$$\underbrace{F(s,\mu)}_{\text{free-energy}} = \underbrace{E_q[-\log p(s,\psi \mid m)]}_{\text{energy}} - \underbrace{H[q(\psi \mid \mu)]}_{\text{entropy}} = \underbrace{-\log p(s \mid m)}_{\text{surprise}} + \underbrace{D_{\mathrm{KL}}[q(\psi \mid \mu) \parallel p(\psi \mid s,m)]}_{\text{divergence}} \geq \underbrace{-\log p(s \mid m)}_{\text{surprise}}$$

Free Energy on the Edge of Chaos

Inês Hipólito

MPML Seminar

08.09.2022

BERLIN SCHOOL OF MIND AND BRAIN
HUMBOLDT-UNIVERSITÄT ZU BERLIN
UCL
E-GLOW
European Institute for Global Well-being

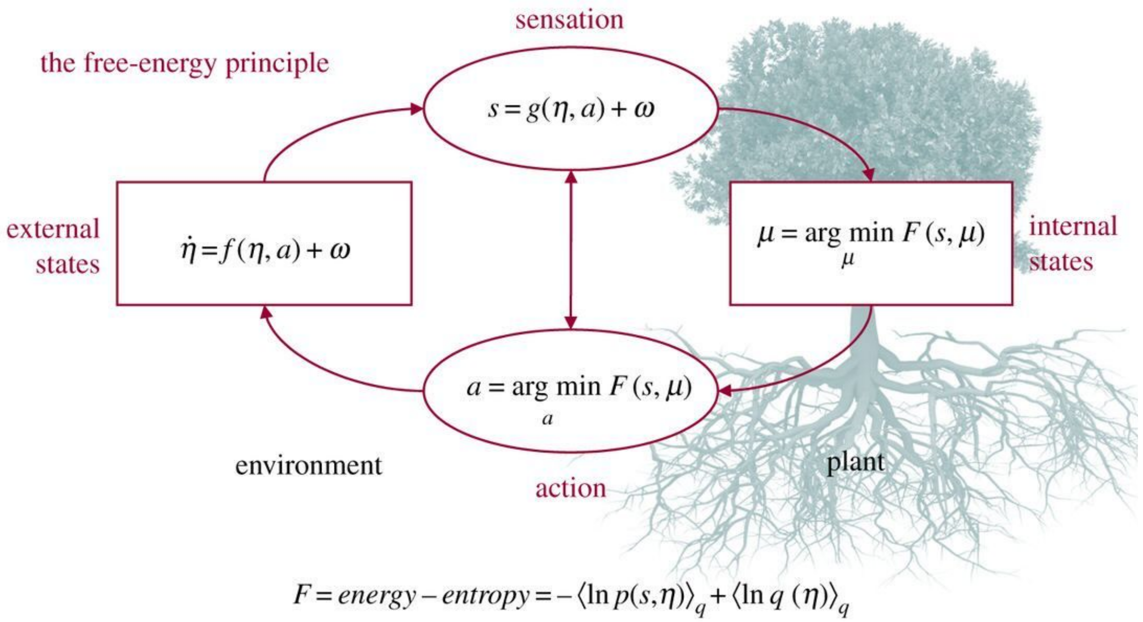

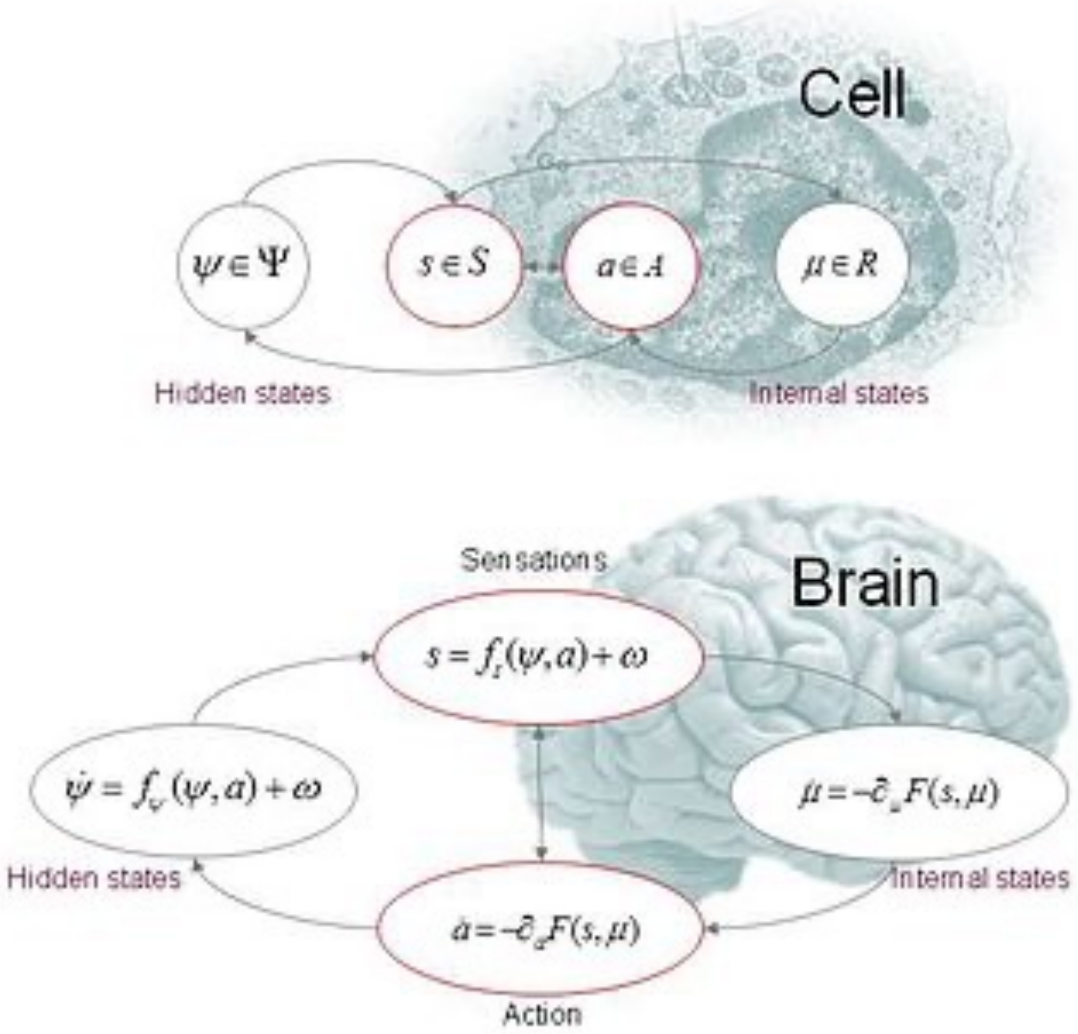

- First theory that **unifies observations made by different fields of research** (biology/physics/neuroscience/philosophy/computer science)

How do living organisms **persist** while engaging in **adaptive exchanges** with their environments?

The second law of thermodynamics

Living beings do an extraordinary thing. If they are alive it means that they are (seemly) **resisting the second law of thermodynamics.**

This law stipulates that **open systems** - just like living beings - **tend to dissipate** by the increase of chaos or **entropy.**

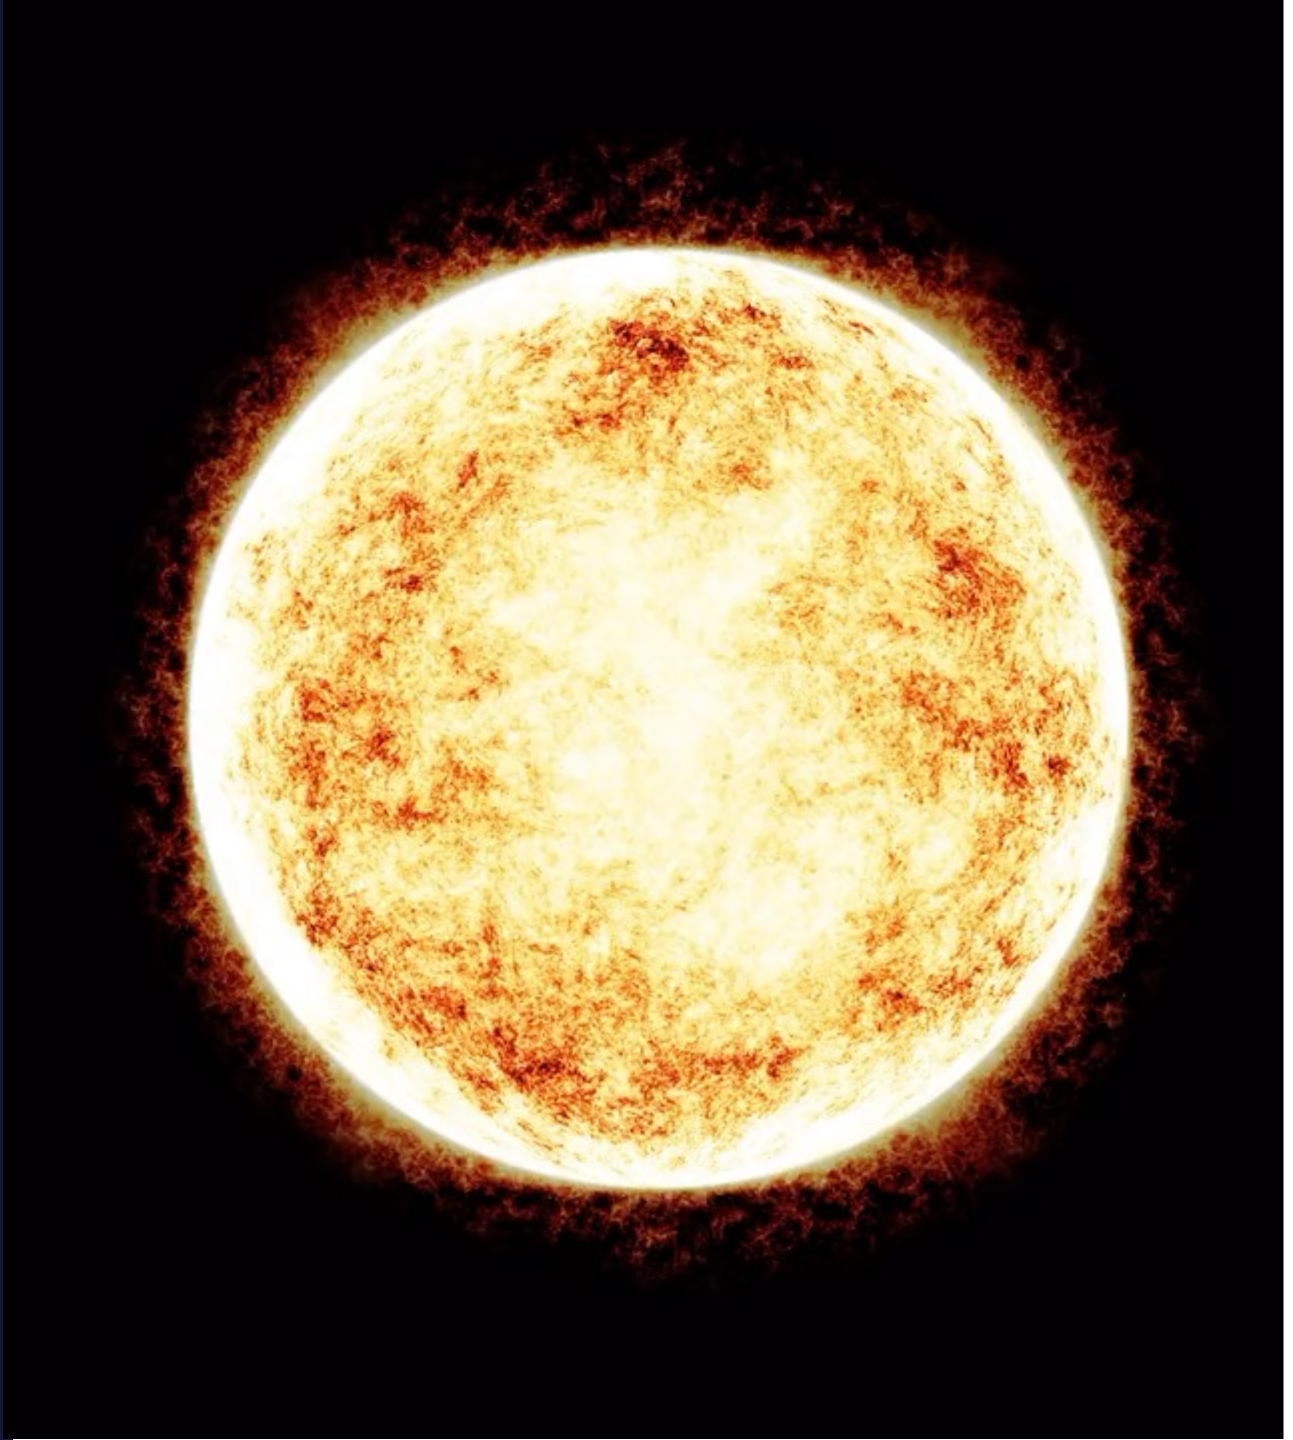

Living beings

Strikingly however, **all living beings**,

from **plants** to more **complex organisms equipped with nervous systems**,

defy and resist the **second law**

How?

Living beings

by **coupling** with
and **adapting** to
their environments

Living beings

Because all living beings defy the second law by adjusting and engaging with the environment,

a prominent question is how do living organisms persist while engaging in adaptive exchanges with their environments?

The high road to the FEP

The Free Energy Principle (FEP) suggests that living systems maintain themselves by remaining in **non-equilibrium steady states** by restricting themselves to a **limited number of states**

The high road to the FEP

it starts from **first principles in statistical physics** and the **central imperative that organisms must maintain their existence**

that is, avoid surprising states—and then introduces the **minimization of free energy** as a computationally tractable **solution to this problem**.

The high road to the FEP

(P1) to survive, any living organism has to maintain itself in a suitable set of **preferred states,** while avoiding other, dis-preferred states of the environment.

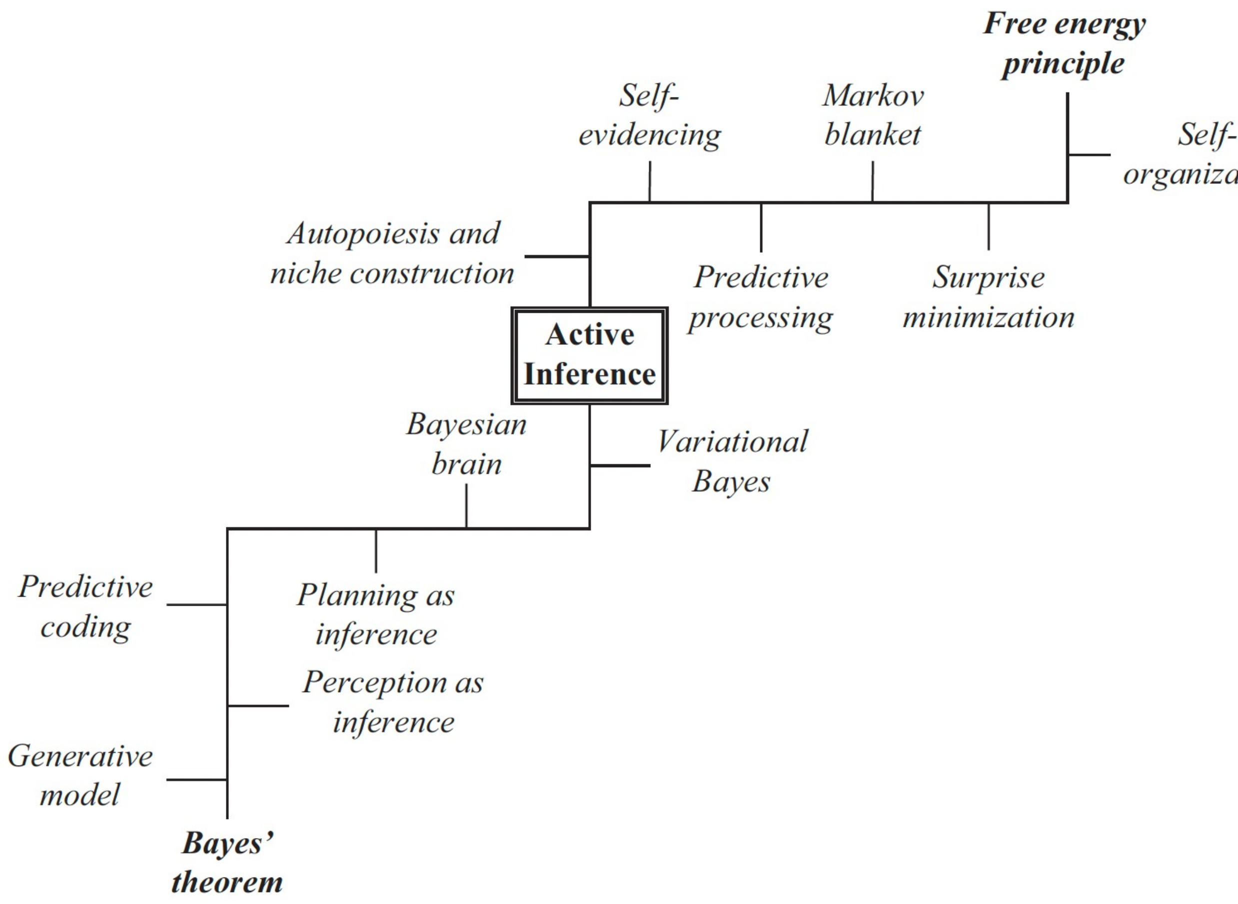

Prefered states

These preferred states are

- defined by **niche-specific evolutionary adaptations**.
- extend to **learned cognitive goals**

- Living organisms resolve this fundamental biological problem
- by exerting active control over their states (e.g., of body temperature) at many levels,
- which range from automatic regulatory mechanisms such as sweating (physiology)
- to cognitive mechanisms such as buying and consuming a drink (psychology)
- to cultural practices such as distributing air conditioning systems (social sciences).

The FEP

organisms follow a unique imperative:

minimizing the surprise of their sensory observations.

Active Inference

Surprise has to be interpreted in a technical sense: it measures **how much an agent's current sensory observations differ** from its preferred sensory observations—that is, those that preserve its integrity

e.g., for a fish, being in the water.

Minimising surprise

minimizing surprise **is not something that can be done by passively** observing the environment:

rather, agents must **adaptively control** their **action-perception loops** to solicit **desired sensory observations.**

This is the **active** bit of Active Inference.

Surprise

I get this claim from the Alius interview, where Friston says:

> *If you subscribe to the premise that that creatures like you and me act to minimize their expected free energy, then we act to reduce expected surprise or, more simply, resolve uncertainty. So what's the first thing that we would do on entering a dark room — we would turn on the lights. Why? Because this action has epistemic affordance; in other words, it resolves uncertainty (expected free energy). This simple argument generalizes to our inferences about (hidden or latent) states of the world — and the contingencies that underwrite those states of affairs.*

Minimising surprise

Active Inference offers a solution to this problem.

It assumes that even if living organisms cannot directly minimize their surprise, they can minimize a proxy–called **(variational) free energy.**

Minimising surprise

This quantity can be minimized through **neural computation** in response to (and in anticipation of) sensory observations.

This emphasis on **free energy minimization** discloses the relation between **Active Inference** and the (first) principle that motivates it: **the free energy principle** (Friston 2009).

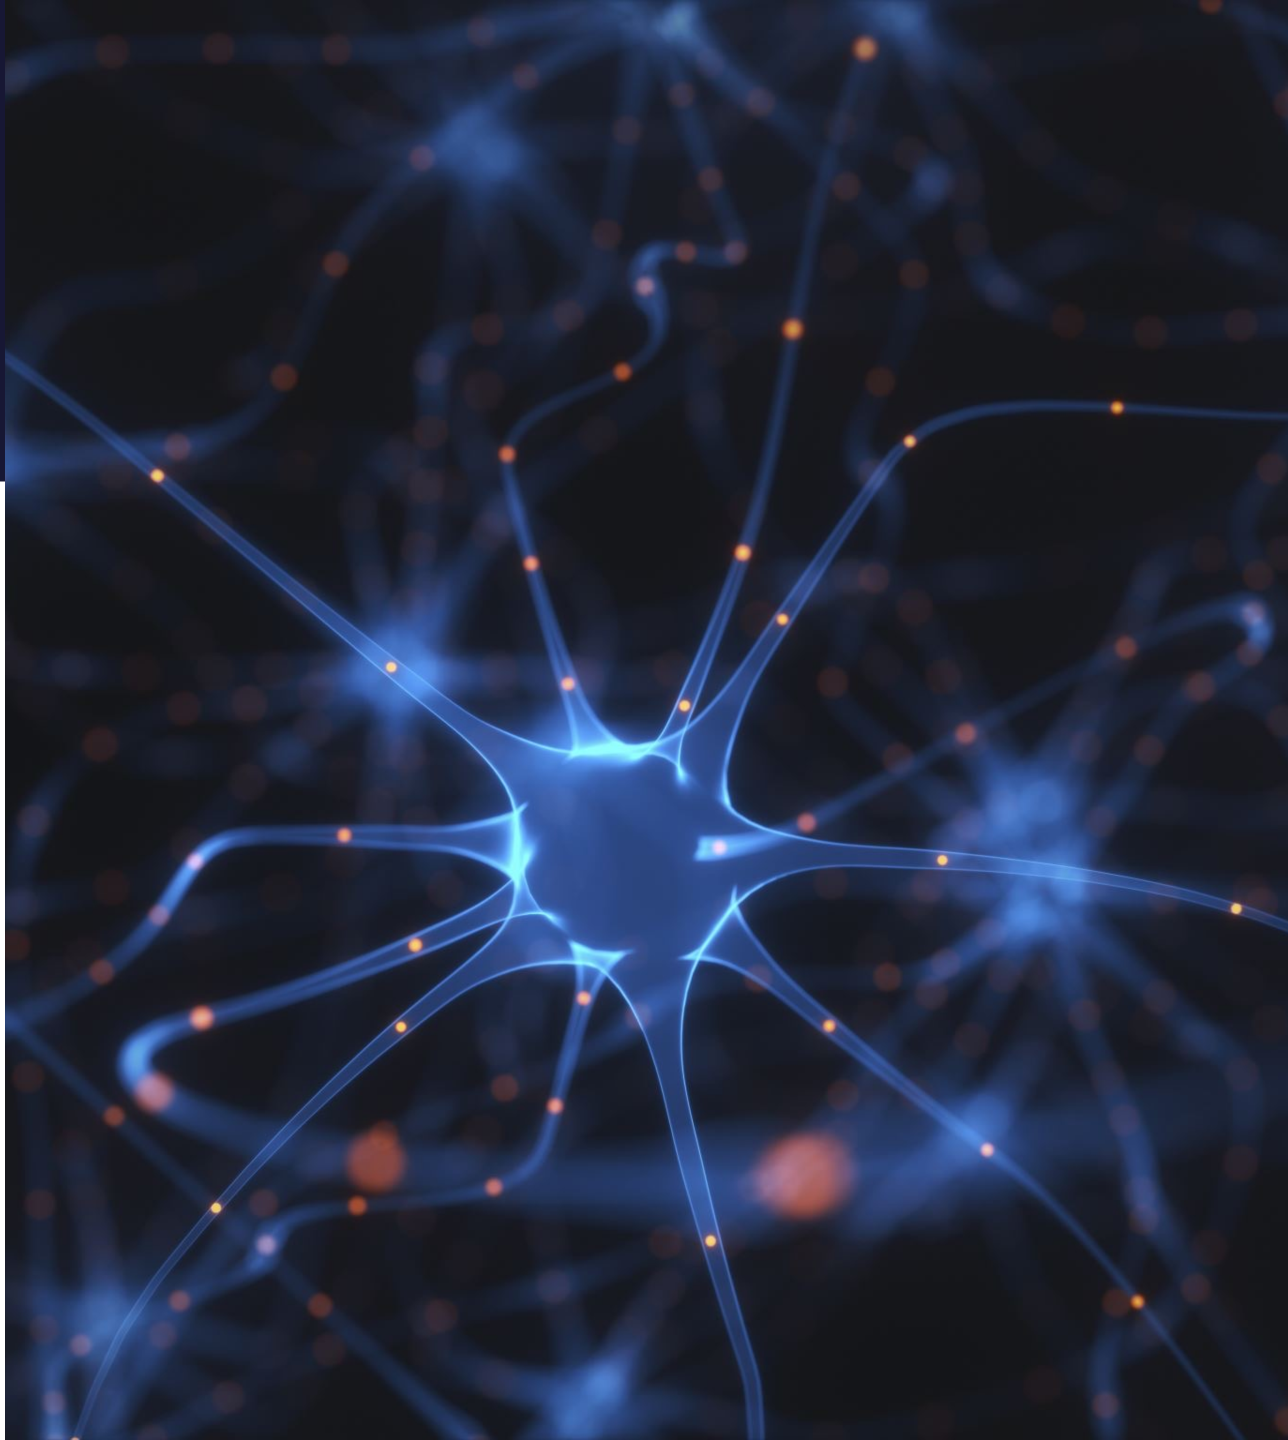

Free energy minimization

Free energy minimization seems a very abstract starting point to explain biological phenomena. However, it is possible to derive a number of formal and empirical implications:

- how the variables involved in free energy minimization may be encoded in **neuronal populations**;
- how the computations of minimized free energy map to specific **cognitive processes**, such as perception, action selection, and learning;
- what kind of **behaviors** emerge when an Active Inference agent minimizes its free energy

Active inference

Active Inference and **free energy minimization** at the level of

- living organisms–simpler (e.g., bacterial) or more complex (e.g., human)– and their **behavioral, cognitive, social, and neural** processes.

- Varied **biological phenomena** and timescales from evolutionary to **cellular** and **cultural** (Friston, Levin et al. 2015; Isomura and Friston 2018; Palacios, Razi et al. 2020; Veissi.re et al. 2020)

Active inference

Active Inference is a theory of how **living organisms maintain their existence** by minimizing surprise—or a tractable proxy to surprise, variational Free energy—via **perception and action**

Active inference

01

Active Inference casts the biological Problem of—or explanation for—survival as surprise minimization.

02

This formulation rests on a technical definition of surprising states from information theory—essentially, surprising states index those outside the comfort zone of living organisms.

03

It then proposes free energy minimization as a practical and biologically grounded way for organisms or adaptive systems to minimize the surprise of sensory encounters.

How to delineate the dependencies between an adpative system and its changing environment (FE minimisation)?

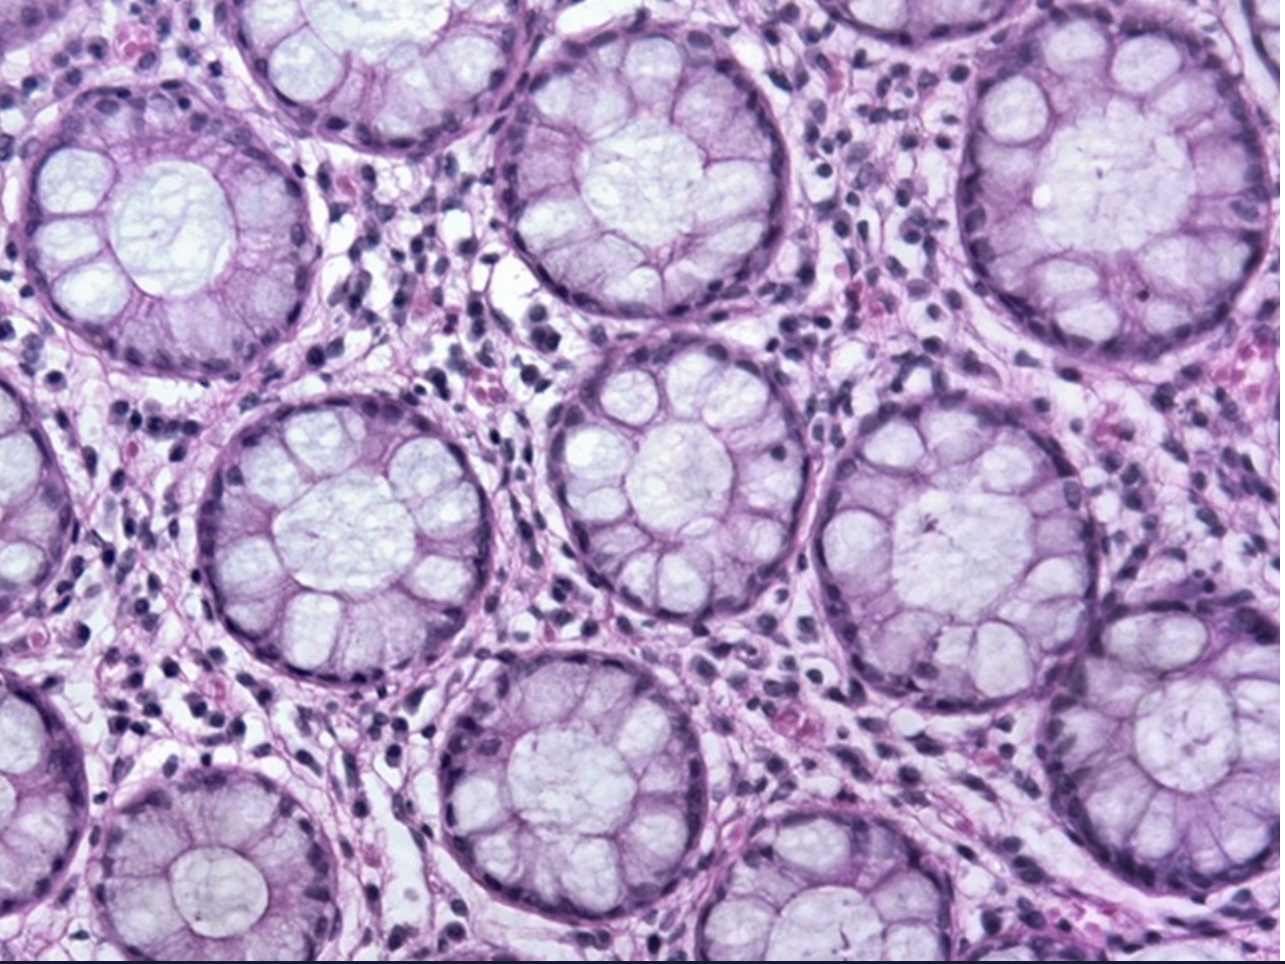

An important precondition for any adaptive system is that it must enjoy some separation and autonomy from the environment

Without which it would simply **dissipate, dissolve, and thereby succumb to environmental dynamics**. In the absence of this separation, there would be no surprise to minimize

Markov blankets

What is a Markov blanket?

Markov blankets

1) Conditions for the present event / state of affairs
2) Prediction of a future state given the present conditions

Markov blanket

Formal way to express a **separation and coupling** between a system and the rest of the environment- statistical construct (Pearl 1988)

$$\mu \perp x | b \Leftrightarrow p(\mu, x | b) = p(\mu | b) p(x | b)$$

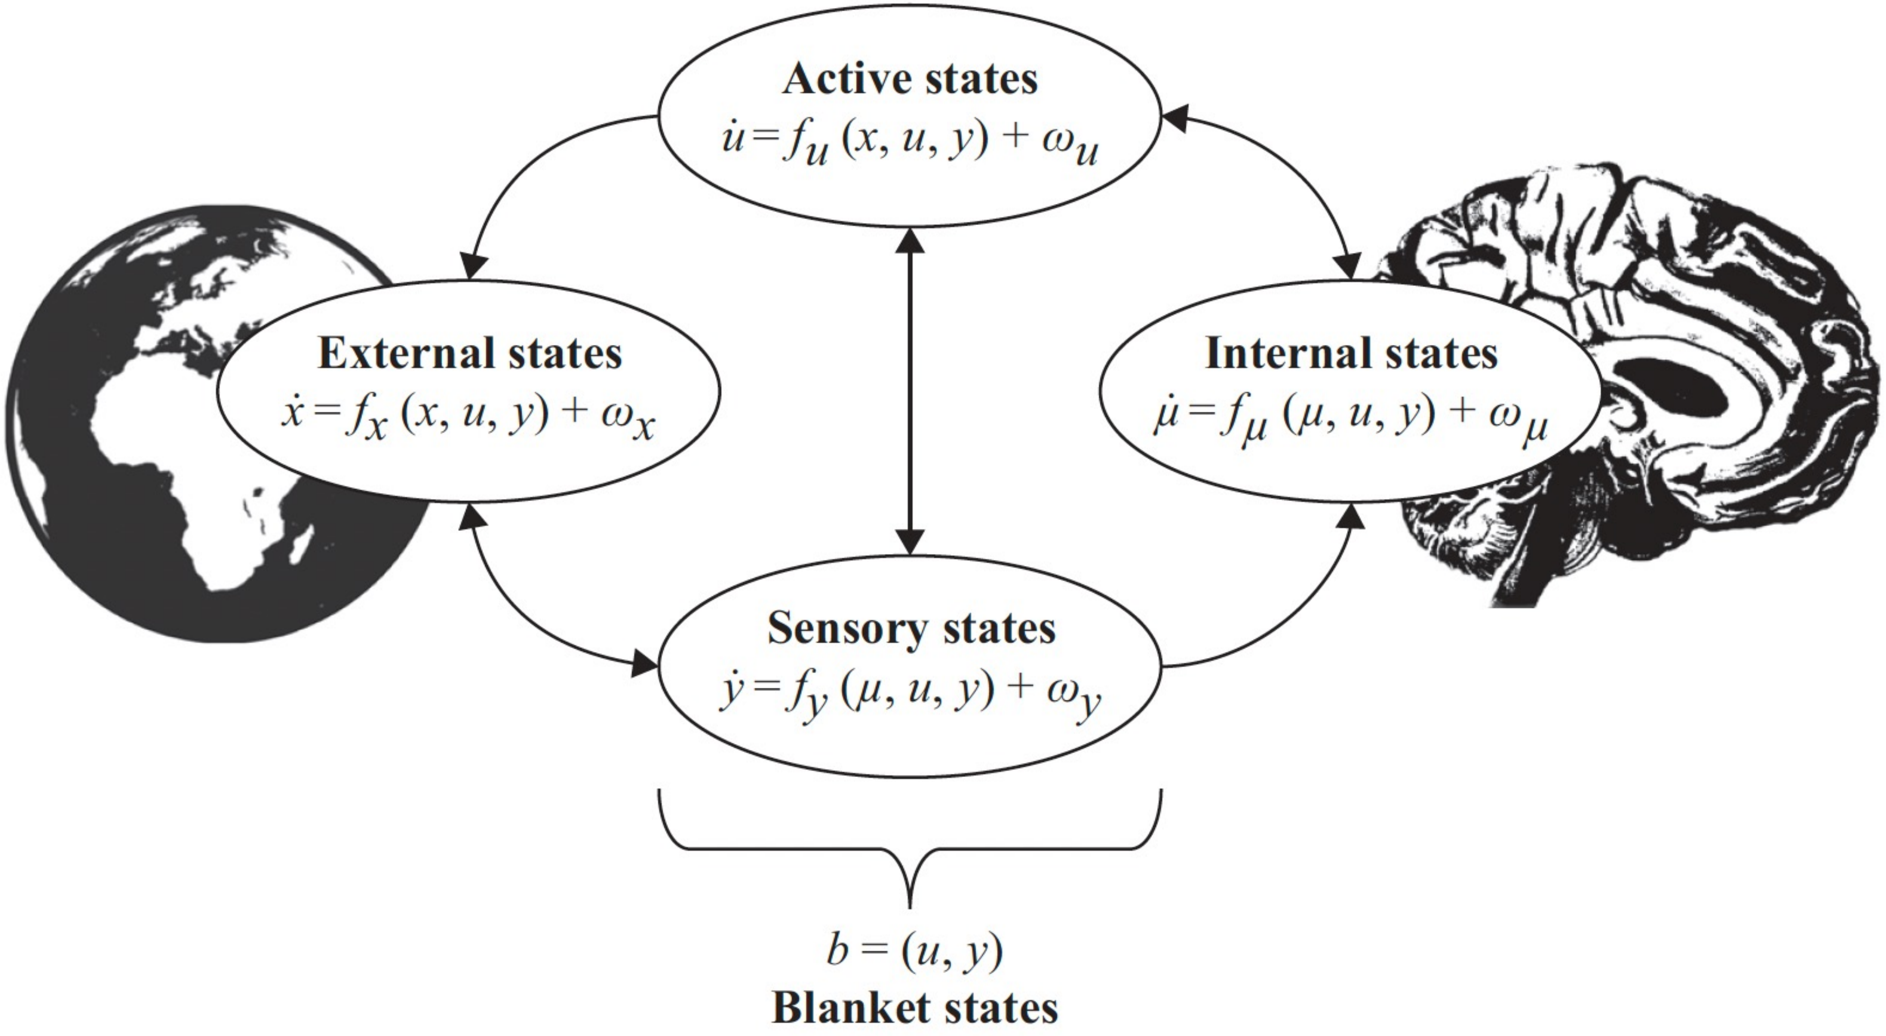

Active states
$\dot{u} = f_u(x, u, y) + \omega_u$
External states
$\dot{x} = f_x(x, u, y) + \omega_x$
Internal states
$\dot{\mu} = f_\mu(\mu, u, y) + \omega_\mu$
Sensory states
$\dot{y} = f_y(\mu, u, y) + \omega_y$
$b = (u, y)$
Blanket states

Markov blanket

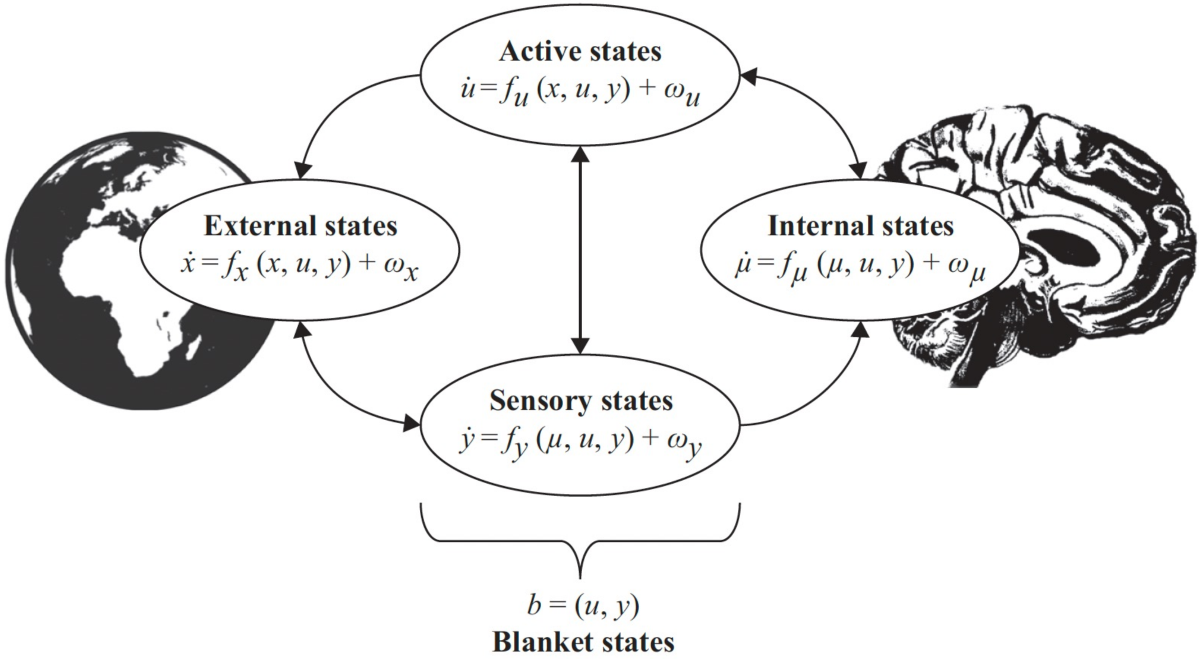

one can use a Markov blanket to separate an entire organism from the environment or **nest multiple Markov blankets** within one another.

Markov blanket

(1) formalizes the fact that an **adaptive system's internal states are autonomous from environmental dynamic**s and can therefore resist their influences.

(2) it scaffolds the way in which adaptive systems minimize their surprise: it highlights the **internal, sensory, and active states** they have access to.

Markov blanket

Internal states of an adaptive system bear a formal relation to **external states.**

This is due to a kind of symmetry across the Markov blanket as both **influence and are influenced by blanket states**.

Surprise Minimization and Self-Evidencing

The **dynamics of internal states** correspond to a form of (approximate) Bayesian inference of external states, as their motion changes the associated probability distribution,

which is afforded by an **implicit generative model** of how sensations (or sensory states in the Markov blanket jargon) are generated.

If we reinstate the notion of an agent as constituted by internal and blanket states, we can talk about an **agent's generative model**.

Agent’s generative model

the **agent’s generative model** cannot simply mimic external dynamics (otherwise the agent would simply follow external dissipative dynamics).

Rather, the model must also **specify the preferred conditions for the agent’s existence**, or the regions of states that the agent has to visit to maintain its existence, or satisfy the criteria for its existence in terms of occupying characteristic states.

Priors of the model

These **preferred states (or observations)** can be specified as the **priors of the model**–which implies that the model implicitly assumes that its preferred (prior) sensations are more likely to occur (i.e., are less surprising) if it satisfies the criteria for existence.

optimism bias is necessary for the agent to go beyond the mere duplication of external dynamics to prescribe active states that underwrite its preferred or characteristic states.

Priors of the model

optimal behavior (with respect to prior preferences) as the maximization of model evidence by perception and action.

model evidence summarizes how well the generative model fits or explains sensations.

Relations between Inference, Cognition, and Stochastic Dynamics

The physicist E. T. Jaynes famously argued **that inference, information theory, and statistical physics** are different perspectives on the same thing (Jaynes 1957).

Relations between Inference, Cognition, and Stochastic Dynamics

Bayesian and statistical physics perspectives offer two equivalent ways to understand **surprise minimization and optimal behavior**:

interpretations

Statistical physics	Bayesian inference and information theory	Cognitive interpretation
Minimize variational free energy	Maximize model evidence (or marginal likelihood); minimize surprisal (or self-information)	Perception and action
Minimize expected free energy; Hamiltonian principle of least Action	Infer the most likely (or less surprising) course of action	Planning as inference
Attain nonequilibrium steady-state	Perform approximate Bayesian inference	Self-evidencing
Gradient flows on energy functions; gradient	Gradient ascent on model evidence; gradient	Neuronal dynamics

Notes

Free energy is widely used in statistical physics to characterize (for example) thermodynamic systems;

Active Inference uses exactly the same equations, it applies them to characterize the belief state of an agent (in relation to a generative model); referring to processes that change its belief state, not (for example) the particles of its body

Variational Free Energy, Model Evidence, and Surprise

maximization of model evidence

(or marginal likelihood) in Bayesian inference

the minimization of variational free energy

Both minimise surprise

Active Inference: A Novel Foundation to Understand Behavior and Cognition

Active Inference: A Novel Foundation to Understand Behavior and Cognition

Behavior is the result of inference and its optimization is a function of beliefs. This formulation unites notions of (prior) belief and preference

agent's preference for a course of action becomes simply a belief about what it expects to do, and to encounter, in the future– or a belief about future trajectories of states that it will visit.

Advantages

self-consistent process model of purposive (or teleological) behavior, which is akin to cybernetic formulations.

#behavior as a functional of beliefs (probability distributions) automatically entails notions such as degree of belief and uncertainty.

#optimal behavior comes to follow a Hamiltonian principle of least Action

becomes an energy function—and the most likely course of action of an Active Inference agent is the one that minimizes free energy.

Consequences

living organisms behave according to Hamilton’s principle of least Action:

they follow a path of least resistance until they reach a steady state (or a trajectory of states),

Enactivism

self-organization of behavior and **autopoietic interactions** with the environment, which ensure that living organisms remain within acceptable bounds (Maturana and Varela 1980).

Via formal framework explaining how living organisms manage to resist the dispersion of their states by self-organizing a statistical structure—the Markov blanket

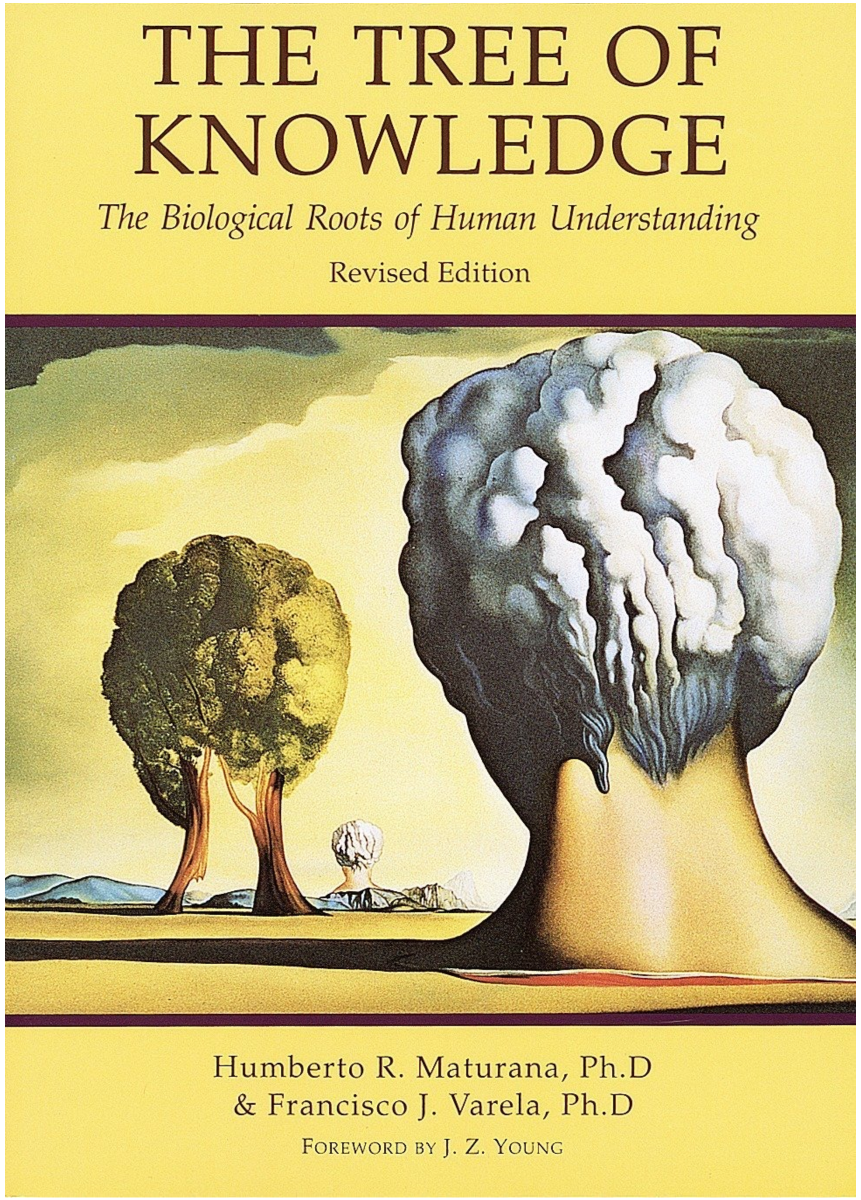

Cognition & Behaviour under FEP

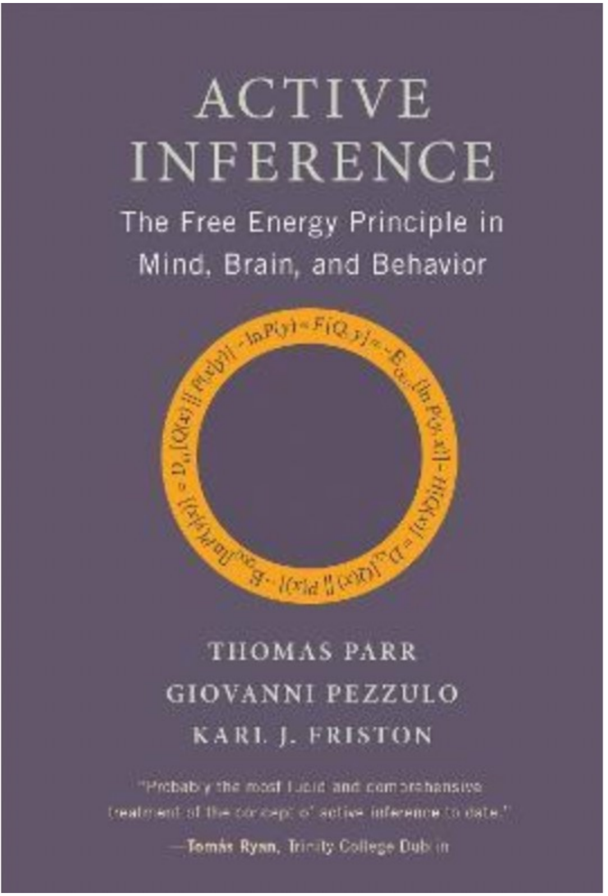

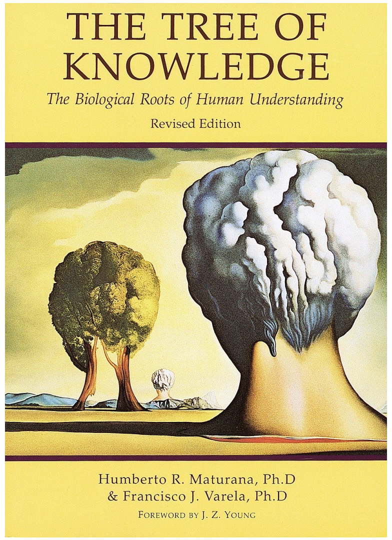

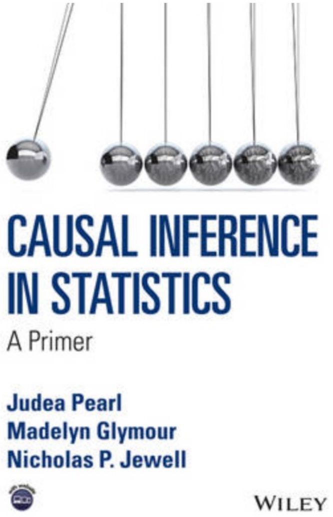

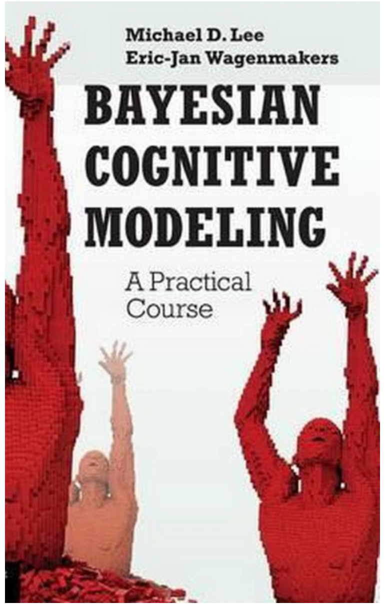

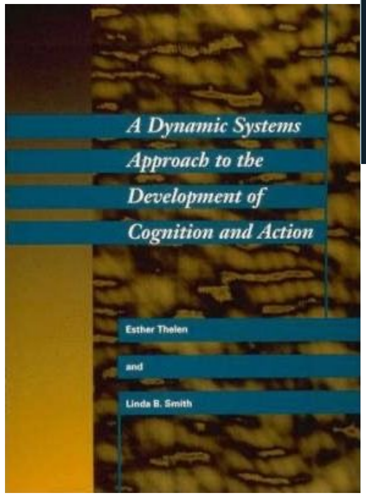

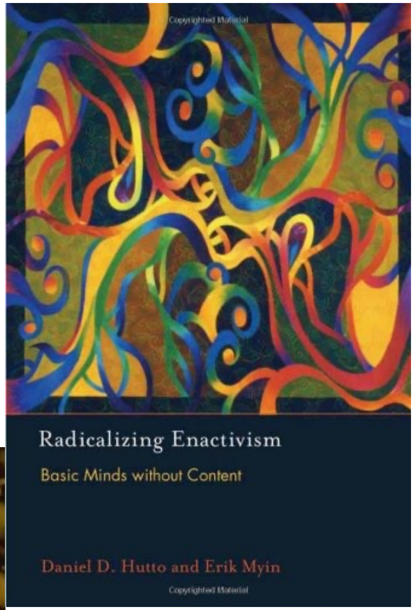

Thank you

Ines.hipolito@hu-berlin.de

www.ineshipolito.com

BERLIN SCHOOL OF MIND AND BRAIN

HUMBOLDT-UNIVERSITÄT ZU BERLIN

UCL

E-GLOW

European Institute for Global Well-being