Jour Fixe

28 July 2020

The Cybernetic Bayesian Brain:

From Interoceptive Inference to Sensorimotor Contingencies

A. Seth,

in T. Metzinger & J. M. Windt (Eds). Open MIND: 35(T). Frankfurt am Main: MIND Group,1–24, 2015.
doi: 10.15502/9783958570108

Cognitive Behavior of Humans, Animals, and Machines: Situation Model Perspectives

ZiF Zentrum für interdisziplinäre Forschung
Center for Interdisciplinary Research

A single principle
to account for

perception
action
cognition
consciousness

based on
neural operations

Internal models

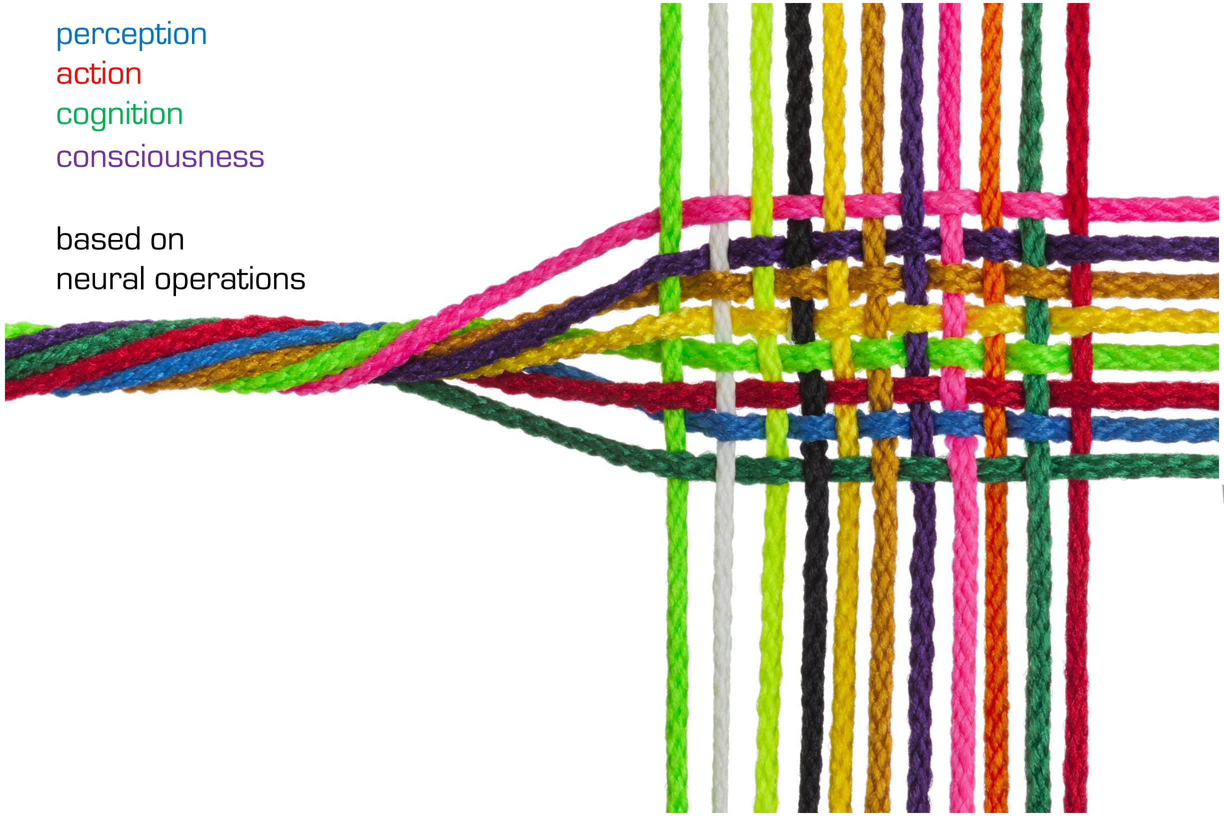

https://fashion-history.lovetoknow.com/fabrics-fibers/weave-types

Bayesian brain: predictive processing (PP) and active inference

Free energy principle: minimization of average surprisal

Cybernetics: self-regulation, adaptive self-regulation, ultrastability

Interoception and emotion

Embodiment

Enactive cognitive science: weak and strong

Sensorimotor contingency (SMC) theory & PPSMC

Counterfactual PP and active inference

A single principle
to account for

perception
action
cognition
consciousness

based on
neural operations

Internal models

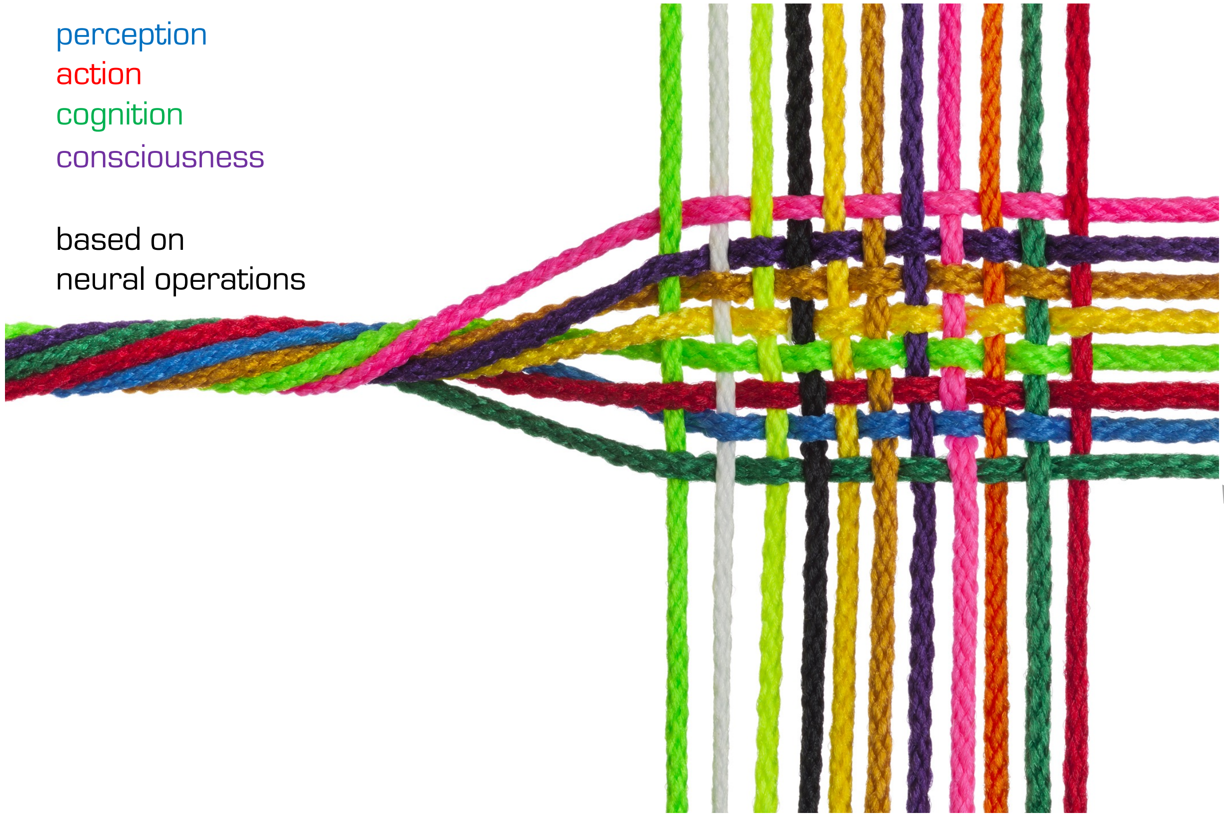

https://fashion-history.lovetoknow.com/fabrics-fibers/weave-types

Bayesian brain: predictive processing (PP) and active inference

Free energy principle: minimization of average surprisal

Cybernetics: self-regulation, adaptive self-regulation, ultrastability

Interoception and emotion

Embodiment

Enactive cognitive science: weak and strong

Sensorimotor contingency (SMC) theory & PPSMC

Counterfactual PP and active inference

The Bayesian Brain

Infer likely causes of sensory inputs

A prediction machine

aka predictive processing (PP)

Hidden Markov Model

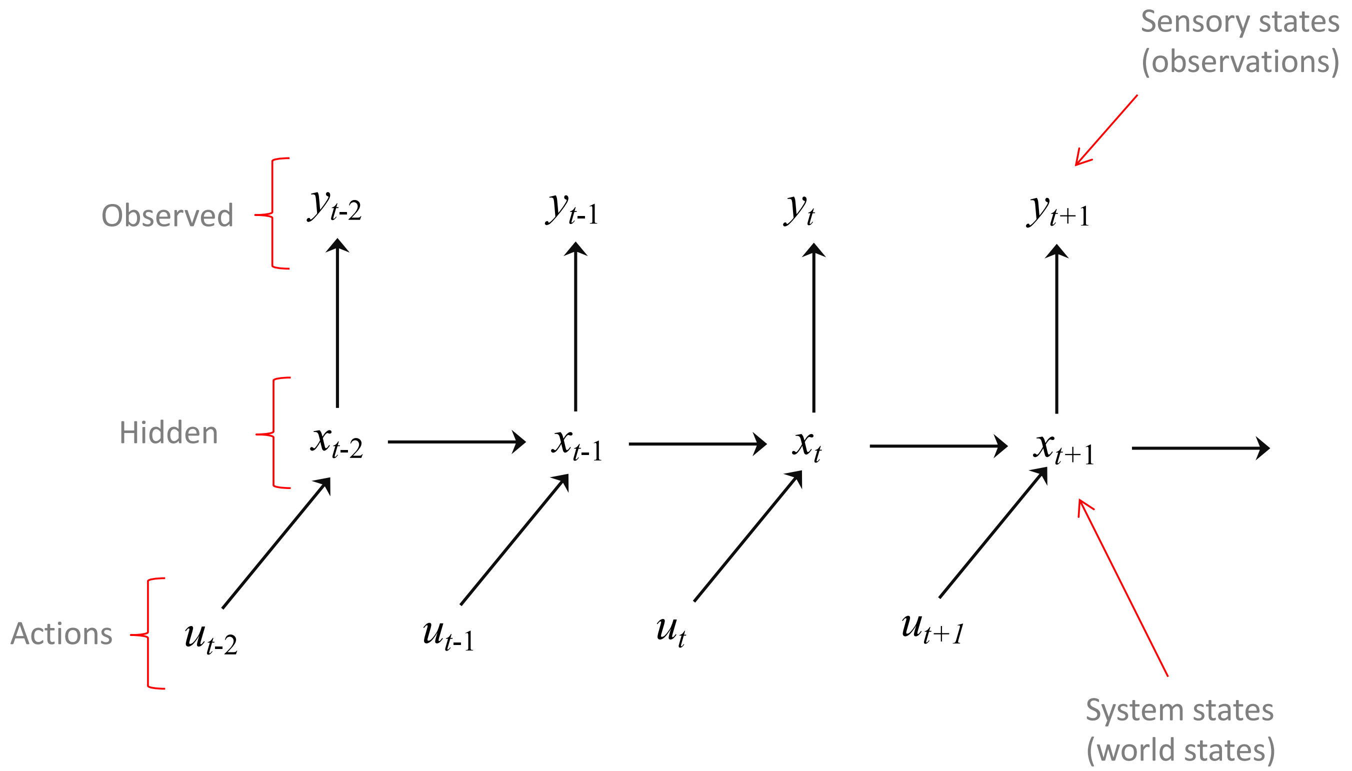

Joint distribution of time series of sensory inputs $\mathbf{y}$,
latent states $\mathbf{z}$, internal states $\mathbf{x}$, and motor signals $\mathbf{u}$.

Internal states

Sensory inputs

$$p\,(\mathbf{x}, \mathbf{z}, \mathbf{y}, \mathbf{u})$$

Latent states

Motor signals

D. McNamee and D. M. Wolpert. Internal models in biological control. Annual Review of Control, Robotics, and Autonomous Systems, 2:339–364, 2019.

Joint distribution of time series of sensory inputs $\mathbf{y}$,
latent states $\mathbf{z}$, internal states $\mathbf{x}$, and motor signals $\mathbf{u}$.

$$p\,(\underbrace{\mathbf{x},\ \mathbf{z}}_{\text{Infer these ...}},\ \underbrace{\mathbf{y},\ \mathbf{u}}_{\text{... from these, for example}})$$

Infer these ...

... from these, for example

D. McNamee and D. M. Wolpert. Internal models in biological control. Annual Review of Control, Robotics, and Autonomous Systems, 2:339–364, 2019.

Generative Model

describes the process that generate sensory data
(c.f. predictive processing)

$$p\,(\mathbf{y}, \mathbf{z})$$

Joint distribution of between sensory input and latent variables

D. McNamee and D. M. Wolpert. Internal models in biological control. Annual Review of Control, Robotics, and Autonomous Systems, 2:339–364, 2019.

Perceptual Inference Model

$$p\,(\mathbf{z} \mid \mathbf{y})$$

Infer world latent states ...

... given the sensory input

D. McNamee and D. M. Wolpert. Internal models in biological control. Annual Review of Control, Robotics, and Autonomous Systems, 2:339–364, 2019.

Bayesian Inference

$$p(z|y) = \frac{p(y|z)p(z)}{p(y)}$$

Infer world latent states
given the sensory input ...

... by combining a prior with a likelihood
in a statistically optimal manner

D. McNamee and D. M. Wolpert. Internal models in biological control. Annual Review of Control, Robotics, and Autonomous Systems, 2:339–364, 2019.

Key Idea: minimize prediction errors

Key Idea: minimize prediction errors

Continually update the model
to minimize the difference between
actual sensory signals

and

signals produced by probabilistic predictive models

"Passively"

Change the model to fit the incoming data:

Perceptual inference

Key Idea: minimize prediction errors

Perform actions to confirm sensory predictions

"Actively"

Change the sampling to suit the prediction:

Active inference

Precision Weighting

Predictions and prediction errors in a Bayesian framework have associated precisions (inverse variances)

An indicator of the reliability of the prediction error

Used to determine the influence of the prediction error when updating the predictive model

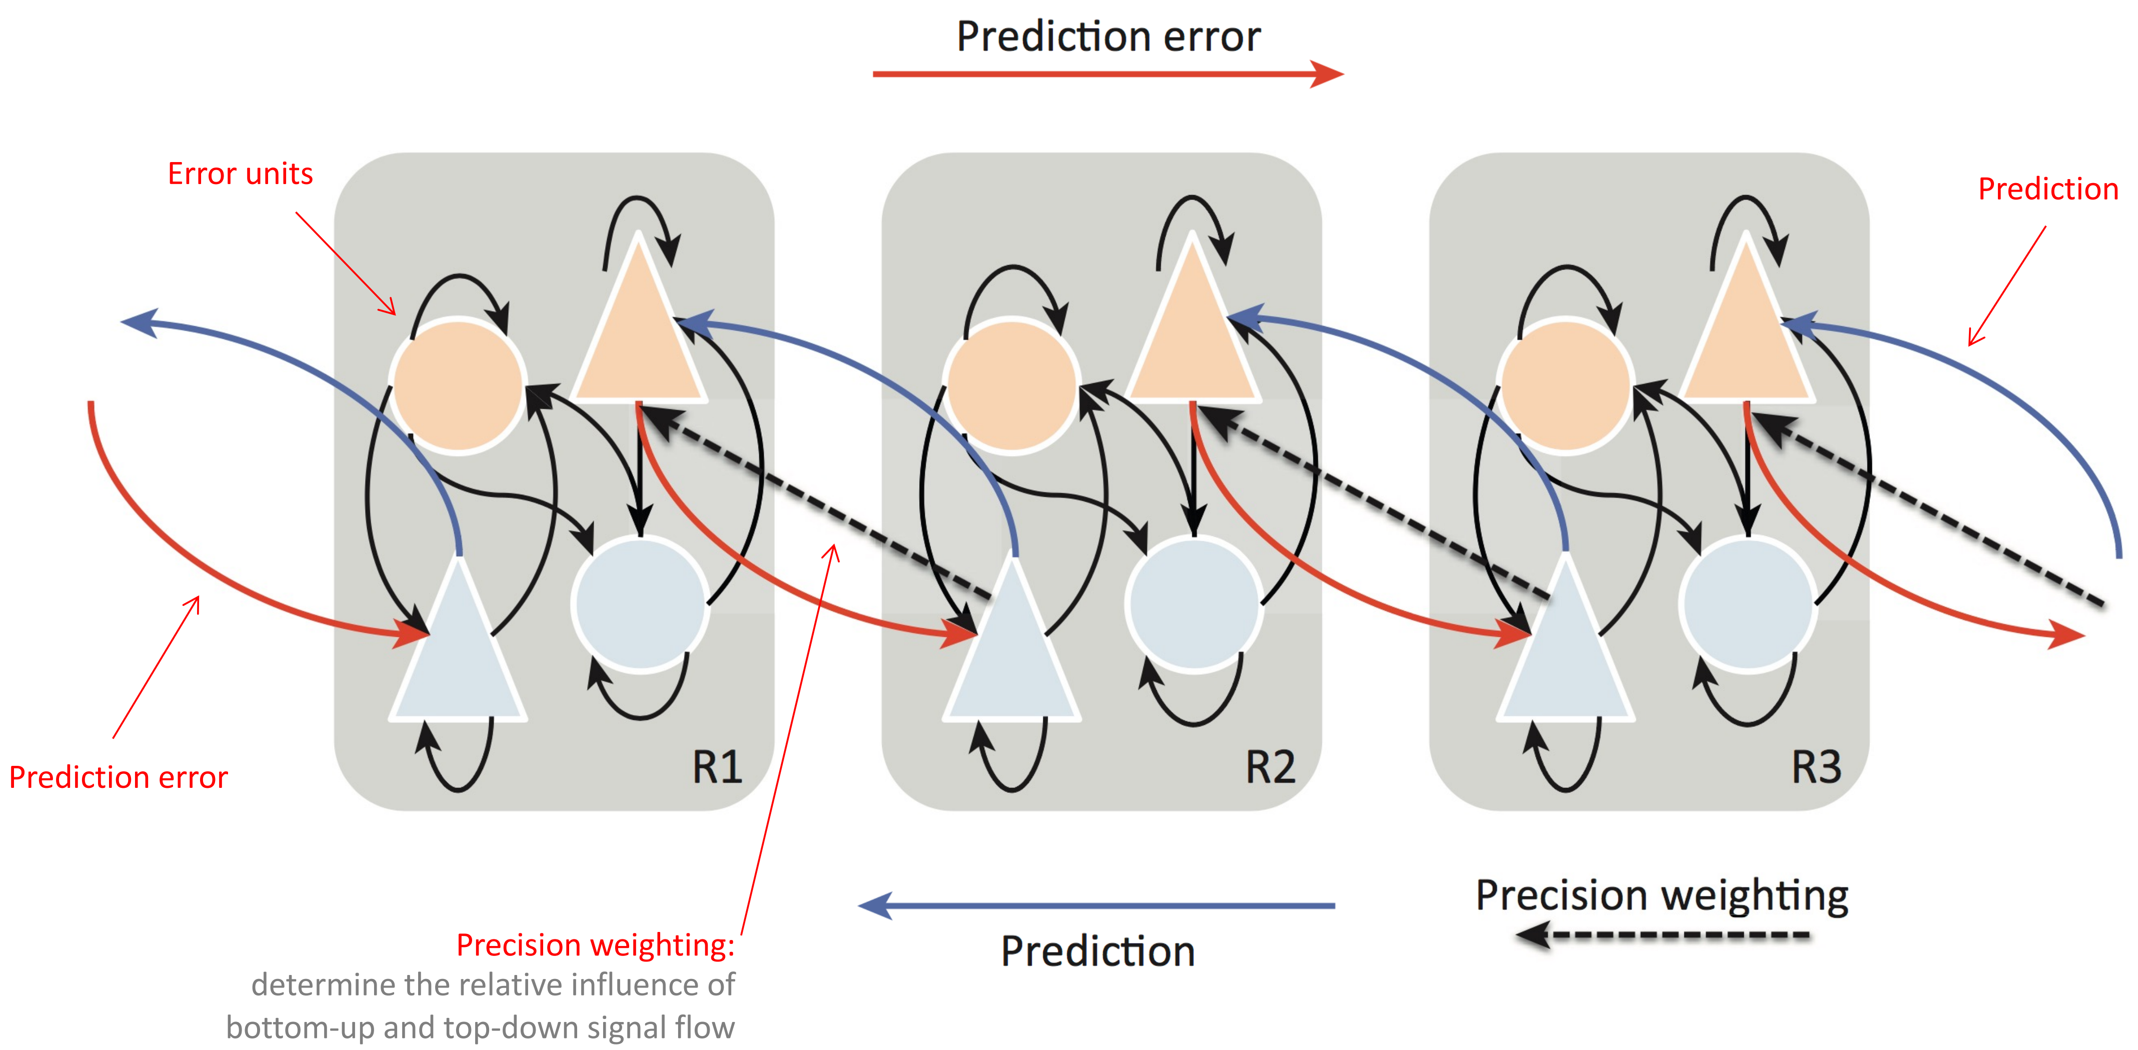

A. K. Seth. Interoceptive inference, emotion, and the embodied self. Trends in Cognitive Sciences, 17(11):565–573, November 2013.

$$\xi^{(i)v} = \tilde{\mu}^{(i-1)v} - \tilde{g}(\tilde{\mu}^{(i)}) - \Lambda^{(i)z}\,\xi^{(i)v}$$

$$\xi^{(i)x} = D\tilde{\mu}^{(i)x} - \tilde{f}(\tilde{\mu}^{(i)}) - \Lambda^{(i)w}\,\xi^{(i)x}$$

forward prediction error

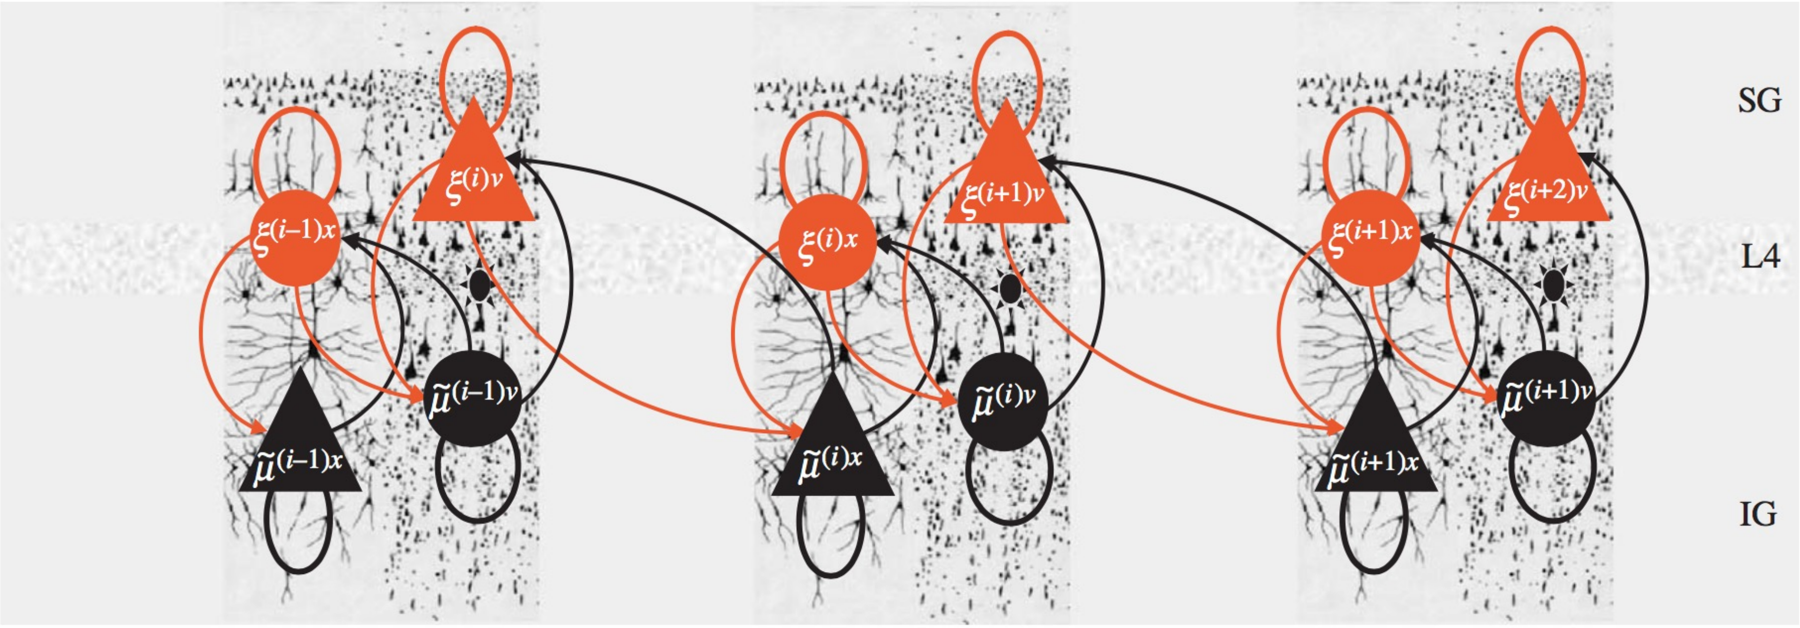

backward predictions

$$\dot{\tilde{\mu}}^{(i)v} = D\tilde{\mu}^{(i)v} - \tilde{\varepsilon}_v^{(i)T}\,\xi^{(i)} - \xi^{(i+1)v}$$

$$\dot{\tilde{\mu}}^{(i)x} = D\tilde{\mu}^{(i)x} - \tilde{\varepsilon}_x^{(i)T}\,\xi^{(i)}$$

K. Friston and S. Kiebel. Predictive coding under the free-energy principle. Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences, 364(1521)):1211–1221, 2009

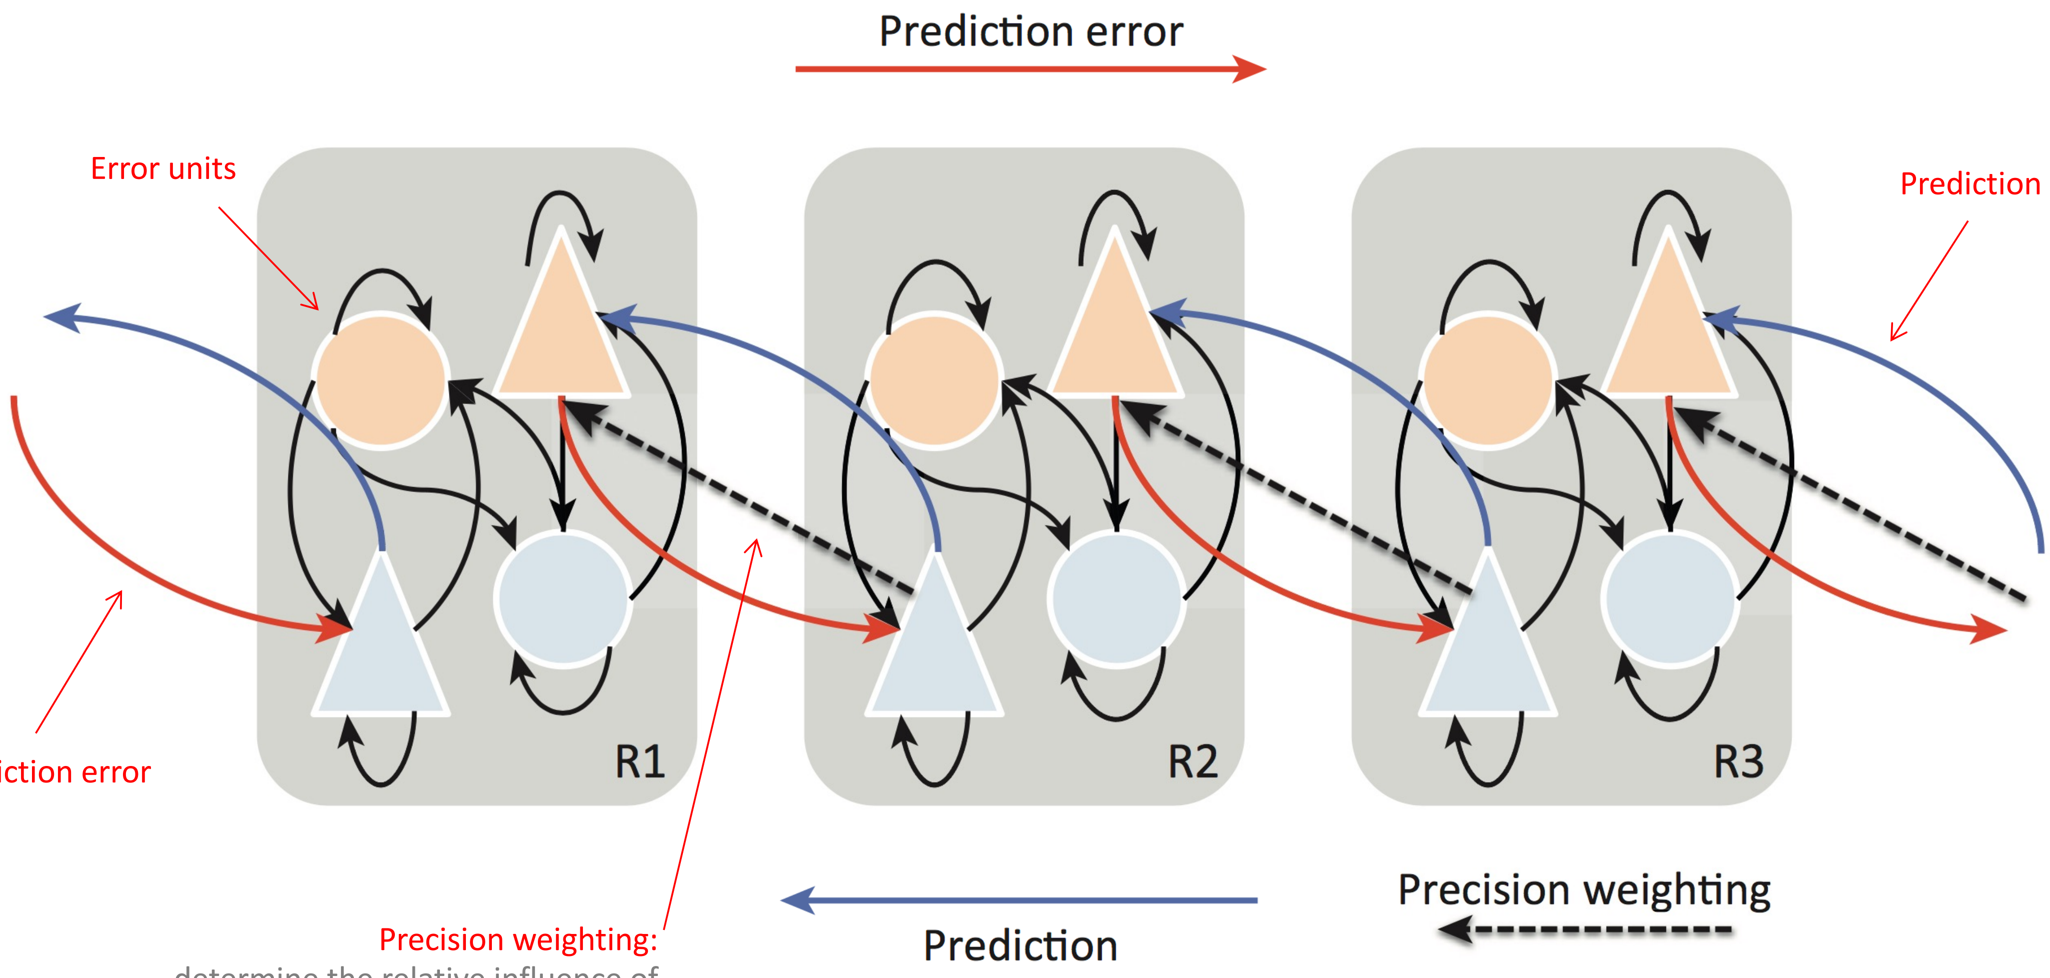

Precision weighting:
determine the relative influence of
bottom-up and top-down signal flow

A. K. Seth. Interoceptive inference, emotion, and the embodied self. Trends in Cognitive Sciences, 17(11):565–573, November 2013.

Increasing
Level of
Abstraction
(modal & amodal)

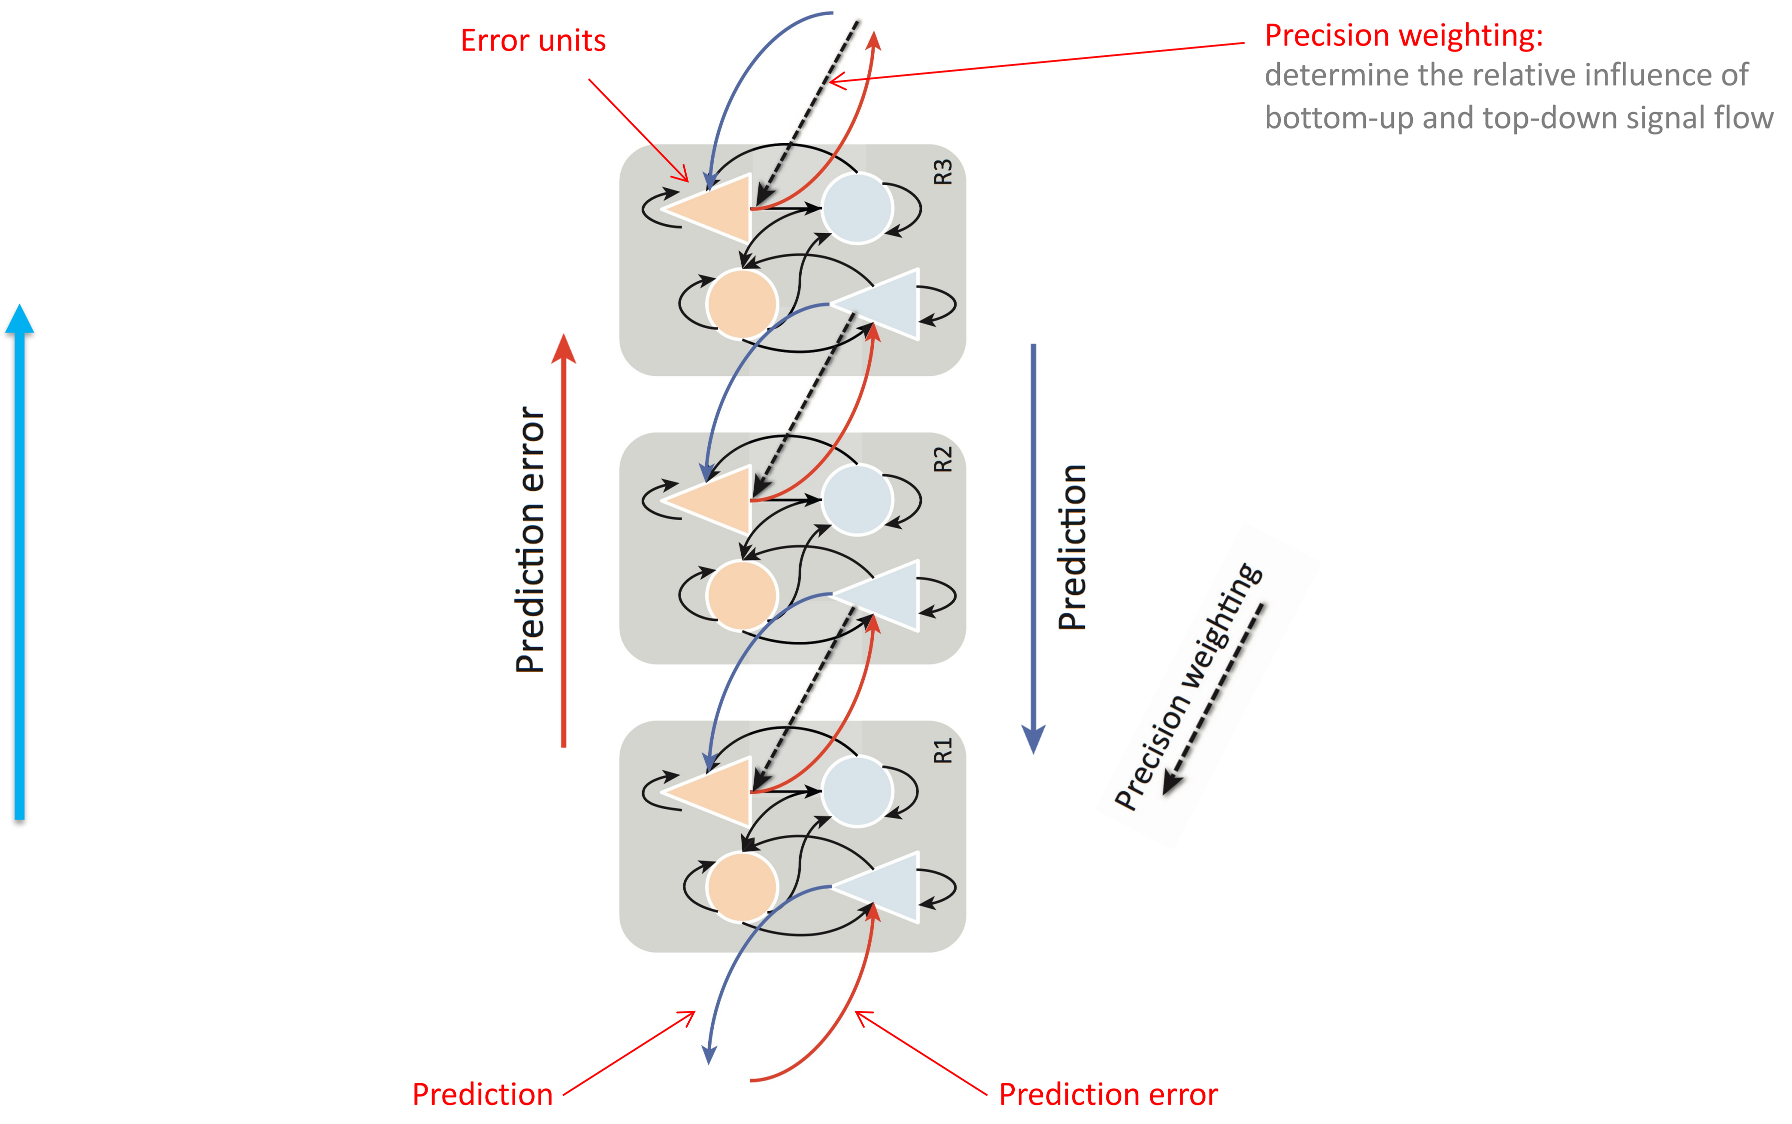

A. K. Seth. Interoceptive inference, emotion, and the embodied self. Trends in Cognitive Sciences, 17(11):565–573, November 2013.

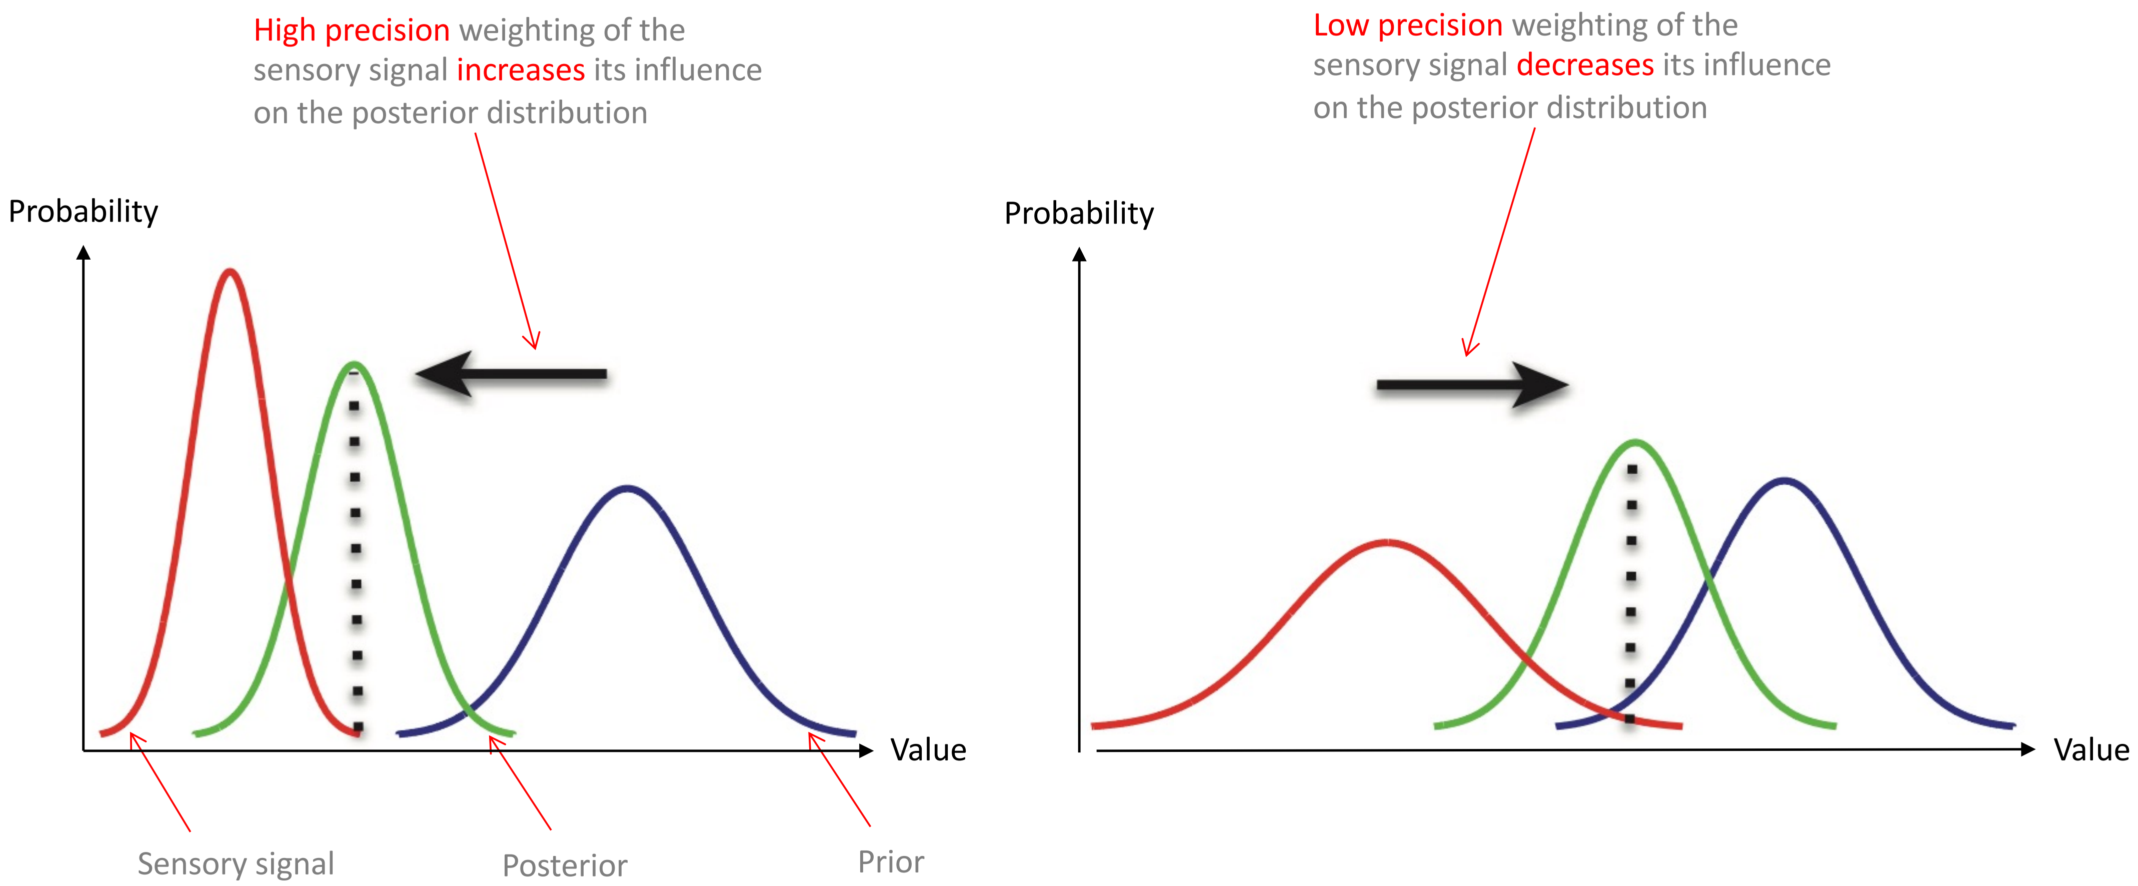

A. K. Seth. Interoceptive inference, emotion, and the embodied self. Trends in Cognitive Sciences, 17(11):565–573, November 2013.

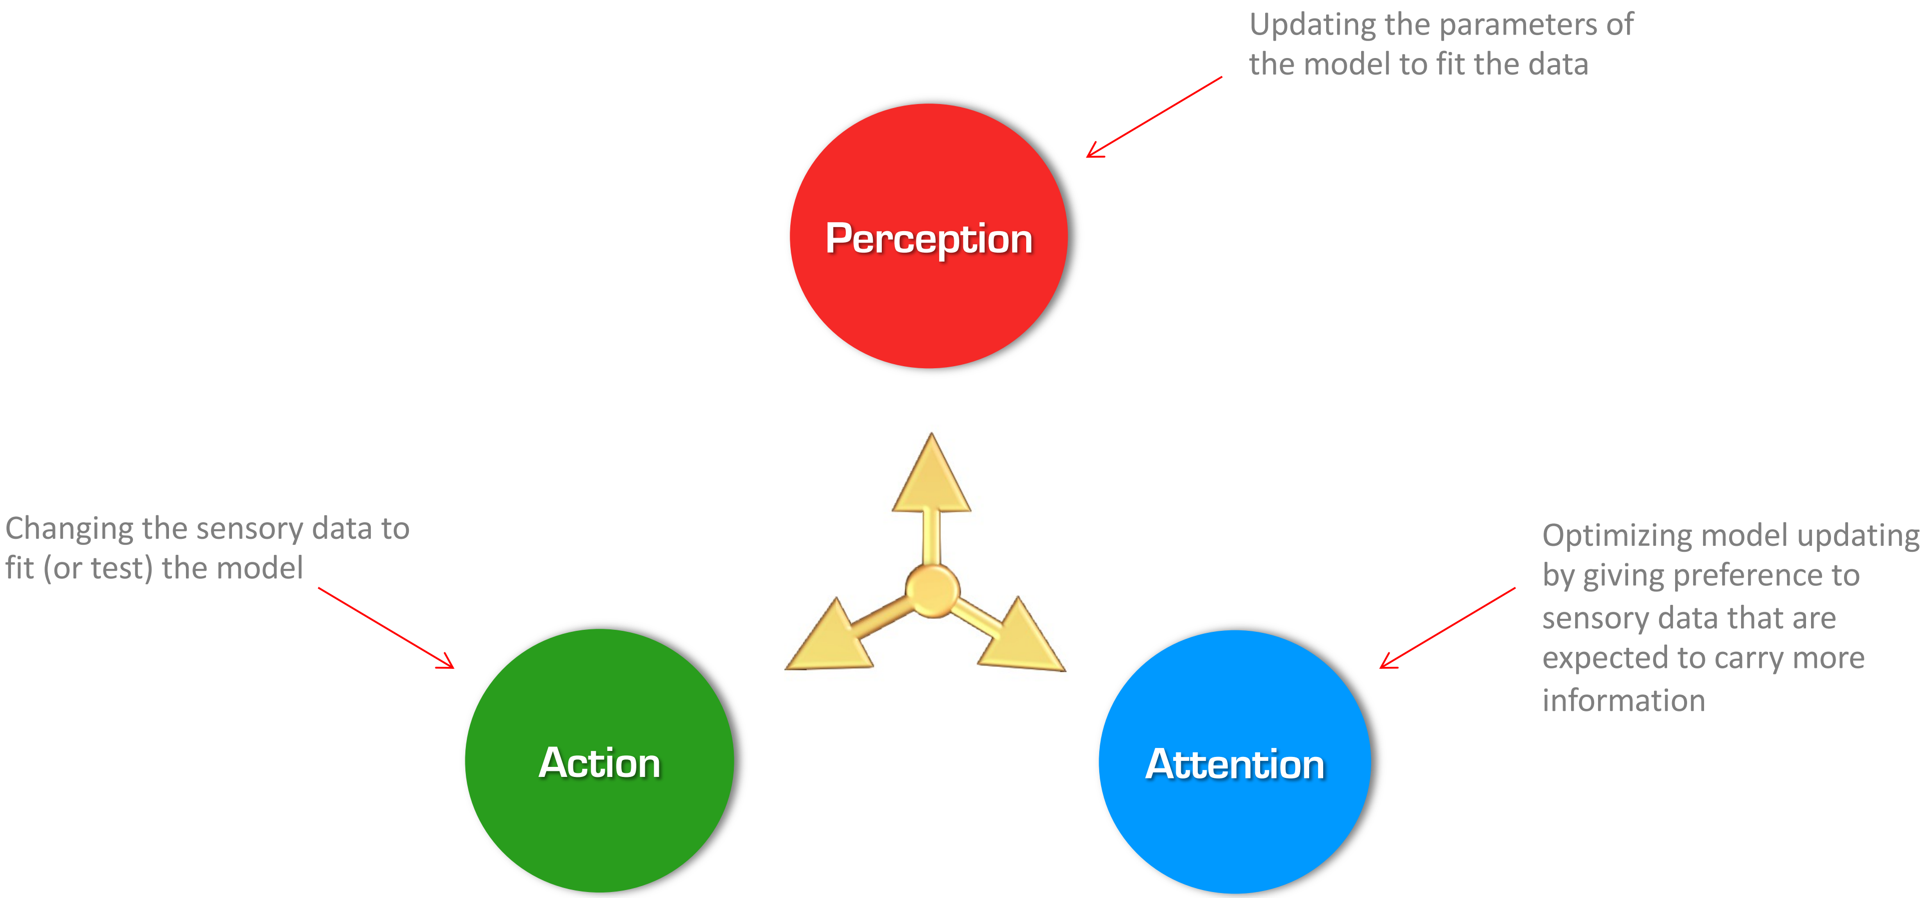
Updating the parameters of the model to fit the data
Perception
Changing the sensory data to fit (or test) the model
Action
Optimizing model updating by giving preference to sensory data that are expected to carry more information
Attention

A single principle
to account for

perception
action
cognition
consciousness

based on
neural operations

Internal models

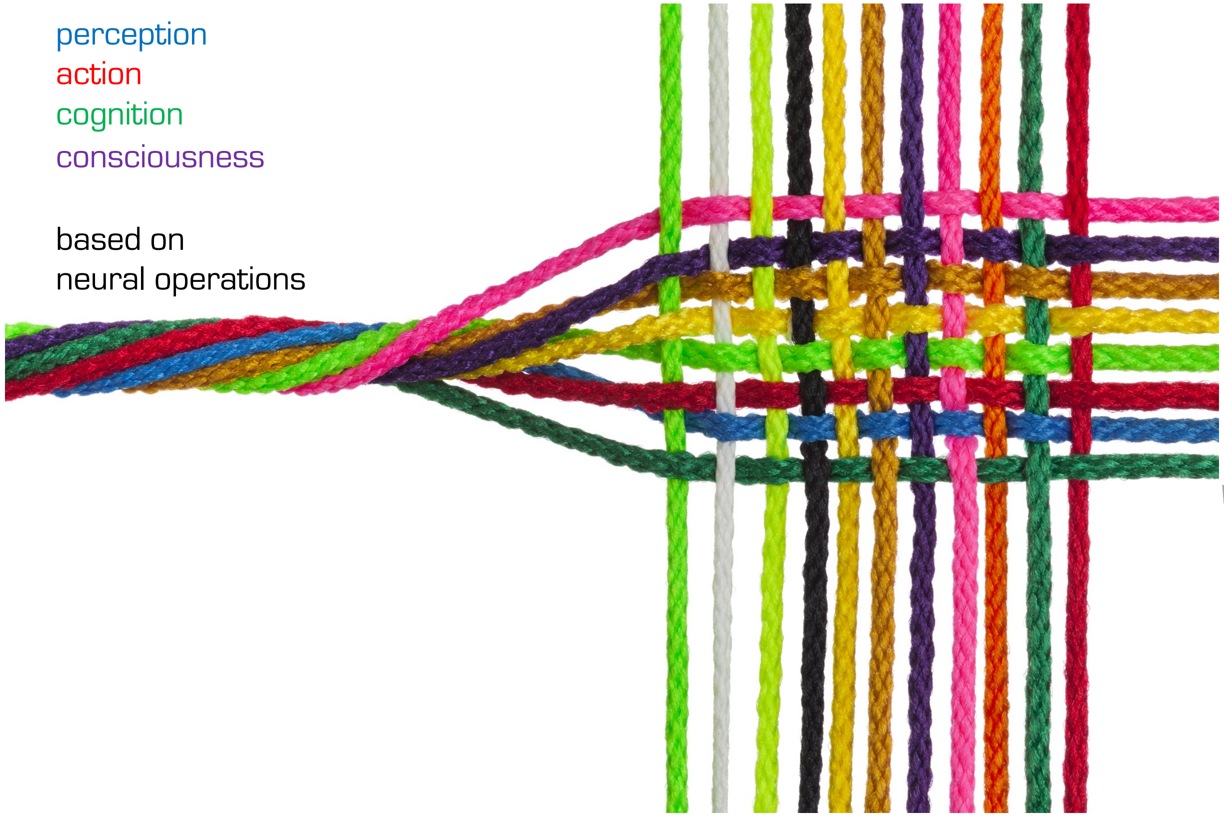

https://fashion-history.lovetoknow.com/fabrics-fibers/weave-types

Free Energy Principle

"… according to which perceptual inference and action emerge as a consequence of a more fundamental imperative towards the avoidance of "surprising" events"

"Organisms must minimize the long-run average surprise of sensory states, since surprising sensory states are likely to reflect conditions incompatible with continued existence"

Free Energy Principle

"… according to which perceptual inference and action emerge as a consequence of a more fundamental imperative towards the avoidance of "surprising" events"

"Surprising" in an information-theoretic sense:
surprisal, a measure of uncertainty of an event
(or "unlikeliness of the occurrence of an event")

$$u_i = -\log_2(P_i) \quad [= \log_2(1/P_i)]$$

The average surprisal is entropy $H = -\sum_{i=1}^{M} P_i \log_2 P_i$

Cannot evaluate surprisal of world state or average long-run surprisal, i.e. entropy

World states are hidden & entropy requires averaging over all possible occurrences of events

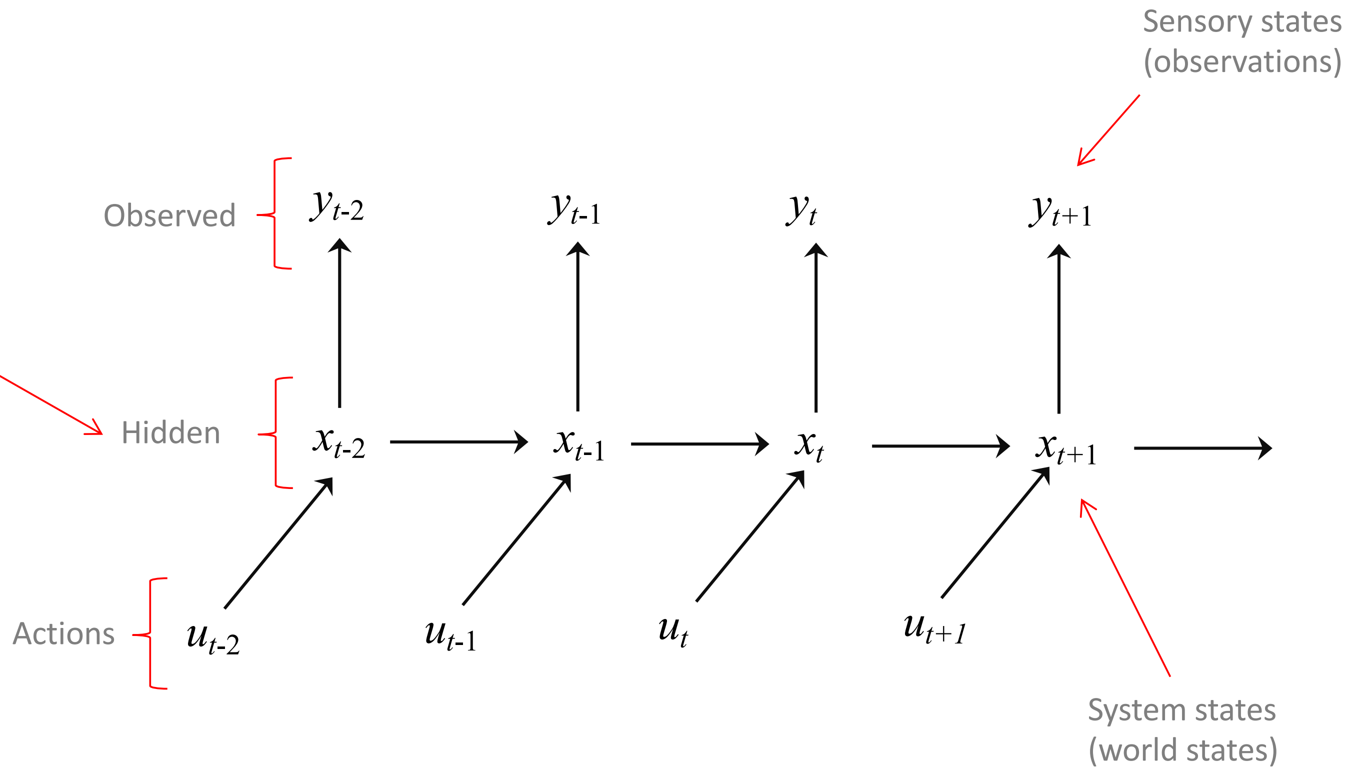

Free Energy Principle

Instead, the agent maintains a lower limit on surprisal by minimizing

the difference between actual sensory signals and the signals produced by a predictive model

This difference is **free energy**

Under some general assumptions, this is the long-run sum of prediction error

This links the free energy principle with predictive processing

Entropy is the average surprisal $H = -\sum_{i=1}^{M} P_i \log_2 P_i$

Free Energy Principle

"Organisms minimize an upper bound on the entropy of sensory signals (the free energy).

Under specific assumptions, free energy translates to prediction error."

[Seth 2013]

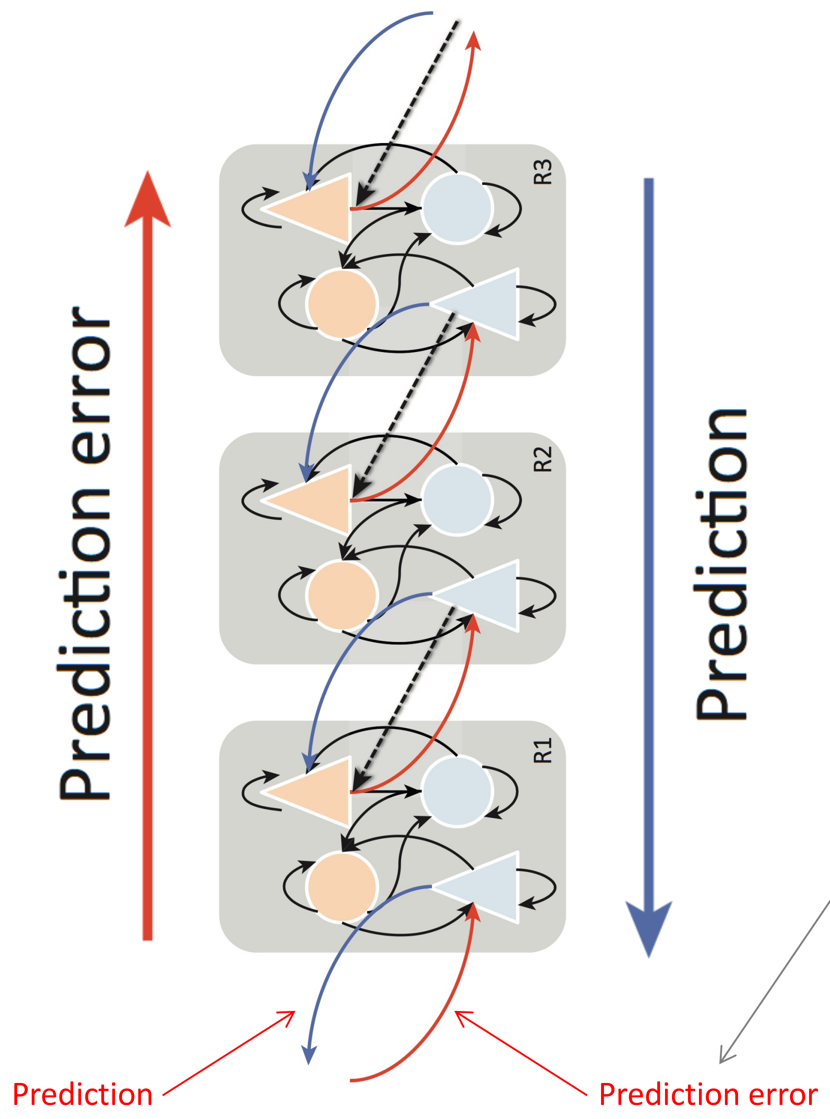

Free energy principle
Minimize the long-run prediction error

A. K. Seth. Interoceptive inference, emotion, and the embodied self. Trends in Cognitive Sciences, 17(11):565–573, November 2013.

Key Idea: minimize prediction errors

Perform actions to confirm or test sensory predictions

"actively"

Change the sampling (act) "to find evidence against the current hypothesis, and/or efficiently disambiguate between competing hypotheses"

"Finessed sense of" Active inference

Counterfactual active inference

A single principle
to account for

perception
action
cognition
consciousness

based on
neural operations

Internal models

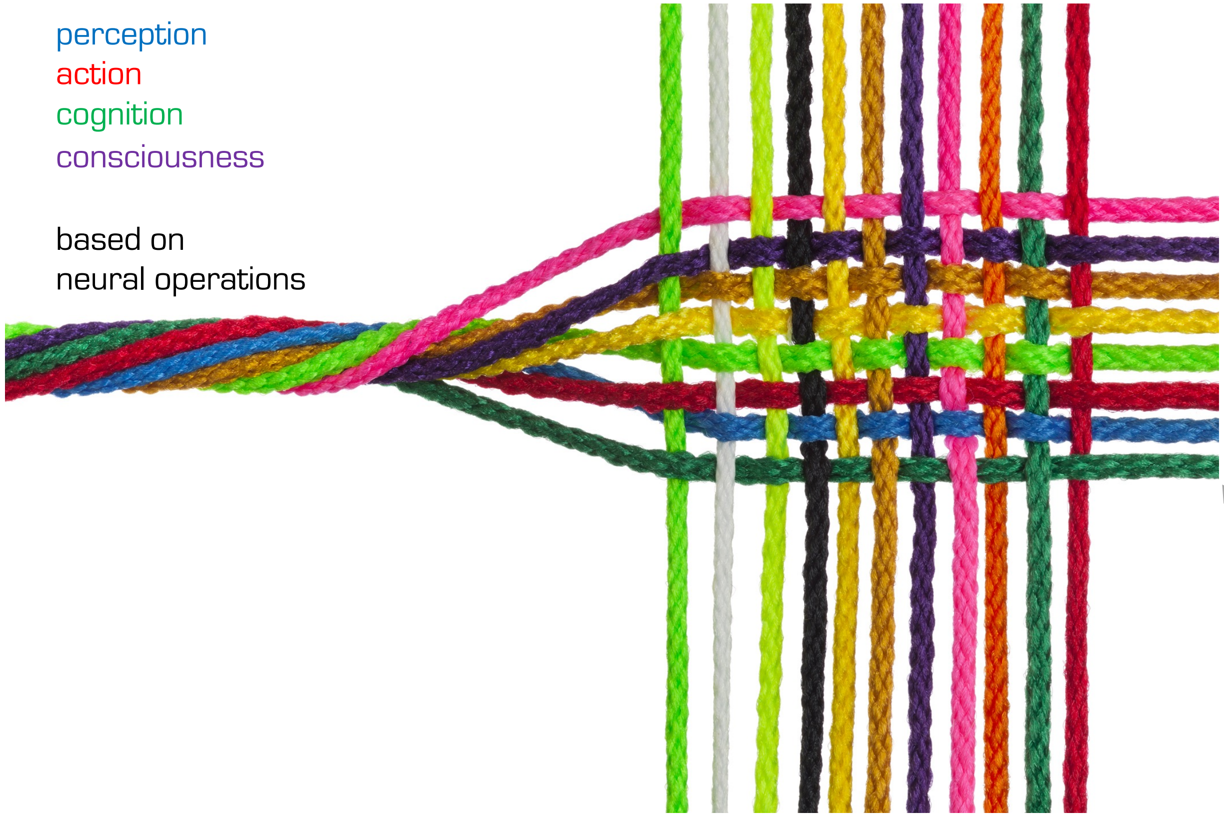

https://fashion-history.lovetoknow.com/fabrics-fibers/weave-types

Bayesian brain: predictive processing (PP) and active inference

Free energy principle: minimization of average surprisal

Cybernetics: self-regulation, adaptive self-regulation, ultrastability

Interoception and emotion

Embodiment

Enactive cognitive science: weak and strong

Sensorimotor contingency (SMC) theory & PPSMC

Counterfactual PP and active inference

Cybernetics

- PP– and specifically perceptual inference – can be traced to Hermann von Helmholtz
  - Passive
  - Not very much concerned with behaviour
- Free energy principle focusses on close coupling of perception and action
- Suggesting a deep connection between PP and cybernetics

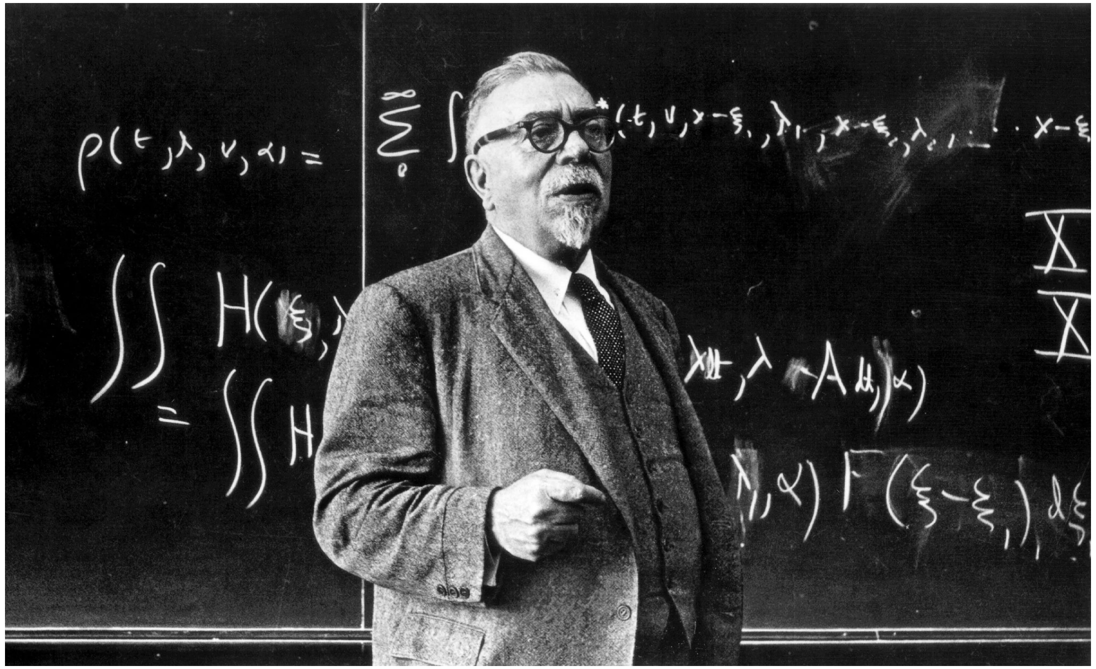

Norbert Wiener

https://www.nytimes.com/2013/05/21/science/mit-scholars-1949-essay-on-machine-age-is-found.html

W. Ross Ashby

http://www.rossashby.info/index.html

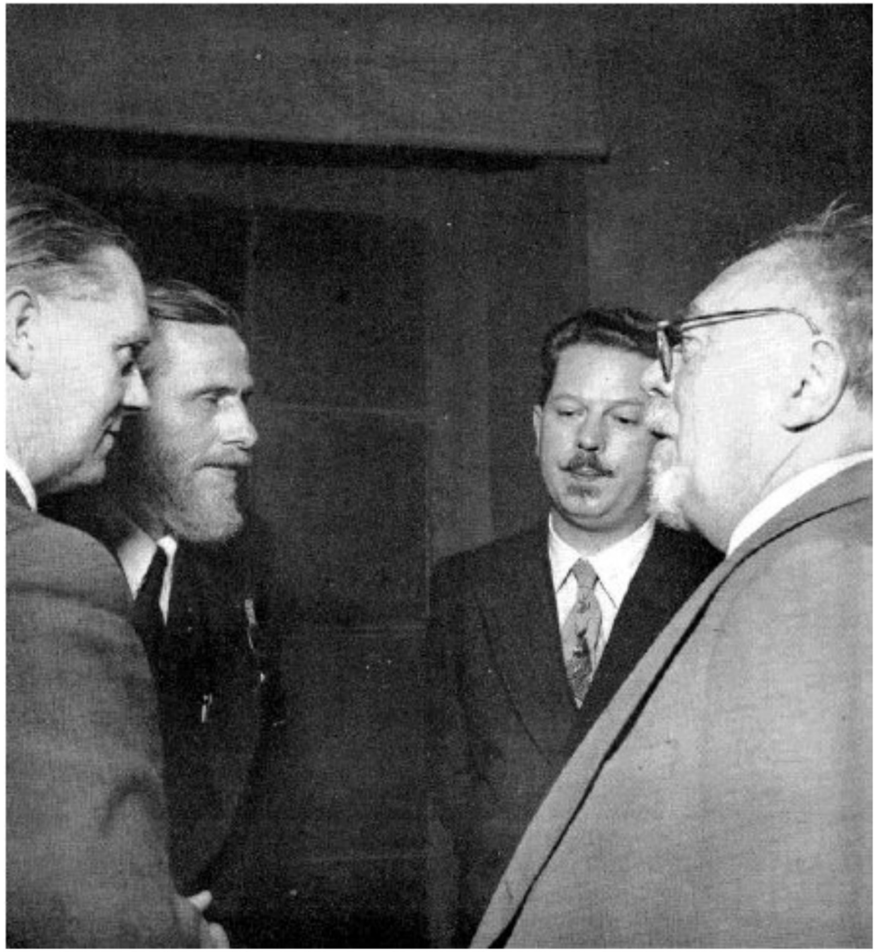

W. Ross Ashby, Warren McCulloch, Grey Walter, Norbert Wiener
at the 1951 Congress on Cybernetics, Paris

https://www.researchgate.net/publication/287293010_Warren_McCulloch_and_the_British_Cyberneticians/figures?lo=1

N. Wiener
Cybernetics: or the Control and Communication in the Animal and the Machine, 1948.

(*κυβερνητης* or kybernetes: steersman)

W. Ross Ashby
Design for a Brain, first edition, 1952 ... 1956, 1960.
Introduction to Cybernetics, 1957

Cybernetics

Louis Kauffman
President of the American Society for Cybernetics

"Cybernetics is the study of systems and processes that interact with themselves and produce themselves from themselves."

Cybernetics

Central focus

- Prediction and control of behaviour
- In telelogical, purposeful, goal-directed machines
- Circular causal chains involving feedback coupling goal-directed sensation and action

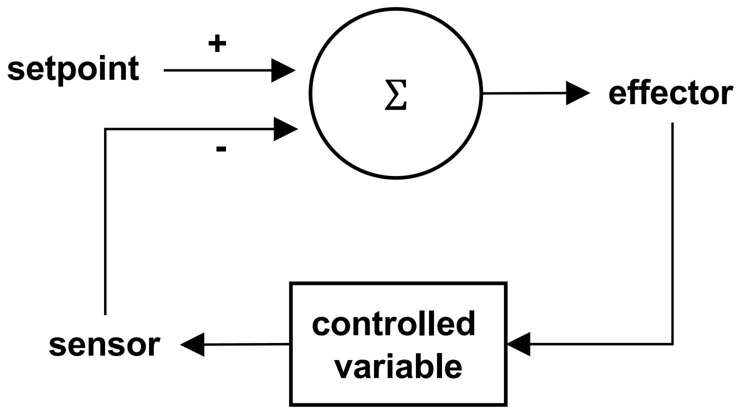

Homeostasis: the automatic regulation of physiological functions … essential variables (EV)

Walter Cannon in 1929: “Organization for Physiological Homeostasis”

Goal: **stability through constancy**: set point & negative feedback

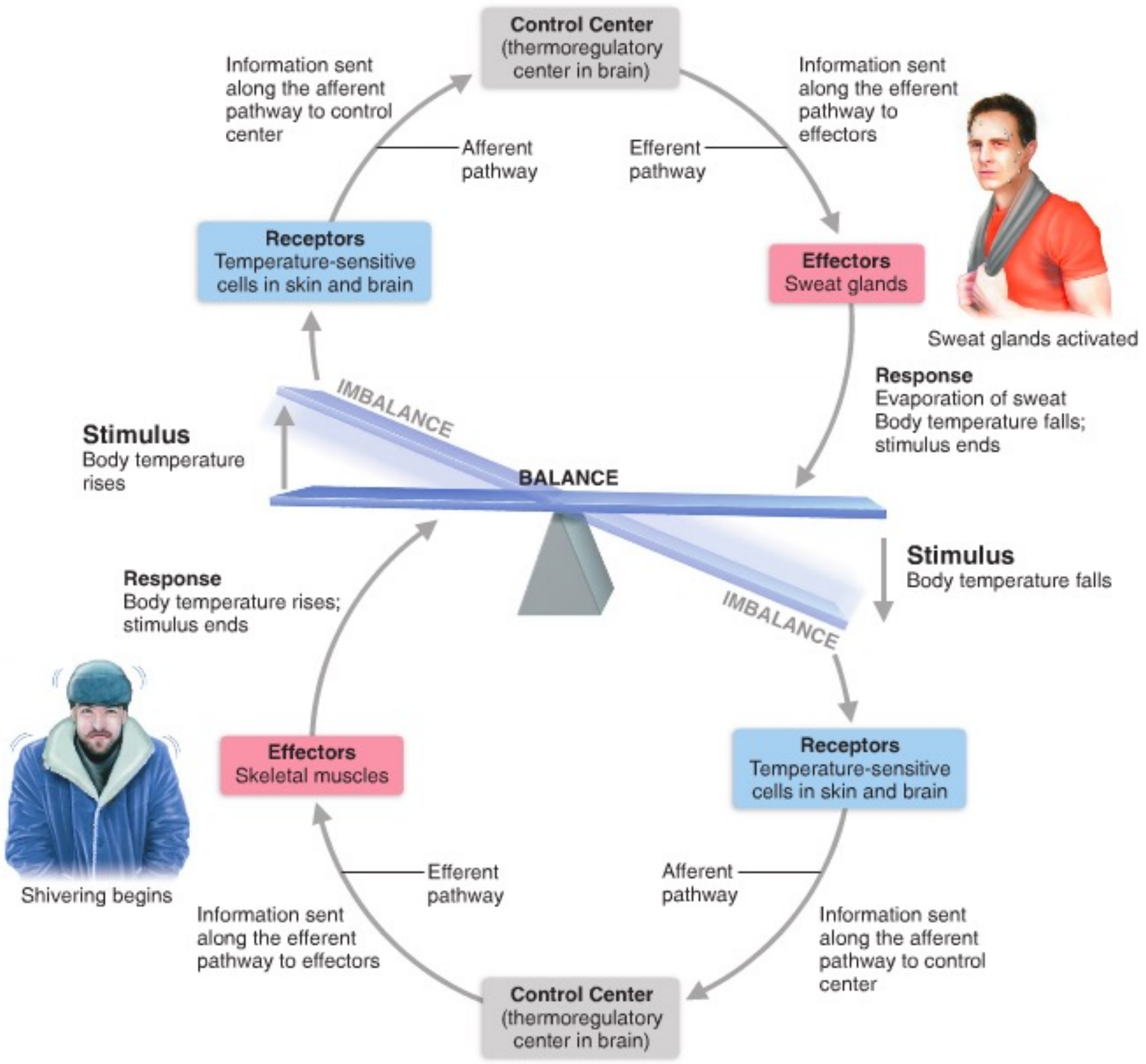

https://www.studyblue.com/notes/note/n/homeostasis/deck/6151960

Feedback Control

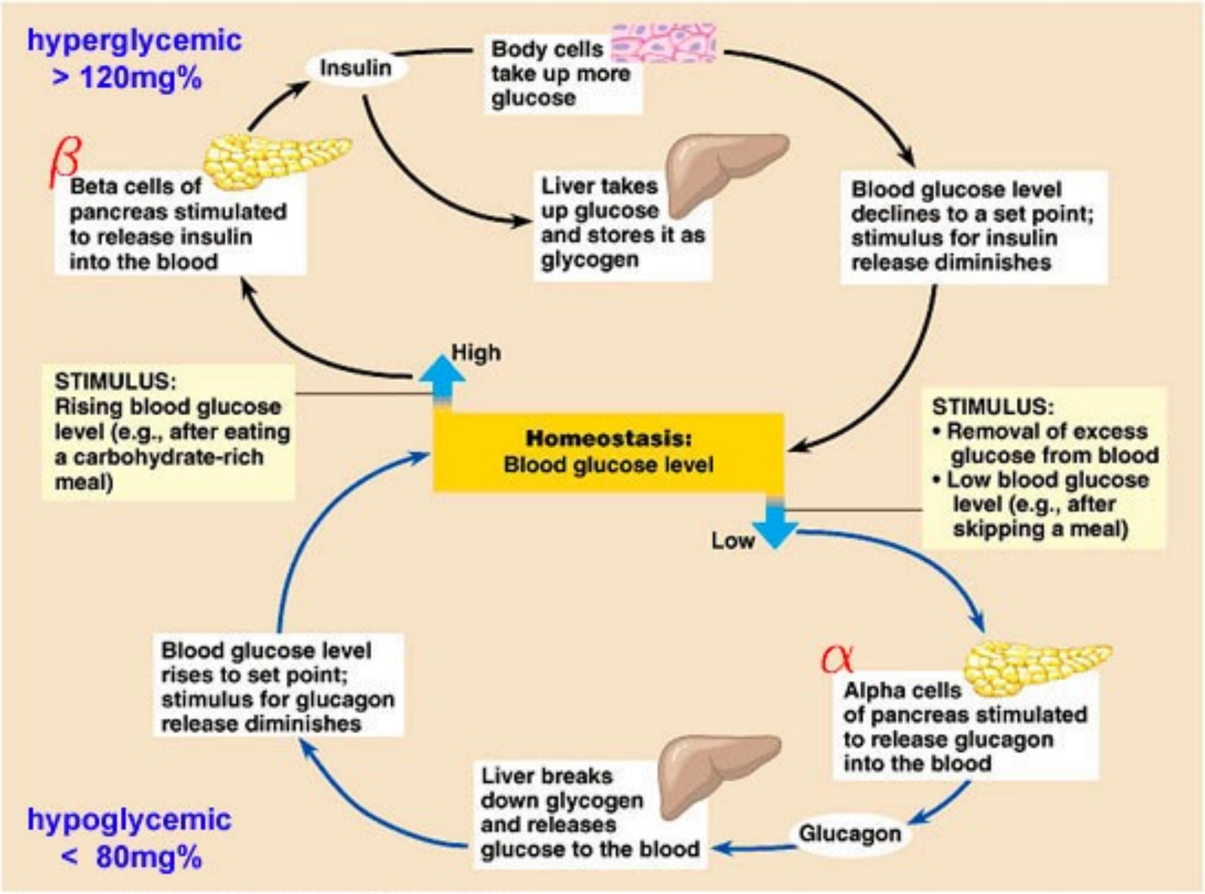

http://biochemgroup.blogspot.com/2010/07/diabetes-mellitus.html

Insight 1: Adaptive homeostasis

First-order feedback

- Homeostasis keeps essential variables within viable bounds
- When essential variables move beyond viability limits ...
- Adaptive processes re-parameterize the system ...

Second-order feedback: allostasis
(random changes in Ashby's Homeostat)

- Until it reaches a new equilibrium (& homeostasis is restored)
- These are called ultrastable systems
- Allostasis: "the process of achieving homeostasis"

W. Ross Ashby

http://www.rossashby.info/index.html

Allostasis

"The brain monitors a very large number of external and internal parameters to anticipate changing needs, evaluate priorities, and prepare the organism to satisfy them before they lead to errors. The brain even anticipates its own local needs, increasing flow to certain regions — before there is an error signal"

Peter Sterling

P. Sterling. Principles of allostasis: optimal design, predictive regulation, pathophysiology and rational therapeutics. In J. Schulkin, editor, Allostasis, Homeostasis, and the Costs of Adaptation, pages 17–64. Cambridge University Press., Cambridge, England, 2004.

P. Sterling. Allostasis: A model of predictive regulation. Physiology and Behaviour, 106 (1): 5–15, 2012.

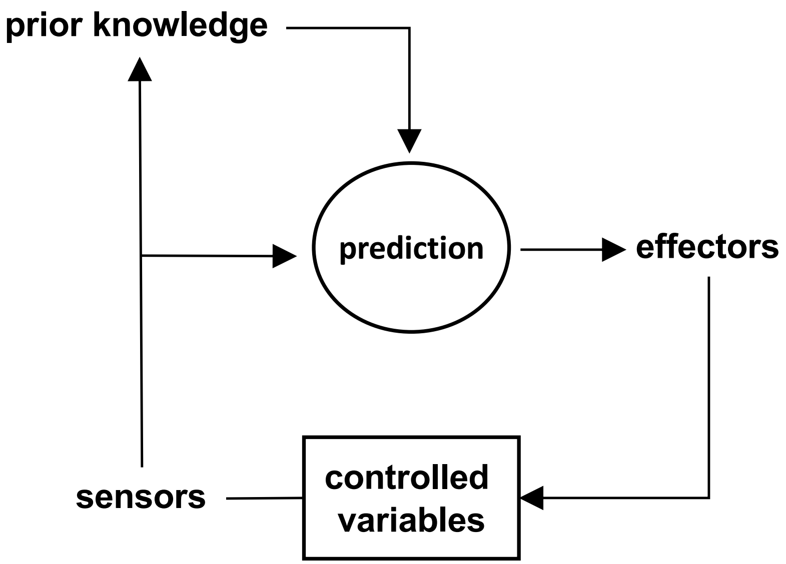

Homeostasis: adjusting to an event

Allostasis: adjusting before an event occurs

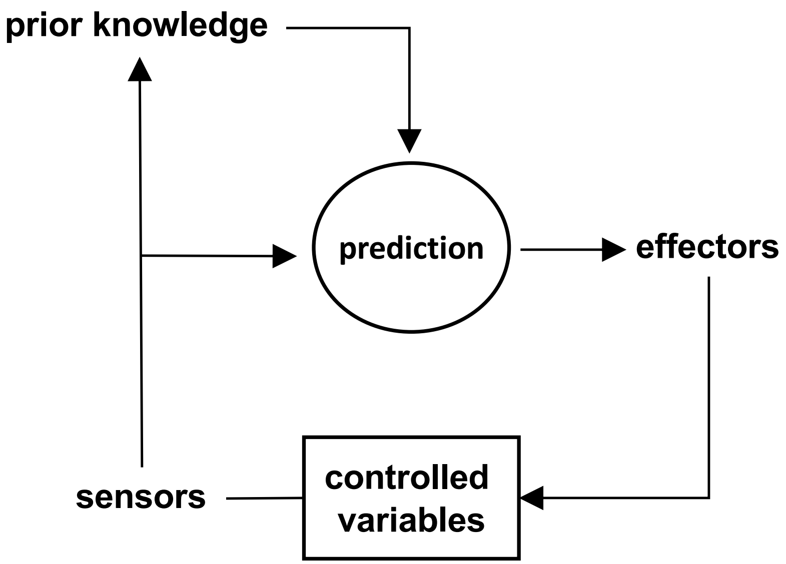

Predictive self-regulation vs. reactive self-regulation

Ready themselves for multiple contingencies

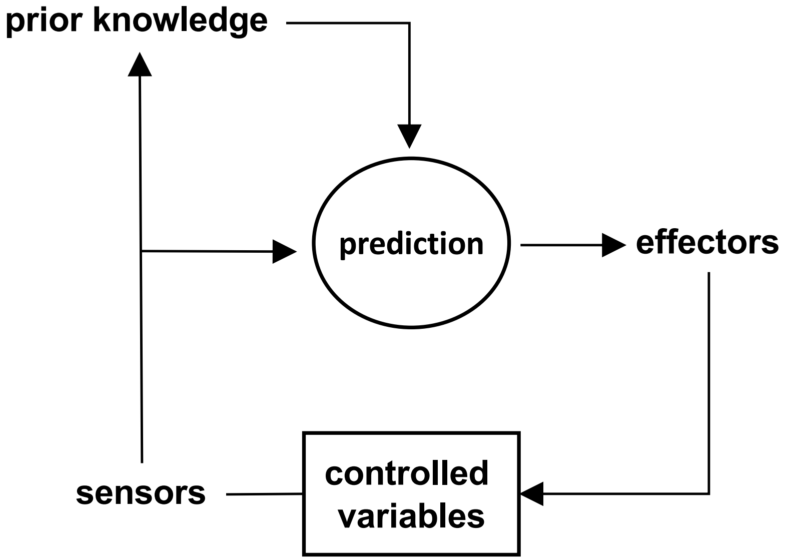

Use priors to anticipate the likely demands that will be placed on the system

Pre-emptively adjust all the parameters to meet this demand

Change the controlled variable by predicting what value will be needed and **overriding** local feedback to meet anticipated demand

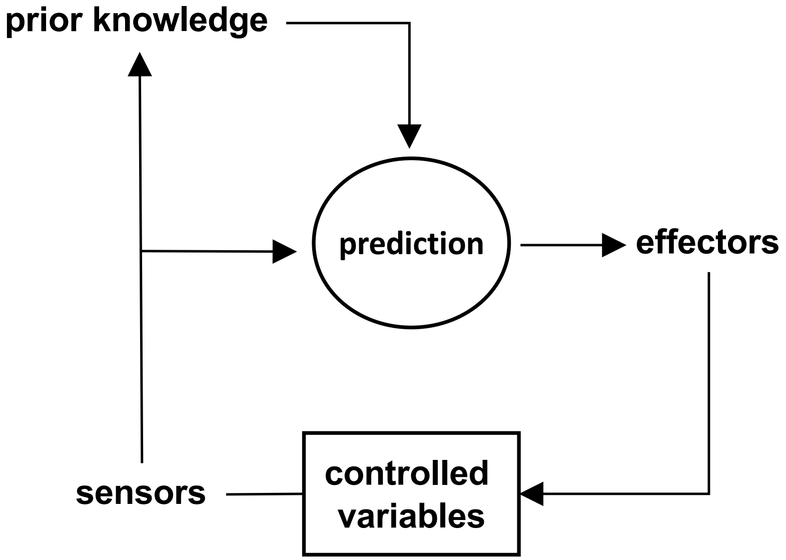

Allostatic systems adapt to change rather than resist it

Allostasis is effected at a higher level of organization, involving greater number of sub-systems acting together in a coordinated manner

(Homeostasis operates at a simpler level of negative feedback control)

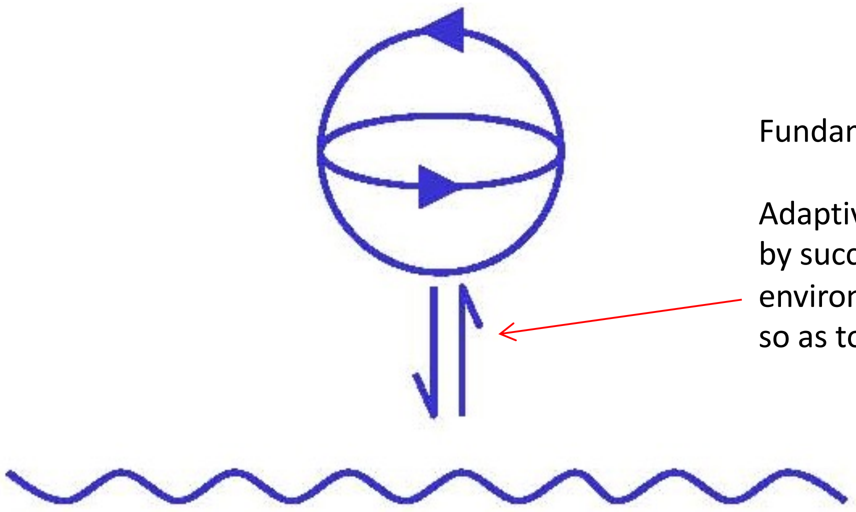

Fundamental principle of cybernetics

Adaptive systems ensure their continued existence
by successfully responding to
environmental perturbations
so as to maintain their internal organization

Insight 2: Law of Requisite Variety

- A successful control system must be capable of entering at least as many states as the system being controlled
- The complexity of the environment-system perturbations determines the minimum level of complexity of the controller
- "Every good regulator of a system must be a model of that system" Conant & Ashby (1970)

Good regulator theorem

Links to the free energy principle:
adaptive systems minimize a limit on free energy (long-run average surprisal)
by inducing and refining a generative model of the causes of sensory signal ...
more powerful than random variation of higher-order parameters

W. Ross Ashby

http://www.rossashby.info/index.html

"The purpose of cognition (including perception and action) is to maintain the homeostasis of essential variables and of internal organization (ultrastability)"

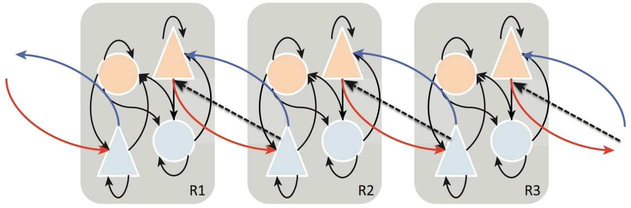

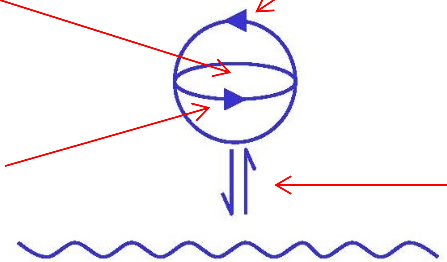

"The purpose of cognition (including perception and action) is to maintain the homeostasis of essential variables and of internal organization (ultrastability)"

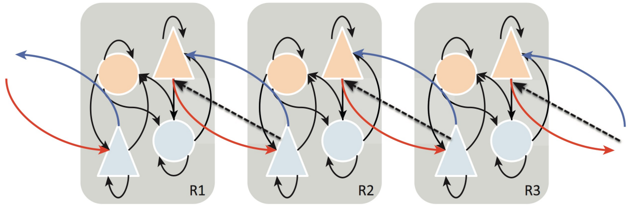

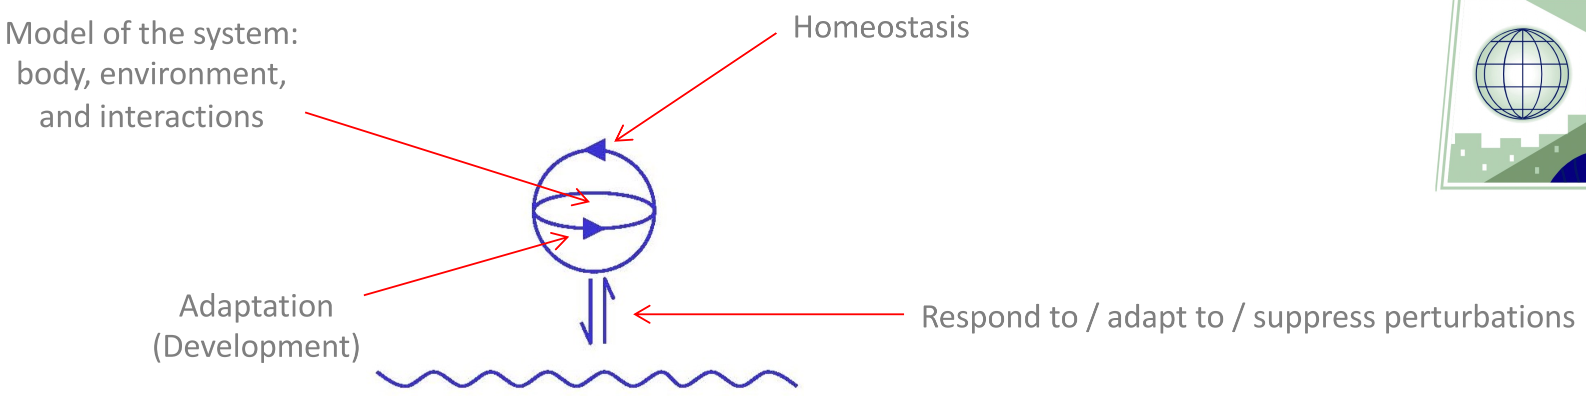

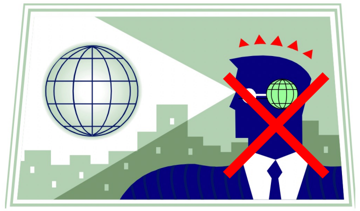

"Perception emerges as a consequence of a more fundamental imperative towards organizational homeostasis, and not as a stage in some process of internal world-model construction"

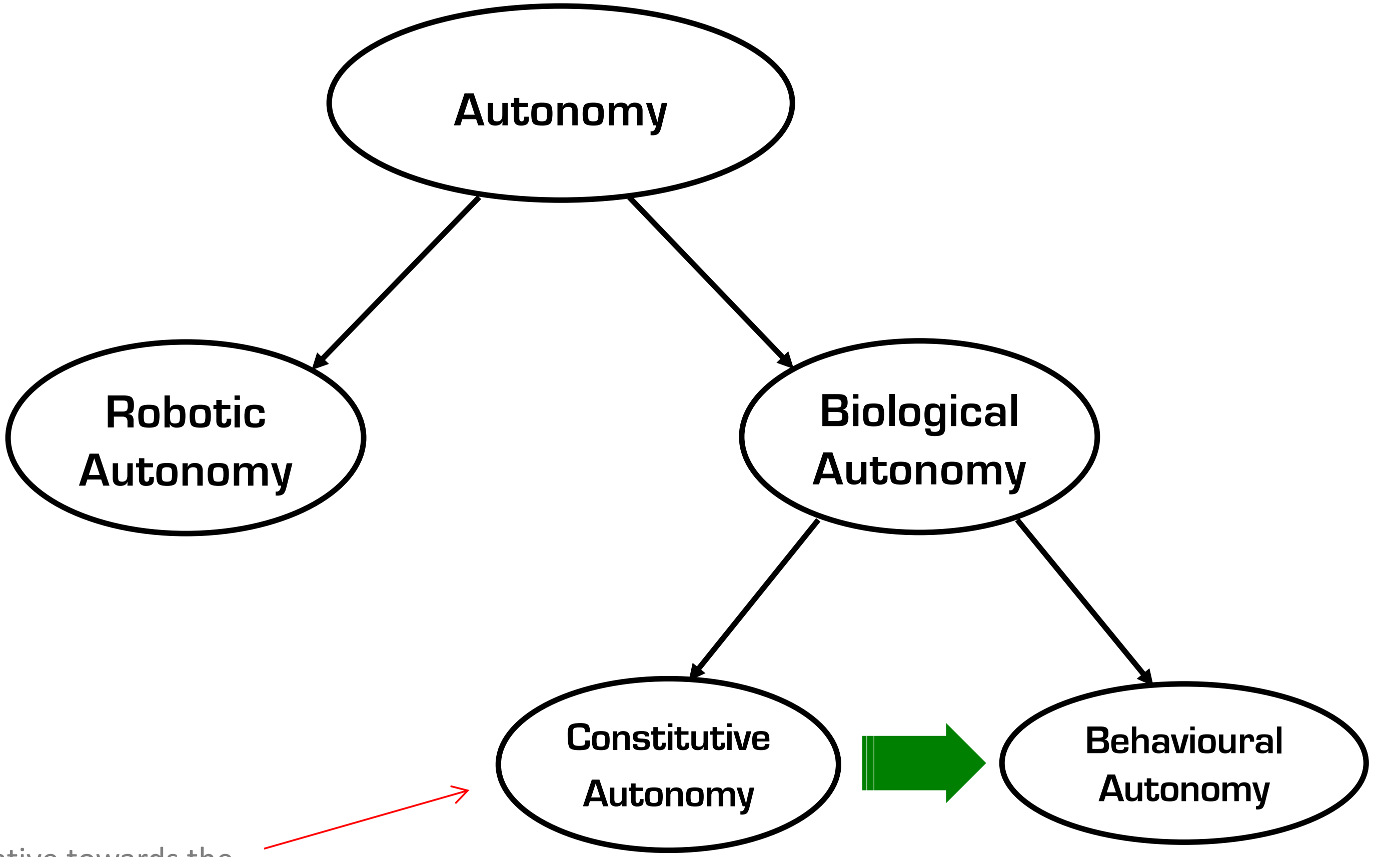

Primary imperative towards the continued survival of the organism

Stable environment and simple mapping between perturbations and homeostatic response

⇒

good regulator only needs to implicitly model the environment

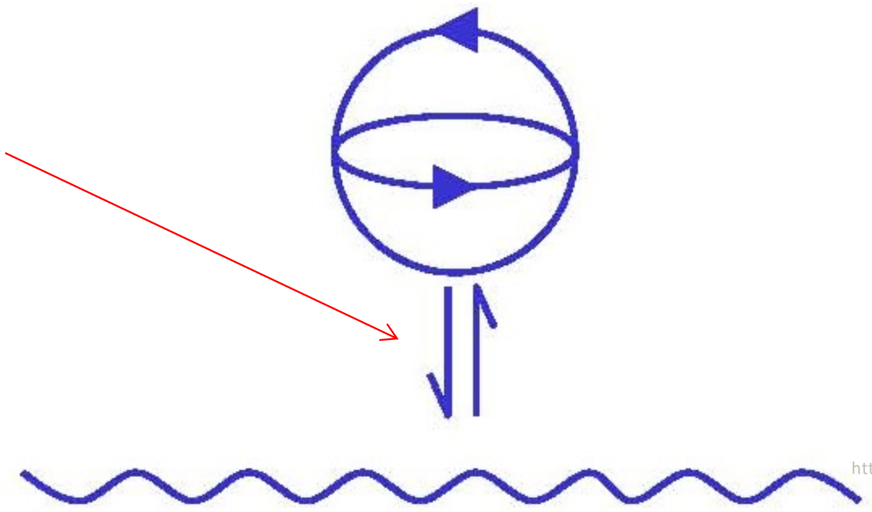

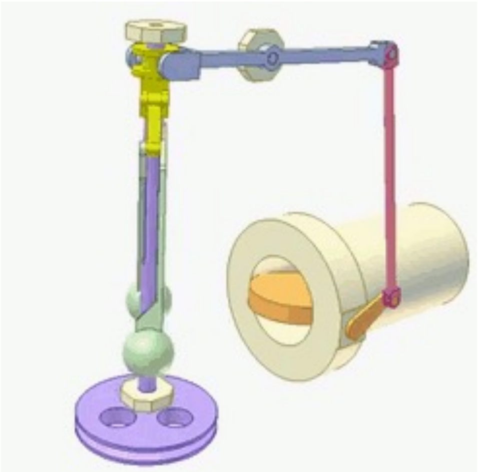

https://www.mech.kuleuven.be/en/tme/thermotechnisch-instituut/basisprincipes/Watt-regulator

Many-to-many mappings between
sensory states and probable causes

⇒

good regulator needs explicit inferential processes to model the environment and extract the most probable causes

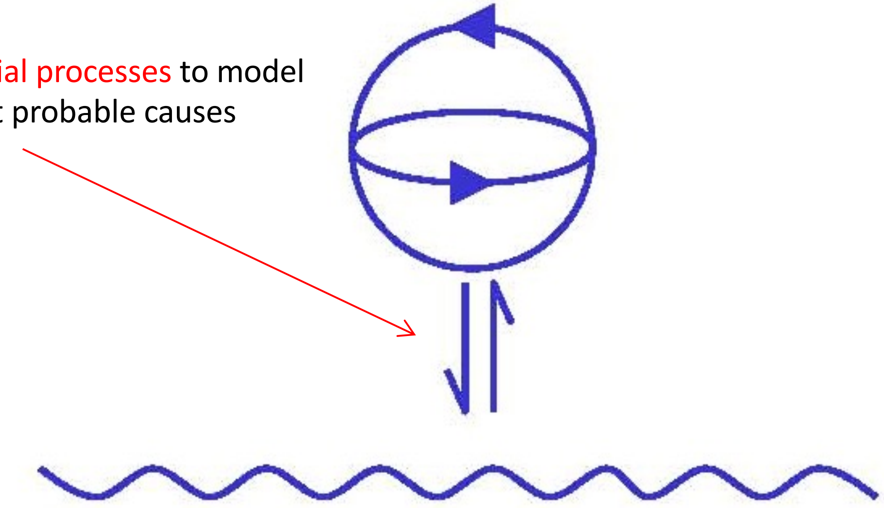

Shift in Perspective

Exteroception

Predictive processing:

from Helmholtzian "perception as inference"

to

model-based predictive control entailed by a fundamental imperative towards internal homeostasis

(and perception as inference)

Interception

https://unsplash.com/photos/-TQUERQGUZ8?utm_source=unsplash&utm_medium=referral&utm_content=creditCopyText

A single principle
to account for

perception
action
cognition
consciousness

based on
neural operations

Internal models

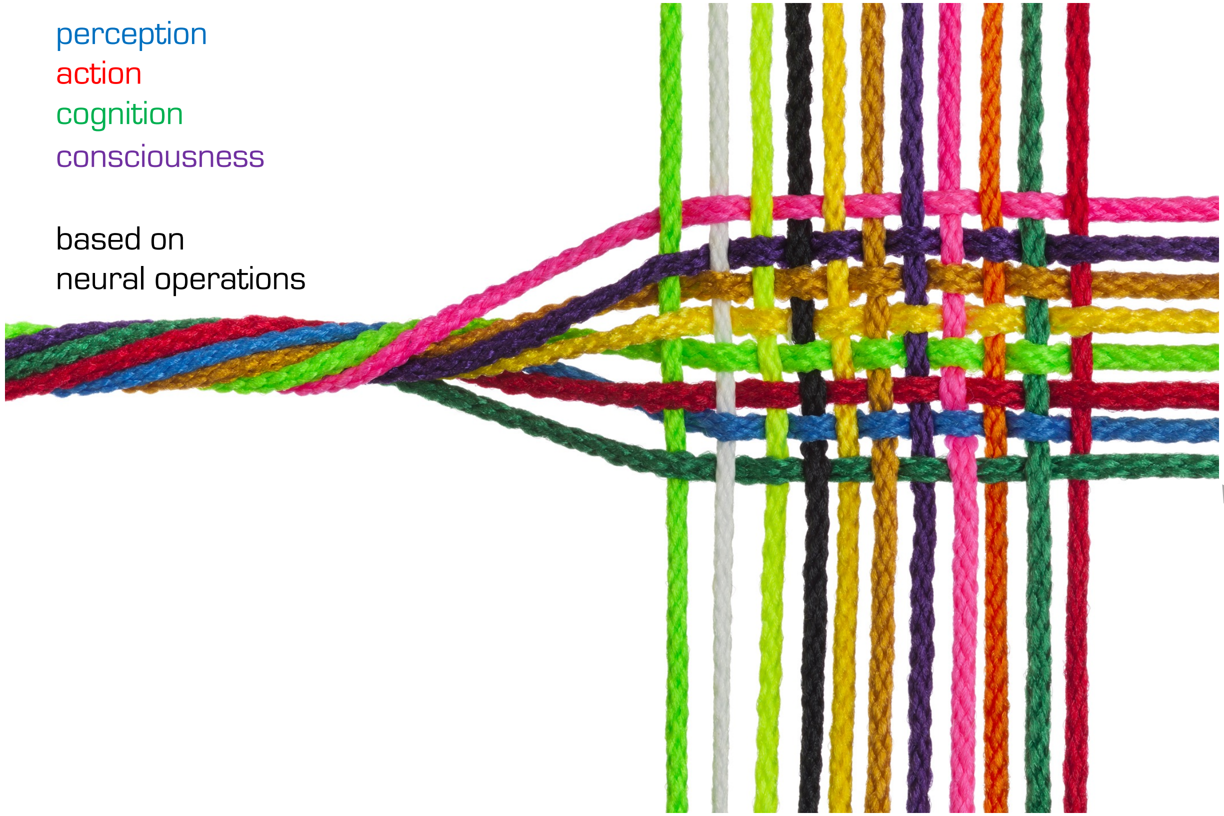

https://fashion-history.lovetoknow.com/fabrics-fibers/weave-types

Bayesian brain: predictive processing (PP) and active inference

Frees energy principle: minimization of average surprisal

Cybernetics: self-regulation, adaptive self-regulation, ultrastability

Interoception and emotion

Embodiment

Enactive cognitive science: weak and strong

Sensorimotor contingency (SMC) theory & PPSMC

Counterfactual PP and active inference

Key idea:
Predictive processing is more naturally applied to interoception

It's more important to avoid unexpected (surprising) internal states that to avoid encountering unexpected exteroceptive states

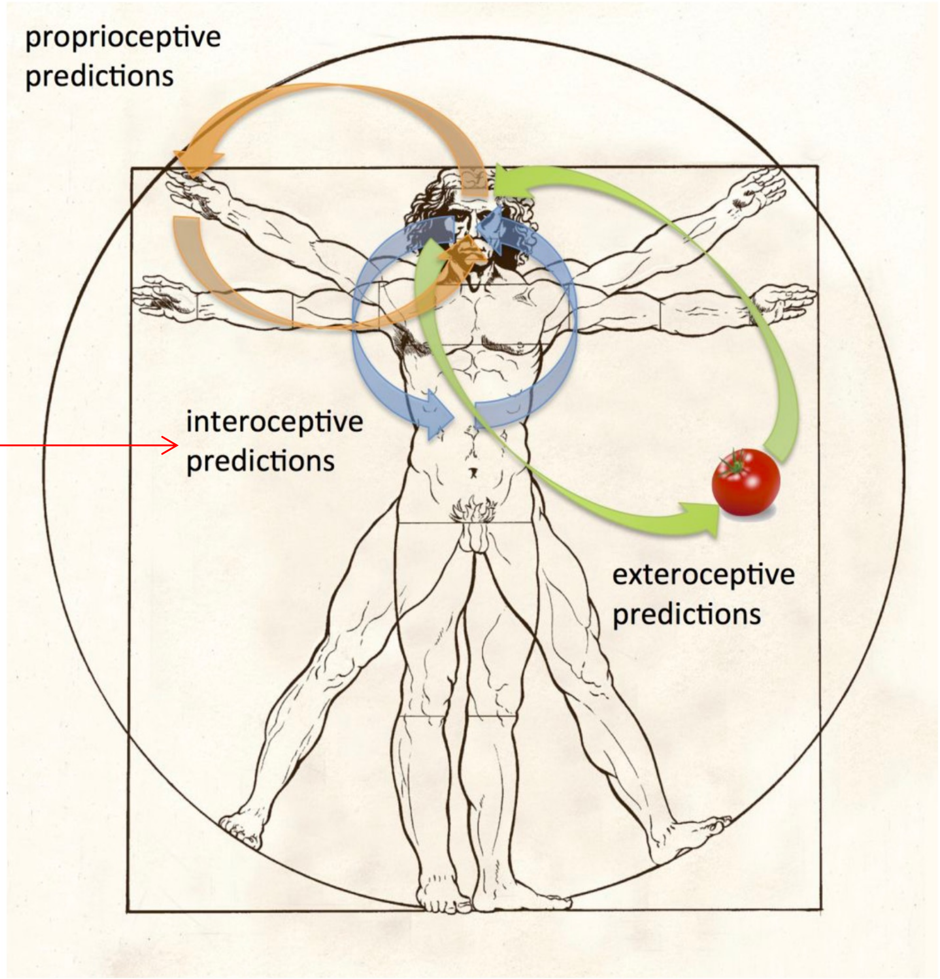

Key idea:
Emotional states arise from top-down predictive inference of the causes of interceptive sensory signals

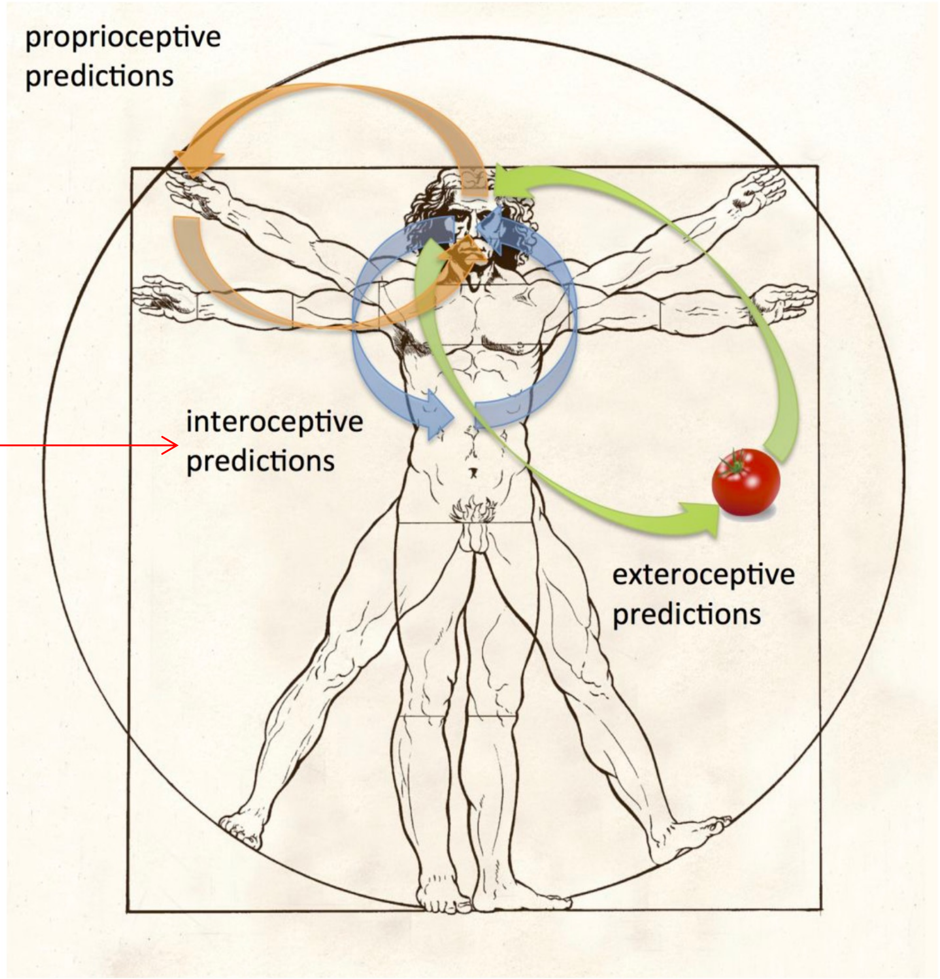

Interoceptive prediction errors can be minimized by:

1. Updating predictive models: interoceptive predictive processing (perception, corresponding to new emotional contents),

2. Changing the interoceptive signals through autonomic reflexes: interoceptive active inference (autonomic control or), or

3. Performing behaviour so as to alter external conditions that impact on internal homeostasis (allostasis)

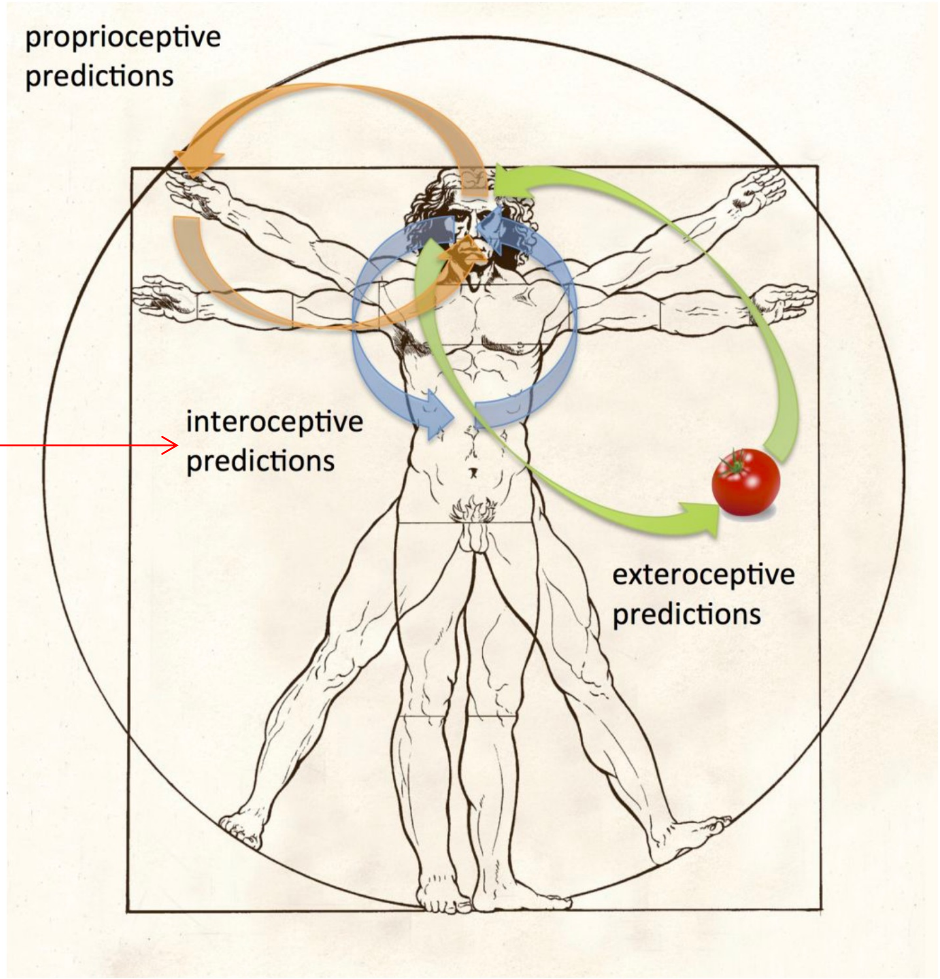

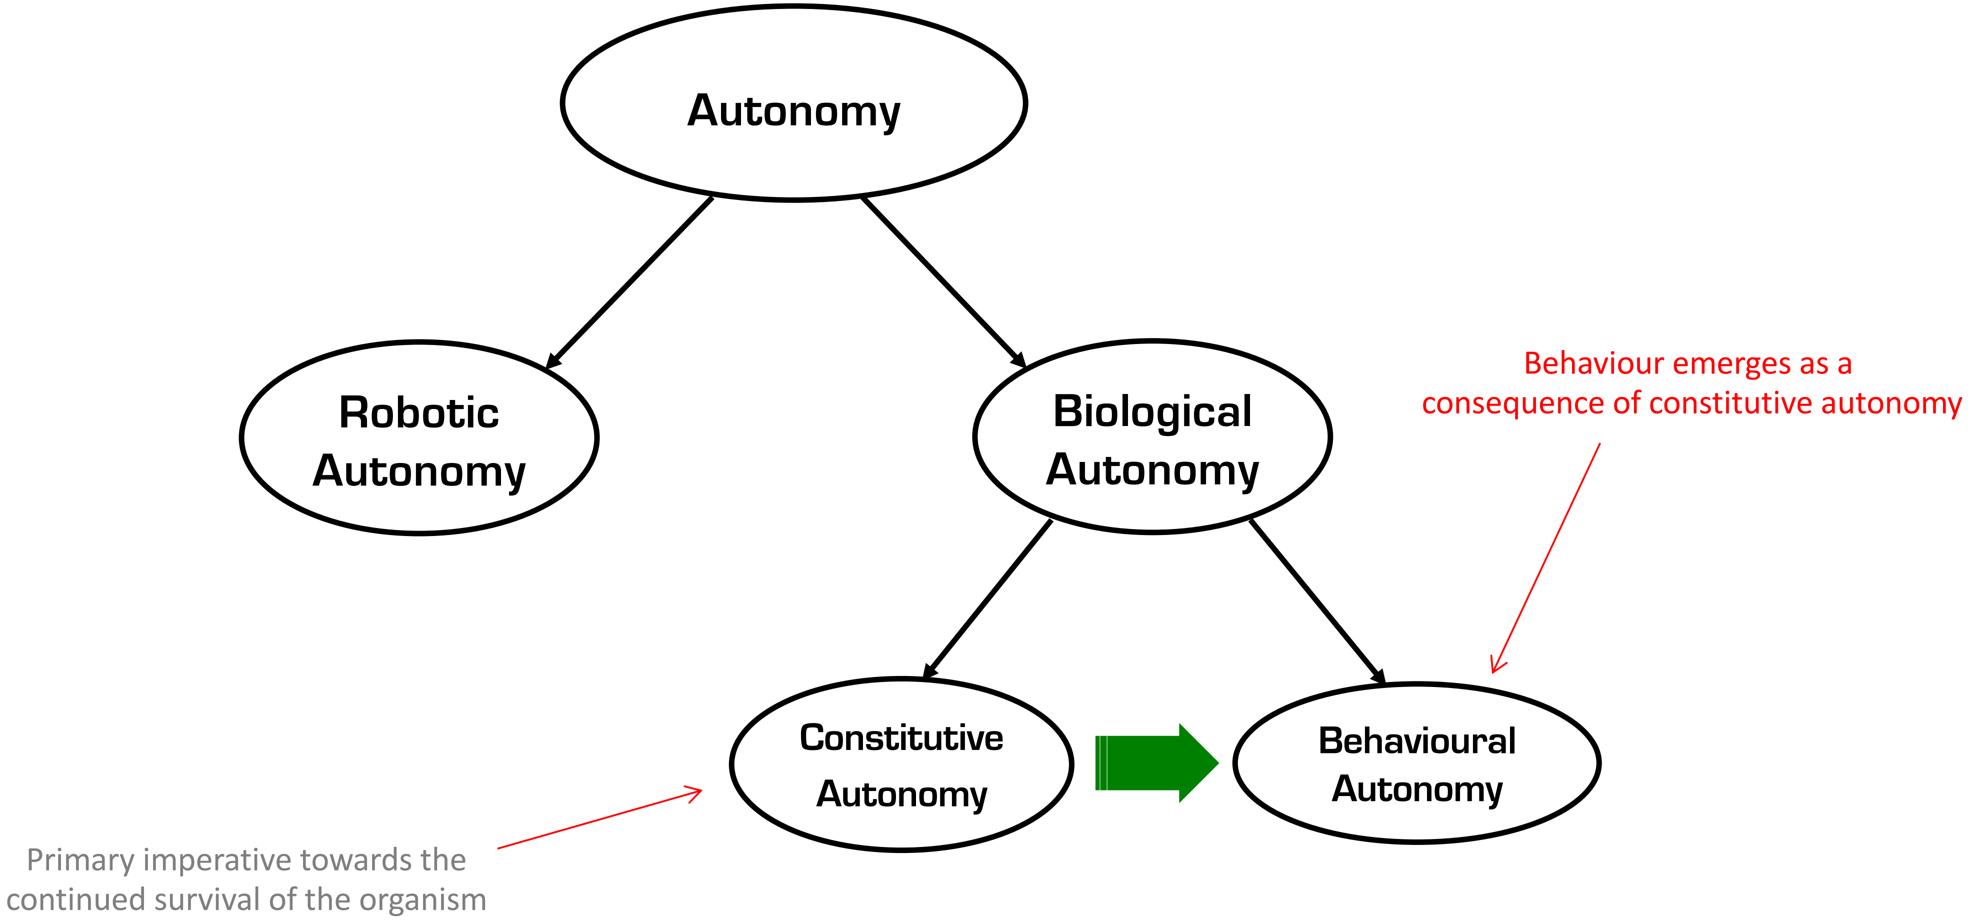
Autonomy
Robotic Autonomy
Biological Autonomy
Behaviour emerges as a consequence of constitutive autonomy
Constitutive Autonomy
Behavioural Autonomy
Primary imperative towards the continued survival of the organism

"Hierarchically higher levels will deploy multimodal
and even amodal predictive models spanning these domains

which are capable of generating multimodal predictions of afferent signals"

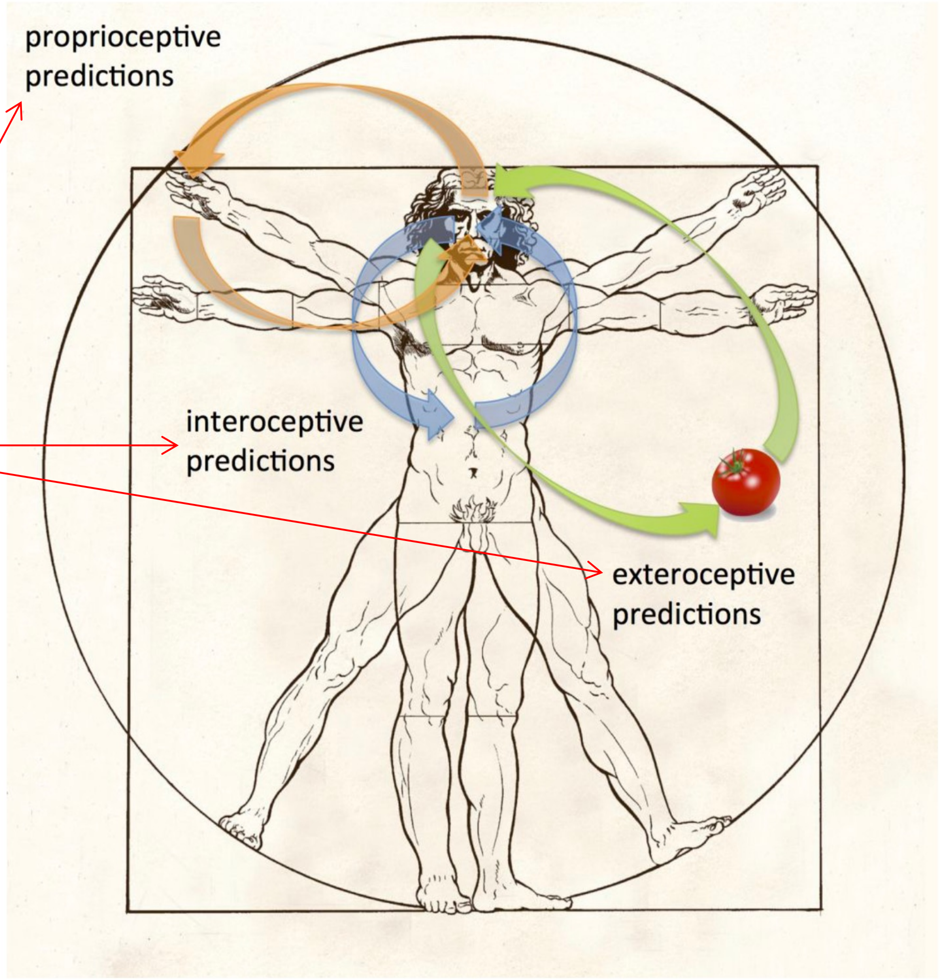

"Interoceptive inference integrates cognition and emotion within the powerful setting of PP"

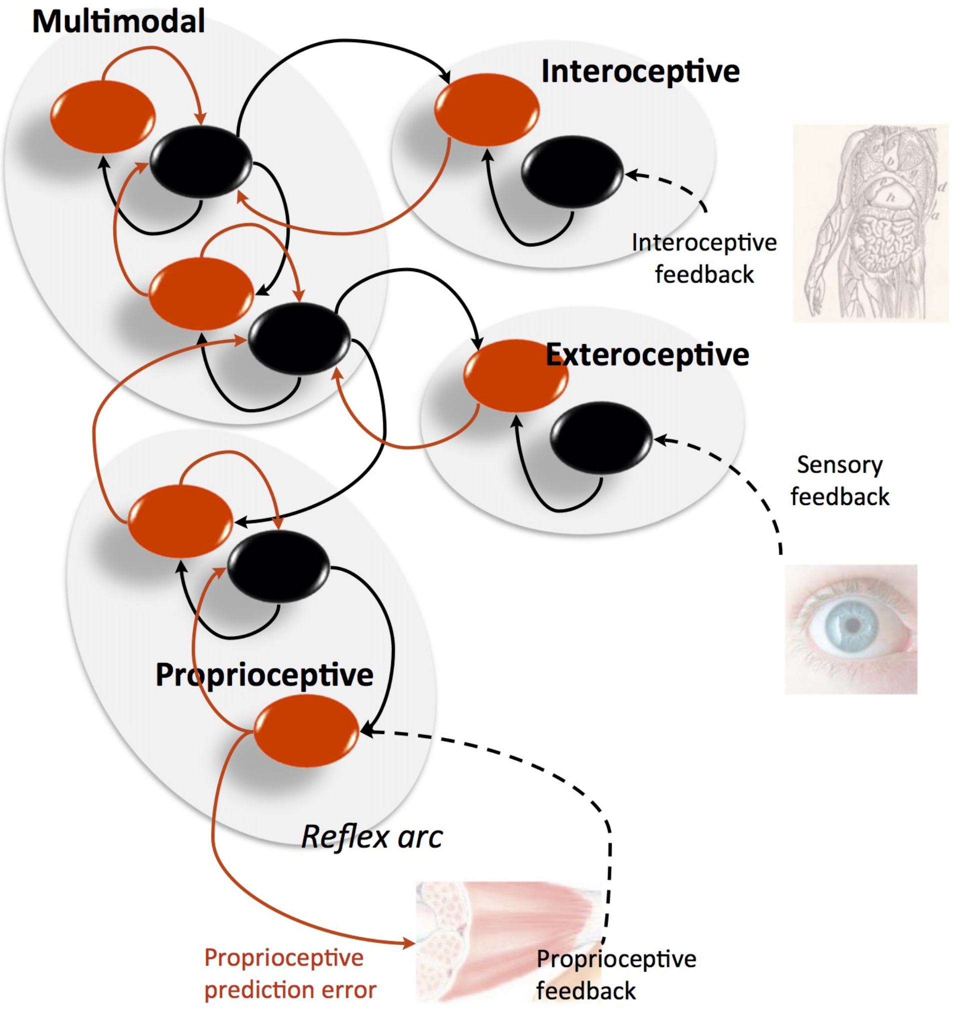

G. Pezzulo and P. Cisek. Navigating the affordance landscape: feedback control as a process model of behaviour and cognition. Trends in Cognitive Sciences, 20(6):414–424, 2016.

Aside: Internal Interaction

- Interoception & Internal Robotics [Parisi 2004]
    - CNS
    - Endocrinal System
- Being “properly embodied” [Stapleton 2013]

D. Parisi. Internal Robotics. Connection Science, 16(4):325– 338, 2004.

M. Stapleton. Steps to a “Properly Embodied” cognitive science. Cognitive Systems Research, 22—23:1–11, 2013.

A single principle
to account for

perception
action
cognition
consciousness

based on
neural operations

Internal models

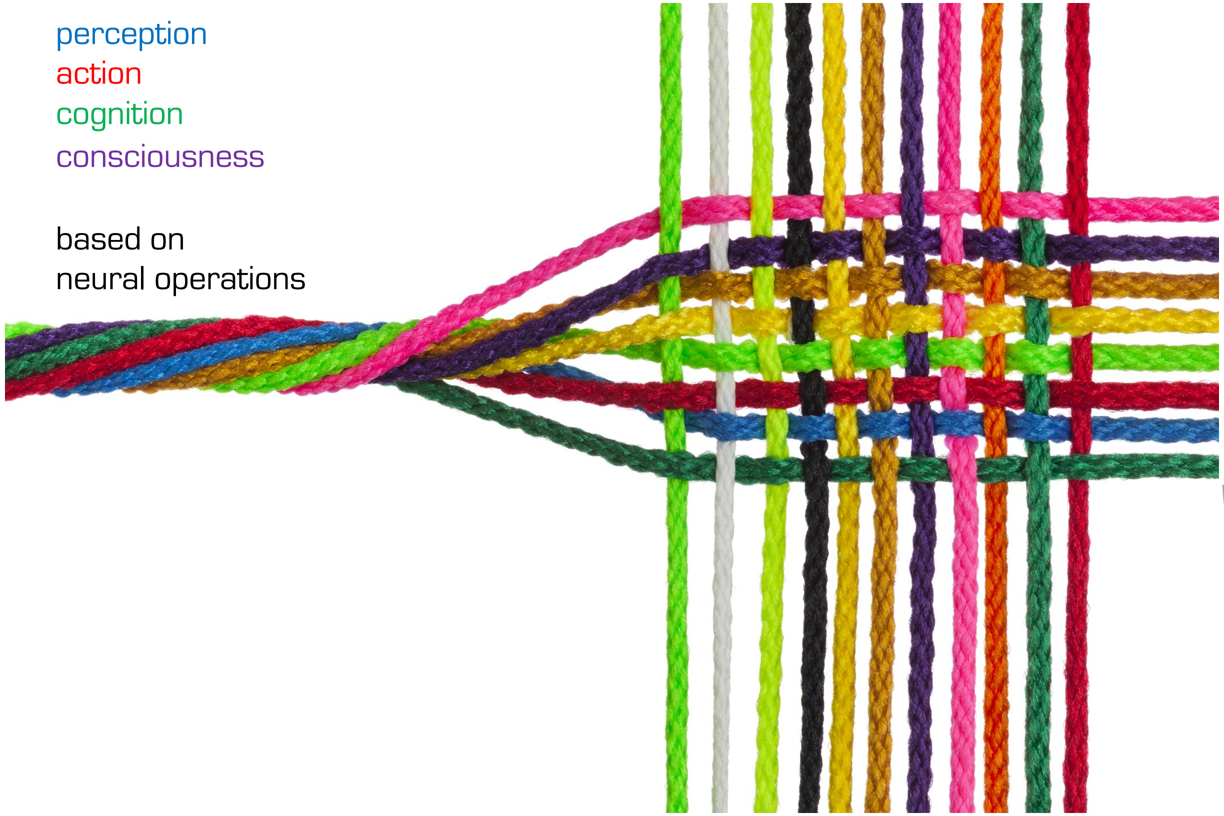

https://fashion-history.lovetoknow.com/fabrics-fibers/weave-types

Bayesian brain: predictive processing (PP) and active inference

Frees energy principle: minimization of average surprisal

Cybernetics: self-regulation, adaptive self-regulation, ultrastability

Interoception and emotion

Embodiment

Enactive cognitive science: weak and strong

Sensorimotor contingency (SMC) theory & PPSMC

Counterfactual PP and active inference

Predictive Processing and Embodiment

"Just as the brain has no direct access to causal structures in the external environment, it also lacks direction access to it's own body"

Predictive Processing and Embodiment

"The experienced body (and self) depends on the brain's best guess of the causes of those sensory signals most likely to be "me"... across interoceptive, proprioceptive, and exteroceptive domains"

The experience of body ownership
is highly plastic

If the (hidden) real hand and the (attended) rubber hand are both stroked synchronously

The rubber hand becomes part of the subject's body schema

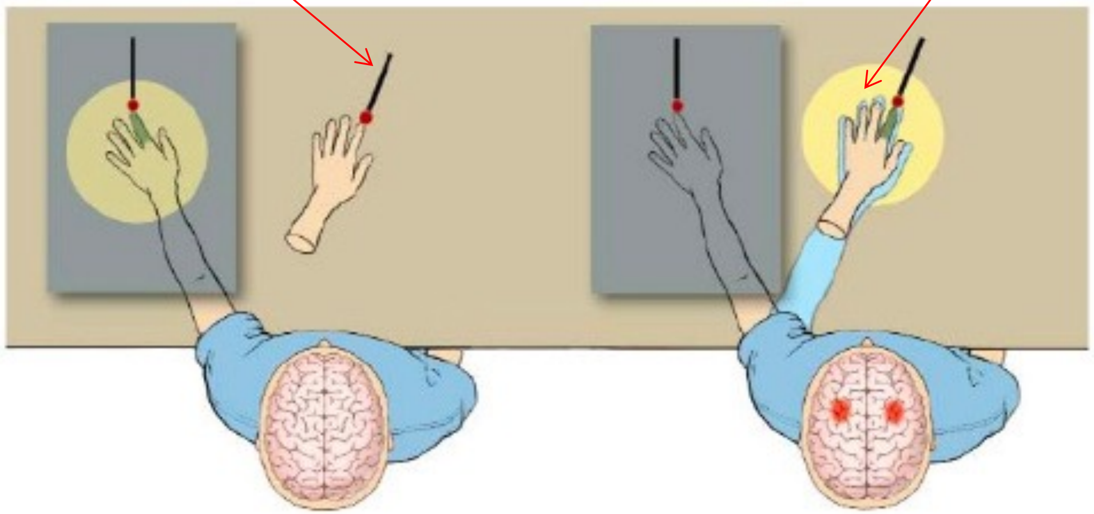

http://embodiedknowledge.blogspot.com/2010/04/rubber-hand-illusion.html

Multisensory integration model: change in experience of body ownership is due to correlation between vision and touch overriding conflicting proprioception

If the (hidden) real hand and the (attended) rubber hand are both stroked synchronously

The rubber hand becomes part of the subject's body schema

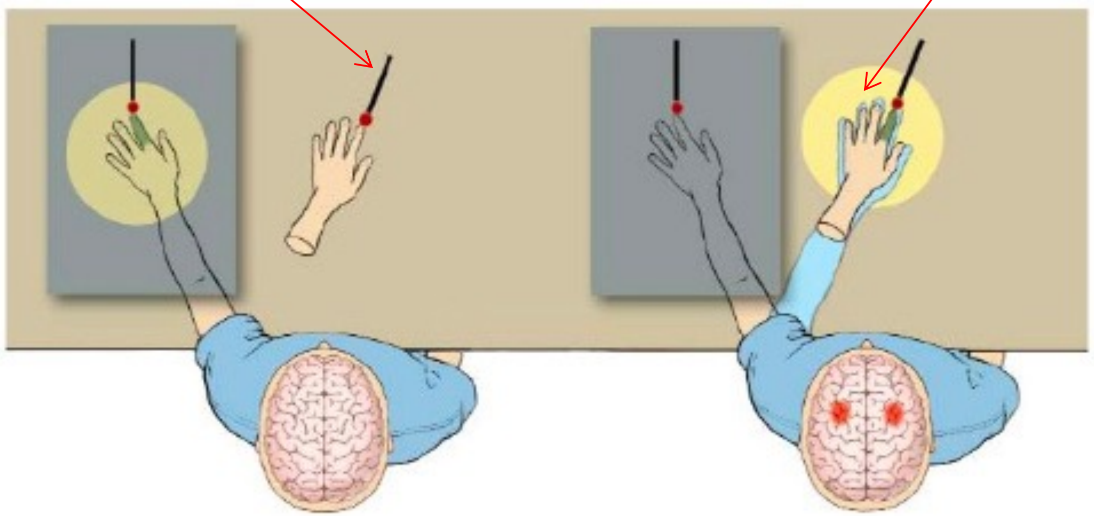

http://embodiedknowledge.blogspot.com/2010/04/rubber-hand-illusion.html

PP perspective: prediction errors induced by multisensory conflicts will over time update self-related priors, with different sensory modalities precision-weighted according to their expected reliability, and strong prior expectations for correlated input

Active inference and Embodiment

Moving a finger strengthens the illusion when a virtual rubber hand mimics the movement: the predicted visual signals are confirmed

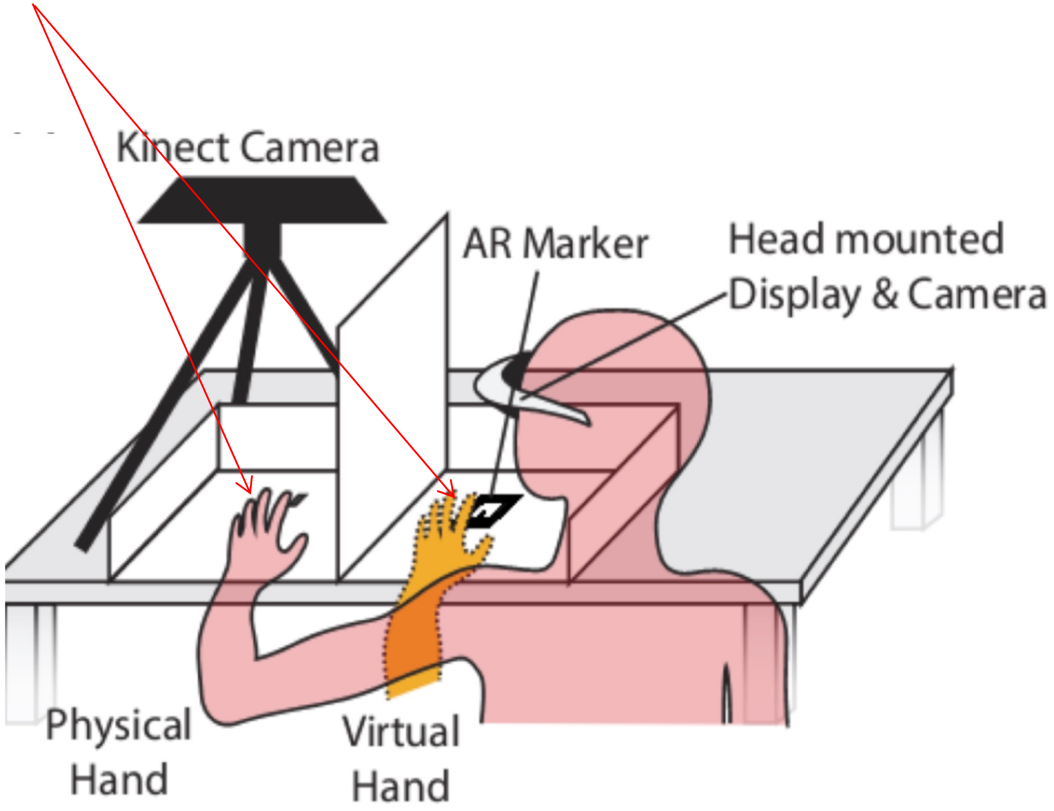

Active Inference and Embodiment

Active inference:

- Confirming sensory predictions ← The usual view

- Seeking "disruptive" actions that test current predictions

  and/or

  Disambiguating competing predictions

New view: counterfactually-equipped predictive models

Self-Modelling in Evolutionary Robotics

Prior potential self-models

Perform random actions

Evaluate self-model based on ability to predict resulting proprioceptive afference signals

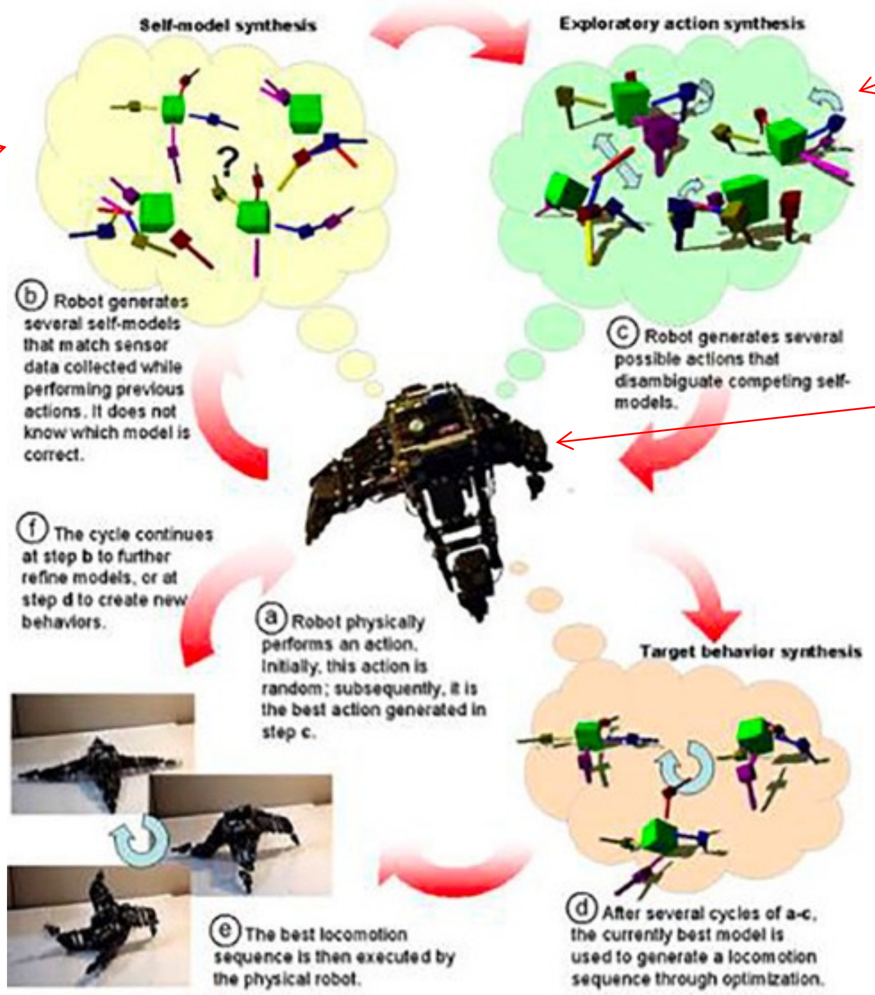

Key step: disambiguate

Evaluate new candidate actions based on the degree to which the current best self-models make different predictions to their (proprioceptive) consequences

Develop a controller to generate forward movement

Unknown morphology

Use the current best model

http://creativemachines.cornell.edu/emergent_self_models

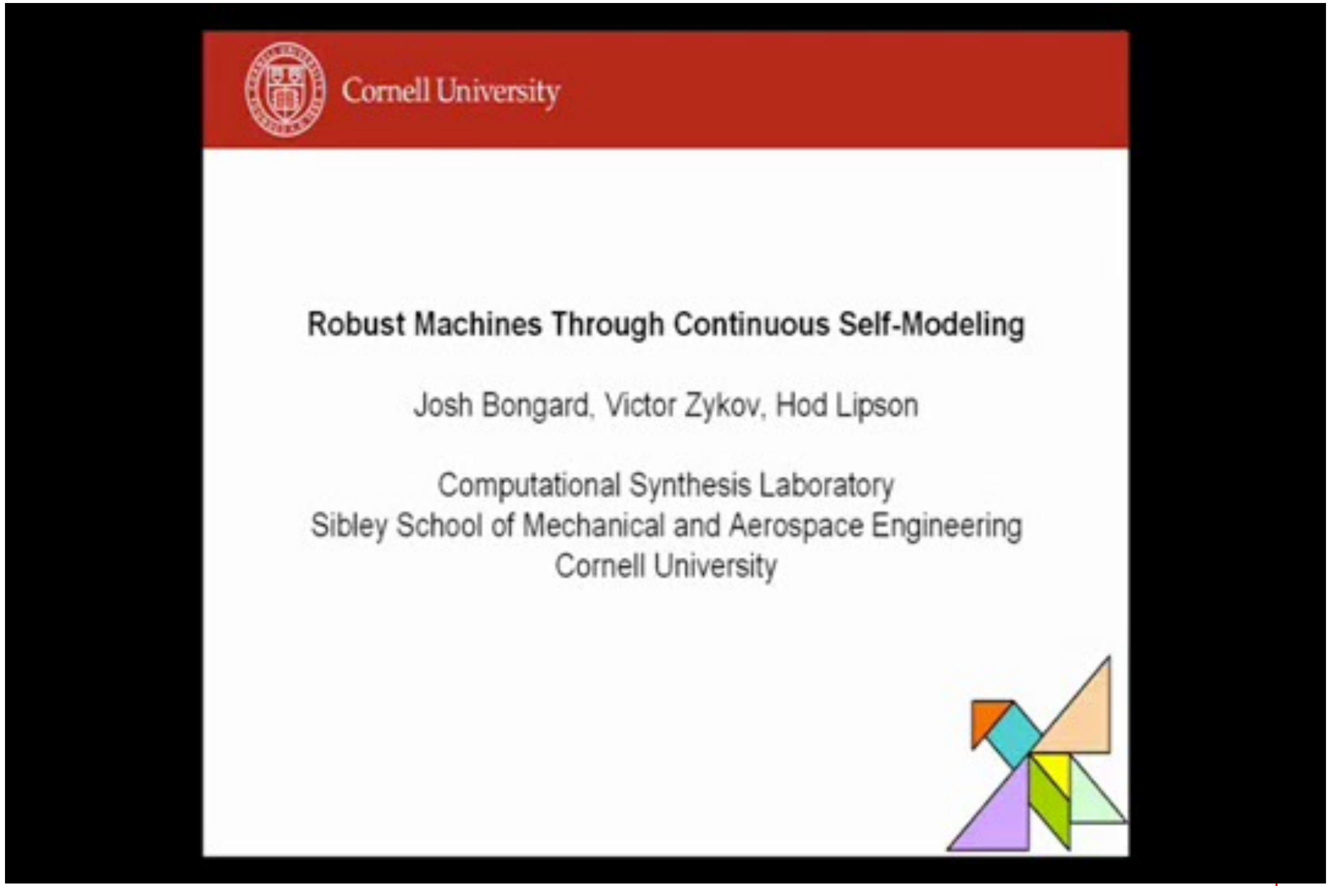

https://www.youtube.com/watch?v=x579QKA6fkY&feature=youtu.be

http://creativemachines.cornell.edu/emergent_self_models

Self-Modelling in Evolutionary Robotics

Operational criterion for a successful self-model:

- Not its fidelity to the physical robot
- Ability to predict sensory inputs under a repertoire of actions

Self-Modelling in Evolutionary Robotics

Agents encode predictions about likely sensory consequences of a range of potential actions

Use actions that are likely to be most informative in disambiguating

$\Rightarrow$ counterfactually-equipped predictive model

Self-Modelling in Evolutionary Robotics

cf. attention based on maximization of Bayesian surprise (Itti & Baldi 2009)

Also, attention based on maximization of information-theory surprisal or self-information (Bruce and Tsotsos 2009)

A single principle
to account for

perception
action
cognition
consciousness

based on
neural operations

Internal models

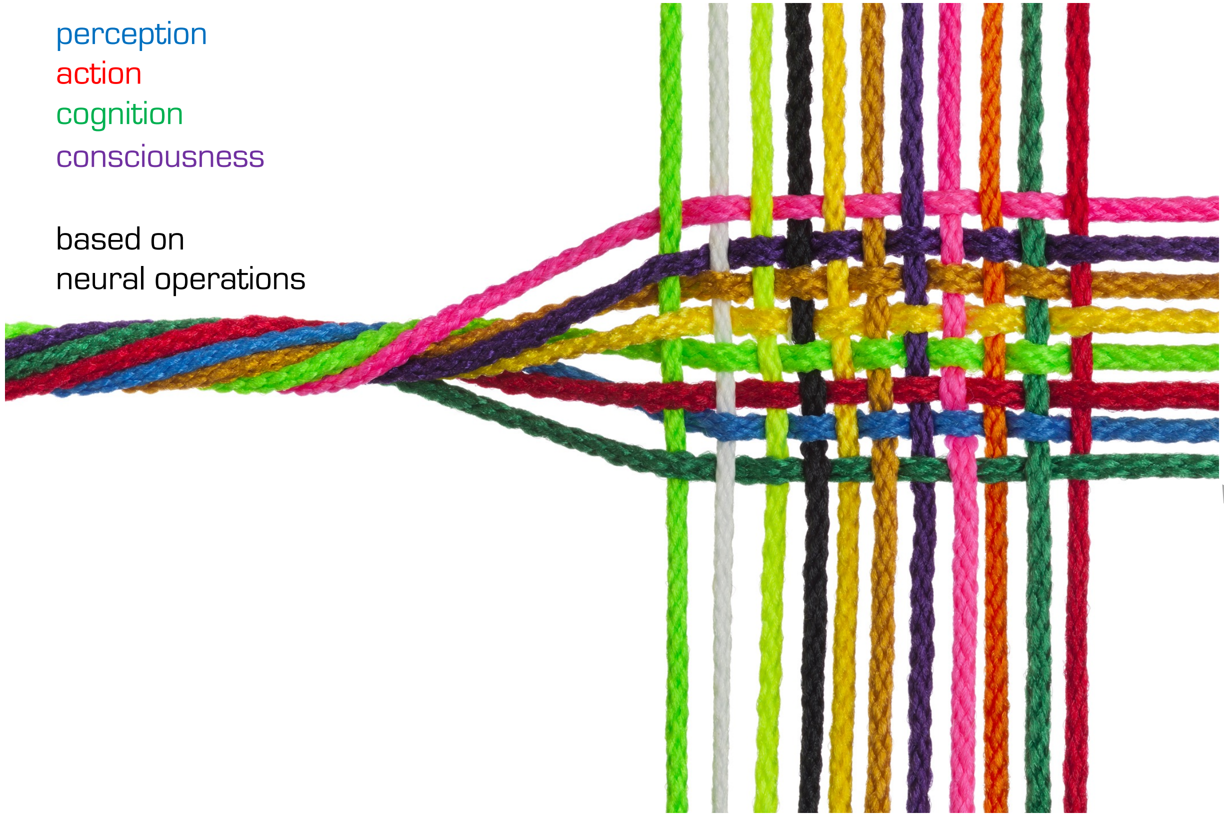

https://fashion-history.lovetoknow.com/fabrics-fibers/weave-types

Bayesian brain: predictive processing (PP) and active inference

Frees energy principle: minimization of average surprisal

Cybernetics: self-regulation, adaptive self-regulation, ultrastability

Interoception and emotion

Embodiment

Enactive cognitive science: weak and strong

Sensorimotor contingency (SMC) theory & PPSMC

Counterfactual PP and active (interoceptive) inference

Enactive Theories, Weak and Strong

"Within cognitive science, ... anti-representationalism is most vociferously defended by the movement variously known as "enactive", "embodied", or "extended" cognitive science"

Enactive Theories, Weak and Strong

Embodiment $\not\Leftrightarrow$ Enaction

Having a body $\not\Rightarrow$ Embodied

Embodied $\not\Rightarrow$ Enactive

Extended cognition $\not\Rightarrow$ Enactive

Enactive $\Rightarrow$ Embodied

Cognition	Necessary Constituents	Typical Characteristics
Embodied	Depends on interpretation	Body and brain are both constitutive elements of the cognitive process
Situated	Brain	Real-time interaction with the environment
Embedded	Brain, body	Exploit the environment and other agents to assist with cognitive activities
Grounded	Brain and body	Experiential modal representations and internal simulation
Extended	Brain, body, environment	Environment is part of the cognitive system
Distributed	Brain, body, environment	Cognitive systems include environmental systems

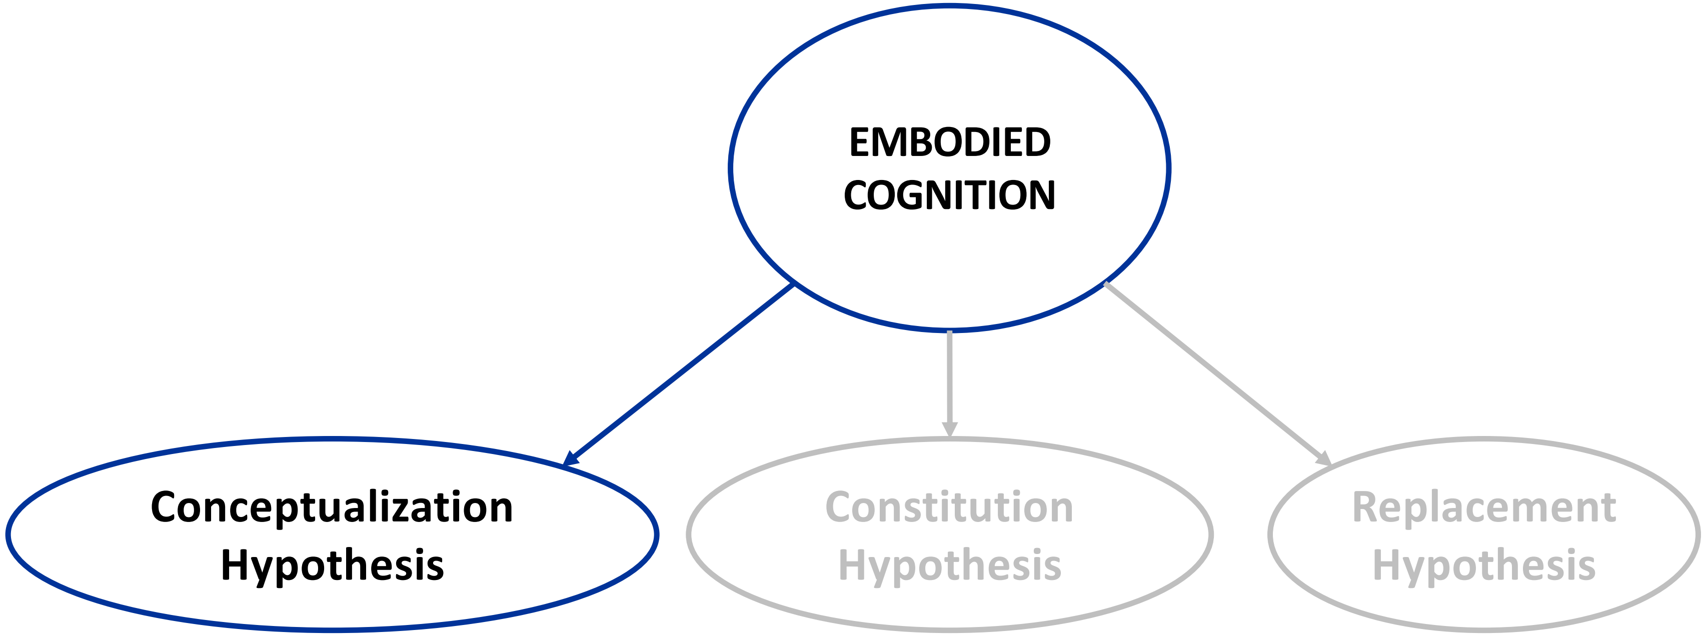

The body conditions / constrains cognition

L. Shapiro. *Embodied Cognition*. Routledge, 2011.

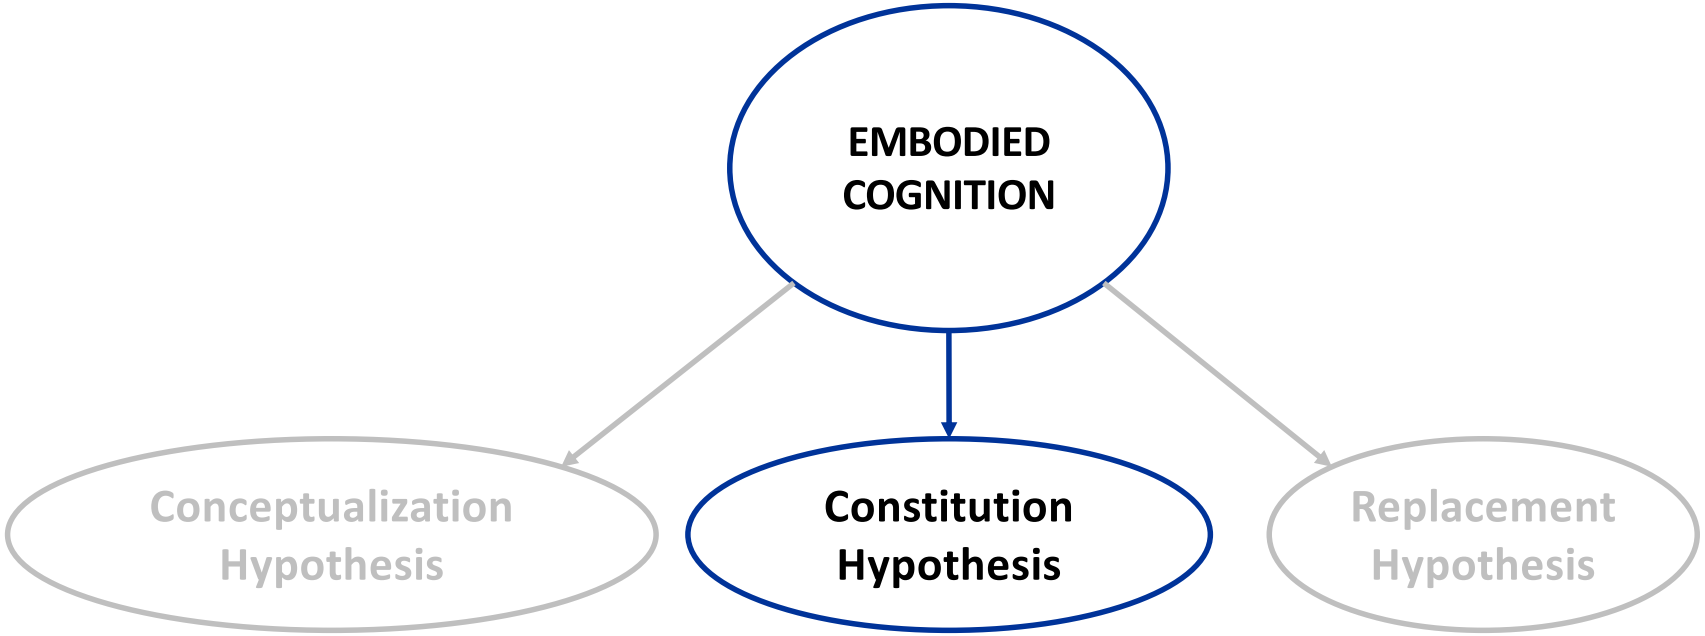

The body is itself an integral part of the process of cognition

The way the body is shaped and moves augments brain-centred neural processing

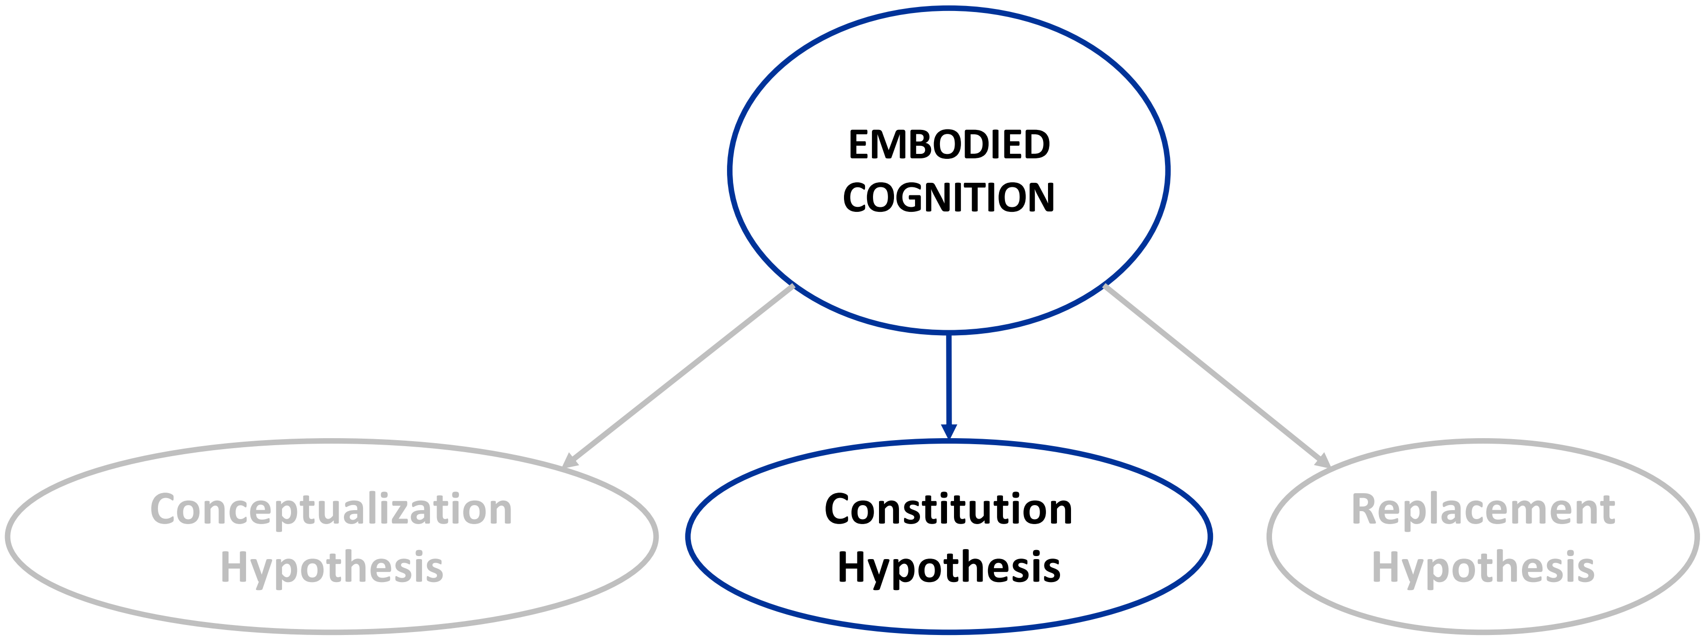

Cognition is distributed between the neural and the non-neural
The body simplifies what the brain has to do or takes over responsibility for it completely

L. Shapiro. Embodied Cognition. Routledge, 2011.

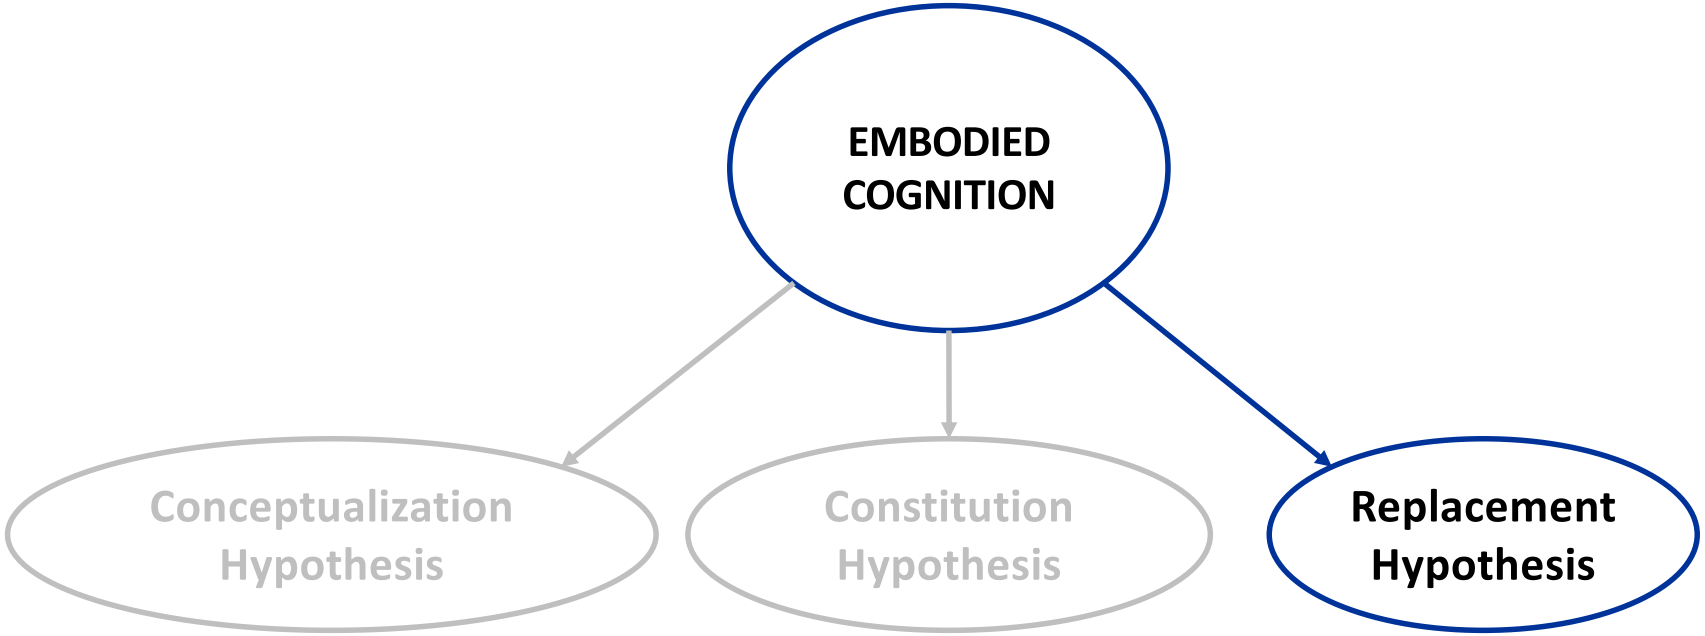

An agent's body in real-time interaction with its environment **replaces** the need for representational processes

L. Shapiro. Embodied Cognition. Routledge, 2011.

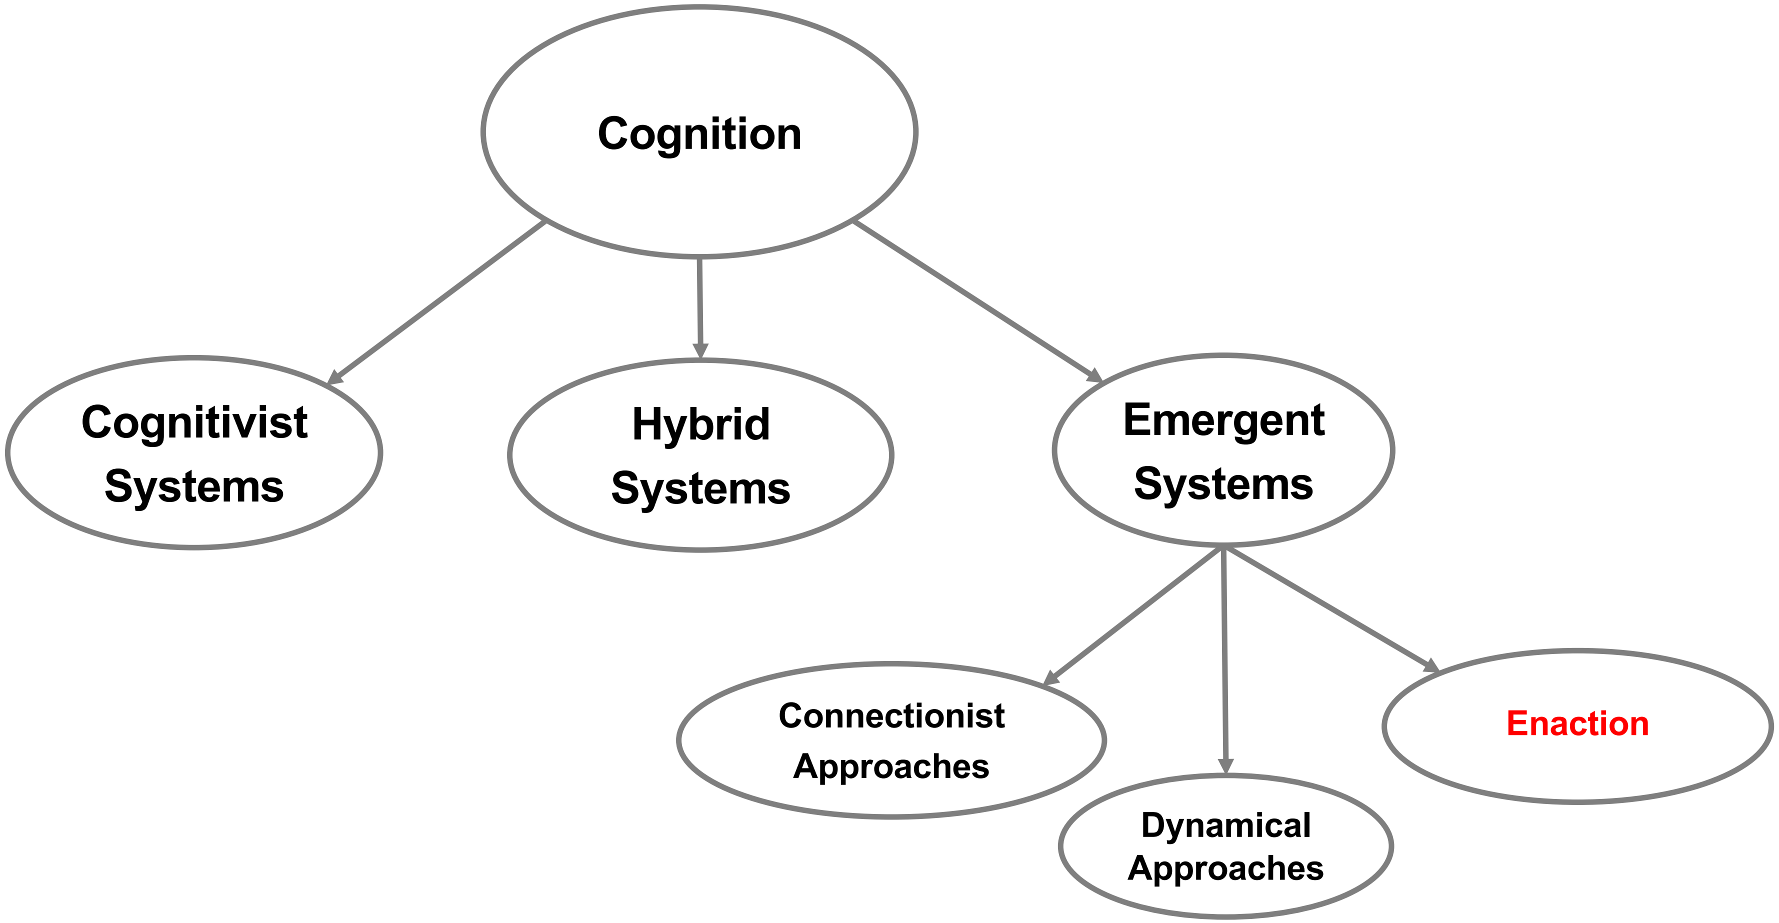
Cognition
Cognitivist Systems
Hybrid Systems
Emergent Systems
Connectionist Approaches
Dynamical Approaches
Enaction

Enaction

Orthodoxy (cognitivist)

World as the system experiences it is independent of the cognitive system (knower)

Enactive view

Known and knower "stand in relation to each other as mutual specification: they arise together"

Closely linked to phenomenology

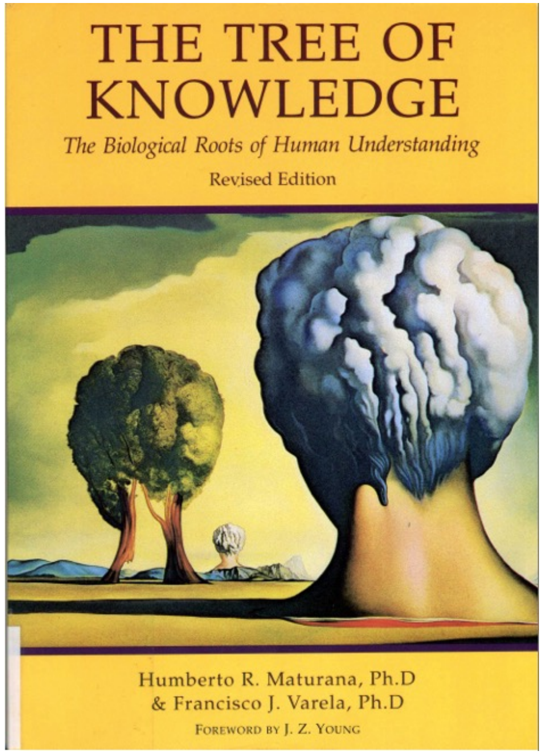

Enaction

Five key elements to enactive systems

- Autonomy
- Embodiment
- Emergence
- Experience
- Sense-making

Enaction

Sense-making

- Knowledge is generated by the system itself
- Captures some regularity or lawfulness in the interactions
- The ‘sense’ is dependent on the way interaction can take place
  - Perception & Action
- Modify its own state (CNS) to enhance
  - Predictive capacity
  - Action capabilities

Enaction

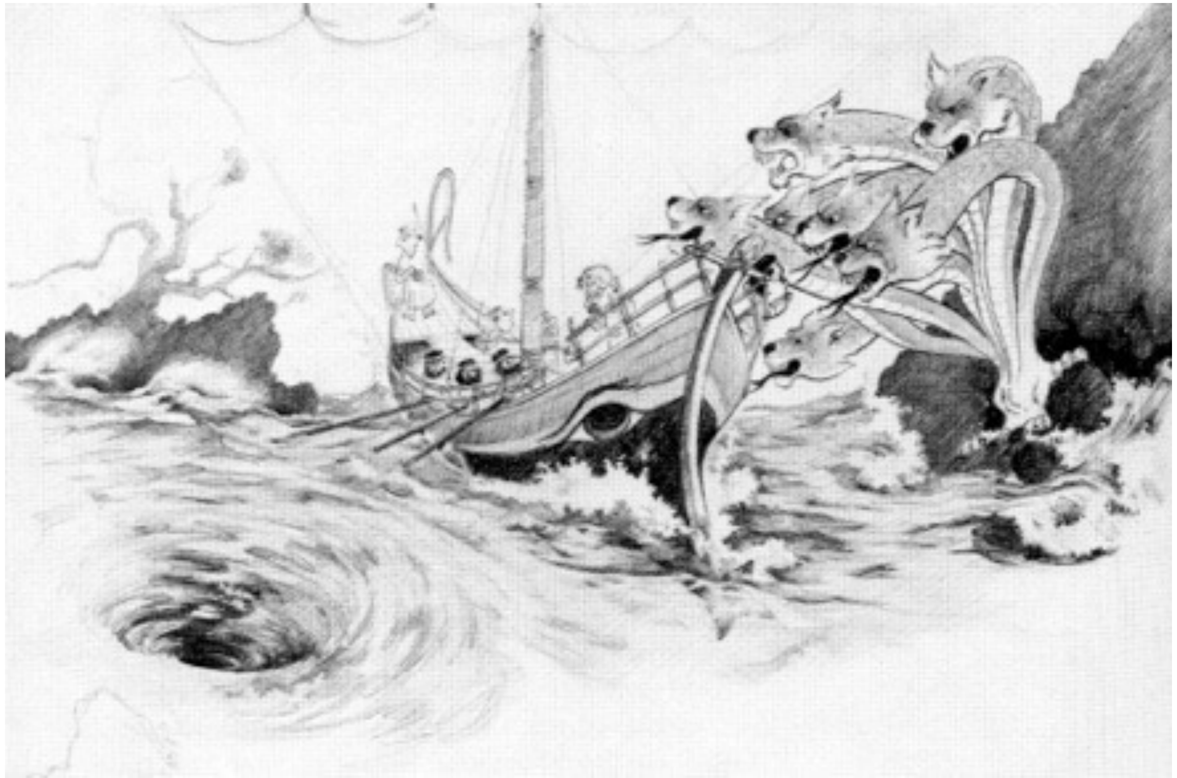

The epistemological Odyssey: sailing between the Scylla monster of representationalism and the Charybdis whirlpool of solipsism.
From Maturana and Varela, Tree of Knowledge, p. 134

A single principle
to account for

perception
action
cognition
consciousness

based on
neural operations

Internal models

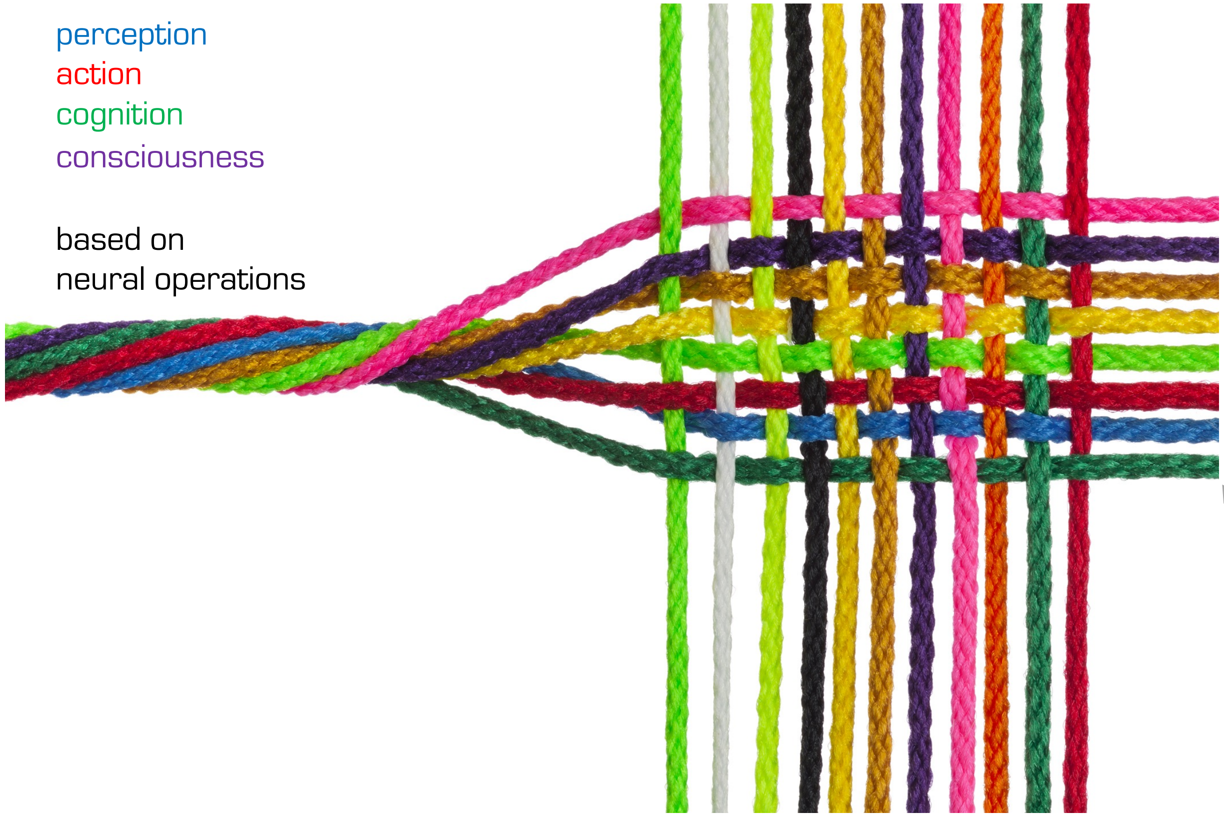

https://fashion-history.lovetoknow.com/fabrics-fibers/weave-types

Bayesian brain: predictive processing (PP) and active inference

Frees energy principle: minimization of average surprisal

Cybernetics: self-regulation, adaptive self-regulation, ultrastability

Interoception and emotion

Embodiment

Enactive cognitive science: weak and strong

Sensorimotor contingency (SMC) theory & PPSMC

Counterfactual PP and active inference

Sensorimotor contingency (SMC) theory & PPSMC

"SMC theory claims that experience and perception are not things that are "generated" by the brain ... but are, rather, "skills" consisting of fluid patterns of on-going interaction with the environment"

Close coupling of perception and action

Interaction structures perception, action, and behaviour

Structural-coupling, in the language of enaction

Sensorimotor contingency (SMC) theory & PPSMC

"SMC theory claims that experience and perception are not things that are "generated" by the brain ... but are, rather, "skills" consisting of fluid patterns of on-going interaction with the environment"

Mastery of an SMC requires and essentially counterfactual knowledge of relations between particular actions and the resulting sensations

SMCs ... relate potential actions to their likely effects

Sensorimotor contingency (SMC) theory & PPSMC

1. "Hierarchical active inference implies the existence of predictive models encoding information very much like that required by SMC"

2. Counterfactually-rich predictive models

Key idea: encoding of how sensory inputs change based on a repertoire of possible actions, even if those actions are not performed

Sensorimotor contingency (SMC) theory & PPSMC

1. "Hierarchical active inference implies the existence of predictive models encoding information very much like that required by SMC"

2. Counterfactually-rich predictive models

Key idea: encoding of how sensory inputs change based on a repertoire of possible actions, even if those actions are not performed

Finesses the problem of representation:
by asserting only a collection (repertoire) of internal models that encapsulate associations between sensory data and predicted proprioceptive and exteroceptive consequences

Finesses the problem of representation:

By asserting only a collection of internal models:

Associations between predictions of sensory data
(i.e. proprioceptive and exteroceptive consequences)

and possible actions
(i.e. resolution of proprioceptive prediction errors)

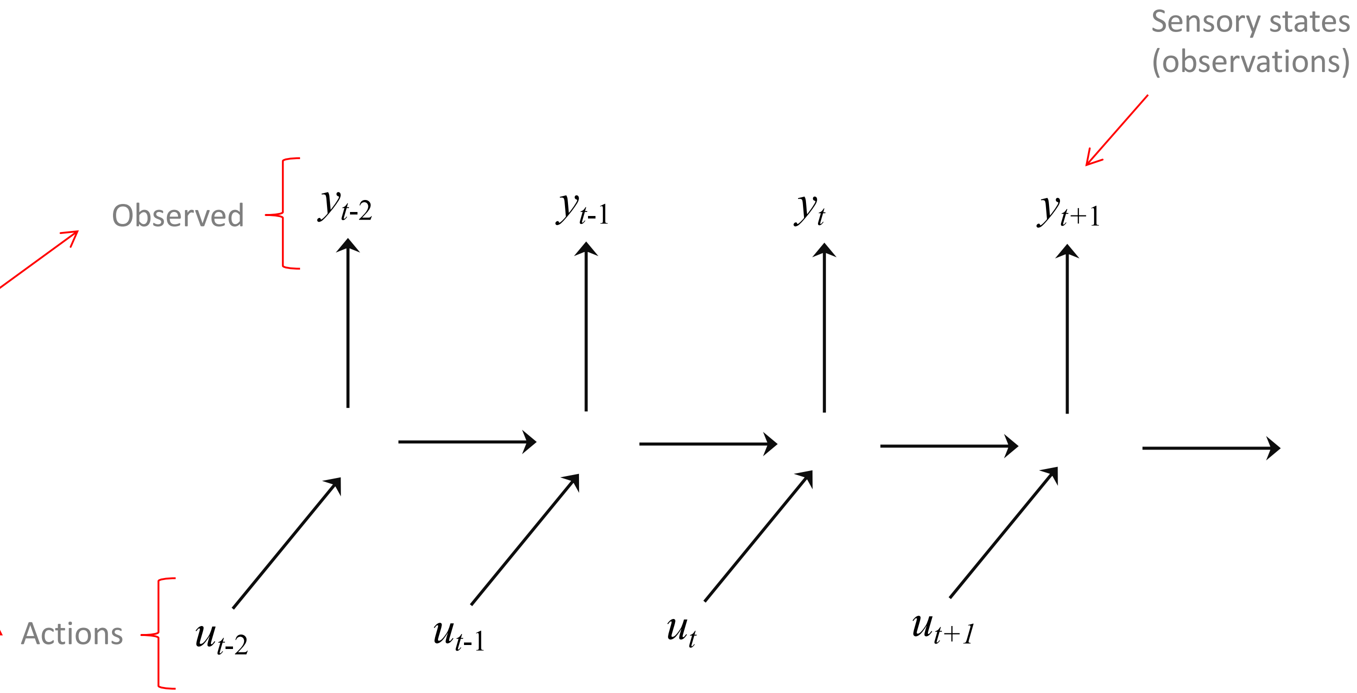

Sensorimotor contingency (SMC) theory & PPSMC

"We experience normal perception as world-revealing precisely because the predictive models underlying perceptional content specify a rich repertoire of counterfactually explicit probability densities encoding the mastery of SMSs"

A. Seth, "The Cybernetic Bayesian Brain: From Interoceptive Inference to Sensorimotor Contingencies", in T. Metzinger & J. M. Windt (Eds). Open MIND: 35(T). Frankfurt am Main: MIND Group,1–24, 2015.

A single principle
to account for

perception
action
cognition
consciousness

based on
neural operations

Internal models

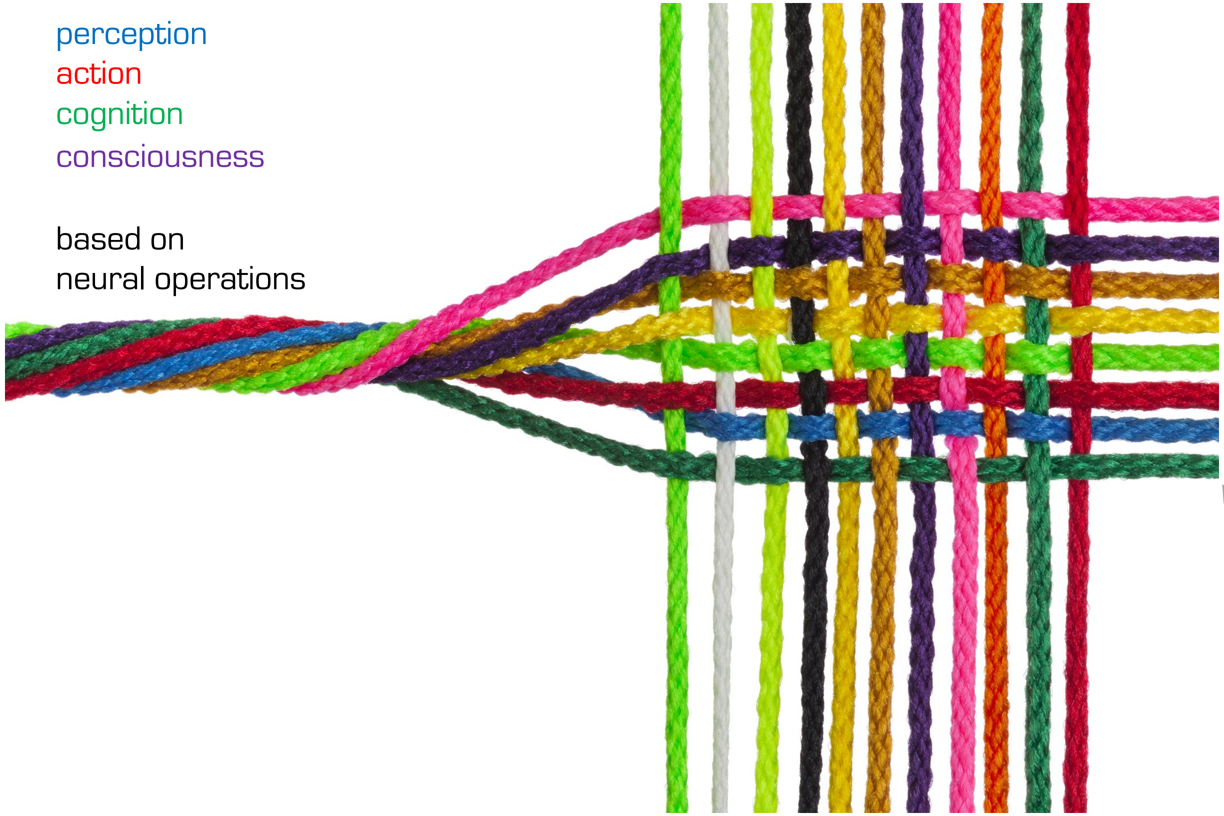

https://fashion-history.lovetoknow.com/fabrics-fibers/weave-types

Bayesian brain: predictive processing (PP) and active inference

Frees energy principle: minimization of average surprisal

Cybernetics: self-regulation, adaptive self-regulation, ultrastability

Interoception and emotion

Embodiment

Enactive cognitive science: weak and strong

Sensorimotor contingency (SMC) theory & PPSMC

Counterfactual PP and active inference

Counterfactual PP and Active Inference

Several forms of active inference, tied to:

1. Proprioception

2. Exteroception

3. Interoception

Counterfactual PP and Active Inference

Active inference tied to

1. Proprioception

   - Actions emerge from the minimization of proprioceptive prediction errors
   - By engaging classical reflex arcs
   - Requires generative models that predict time-varying flows of proprioceptive input

Counterfactual PP and Active Inference

Active inference tied to

2. Exteroception

- Actions are engaged to generate new sensory samples
- To minimize the perceptual prediction errors
- Actions can be selected in three ways:
    1. To confirm current perceptual hypotheses
    2. To disconfirm current perceptual hypotheses
    3. To disambiguate between competing hypotheses

Necessitates counterfactual PP with multiple predictive models

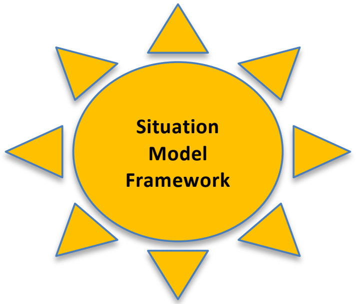

Counterfactual PP and Active Inference

Active inference tied to

3. Interoception

- Works for autonomic and allostatic regulation
- But what of counterfactual predictive processing?

    1. To confirm current interoceptive hypotheses
    2. To disconfirm current interoceptive hypotheses
    3. To disambiguate between competing hypotheses

Linking fictive interoceptive signals and likely causes to autonomic or allostatic regulation

"We do not want to drive our essential variables continually close to viability limits"

Maybe we do, e.g., during training?

Expanded view of the Free Energy Principle

The long-run survival of a system

not just by by minimizing prediction errors

but by inducing the most predictive model of the causes of sensory signals, requiring
counterfactually-rich predictive processing and
disruptive and/or disambiguating active inference
"in order to always put the current-best model to the test"

A single principle
to account for

perception
action
cognition
consciousness

based on
neural operations

Internal models

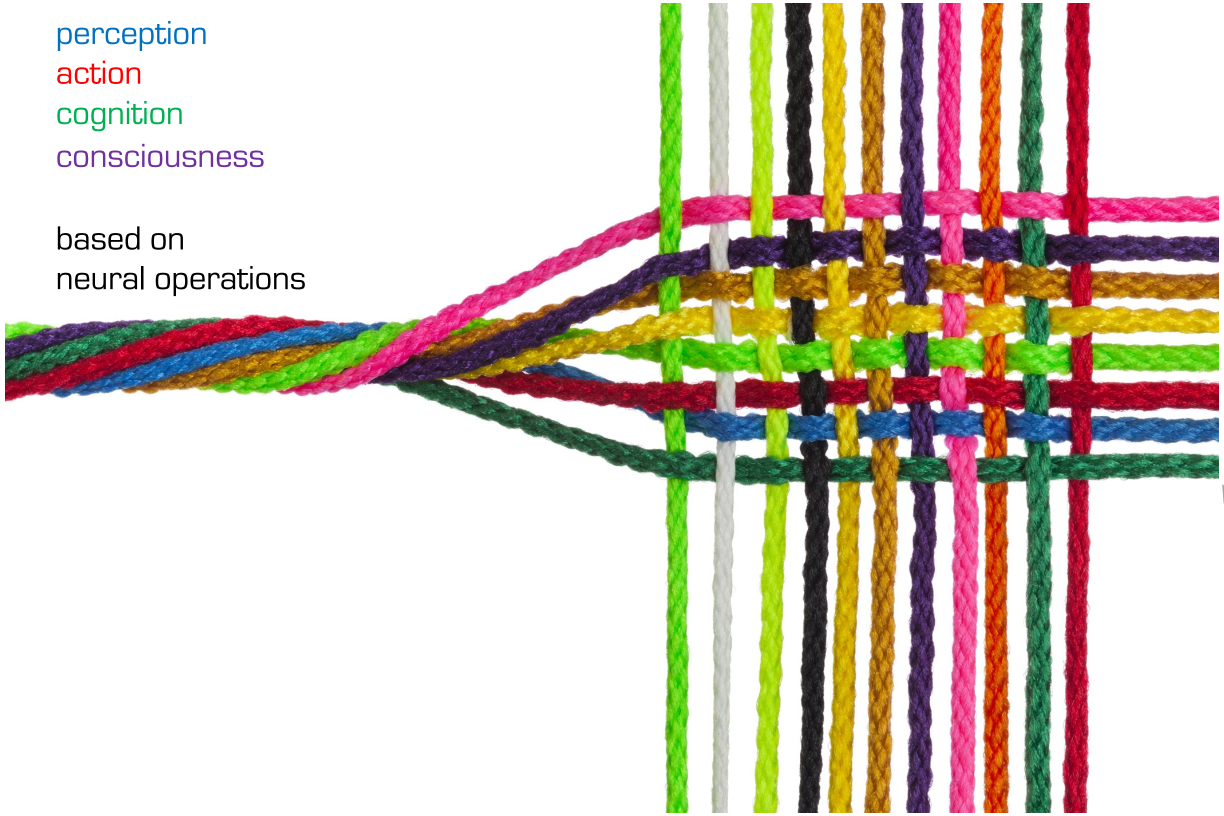

https://fashion-history.lovetoknow.com/fabrics-fibers/weave-types

Bayesian brain: predictive processing (PP) and active inference

Frees energy principle: minimization of average surprisal

Cybernetics: self-regulation, adaptive self-regulation, ultrastability

Interoception and emotion

Embodiment

Enactive cognitive science: weak and strong

Sensorimotor contingency (SMC) theory & PPSMC

Counterfactual PP and active inference

"A distinctive integration of
predictive processing, cybernetics, and enactivism"

A shift from the perspective of perceptual inference as furnishing representations of the external world for the consumption of general-purpose cognitive mechanisms, towards model-based predictive control as a primary survival imperative from which perception, action, and cognition ensue"

https://unsplash.com/photos/-TQUERQGUZ8?utm_source=unsplash&utm_medium=referral&utm_content=creditCopyText

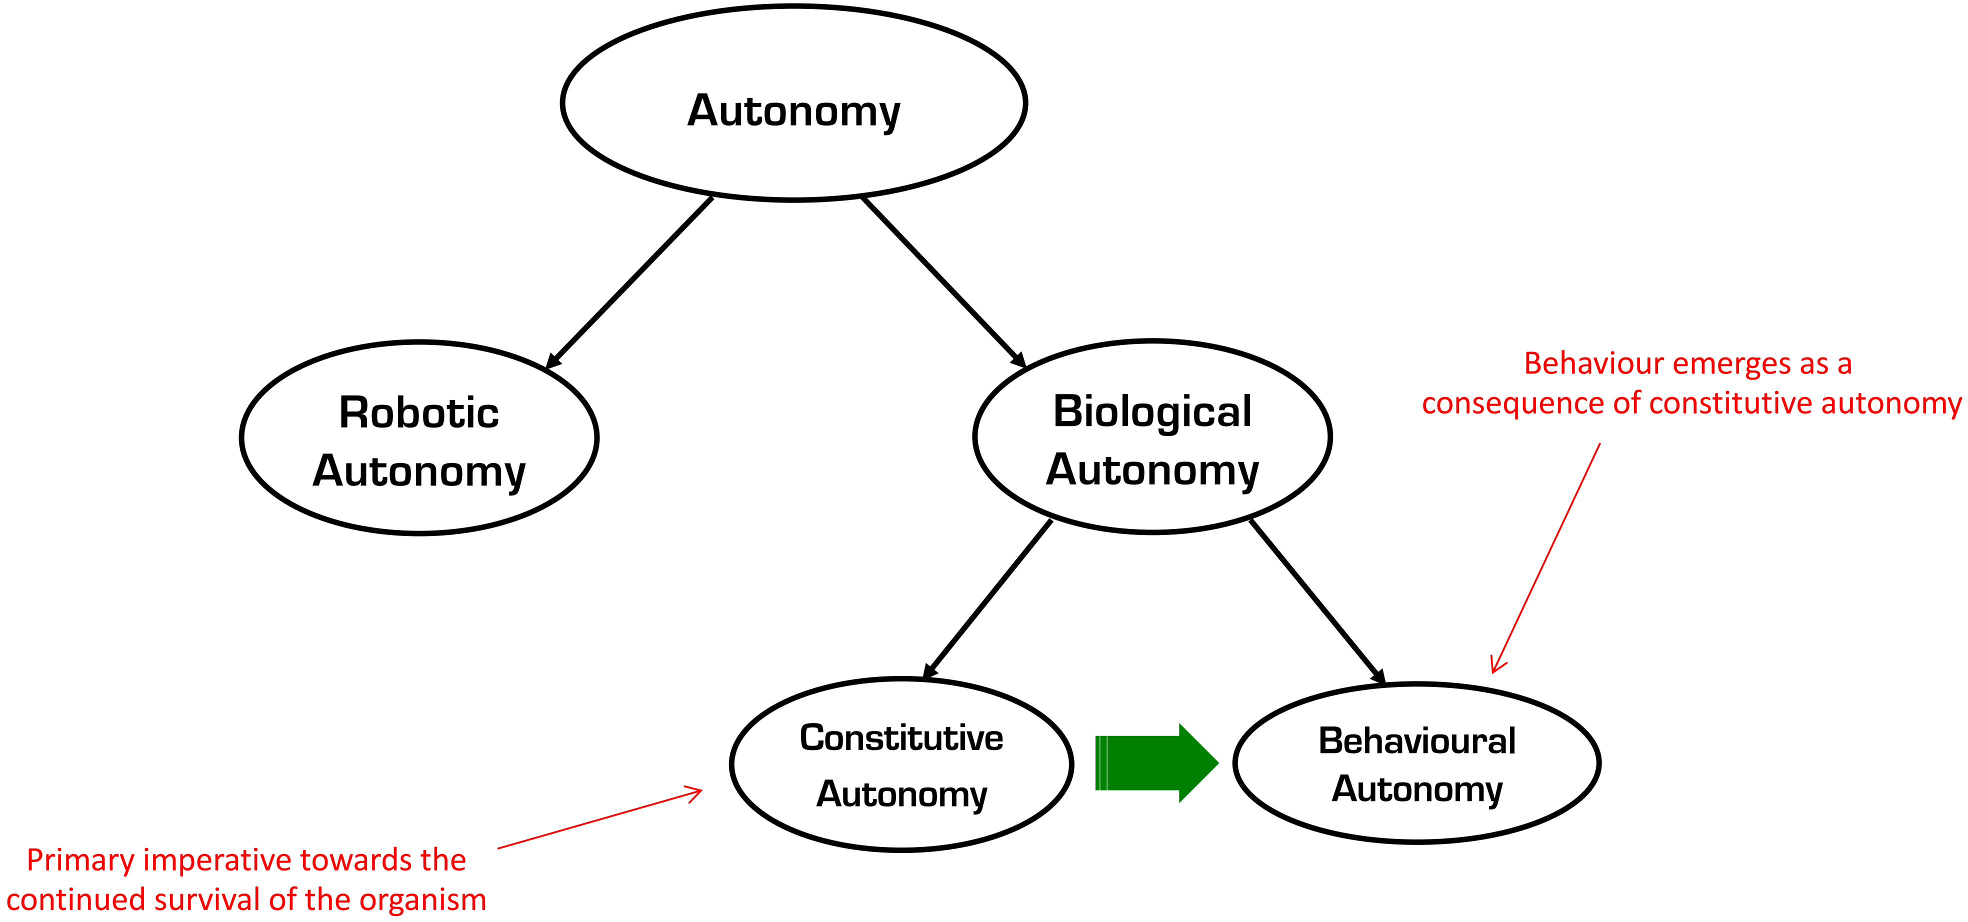
Autonomy
Robotic Autonomy
Biological Autonomy
Behaviour emerges as a consequence of constitutive autonomy
Constitutive Autonomy
Behavioural Autonomy
Primary imperative towards the continued survival of the organism

Caveat

"Even an exhaustive treatment would reveal that this literature so far provides only circumstantial support for the basics of PP, let alone the extensions described here"

Box 2. Active Inference as a Modern Version of Cybernetic Theory

Active inference is essentially a (Bayesian) predictive coding architecture extended with **reflexes** [14,37,78]. Predictive coding was first proposed as a model of visual perception, in which the hierarchical layers are coupled through top-down and bottom-up signals, encoding predictions and prediction errors, respectively, and weighted by their precision (inverse variance). Top-down and bottom-up dynamics serve to suppress prediction errors (or free energy [75]); sensory mismatches at the lowest layer propagate upward and help revise (higher) perceptual hypotheses. In contrast to predictive coding, active inference can also minimize prediction error by acting: by engaging reflexes that suppress residual (proprioceptive) errors. For example, if one expects to see a berry but does not see it, not only can he revise the perceptual hypothesis ('there is no berry') but he can also put the berry in front of him or search for the berry by moving the eyes, until there is no more prediction error.

Active inference implements planning in a way that resembles the idea that distal affordances (e.g., apple reachability) can influence the competition between proximal actions (picking the berry versus walking) [14]. It uses a hierarchical generative (forward) model to predict action consequences, and the ensuing 'value' of possible action sequences (plans) by considering – iteratively – whether the distal states they make accessible approximate the goal states (encoded as prior preferences). These plan values enable the selection of immediate actions: the greater the plan's value, the more likely it is to specify the next action.

As these examples illustrate, active inference can be considered a biologically grounded synthesis of cybernetic ideas (on homeostasis and control) and the Bayesian brain hypothesis. This might seem odd – because cybernetic theory often dispenses with an 'inner model', while according to the Bayesian brain hypothesis, the brain is a statistical machine that learns world models. However, in active inference the necessity of models stems from control principles (e.g., the 'good regulator theorem' that the best regulator requires a model [79]). Furthermore, there is an essential contribution of the body and environment in structuring the content of generative models, because it needs to embody the structure of sensorimotor interactions. Although the representational aspects of active inference seem odd to 'radical' embodied theories, it is possible that within this scheme one can understand how representational abilities emerge that are relevant for interactive behavior [49,80].

G. Pezzulo and P. Cisek. Navigating the affordance landscape: feedback control as a process model of behaviour and cognition. Trends in Cognitive Sciences, 20(6):414–424, 2016.

Increasing
Level of
Abstraction
(modal & amodal)

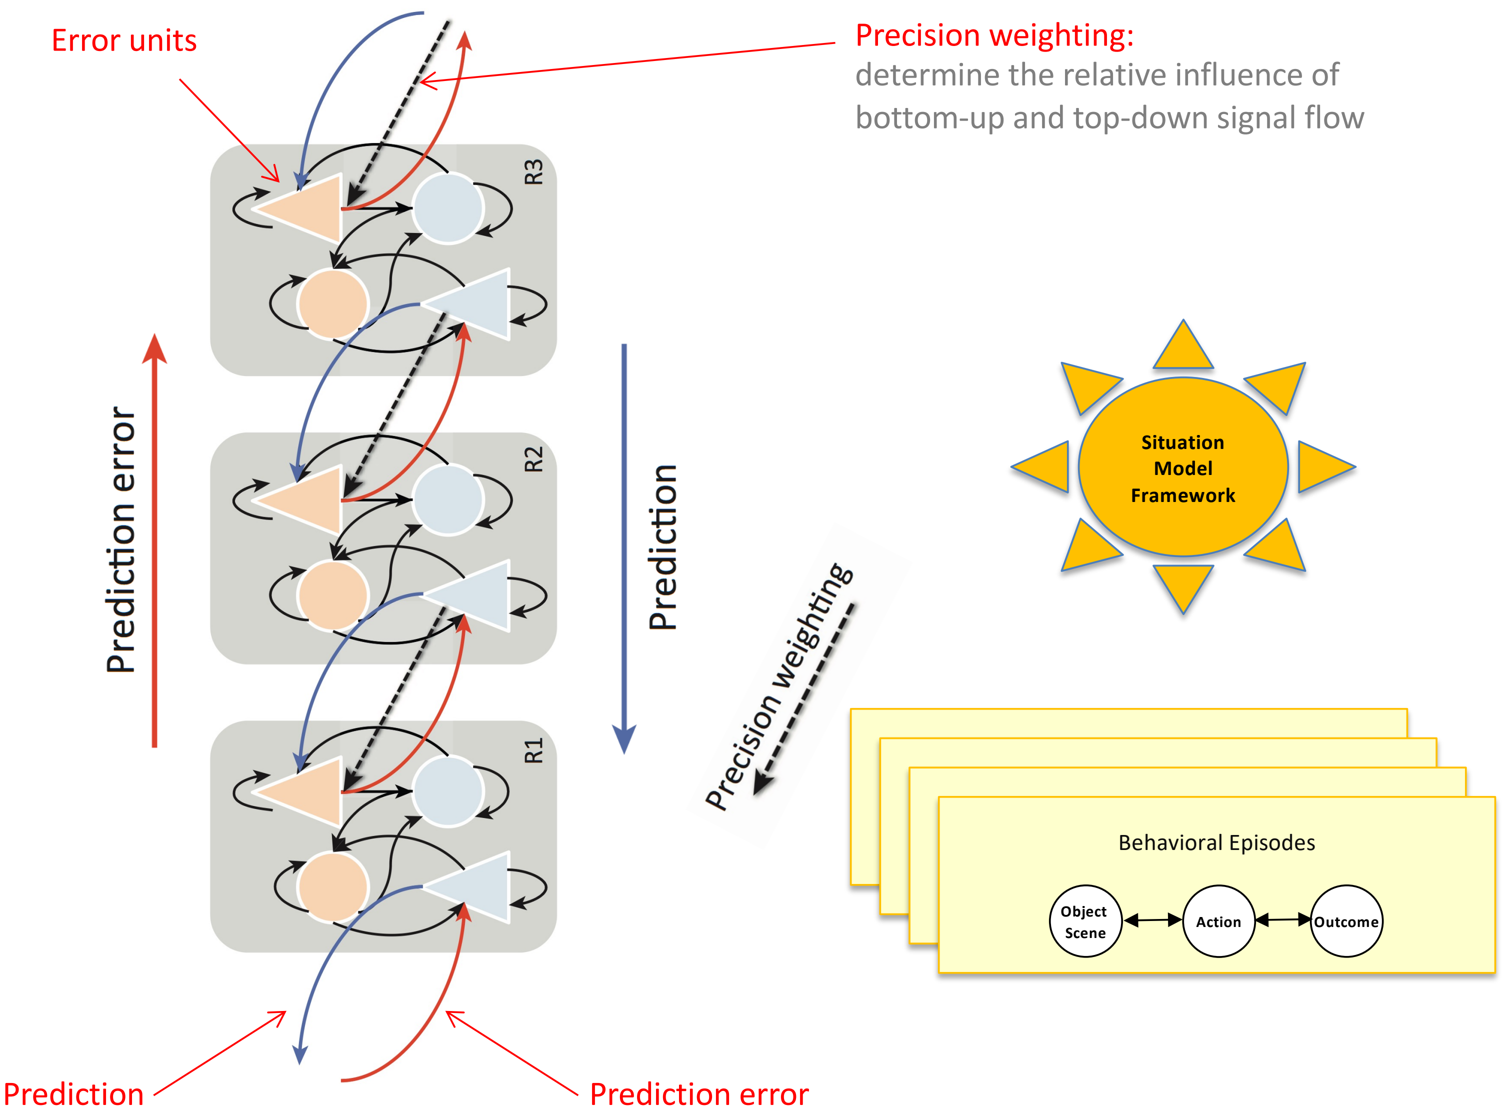

A. K. Seth. Interoceptive inference, emotion, and the embodied self. Trends in Cognitive Sciences, 17(11):565–573, November 2013.

Increasing
Level of
Abstraction
(modal & amodal)

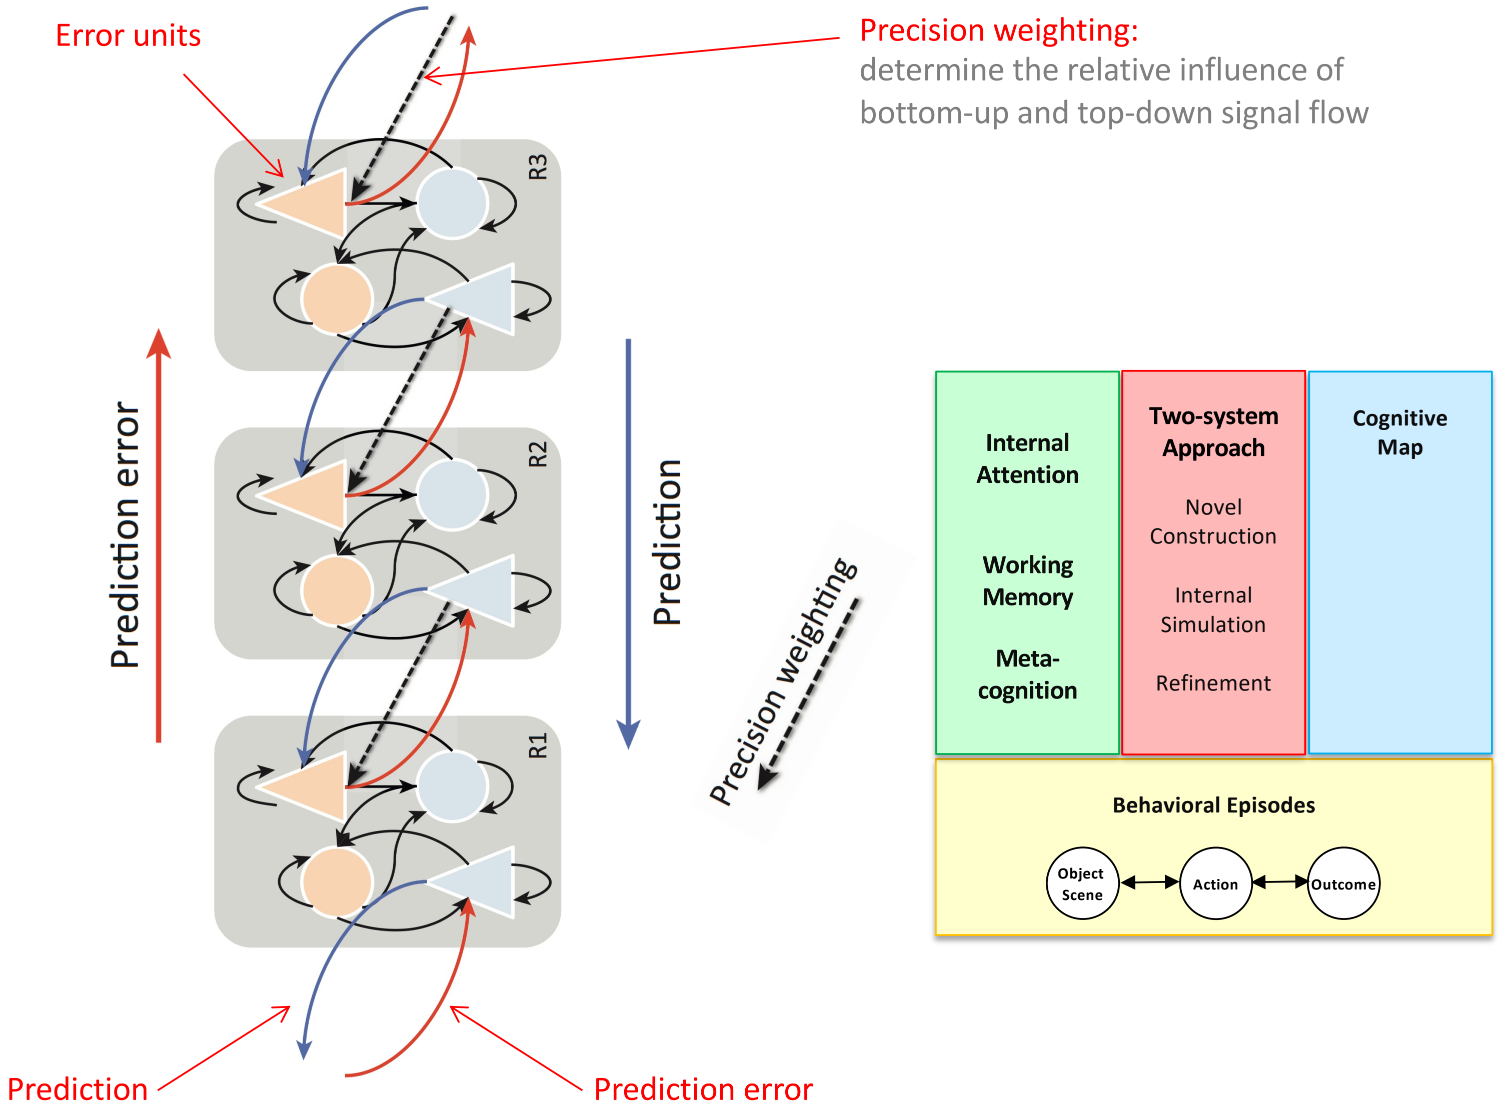

A. K. Seth. Interoceptive inference, emotion, and the embodied self. Trends in Cognitive Sciences, 17(11):565–573, November 2013.

Caveat

"Many open questions remain.

A key challenge is to detail the underlying neural operations"

Other Questions …

- What process selects between counterfactuals?
- What process governs the diversity of the repertoire of counterfactuals?
- What process selects between predictive processing and active inference?
- What process selects between different forms of active inference?
- How is the hierarchy structured and organized?
- What process drives allostasis?
- How is goal-directed behaviour to be addressed?
- What are the drives and motivations?

Jour Fixe

28 July 2020

The Cybernetic Bayesian Brain:

From Interoceptive Inference to Sensorimotor Contingencies

A. Seth,

in T. Metzinger & J. M. Windt (Eds). Open MIND: 35(T). Frankfurt am Main: MIND Group,1–24, 2015.
doi: 10.15502/9783958570108

Cognitive Behavior of Humans, Animals, and Machines: Situation Model Perspectives

ZiF Zentrum für interdisziplinäre Forschung
Center for Interdisciplinary Research