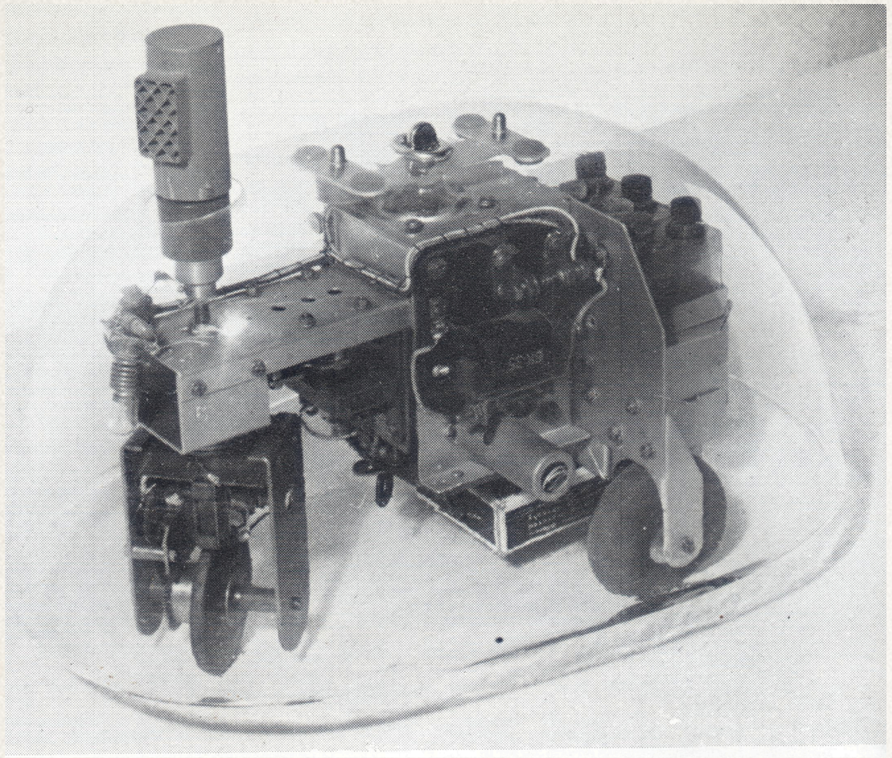

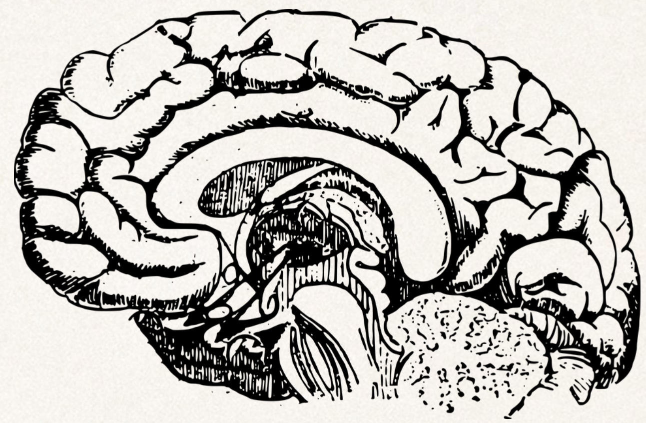

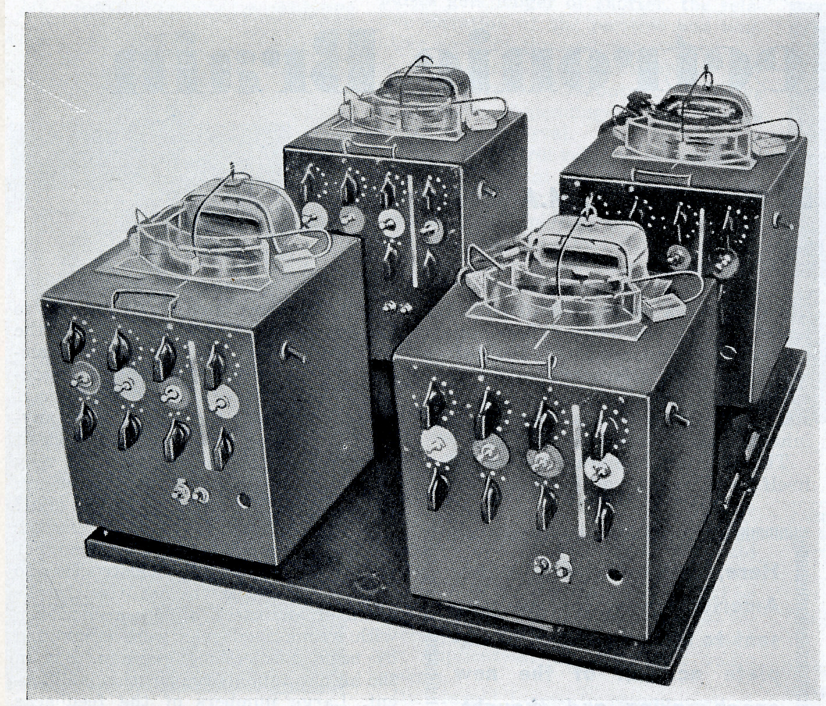

Fig. I—The homeostat, with its four units, each one of which reacts on all the others.

Variational inference in agents, with connections to control theory and cognitive (neuro)science

Manuel Baltieri

- Lab for Neural Computation and Adaptation, RIKEN CBS, Wako-shi (Japan)

- EASY - CCNR - Sussex Neuroscience, University of Sussex, Brighton (UK)

RIKEN CBS
Center for Brain Science

Twitter: @manuelbaltieri

19/07/2021 - Tokyo (online)

Roadmap

- World models? Reconstructing vs controlling
  - Generative models of behaviour
- A Bayesian angle on classical control
  - PID controllers and their design process
- Variational inference in cognitive (neuro)science
  - Duality of inference (perception) and control (action) and dual effects of control (action)
- Current directions

Initial motivation: understand if what Friston proposes in neuroscience makes sense

Actual motivation: understand if variational updates in belief space can describe life and cognition at their core

But first… who am I?

- **BEng** - Computer and software engineering, business administration
- **MSc** - Cybernetics, evolutionary computation, computational modelling (neuroscience, biology, behaviour), artificial life
- **PhD** - Theoretical neuroscience, cognitive science, motor control / control theory / cybernetics, stochastic processes and filtering, artificial life
- (Mini) **Postdoc** - Bayesian neural networks, robotics + uncertainty modelling in psychophysics
- **Postdoc** (now) - Theoretical neuroscience (motor control and behaviour), filtering, (some) category theory, (some) non-equilibrium physics

My interests

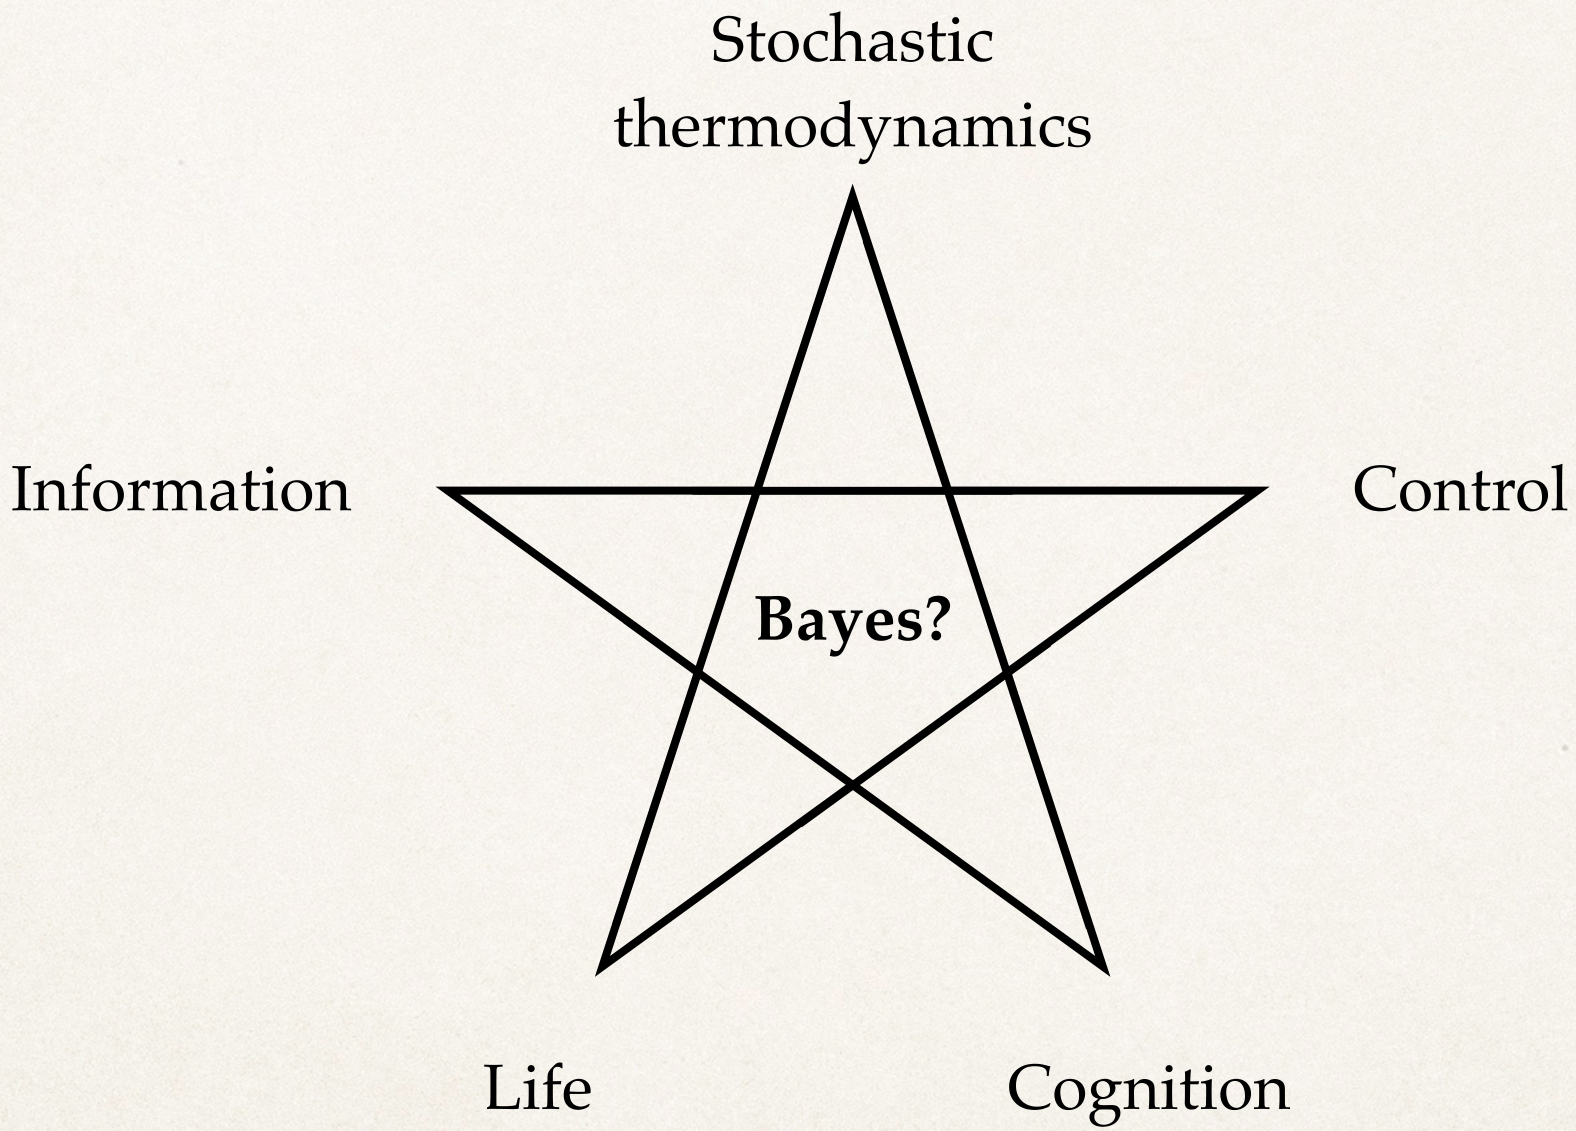

Disclaimer

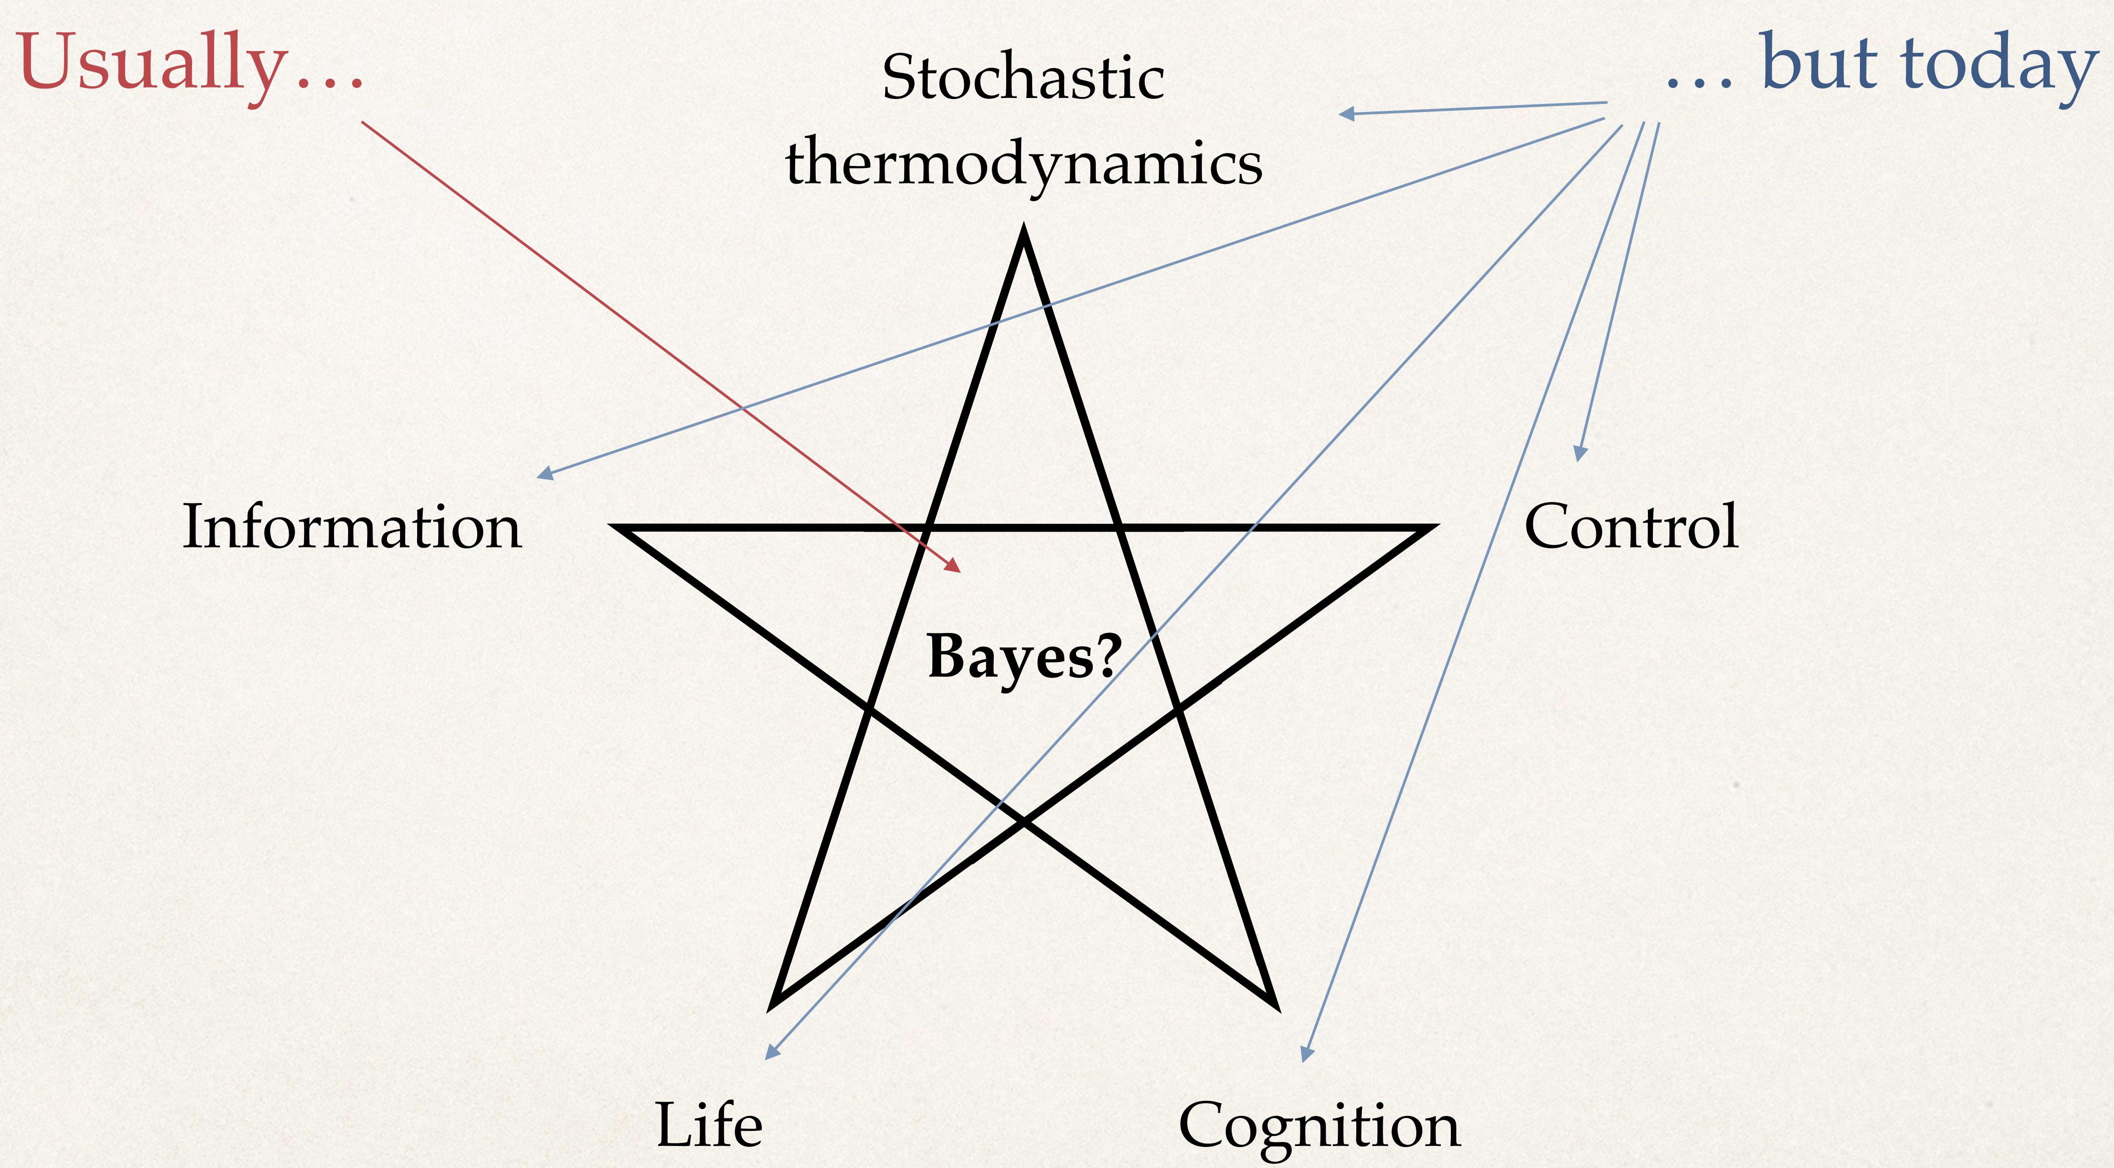
Usually…
Stochastic thermodynamics
… but today
Information
Control
Bayes?
Life
Cognition

Part 1.

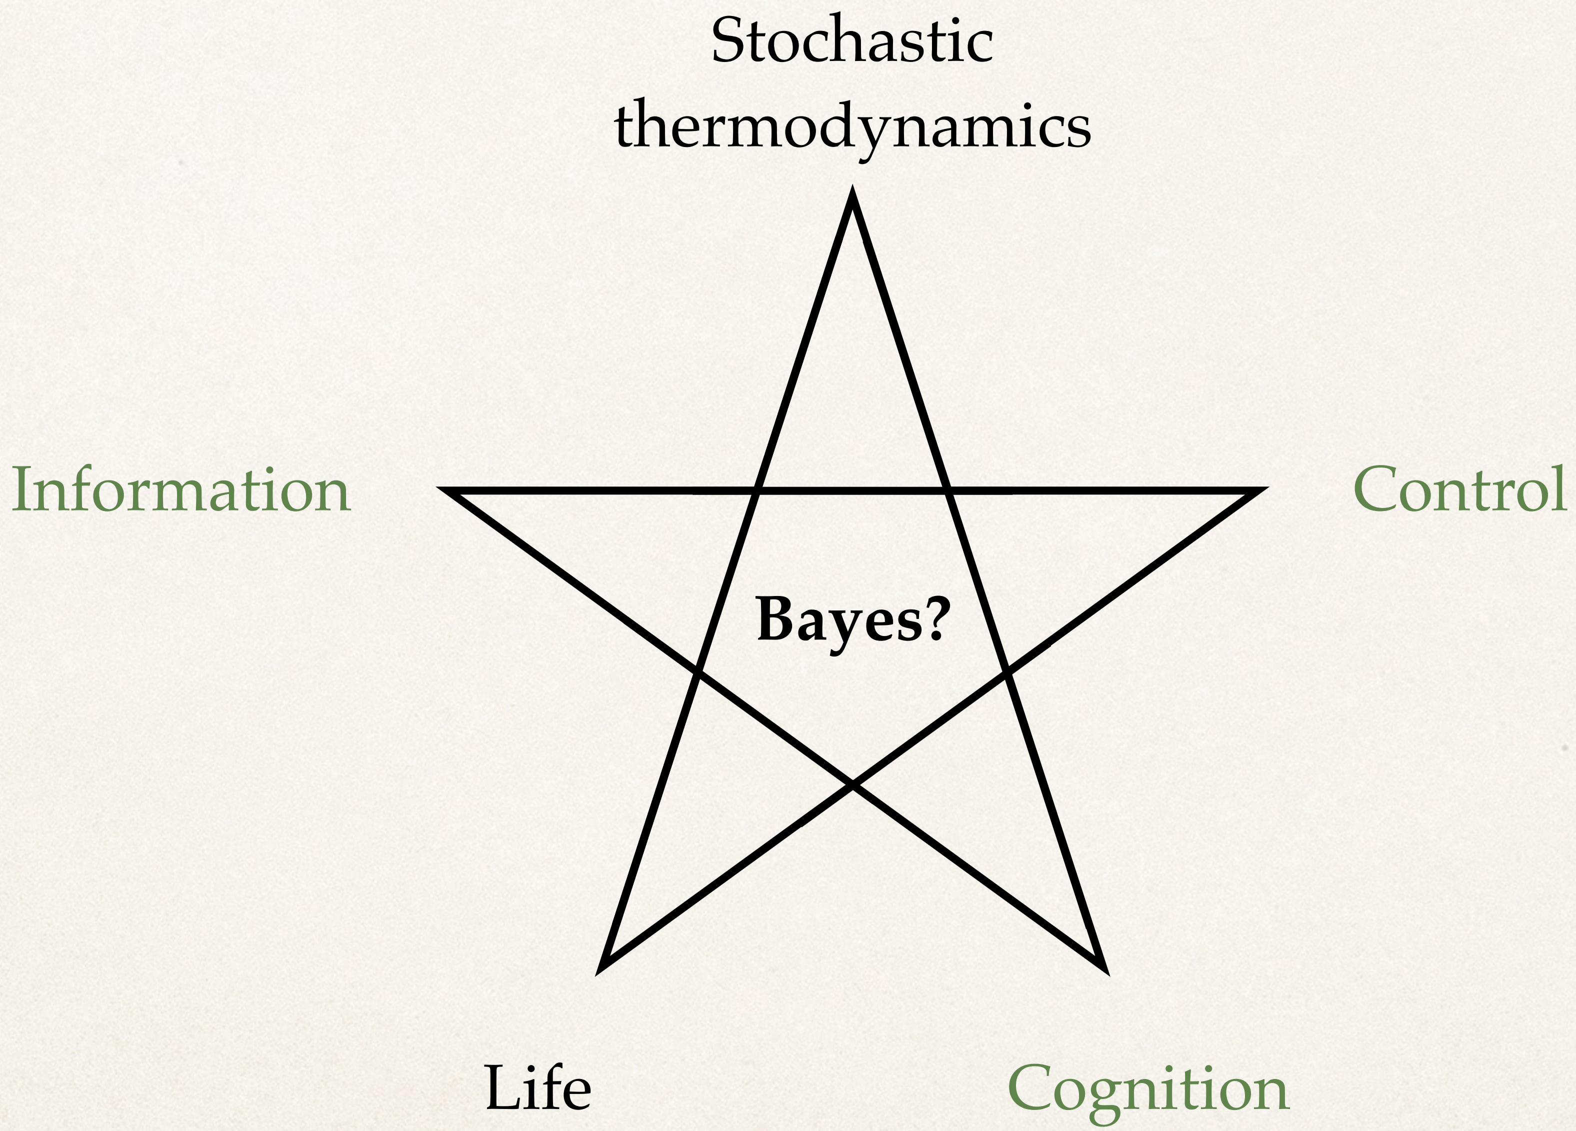
Stochastic
thermodynamics
Information
Control
Bayes?
Life
Cognition

Background - Claim 1

Perception can be described as a process of (Bayesian) inference or estimation

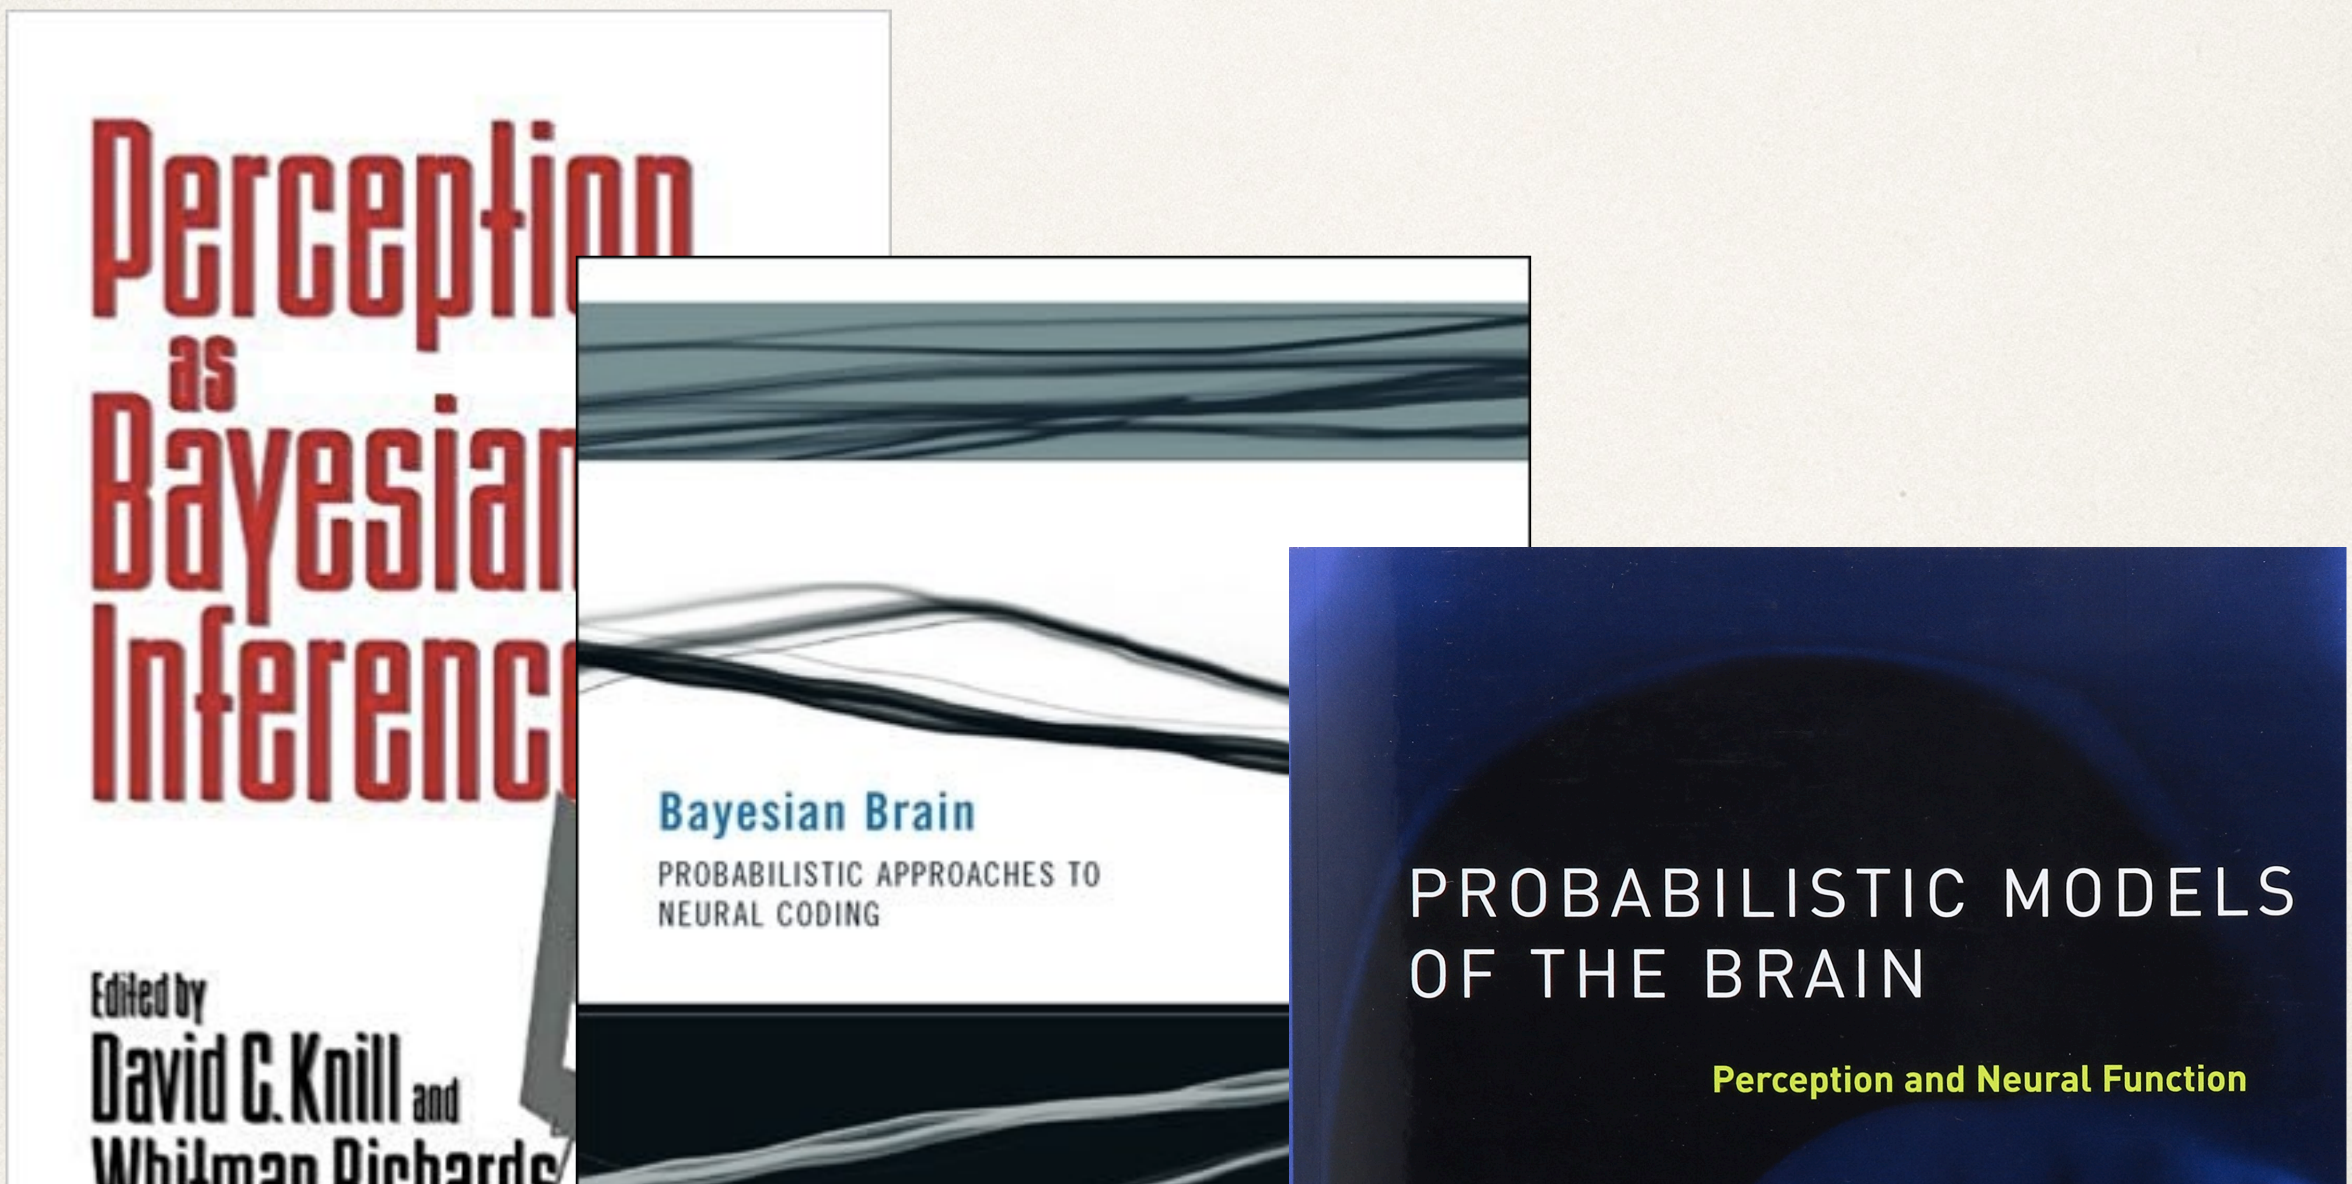

Background - Claim 2

Action can be described as a process of (optimal) control

Internal models for motor control and trajectory planning
Mitsuo Kawato

A number of internal model concepts are nc
neuroscience and cognitive science. These
supported by behavioral, neurophysiologica
furthermore, these models have had their st
functions revealed by such data. In particula
on inverse dynamics model learning is direc
unit recordings from cerebellar Purkinje cell
forward inverse models describing how dive
environments can be controlled and learnec
recently been proposed. The 'minimum varia
another major recent advance in the compu
motor control. This model integrates two fu
approaches on trajectory planning, strongly
both kinematic and dynamic internal models
movement planning and control.

Addresses
ATR Human Information Processing Research La

© 2000 Nature America Inc. • http://neurosci.nature.com

review

Computational principles of movement neuroscience

Daniel M. Wolpert[1] and Zoubin Ghahra

[1] Sobell Department of Neurophysiology, Institute of Neurolog
[2] Gatsby Computational Neuroscience Unit, Queen Square, U
Correspondence should be addressed to D.M.W. (wolpert@he

**Unifying principles of movement have em
review several of these principles and show
control, estimation, prediction and learnin
ing from the computational approach pro**

The computational study of motor control is fund
cerned with the relationship between sensory sig
commands. The transformation from motor com
sensory consequences is governed by the physics
ment, the musculoskeletal system and sensory

://neurosci.nature.com

REVIEW

nature neuroscience

/natureneuroscience

Optimality principles in sensorimotor control

Emanuel Todorov

Background (maths)

State-space models (SSM) formulation

$x' = f(x, v, \theta) + w$: dynamics

$y = g(x, v, \theta) + z$: measurements

$w \sim N(0, \pi_w = h(\lambda))$: fluctuations on dynamics

$z \sim N(0, \pi_z = k(\lambda))$: measurement noise

Probabilistic formulation

$p(y, x, v, \theta, \lambda) = p(y \mid x, v, \theta, \lambda) p(x' \mid x, v, \theta, \lambda)$

(gen. model) (measurements) (dynamics)

Variational distribution

$q(x, \theta)$

Background (maths)

Active inference in continuous space and time (Friston's framework, and what I used in Part 1.):

- fixed-form Gaussian variational inference (+ hierarchical models, here not used)
- separation of timescales for hidden states/inputs (fast) and parameters/hyperparameters (slow, fixed), via explicit mean-field or other assumptions
- fast variables updated via free energy, slow variables via path integral of free energy (i.e. free energy of trajectories, see Archambeau and Opper (2008), but in practice approximated locally)
- <u>actions unknown to agents and treated as hidden inputs</u> (although some clever tricks are implemented to calculate dF/da)

Variational updates

y	: observations	= action, assuming that $y = y(a)$
x	: (hidden) states	= perception/estimation/inference
v	: (hidden) inputs	= perception/estimation/inference
θ	: (hidden) parameters	= learning
λ	: (hidden) hyperparameters	= attention

Archambeau, Cédric, et al. "Variational inference for diffusion processes." (2008): 17-24.

Limitations

- Stationary (time-independent) policies, but wait for the end of the talk
- No learning of SSM parameters (but see Tschanz et al. 2020)
- Fixed-form **Gaussian** VI

Time-independent vs. time-dependent policies

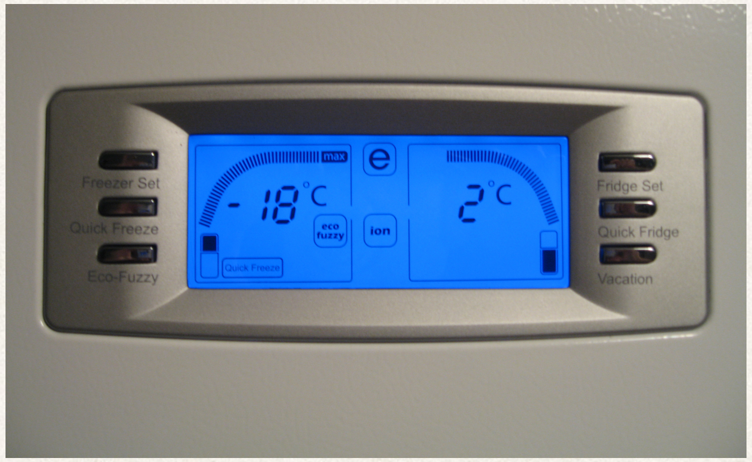

vs.

https://www.freeimages.com/photo/fridge-1325918

https://unsplash.com/photos/3GbcPmYXVwQ

Tschantz, Alexander, Anil K. Seth, and Christopher L. Buckley. "Learning action-oriented models through active inference." PLoS computational biology 16.4 (2020): e1007805.

The ‘usual’ generative models

In statistics / ML:

given observations ‘y’ and labels (categories, classes, states, etc.) ‘x’, find the joint distribution that best represents the data.

Discriminative model:
create a decision boundary

$$P(x \mid y)$$

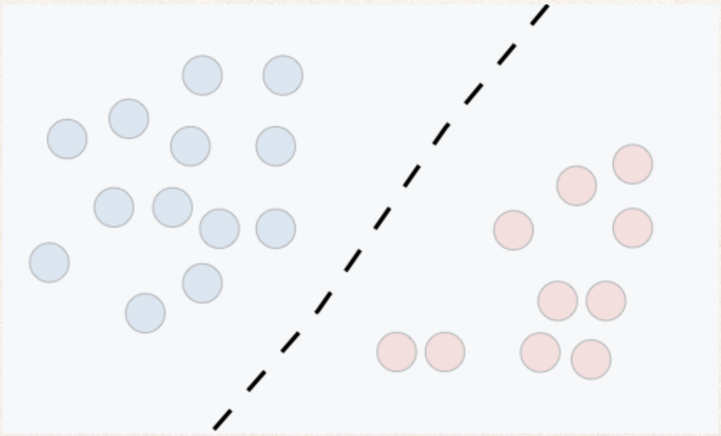

Regression(s), SVMs, etc.

Generative model:
generate a distribution of the data

$$P(y, x)$$

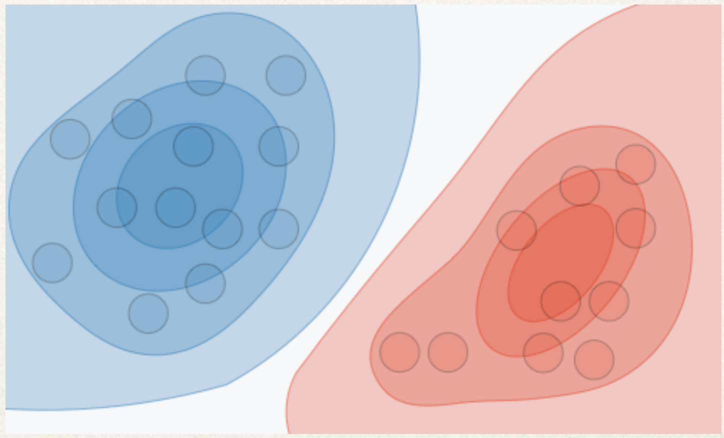

Naive Bayes, HMMs, AR models, etc.

https://stanford.edu/~shervine/teaching/cs-229/cheatsheet-supervised-learning

Example: a generative model in robotics

Goal: (e.g., find a light/phototaxis)

"place a wheeled robot in a random environment, provide it with (at least) light sensors, get it to approach the light source (for simplicity, let's assume there's only one)"

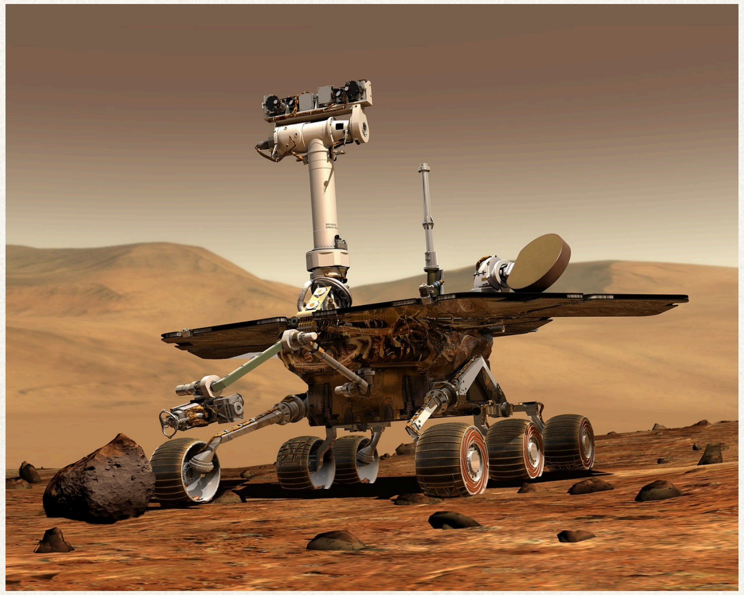

https://pixabay.com/photos/mars-mars-rover-space-travel-rover-67522/

Y - Observations/measurements: light sensors + …

X - States: light's location + commands to reach it + …

Thrun, Sebastian, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. MIT press, 2005.

Standard solution: SLAM

Simultaneous Localisation And Mapping (SLAM)

TL;DR: a robot (iteratively) building an estimate of its pose (position + orientation) on a map while building an estimate of the map itself

Example: World models

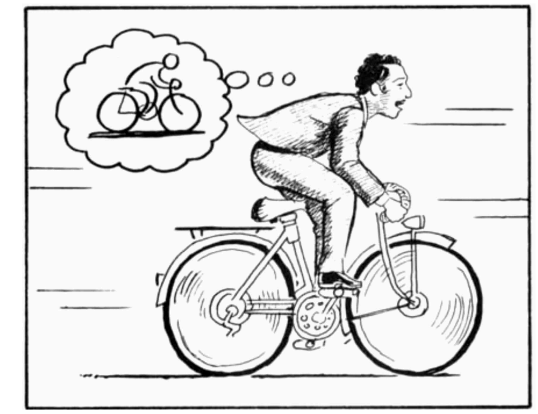

Figure 1. A World Model, from Scott McCloud's *Understanding Comics.* (McCloud, 1993; E, 2012)

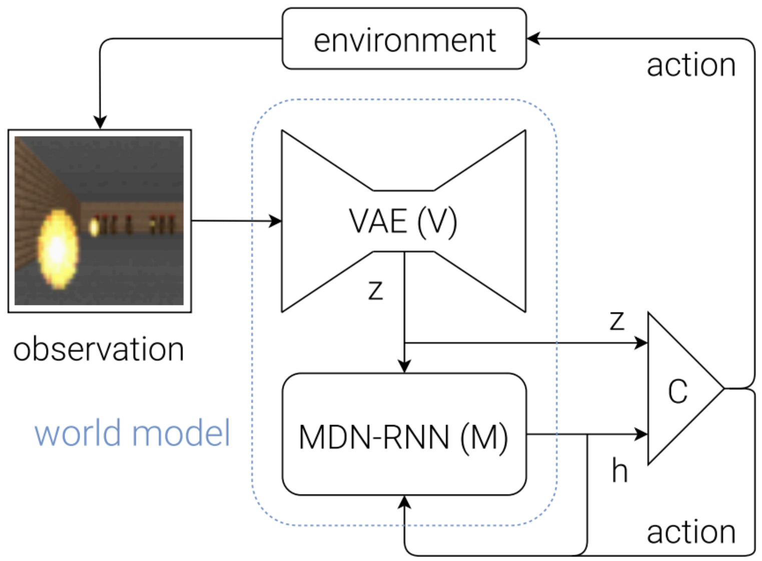

Figure 8. Flow diagram of our Agent model. The raw observation is first processed by V at each time step t to produce z_t. The input into C is this latent vector z_t concatenated with M's hidden state h_t at each time step. C will then output an action vector a_t for motor control, and will affect the environment. M will then take the current z_t and action a_t as an input to update its own hidden state to produce h_{t+1} to be used at time $t + 1$.

World models =>
Replicating a model of
the world inside an agent

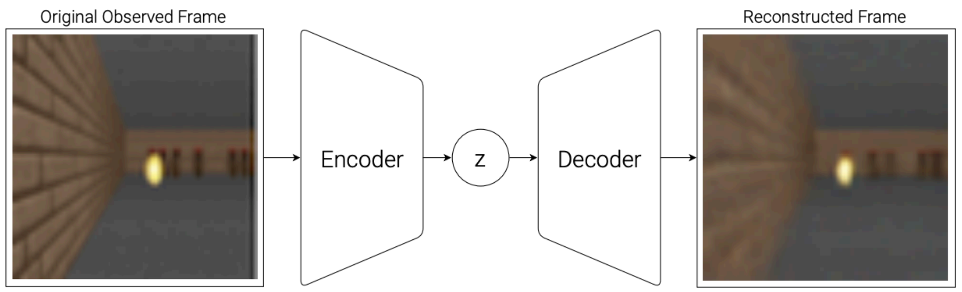

Figure 5. Flow diagram of a Variational Autoencoder (VAE).

Ha, David, and Jürgen Schmidhuber. "World models." arXiv preprint arXiv:1803.10122 (2018).

however…

"…the rule "collect truth for truth's sake" may be justified when the truth is unchanging; but **when the system is not completely isolated from its surroundings, and is undergoing secular changes, the collection of truth is futile, for it will not keep.**"

– Ashby W. R. (1958)

Example: Braitenberg vehicles

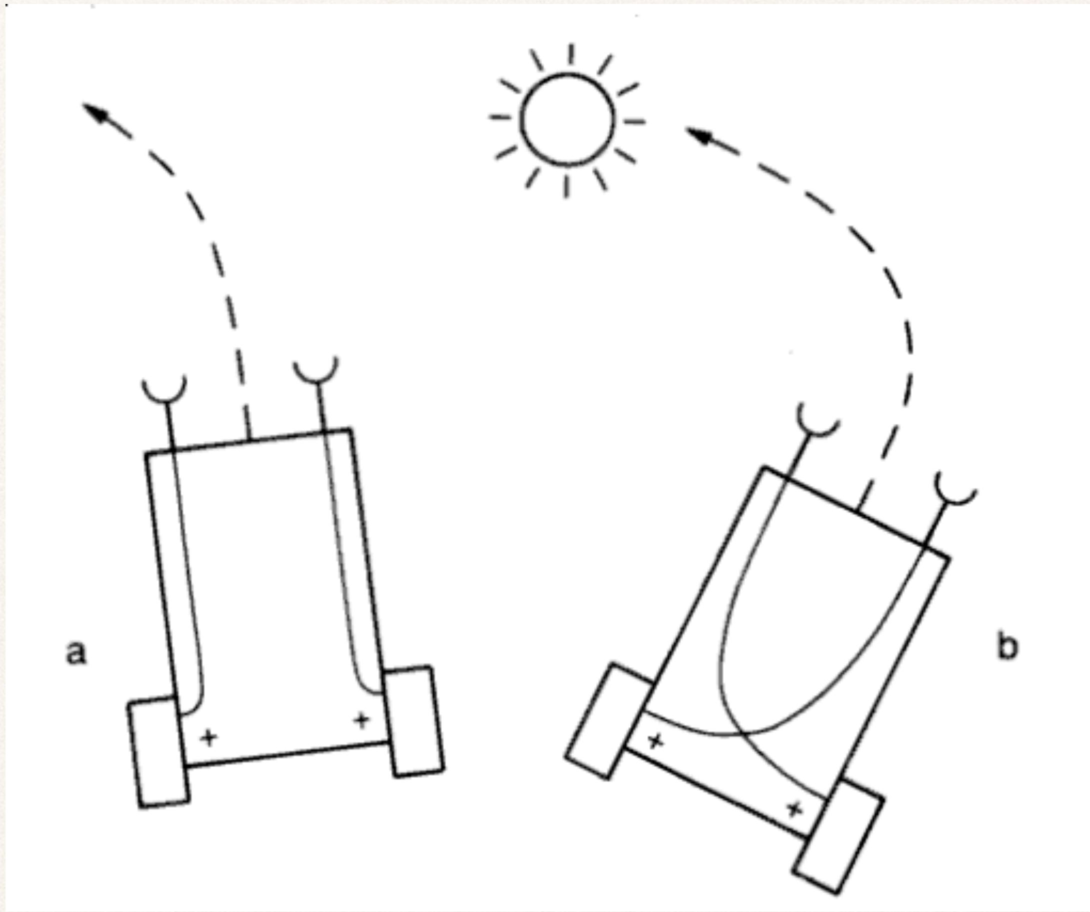

Braitenberg, Valentino. Vehicles: Experiments in synthetic psychology. MIT press, 1986.

Phototaxis in active inference

Generative model

$$y_{l_1} = x_{l_1} + z_{l_1} \qquad y_{l_2} = x_{l_2} + z_{l_2}$$

$$y_{m_1} = x_{m_1} + z_{m_1} \qquad y_{m_2} = x_{m_2} + z_{m_2}$$

$$x_{m_1} = x_{l_2} + w_{m_1} \qquad x_{m_2} = x_{l_1} + w_{m_2}$$

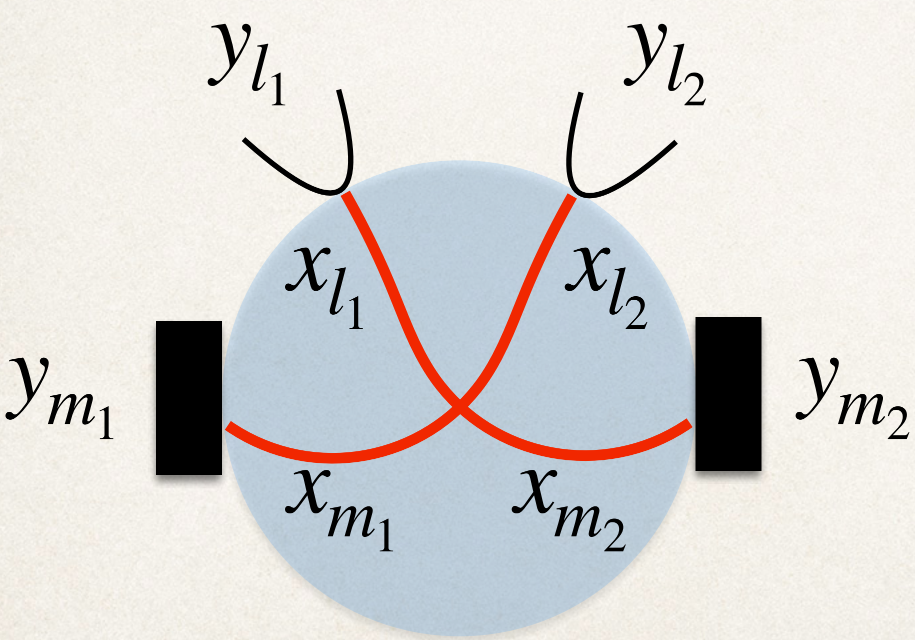

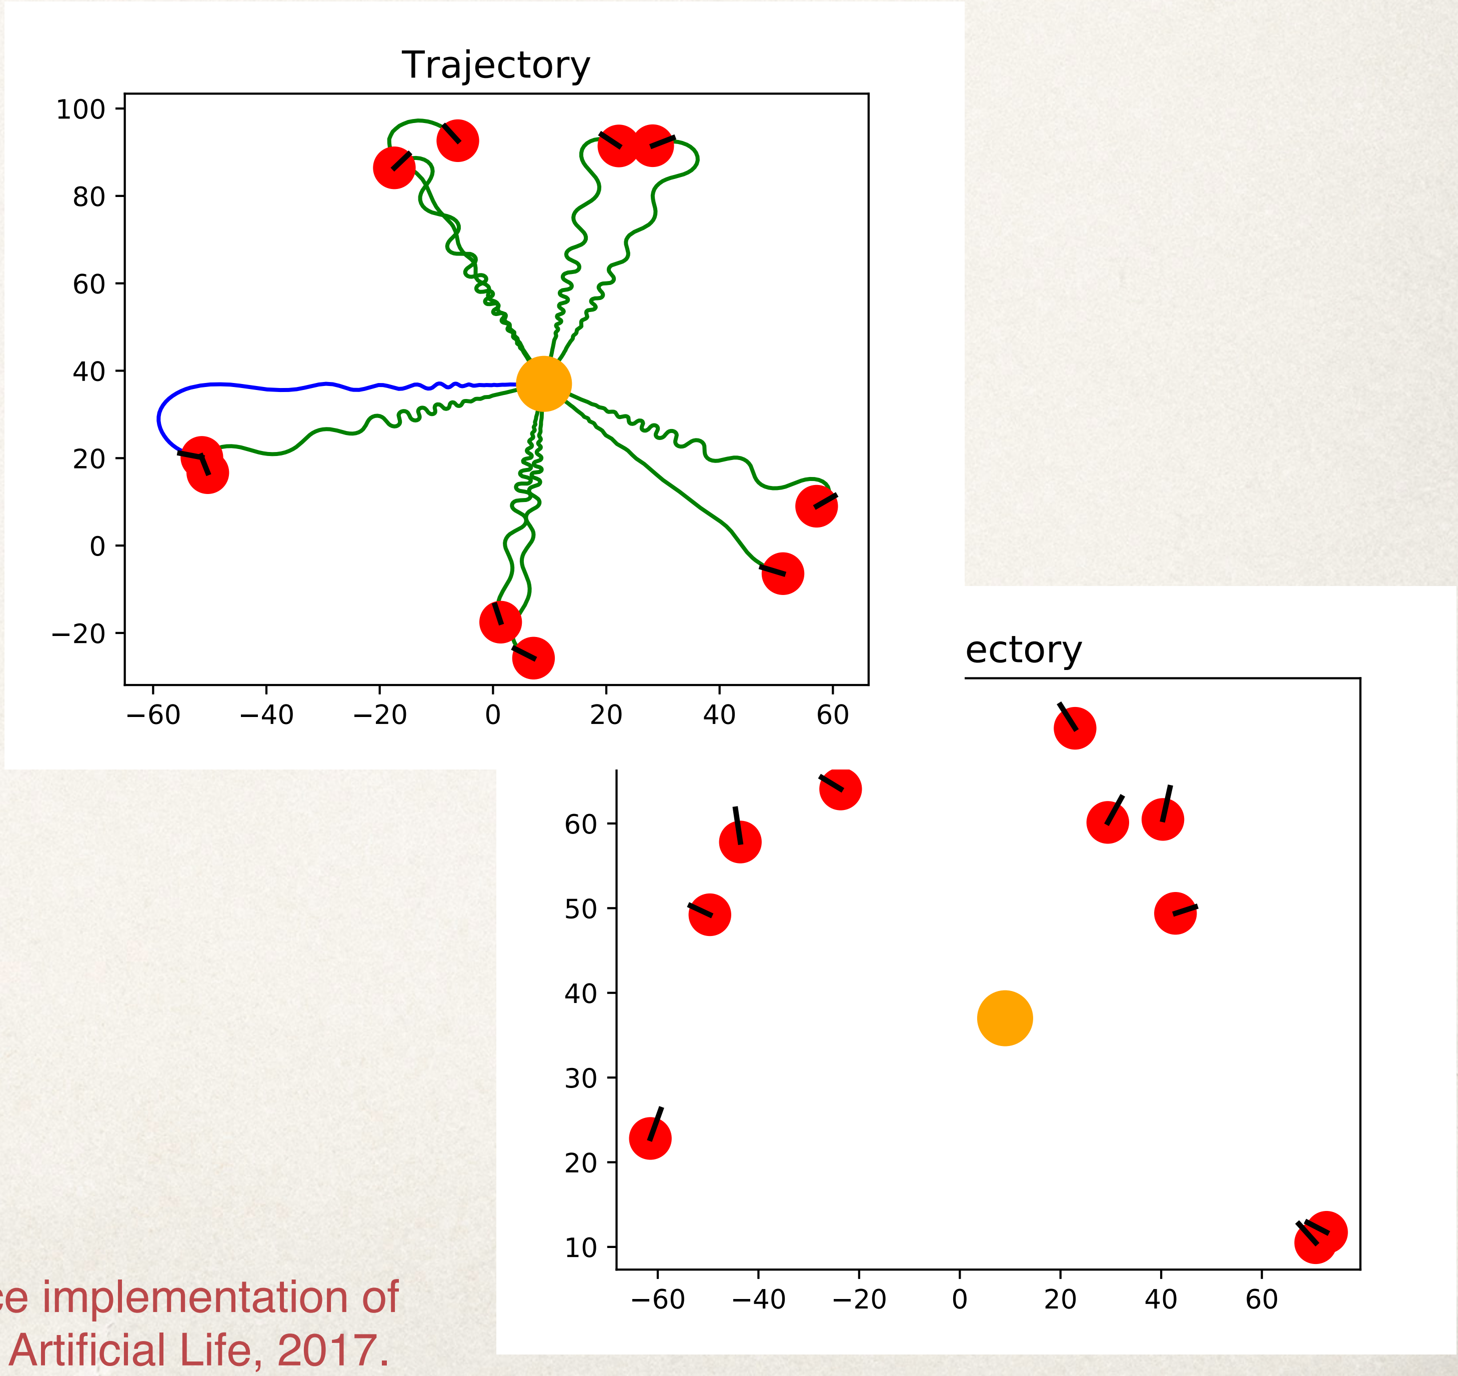

Baltieri, M. and Buckley, C. L. (2017). "An active inference implementation of phototaxis." Proceedings of the European Conference on Artificial Life, 2017.

Braitenberg vehicles-like agents in active inference

Variational free energy for fixed-form VI

$$F \approx \frac{1}{2}\Big(\pi_{z_{l_1}}(y_{l_1} - \mu_{l_1})^2 + \pi_{z_{l_2}}(y_{l_2} - \mu_{l_2})^2 + \pi_{z_{m_1}}(y_{m_1} - \mu_{m_1})^2 + \pi_{z_{m_2}}(y_{m_2} - \mu_{m_2})^2$$
$$+\pi_{w_{m_1}}(\mu_{m_1} - \mu_{l_2})^2 + \pi_{w_{m_2}}(\mu_{m_2} - \mu_{x_1})^2 - \ln(\pi_{z_{l_1}}\pi_{z_{l_2}}\pi_{z_{m_1}}\pi_{z_{m_2}}\pi_{w_{m_1}}\pi_{w_{m_2}})\Big)$$

Variational updates

Perception

$$\dot{\mu}_{l_1} = -k\Big(\pi_{z_{l_1}}(\mu_{l_1} - y_{l_1}) + \pi_{w_{m_2}}(\mu_{l_1} - \mu_{m_2})\Big)$$
$$\dot{\mu}_{l_2} = -k\Big(\pi_{z_{l_2}}(\mu_{l_2} - y_{l_2}) + \pi_{w_{m_1}}(\mu_{l_2} - \mu_{m_1})\Big)$$
$$\dot{\mu}_{m_1} = -k\Big(\pi_{z_{m_1}}(\mu_{m_1} - y_{m_1}) + \pi_{w_{m_1}}(\mu_{m_1} - \mu_{l_2})\Big)$$
$$\dot{\mu}_{m_2} = -k\Big(\pi_{z_{m_1}}(\mu_{m_2} - y_{m_2}) + \pi_{w_{m_2}}(\mu_{m_2} - \mu_{l_1})\Big)$$

Action

$$\dot{a}_1 = -k\Big(\pi_{z_{m_1}}(y_{m_1} - \mu_{m_1})\frac{\partial y_{m_1}}{\partial a_1} + \pi_{z_{m_2}}(y_{m_2} - \mu_{m_2})\frac{\partial y_{m_2}}{\partial a_1}\Big)$$
$$\dot{a}_2 = -k\Big(\pi_{z_{m_1}}(y_{m_1} - \mu_{m_1})\frac{\partial y_{m_1}}{\partial a_2} + \pi_{z_{m_2}}(y_{m_2} - \mu_{m_2})\frac{\partial y_{m_2}}{\partial a_2}\Big)$$

The physics of the problem

Forces

Torques

Agent's body

…

The belief space of the problem?

Forces	=	Forces
Torques	=	Torques
Generative Process	=	Generative Model
Agent's body	=	Agent's body
…	=	…

The belief space of the agent

Forces	=	~~Forces~~
Torques	=	~~Torques~~
Agent's body	=	~~Agent's body~~
…	=	…

See also:

- **Baltieri, M.** and Buckley, C. L. (2019). "Generative models as parsimonious descriptions of sensorimotor loops." (Commentary to Brette (2019): Is coding a relevant metaphor for the brain? Behavioral and Brain Sciences.)
- **Baltieri M.,** Buckley C.L. and Bruineberg J., "Predictions in the eye of the beholder: an active inference account of Watt governors." Proceedings of the International Conference on Artificial Life, Montreal, Canada, 2020
- Mannella F., Maggiore F., **Baltieri M.** and Pezzulo G. (2021), "Active inference through whiskers" (accepted at Neural Networks)

Generative models, a spectrum

Reconstructing a copy of the world

Controlling the world with approximate models

Where our model(s) are

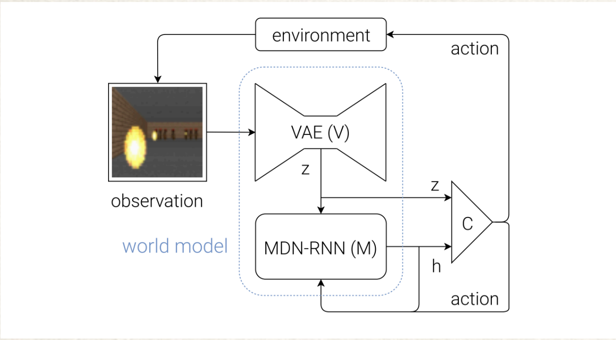

World models

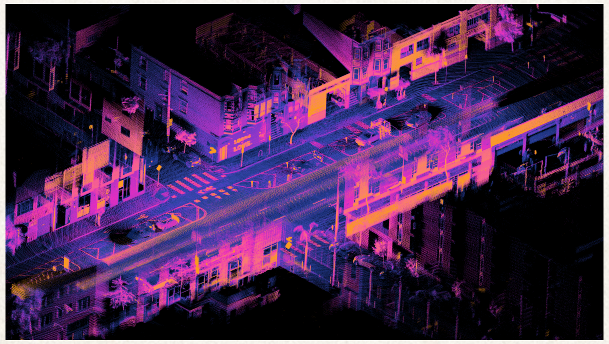

SLAM

https://en.wikipedia.org/wiki/Simultaneous_localization_and_mapping#/media/File:Ouster_OS1-64_lidar_point_cloud_of_intersection_of_Folsom_and_Dore_St,_San_Francisco.png

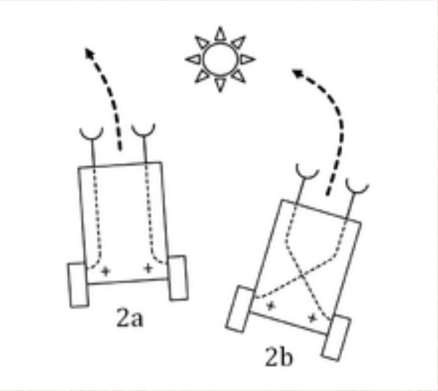

Braitenberg vehicles

More traditional example:PID

Set-point control where:

- P term (negative feedback, delta rule, Rescorla-Wagner)
- D term dampens oscillations
- I term deals with step changes, e.g., external unexpected inputs

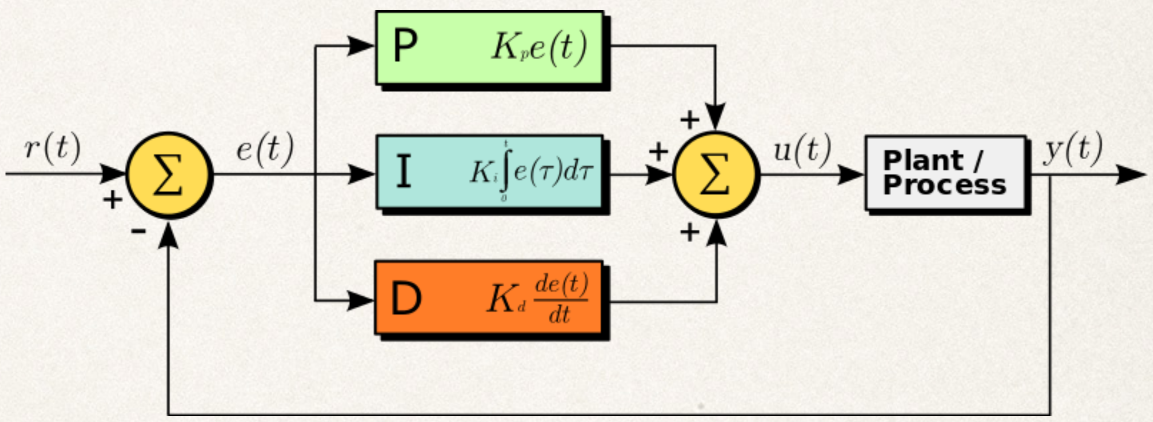

https://en.wikipedia.org/wiki/Proportional%E2%80%93integral%E2%80%93derivative_controller#/media/File:PID_en.svg

Applications

- Engineering (everywhere really, e.g., cruise controllers, thermostats)

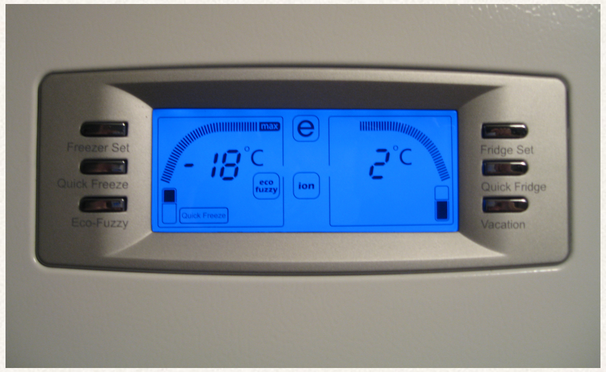

https://www.freeimages.com/photo/fridge-1325918

- Biology (e.g., chemotaxis in E. Coli, gene regulatory networks)

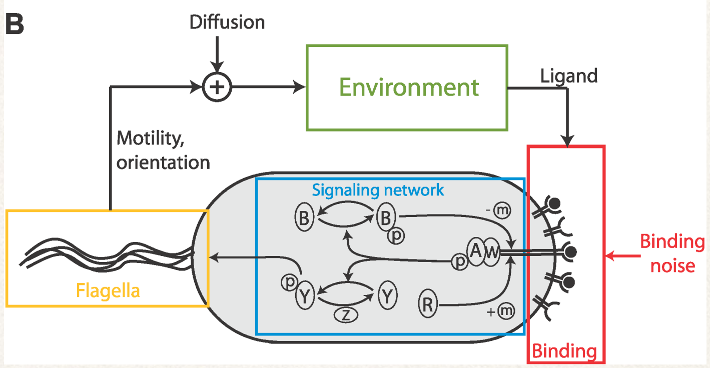

Andrews, Burton W., Tau-Mu Yi, and Pablo A. Iglesias. "Optimal noise filtering in the chemotactic response of Escherichia coli." PLoS computational biology 2.11 (2006): e154.

- Psychology (e.g., adaptive behaviour beyond delta rule)

PID controllers as linear generative models

Equation of motion (example)

$$m\frac{d^2s}{dt^2} = F - F_d$$

(disturbances)

$$F_d = F_g + F_r + F_a$$

$$F = r_g a(t) T_m \left(1 - \beta\left(\frac{\omega}{\omega_m} - 1\right)^2\right)$$

$$F_g = mg \sin \lambda$$

$$F_r = mg\, C_r\, \mathrm{sgn}(\dot{s})$$

$$F_a = \frac{1}{2}\rho\, C_d A \dot{s}^2$$

Generative model

$$\begin{aligned} y &= x + z & \dot{x} &= x' = -\alpha(x + v) + w \\ y' &= x' + z' & \dot{x}' &= x'' = -\alpha(x' + v') + w' \\ y'' &= x'' + z'' & \dot{x}'' &= x''' = -\alpha(x'' + v'') + w'' \end{aligned}$$

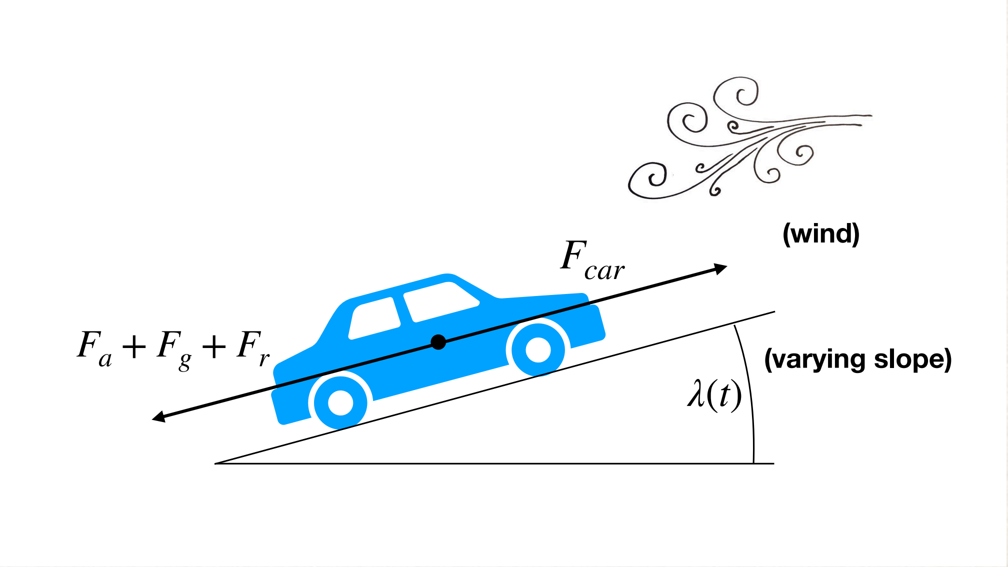

Baltieri, M. and Buckley, C. L. (2019). "PID Control as a Process of Active Inference with Linear Generative Models." Entropy, 21(3), 257.

A problem with PID parameters

$$a(t) = k_p e(t) + k_i \int_0^t e(\tau)d\tau + k_d \frac{de(t)}{dt}$$

$$e(t) = r - y(t)$$

(Standard PID control)

- How are (free) parameters k_p, k_i, k_d determined? Not even obvious what they mean.
- Huge (really massive) literature but, so far, mostly based on trial-and-error, look-up tables, heuristics, experience, etc.

Åström, Karl Johan, Tore Hägglund. Advanced PID control. 2006.

> 2000 citations (first edition, > 6000)
> <u>100 pages on how to find k's</u>

Franklin, Gene F., et al. Feedback control of dynamic systems. 2014.

> 6000 citations
> <u>300 pages on how to find k's</u>

A solution … ference

Gains k_p, k_i, k_d a … n higher embedding ord … ontinuous time).

They can be opt … sation scheme

AALBORG UNIVERSITY

Department of Electronic Systems
Fredrik Bajers Vej 7B
9220 Aalborg
www.es.aau.dk

Title:
PID Control as a Process of Active Inference Applied to a Refrigeration System

Project:
Master's Thesis

Semester:
Fourth

Project Period:
01/02/2021 - 03/06/2021

Project Group:
1034

Group Members:
Adrián Rocandio

Supervisors:
Henrik Schiøler
Roozbeh Izadi-Zamanabadi
Basil M. Al-Hadithi

Pages:
50

Submission:
03/06/2021

The content of this report is confidential.
© Aalborg University, 2021

This report is made with LaTeX

Abstract:

Classical PID control is a widely used technique in many industrial applications due to its good performance and relatively low complexity. Nevertheless, these regulators are not sufficient in some cases. This project investigates a novel probabilistic interpretation of PID control. Under this framework, it is assumed that only sensed variables are accessible. That is, no prior information of the process is available (i.e., plant model). Thus, the controller is furnished with a simple generative model that tries to deduce the measurement causes. This model, which is refined with every new measurement, permits designing the PID regulator. The innovation with respect to the classical approach is that here the controller gains encode measurement noise properties that can be inferred. The model enhancement and the applied control law obey a biological principle known as *free energy*.

The thesis proposes to implement this PID regulator in a refrigeration process. Specifically, it is aimed to control the evaporator outlet temperature. Simulation results prove good performance when dealing with changes in the set-point. The robustness test, however, shows poor outcomes as the system's response is not able to recover from a small input disturbance. Furthermore, the controller is sensitive to subtle changes in certain parameters when tuning, thus leading to instability.

Recently used with mixed results, so more tests will be needed!

Baltieri, M. and Buckley, C. L. (2019). "PID Control as a Process of Active Inference with Linear Generative Models." Entropy, 21(3), 257.

Learning gains with active inference

Real (log-)precisions
or (log-)gains

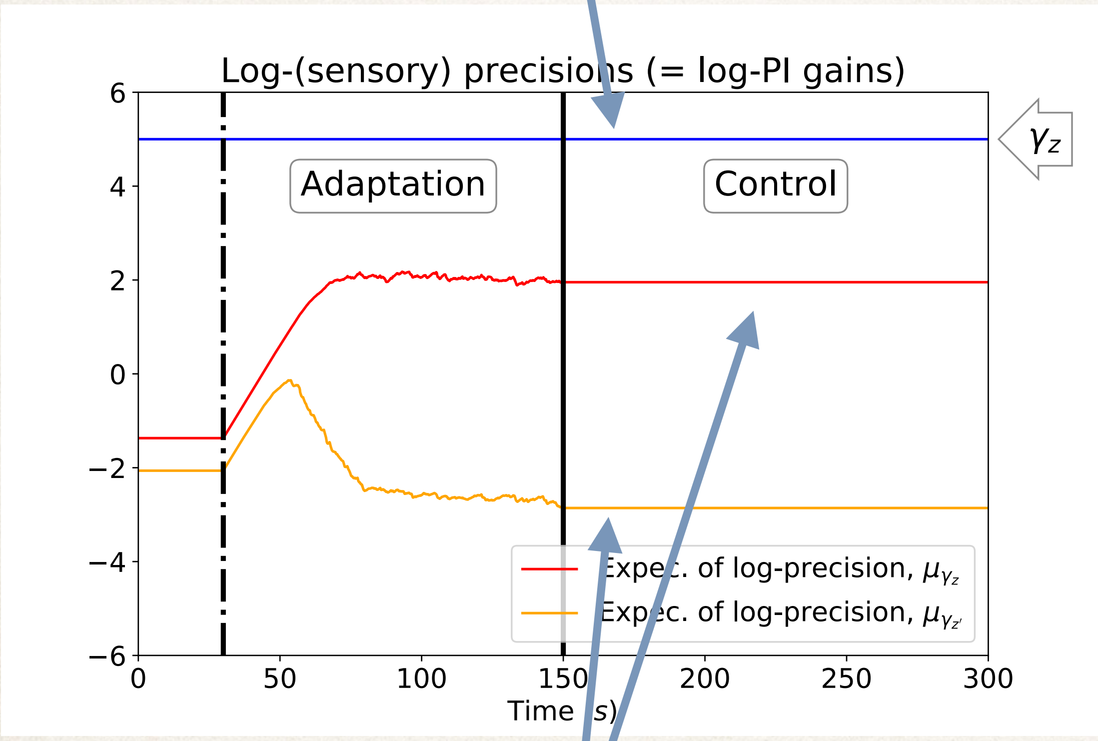

(target)

Agent's estimates of (log-)precisions
or (log-)gains <u>=/= real (log-)precisions</u>

Integral absolute error (IAE)
between two zero-crossings
(~ oscillations):

$$IAE = \int_{t}^{t+\tau} \left| e(t) \right| \, \mathrm{d}t$$

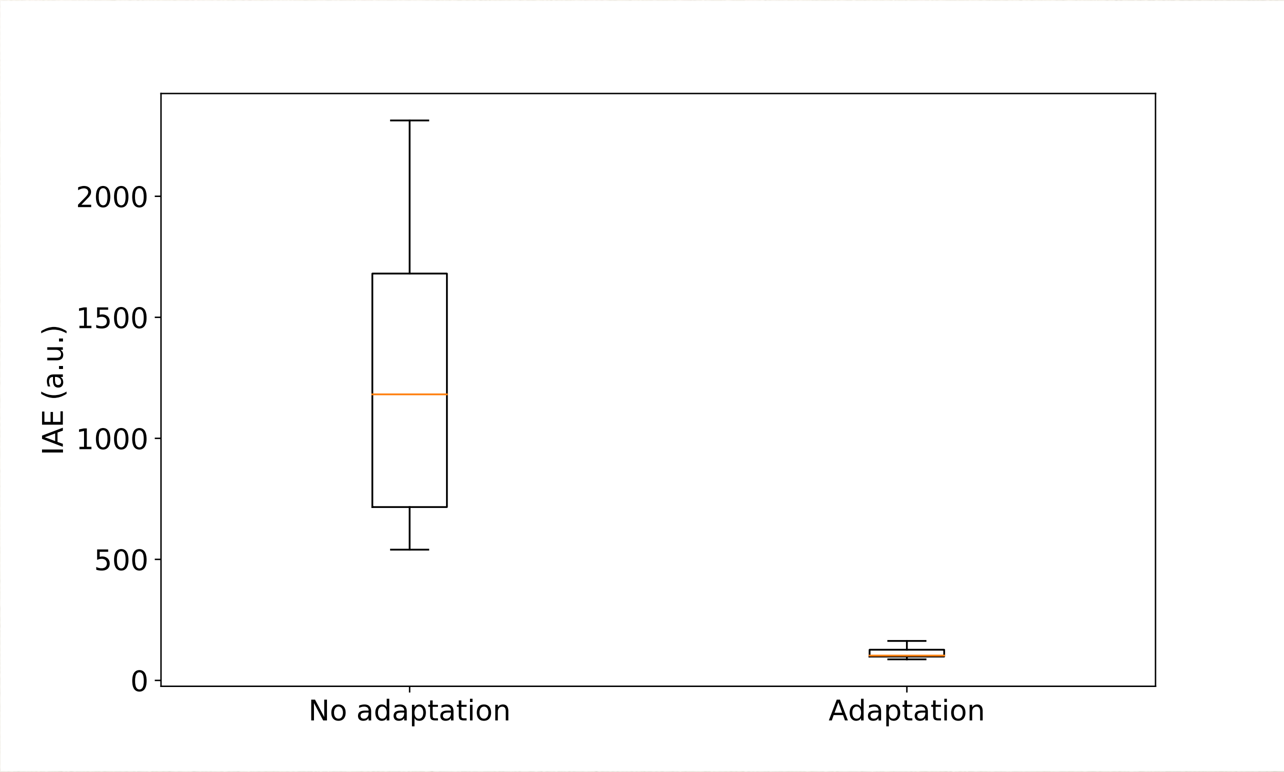

Bayes in classical control

Duality of inference and control (more later):

- Integral control is equivalent to inference on hidden (constant) inputs

Great! But got scooped by another paper…

- … in 1971 (although relatively ignored),
- generalising this to polynomial hidden inputs of arbitrary order (*n* orders —> *n+1* integrations),
- see review: Johnson, Carroll D. "On observers for systems with unknown and inaccessible inputs." International journal of control 21.5 (1975): 825-831.

Bayes for PID design

(Following Åstrom and Hägglund (2001))

Performance:

- load disturbance response, how a controller reacts to changes in external inputs, e.g. a step input
- set-point response, how a controller responds to different set-points over time
- measurement noise response, how noise on the observations impacts the regulation process

Robustness:

- robustness to model uncertainty, how uncertainty on the plant/environment dynamics affects the controller

Åström, Karl Johan, and Tore Hägglund. "The future of PID control." Control engineering practice 9.11 (2001): 1163-1175.

Bayes as a design framework

$$F \approx \frac{1}{2}\Bigg[\underbrace{\mu_{\pi_z}\left(y-\mu_x\right)^2+\mu_{\pi_{z'}}\left(y'-\mu_x'\right)^2+\mu_{\pi_{z''}}\left(y''-\mu_x''\right)^2}_{\text{Load disturbance response}}$$

$$\underbrace{+\mu_{\pi_w}\left(\mu_x'+\alpha\left(\mu_x-\eta_x\right)\right)^2+\mu_{\pi_{w'}}\left(\mu_x''+\alpha\left(\mu_x'-\eta_x'\right)\right)^2+\pi_{w''}\left(\mu_x'''+\alpha\left(\mu_x''-\eta_x''\right)\right)^2}_{\text{Set-point response}}$$

$$\text{Contexts}\left\{\begin{array}{l}\underbrace{+p_{\gamma_z}\left(\mu_{\gamma_z}-\eta_{\gamma_z}\right)^2+p_{\gamma_{z'}}\left(\mu_{\gamma_{z'}}-\eta_{\gamma_{z'}}\right)^2+p_{\gamma_{z''}}\left(\mu_{\gamma_{z''}}-\eta_{\gamma_{z''}}\right)^2}_{\text{Measurement noise response}} \\ \underbrace{+p_{\gamma_w}\left(\mu_{\gamma_w}-\eta_{\gamma_w}\right)^2+p_{\gamma_{w'}}\left(\mu_{\gamma_{w'}}-\eta_{\gamma_{w'}}\right)^2+p_{\gamma_{w''}}\left(\mu_{\gamma_{w''}}-\eta_{\gamma_{w''}}\right)^2-\ln(\varphi)\Bigg]}_{\text{Model uncertainty}}\end{array}\right.$$

Assumptions:

- Unknown (hyper)parameters
- Re-parametrisation for non-negativity

$$\pi_{\tilde{z}} \rightarrow \mu_{\pi_{\tilde{z}}}$$

$$\mu_{\pi_{\tilde{z}}} = \exp\left(\mu_{\gamma_{\tilde{z}}}\right)$$

Baltieri, M., "A Bayesian perspective on classical control", Proceedings of the International Joint Conference on Neural Networks, Glasgow, UK, 2020

From control theory to cognitive agents (Claim 1 + Claim 2)

Perceptionput and appturn architectures (classical) cognitive (neuro) science

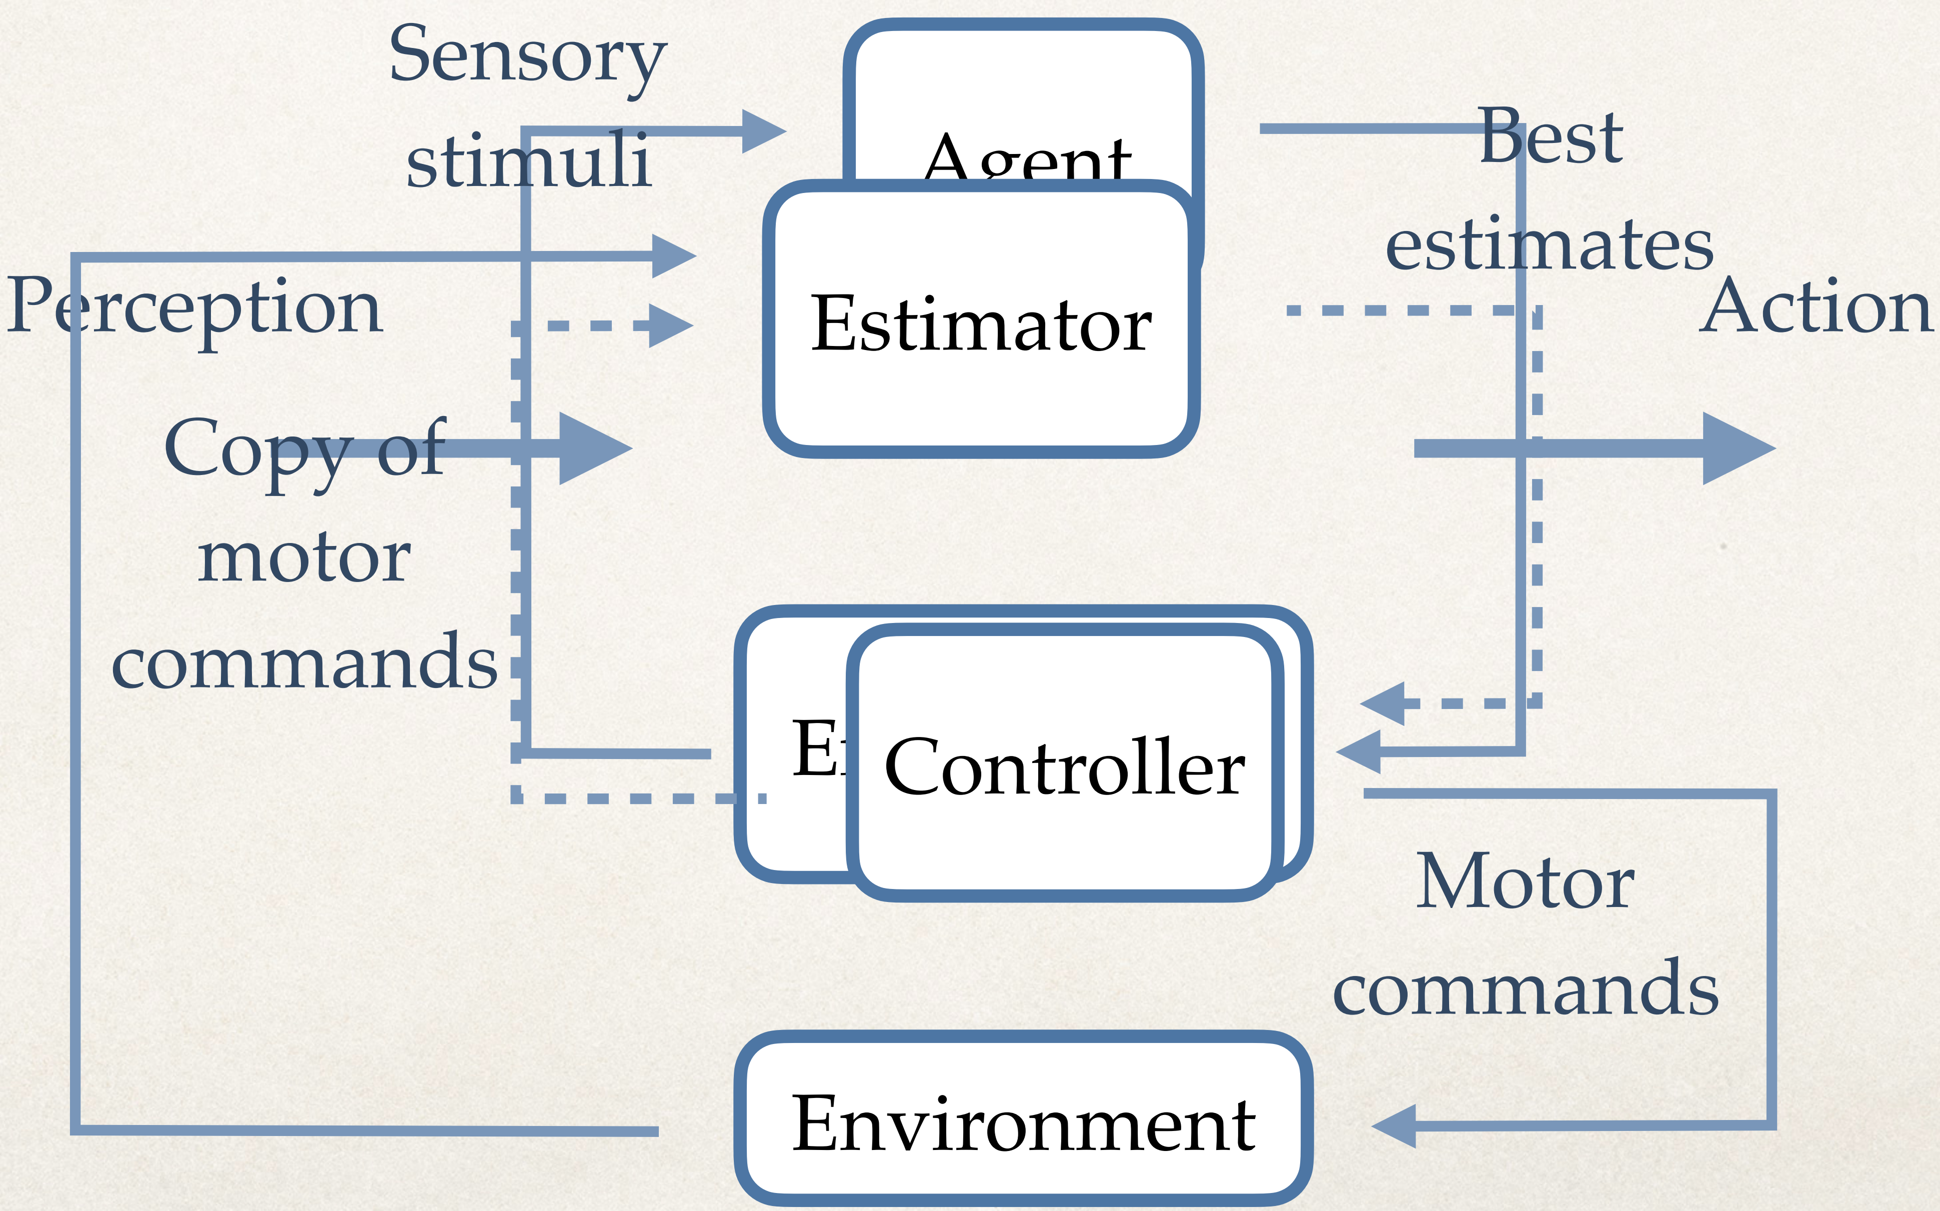

Estimation (perception) and control (action) are separable?

(Environment)

$y(t)$

Perception

$\hat{x}(t_1)$

Cognition

$g(\hat{x}(t_1))$

$h(\hat{x}(t_1))$

$\dots$

Action

$a(t_1) = u(t_1) = f(\hat{x}(t_1))$

"One may separate the problem of physical realization [of a controller] into two stages: computation of the "best approximation" $\hat{x}(t_1)$ of the state from knowledge of $y(t)$ for $t \leq t_1$ and computation of $u(t_1)$ given $\hat{x}(t_1)$."

"Contributions to the Theory of Optimal Control"

– Kalman R. E. (1960)

The sandwich of cognitive science, or sense-model-plan-act architectures in robotics (see also World Models)

https://pixabay.com/photos/toast-vegan-sandwich-vegan-breakfast-7009956/

The separation principle

Classic result in control theory (cf. “certainty equivalence” in econometrics and separation principle in information theory) for **linear systems**:

LQG (Linear Quadratic Gaussian) control =

Kalman filter (estimator) + Linear quadratic regulator (controller)

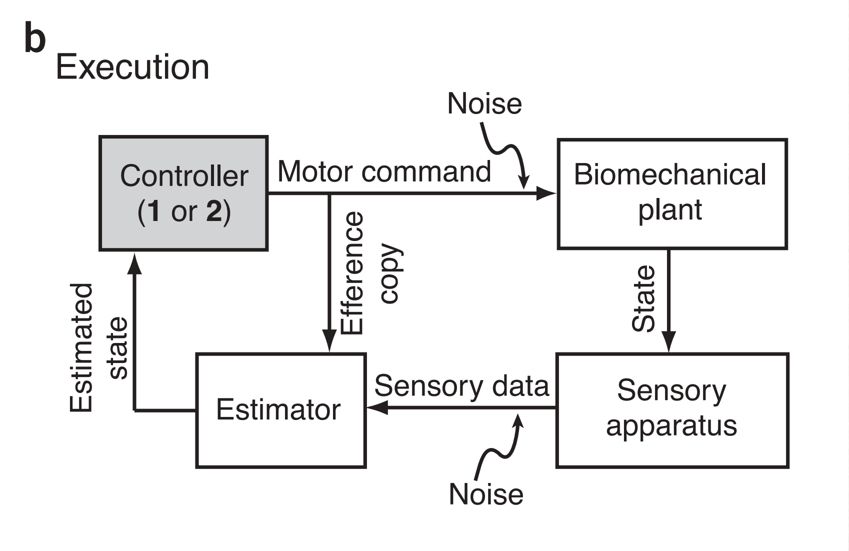

e.g., Todorov (2004)

Todorov, Emanuel. "Optimality principles in sensorimotor control." *Nature neuroscience* 7.9 (2004): 907-915.

The duality of estimation and control

Linear case, Kalman filter (KF) and linear quadratic regulator LQR, for a generalisation see Todorov (2008)

LQE and LQR both solve a Riccati Equation (RE)

$$\dot{y}(x) = q_0(x) + q_1(x)y(x) + q_2(x)y^2(x)$$

- KF $$\dot{P} = CC^T + AP + PA^T - PH^T(DD^T)^{-1}HP$$
- LQR $$-\dot{V} = Q + A^TV + VA - VBR^{-1}B^TV$$

Todorov, Emanuel. "General duality between optimal control and estimation." 2008 47th IEEE Conference on Decision and Control. IEEE, 2008.

The duality of estimation and control - (roughly)

- KF integrates RE forwards in time, LQR backwards.
- Estimation and control seem to solve the same type of (*inference*) problem.
- Techniques from Bayesian inference can be applied to (stochastic) optimal control and vice-versa (e.g. KL-control, path integral control, control as inference, planning as inference, active inference)
- Approximate Bayesian Inference (ABI) appears when exact inference is unfeasible (most of the interesting cases)

The dual role of estimation and control

- Dual role =/= duality
- Usually, estimator and controller are two separate modules (i.e., factorisable generative model, to some extent at least), see LQG
- However many interesting problems involve exploration/exploration problem or *dual control* in control theory, Feldbaum (1960), non-factorisable/non-separable

Modular minds and the separation principle

Robotics and AI, classical sandwich in cog. science	*Cog. (neuro)science*	*Control theory, separation principle*
Perception	Estimation/inference	Kalman(-Bucy) filter
↓	↓	↓
Cognition	(Complicated stuff or "just inference", à la Friston)	(Complicated stuff)
↓	↓	↓
Action	Optimal control	Linear Quadratic Regulator

What about active inference?

Active inference is *biased* inference, i.e. inputs are assumed to be unknown, both external disturbances and internal motor commands

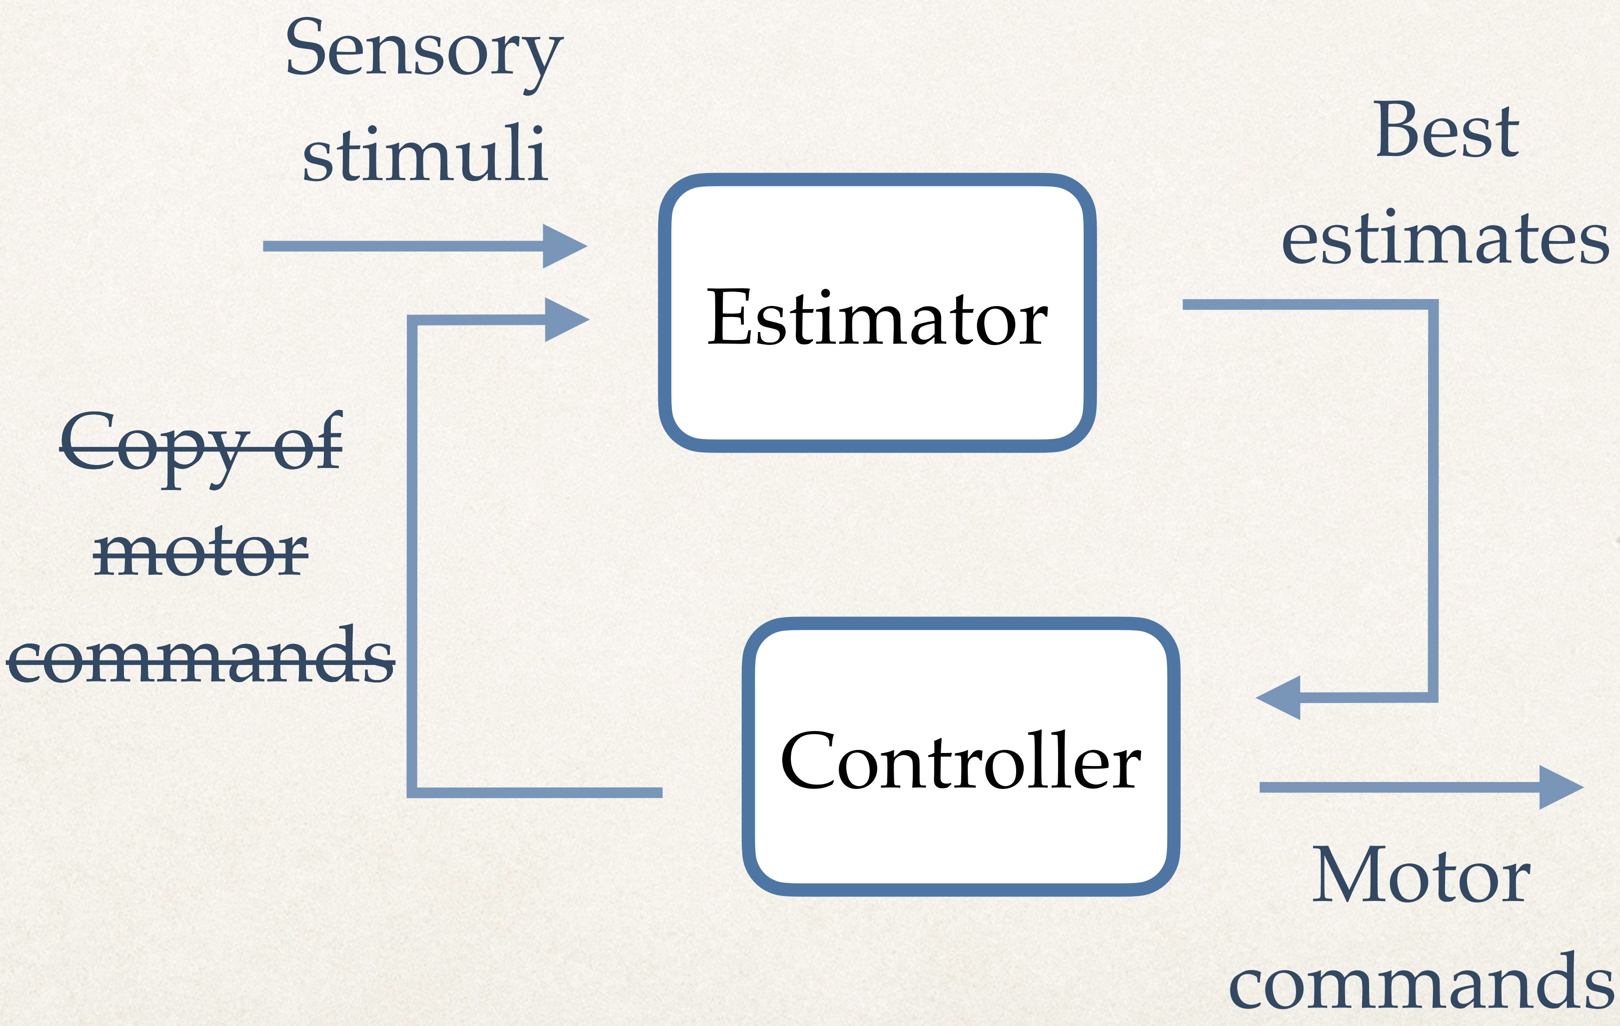

Baltieri, M. and Buckley, C. L. (2018). *"The modularity of action and perception revisited using control theory and active inference."* Proceedings of the International Conference on Artificial Life, Tokyo, Japan, 2018.

LQG vs active inference

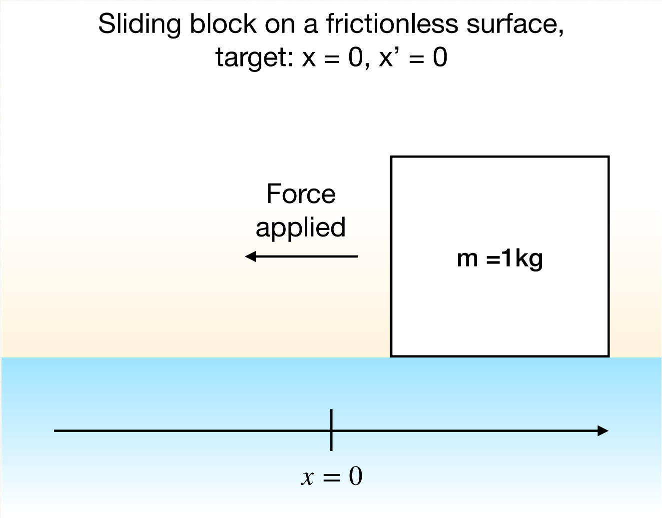

Double integrator ~ model of single joint

$$\dot{\boldsymbol{x}} = A\boldsymbol{x} + B\boldsymbol{a} + \boldsymbol{w} \quad \boldsymbol{y} = C\boldsymbol{x} + \boldsymbol{z}$$

where matrices A, B, C are defined as:

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \quad B = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \quad C = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

and covariance matrices Σ_z, Σ_w as:

$$\Sigma_z = \begin{bmatrix} \exp(0) & 0 \\ 0 & \exp(0) \end{bmatrix} \quad \Sigma_w = \begin{bmatrix} 0 & 0 \\ 0 & \exp(-1) \end{bmatrix}$$

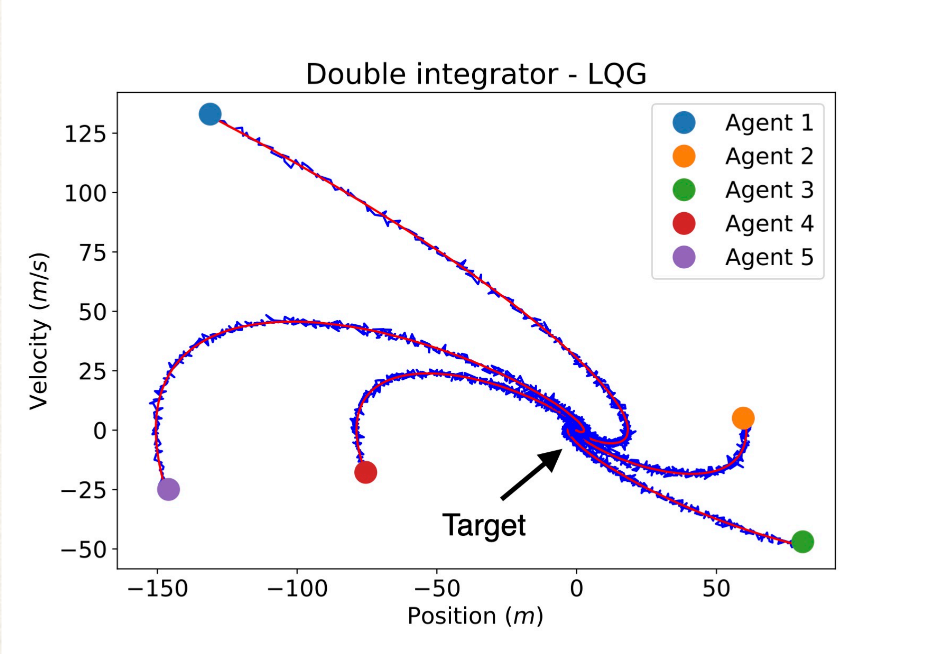

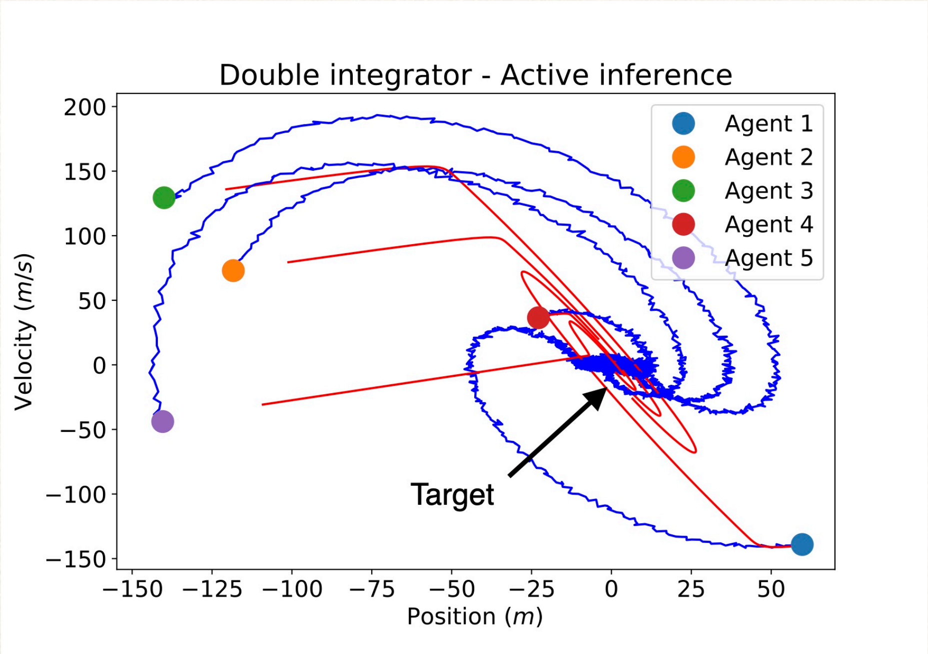

- **Baltieri, M.** and Buckley, C. L. (2019). “Nonmodular architectures of cognitive systems based on active inference.” Proceedings of the International Joint Conference on Neural Networks, Budapest, Humgary, 2019
- **Baltieri M.** and Buckley C. L., "On Kalman-Bucy filters, linear quadratic control and active inference”, arXiv pre-print arXiv:2005.06269 (2020)

LQG vs active inference

- LQG factorises control and inference, active inference doesn't (mostly)
- This leads to a formulation in terms of dual control, which in the more interesting (finite horizon) cases induces time-independent policies
- For a similar account, in discrete time, with less control theory and more RL/ML see also Millidge (2020)

Millidge, Beren, et al. "On the relationship between active inference and control as inference." *International Workshop on Active Inference*. Springer, Cham, 2020.

Part 2. (Work in progress)

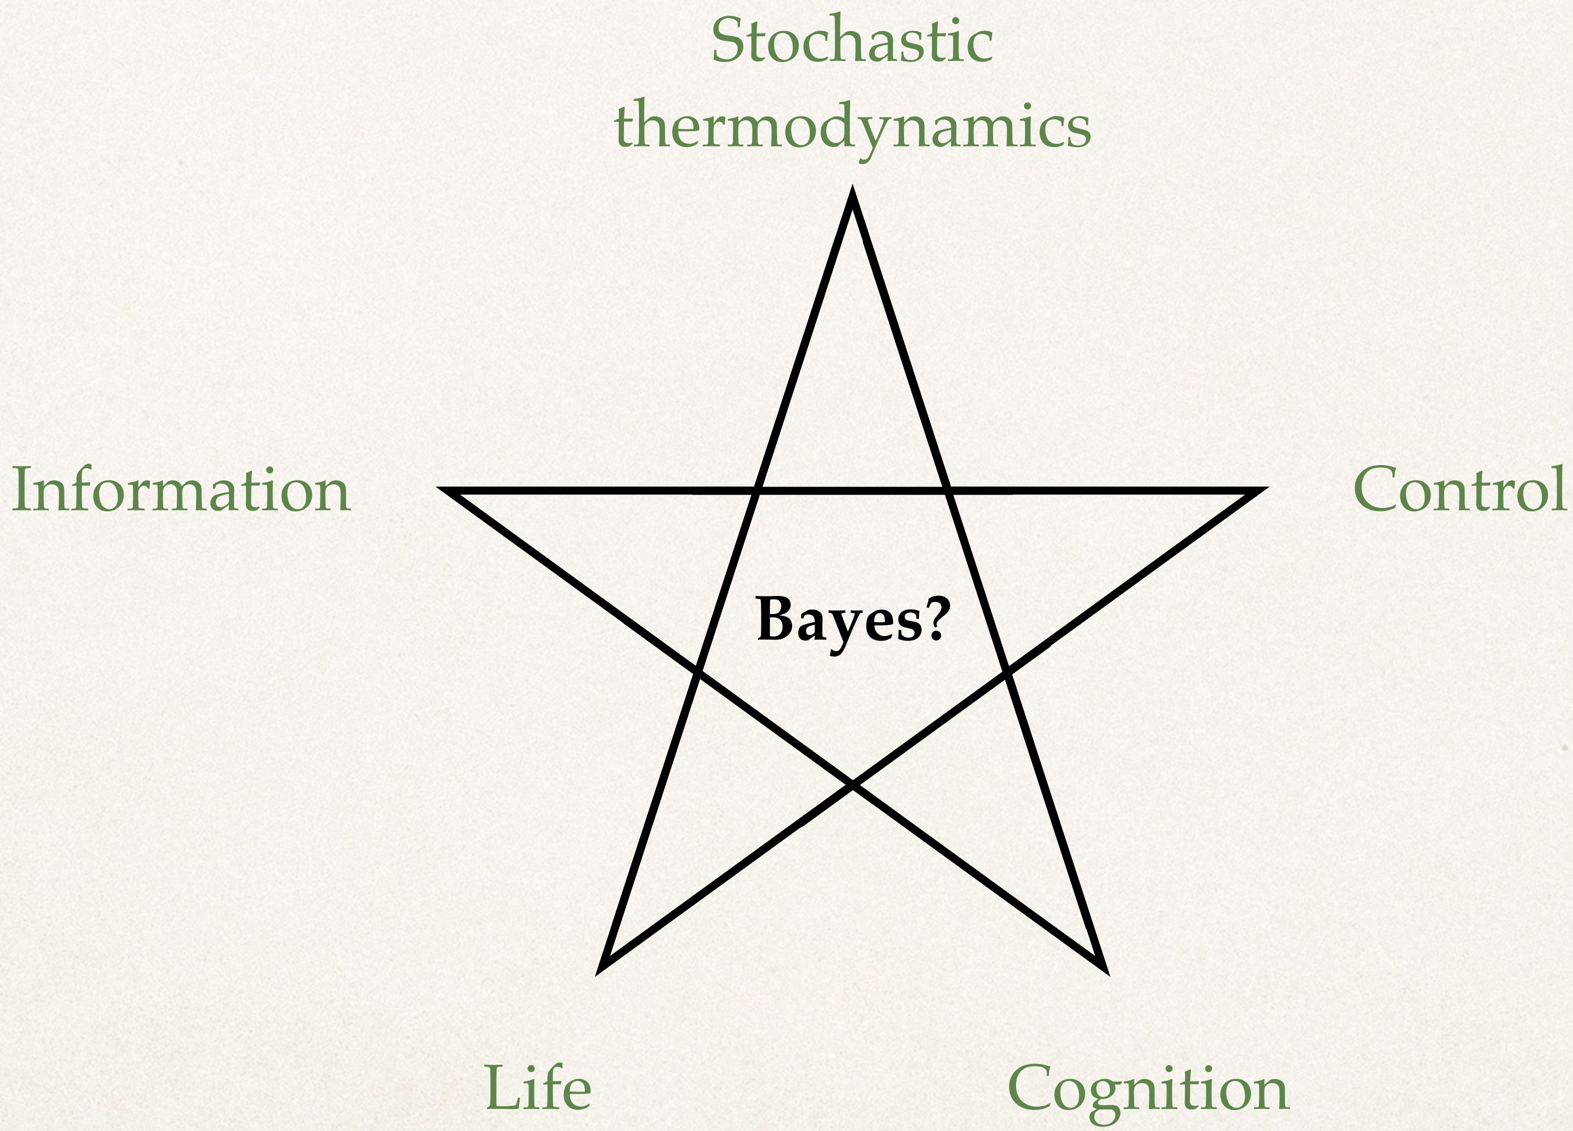

Friston's FEP

The 'free energy principle' (FEP): a framework based on variational inference to (attempt to) model life and cognition. Two lines of research:

- Use VB and derived techniques to model learning, inference, control, etc. (Part 1.)
- Use Bayes to *identify* agents in a stochastic process, given a set of conditional (in)dependences (e.g., a Bayesian network) and use VB to describe what the agent is and does in terms of its beliefs states

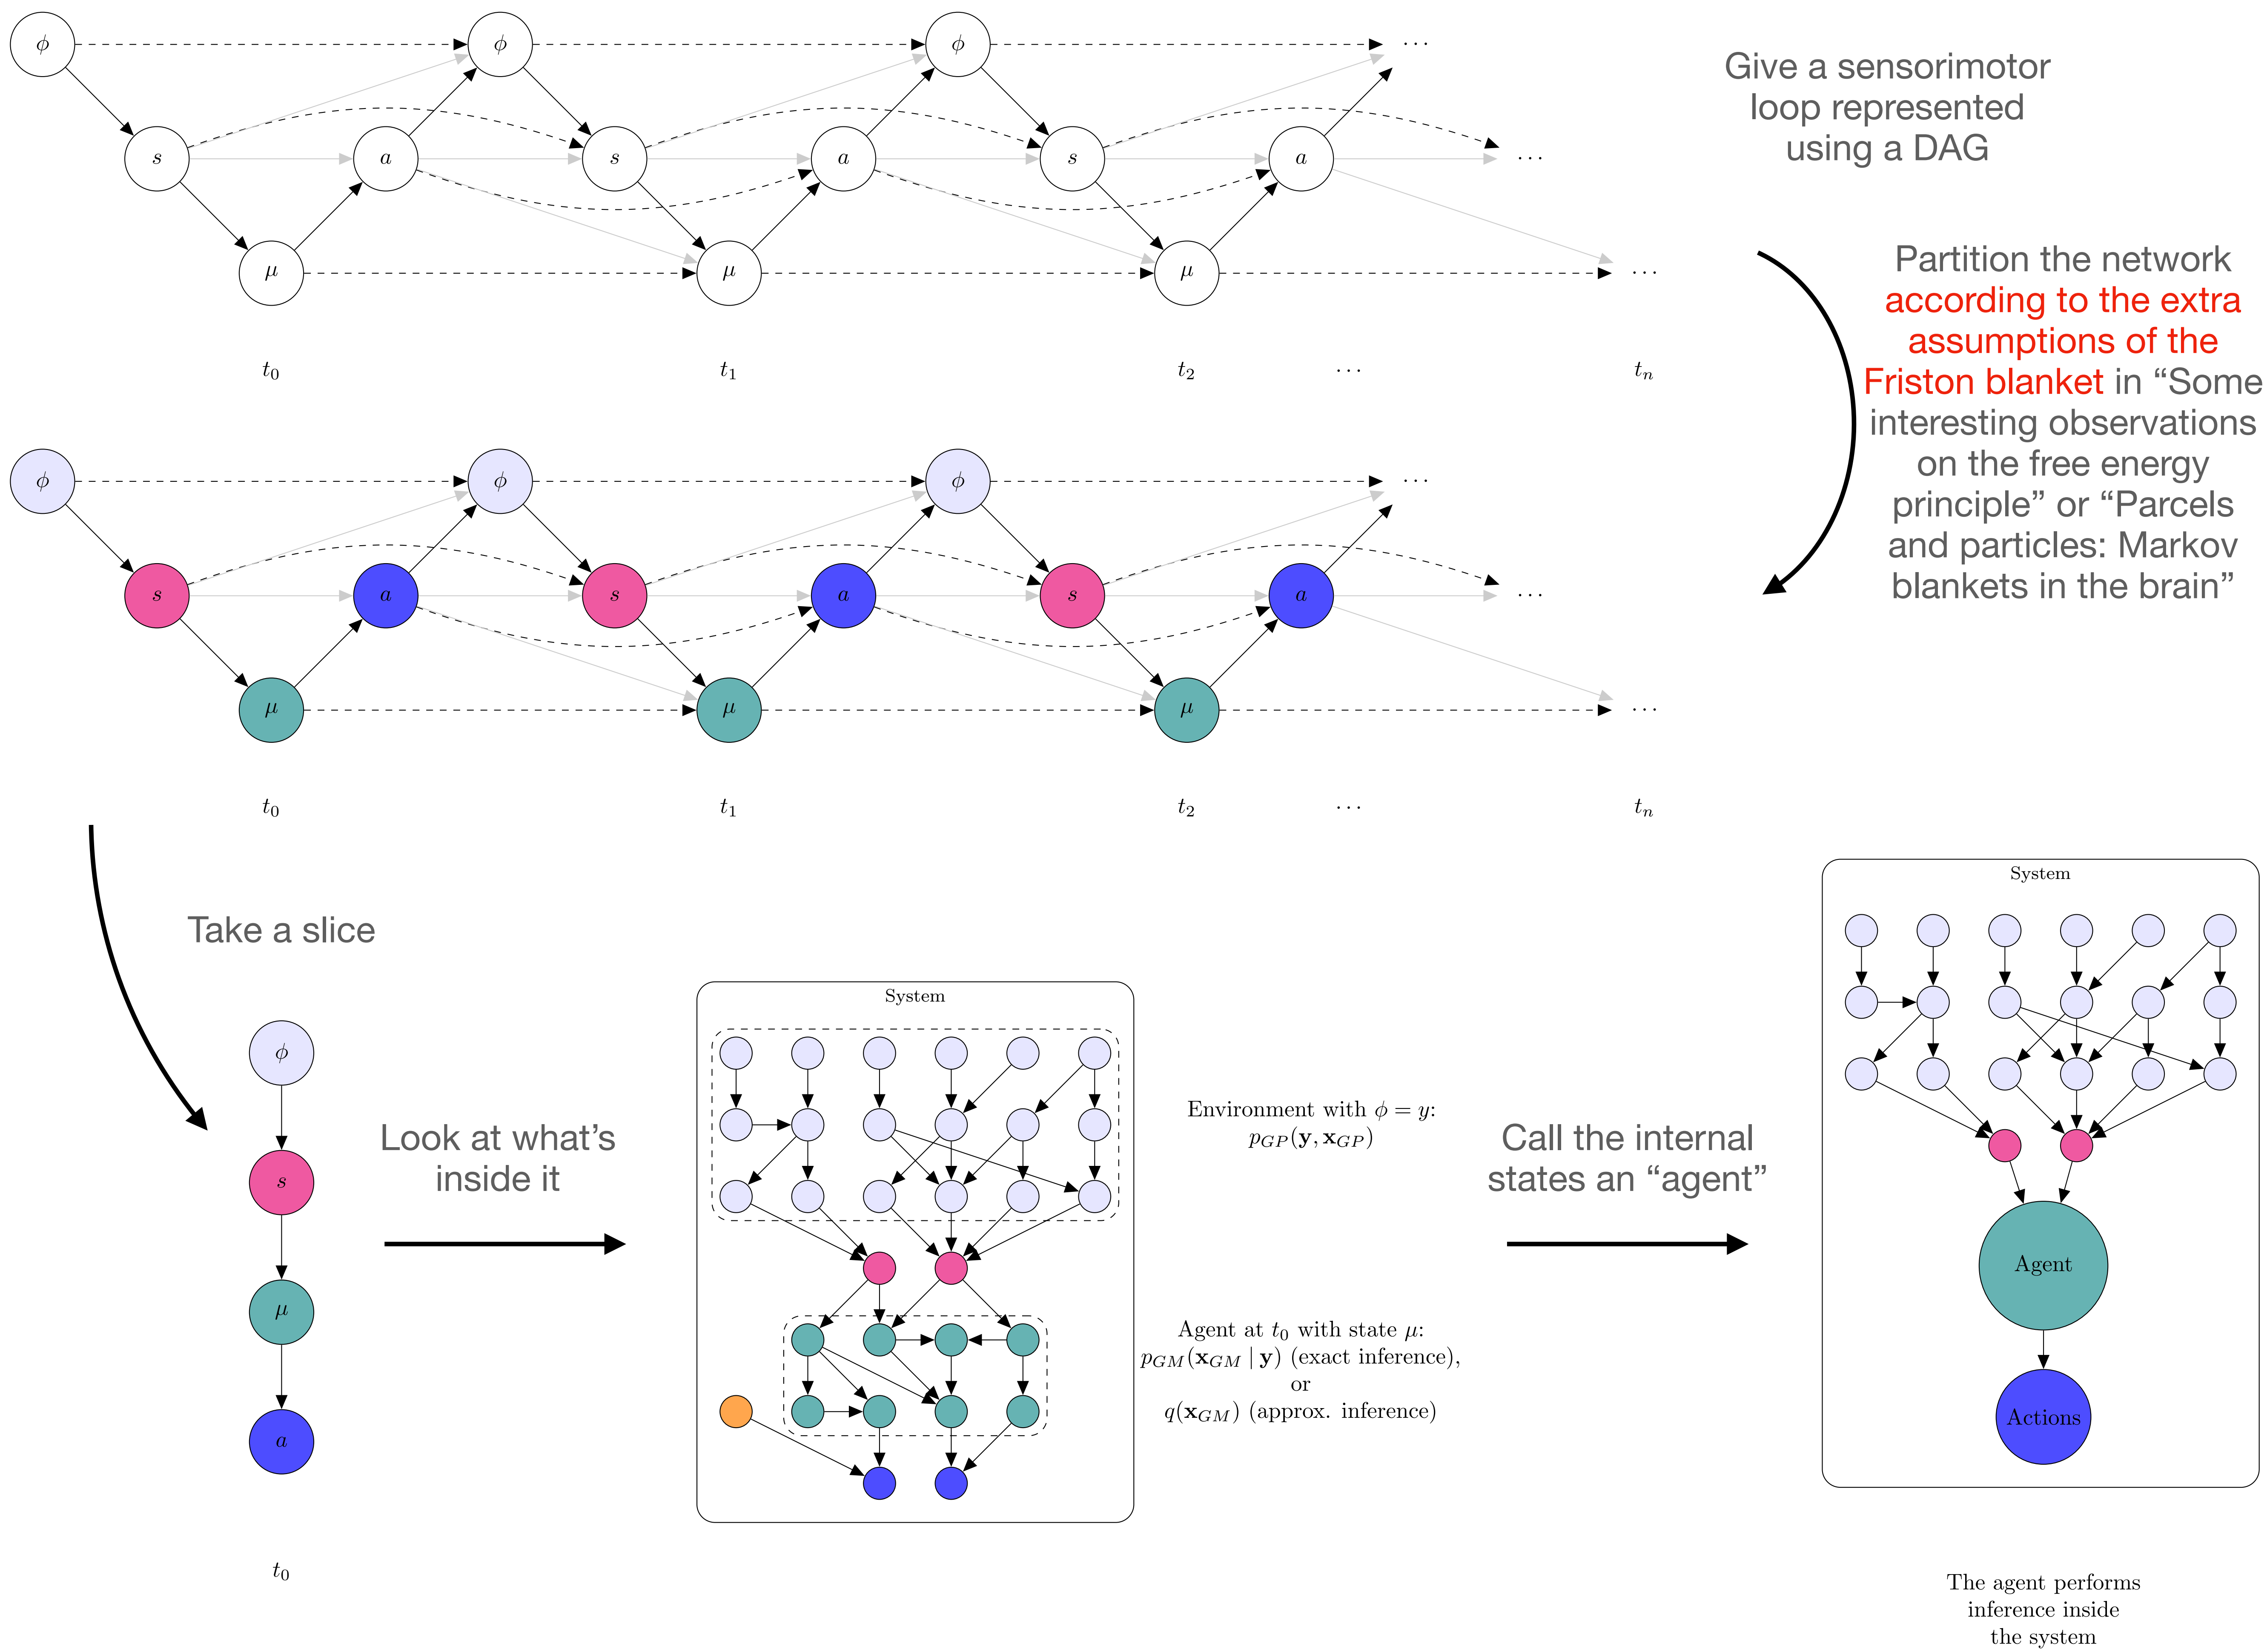

Give a sensorimotor loop represented using a DAG
Partition the network according to the extra assumptions of the Friston blanket in “Some interesting observations on the free energy principle” or “Parcels and particles: Markov blankets in the brain”
t_0
t_1
t_2
...
t_n
Take a slice
Look at what’s inside it
System
Environment with $\phi = y$: $p_{GP}(\mathbf{y}, \mathbf{x}_{GP})$
Agent at t_0 with state μ: $p_{GM}(\mathbf{x}_{GM} \mid \mathbf{y})$ (exact inference), or $q(\mathbf{x}_{GM})$ (approx. inference)
Call the internal states an “agent”
Agent
Actions
The agent performs inference inside the system

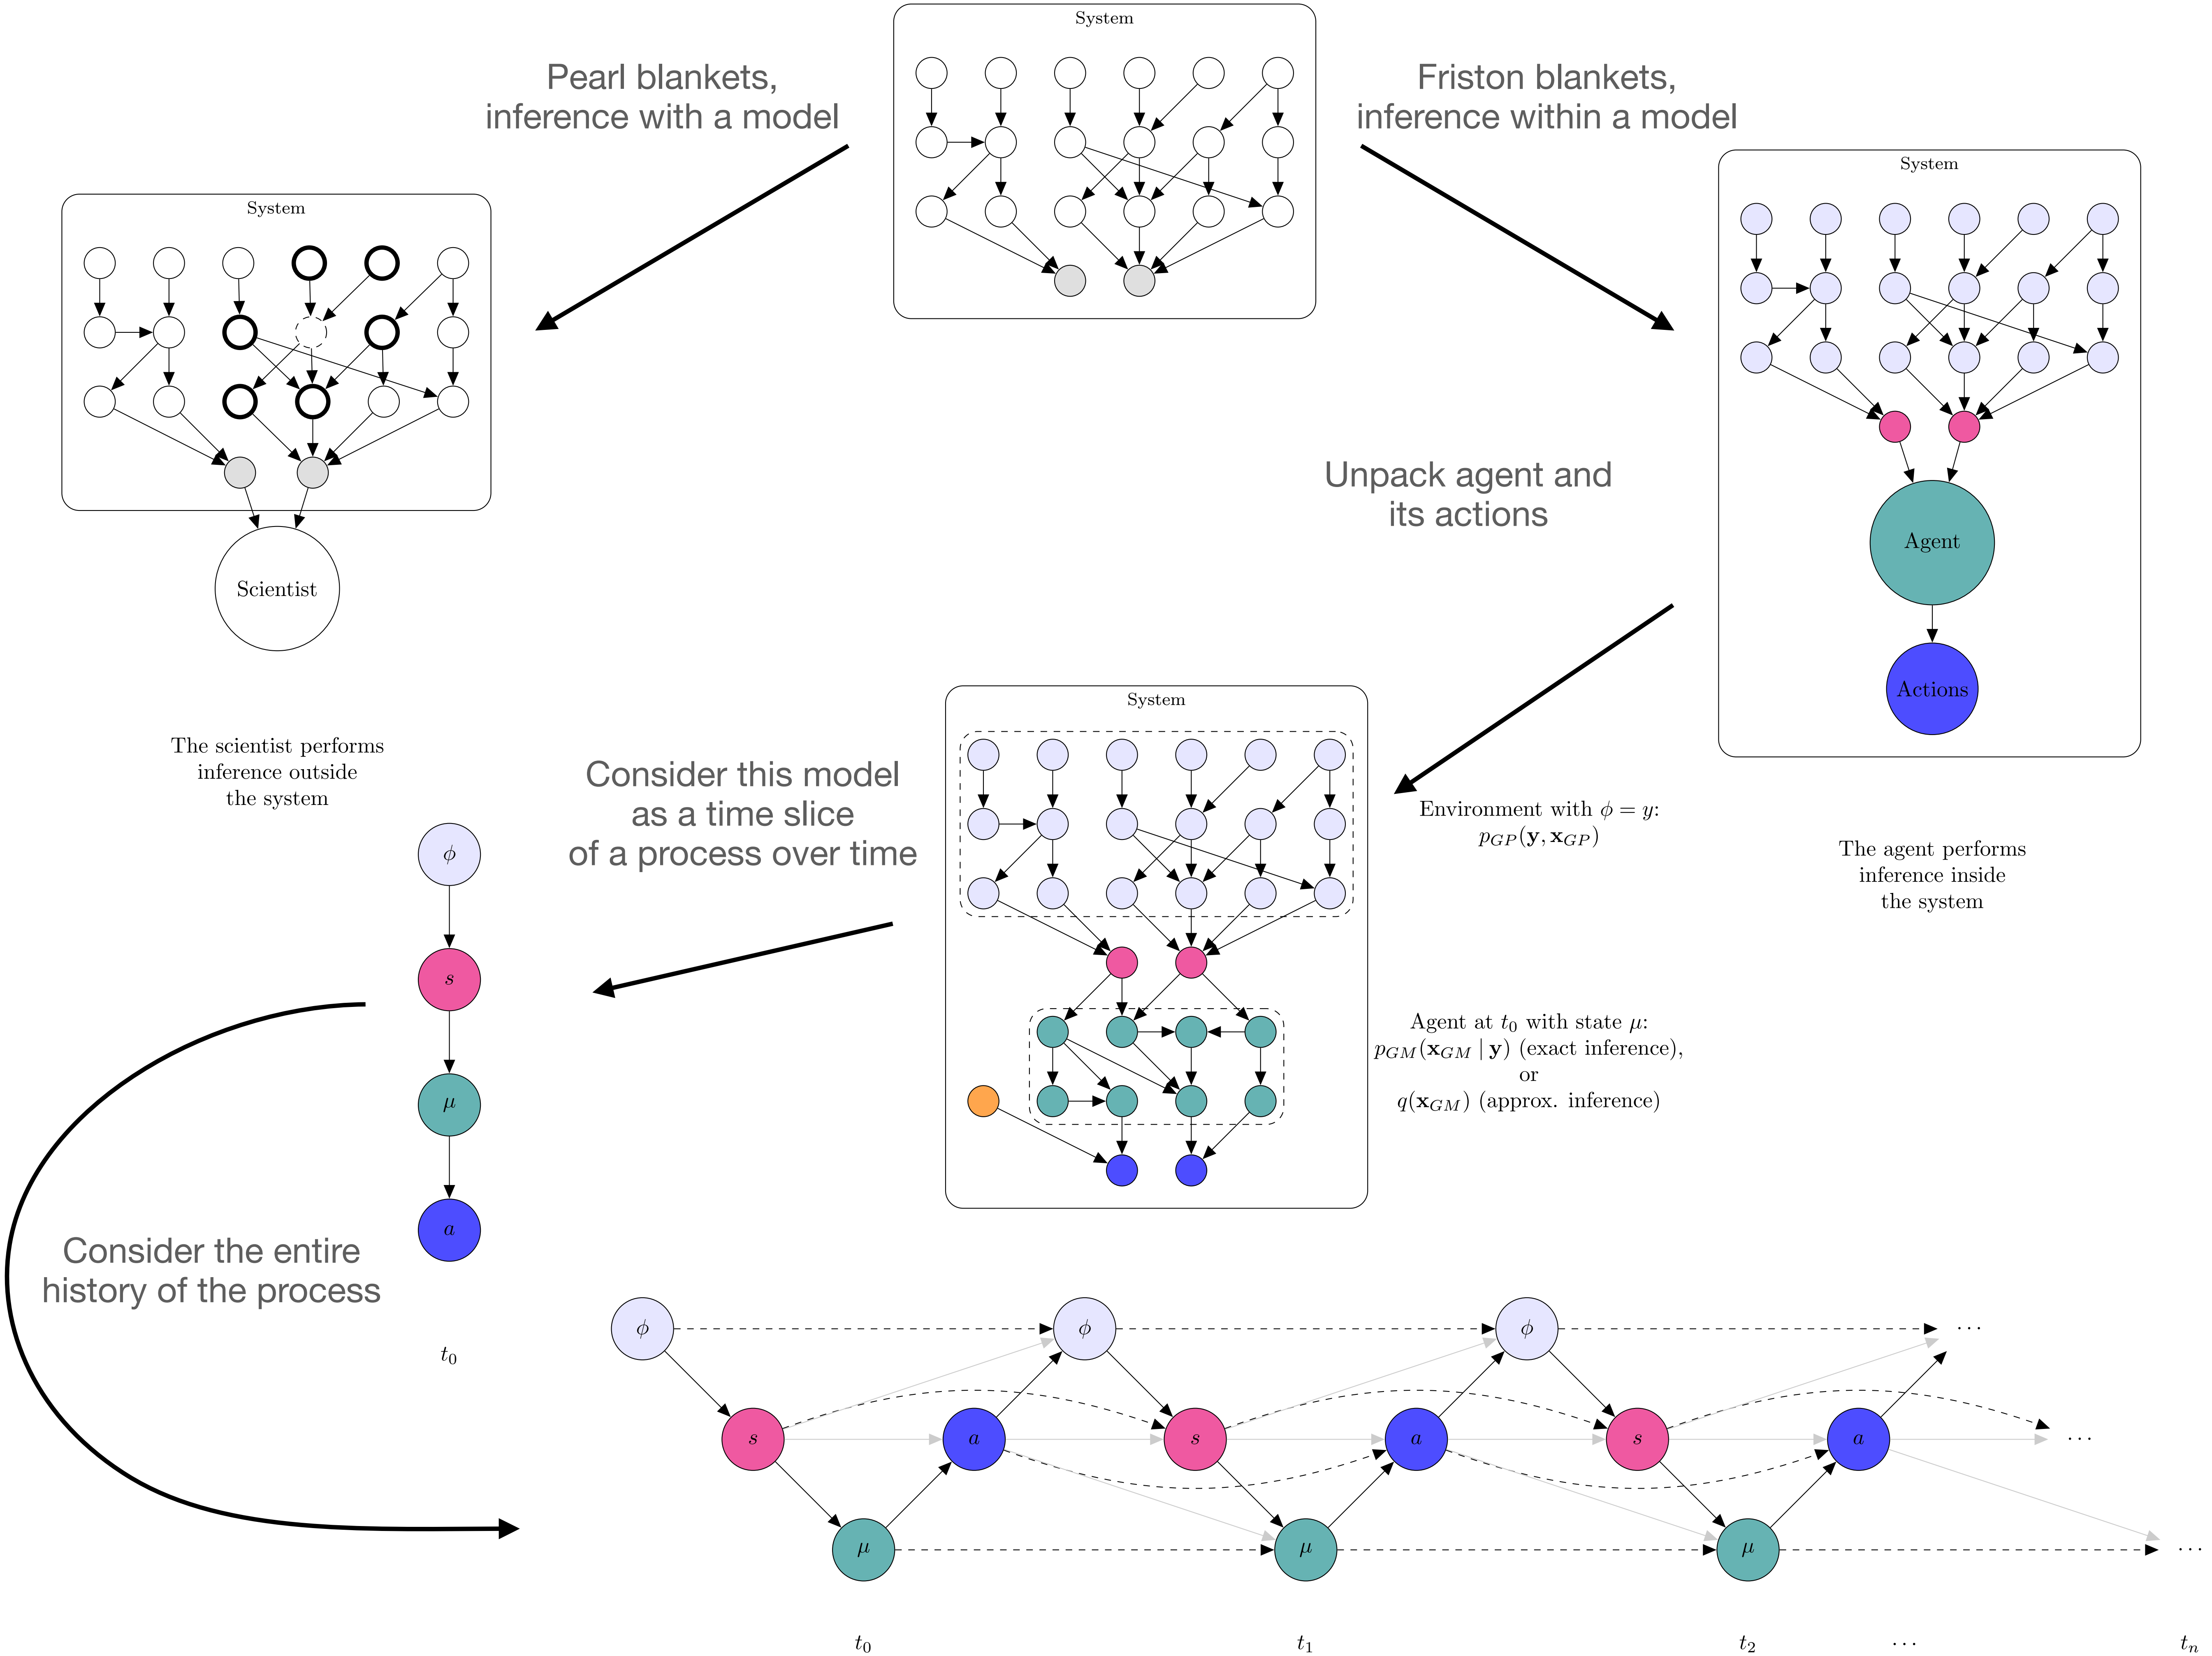
System
Pearl blankets,
inference with a model
Friston blankets,
inference within a model
System
System
Scientist
Unpack agent and
its actions
Agent
Actions
The scientist performs
inference outside
the system
The agent performs
inference inside
the system
Consider this model
as a time slice
of a process over time
System
Environment with $\phi = y$:
$p_{GP}(\mathbf{y}, \mathbf{x}_{GP})$
Agent at t_0 with state μ:
$p_{GM}(\mathbf{x}_{GM} \mid \mathbf{y})$ (exact inference),
or
$q(\mathbf{x}_{GM})$ (approx. inference)
ϕ
s
μ
a
t_0
Consider the entire
history of the process
ϕ s a μ ϕ s a μ ϕ s a μ $\cdots$
t_0
t_1
t_2
$\cdots$
t_n

Issues

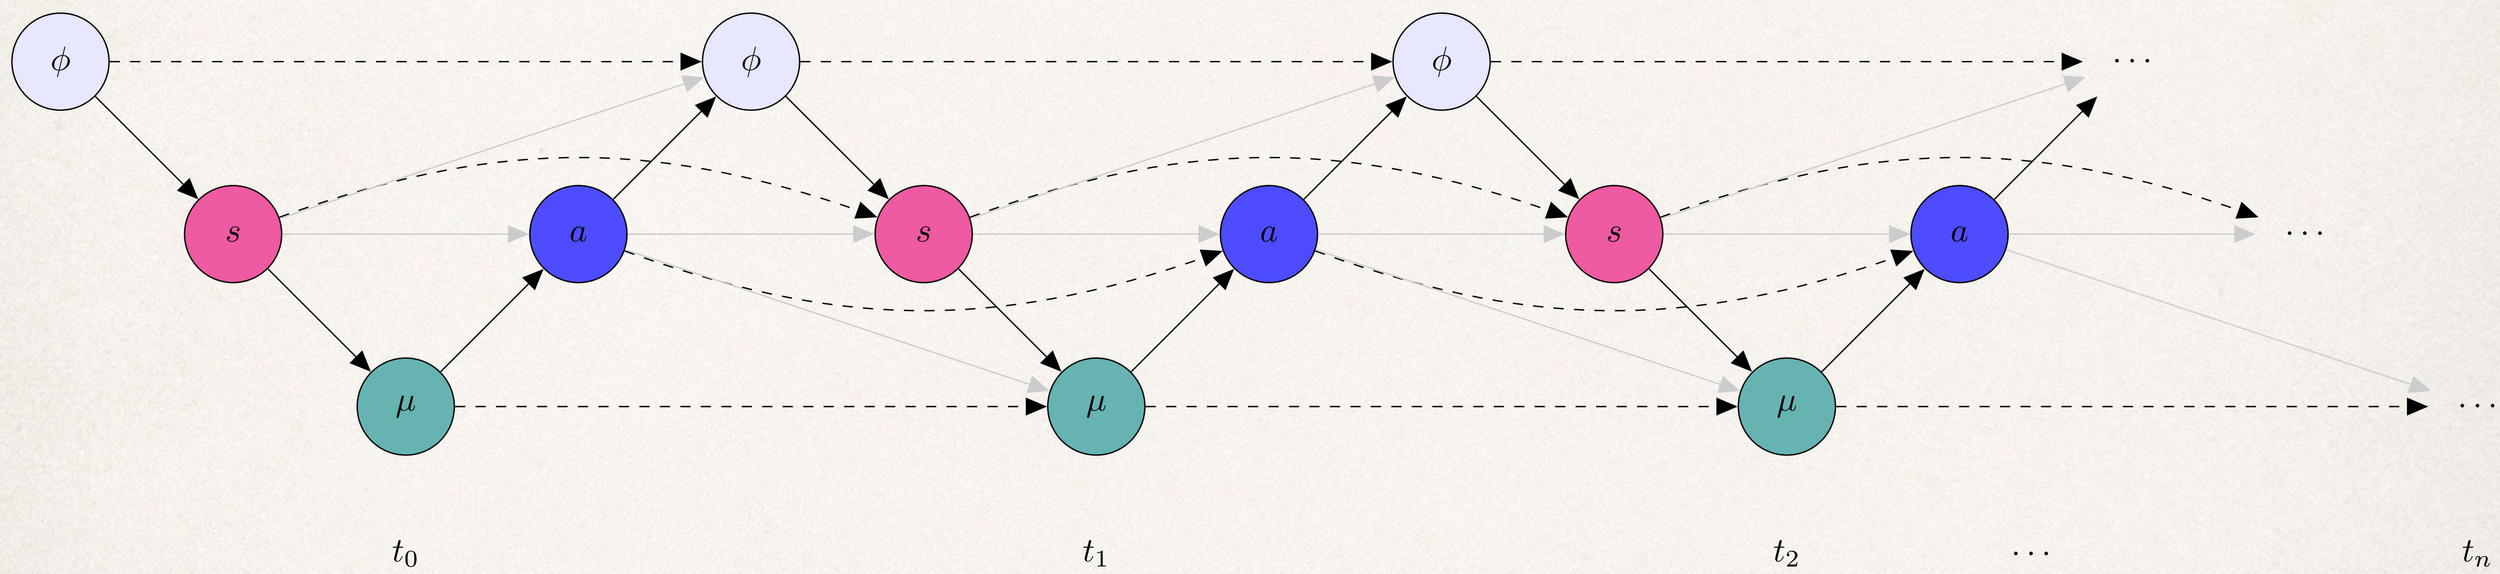

Assumption:

Stationarity of the stochastic process of interest (what's 'conditional independence' otherwise?)

Biehl M., and **Baltieri M.**. "The steady state Kalman filter and its Markov blanket." (In prep.)

Issues x2

- Thresholding of conditional (in)dependencies
- Initial identification of internal states outside of the framework
- Unclear relation between agents and partitions of stochastic process (e.g., role of co-parents)
- Ad-hoc sparsity constraints on non-equilibrium fluxes of steady-steady distribution
- …

Bruineberg J., Dolega K., Dewhurst J. and **Baltieri M.**, *"The Emperor's New Markov Blankets"*, (Accepted at Behavioral and Brain Sciences, IF > 17)

(All) Work in progress

- Context-dependent PID controllers (learning contexts)
- Kalman filters as variational inference (natural gradient) - with Takuya Isomura (RIKEN CBS, Japan)
- Steady-state Kalman filters and their Markov Blankets - with Martin Biehl (Araya Inc., Japan)
- A Bayesian classification of approximate models in psychophysics (based on a correct classification of uncertainties) - with Warrick Roseboom and Anil Seth (University of Sussex, UK)
- More models of whisking in mice - with Giovanni Pezzulo (CNR, Italy)
- Linear quadratic control (cont. time) vs. active inference + applications in neuroscience - with Christopher Buckley (University of Sussex)
- Detailed-balanced exploration in reinforcement learning - with Taro Toyoizumi
- Pytorch (—> JAX?) for continuous control

Summary

- (Approximate) Bayesian inference can be a powerful tool beyond generating accurate descriptions of data (building representations vs. controlling the world)
- This allows a connection to methods in classical control theory, providing a design framework where heuristics otherwise strive
- This also then ties into cognitive (neuro)science, helping articulating cognitive architectures (duality, the problem of dual control, separability, etc.)
- Applications of (A)BI to studies of origins of life (via non-equilibrium physics) are still largely work in progress

Acknowledgements

- Taro Toyoizumi and the lab
- Christopher L. Buckley
- JSPS for funding, Royal Society for selection

- Simon McGregor / Keisuke Suzuki / Anil Seth / Warrick Roseboom / Lionel Barnett (University of Sussex)
- Hideaki Shimazaki (University of Kyoto)
- Takashi Ikegami (University of Tokyo)
- Martin Biehl (Araya Inc.)
- Olaf Witkowsky, Nicholas Guttenberg (Cross Labs)
- Nathaniel Virgo (ELSI)

Contacts:

Email: manuel.baltieri@riken.jp

Twitter: @manuelbaltieri

Thank you

Questions?