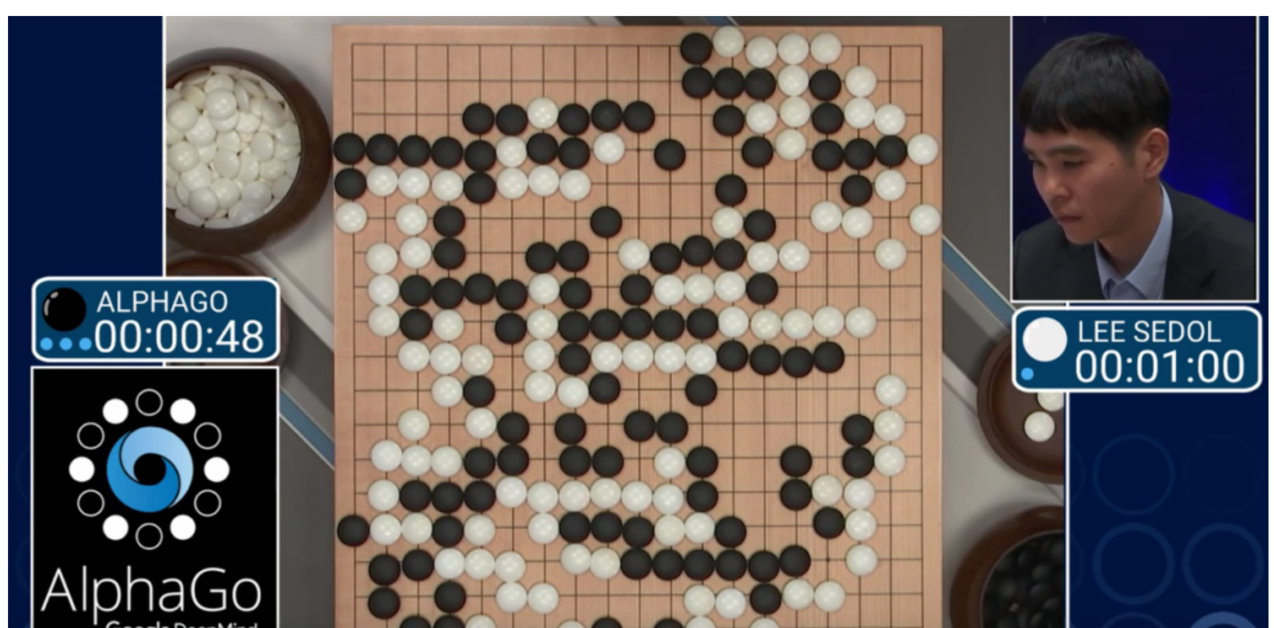

Artificial and biological intelligence
Brain computer analogy

- We can attempt to think of the brain as a digital computer
  - the hardware is the tissue, the software is its behavior
- These days software solve tasks similar to those of the human brain
  - recognition of objects on images

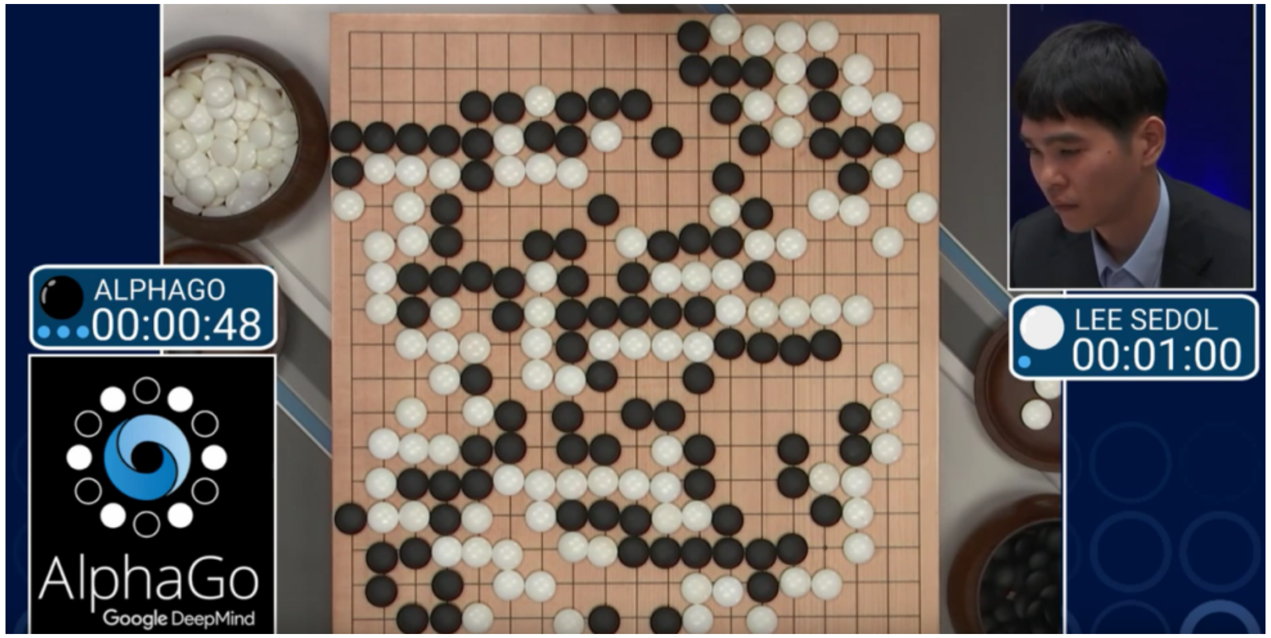

Artificial and biological intelligence

Brain computer analogy

- We can attempt to think of the brain as a digital computer
  - the hardware is the tissue, the software is its behavior
- These days software solve tasks similar to those of the human brain
  - recognition of objects on images
  - learning abstract categories

Artificial and biological intelligence

Brain computer analogy

- We can attempt to think of the brain as a digital computer
  - the hardware is the tissue, the software is its behavior
- These days software solve tasks similar to those of the human brain
  - recognition of objects on images
  - learning abstract categories
  - language understanding and synthesis

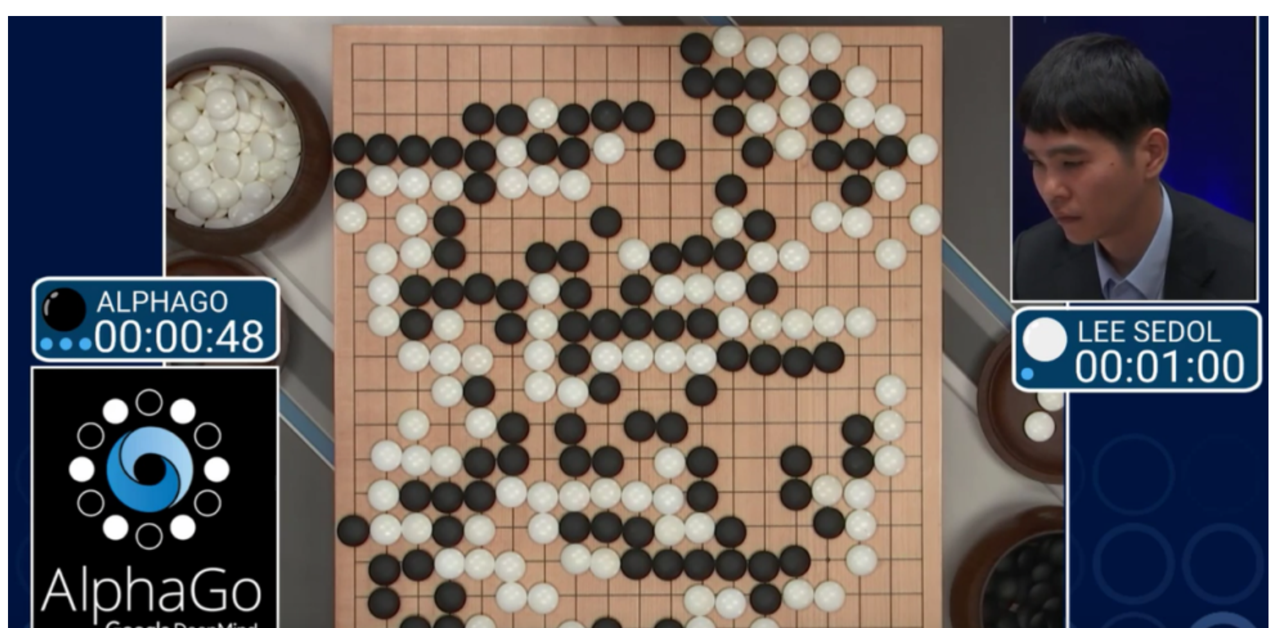

Artificial and biological intelligence
Brain computer analogy

- We can attempt to think of the brain as a digital computer
  - the hardware is the tissue, the software is its behavior
- These days software solve tasks similar to those of the human brain
  - recognition of objects on images
  - learning abstract categories
  - language understanding and synthesis
- Az analógia nem tökéletes, de sokat segít abban, hogy meghatározzuk azt az absztrakciós szintet, amin fel akarjuk tenni a kérdéseinket

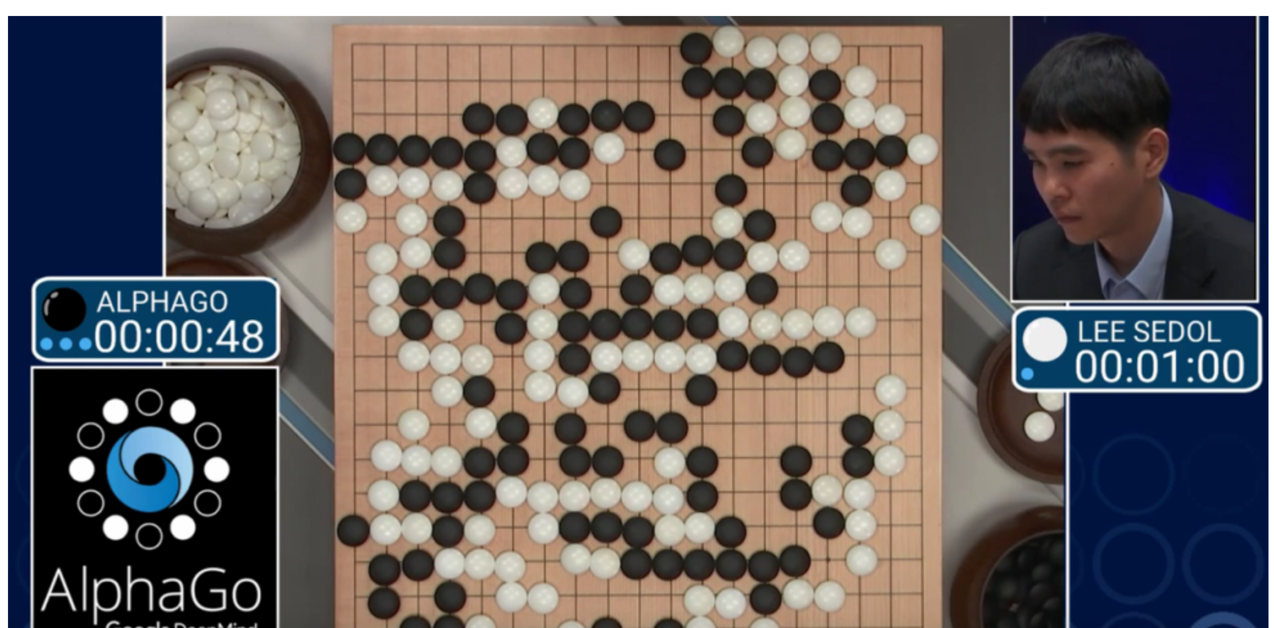

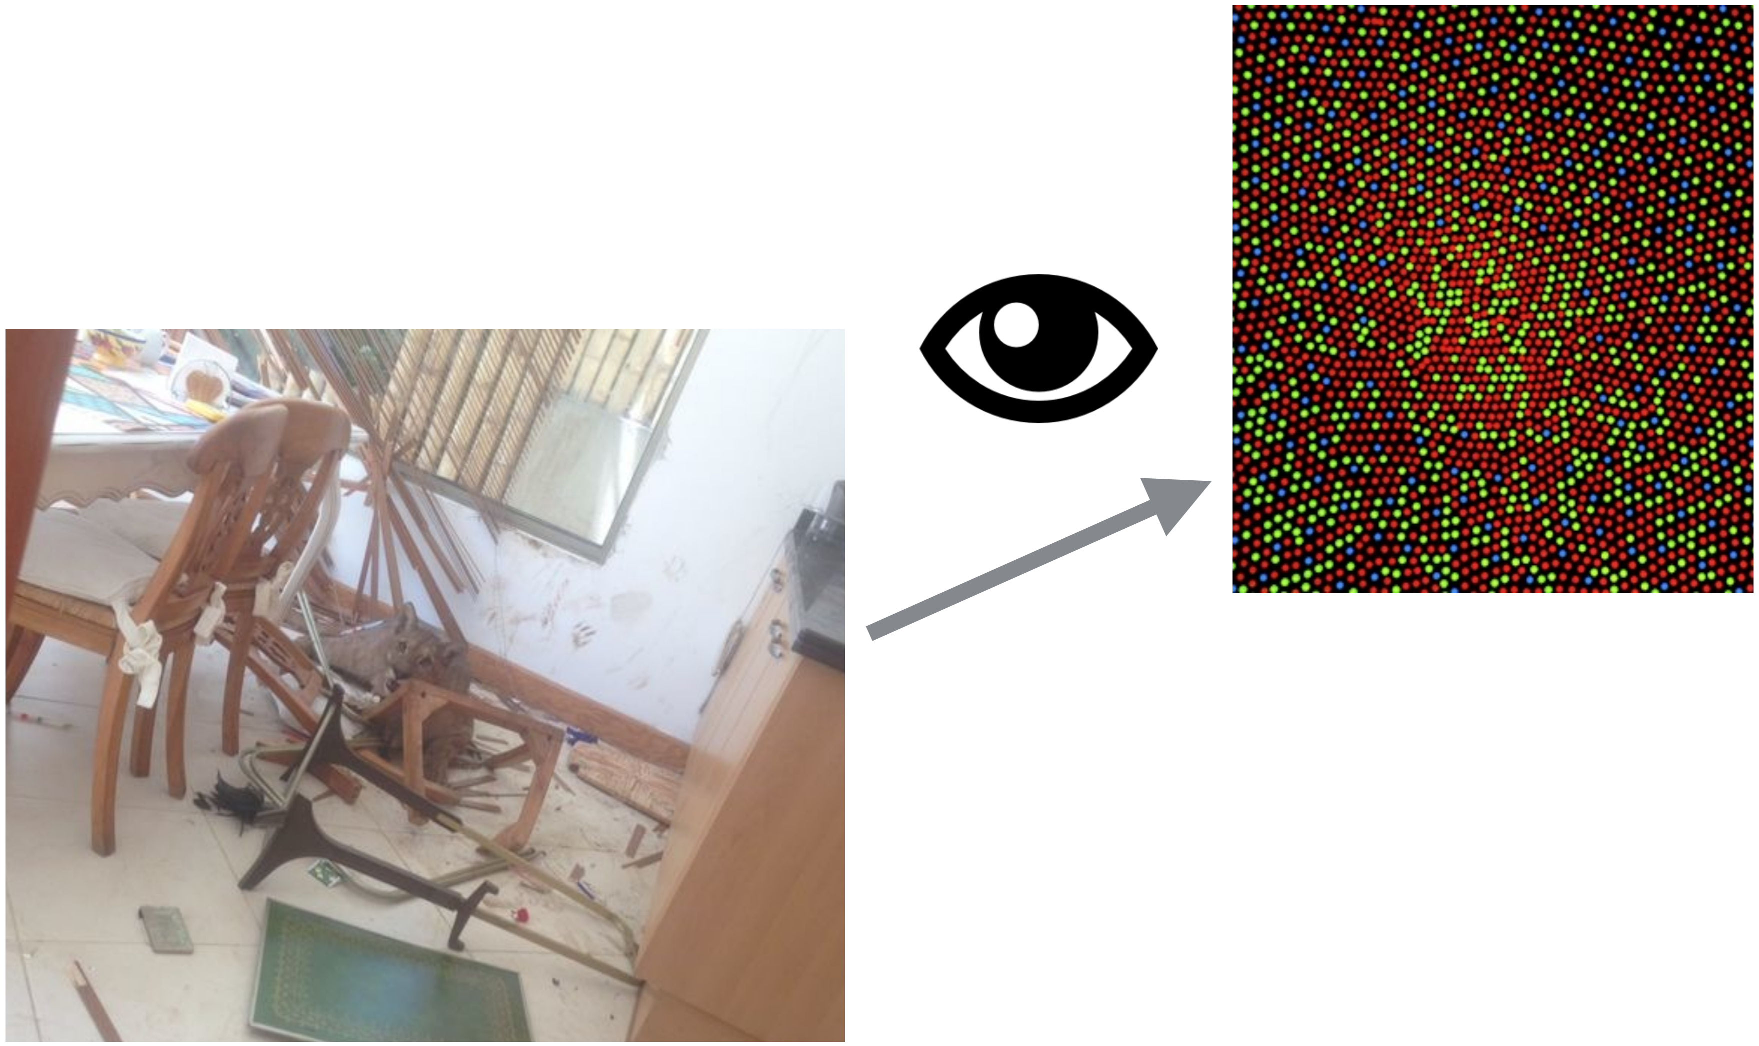

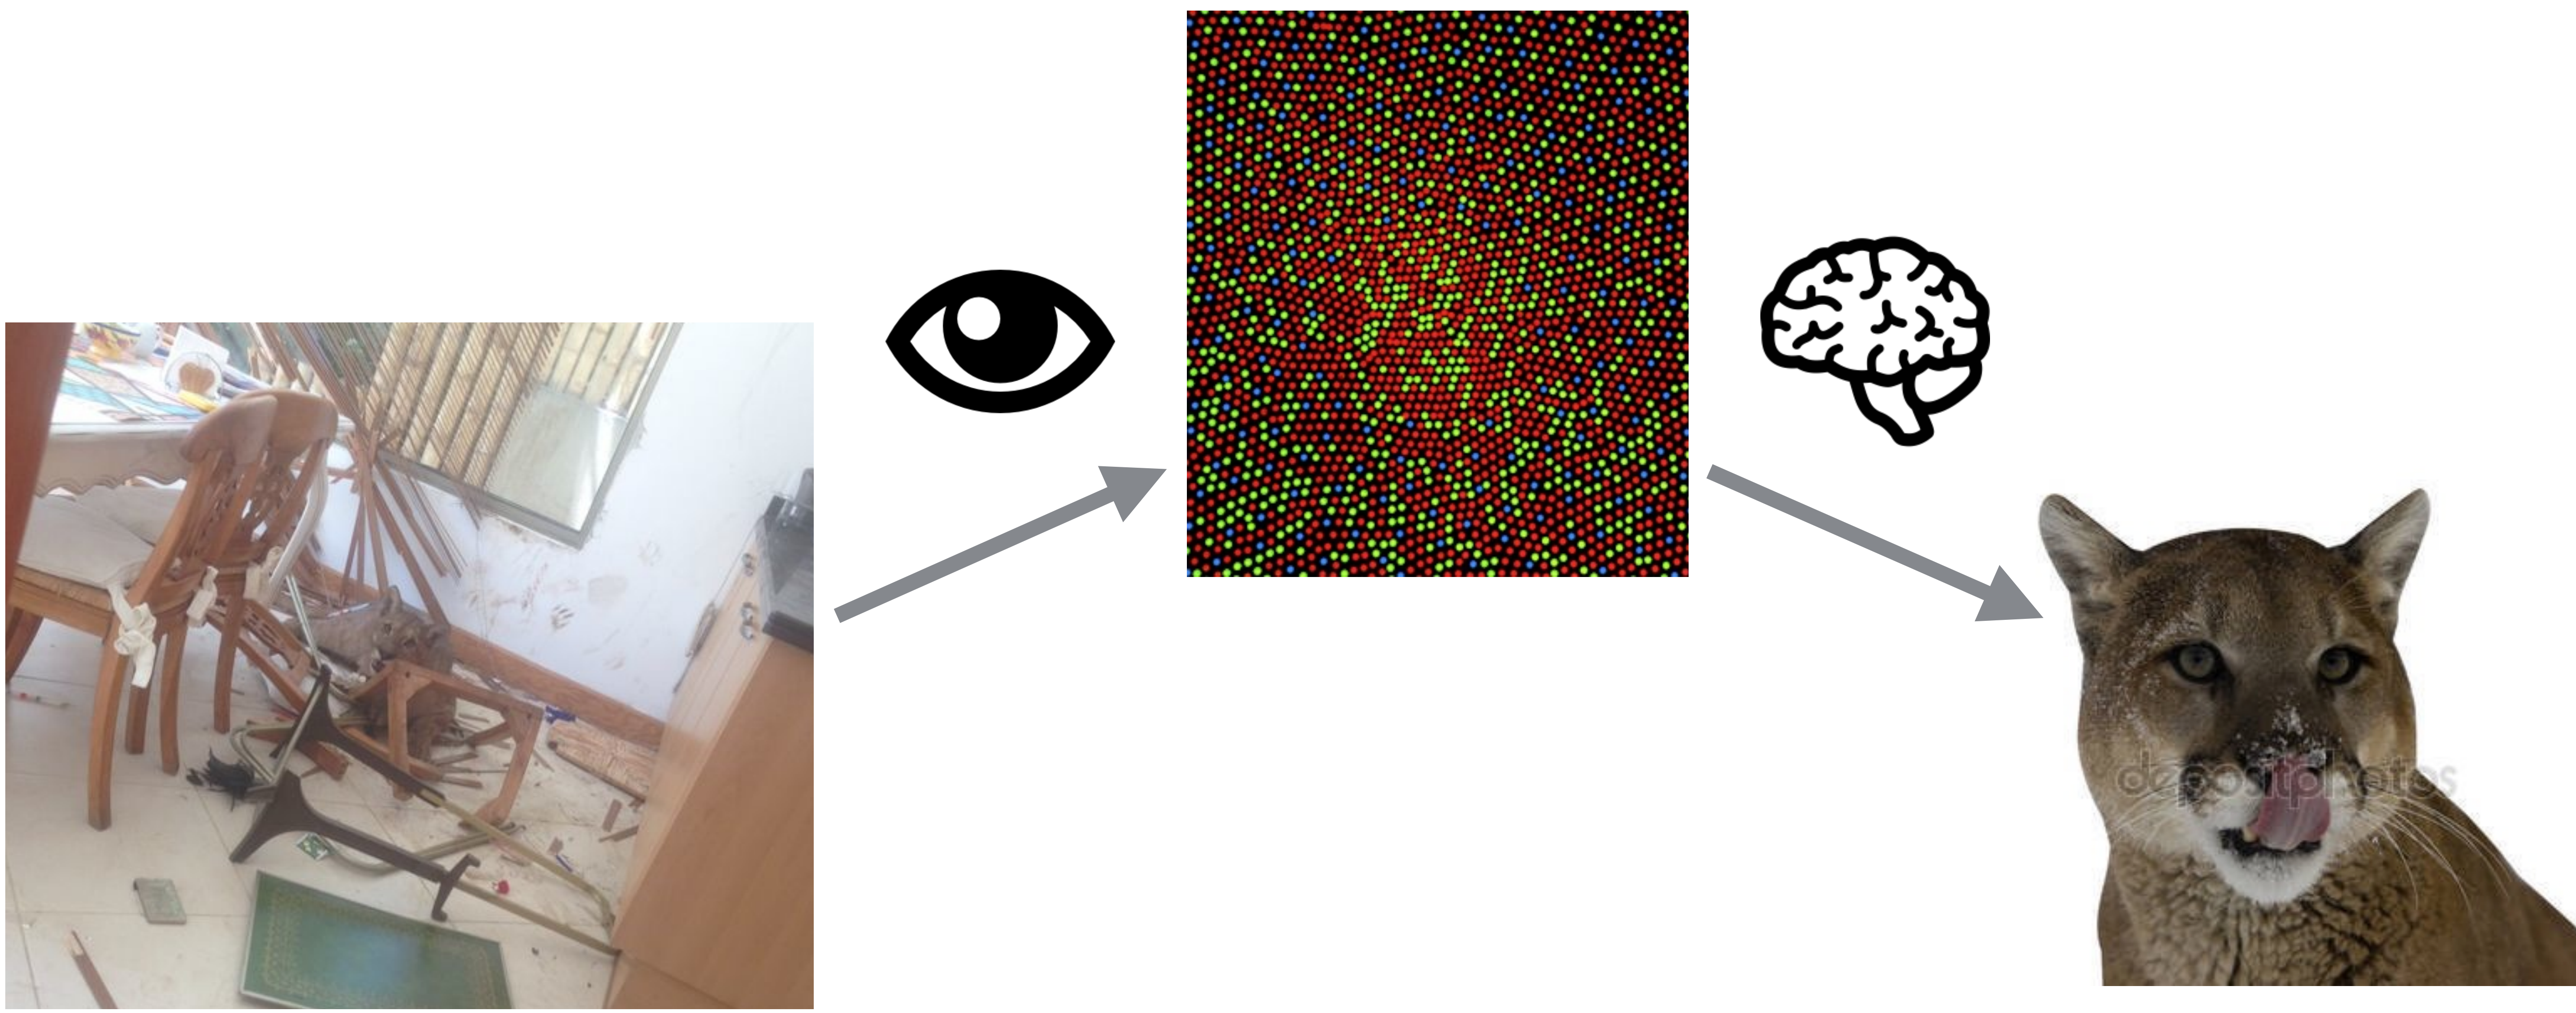

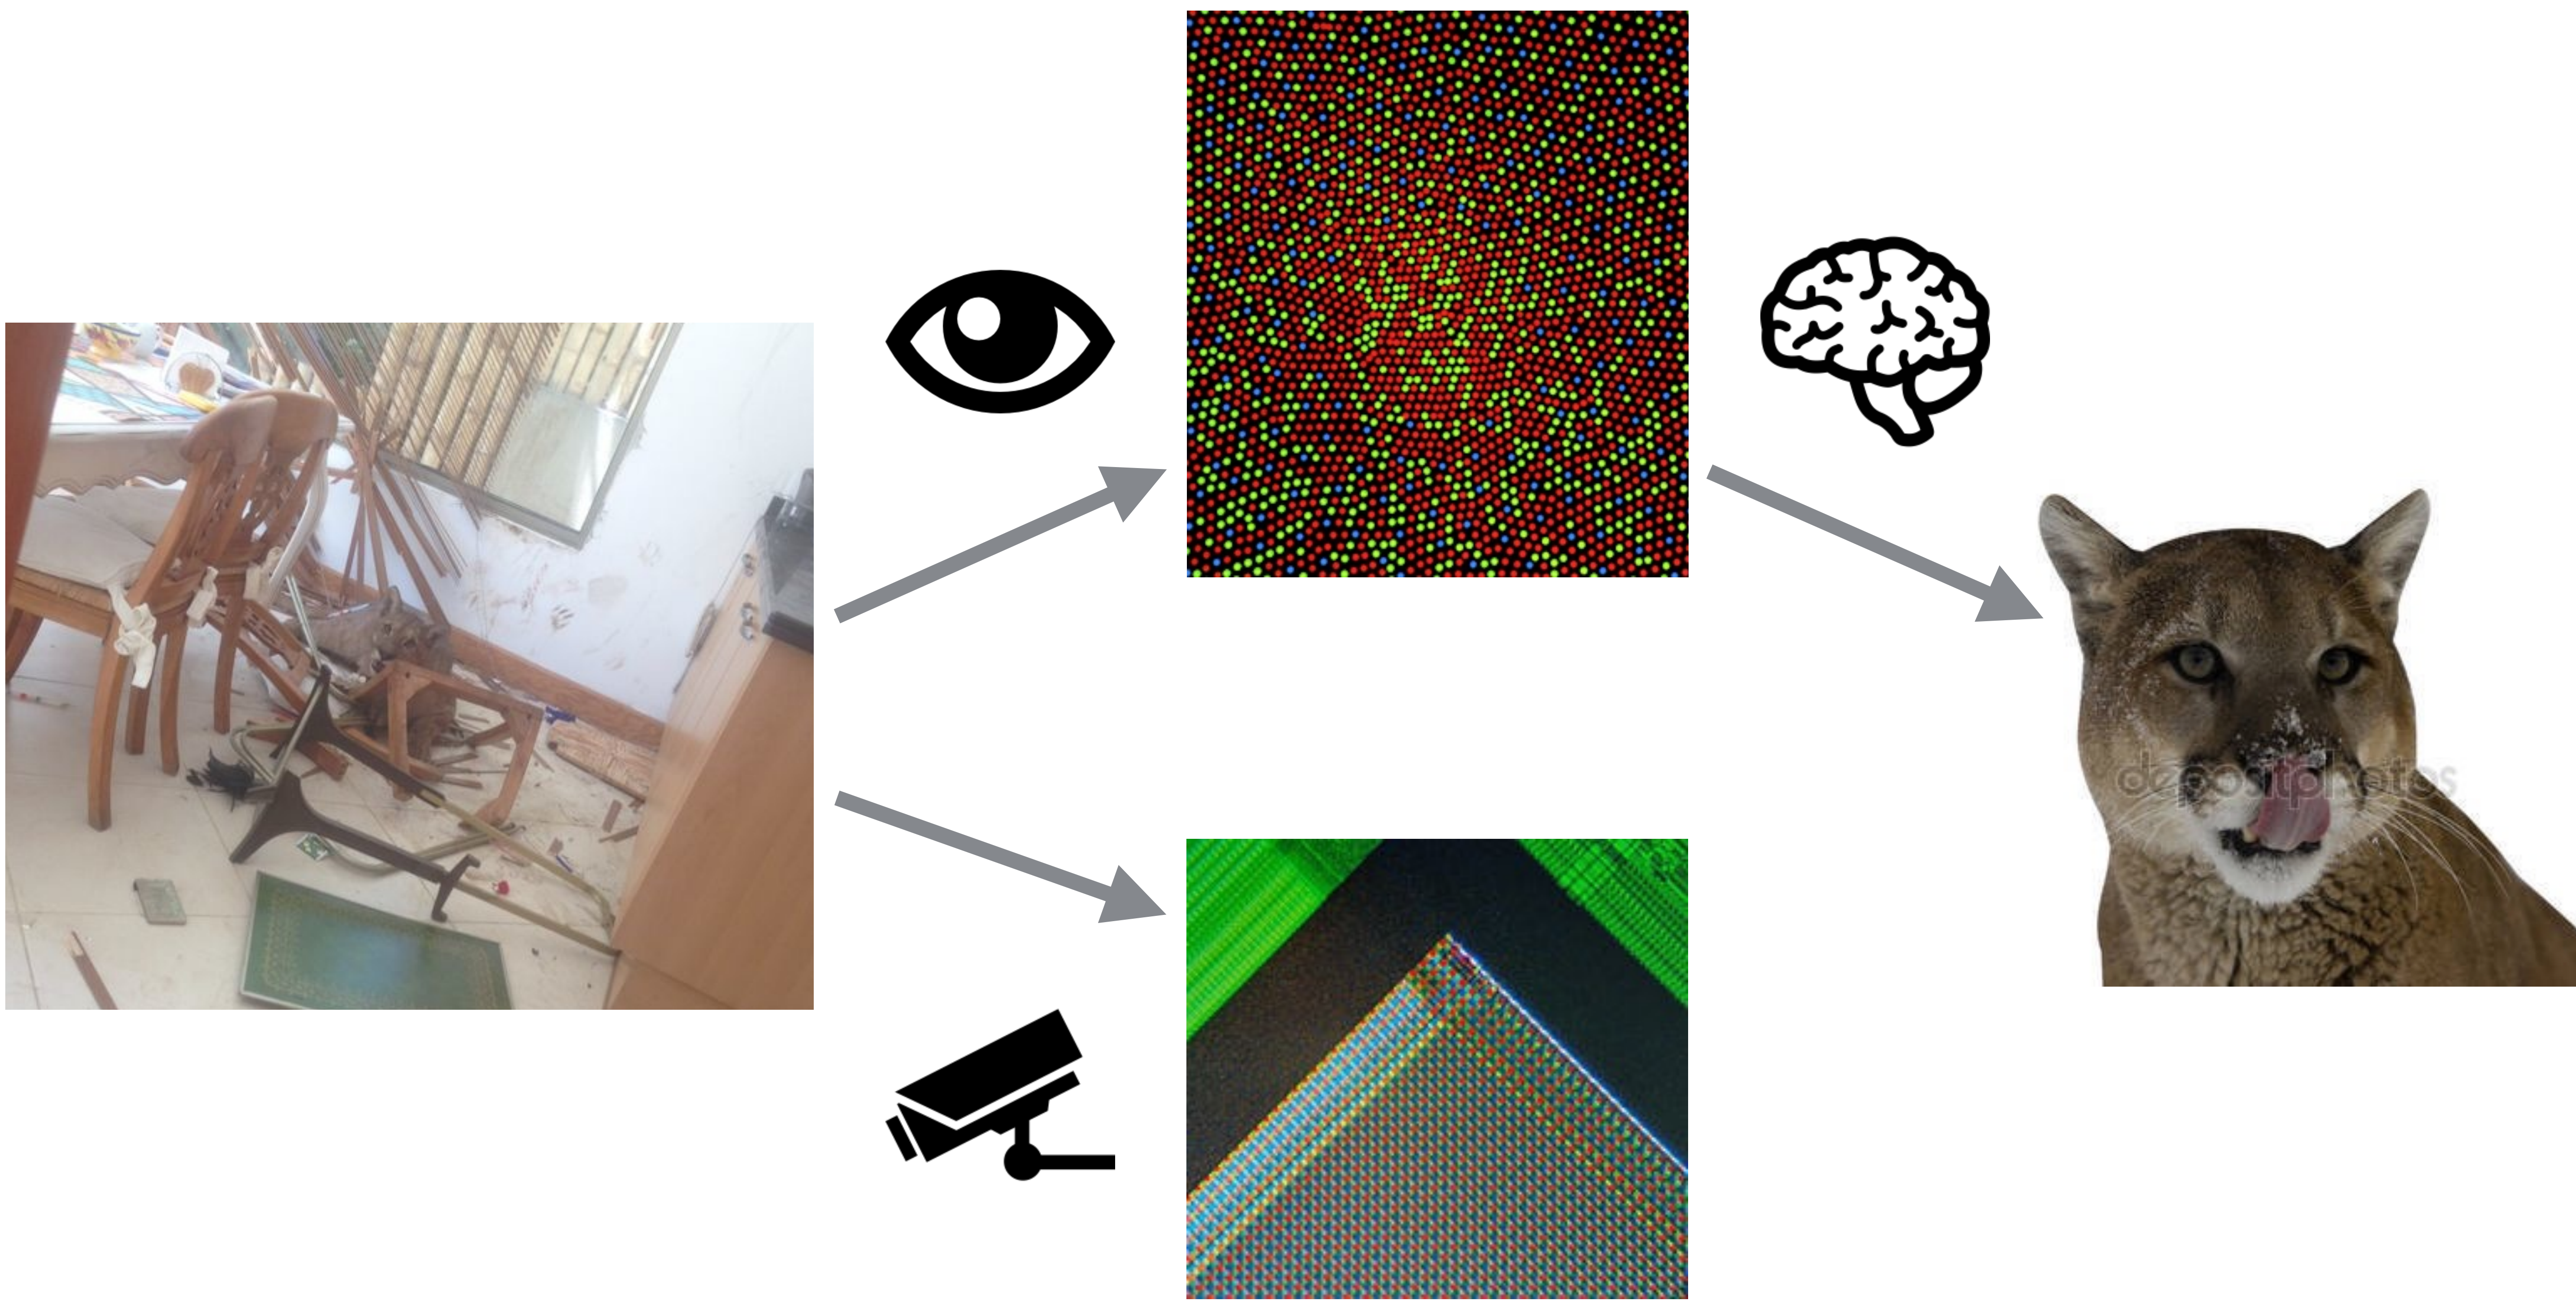

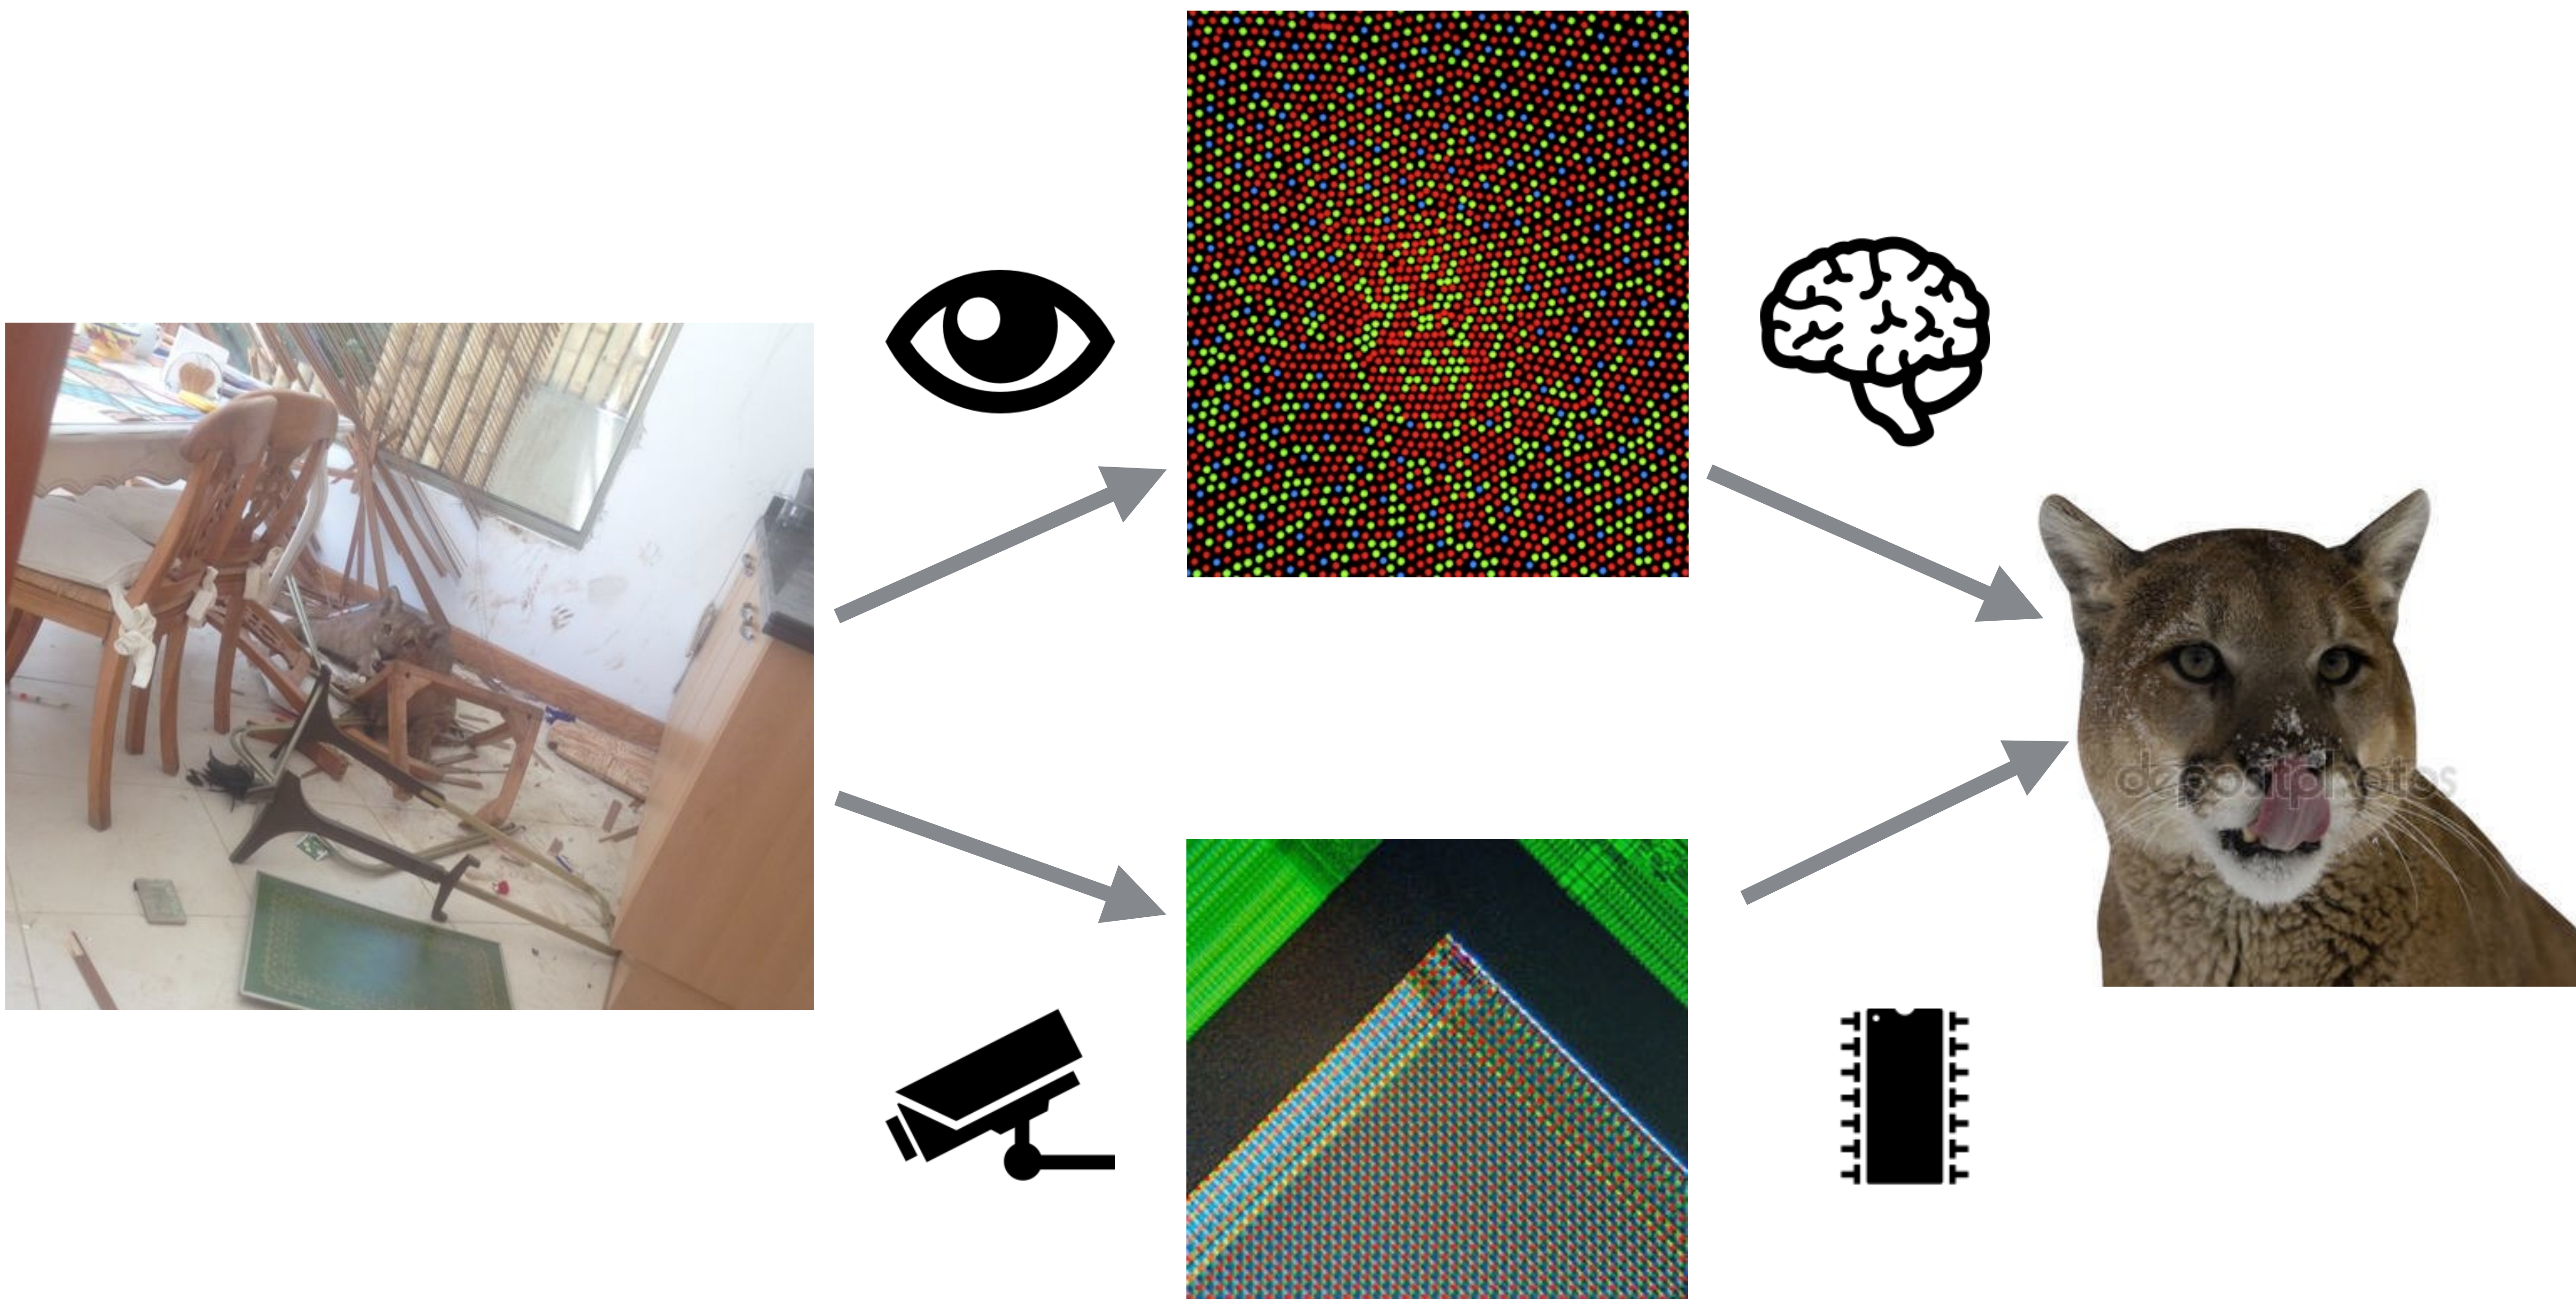

Brain-machine analogy

results in logics

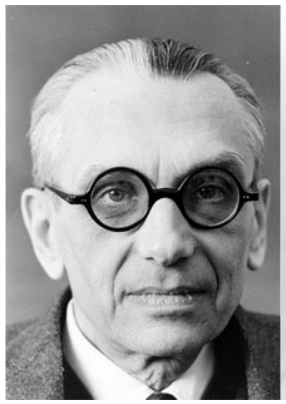 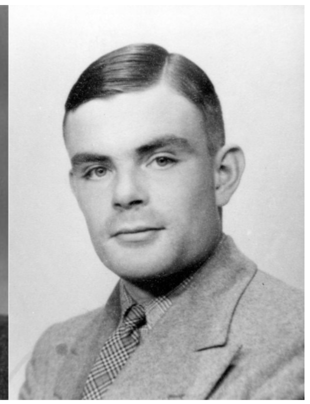

Brain-machine analogy

results in logics

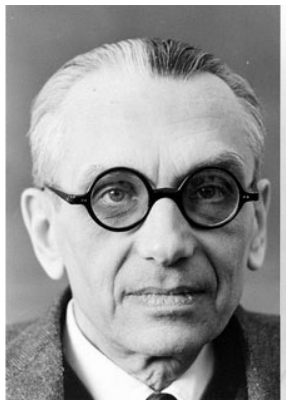

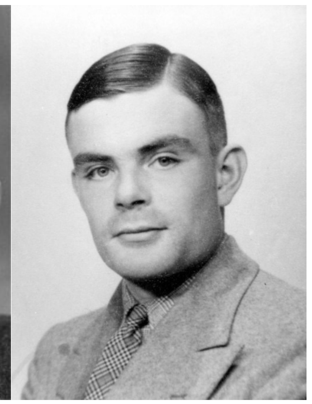

digital computers

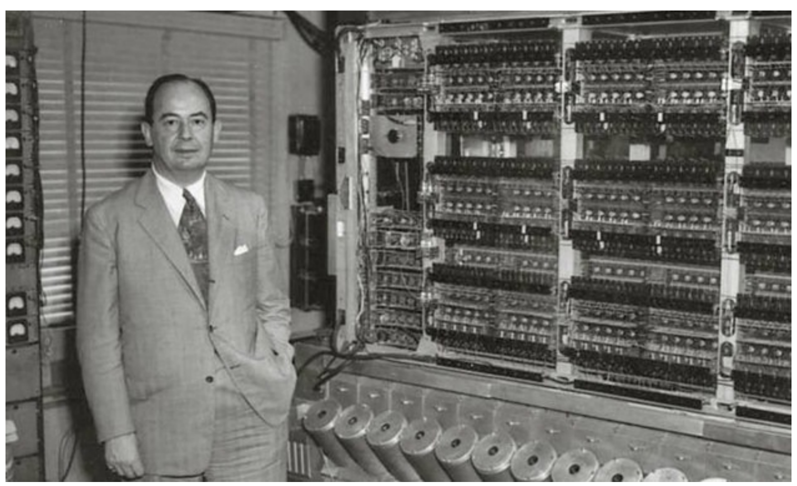

Brain-machine analogy

results in logics

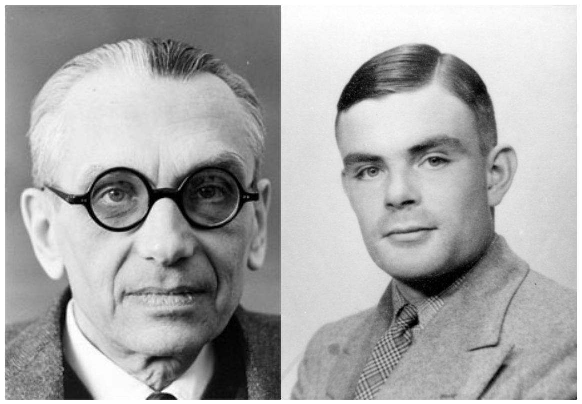

digital computers

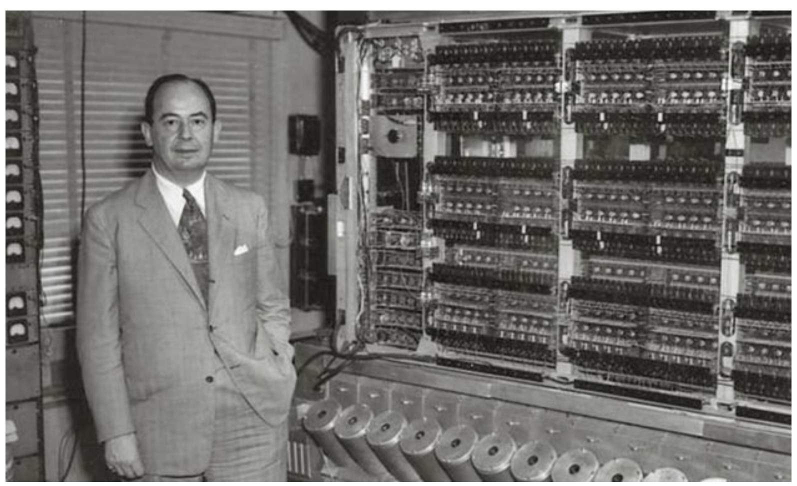

neuroscience experiments

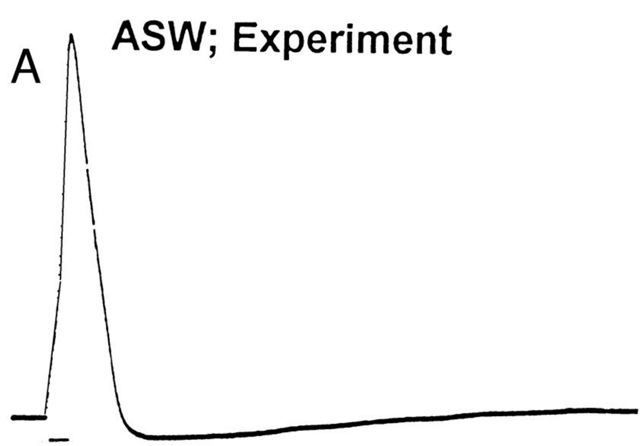

Brain-machine analogy

results in logics

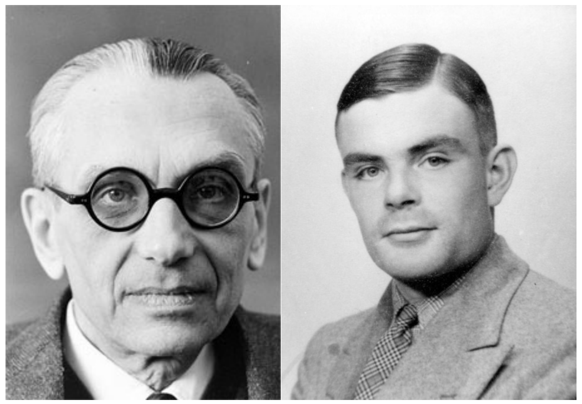

digital computers

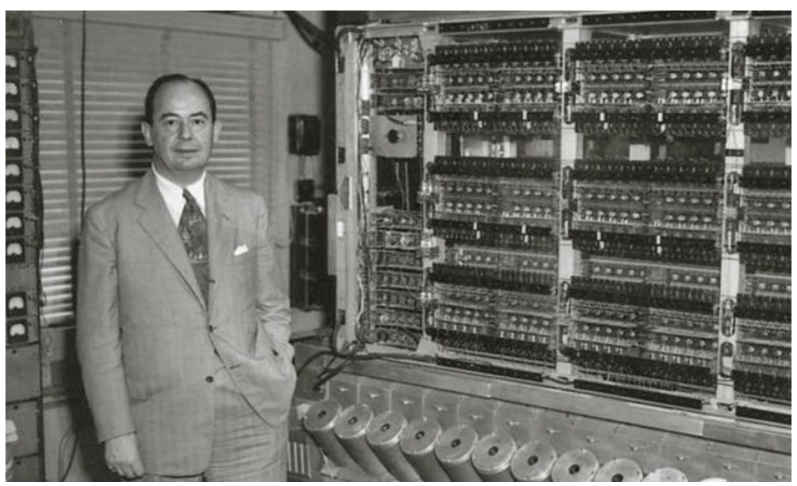

neuroscience experiments

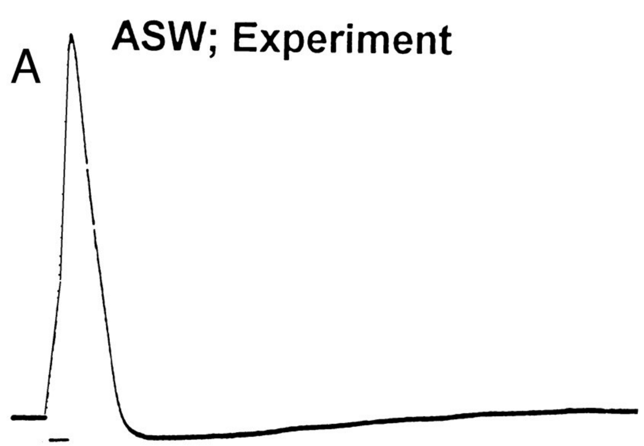

A LOGICAL CALCULUS OF THE IDEAS IMMANENT IN NERVOUS ACTIVITY

WARREN S. MCCULLOCH and WALTER H. PITTS

Because of the "all-or-none" character of nervous activity, neural events and the relations among them can be treated by means of propositional logic. It is found that the behavior of every net can be described in these terms, with the addition of more complicated logical means for nets containing circles; and that for any logical expression satisfying certain conditions, one can find a net behaving in the fashion it describes. It is shown that many particular choices among possible neurophysiological assumptions are equivalent, in the sense that for every net behaving under one assumption, there exists another net which behaves under the other and gives the same results, although perhaps not in the same time. Various applications of the calculus are discussed.

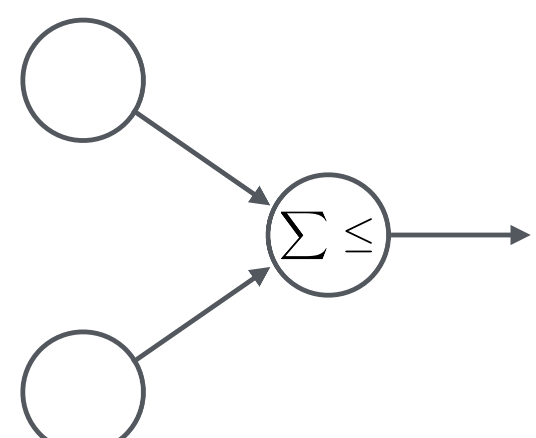

Artificial and biological intelligence

Computational inspiration

- **Computation** - determining the challenge:
  through the goals of the biological agent phrasing a mathematical model that is capable of addressing the challenge

- **Algorithm** - solving the challenge can be achieved in many different ways that can be weighed based on different factors

- **Implementation** - physical realization of the algorithm, in the case of neuroscience, based on neurons, spikes, etc

Artificial and biological intelligence

Computational inspiration

- **Computation** - determining the challenge:
  through the goals of the biological agent phrasing a mathematical model that is capable of addressing the challenge

- **Algorithm** - solving the challenge can be achieved in many different ways that can be weighed based on different factors

- **Implementation** - physical realization of the algorithm, in the case of neuroscience, based on neurons, spikes, etc

Artificial and biological intelligence

Computational inspiration

Artificial and biological intelligence

Computational inspiration

Neuron
Review

CellPress

Neuroscience-Inspired Artificial Intelligence

Demis Hassabis,[1,2,*] Dharshan Kumaran,[1,3] Christopher Summerfield,[1,4] and Matthew Botvinick[1,2]
[1]DeepMind, 5 New Street Square, London, UK
[2]Gatsby Computational Neuroscience Unit, 25 Howland Street, London, UK
[3]Institute of Cognitive Neuroscience, University College London, 17 Queen Square, London, UK
[4]Department of Experimental Psychology, University of Oxford, Oxford, UK
*Correspondence: dhcontact@google.com
http://dx.doi.org/10.1016/j.neuron.2017.06.011

The fields of neuroscience and artificial intelligence (AI) have a long and intertwined history. In more recent times, however, communication and collaboration between the two fields has become less commonplace. In this article, we argue that better understanding biological brains could play a vital role in building intelligent machines. We survey historical interactions between the AI and neuroscience fields and emphasize current advances in AI that have been inspired by the study of neural computation in humans and other animals. We conclude by highlighting shared themes that may be key for advancing future research in both fields.

If deep learning is the answer, what is the question?

Andrew Saxe, Stephanie Nelli and Christopher Summerfield

Abstract | Neuroscience research is undergoing a minor revolution. Recent advances in machine learning and artificial intelligence research have opened up new ways of thinking about neural computation. Many researchers are excited by the possibility that deep neural networks may offer theories of perception, cognition and action for biological brains. This approach has the potential to radically reshape our approach to understanding neural systems, because the computations performed by deep networks are learned from experience, and not endowed by the researcher. If so, how can neuroscientists use deep networks to model and understand biological brains? What is the outlook for neuroscientists who seek to characterize computations or neural codes, or who wish to understand perception, attention, memory and executive functions? In this Perspective, our goal is to offer a road map for systems neuroscience research in the age of deep learning. We discuss the conceptual and methodological challenges of comparing behaviour, learning dynamics and neural representations in artificial and biological systems, and we highlight new research questions that have emerged for neuroscience as a direct consequence of recent advances in machine learning.

NATURE REVIEWS | NEUROSCIENCE

nature neuroscience
FOCUS | PERSPECTIVE
https://doi.org/10.1038/s41593-019-0520-2

A deep learning framework for neuroscience

Blake A. Richards[1,2,3,4,42*], Timothy P. Lillicrap[5,6,42], Philippe Beaudoin[7], Yoshua Bengio[1,4,8], Rafal Bogacz[9], Amelia Christensen[10], Claudia Clopath[11], Rui Ponte Costa[12,13], Archy de Berker[7], Surya Ganguli[14,15], Colleen J. Gillon[16,17], Danijar Hafner[15,18,19], Adam Kepecs[20], Nikolaus Kriegeskorte[21,22], Peter Latham[23], Grace W. Lindsay[22,24], Kenneth D. Miller[22,24,25], Richard Naud[26,27], Christopher C. Pack[3], Panayiota Poirazi[28], Pieter Roelfsema[29], João Sacramento[30], Andrew Saxe[31], Benjamin Scellier[1,8], Anna C. Schapiro[32], Walter Senn[13], Greg Wayne[5], Daniel Yamins[33,34,35], Friedemann Zenke[36,37], Joel Zylberberg[4,38,39], Denis Therien[7,42] and Konrad P. Kording[4,40,41,42]

Systems neuroscience seeks explanations for how the brain implements a wide variety of perceptual, cognitive and motor tasks. Conversely, artificial intelligence attempts to design computational systems based on the tasks they will have to solve. In artificial neural networks, the three components specified by design are the objective functions, the learning rules and the architectures. With the growing success of deep learning, which utilizes brain-inspired architectures, these three designed components have increasingly become central to how we model, engineer and optimize complex artificial learning systems. Here we argue that a greater focus on these components would also benefit systems neuroscience. We give examples of how this optimization-based framework can drive theoretical and experimental progress in neuroscience. We contend that this principled perspective on systems neuroscience will help to generate more rapid progress.

REVIEW
doi:10.1038/nature14541

Probabilistic machine learning and artificial intelligence

Zoubin Ghahramani[1]

How can a machine learn from experience? Probabilistic modelling provides a framework for understanding what learning is, and has therefore emerged as one of the principal theoretical and practical approaches for designing machines that learn from data acquired through experience. The probabilistic framework, which describes how to represent and manipulate uncertainty about models and predictions, has a central role in scientific data analysis, machine learning, robotics, cognitive science and artificial intelligence. This Review provides an introduction to this framework, and discusses some of the state-of-the-art advances in the field, namely, probabilistic programming, Bayesian optimization, data compression and automatic model discovery.

Artificial and biological intelligence

Computational inspiration

Artificial and biological intelligence

Computational inspiration

on-line learning vs.
batch learnin

Artificial and biological intelligence

Computational inspiration

on-line learning vs.
batch learnin

approximate inferencia
(sampling/ variational
approaches)

Artificial and biological intelligence

Computational inspiration

approximate inferencia
(sampling/ variational
approaches)

on-line learning vs.
batch learnin

top-down interactions:
feed forward vs feedback

Artificial and biological intelligence

Computational inspiration

approximate inferencia
(sampling/ variational
approaches)

on-line learning vs.
batch learnin

episodic replay

top-down interactions:
feed forward vs feedback

Artificial and biological intelligence

Computational inspiration

approximate inferencia
(sampling/ variational
approaches)

on-line learning vs.
batch learnin

episodic replay

top-down interactions:
feed forward vs feedback

attention mechanisms

Artificial and biological intelligence

Computational inspiration

approximate inferencia
(sampling/ variational
approaches)

on-line learning vs.
batch learnin

episodic replay

top-down interactions:
feed forward vs feedback

attention mechanisms

generative comprssion

Artificial and biological intelligence

Computational inspiration

approximate inferencia
(sampling/ variational
approaches)

on-line learning vs.
batch learnin

episodic replay

top-down interactions:
feed forward vs feedback

attention mechanisms

generative comprssion

one-shot learning

Artificial and biological intelligence

Computational inspiration

approximate inferencia
(sampling/ variational
approaches)

on-line learning vs.
batch learnin

episodic replay

top-down interactions:
feed forward vs feedback

attention mechanisms

generative comprssion

one-shot learning

continual learning

Problems discussed

- Overview of the relevant functions and basic constituents of the of brain
- We attempt to characterize the computations that the brain needs to implement
- During the course we will mainly focus an challenges in perception
  - ez nagyon jelentős része az agyi funkcióknak: minden, ami az érzékszervi bemeneteket tudásra képezi le, beleértve a jelfeldolgozást, memóriaformációt, tanulást, nyelvi feldolgozást, stb.
  - we might touch upon decision making as well
  - according to normative modelling, what we need now is a mathematical toolset to define the knowledge we acquire from interacting with the environment
- once we have a proper mathematical framework we attempt to do two things
  - we test if it predicts the bahavior of humans and other animals in various settings
  - we build predictive models of biophysical quantities and check them on neuronal data

Artificial intelligence and neuroscience

Artificial intelligence and neuroscience

- Common roots in cybernetics
  - Warren McCulloch, Walter Pitts, Norbert Wiener és Neumann János
  - Building logic based on neuron-like units

Artificial intelligence and neuroscience

- Common roots in cybernetics
  - Warren McCulloch, Walter Pitts, Norbert Wiener és Neumann János
  - Building logic based on neuron-like units
- Analogous problems in the two disciplines
  - AI: hogy how can we build the best tool?
  - Neurosci: how do biological agents do it?

Artificial intelligence and neuroscience

- Common roots in cybernetics
  - Warren McCulloch, Walter Pitts, Norbert Wiener és Neumann János
  - Building logic based on neuron-like units
- Analogous problems in the two disciplines
  - AI: hogy how can we build the best tool?
  - Neurosci: how do biological agents do it?
- Kölcsönös inspirációt nyújtanak egymásnak, például
  - Neurosci -> AI: methods in machine learning building on neural networks
  - AI -> Neurosci: aligning probabilistic models of perception with the computations of neuronal networks

Directions in computational neuroscience

Directions in computational neuroscience

- Biophysical modelling

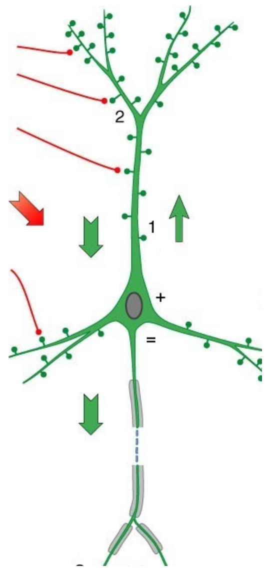

Directions in computational neuroscience

- Biophysical modelling
- Spiking networks
    - Balanced E/I networks

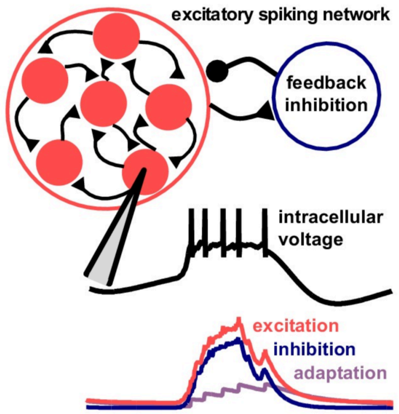

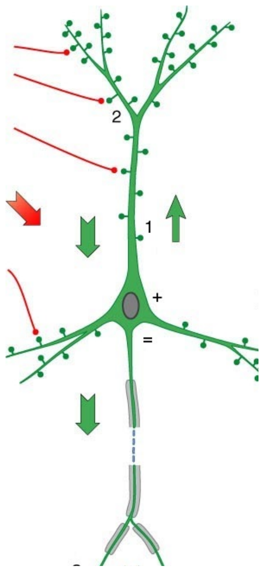

Directions in computational neuroscience

- Biophysical modelling
- Spiking networks
  - Balanced E/I networks
- Connectionist networks
  - Deep learning

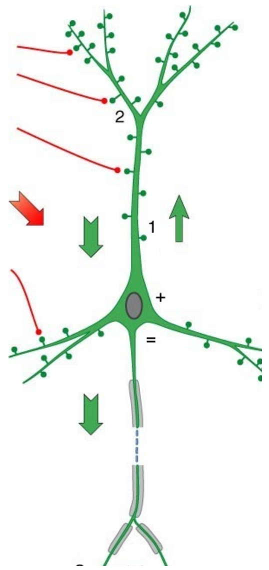

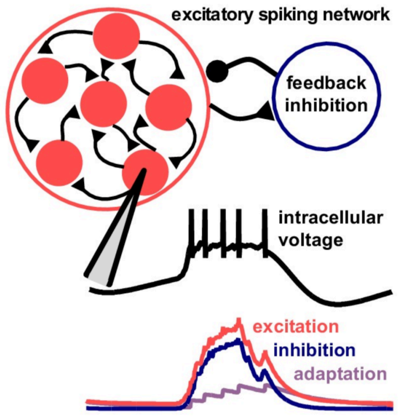

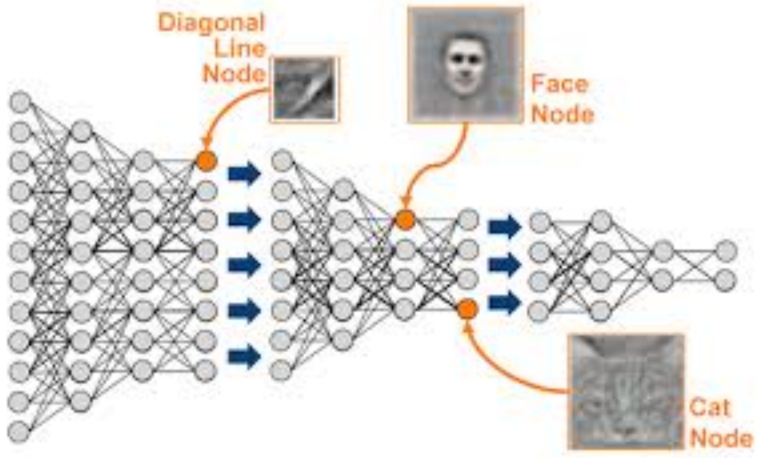

Directions in computational neuroscience

- Biophysical modelling
- Spiking networks
  - Balanced E/I networks
- Connectionist networks
  - Deep learning
- Systems neuroscience

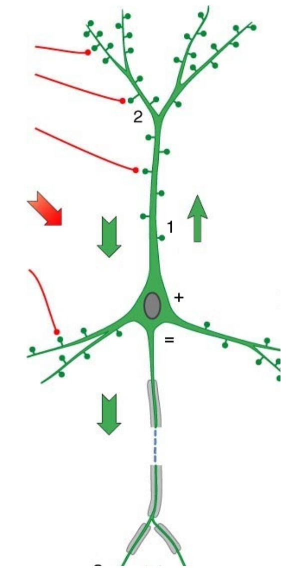

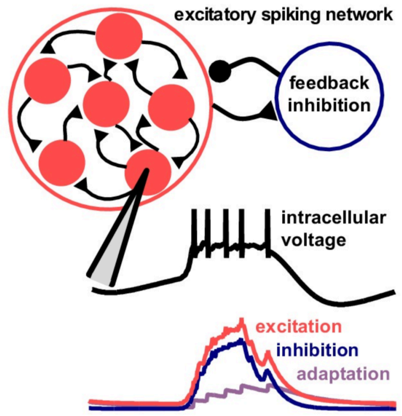

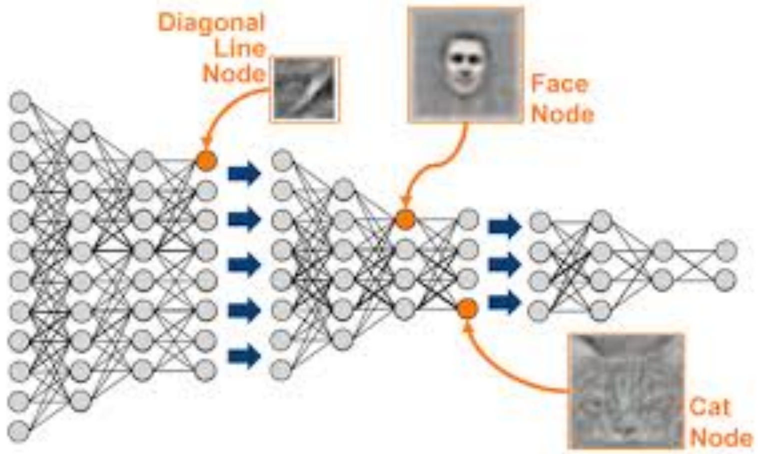

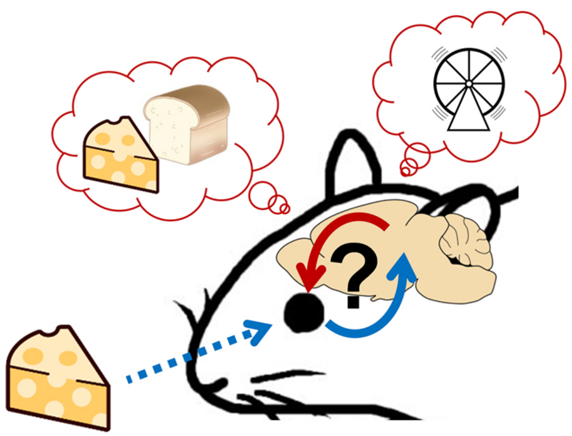

Directions in computational neuroscience

- Biophysical modelling
- Spiking networks
    - Balanced E/I networks
- Connectionist networks
    - Deep learning
- Systems neuroscience
- Cognitive neurosience

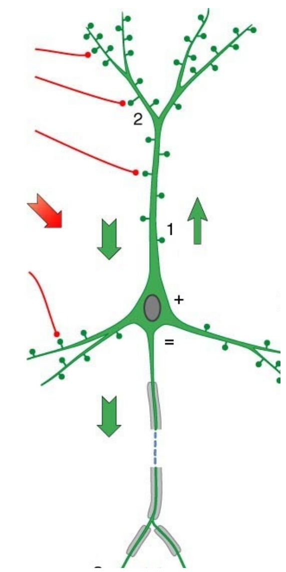

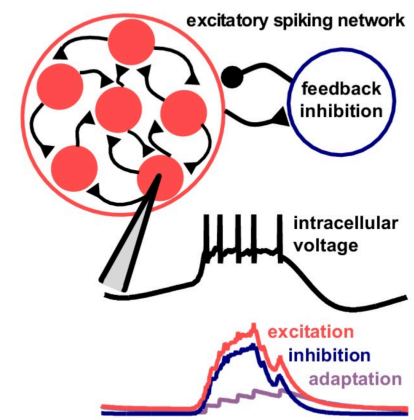

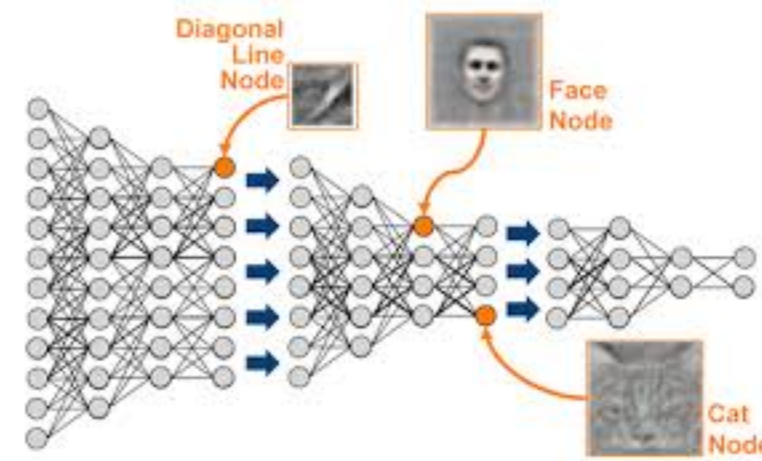

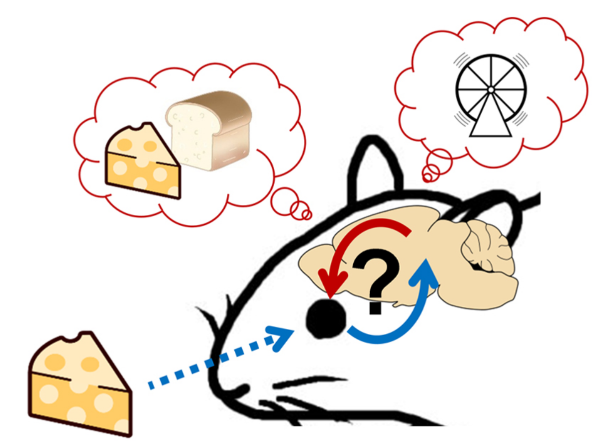

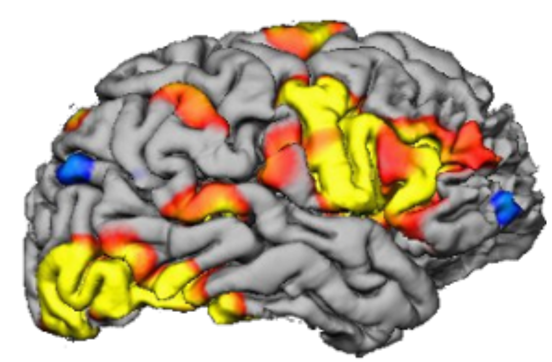

What are the relevant components of the brain?

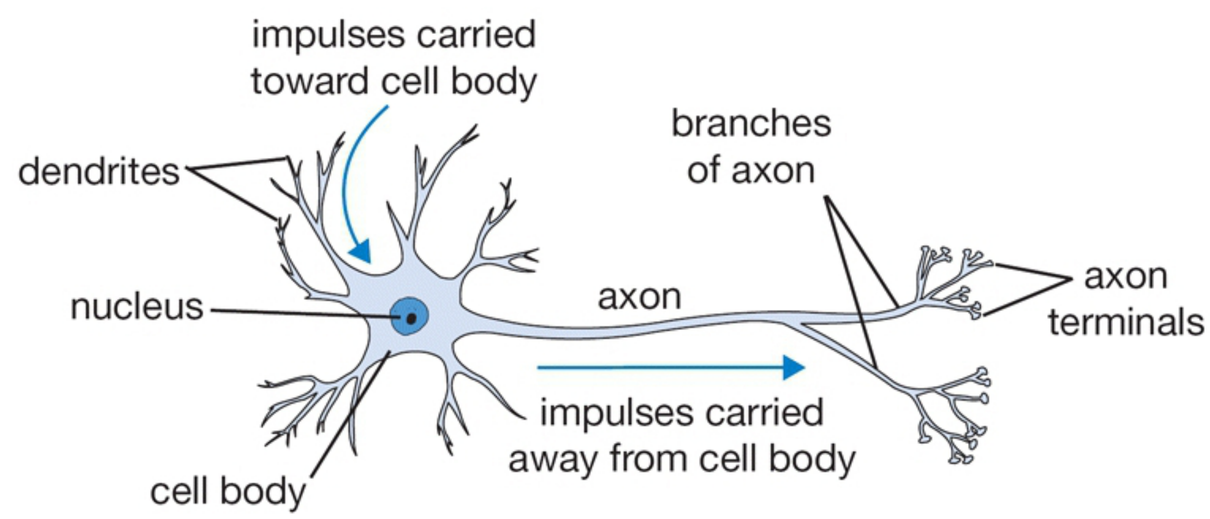

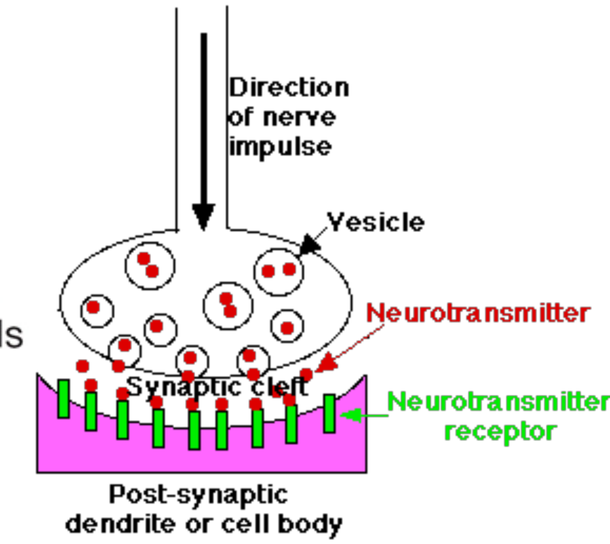

What are the relevant components of the brain?

- Neurons & synapses
- membrane potentials & action potentials
- Functionally specialized brain regions

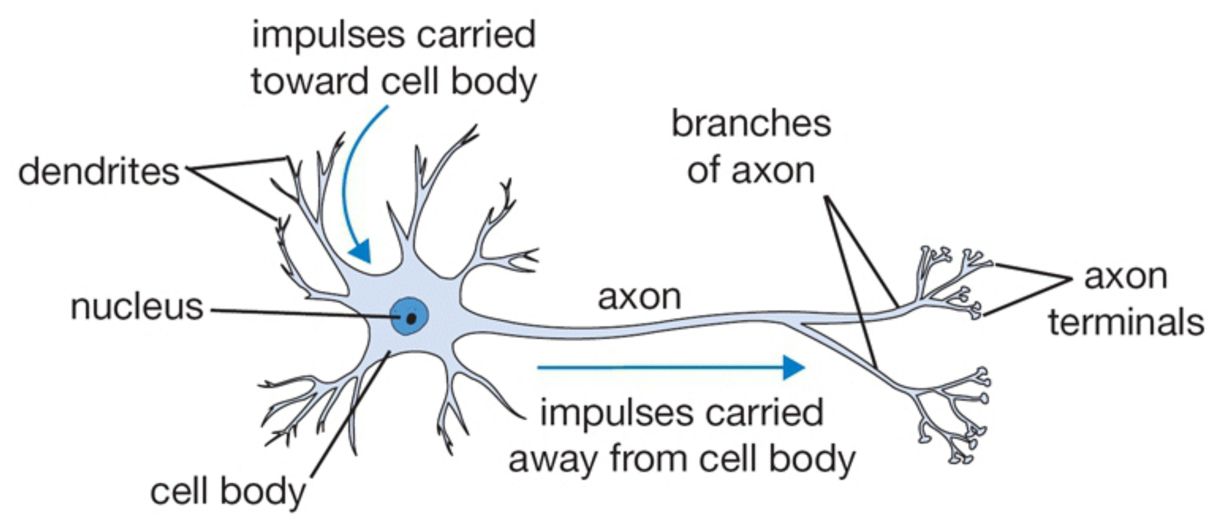

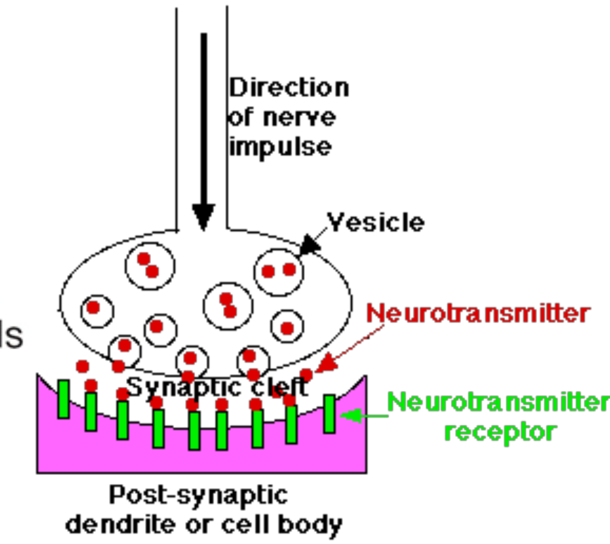

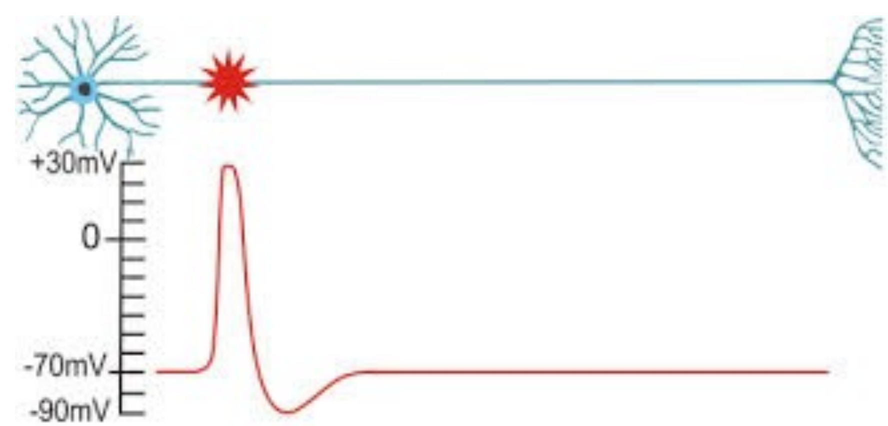

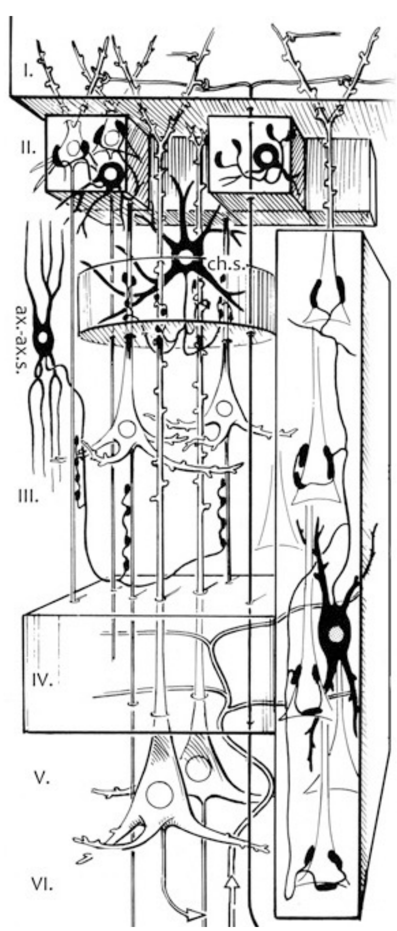

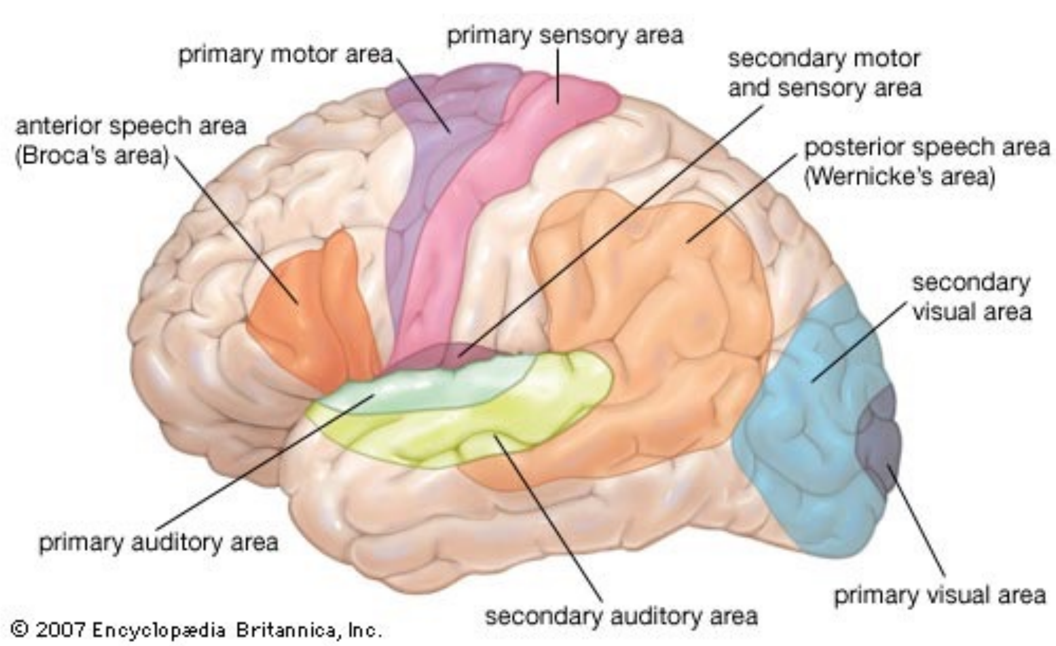

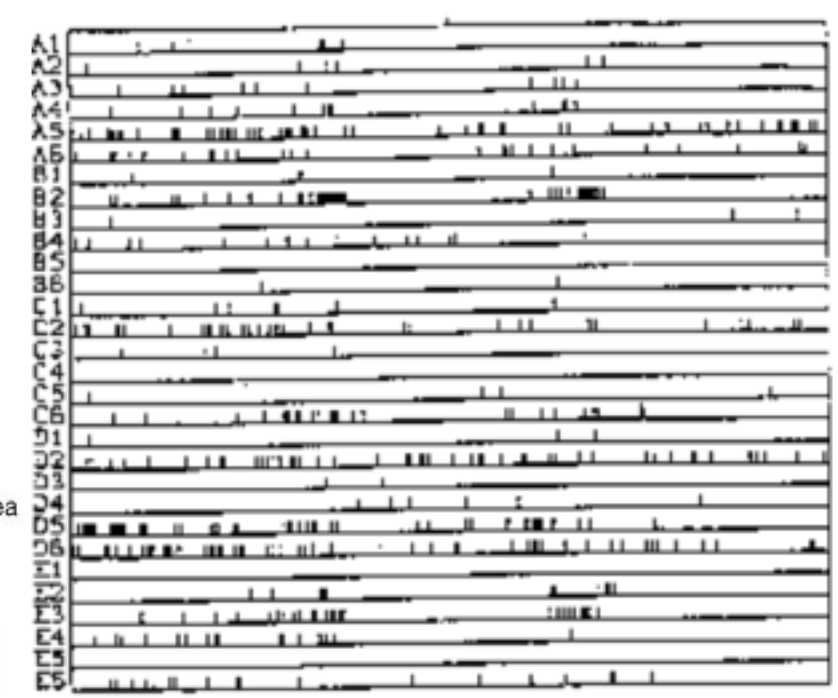

What are the quantities that we can measure?

What are the quantities that we can measure?

- We would like to measure the membrane potential of neurons in the visual cortex of an awake cat while watching visual stimuli: what are the challenges?

What are the quantities that we can measure?

- We would like to measure the membrane potential of neurons in the visual cortex of an awake cat while watching visual stimuli: what are the challenges?
- We cannot hit multiple neurons with an electrode: extracellular measurements — electric/imaging

What are the quantities that we can measure?

- We would like to measure the membrane potential of neurons in the visual cortex of an awake cat while watching visual stimuli: what are the challenges?

- We cannot hit multiple neurons with an electrode: extracellular measurements — electric/imaging

  - We only see action potentials, not membrane potentials

What are the quantities that we can measure?

- We would like to measure the membrane potential of neurons in the visual cortex of an awake cat while watching visual stimuli: what are the challenges?
- We cannot hit multiple neurons with an electrode: extracellular measurements — electric/imaging
  - We only see action potentials, not membrane potentials
  - If we use electrodes, how can we establish which neuron a particular signal belongs to?

What are the quantities that we can measure?

- We would like to measure the membrane potential of neurons in the visual cortex of an awake cat while watching visual stimuli: what are the challenges?
- We cannot hit multiple neurons with an electrode: extracellular measurements — electric/imaging
  - We only see action potentials, not membrane potentials
  - If we use electrodes, how can we establish which neuron a particular signal belongs to?
- How to convince a cat to watch our carefully designed stimuli?

What are the quantities that we can measure?

- We would like to measure the membrane potential of neurons in the visual cortex of an awake cat while watching visual stimuli: what are the challenges?
- We cannot hit multiple neurons with an electrode: extracellular measurements — electric/imaging
  - We only see action potentials, not membrane potentials
  - If we use electrodes, how can we establish which neuron a particular signal belongs to?
- How to convince a cat to watch our carefully designed stimuli?
  - Fixing the animal? Not very natural environment

What are the quantities that we can measure?

- We would like to measure the membrane potential of neurons in the visual cortex of an awake cat while watching visual stimuli: what are the challenges?
- We cannot hit multiple neurons with an electrode: extracellular measurements — electric/imaging
    - We only see action potentials, not membrane potentials
    - If we use electrodes, how can we establish which neuron a particular signal belongs to?
- How to convince a cat to watch our carefully designed stimuli?
    - Fixing the animal? Not very natural environment
    - If you fix it, there will be low incentive to cooperate — end up investigating anaesthesized cats

Neural response variability

- Beyond the receptive field

- Response of neurons to the exact same stimulus displays variability

- The correlation of the responses of pairs of neurons is variable

- Using probabilistic models, we can link these forms of variability to the probabilistic representation of environmental features

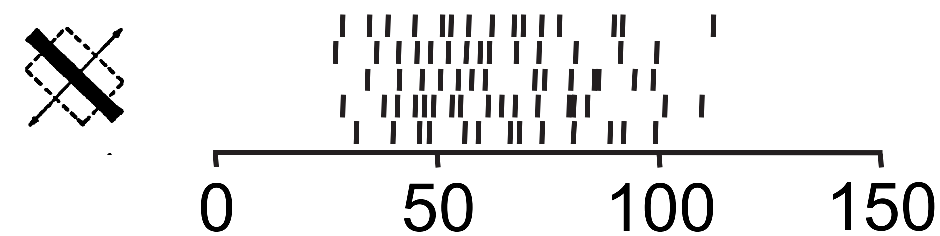

Gur & Snodderly, Cereb Cortex 2006

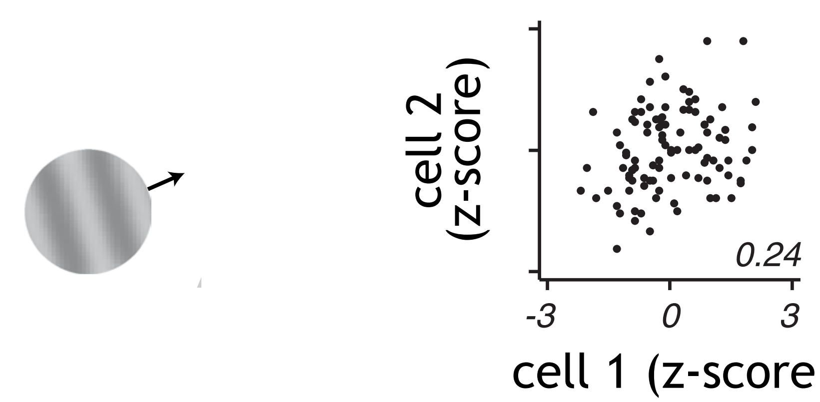

Kohn & Smith, J Neurosci 2005

Layout of the course

http://golab.wigner.mta.hu/teaching/

Intro

Layout of the course

http://golab.wigner.mta.hu/teaching/

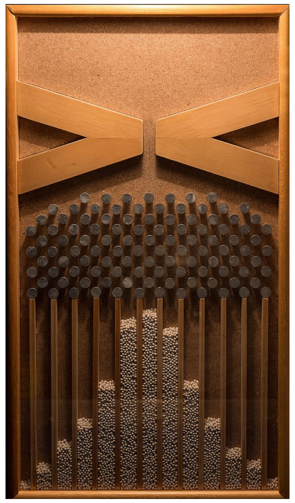

Valószínűségi
modellezés

Intro

Layout of the course

http://golab.wigner.mta.hu/teaching/

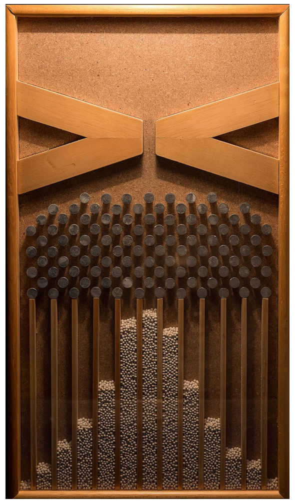

Valószínűségi modellezés

Intro

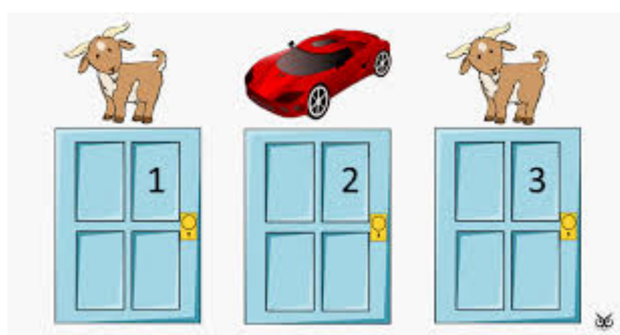

Modelling behavior

Layout of the course

http://golab.wigner.mta.hu/teaching/

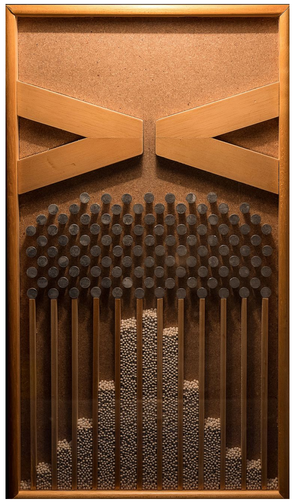

Valószínűségi modellezés

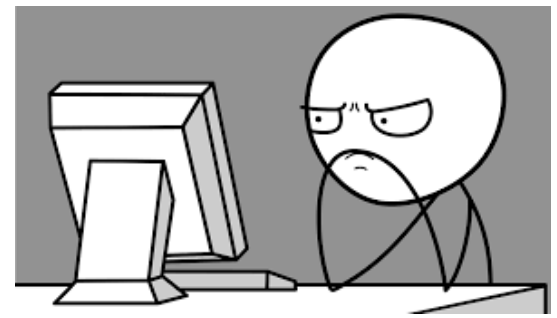

Coding assignment

Intro

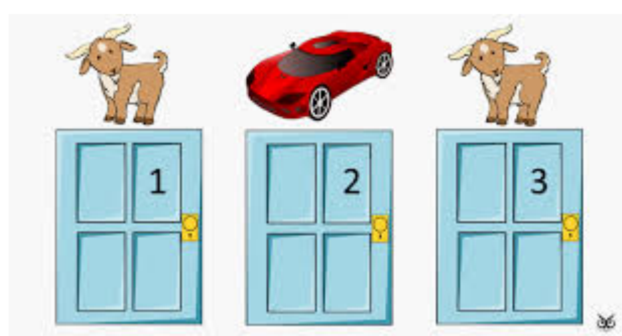

Modelling behavior

Layout of the course

http://golab.wigner.mta.hu/teaching/

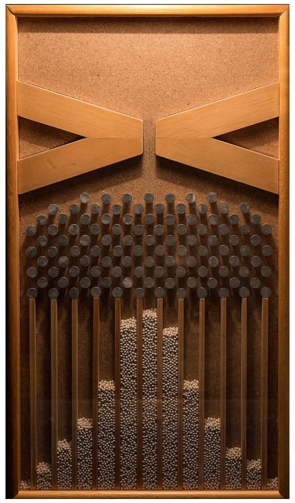

Valószínűségi modellezés

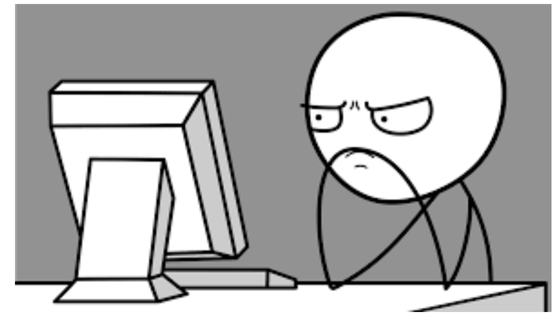

Coding assignment

Intro

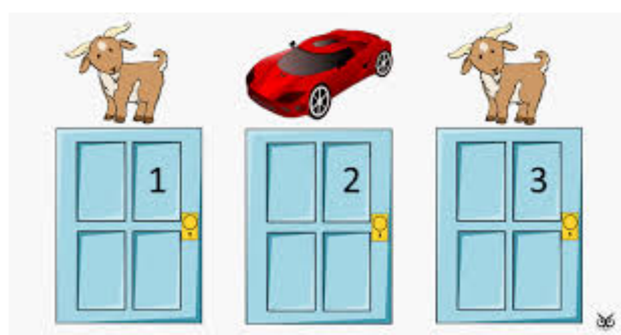

Modelling behavior

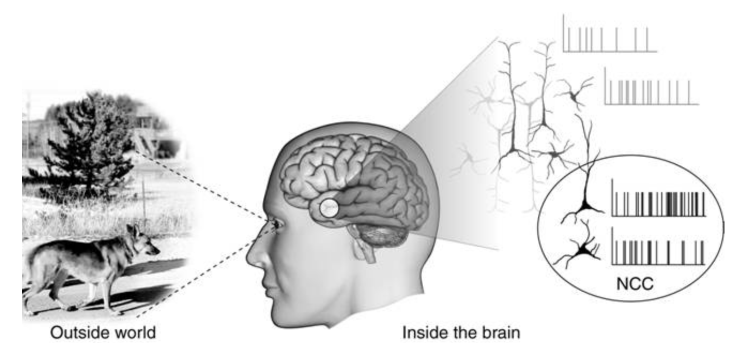

Neural representations

Layout of the course

http://golab.wigner.mta.hu/teaching/

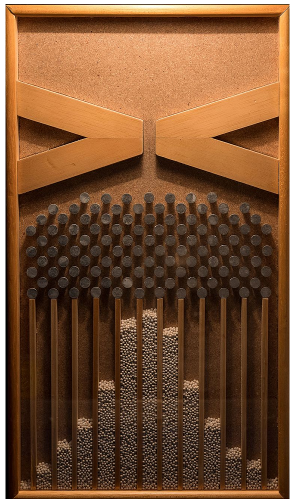

Valószínűségi modellezés

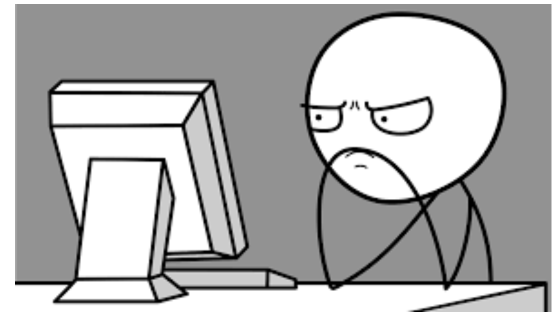

Coding assignment

Theories of visual processing

Intro

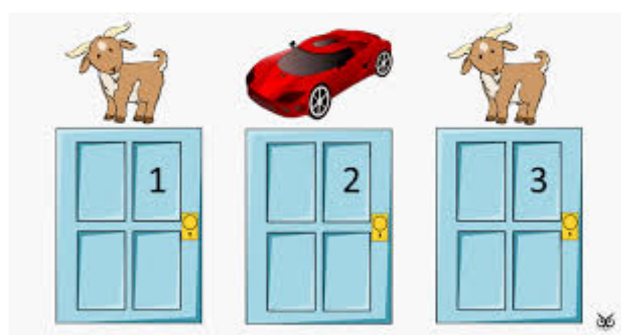

Modelling behavior

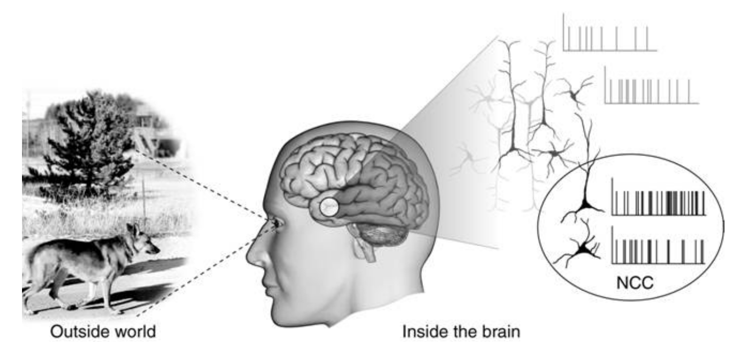

Neural representations

Házi feladat

- Keress egy érdekes illúziót, és határozd meg, hogy mik a lehetséges értelmezései a szenzorikus bemenetnek (kép, hang, szag, tapintás, stb.)

- Keress magyarázatot arra, hogy korábban tanult szabályosságok vagy az aktuális környezet (kontextus) hogyan befolyásolja azt, hogy mit érzékelünk