A Free Energy Principle Approach to Expert Perception in Complex Tasks

Emerging Aspirations Workshop, Sept. 2023

Dr. Michael Harré

Perspective: What is the goal here?

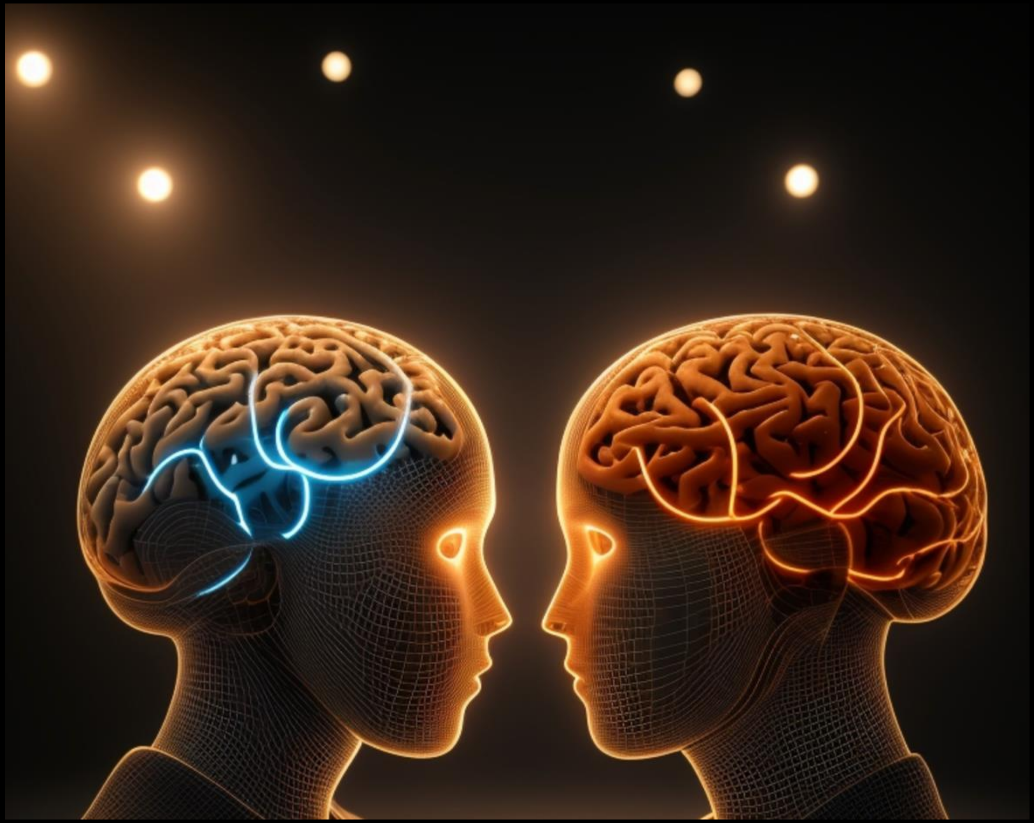

Stable Diffusion

Perspective: What is the goal here?

Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition

Jonathan St. B. T. Evans

Annu. Rev. Psychol. 2008. 59:255–78

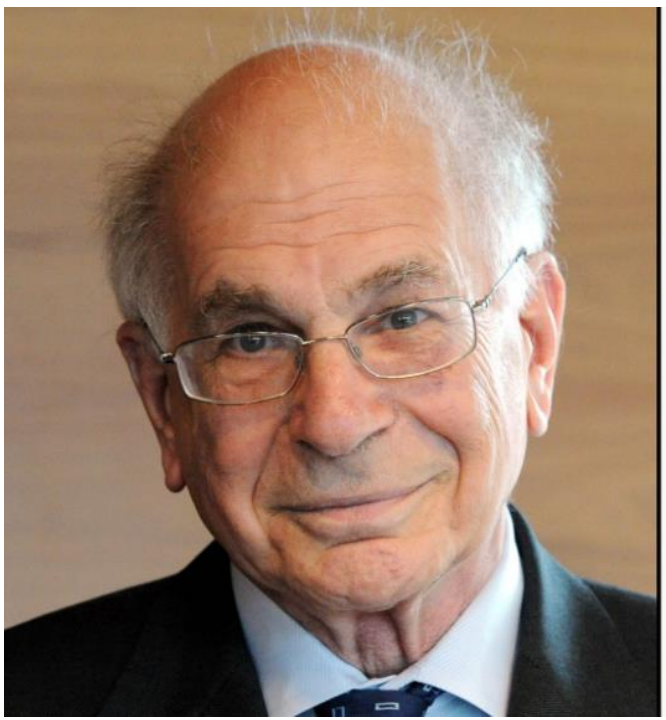

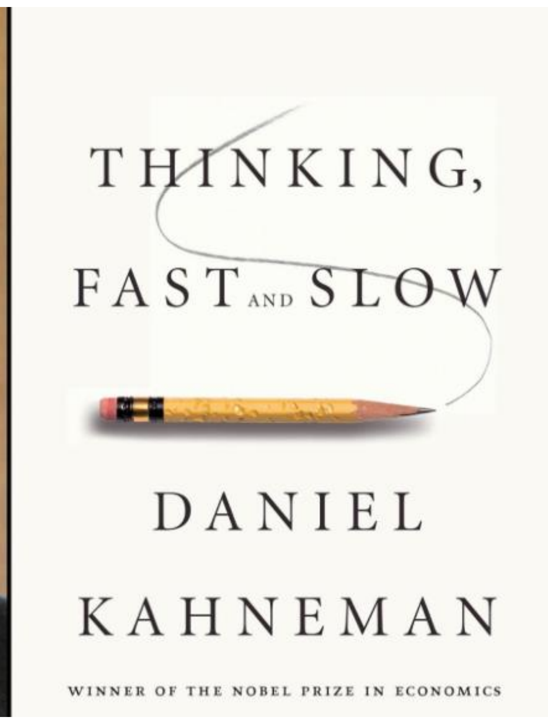

Perspective: What is the goal here?

Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition

Jonathan St. B. T. Evans

Annu. Rev. Psychol. 2008. 59:255–78

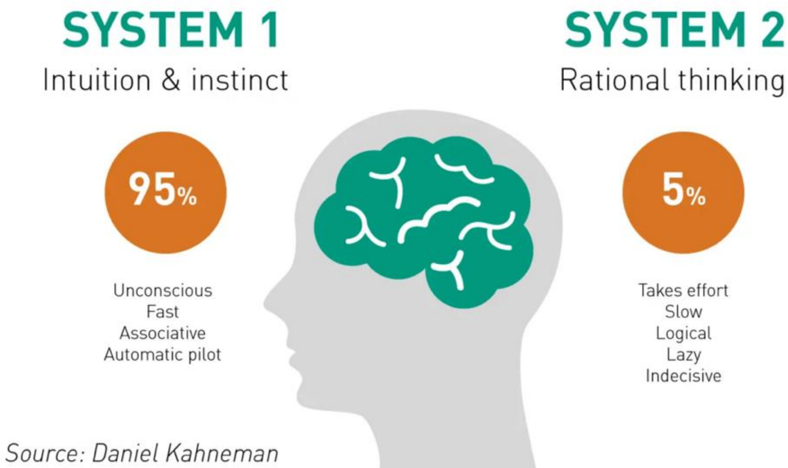

Perspective: What is the goal here?

Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition

Jonathan St. B. T. Evans

Annu. Rev. Psychol. 2008. 59:255–78

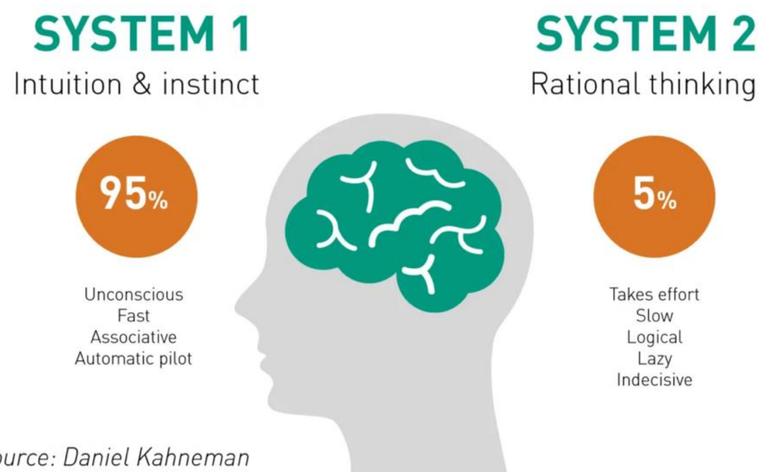

Source: Daniel Kahneman

System 1	System 2
Cluster 1 (Consciousness)	
Unconscious (preconscious)	Conscious
Implicit	Explicit
Automatic	Controlled
Low effort	High effort
Rapid	Slow
High capacity	Low capacity
Default process	Inhibitory
Holistic, perceptual	Analytic, reflective
Cluster 2 (Evolution)	
Evolutionarily old	Evolutionarily recent
Evolutionary rationality	Individual rationality
Shared with animals	Uniquely human
Nonverbal	Linked to language
Modular cognition	Fluid intelligence
Cluster 3 (Functional characteristics)	
Associative	Rule based
Domain specific	Domain general
Contextualized	Abstract
Pragmatic	Logical
Parallel	Sequential
Stereotypical	Egalitarian
Cluster 4 (Individual differences)	
Universal	Heritable
Independent of general intelligence	Linked to general intelligence
Independent of working memory	Limited by working memory capacity

Perspective: What is the goal here?

Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition

Jonathan St. B. T. Evans

Annu. Rev. Psychol. 2008. 59:255–78

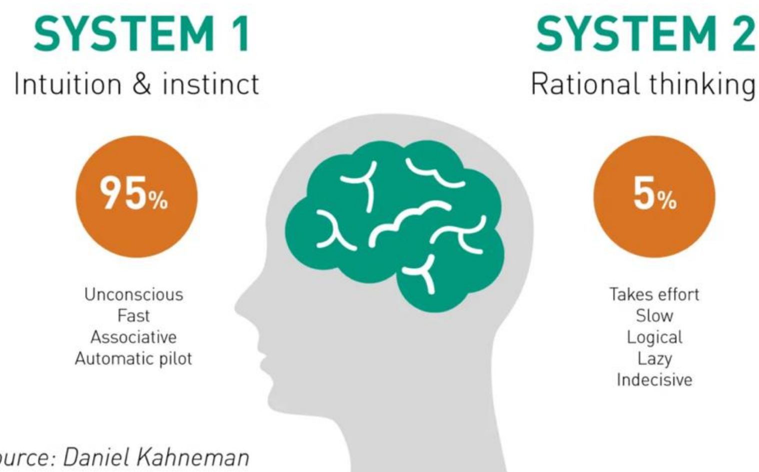

Source: Daniel Kahneman

System 1	System 2
Cluster 1 (Consciousness)	
Unconscious (preconscious)	Conscious
Implicit	Explicit
Automatic	Controlled
Low effort	High effort
Rapid	Slow
High capacity	Low capacity
Default process	Inhibitory
Holistic, perceptual	Analytic, reflective
Cluster 2 (Evolution)	
Evolutionarily old	Evolutionarily recent
Evolutionary rationality	Individual rationality
Shared with animals	Uniquely human
Nonverbal	Linked to language
Modular cognition	Fluid intelligence
Cluster 3 (Functional characteristics)	
Associative	Rule based
Domain specific	Domain general
Contextualized	Abstract
Pragmatic	Logical
Parallel	Sequential
Stereotypical	Egalitarian
Cluster 4 (Individual differences)	
Universal	Heritable
Independent of general intelligence	Linked to general intelligence
Independent of working memory	Limited by working memory capacity

Perspective: What is the goal here?

Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition

Jonathan St. B. T. Evans

Annu. Rev. Psychol. 2008. 59:255–78

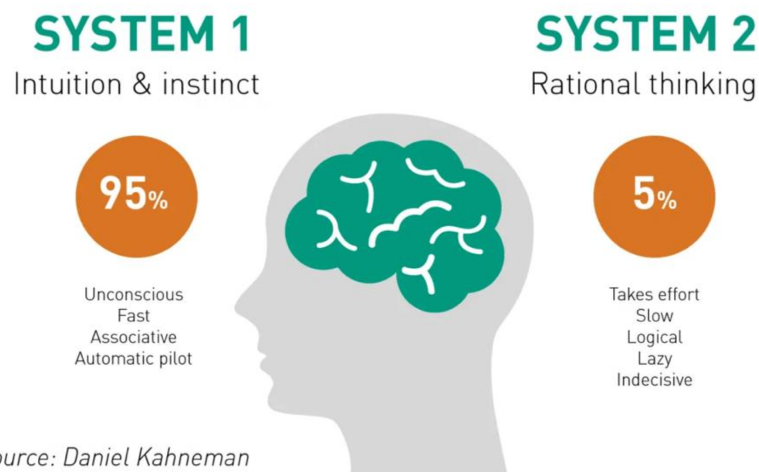

Source: Daniel Kahneman

System 1	System 2
Cluster 1 (Consciousness)	
Unconscious (preconscious)	Conscious
Implicit	Explicit
Automatic	Controlled
Low effort	High effort
Rapid	Slow
High capacity	Low capacity
Default process	Inhibitory
Holistic, perceptual	Analytic, reflective
Cluster 2 (Evolution)	
Evolutionarily old	Evolutionarily recent
Evolutionary rationality	Individual rationality
Shared with animals	Uniquely human
Nonverbal	Linked to language
Modular cognition	Fluid intelligence
Cluster 3 (Functional characteristics)	
Associative	Rule based
Domain specific	Domain general
Contextualized	Abstract
Pragmatic	Logical
Parallel	Sequential
Stereotypical	Egalitarian
Cluster 4 (Individual differences)	
Universal	Heritable
Independent of general intelligence	Linked to general intelligence
Independent of working memory	Limited by working memory capacity

Perspective: What is the goal here?

Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition

Jonathan St. B. T. Evans — Annu. Rev. Psychol. 2008. 59:255–78

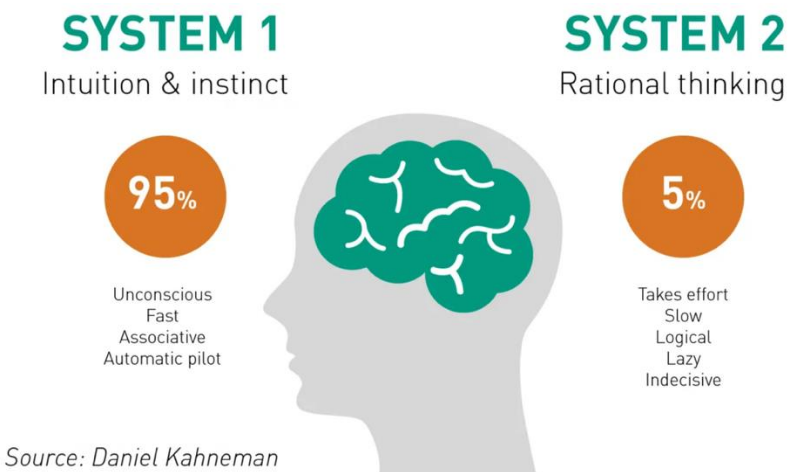

System 1	System 2
Cluster 1 (Consciousness)	
Unconscious (preconscious)	Conscious
Implicit	Explicit
Automatic	Controlled
Low effort	High effort
Rapid	Slow
High capacity	Low capacity
Default process	Inhibitory
Holistic, perceptual	Analytic, reflective
Cluster 2 (Evolution)	
Evolutionarily old	Evolutionarily recent
Evolutionary rationality	Individual rationality
Shared with animals	Uniquely human
Nonverbal	Linked to language
Modular cognition	Fluid intelligence
Cluster 3 (Functional characteristics)	
Associative	Rule based
Domain specific	Domain general
Contextualized	Abstract
Pragmatic	Logical
Parallel	Sequential
Stereotypical	Egalitarian
Cluster 4 (Individual differences)	
Universal	Heritable
Independent of general intelligence	Linked to general intelligence
Independent of working memory	Limited by working memory capacity

Perspective: What is the goal here?

Dual processes are used in complex tasks:

- No ‘normative’ solution
- Optimisation is difficult/impossible
- Noisy and/or ambiguous information
- Non-trivial but sparse structures in the environment
- Little direct access to our cognitive strategies
- A huge amount of data is still insufficient data for exact solutions

Information Theory and Decision-Making

Published: 09 April 2011

The Development of Human Expertise in a Complex Environment

Michael Harré, Terry Bossomaier & Allan Snyder

Minds and Machines 21, 449–464 (2011) | Cite this article

Fig. 6

From: The Development of Human Expertise in a Complex Environment

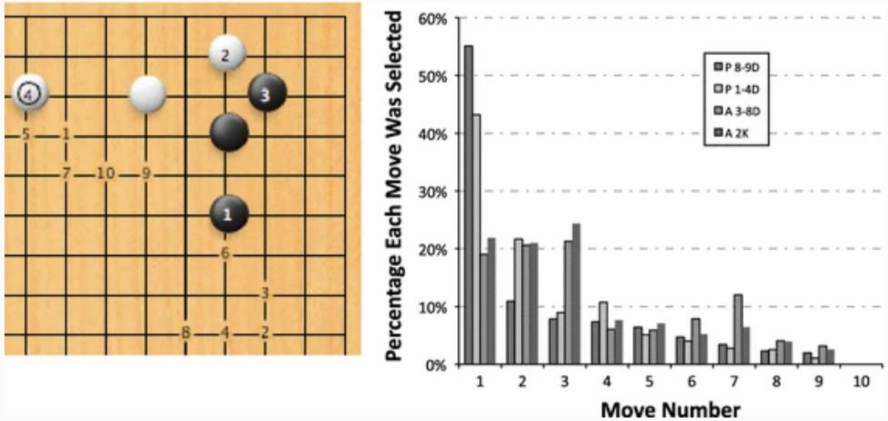

The board state and probability distributions over the next moves. *Top* One of the joseki showing the first six stones played in the local area by 8-9Dan professionals. *Bottom* Four example histograms of the frequency each of the ten moves might follow. Note the order of the moves on the *horizontal axis* in this plot is with respect to the 8-9Dan professionals, i.e. the most preferred move, labelled '1', is most preferred by the 8-9Dan professionals, the move labelled '2' is the second most preferred by these players etc

Information Theory and Decision-Making

Published: 24 March 2011

The aggregate complexity of decisions in the game of Go

M. S. Harré, T. Bossomaier, A. Gillett & A. Snyder

The European Physical Journal B **80**, 555–563 (2011) | Cite this article

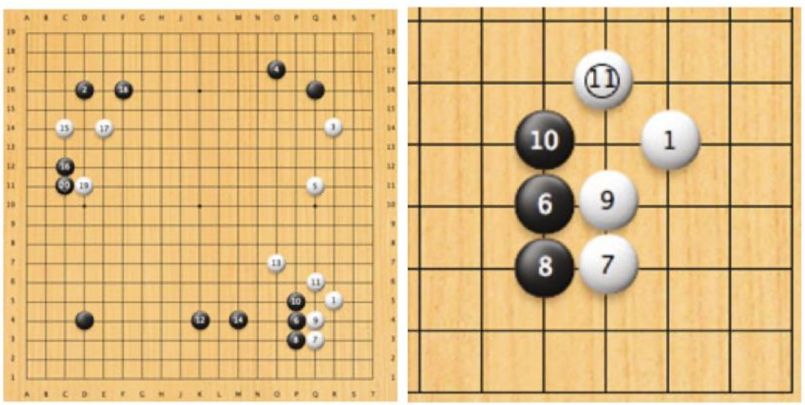

Fig. 1. (Color online) Left: the first 20 moves in a game of Go. Right: Stones played in a 7 × 7 region in the lower right corner, the numbers record the order in which they were played (moves 2 to 5 were played elsewhere on the board).

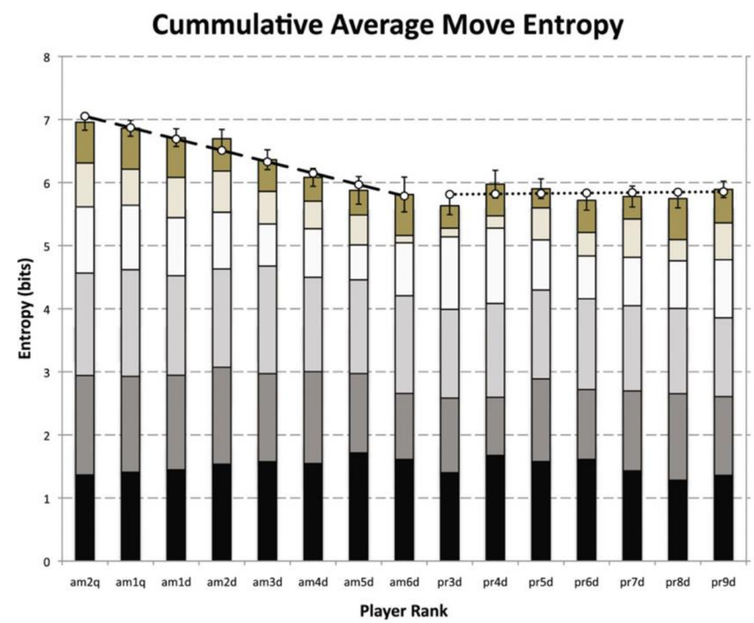

Information Theory and Decision-Making

Published: 09 April 2011

The Development of Human Expertise in a Complex Environment

Michael Harré, Terry Bossomaier & Allan Snyder

Minds and Machines **21**, 449–464 (2011) | Cite this article

Fig. 1

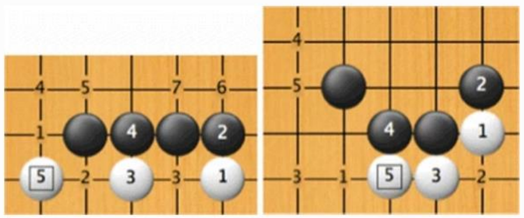

The two smallest patterns used. The starting pattern (the *two black un-numbered stones*) on the *left* is called 'skip one' and the starting pattern on the *right* is called 'knight's move' in Reitman's study

Fig. 2

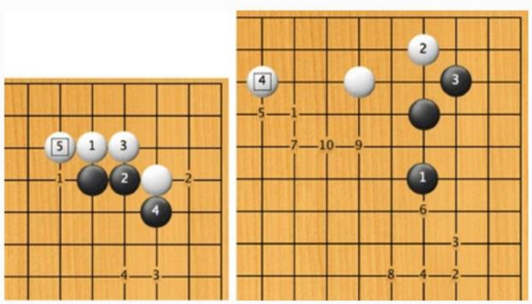

The two mid-sized patterns used. The pattern on the *left* is a variation of the 'avalanche' joseki and the pattern on the *right* is a variation of the '4–4 point low approach high extension' joseki. Note that the board is bounded by the corner in these patterns so that they only ever occur in the corner of the board, unlike the smaller patterns in Fig. 1 that may occur anywhere on the board, including the corners

Fig. 3

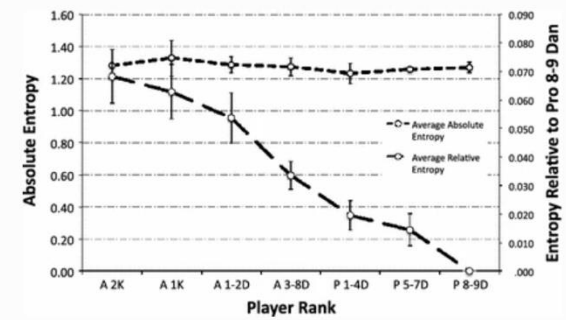

$$\widetilde{K}(p_1(x), p_2(x)) = \sum_{x_i \in X} p_1(x_i) \log\left[\frac{2p_1(x_i)}{[p_1(x_i) + p_2(x_i)]}\right]$$

Fig. 4

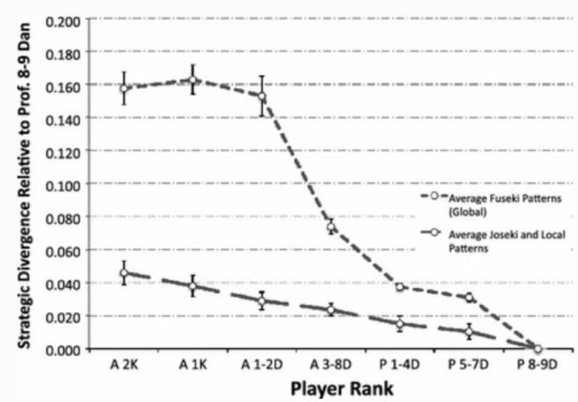

The strategic divergence for global and local patterns. Errors are ± s.e.m

Information Theory and Decision-Making

Published: 09 April 2011

The Development of Human Expertise in a Complex Environment

Michael Harré, Terry Bossomaier & Allan Snyder

Minds and Machines **21**, 449–464 (2011) | Cite this article

Fig. 1

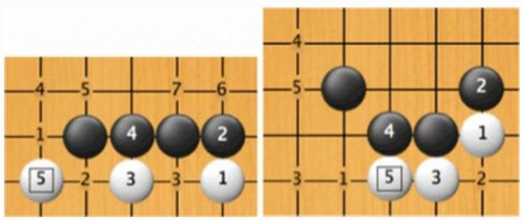

The two smallest patterns used. The starting pattern (the *two black un-numbered stones*) on the *left* is called 'skip one' and the starting pattern on the *right* is called 'knight's move' in Reitman's study

Fig. 2

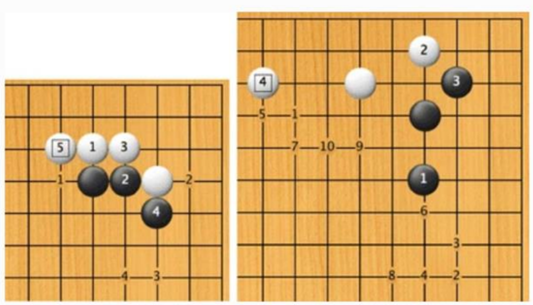

The two mid-sized patterns used. The pattern on the *left* is a variation of the 'avalanche' joseki and the pattern on the *right* is a variation of the '4–4 point low approach high extension' joseki. Note that the board is bounded by the corner in these patterns so that they only ever occur in the corner of the board, unlike the smaller patterns in Fig. 1 that may occur anywhere on the board, including the corners

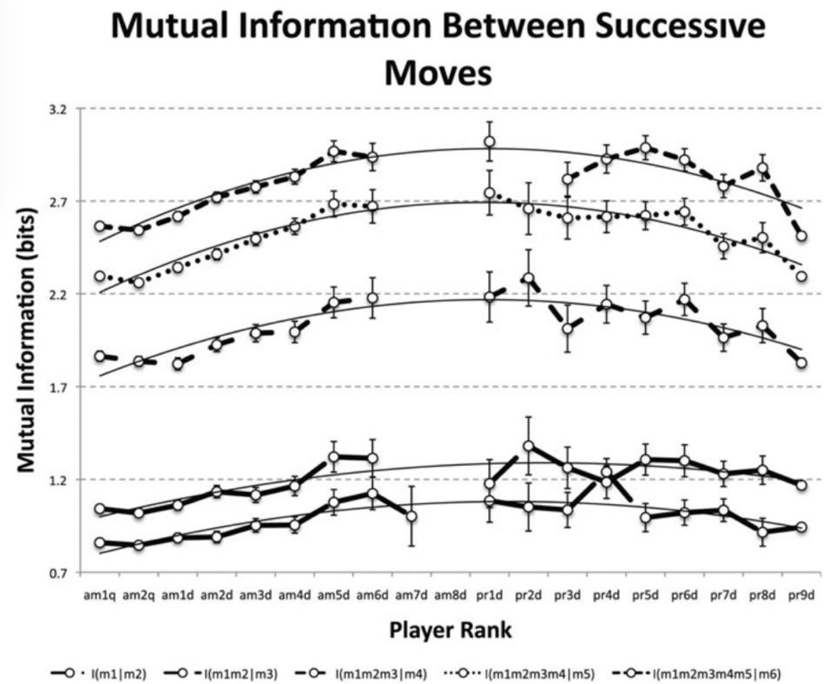

Fig. 4. Mutual information between successive moves.

$$I(x_1; x_2, \rho) = -0.0026\rho^2 + 0.0587\rho + 0.7445$$
$$I(\{x_1, x_2\}; x_3, \rho) = -0.0029\rho^2 + 0.0657\rho + 1.0317$$
$$I(\{x_1, \ldots, x_3\}; x_4, \rho) = -0.0037\rho^2 + 0.0765\rho + 2.0085$$
$$I(\{x_1, \ldots, x_4\}; x_5, \rho) = -0.0039\rho^2 + 0.0792\rho + 2.3886$$
$$I(\{x_1, \ldots, x_5\}; x_6, \rho) = -0.0037\rho^2 + 0.0794\rho + 2.6319$$

[7] For this purpose we set the following integer values for ρ (rank): am2q = 1, am1q = 2, ... , pr9d = 19.

Information Theory and Decision-Making

Published: 09 April 2011

The Development of Human Expertise in a Complex Environment

Michael Harré, Terry Bossomaier & Allan Snyder

Minds and Machines **21**, 449–464 (2011) | Cite this article

Fig. 1

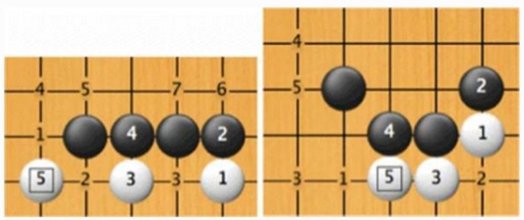

The two smallest patterns used. The starting pattern (the *two black un-numbered stones*) on the *left* is called 'skip one' and the starting pattern on the *right* is called 'knight's move' in Reitman's study

Fig. 2

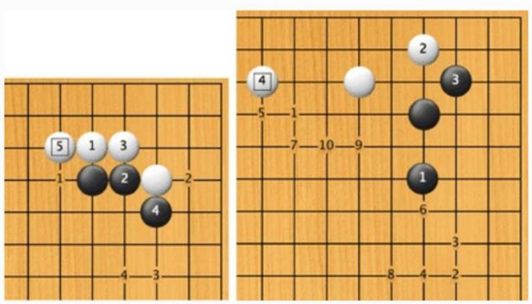

The two mid-sized patterns used. The pattern on the *left* is a variation of the 'avalanche' joseki and the pattern on the *right* is a variation of the '4–4 point low approach high extension' joseki. Note that the board is bounded by the corner in these patterns so that they only ever occur in the corner of the board, unlike the smaller patterns in Fig. 1 that may occur anywhere on the board, including the corners

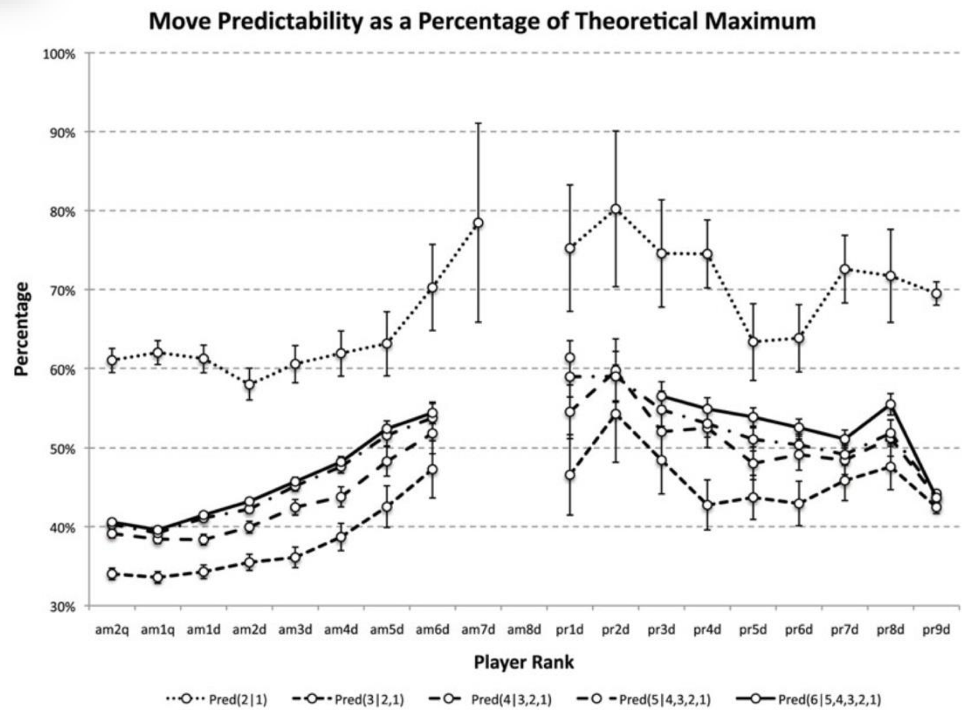

Information Theory and Decision-Making

Experimentally derived
“Chunks” in Go

COGNITIVE PSYCHOLOGY 8, 336–356 (1976)

Skilled Perception in Go: Deducing Memory Structures
from Inter-Response Times

JUDITH S. REITMAN

PERCEPTION IN GO

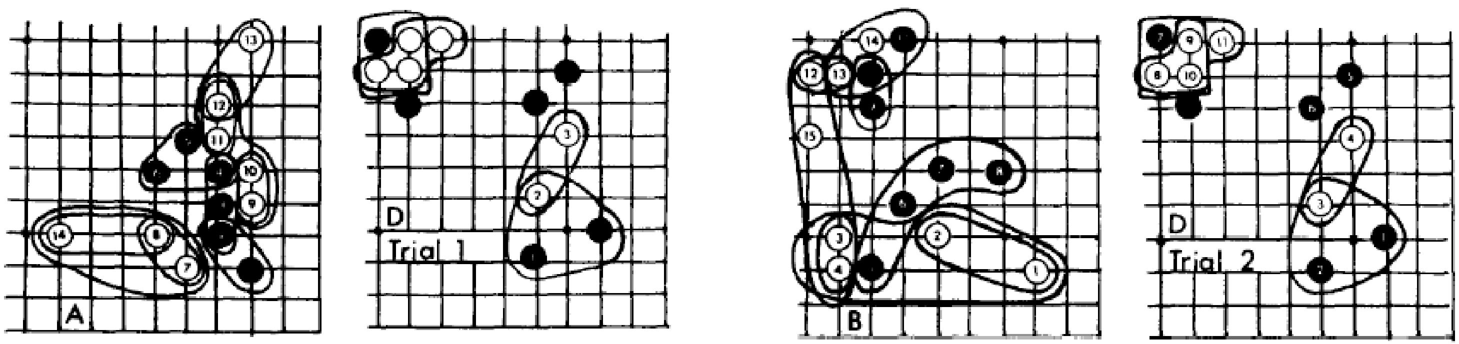

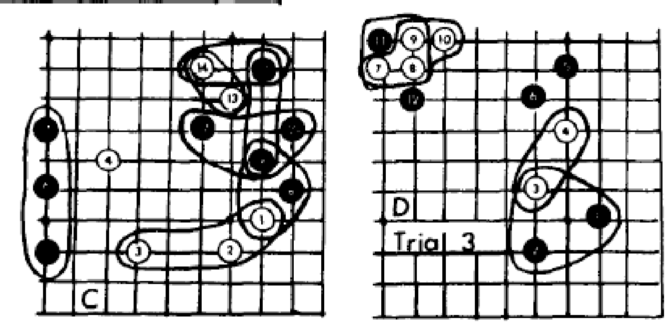

FIG. 6. Examples of the Go Master's penciled partitioning of meaningful patterns and the orders in which he recalled the elements.

Information Theory and Decision-Making

The general trends observed on the way to being "world class"

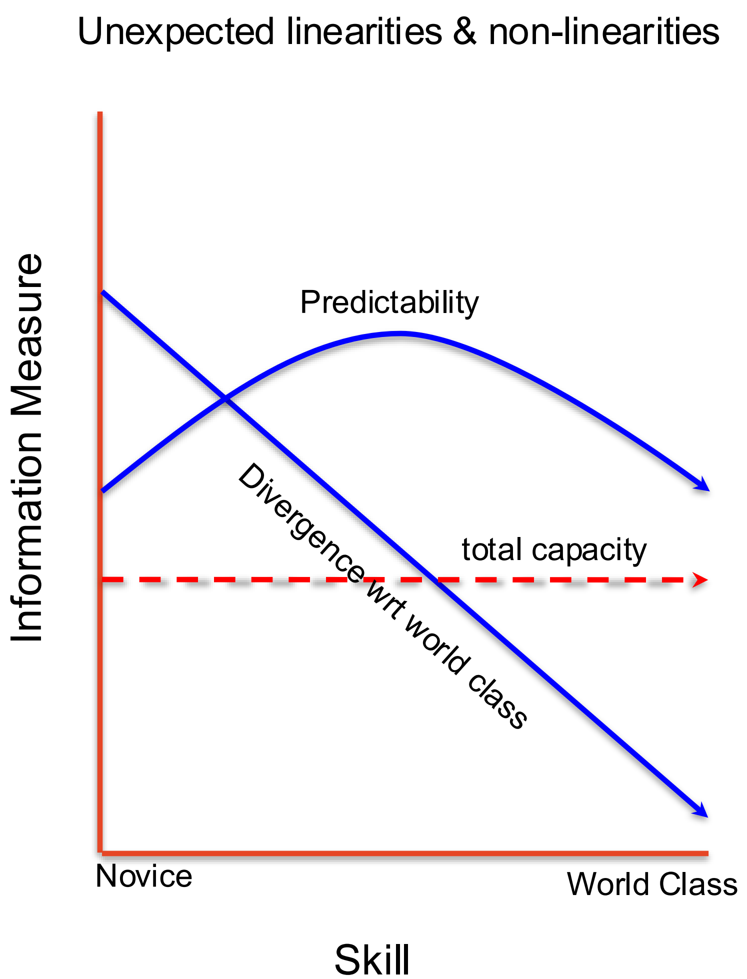

The Psychology of Perception in Games

Other factors that seemed relevant:

- Large span memory: So many details remembered so well

Research Article

VISUAL SPAN IN EXPERT CHESS PLAYERS: Evidence From Eye Movements

Eyal M. Reingold,[1] Neil Charness,[2] Marc Pomplun,[1] and Dave M. Stampe[1]

[1]University of Toronto, Toronto, Ontario, Canada, and [2]Florida State University

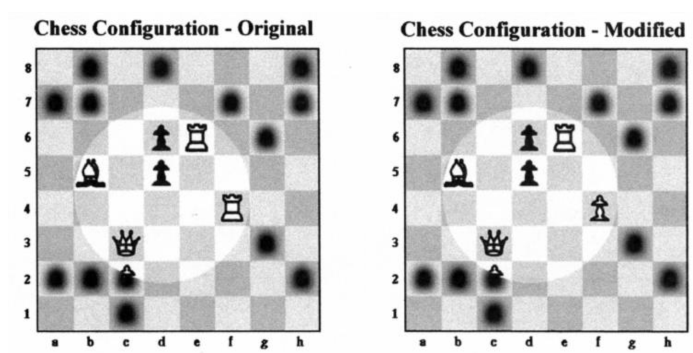

Fig. 1. Illustration of the flicker paradigm. The top row displays an original and a modified (the changed piece is in square f4) chess configuration taken from an actual game. The bottom row displays an original and a modified (the changed piece is in square b2) random configuration obtained by scrambling an actual game configuration. In all four displays, a gaze-contingent window is present, with chess pieces outside the window being replaced by blobs masking their identity and color. (The difference in luminance between the regions inside and outside the window was not present in actual experimental displays and was added here for illustrative purposes.)

The Psychology of Perception in Games

Visual foraging: amateurs yes, experts less so

- Experts: didn't saccade ~16% of the time
- Intermediates: didn't saccade ~3% of the time
- Novices: didn't saccade ~2% of the time

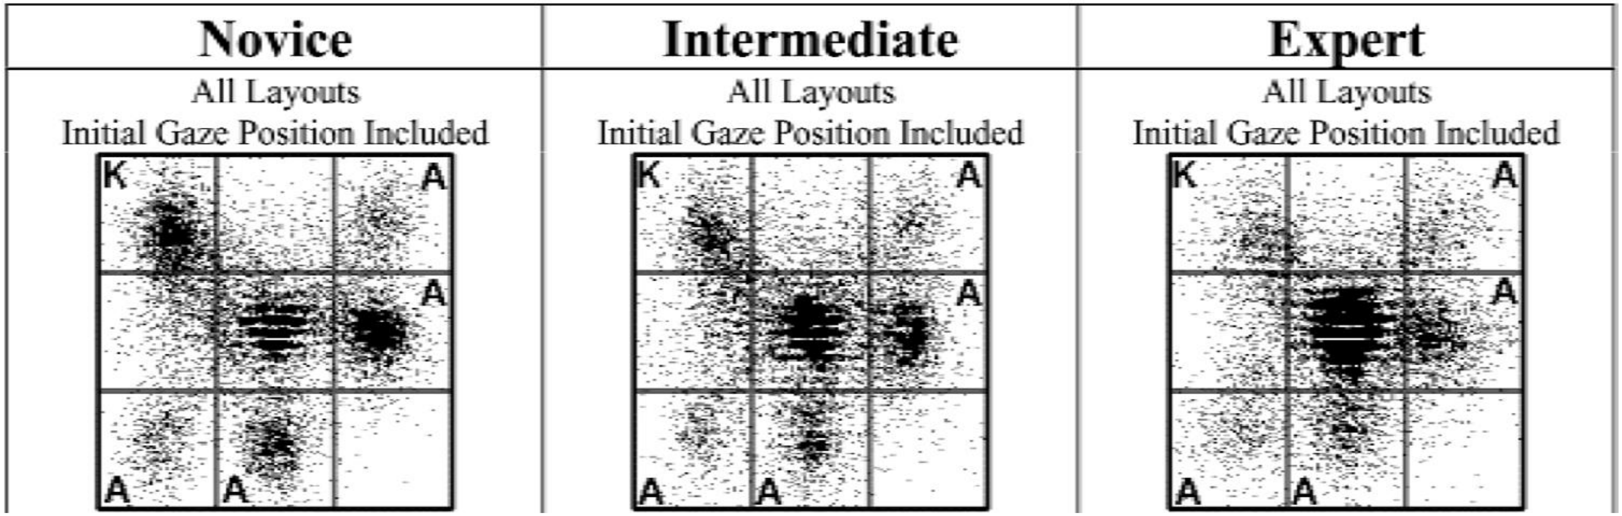

The Psychology of Perception in Games

Visual foraging: amateurs yes, experts less so

- Experts: didn't saccade ~16% of the time
- Intermediates: didn't saccade ~3% of the time
- Novices: didn't saccade ~2% of the time

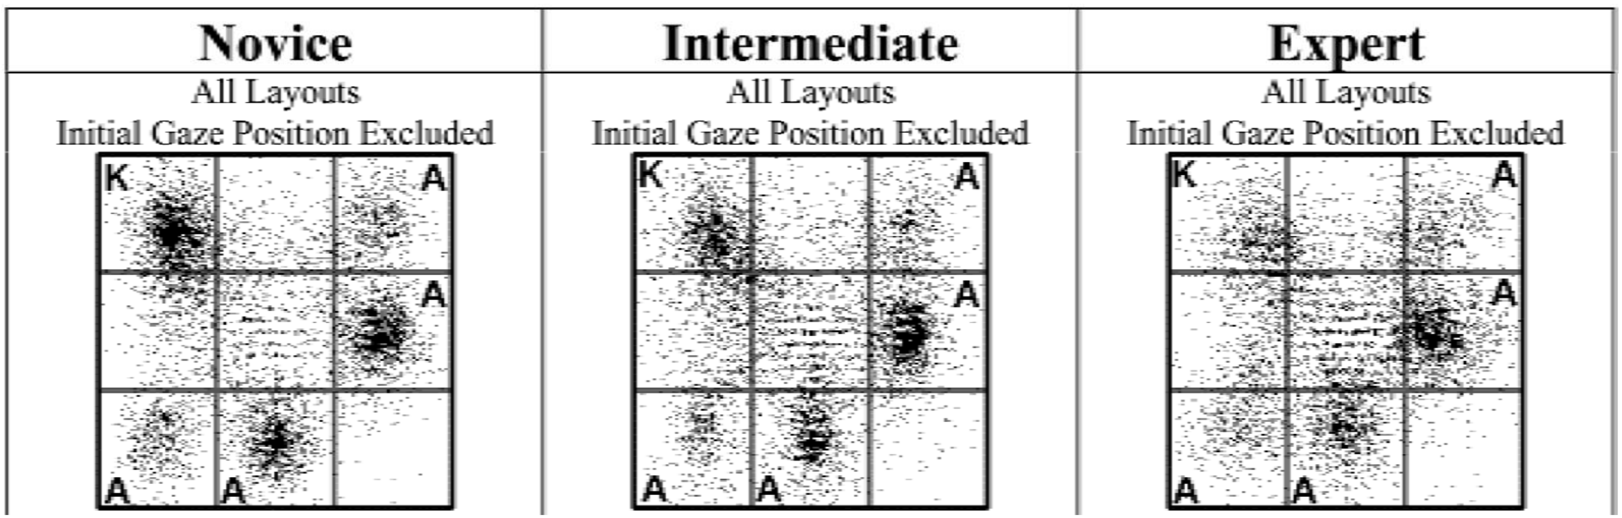

The Psychology of Perception in Games

Visual foraging: amateurs yes, experts less so

- Experts: didn't saccade ~16% of the time
- Intermediates: didn't saccade ~3% of the time
- Novices: didn't saccade ~2% of the time

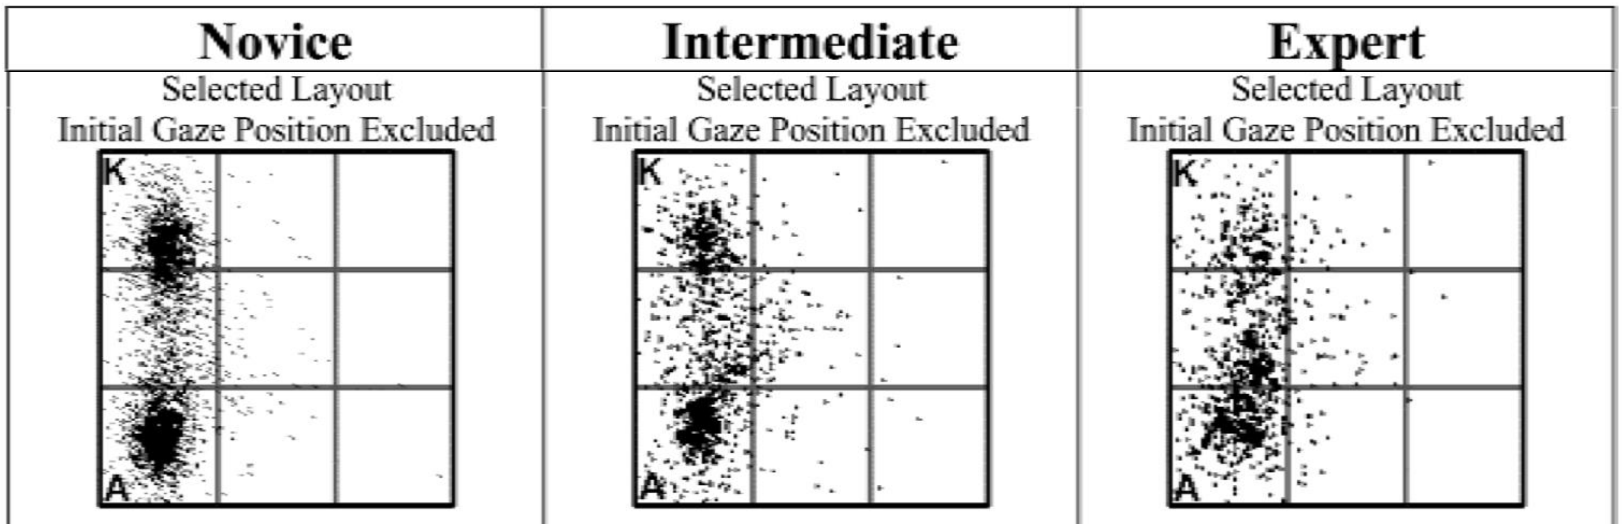

The Psychology of Perception in Games

Eye saccades change with expertise and goals:

Experts perceive the environment differently: Expert versus novice eye saccades[1]

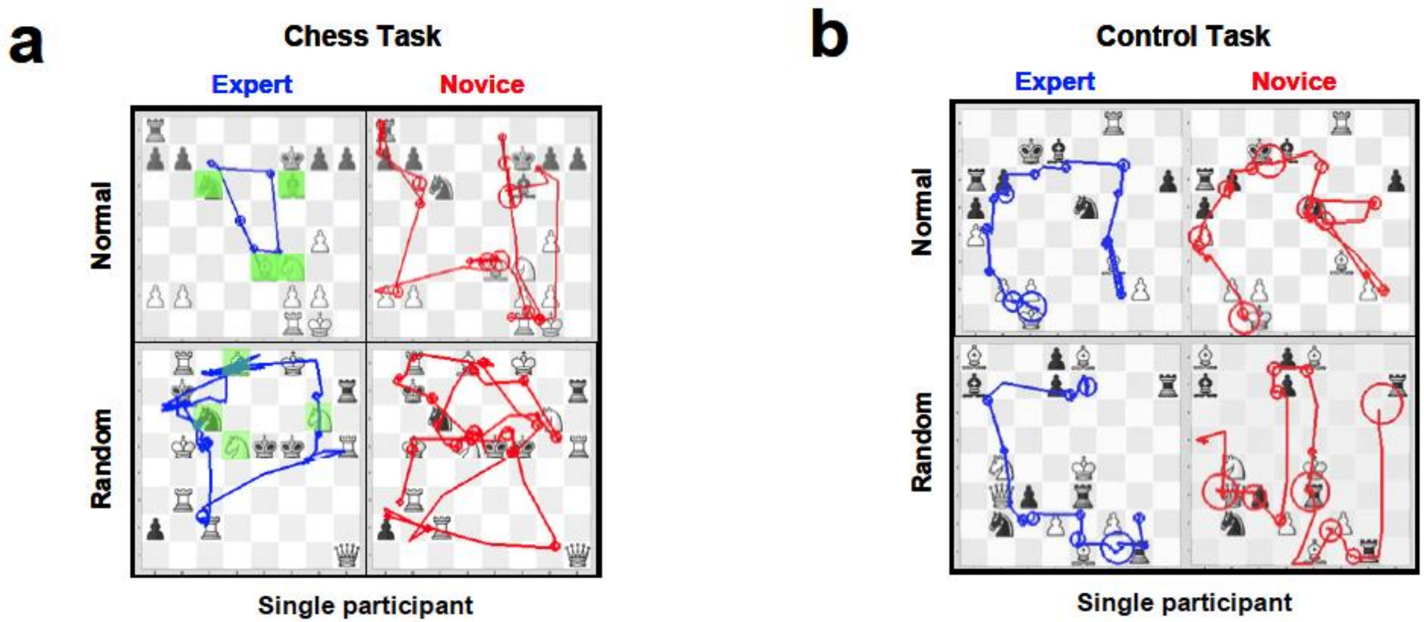

[1] M. Bilalic et al. "Mechanisms and Neural Basis of Object and Pattern Recognition: A Study with Chess Experts" (2010)

The Psychology of Perception in Games

Eye saccades change with expertise and goals:

Experts perceive the environment differently: Expert versus novice eye saccades[1]

"Experts' extensive knowledge facilitates <u>immediate pattern recognition by directing experts toward the relevant objects</u> and allowing them to ignore irrelevant ones."

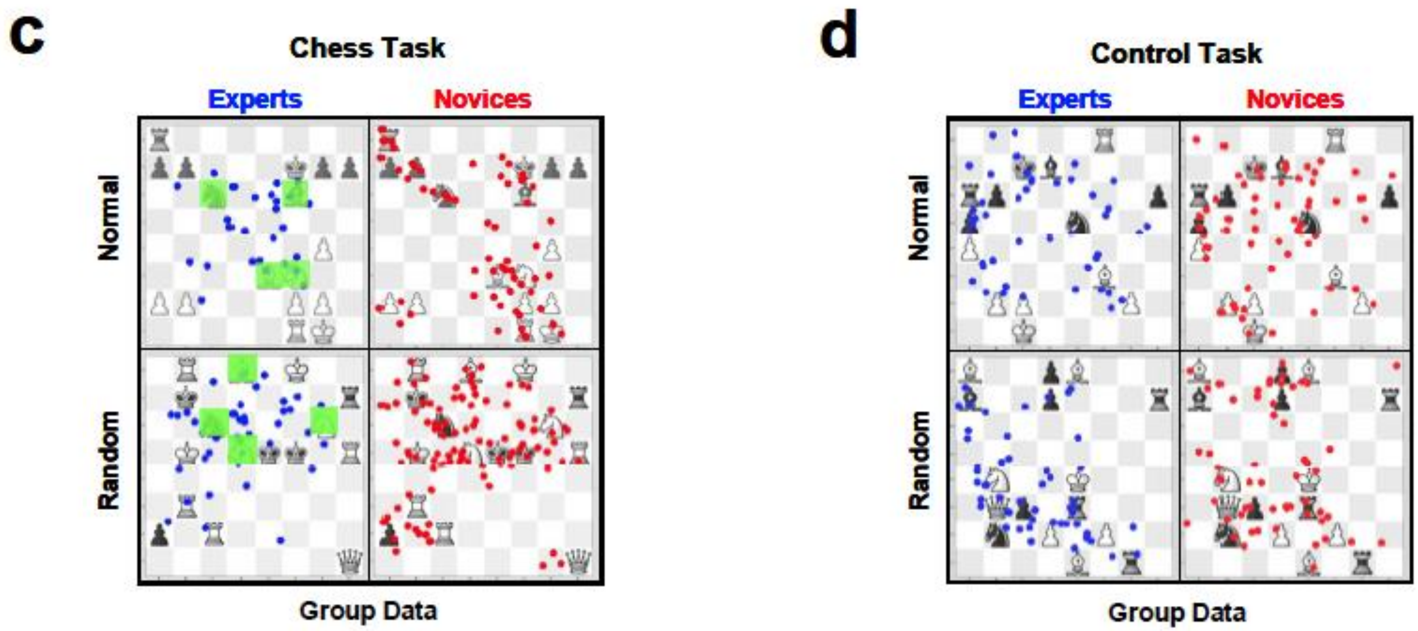

[1] M. Bilalic et al. "Mechanisms and Neural Basis of Object and Pattern Recognition: A Study with Chess Experts" (2010)

Visual Foraging and Whole of Scene Perception

Object Perception as Bayesian Inference

Annual Review of Psychology
Vol. 55:271-304 (Volume publication date 4 February 2004)
First posted online on October 6, 2003
https://doi.org/10.1146/annurev.psych.55.090902.142005

Daniel Kersten
Pascal Mamassian
Alan Yuille

Annual Review of Neuroscience

Integration of Feedforward and Feedback Information Streams in the Modular Architecture of Mouse Visual Cortex

Andreas Burkhalter,[1] Rinaldo D. D'Souza,[1] Weiqing Ji,[1] and Andrew M. Meier[2]

Annu. Rev. Neurosci. 2023. 46:259–80

Perceptual Learning: Toward a Comprehensive Theory

Takeo Watanabe and Yuka Sasaki

Department of Cognitive, Linguistic, and Psychological Sciences, Brown University, Providence, Rhode Island 02912; email: Takeo_Watanabe@Brown.edu

Annu. Rev. Psychol. 2015. 66:197–221

Visual Foraging and Whole of Scene Perception

Eye saccades change with expertise and goals:

The Active Inference solution to this problem

frontiers in Computational Neuroscience

HYPOTHESIS AND THEORY
published: 14 June 2016
doi: 10.3389/fncom.2016.00056

Scene Construction, Visual Foraging, and Active Inference

M. Berk Mirza[1], Rick A. Adams[2,3], Christoph D. Mathys[1,4,5] and Karl J. Friston[1]*

Visual Foraging and Whole of Scene Perception

Eye saccades change with expertise and goals:

The Active Inference solution to this problem

frontiers in Computational Neuroscience

HYPOTHESIS AND THEORY
published: 14 June 2016
doi: 10.3389/fncom.2016.00056

Scene Construction, Visual Foraging, and Active Inference

M. Berk Mirza[1], Rick A. Adams[2,3], Christoph D. Mathys[1,4,5] and Karl J. Friston[1]*

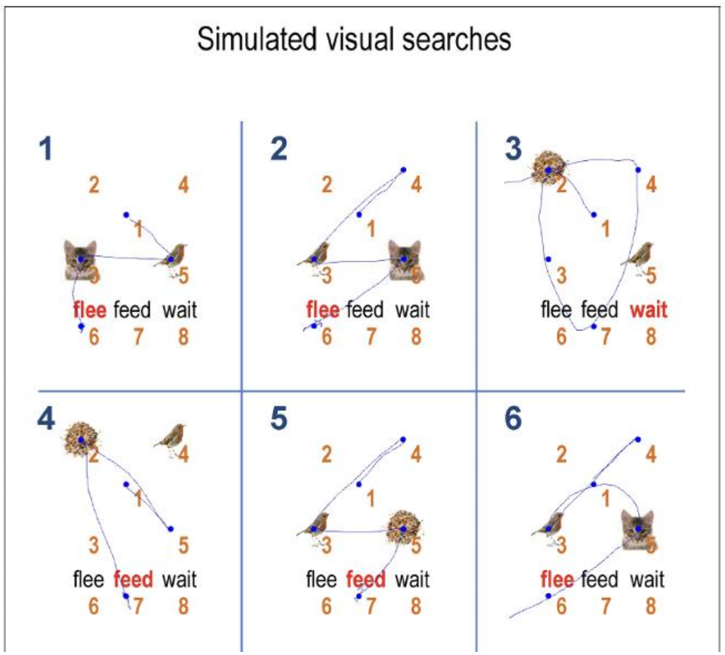

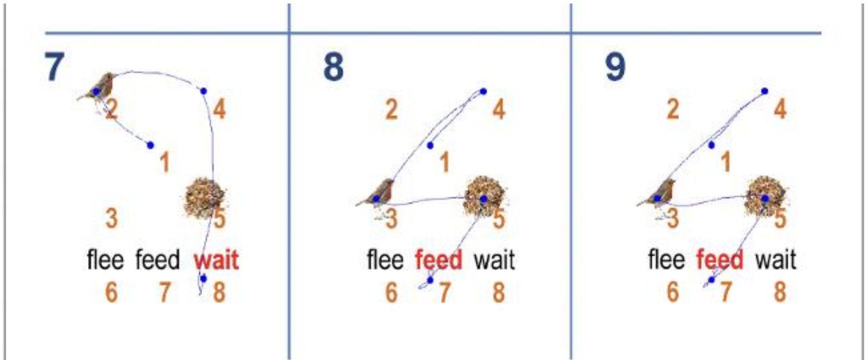

FIGURE 7 | Sequences of saccades: this figure illustrates the behavior for the first nine trials shown in the previous figure using the same format as **Figure 4** (upper right panel). The numbers on the top left in each cell show the trial number. With the exception of the third trial, the agent is able to recognize or categorize the scene after a small number of epistemically efficient saccades.

Visual Foraging and Whole of Scene Perception

Making a "decision": how many cars are in this picture?

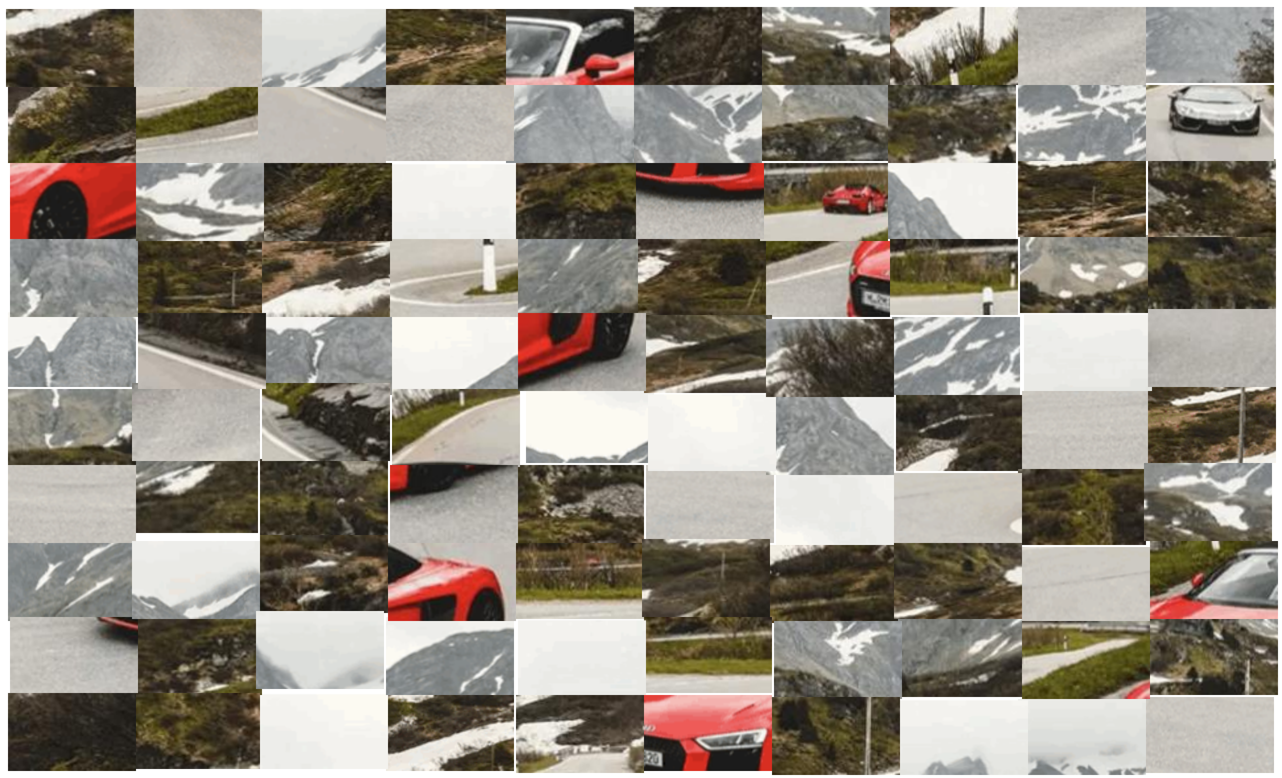

Visual Foraging and Whole of Scene Perception

Making a “decision”: how many cars are in this picture?

Visual Foraging and Whole of Scene Perception

Making a “decision”: how many cars are in this picture?

Visual Foraging and Whole of Scene Perception

Visual Functions of Primate Area V4

Annual Review of Vision Science

Vol. 6:363-385 (Volume publication date September 2020)
First published as a Review in Advance on June 24, 2020
https://doi.org/10.1146/annurev-vision-030320-041306

Anitha Pasupathy,[1,2] Dina V. Popovkina,[3] and Taekjun Kim[1,2]

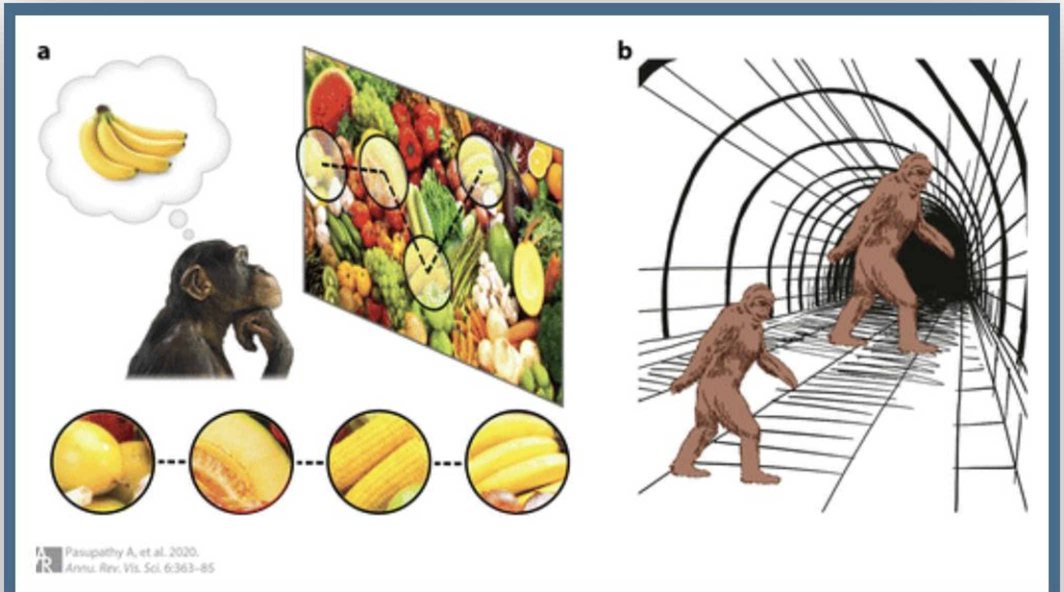

Figure 7 Goal-oriented representations. (*a*) When confronted with the challenge of spotting bananas in a cluttered produce aisle, the subject may saccade to different locations with yellow objects (*dashed trajectory*) and compare the shape of the object at the attentional focus (*circles*) with a remembered object. Area V4 is thought to be important for all aspects of this process. (*b*) Size illusion. The retinal sizes of the two sasquatches in this image are identical, but the perceived sizes are dramatically different. This is because the surrounding context suggests that the sasquatch at right is farther away from the observer; thus, the same retinal size would imply a much larger sasquatch farther away.

Visual Foraging and Whole of Scene Perception

Visual foraging is enough in unstructured environs

1. Local pattern matching is as good as possible in unstructured environs;
2. Consequently, foraging is necessary in such situations.

But for “complex” decisions:

1. Experts’ saccades systematically vary in task relevant contexts;
2. This largely vanishes if the structure is removed;
3. Sometimes experts (natural or trained) don’t saccade;
4. “Whole of scene” perception can be enough.

Visual Foraging and Whole of Scene Perception

There is a data problem for human (and machine) learning

The problem: pattern matching doesn't scale up to a whole board

Local: very common

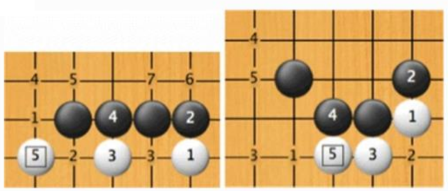

Global: very(!) uncommon

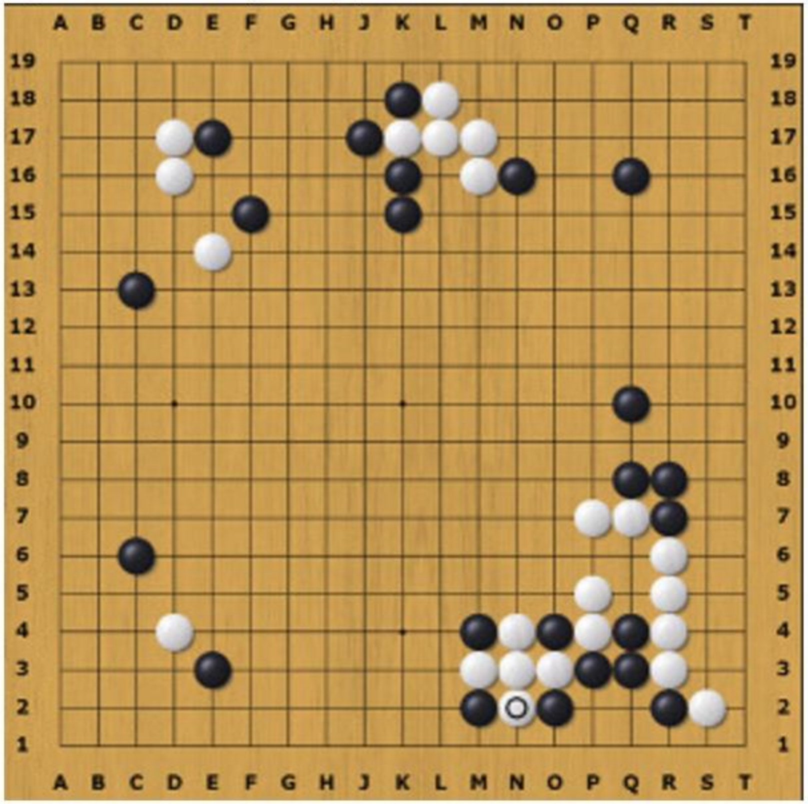

Visual Foraging and Whole of Scene Perception

There is a data problem for human (and machine) learning

Serre et al's hierarchical model of vision:

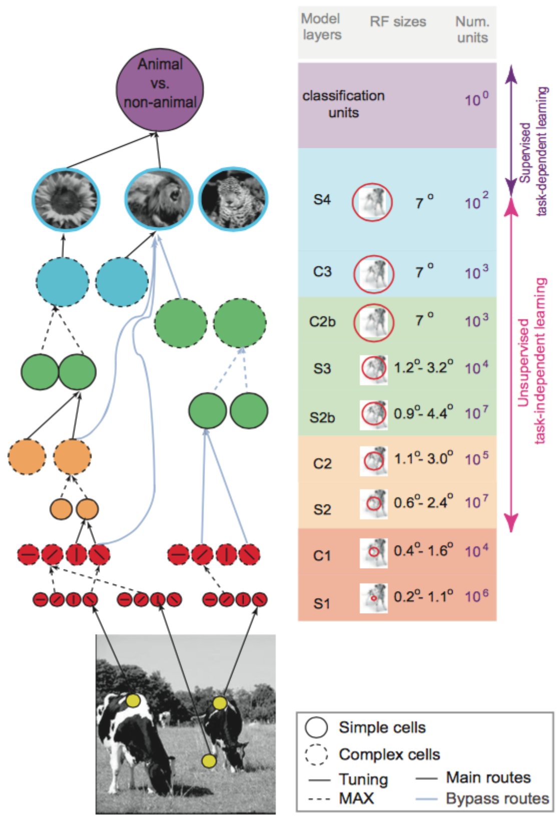

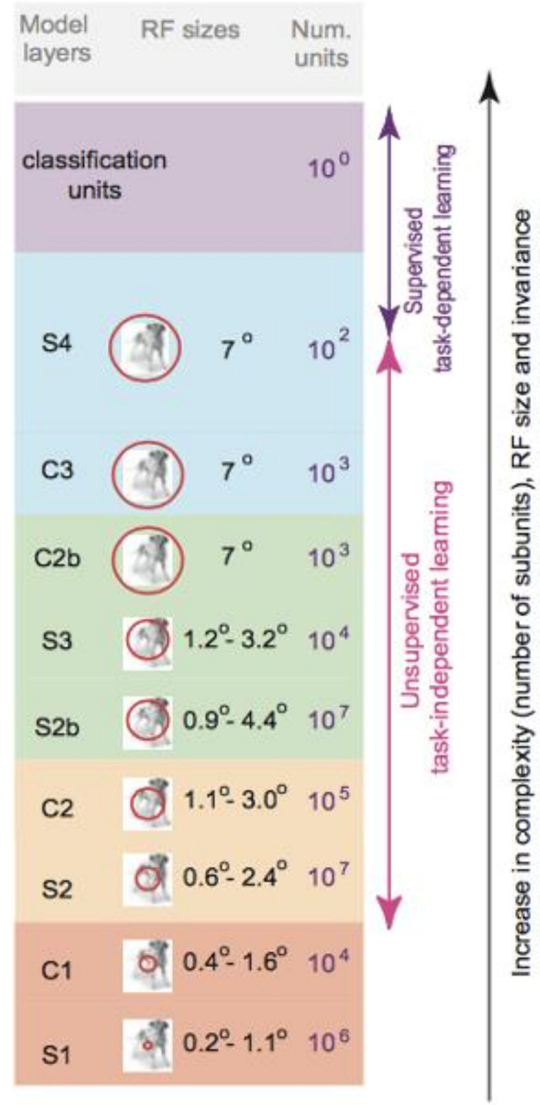

"A feedforward architecture accounts for rapid categorization", T. Serre, A. Oliva & T. Poggio PNAS (2007)
"The Receptive Fields of the Inferior Temporal Cortex Neurons in Natural Scenes", E. Rolls, N. Aggelopoulos & F. Zheng, Journal of Neuroscience (2003)
"Visual cortical mechanisms of perceptual grouping: interacting layers, networks, columns, and maps", W.D. Ross, S. Grossberg, E. Mingolla

Visual Foraging and Whole of Scene Perception

There is a data problem for human (and machine) learning

Serre et al's hierarchical model of vision:

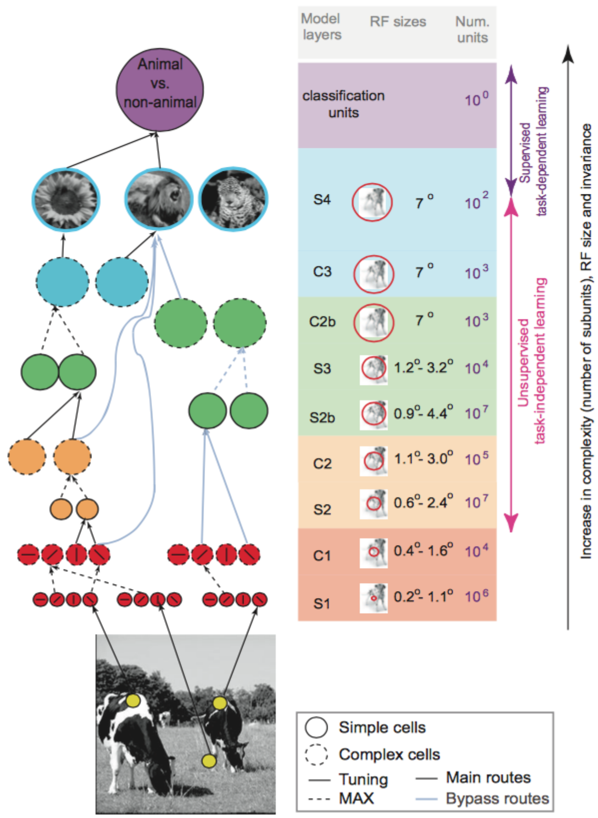

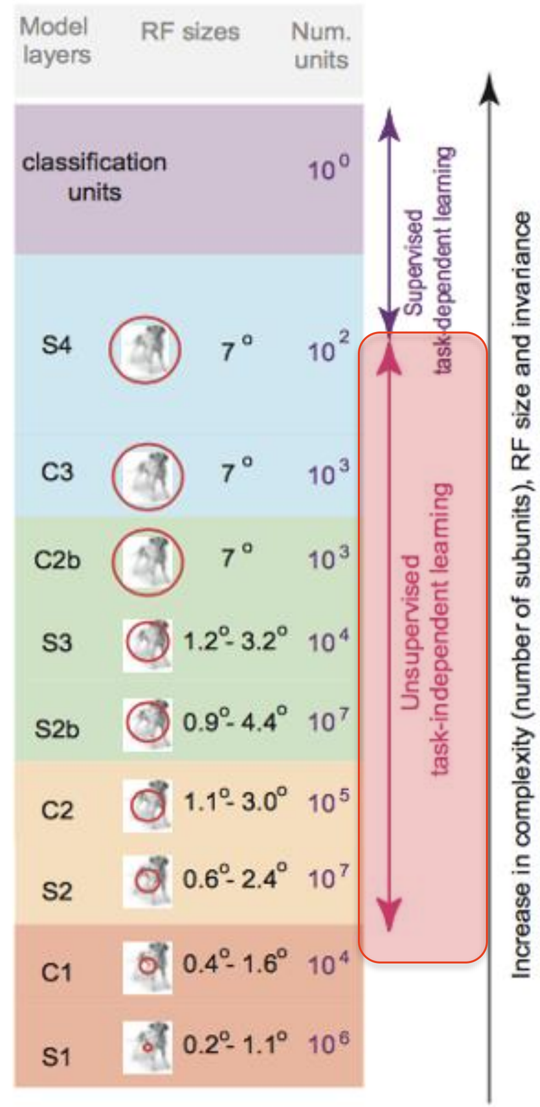

"A feedforward architecture accounts for rapid categorization", T. Serre, A. Oliva & T. Poggio PNAS (2007)
"The Receptive Fields of the Inferior Temporal Cortex Neurons in Natural Scenes", E. Rolls, N. Aggelopoulos & F. Zheng, Journal of Neuroscience (2003)
"Visual cortical mechanisms of perceptual grouping: interacting layers, networks, columns, and maps", W.D. Ross, S. Grossberg, E. Mingolla

Visual Foraging and Whole of Scene Perception

There is a data problem for human (and machine) learning

Serre et al's hierarchical model of vision:

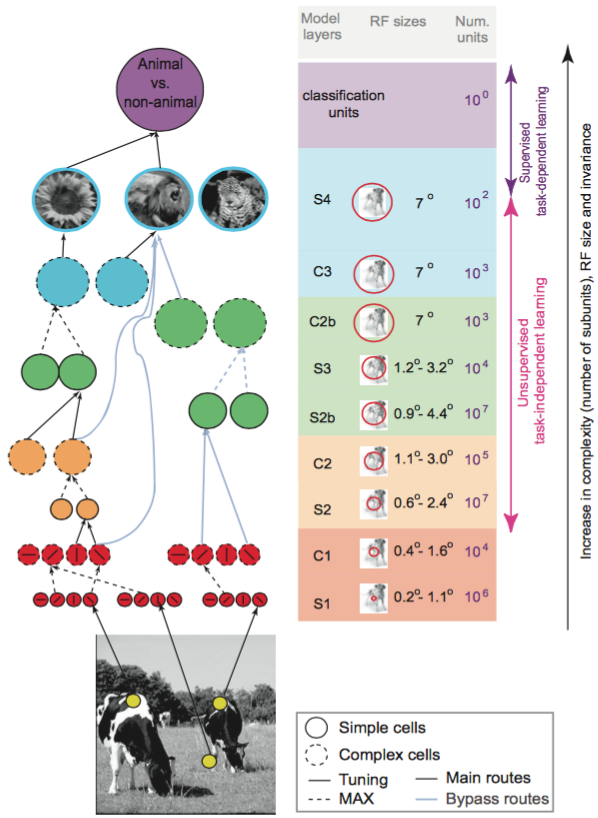

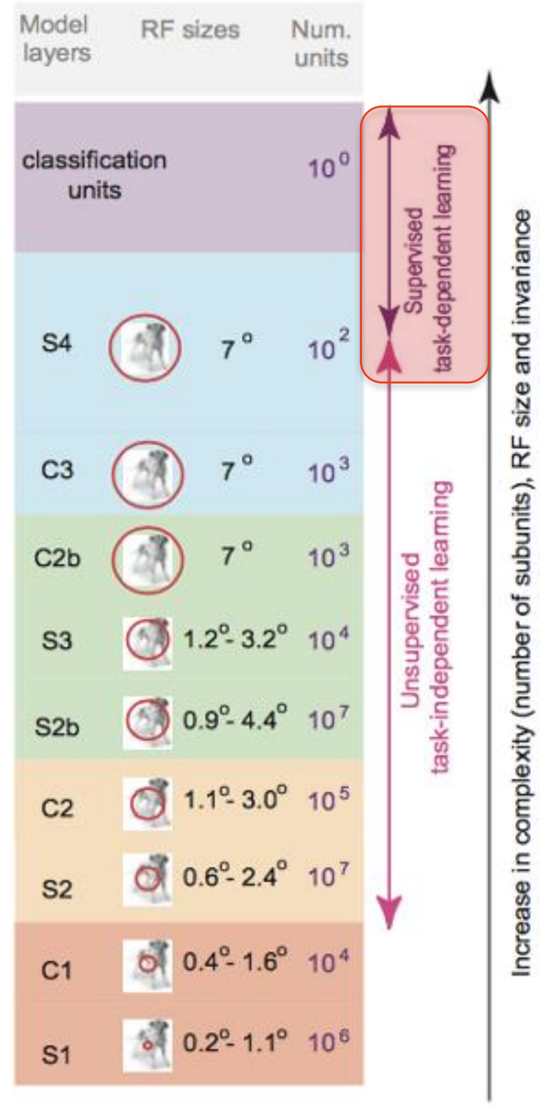

"A feedforward architecture accounts for rapid categorization", T. Serre, A. Oliva & T. Poggio PNAS (2007)
"The Receptive Fields of the Inferior Temporal Cortex Neurons in Natural Scenes", E. Rolls, N. Aggelopoulos & F. Zheng, Journal of Neuroscience (2003)
"Visual cortical mechanisms of perceptual grouping: interacting layers, networks, columns, and maps", W.D. Ross, S. Grossberg, E. Mingolla

Visual Foraging and Whole of Scene Perception

There is a data problem for human (and machine) learning

Serre et al's hierarchical model of vision:

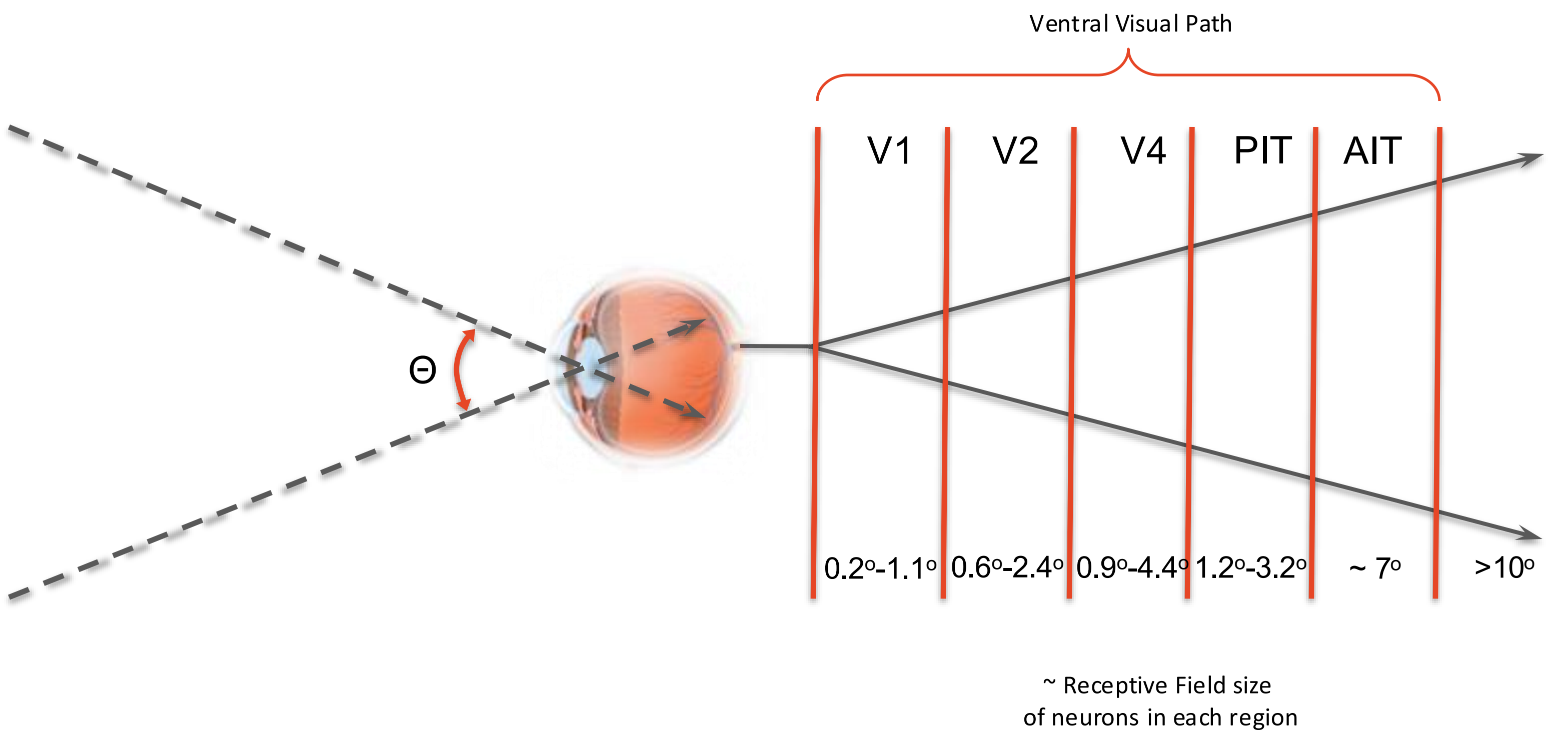

Visual Foraging and Whole of Scene Perception

There is a data problem for human (and machine) learning

Serre et al's hierarchical model of vision:

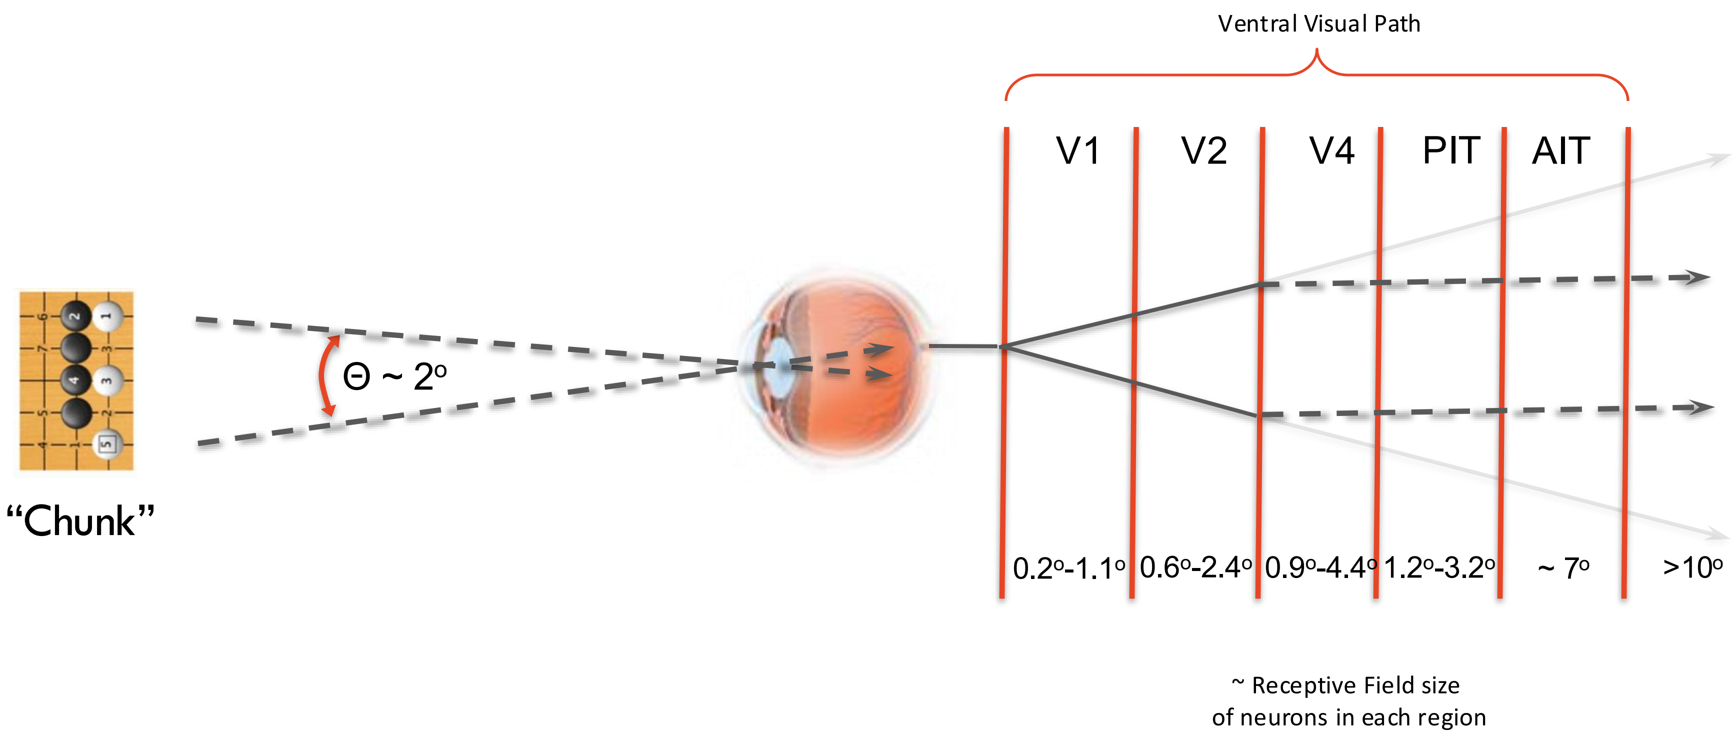

Visual Foraging and Whole of Scene Perception

There is a data problem for human (and machine) learning

Serre et al's hierarchical model of vision:

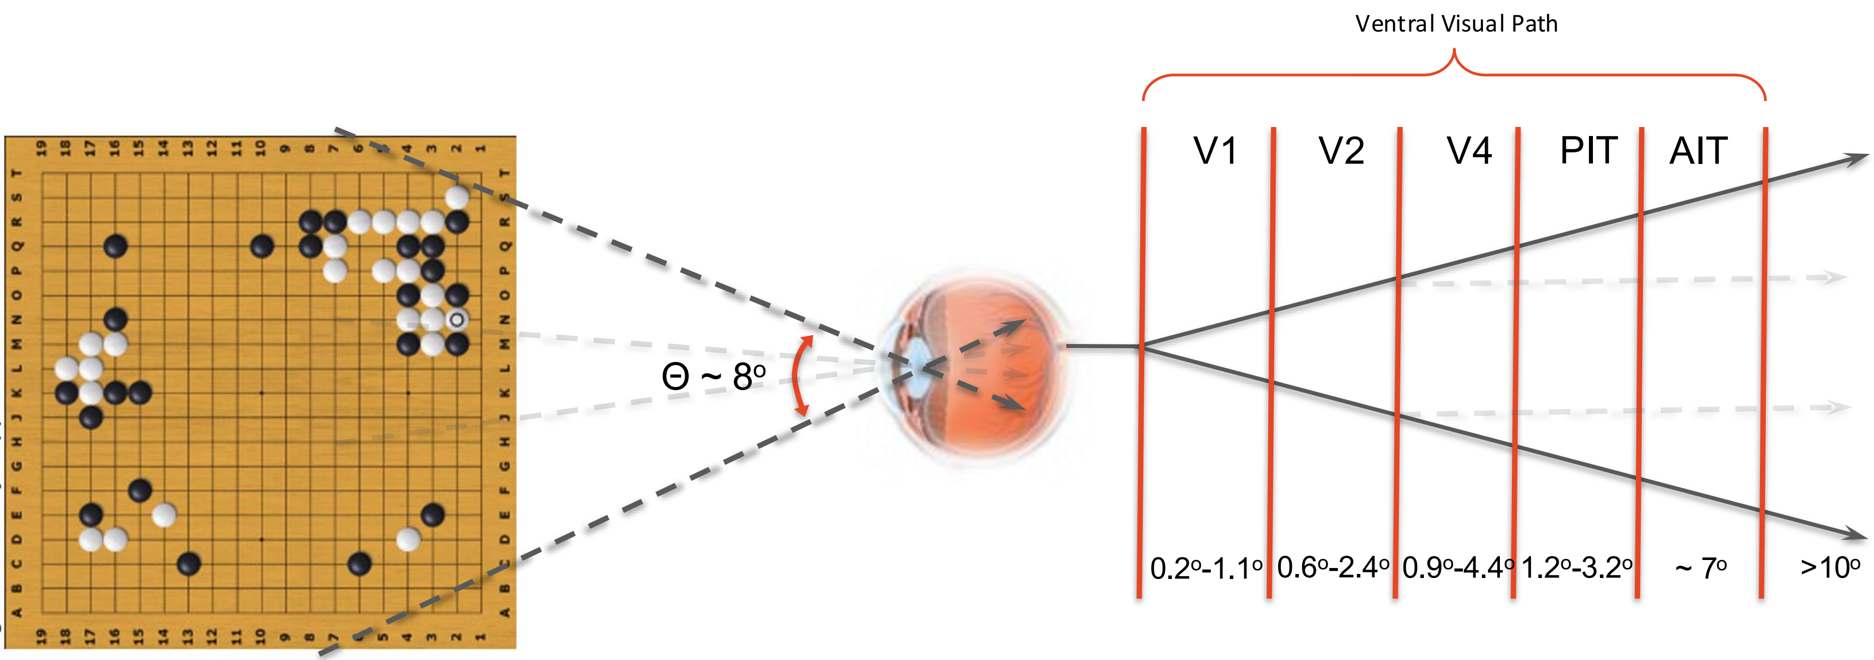

Whole scene

Visual Foraging and Whole of Scene Perception

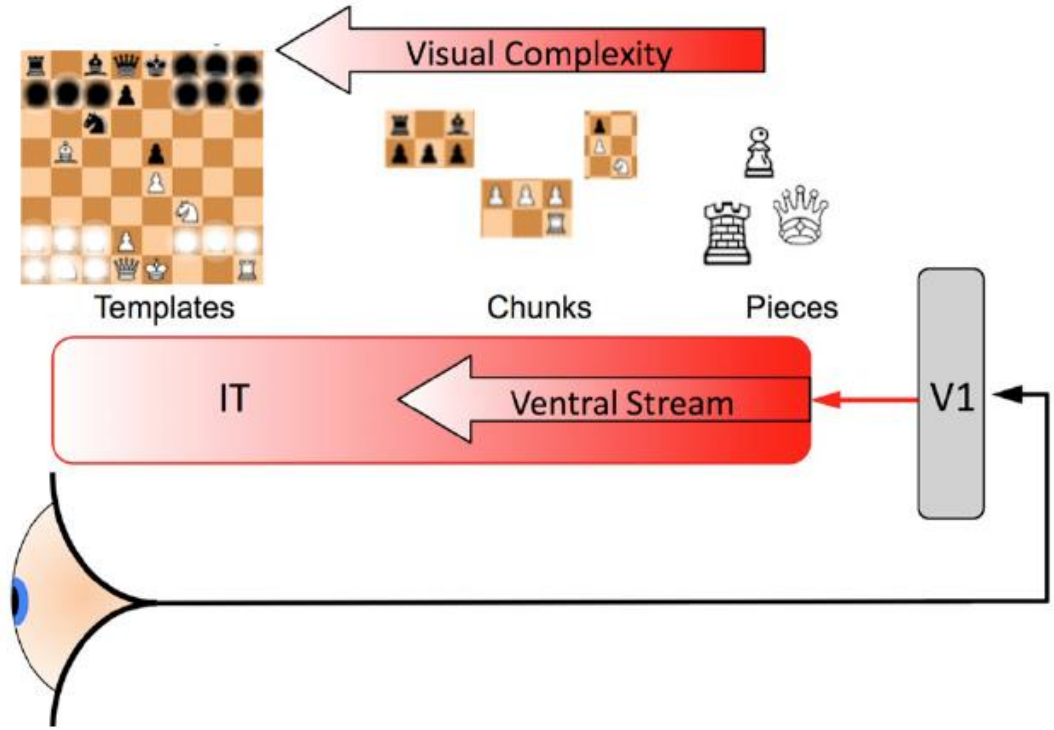

FIGURE 2 | A representation of the category formation mechanism.

"The neural circuitry of expertise: perceptual learning and social cognition", M. Harré, Frontiers in Human Neuroscience (2013)

Visual Foraging and Whole of Scene Perception

Fig. 4

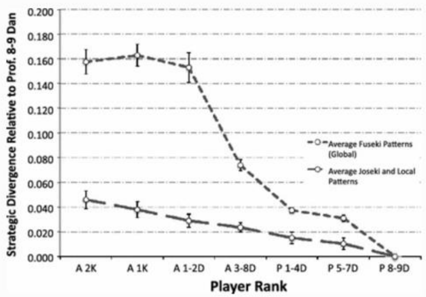

The strategic divergence for global and local patterns. Errors are ± s.e.m

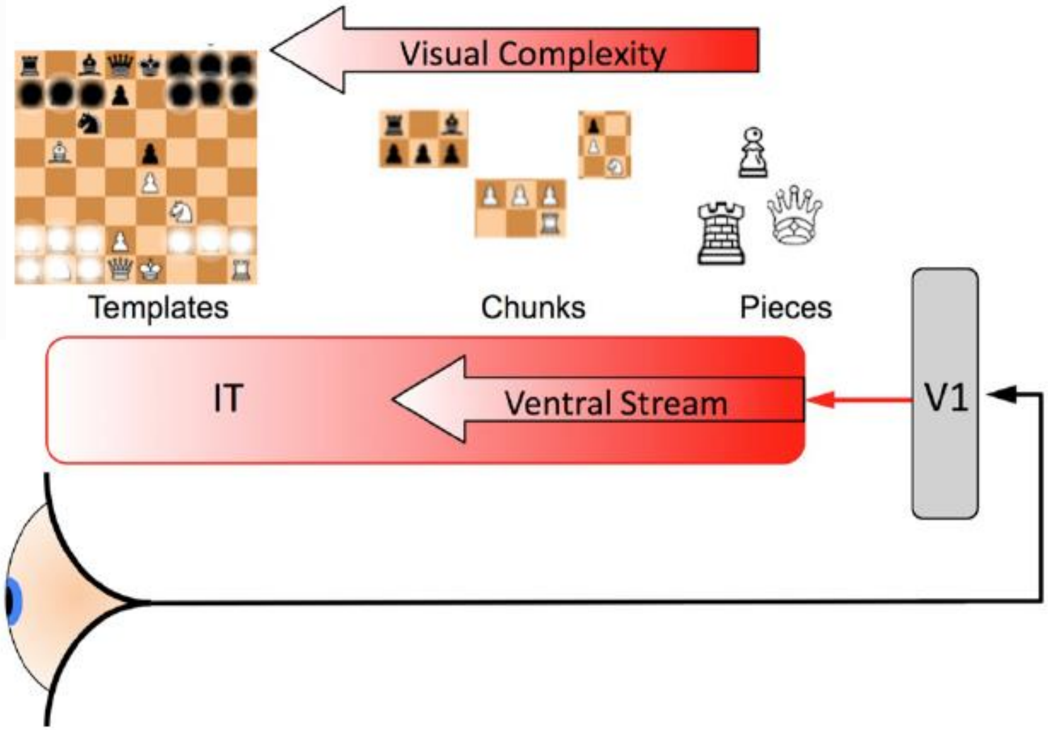

FIGURE 2 | A representation of the category formation mechanism.

“The neural circuitry of expertise: perceptual learning and social cognition”, M. Harré, Frontiers in Human Neuroscience (2013)

Visual Foraging and Whole of Scene Perception

There is a data problem for human (and machine) learning

How this looks in practice: novices forage for information in the environment

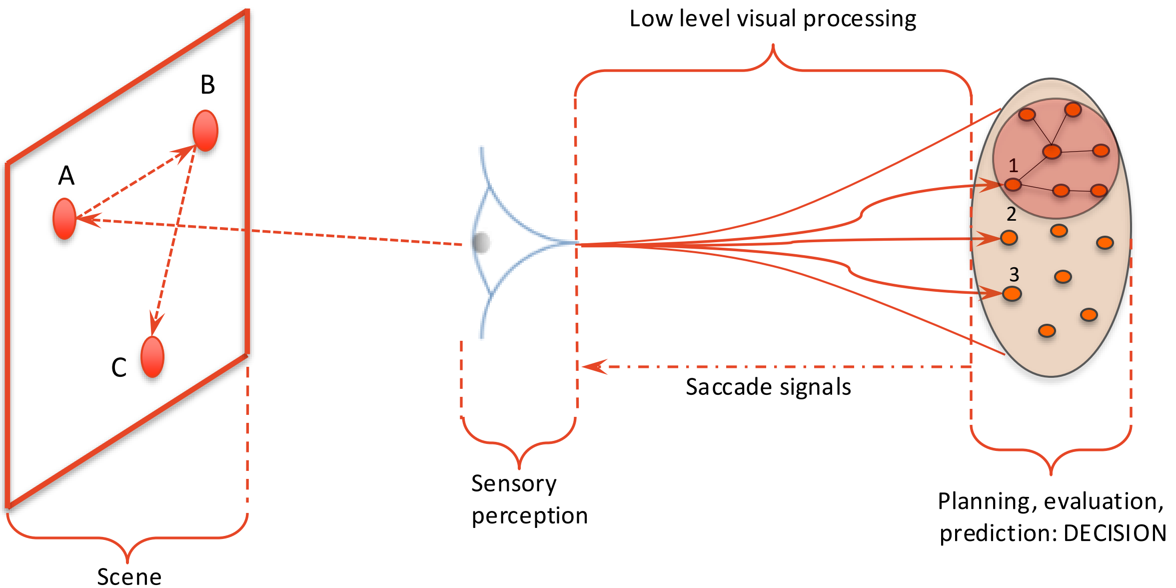

Visual Foraging and Whole of Scene Perception

There is a data problem for human (and machine) learning

How this looks in practice: experts contextualise their foraging using visual cues ("templates")

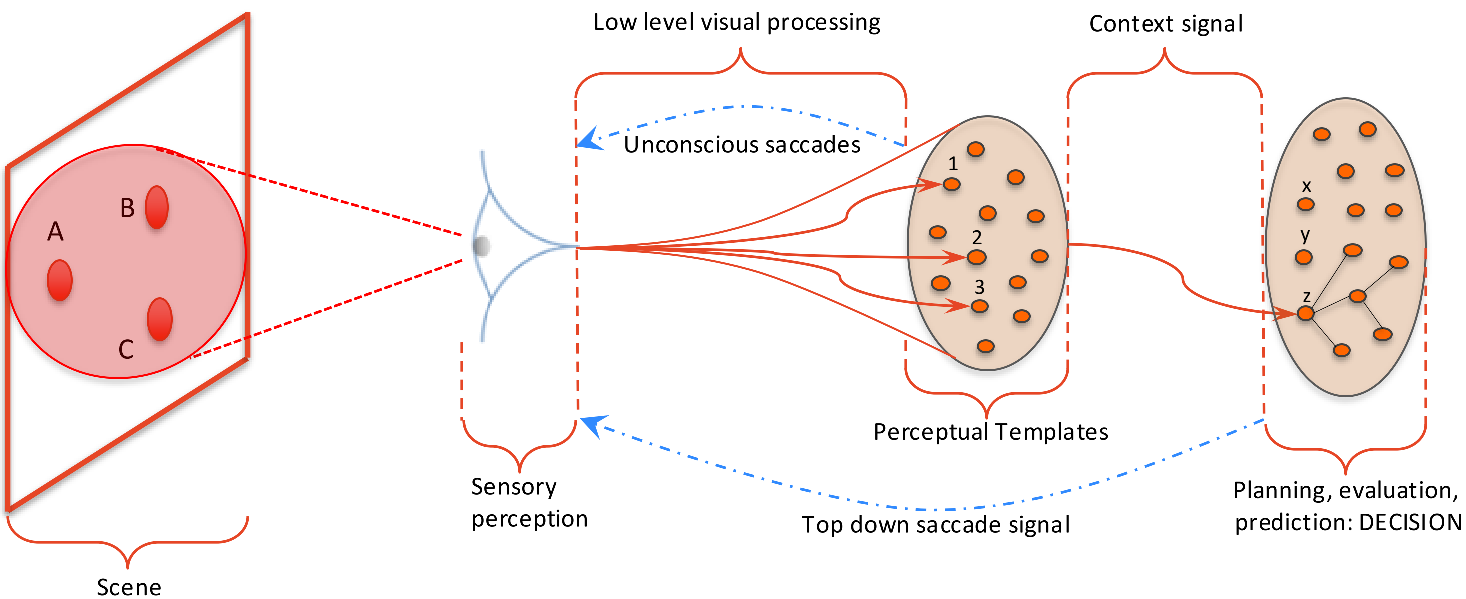

Originally developed in M. Harré T. Bossoamier & A. Snyder, "The perceptual cues that reshape expert reasoning", Scientific Reports (2012)

Extracting Context From Complex Data

A 1.

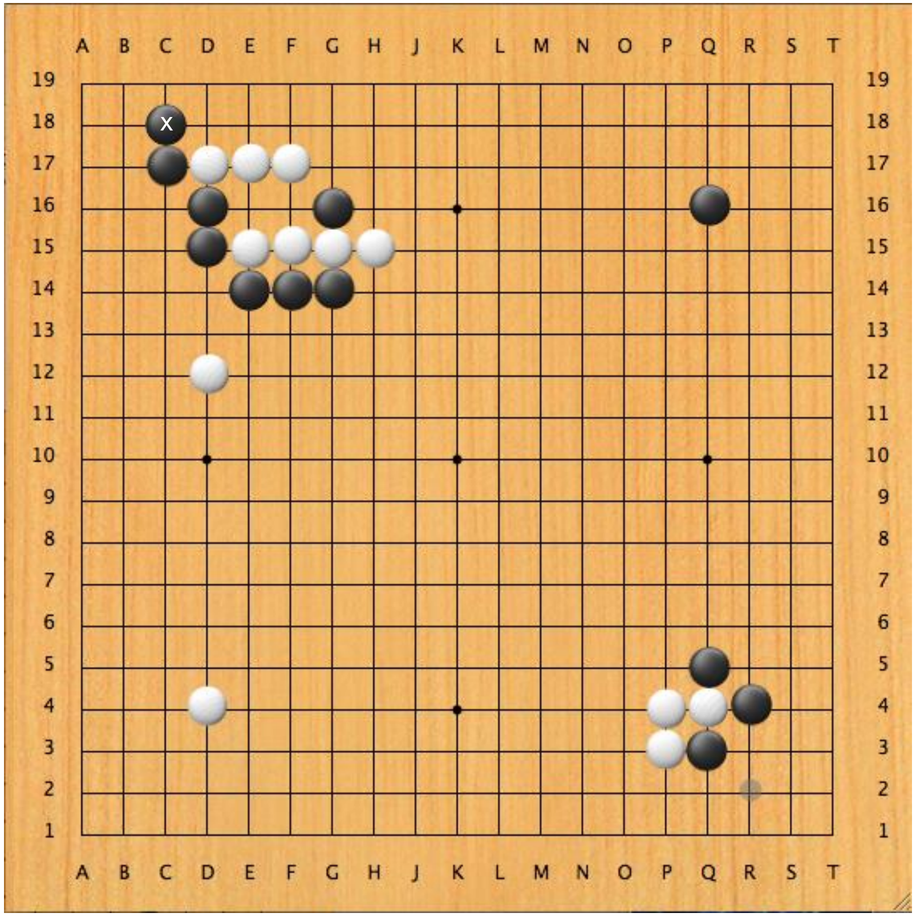

Originally developed in M. Harré T. Bossoamier & A. Snyder, "The perceptual cues that reshape expert reasoning", Scientific Reports (2012)

Extracting Context From Complex Data

A 2.

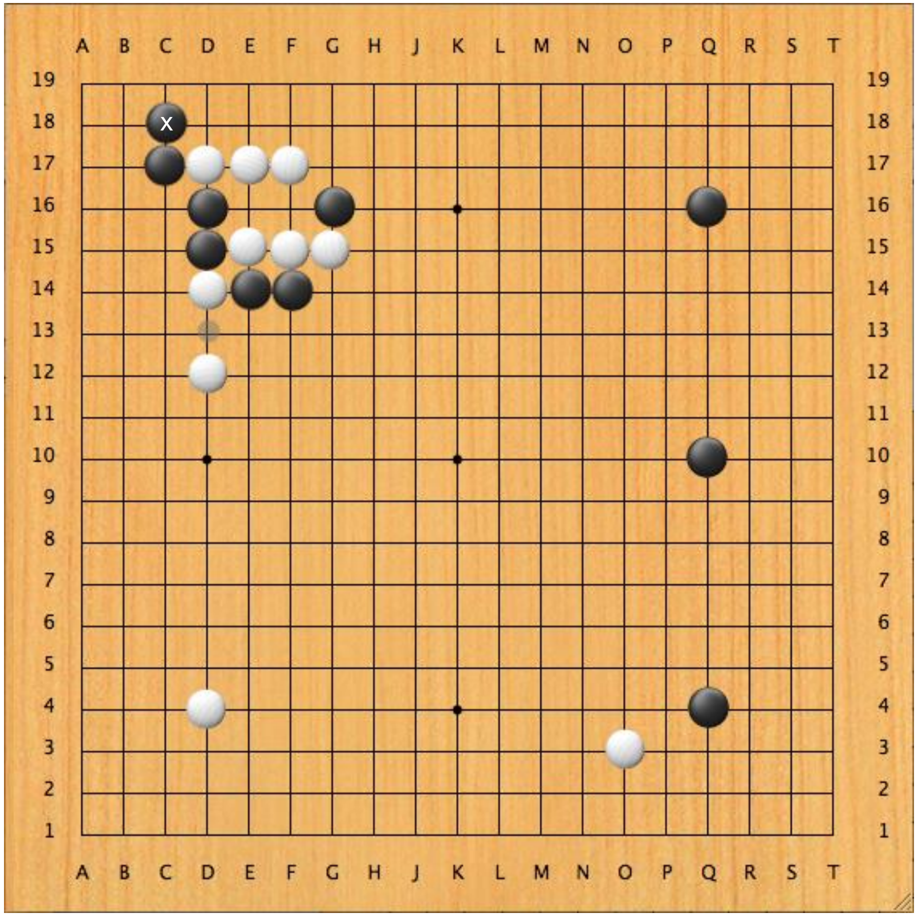

Originally developed in M. Harré T. Bossoamier & A. Snyder, “The perceptual cues that reshape expert reasoning”, Scientific Reports (2012)

Extracting Context From Complex Data

A 3.

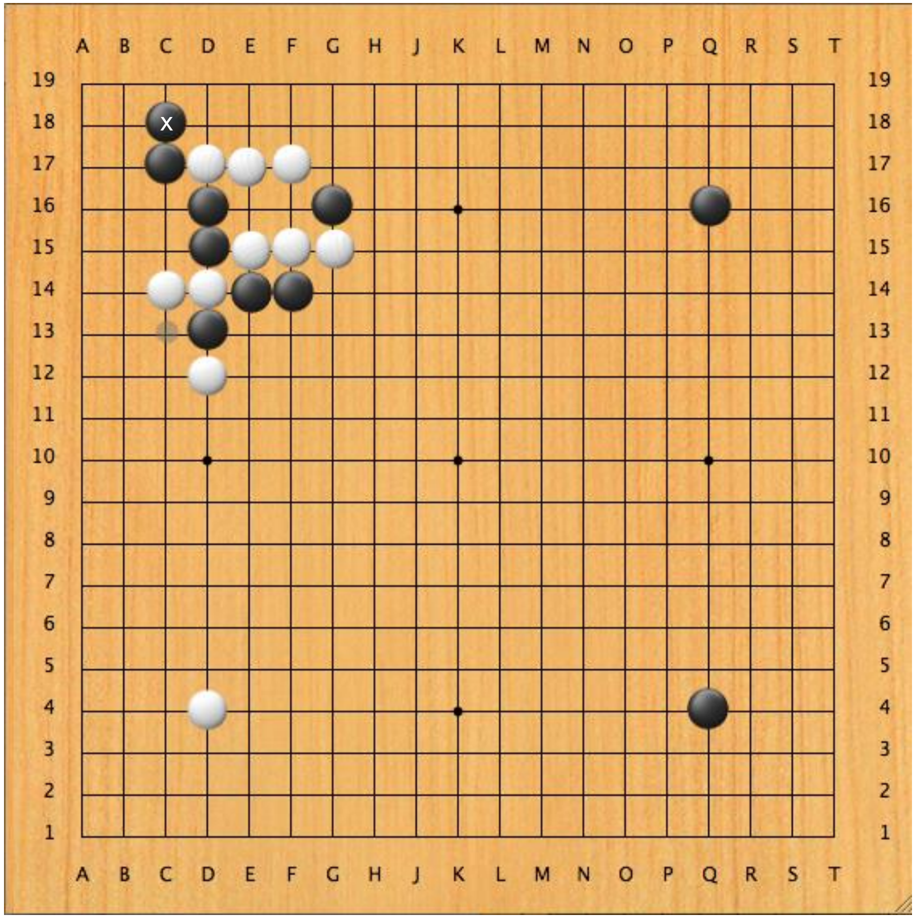

Originally developed in M. Harré T. Bossoamier & A. Snyder, “The perceptual cues that reshape expert reasoning”, Scientific Reports (2012)

Extracting Context From Complex Data

B 1.

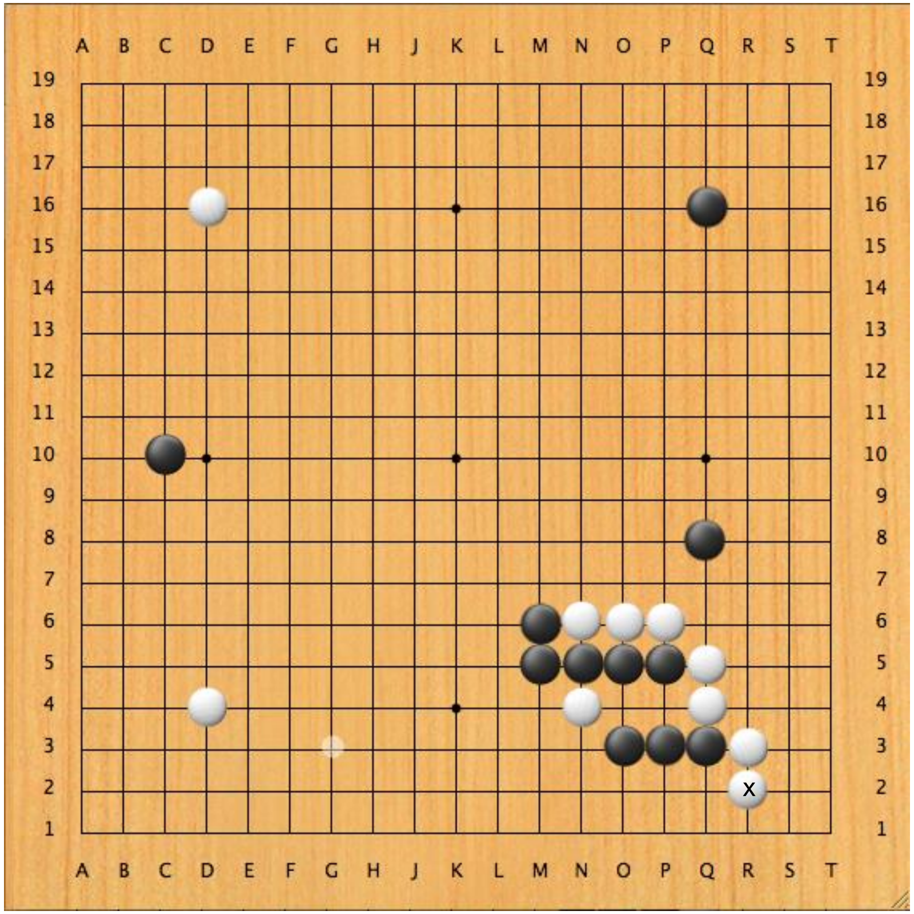

 Originally developed in M. Harré T. Bossoamier & A. Snyder, “The perceptual cues that reshape expert reasoning”, Scientific Reports (2012)

Extracting Context From Complex Data

B 2.

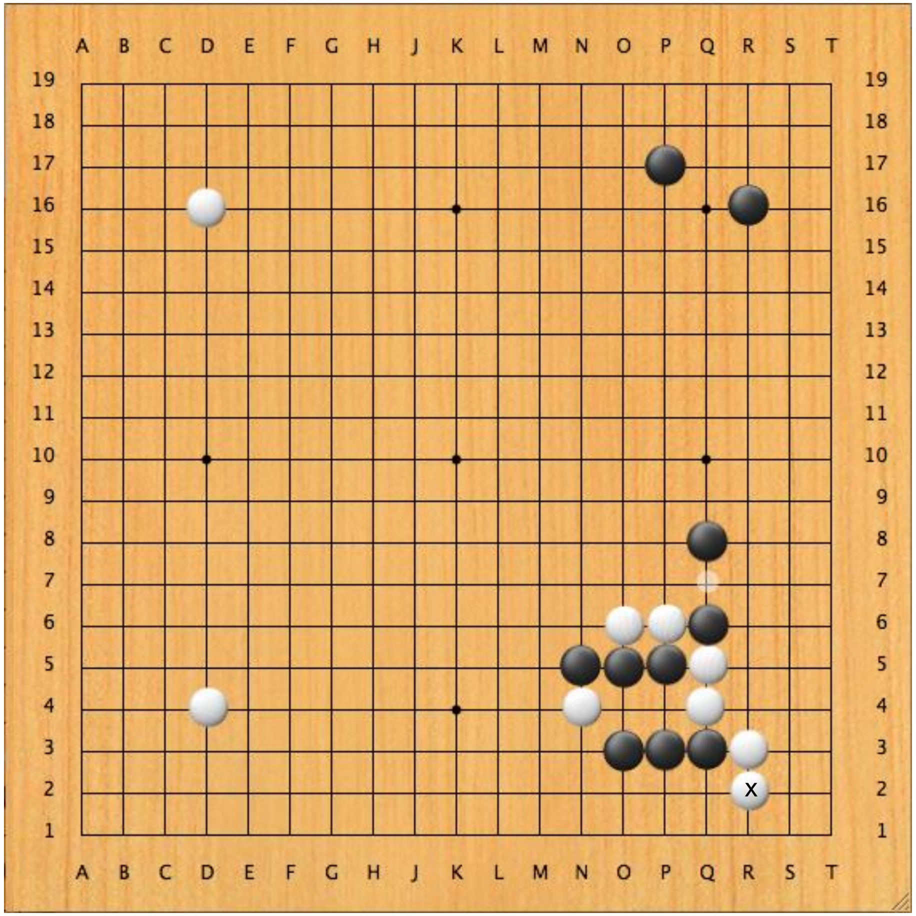

 Originally developed in M. Harré T. Bossoamier & A. Snyder, “The perceptual cues that reshape expert reasoning”, Scientific Reports (2012)

Extracting Context From Complex Data

B 3.

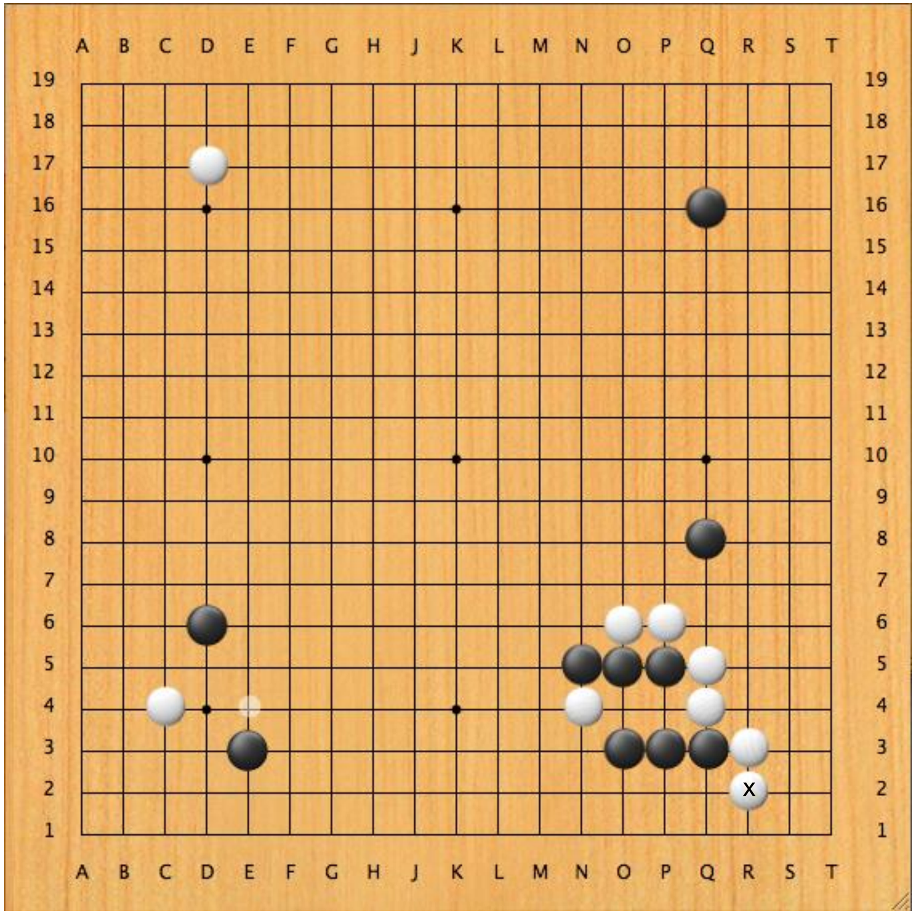

 Originally developed in M. Harré T. Bossoamier & A. Snyder, “The perceptual cues that reshape expert reasoning”, Scientific Reports (2012)

Extracting Context From Complex Data

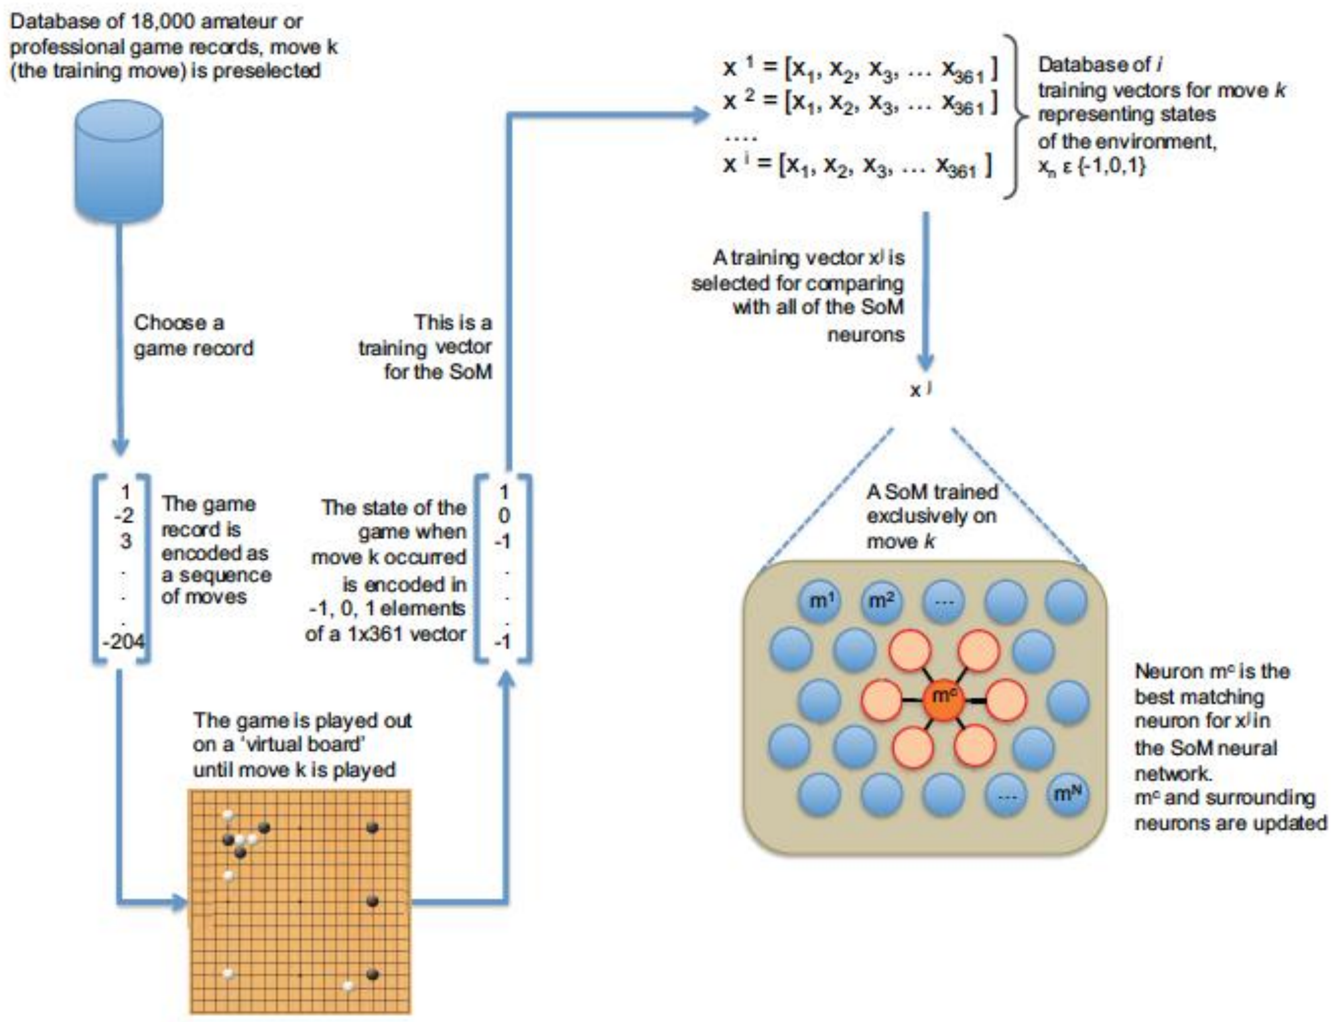

 Originally developed in M. Harré T. Bossoamier & A. Snyder, "The perceptual cues that reshape expert reasoning", Scientific Reports (2012)

Extracting Context From Complex Data

Learning the co-occurring elements results in *Perceptual Templates.*

This is just one such template encoded in a single SoM neuron:

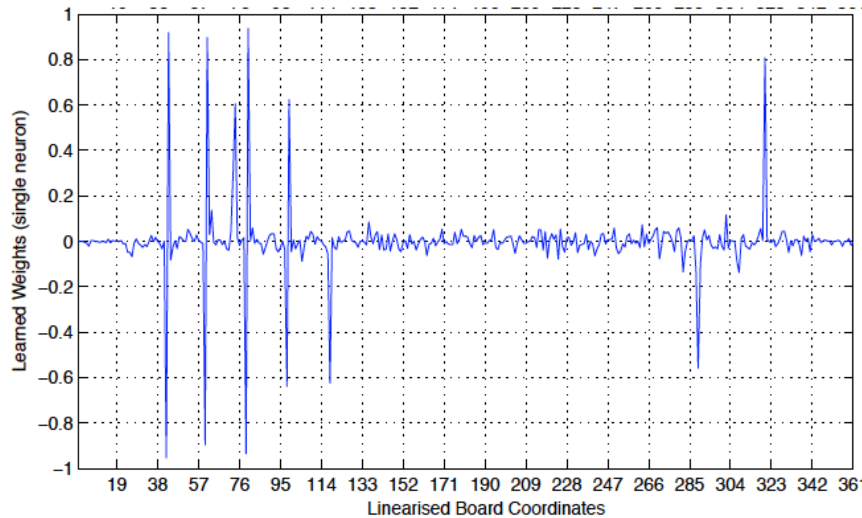

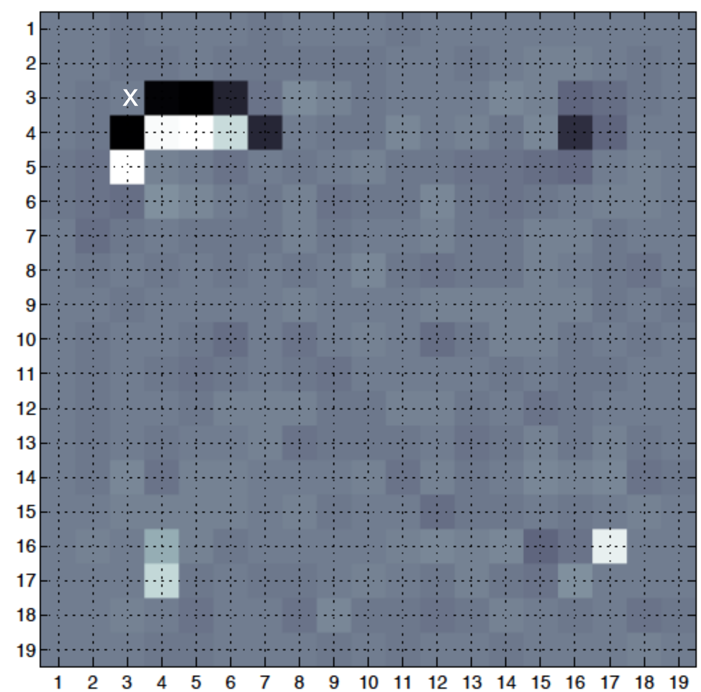

M. Harré & A. Snyder, "Intuitive Expertise and Perceptual Templates" (2011)

Extracting Context From Complex Data

The brain as a ***stochastic*** evidence accumulator for two alternatives

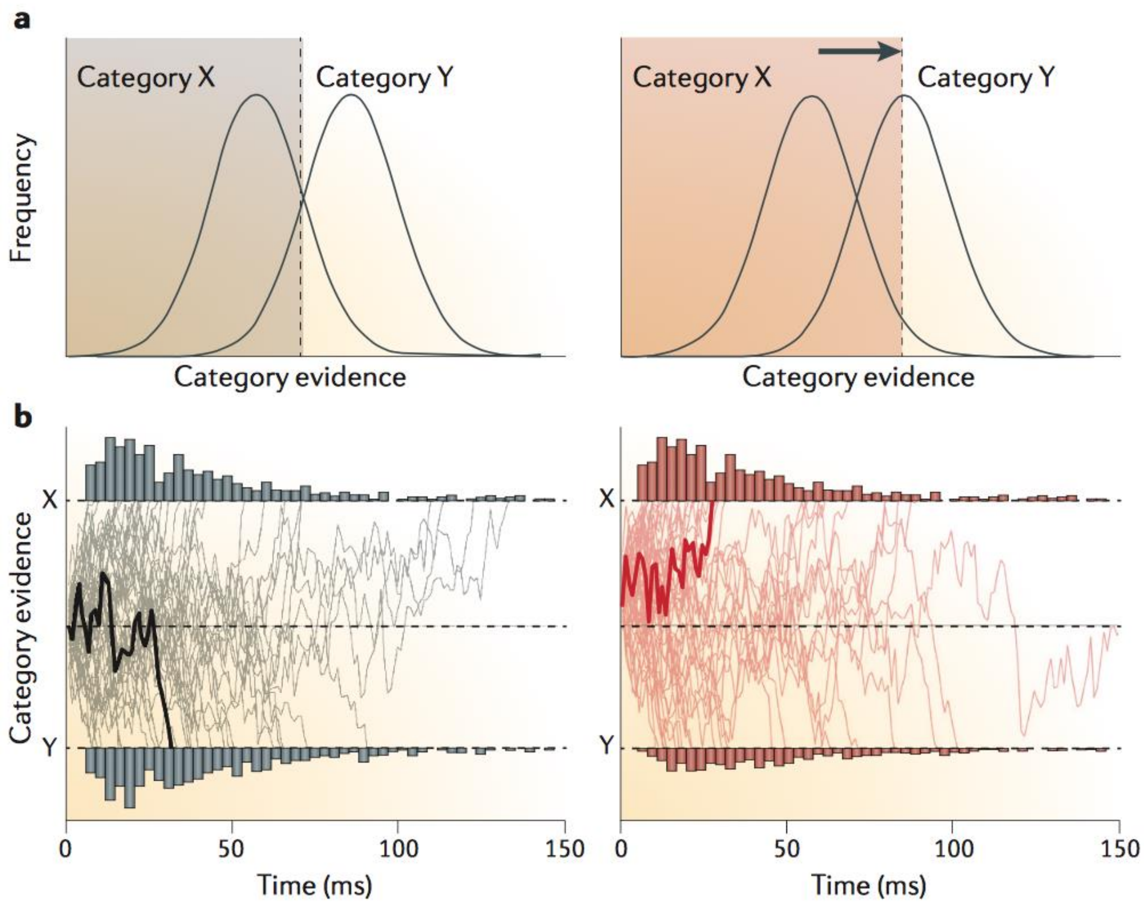

Figure 2 | **Decision-theoretic approaches to understanding expectation.**

"Expectation in Neural Decision Making: Neural and Computational Mechanisms",
C. Summerfield and F. de Lange, Nature Reviews | Neuroscience (2014)

Extracting Context From Complex Data

$\left\{b_i^c\right\}$ = a set of individual stones on the board in positions i and each of colour $c \in \{\text{black,white}\}$

fix(a_x) = hold stone at position i fixed, this is the *action* for which we want the *contex*

Using a custering algorithm, Self-organising map (SoM) in this case, compute the neurons belonging to each cluster for a given a_x, producing a SoM with k neurons in each of y clusters:

$$\text{Clust}^y\left[\left\{b_i^c\right\} \middle| \text{fix}(a_x)\right] = \left\{b_i^c\right\}_{j=1\ldots k}^{y}$$

The number of SoM neurons in each cluster is the "strength" of the cluster, i.e. how likely it is to occur.

Extracting Context From Complex Data

$$\mathrm{Clust}^{y}\left[\left\{b_i^c\right\}\middle|\,\mathrm{fix}(a_x)\right] = \left\{b_i^c\right\}_{j=1\,\ldots k}^{y}$$

$$\frac{\int_j \left\{b_i^c\right\}_{j=1\,\ldots k}^{y}\,\mathrm{d}j}{\int_y \int_j \left\{b_i^c\right\}_{j=1\,\ldots k}^{y}\,\mathrm{d}j\,\mathrm{d}y} = p\left(\left\{b_i^c\right\}^{y}\middle|\,a_x\right)$$

$$= p\left(\mathrm{context}^{y}\middle|\,a_x\right)$$

where integrals count the neurons. This is the probability of a context (whole of scene perception) $\left\{b_i^c\right\}^{y}$ being present when move a_x is made on the board.

$$\int_x p\left(\mathrm{context}^{y}\middle|\,a_x\right) = p\left(\mathrm{context}^{y},\, a_x\right)$$ i.e. the joint probability of a move and a context.

Extracting Context From Complex Data

Similarity matrices of Go templates for Go: Amateurs vs. Experts

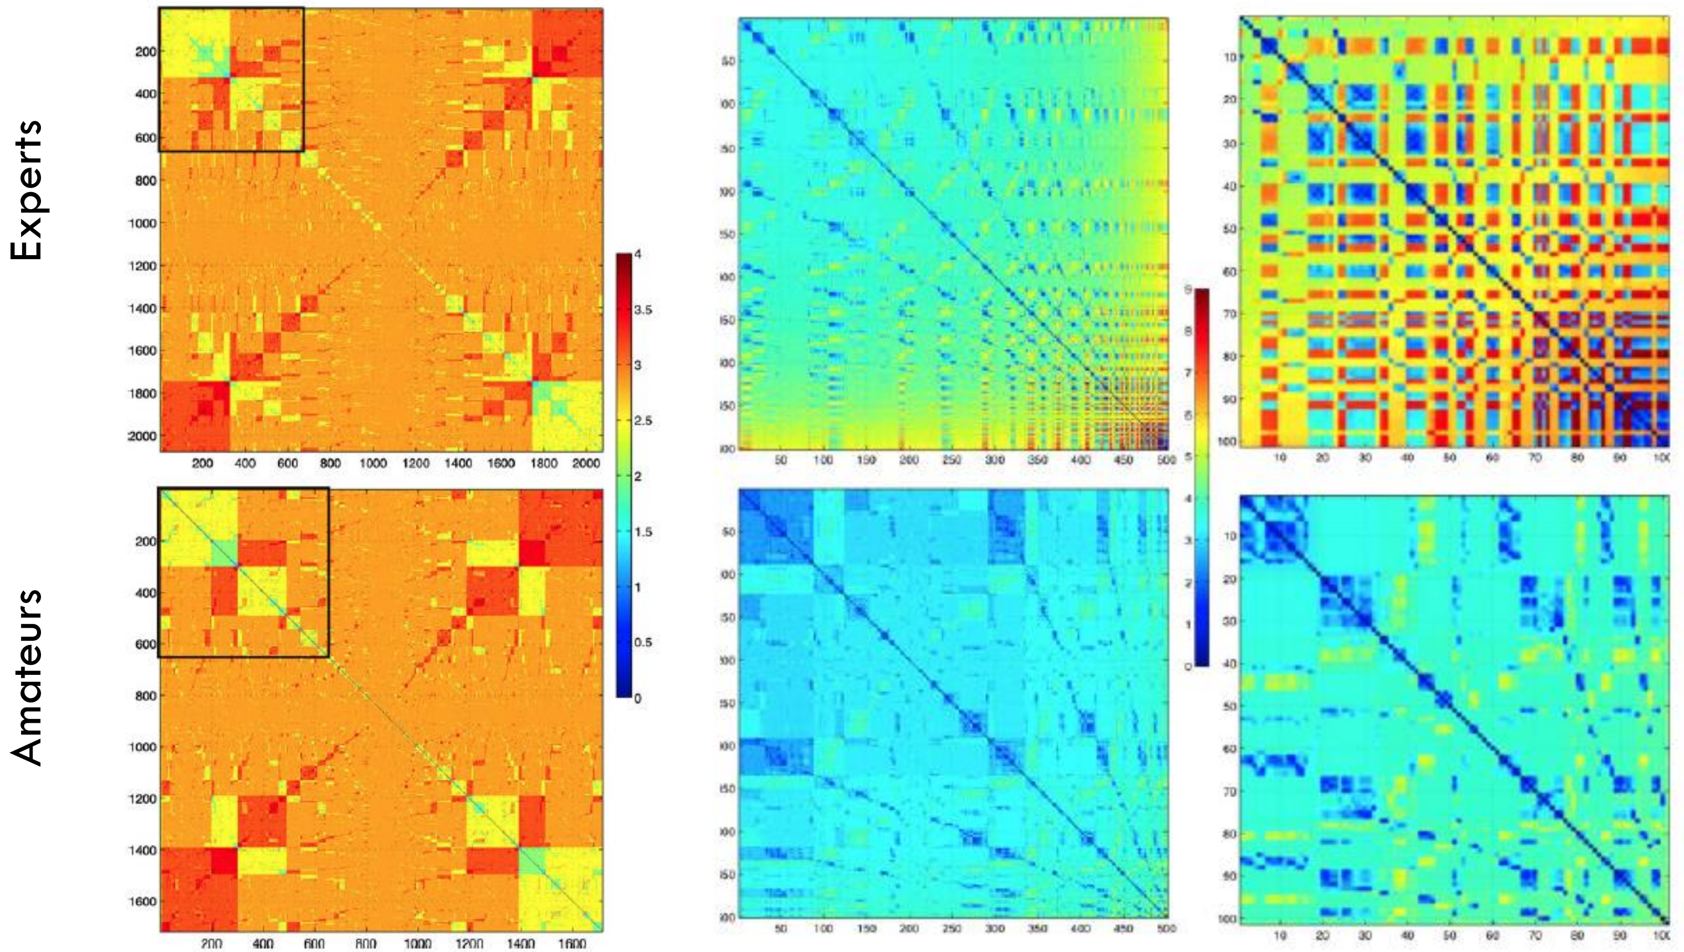

M. Harré, “From Amateur to Professional: A Neuro-cognitive Model of Categories and Expert Development”, Minds and Machines (2013)

Extracting Context From Complex Data

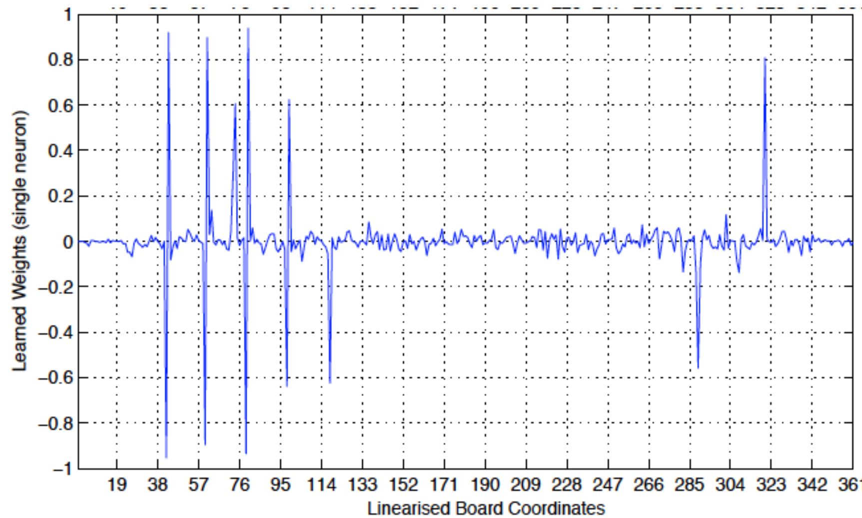

Learning the co-occurring elements results in *Perceptual Templates*:

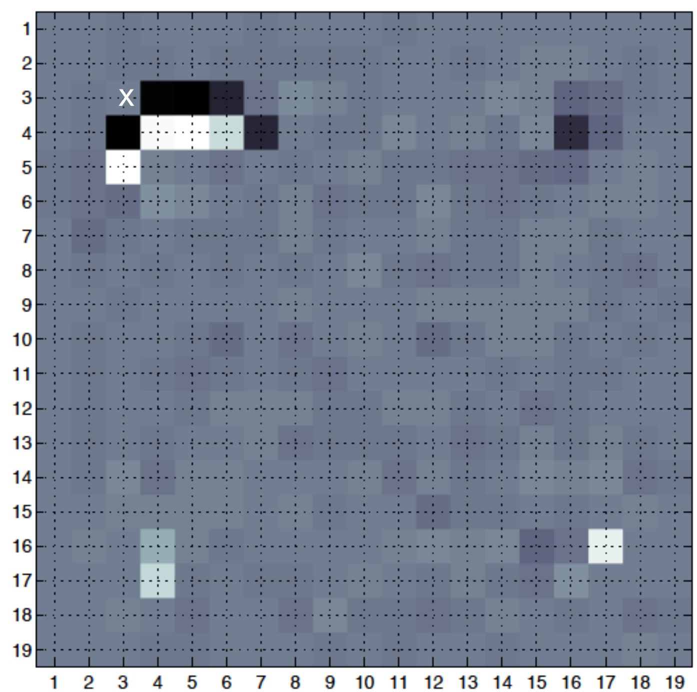

M. Harré & A. Snyder, "Intuitive Expertise and Perceptual Templates" (2011)

Utility value and the Free Energy Principle

Free Energy, Free Utility:

$$\left.\begin{aligned} H(p) &= -\sum_x p(x)\log(p(x)) \\ V(p) &= \text{internal energy (potential function)} \\ \mathcal{F}(p) &= V(p) - T\,H(p) \quad \text{(where } T \text{ is temperature)} \end{aligned}\right\} \text{Helmholtz Free Energy} \qquad (1)$$

Utility value and the Free Energy Principle

Free Energy, Free Utility:

$$\left.\begin{aligned}
H(p) &= -\sum_{x} p(x)\log(p(x)) \\
V(p) &= \text{internal energy (potential function)} \\
\mathcal{F}(p) &= V(p) - T\,H(p) \quad \text{(where } T \text{ is temperature)}
\end{aligned}\right\}\ \text{Helmholtz Free Energy} \quad (1)$$

$$\left.\begin{aligned}
H(P) &= -\sum_{x,y} P(x,y)\log(P(x,y)) \\
U_a(P) &= E_p[\,U_a(x,y)\,]\ \text{(expected utility for } a\text{)} \\
\mathcal{F}(P) &= U_a(P) - T\,H(P) \quad (T \text{ is uncertainty or error})
\end{aligned}\right\}\ \text{Free Utility (2 players)} \quad (2)$$

Strategic Choice of Preferences: the Persona Model

David Wolpert, *Julian Jamison*, *David Newth* and *Michael Harre*

From the journal The B.E. Journal of Theoretical Economics

https://doi.org/10.2202/1935-1704.1593

Utility value and the Free Energy Principle

$$\left.\begin{aligned}
H(P) &= -\sum_{x,y} P(x,y)\log(P(x,y)) \\
U_a(P) &= E_p[\; U_a(x,y)\;] \text{ (expected utility for } a\text{)} \\
\mathcal{F}(P) &= U_a(P) - T\,H(P) \quad (T \text{ is uncertainty or error})
\end{aligned}\right\} \text{Free Utility (2 players)} \quad (2)$$

Optimising these "free functionals" leads to standard exponential solutions:

$$P(x) \propto \exp(-\beta\, V(P))$$

$$\begin{aligned}
P(x \mid x = x_i) &\propto \exp(\,\beta\, U_a(\,P(y) \mid x = x_i\,)\,) \\
P(y \mid y = y_i) &\propto \exp(\,\beta\, U_b(\,P(x) \mid y = y_i\,)\,)
\end{aligned}$$

Strategic Choice of Preferences: the Persona Model

David Wolpert, *Julian Jamison*, *David Newth* and *Michael Harre*

From the journal The B.E. Journal of Theoretical Economics
https://doi.org/10.2202/1935-1704.1593

Utility value and the Free Energy Principle

Published: 13 January 2010

The free-energy principle: a unified brain theory?

Karl Friston

Nature Reviews Neuroscience **11**, 127–138 (2010) | Cite this article

Friston's Free Energy Principle:

One of the goals of Friston's work is to estimate the joint probability of states observed o and actual states s via Bayes Theorem:

$$P(s, o) = P(o|s)P(s)$$

Utility value and the Free Energy Principle

Published: 13 January 2010

The free-energy principle: a unified brain theory?

Karl Friston

Nature Reviews Neuroscience **11**, 127–138 (2010) | Cite this article

Friston's Free Energy Principle:

One of the goals of Friston's work is to estimate the joint probability of states observed o and actual states s via Bayes Theorem:

$$P(s,o) = P(o|s)P(s)$$

This calculation is often too difficult to compute directly so Friston's "Free Energy Principle" for the brain addresses this by estimating an alternative probability $Q(s)$ via opitmisation:

$$Q^*(s) = \underset{Q(s)}{\operatorname{argmin}} \mathcal{F}(\, Q \,) \tag{3}$$

$$Q^*(s) \simeq P(s|o) \tag{4}$$

Utility value and the Free Energy Principle

Published: 13 January 2010

The free-energy principle: a unified brain theory?

Karl Friston

Nature Reviews Neuroscience **11**, 127–138 (2010) | Cite this article

Friston's Free Energy Principle:

One of the goals of Friston's work is to estimate the joint probability of states observed o and actual states s via Bayes Theorem:

$$P(s,o) \; = \; P(o|s)P(s)$$

This calculation is often too difficult to compute directly so Friston's "Free Energy Principle" for the brain addresses this by estimating an alternative probability $Q(s)$ via opitmisation:

$$Q^{*}(s) \; = \; \underset{Q(s)}{\operatorname{argmin}} \, \mathcal{F}(\, Q \,) \tag{3}$$

$$Q^{*}(s) \; \simeq \; P(s|o) \tag{4}$$

$$\mathcal{F}(\, Q \,) \; = \; E_Q[\, \log(Q(s)) \; - \; \log(\, P(s|o) \,) \,] \tag{5}$$

$$= \; \underbrace{E_Q(\, -\log(P(s|o)) \,)}_{\substack{\text{cross entropy} \\ = \text{expected log loss}}} \; - \; \underbrace{H(\, Q(s) \,)}_{\text{entropy}} \tag{6}$$

Utility value and the Free Energy Principle

Information, Utility and Bounded Rationality

Pedro Alejandro Ortega and Daniel Alexander Braun

Department of Engineering, University of Cambridge
Trumpington Street, Cambridge, CB2 1PZ, UK
{dab54,pao32}@cam.ac.uk

Path Integral Control and Bounded Rationality

Daniel A. Braun, Univ. Southern California, Los Angeles, USA, dab54@cam.ac.uk
Pedro A. Ortega, Univ. Cambridge, Cambridge, UK, peortega@dcc.uchile.cl
Evangelos Theodorou, Univ. Southern California, Los Angeles, USA, etheodor@usc.edu
Stefan Schaal, Univ. Southern California, Los Angeles, USA, sschaal@usc.edu

1. **Control.** Given an initial policy represented by the probability measure $\mathbf{P}_i$ and the constraint utilities $\mathbf{U}_*$, we are looking for the final system $\mathbf{P}_f$ that optimizes the trade-off between utility and resource costs. That is,

$$\mathbf{P}_f = \arg\max_{\mathbf{Pr}} \sum_{x\in\mathcal{X}} \mathbf{Pr}(x)\mathbf{U}_*(x) - \alpha \sum_{x\in\mathcal{X}} \mathbf{Pr}(x) \log \frac{\mathbf{Pr}(x)}{\mathbf{P}_i(x)}. \tag{2}$$

$$\arg\max_{\mathbf{P}_f(x)} \left(\sum_{x\in\mathcal{X}} \mathbf{P}_f(x)\mathbf{U}_*(x) - \alpha \sum_{x\in\mathcal{X}} \mathbf{P}_f(x) \log \frac{\mathbf{P}_f(x)}{\mathbf{P}_i(x)} \right). \tag{5}$$

The solution to this variational problem is given by

$$\mathbf{P}_f(x) \propto \mathbf{P}_i(x) \exp\left(\frac{1}{\alpha}\mathbf{U}_*(x)\right).$$

In particular, at very low temperature $\alpha \approx 0$, the maximum expected utility principle is recovered as

$$\mathbf{J}_f - \mathbf{J}_i \approx \sum_{x\in\mathcal{X}} \mathbf{P}_f(x)\mathbf{U}_*(x),$$

and hence resource costs are ignored in the choice of $\mathbf{P}_f$, leading to $\mathbf{P}_f \approx \delta_{x^*}(x)$, where $x^* = \arg\max_x \mathbf{U}_*(x)$. Similarly, at a high temperature, the difference is

$$\mathbf{J}_f - \mathbf{J}_i \approx -\alpha \sum_{x\in\mathcal{X}} \mathbf{P}_f(x) \log \frac{\mathbf{P}_f(x)}{\mathbf{P}_i(x)},$$

Utility value and the Free Energy Principle

What is epistemic value in free energy models of learning and acting? A bounded rationality perspective

Pedro A. Ortega[1]
and Daniel A. Braun[2,3]

$$\max_{P(a)} \operatorname*{ext}_{P(s|a)} \sum_a P(a)\left(-\frac{1}{\beta}D_{KL}[P(s|a)||P_{des}(s|a)] - \frac{1}{\alpha}\log\frac{P(a)}{P_0(a)} - \frac{1}{\beta}\log Z_\beta(a)\right)$$

with the solution $P^*(s|a) = P_0(s|a)\exp\{\beta U(s,a)\}/Z_\beta(a)$

and $P^*(a) = P_0(a)\exp\{\frac{\alpha}{\beta}\log Z_\beta(a)\}/Z_\alpha$

Assuming: $P_{des}(s|a) = P_{des}(s)$ and $\alpha \to \infty$

$$-D_{KL}[P(s|a)||P_{des}(s)] = \underbrace{\mathrm{E}_{P(s|a)}[\log P_{des}(s)]}_{extrinsic-value} + \underbrace{H[P(s|a)]}_{intrinsic-value}$$

Utility value and the Free Energy Principle

What is epistemic value in free energy models of learning and acting? A bounded rationality perspective

Pedro A. Ortega[1]
and Daniel A. Braun[2,3]

$$\max_{P(a)} \operatorname*{ext}_{P(s|a)} \sum_{a} P(a)\left(-\frac{1}{\beta}D_{KL}[P(s|a)||P_{des}(s|a)] - \frac{1}{\alpha}\log\frac{P(a)}{P_0(a)} - \frac{1}{\beta}\log Z_\beta(a)\right)$$

with the solution $P^*(s|a) = P_0(s|a)\exp\{\beta U(s,a)\}/Z_\beta(a)$

and $P^*(a) = P_0(a)\exp\{\frac{\alpha}{\beta}\log Z_\beta(a)\}/Z_\alpha$

Note: $P_{des}(s|a) = P_0(s|a)\exp\{\beta U(s,a)\}/Z_\beta(a)$ by definition (i.e. a prior)

$$-D_{KL}[P(s|a)||P_{des}(s)] = \underbrace{\mathrm{E}_{P(s|a)}[\log P_{des}(s)]}_{extrinsic-value} + \underbrace{H[P(s|a)]}_{intrinsic-value}$$

Utility value and the Free Energy Principle

What is epistemic value in free energy models of learning and acting? A bounded rationality perspective

Pedro A. Ortega[1]
and Daniel A. Braun[2,3]

$$\mathcal{F}(Q) = \underbrace{E_Q(-\log(P(s|o)))}_{\substack{\text{cross entropy} \\ = \text{expected log loss}}} - \underbrace{H(Q(s))}_{\text{entropy}}$$

$$\mathbf{Q}_\tau(\pi) = -\underbrace{D[Q(s_\tau|\pi)||P(s_\tau|\pi)]}_{\text{KL divergence}} = \underbrace{E_{Q(s_\tau|\pi)}[\ln P(s_\tau|m)]}_{\text{Extrinsic value}} + \underbrace{H[Q(s_\tau|\pi)]}_{\text{Epistemic value}}$$

$$-D_{KL}[P(s|a)||P_{des}(s)] = \underbrace{\mathrm{E}_{P(s|a)}[\log P_{des}(s)]}_{extrinsic-value} + \underbrace{H[P(s|a)]}_{intrinsic-value}$$

Utility value and the Free Energy Principle

What is epistemic value in free energy models of learning and acting? A bounded rationality perspective

Pedro A. Ortega[1]
and Daniel A. Braun[2,3]

$$\mathcal{F}(Q) = \underbrace{E_Q(-\log(P(s|o)))}_{\substack{\text{cross entropy} \\ = \text{expected log loss}}} - \underbrace{H(Q(s))}_{\text{entropy}}$$

Friston: no utility

$$\mathbf{Q}_\tau(\pi) = -\underbrace{D[Q(s_\tau|\pi)||P(s_\tau|\pi)]}_{\text{KL divergence}} = \underbrace{E_{Q(s_\tau|\pi)}[\ln P(s_\tau|m)]}_{\text{Extrinsic value}} + \underbrace{H[Q(s_\tau|\pi)]}_{\text{Epistemic value}}$$

$$-D_{KL}[P(s|a)||P_{des}(s)] = \underbrace{\mathrm{E}_{P(s|a)}[\log P_{des}(s)]}_{extrinsic-value} + \underbrace{H[P(s|a)]}_{intrinsic-value}$$

Ortega and Braun: with utility

Utility value and the Free Energy Principle

RESEARCH ARTICLE

Human-level play in the game of *Diplomacy* by combining language models with strategic reasoning

META FUNDAMENTAL AI RESEARCH DIPLOMACY TEAM (FAIR)†, ANTON BAKHTIN, NOAM BROWN, EMILY DINAN, GABRIELE FARINA, COLIN FLAHERTY, DANIEL FRIED, ANDREW GOFF, JONATHAN GRAY, [...], AND MARKUS ZIJLSTRA +17 authors Authors Info & Affiliations

SCIENCE • 22 Nov 2022 • First Release • DOI: 10.1126/science.ade9097

piKL: KL-regularized planning

piKL assumes player i seeks a policy π_i that maximizes the modified utility function

$$U_i(\pi_i, \pi_{-i}) = u_i(\pi_i, \pi_{-i}) - \lambda D_{KL}(\pi_i \parallel \tau_i) \tag{1}$$

where π_{-i} represents the policies of all players other than i, and $u_i(\pi_i, \pi_{-i})$ is the expected value of π_i given that other players play π_{-i}. Specifically, let $Q_i^{t-1}(a_i) = u_i(a_i, \pi_{-i}^{t-i})$ and let

$$\pi_i^{\Delta t}(a_i) \propto \tau_i(a_i) \exp\left[\frac{Q_i^{t-1}(a_i)}{\lambda}\right] \tag{2}$$

On each iteration t, piKL updates its prediction of the players' joint policies to be

$$\pi^t = \left(\frac{t-1}{t}\right)\pi^{t-1} + \left(\frac{1}{t}\right)\pi^{\Delta t} \tag{3}$$

Utility value and the Free Energy Principle

RESEARCH ARTICLE

Human-level play in the game of *Diplomacy* by combining language models with strategic reasoning

META FUNDAMENTAL AI RESEARCH DIPLOMACY TEAM (FAIR)†, ANTON BAKHTIN, NOAM BROWN, EMILY DINAN, GABRIELE FARINA, COLIN FLAHERTY, DANIEL FRIED, ANDREW GOFF, JONATHAN GRAY, [...], AND MARKUS ZIJLSTRA +17 authors Authors Info & Affiliations

SCIENCE • 22 Nov 2022 • First Release • DOI: 10.1126/science.ade9097

Cisero's equivalent of Free Energy is:

$$U(\pi_i, \pi_{-i}) = u(\pi_i, \pi_{-i}) + \lambda\, D_{KL}(\pi_i \mid \tau_i) \tag{12}$$

$$= u(\pi_i, \pi_{-i}) + \lambda \underbrace{\left(-\underbrace{H(\pi_i)}_{\text{Entropy}} + \underbrace{H(\pi_i, \tau_i)}_{\text{cross entropy}}\right)}_{\text{Log–loss game}} \tag{13}$$

$$= \overbrace{u(\pi_i, \pi_{-i}) + \lambda(\underbrace{-H(\pi_i)}_{\text{Entropy}}}^{\text{(local) MaxEnt Game Theory}} + \underbrace{H(\pi_i, \tau_i)}_{\text{cross entropy}}) \tag{14}$$

Bringing These Together in a Dual Process Model

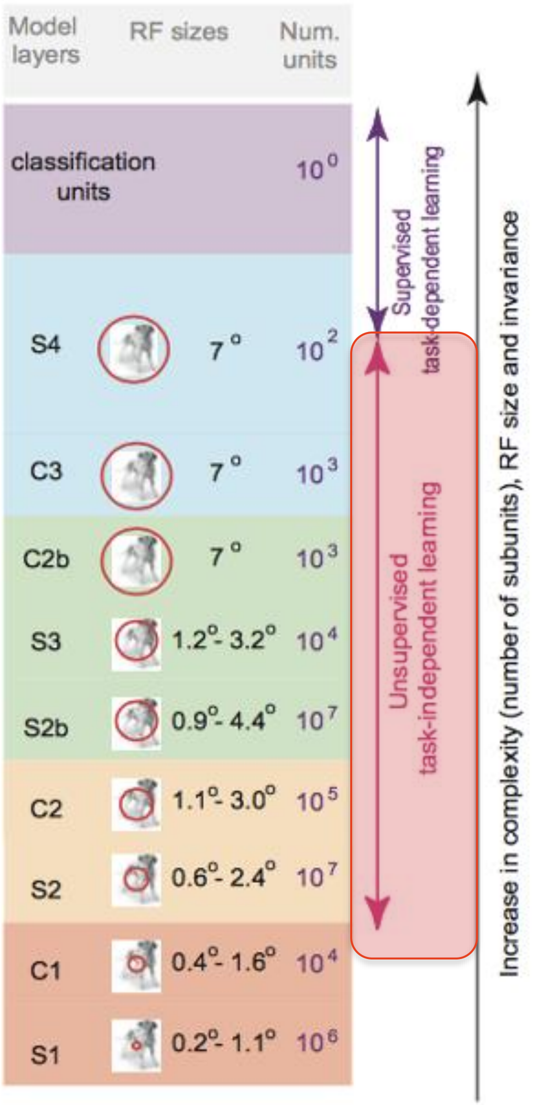

Bringing These Together in a Dual Process Model

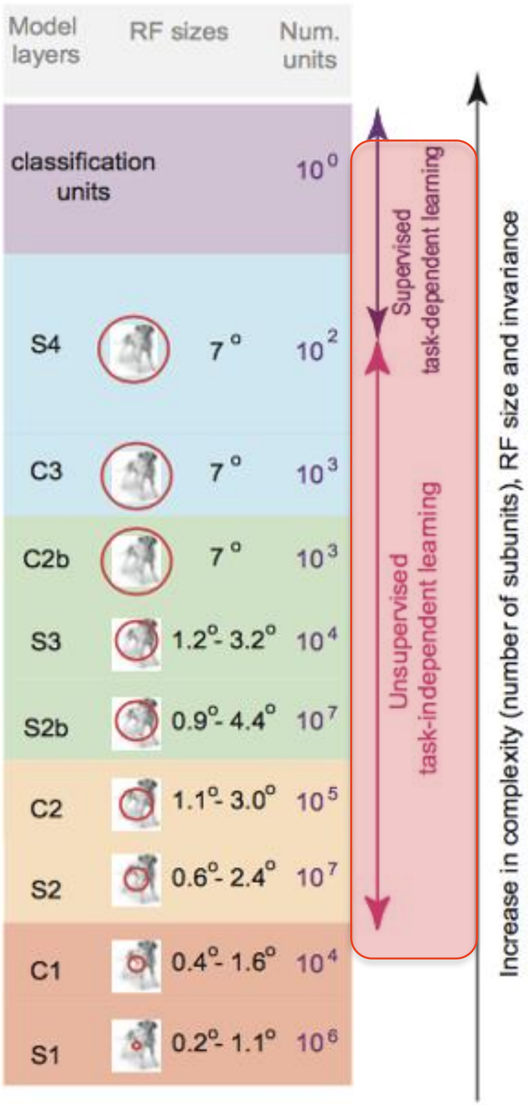

Bringing These Together in a Dual Process Model

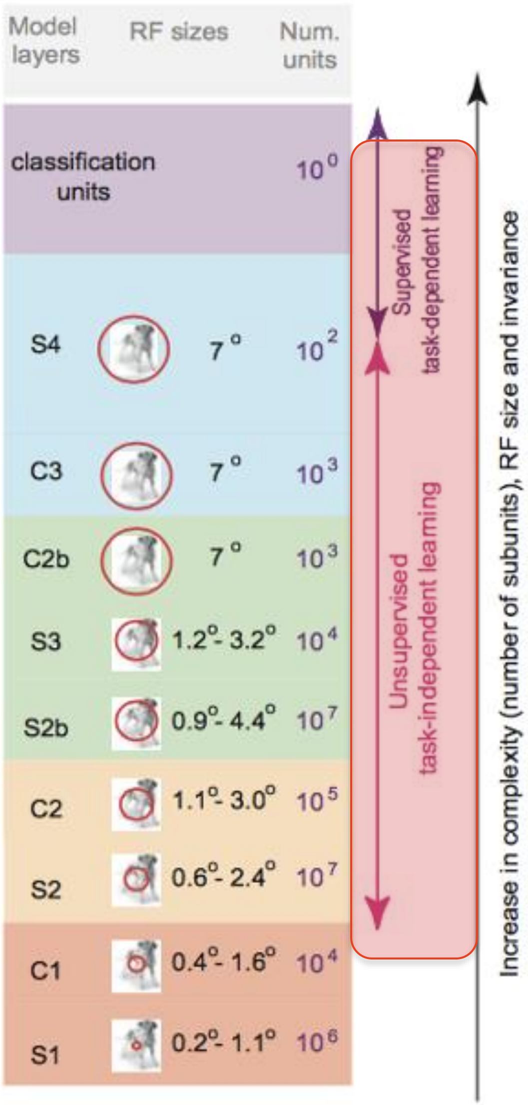

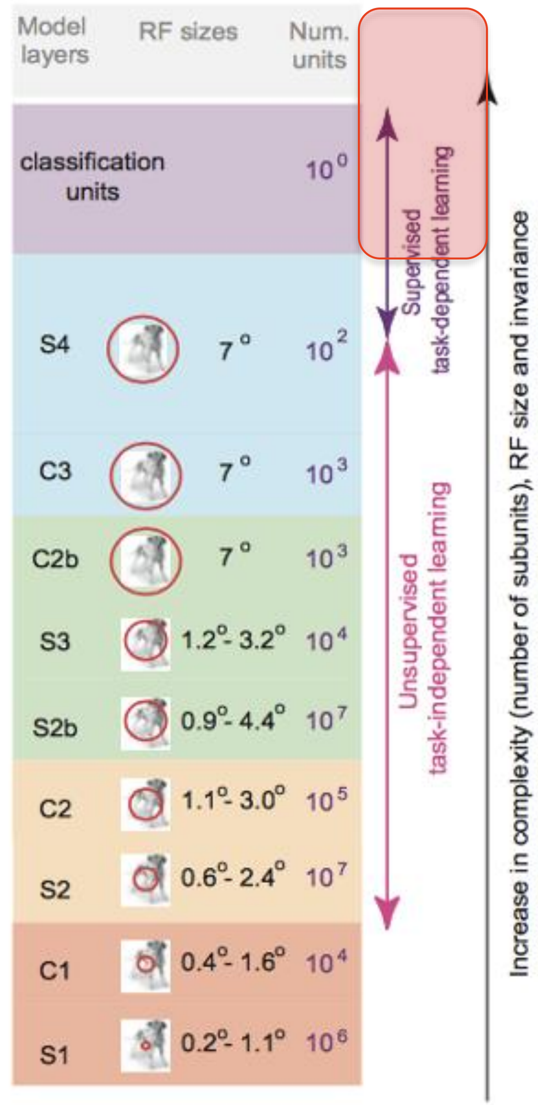

Final word from Herb Simon

In an information-rich world, the wealth of information means a dearth of something else: a scarcity of whatever it is that information consumes. What information consumes is rather obvious: it consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention and a need to allocate that attention efficiently among the overabundance of information sources that might consume it.

Designing Organizations for an Information-Rich World (1971)

Thank you

Questions

THE UNIVERSITY OF
SYDNEY

F