Morphogenesis : Basal Cognition : : Self-Organisation : Maximum Entropy

Dalton A R Sakthivadivel

Stony Brook University; VERSES Research Lab

31st May 2022

Outline

Slides can be found later at darsakthi.github.io/talks

Based on 2203.08119, 2204.05084, 2205.11543

Inference and self-organisation

Modelling embedded systems that carry beliefs: a three-part problem

- Given a system in an agent-environment loop, how do we model the system as having beliefs?
- Given those beliefs, how do we model the system as applying them to self-organise?
- And what does that have to do with basal cognition?

Inference and self-organisation

Modelling embedded systems that carry beliefs: a three-part problem

- Given a system in an agent-environment loop, how do we model the system as having beliefs?
- Given those beliefs, how do we model the system as applying them to self-organise?
- And what does that have to do with basal cognition?

Inference and self-organisation

Modelling embedded systems that carry beliefs: a three-part problem

- Given a system in an agent-environment loop, how do we model the system as having beliefs?
- Given those beliefs, how do we model the system as applying them to self-organise?
- And what does that have to do with basal cognition?

Inference and self-organisation

Modelling embedded systems that carry beliefs: a three-part problem

- Given a system in an agent-environment loop, how do we model the system as having beliefs?
- Given those beliefs, how do we model the system as applying them to self-organise?
- And what does that have to do with basal cognition?

Inference and self-organisation

Modelling embedded systems that carry beliefs: a three-part problem

- Given a system in an agent-environment loop, how do we model the system as having beliefs?
- Given those beliefs, how do we model the system as applying them to self-organise?
- And what does that have to do with basal cognition?

In a phrase—how is *pattern formation* driven by the combination of inference over *self* (individuation) and *other* (cognition)?

Maximum entropy

What is inference?

- Given some data points $[x_1, x_2, \ldots, x_n]$, what is the nature of the process generating that data? In other words, how likely are we to observe some x_i under a given model?
- Usually followed by asking if $p(x_i)$ matches the (i) evidence that x_i is true, *or* (ii) the empirically observed likelihood of x_i
- Usually described by some loss function(al), i.e., the ideal classification $p(x)$ minimises or maximises some quantity

Maximum entropy

What is inference?

- Given some data points $[x_1, x_2, \ldots, x_n]$, what is the nature of the process generating that data? In other words, how likely are we to observe some x_i under a given model?
- Usually followed by asking if $p(x_i)$ matches the (i) evidence that x_i is true, *or* (ii) the empirically observed likelihood of x_i
- Usually described by some loss function(al), i.e., the ideal classification $p(x)$ minimises or maximises some quantity

Maximum entropy

What is inference?

- Given some data points $[x_1, x_2, \ldots, x_n]$, what is the nature of the process generating that data? In other words, how likely are we to observe some x_i under a given model?
- Usually followed by asking if $p(x_i)$ matches the (i) evidence that x_i is true, *or* (ii) the empirically observed likelihood of x_i
- Usually described by some loss function(al), i.e., the ideal classification $p(x)$ minimises or maximises some quantity

Maximum entropy

What is inference?

- Given some data points $[x_1, x_2, \ldots, x_n]$, what is the nature of the process generating that data? In other words, how likely are we to observe some x_i under a given model?
- Usually followed by asking if $p(x_i)$ matches the (i) evidence that x_i is true, *or* (ii) the empirically observed likelihood of x_i
- Usually described by some loss function(al), i.e., the ideal classification $p(x)$ minimises or maximises some quantity

Maximum entropy

What is inference?

- Given some data points $[x_1, x_2, \ldots, x_n]$, what is the nature of the process generating that data? In other words, how likely are we to observe some x_i under a given model?
- Usually followed by asking if $p(x_i)$ matches the (i) evidence that x_i is true, *or* (ii) the empirically observed likelihood of x_i
- Usually described by some loss function(al), i.e., the ideal classification $p(x)$ minimises or maximises some quantity

In general, inference is the *probabilistic reconstruction* of data

Maximum entropy

What is inference?

- Given some data points $[x_1, x_2, \ldots, x_n]$, what is the nature of the process generating that data? In other words, how likely are we to observe some x_i under a given model?
- Usually followed by asking if $p(x_i)$ matches the (i) evidence that x_i is true, *or* (ii) the empirically observed likelihood of x_i
- Usually described by some loss function(al), i.e., the ideal classification $p(x)$ minimises or maximises some quantity

In general, inference is the *probabilistic reconstruction* of data

Examples: MAP estimation, evidence lower bound, etc

Maximum entropy

Maximum entropy is a general method for inference, originally due to Jaynes

The maximum of the entropy functional is a probability density $p(x)$

Claim: (constrained) diffusion maximises (constrained) entropy

Given constraints $J(x)$ on the likelihood of states, $p(x)$ maximises entropy when it equals $\exp\{-\lambda J(x)\}$.

For computational implementations, see work by Ken Dill

Maximum entropy

Maximum entropy is a general method for inference, originally due to Jaynes

The maximum of the entropy functional is a probability density $p(x)$

Claim: (constrained) diffusion maximises (constrained) entropy

Given constraints $J(x)$ on the likelihood of states, $p(x)$ maximises entropy when it equals $\exp\{-\lambda J(x)\}$. Example: in a system where the penalty on x increases with the distance away from some central k, $p(x)$ ought to be symmetric about k and decrease away from it.

For computational implementations, see work by Ken Dill

Maximum entropy

Maximum entropy is a general method for inference, originally due to Jaynes

The maximum of the entropy functional is a probability density $p(x)$

Claim: (constrained) diffusion maximises (constrained) entropy

Given constraints $J(x)$ on the likelihood of states, $p(x)$ maximises entropy when it equals $\exp\{-\lambda J(x)\}$. Example: in a system where the penalty on x increases with the distance away from some central k, $p(x)$ ought to be symmetric about k and decrease away from it. Indeed, for $J(x) = (x-k)^2$, $\exp\{-\lambda(x-k)^2\}$ is a Gaussian density.

For computational implementations, see work by Ken Dill

Maximum entropy

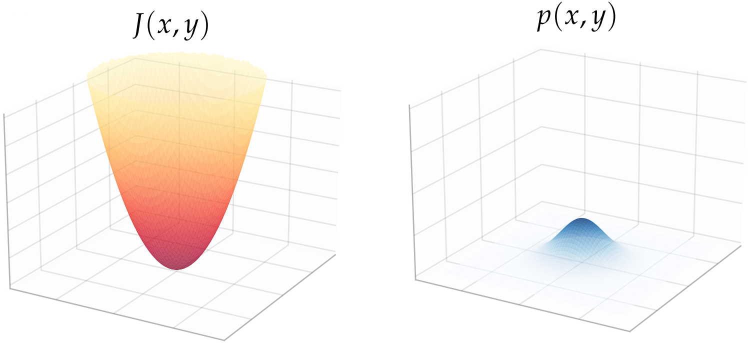

The probability of occupying a location is inversely proportional to the penalty on that location, such that $p(x, y) \propto \exp\{-J(x, y)\}$. Adapted from *On Bayesian Mechanics*. Credit to Brennan Klein.

Maximum entropy

Idea: entropy is like the spread of diffusion, or noisiness of a noise process. Unconstrained diffusion spreads out; maximises entropy

Maximum entropy

Idea: entropy is like the spread of diffusion, or noisiness of a noise process. Unconstrained diffusion spreads out; maximises entropy

When diffusion is constrained to localise around k, the diffusing system maximises entropy under a distance constraint

Maximum entropy

Idea: entropy is like the spread of diffusion, or noisiness of a noise process. Unconstrained diffusion spreads out; maximises entropy

When diffusion is constrained to localise around k, the diffusing system maximises entropy under a distance constraint

A body of diffusing particles can be modelled as performing inference against constraints, i.e., over its own probable locations in space

Maximum entropy

Idea: entropy is like the spread of diffusion, or noisiness of a noise process. Unconstrained diffusion spreads out; maximises entropy

When diffusion is constrained to localise around k, the diffusing system maximises entropy under a distance constraint

A body of diffusing particles can be modelled as performing inference against constraints, i.e., over its own probable locations in space

An ensemble of particles occupies locations in space according to constraints (penalties) on those locations. As a crowd, the entire mass of particles drifts towards k

A brief overview of the FEP

What is the free energy principle?

- One answer: a control systems perspective on self-organisation and allostasis
- The FEP dictates that a self-organising system controls its own set-points, like intended values of existential[1] or essential[2] variables
- In doing so, it remains independent of the environment
- The link between these two statements is *surprisal*

[1] Mel Andrews. The math is not the territory. *Biol Philos*. 2021.
[2] W Ross Ashby. Design for a Brain. 1952.

A brief overview of the FEP

What is the free energy principle?

- One answer: a control systems perspective on self-organisation and allostasis
- The FEP dictates that a self-organising system controls its own set-points, like intended values of existential[1] or essential[2] variables
- In doing so, it remains independent of the environment
- The link between these two statements is *surprisal*

[1] Mel Andrews. The math is not the territory. *Biol Philos*. 2021.
[2] W Ross Ashby. Design for a Brain. 1952.

A brief overview of the FEP

What is the free energy principle?

- One answer: a control systems perspective on self-organisation and allostasis
- The FEP dictates that a self-organising system controls its own set-points, like intended values of existential[1] or essential[2] variables
- In doing so, it remains independent of the environment
- The link between these two statements is *surprisal*

[1] Mel Andrews. The math is not the territory. *Biol Philos*. 2021.
[2] W Ross Ashby. Design for a Brain. 1952.

A brief overview of the FEP

What is the free energy principle?

- One answer: a control systems perspective on self-organisation and allostasis
- The FEP dictates that a self-organising system controls its own set-points, like intended values of existential[1] or essential[2] variables
- In doing so, it remains independent of the environment
- The link between these two statements is *surprisal*

[1] Mel Andrews. The math is not the territory. *Biol Philos*. 2021.
[2] W Ross Ashby. Design for a Brain. 1952.

A brief overview of the FEP

What is the free energy principle?

- One answer: a control systems perspective on self-organisation and allostasis
- The FEP dictates that a self-organising system controls its own set-points, like intended values of existential[1] or essential[2] variables
- In doing so, it remains independent of the environment
- The link between these two statements is *surprisal*

[1] Mel Andrews. The math is not the territory. *Biol Philos*. 2021.
[2] W Ross Ashby. Design for a Brain. 1952.

A brief overview of the FEP

Systems which individuate themselves and persist as individuals must both

- maintain a statistical boundary between agent and environment, and
- control the values of the existential variables definitional of the agent.

A brief overview of the FEP

Systems which individuate themselves and persist as individuals must both

- maintain a statistical boundary between agent and environment, and
- control the values of the existential variables definitional of the agent.

A brief overview of the FEP

Systems which individuate themselves and persist as individuals must both

- maintain a statistical boundary between agent and environment, and
- control the values of the existential variables definitional of the agent.

A brief overview of the FEP

Systems which individuate themselves and persist as individuals must both

- maintain a statistical boundary between agent and environment, and
- control the values of the existential variables definitional of the agent.

Control is equivalent to blanket maintenance:

A brief overview of the FEP

Systems which individuate themselves and persist as individuals must both

- maintain a statistical boundary between agent and environment, and
- control the values of the existential variables definitional of the agent.

Control is equivalent to blanket maintenance: a system remains distinguished from its environment when it maintains the set-points determining the 'sort of system that it is.'

A brief overview of the FEP

A 5 minute recipe for self-organisation: just add Markov blankets

- Partition a system into an agent μ, an environment η, and a blanket b separating the two
- The statistics of μ and b ought to be *conditionally independent* of η
- Define a coupling between agent and environment reaching across b, $\sigma(\hat{\mu}, b) = \hat{\eta}$
- By σ, the agent does inference over parameters of $p(\eta)$
- Two consequences:
  - The system's ideal set-point state encodes this belief, i.e., is $k = \sigma^{-1}(\hat{\eta}_b)$
  - When $\hat{\mu}_b = k$, the surprisal of blanket states is minimised; i.e., the system's boundary is maintained (see Prop 4.2, *Towards a Geometry and Analysis for Bayesian Mechanics* (DARS))

A brief overview of the FEP

A 5 minute recipe for self-organisation: just add Markov blankets

- Partition a system into an agent μ, an environment η, and a blanket b separating the two
- The statistics of μ and b ought to be *conditionally independent* of η
- Define a coupling between agent and environment reaching across b, $\sigma(\hat{\mu}, b) = \hat{\eta}$
- By σ, the agent does inference over parameters of $p(\eta)$
- Two consequences:
  - The system's ideal set-point state encodes this belief, i.e., is $k = \sigma^{-1}(\hat{\eta}_b)$
  - When $\hat{\mu}_b = k$, the surprisal of blanket states is minimised; i.e., the system's boundary is maintained (see Prop 4.2, *Towards a Geometry and Analysis for Bayesian Mechanics* (DARS))

A brief overview of the FEP

A 5 minute recipe for self-organisation: just add Markov blankets

- Partition a system into an agent μ, an environment η, and a blanket b separating the two
- The statistics of μ and b ought to be *conditionally independent* of η
- Define a coupling between agent and environment reaching across b, $\sigma(\hat{\mu}, b) = \hat{\eta}$
- By σ, the agent does inference over parameters of $p(\eta)$
- Two consequences:
  - The system's ideal set-point state encodes this belief, i.e., is $k = \sigma^{-1}(\hat{\eta}_b)$
  - When $\hat{\mu}_b = k$, the surprisal of blanket states is minimised; i.e., the system's boundary is maintained (see Prop 4.2, *Towards a Geometry and Analysis for Bayesian Mechanics* (DARS))

A brief overview of the FEP

A 5 minute recipe for self-organisation: just add Markov blankets

- Partition a system into an agent μ, an environment η, and a blanket b separating the two
- The statistics of μ and b ought to be *conditionally independent* of η
- Define a coupling between agent and environment reaching across b, $\sigma(\hat{\mu}, b) = \hat{\eta}$
- By σ, the agent does inference over parameters of $p(\eta)$
- Two consequences:
  - The system's ideal set-point state encodes this belief, i.e., is $k = \sigma^{-1}(\hat{\eta}_b)$
  - When $\hat{\mu}_b = k$, the surprisal of blanket states is minimised; i.e., the system's boundary is maintained (see Prop 4.2, *Towards a Geometry and Analysis for Bayesian Mechanics* (DARS))

A brief overview of the FEP

A 5 minute recipe for self-organisation: just add Markov blankets

- Partition a system into an agent μ, an environment η, and a blanket b separating the two
- The statistics of μ and b ought to be *conditionally independent* of η
- Define a coupling between agent and environment reaching across b, $\sigma(\hat{\mu}, b) = \hat{\eta}$
- By σ, the agent does inference over parameters of $p(\eta)$
- Two consequences:
  - The system's ideal set-point state encodes this belief, i.e., is $k = \sigma^{-1}(\hat{\eta}_b)$
  - When $\hat{\mu}_b = k$, the surprisal of blanket states is minimised; i.e., the system's boundary is maintained (see Prop 4.2, *Towards a Geometry and Analysis for Bayesian Mechanics* (DARS))

A brief overview of the FEP

A 5 minute recipe for self-organisation: just add Markov blankets

- Partition a system into an agent μ, an environment η, and a blanket b separating the two
- The statistics of μ and b ought to be *conditionally independent* of η
- Define a coupling between agent and environment reaching across b, $\sigma(\hat{\mu}, b) = \hat{\eta}$
- By σ, the agent does inference over parameters of $p(\eta)$
- Two consequences:
  - The system's ideal set-point state encodes this belief, i.e., is $k = \sigma^{-1}(\hat{\eta}_b)$
  - When $\hat{\mu}_b = k$, the surprisal of blanket states is minimised; i.e., the system's boundary is maintained (see Prop 4.2, *Towards a Geometry and Analysis for Bayesian Mechanics* (DARS))

A brief overview of the FEP

A 5 minute recipe for self-organisation: just add Markov blankets

- Partition a system into an agent μ, an environment η, and a blanket b separating the two
- The statistics of μ and b ought to be *conditionally independent* of η
- Define a coupling between agent and environment reaching across b, $\sigma(\hat{\mu}, b) = \hat{\eta}$
- By σ, the agent does inference over parameters of $p(\eta)$
- Two consequences:
  - The system's ideal set-point state encodes this belief, i.e., is $k = \sigma^{-1}(\hat{\eta}_b)$
  - When $\hat{\mu}_b = k$, the surprisal of blanket states is minimised; i.e., the system's boundary is maintained (see Prop 4.2, *Towards a Geometry and Analysis for Bayesian Mechanics* (DARS))

A brief overview of the FEP

A 5 minute recipe for self-organisation: just add Markov blankets

- Partition a system into an agent μ, an environment η, and a blanket b separating the two
- The statistics of μ and b ought to be *conditionally independent* of η
- Define a coupling between agent and environment reaching across b, $\sigma(\hat{\mu}, b) = \hat{\eta}$
- By σ, the agent does inference over parameters of $p(\eta)$
- Two consequences:
  - The system's ideal set-point state encodes this belief, i.e., is $k = \sigma^{-1}(\hat{\eta}_b)$
  - When $\hat{\mu}_b = k$, the surprisal of blanket states is minimised; i.e., the system's boundary is maintained (see Prop 4.2, *Towards a Geometry and Analysis for Bayesian Mechanics* (DARS))

A brief overview of the FEP

To summarise: the internal states of a system reflect probability densities over likely external states; we can read this as inference

- An agent will occupy states given a particular regime of external states
- We can understand such systems as encoding models of their environments, either stored representations or trivial statistical models, by being the preimage of some σ
- Ultimately a statement that controlled systems embody responses to perturbations from an environment
- Likewise, the environment leaves a trace on the 'self,' and the self is informed by the environment

A brief overview of the FEP

To summarise: the internal states of a system reflect probability densities over likely external states; we can read this as inference

- An agent will occupy states given a particular regime of external states
- We can understand such systems as encoding models of their environments, either stored representations or trivial statistical models, by being the preimage of some σ
- Ultimately a statement that controlled systems embody responses to perturbations from an environment
- Likewise, the environment leaves a trace on the 'self,' and the self is informed by the environment

A brief overview of the FEP

To summarise: the internal states of a system reflect probability densities over likely external states; we can read this as inference

- An agent will occupy states given a particular regime of external states
- We can understand such systems as encoding models of their environments, either stored representations or trivial statistical models, by being the preimage of some σ
- Ultimately a statement that controlled systems embody responses to perturbations from an environment
- Likewise, the environment leaves a trace on the ‘self,’ and the self is informed by the environment

A brief overview of the FEP

To summarise: the internal states of a system reflect probability densities over likely external states; we can read this as inference

- An agent will occupy states given a particular regime of external states
- We can understand such systems as encoding models of their environments, either stored representations or trivial statistical models, by being the preimage of some σ
- Ultimately a statement that controlled systems embody responses to perturbations from an environment
- Likewise, the environment leaves a trace on the 'self,' and the self is informed by the environment

A brief overview of the FEP

To summarise: the internal states of a system reflect probability densities over likely external states; we can read this as inference

- An agent will occupy states given a particular regime of external states
- We can understand such systems as encoding models of their environments, either stored representations or trivial statistical models, by being the preimage of some σ
- Ultimately a statement that controlled systems embody responses to perturbations from an environment
- Likewise, the environment leaves a trace on the 'self,' and the self is informed by the environment

A brief overview of the FEP

To summarise: the internal states of a system reflect probability densities over likely external states; we can read this as inference

- An agent will occupy states given a particular regime of external states
- We can understand such systems as encoding models of their environments, either stored representations or trivial statistical models, by being the preimage of some σ
- Ultimately a statement that controlled systems embody responses to perturbations from an environment
- Likewise, the environment leaves a trace on the 'self,' and the self is informed by the environment

See Example 4.1 in *Towards a Geometry and Analysis for Bayesian Mechanics* (DARS).

A brief overview of the FEP

A key takeaway:

A brief overview of the FEP

A key takeaway: any notion of organisation (what the system wants to be) lives in the precise definition of σ,

A brief overview of the FEP

A key takeaway: any notion of organisation (what the system wants to be) lives in the precise definition of σ, and *self*-organisation is a system constraining it*self* to occupy states around $k = \sigma^{-1}(\hat{\eta}_b)$

A brief overview of the FEP

A key takeaway: any notion of organisation (what the system wants to be) lives in the precise definition of σ, and *self*-organisation is a system constraining it*self* to occupy states around $k = \sigma^{-1}(\hat{\eta}_b)$

Under the FEP, σ also makes systems do inference over environmental states η

A brief overview of the FEP

A key takeaway: any notion of organisation (what the system wants to be) lives in the precise definition of σ, and *self*-organisation is a system constraining it*self* to occupy states around $k = \sigma^{-1}(\hat{\eta}_b)$

Under the FEP, σ also makes systems do inference over environmental states η

So, systems do inference over what their environment is, in order to understand how to get to a characteristic state

FEP ⊣ CMEP

The FEP is about a system's beliefs about the environment.

Can we think of the FEP as a statement that systems constrain their states around optimal values k of certain control parameters?

FEP ⊣ CMEP

The FEP is about a system's beliefs about the environment. What is the mirror image of the FEP, modelling *our* beliefs about a system, or a system's beliefs *about itself*?

Can we think of the FEP as a statement that systems constrain their states around optimal values k of certain control parameters?

FEP $\dashv$ CMEP

The FEP is about a system's beliefs about the environment. What is the mirror image of the FEP, modelling *our* beliefs about a system, or a system's beliefs *about itself*? Relational symmetry

Can we think of the FEP as a statement that systems constrain their states around optimal values k of certain control parameters?

FEP ⊣ CMEP

The FEP is about a system's beliefs about the environment. What is the mirror image of the FEP, modelling *our* beliefs about a system, or a system's beliefs *about itself*? Relational symmetry

Can we think of the FEP as a statement that systems constrain their states around optimal values k of certain control parameters? Morphological constraints

FEP $\dashv$ CMEP

We can formulate the FEP as the mirror image to maximum entropy

- What is the likelihood of a system's states given some constraints on those states, like distance from some k?

FEP ⊣ CMEP

We can formulate the FEP as the mirror image to maximum entropy

- What is the likelihood of a system's states given some constraints on those states, like distance from some k?

This allows us to apply the FEP to model systems based on the regimes of states they occupy, rather than the likelihoods of external states parameterised by those states

FEP $\dashv$ CMEP

We can formulate the FEP as the mirror image to maximum entropy

- What is the likelihood of a system's states given some constraints on those states, like distance from some k?

This allows us to apply the FEP to model systems based on the regimes of states they occupy, rather than the likelihoods of external states parameterised by those states

At a higher level: we can formulate the FEP (systems maintain their 'selves' by knowing their environment) equivalently as constrained maximum entropy (we can know systems as maintaining their 'selves')

Morphological constraints

In *Towards a Geometry and Analysis for Bayesian Mechanics* (DARS) and *On Bayesian Mechanics* (Maxwell Ramstead, DARS, et al), the set of constraints on what states a system occupies is referred to as an ontological potential, ontologically definitional of what the system is

Morphological constraints

We could also call these morphological constraints.

Morphological constraints

We could also call these morphological constraints. Suppose there exists some method for encoding and enforcing preferences about states.

Morphological constraints

We could also call these morphological constraints. Suppose there exists some method for encoding and enforcing preferences about states.

- The encoding becomes a particular functional form for σ, and the enforcing is a penalty $(\mu_b - k)^2$
- This encoding affects the development of the system, such that the identity of the system is contingent on those preferences being expressed (i.e., $\hat{\mu}_b = k$)
- This morphological constraint is an ontological potential, and it leads to a system occupying states that parameterise beliefs about external states.

Morphological constraints

We could also call these morphological constraints. Suppose there exists some method for encoding and enforcing preferences about states.

- The encoding becomes a particular functional form for σ, and the enforcing is a penalty $(\mu_b - k)^2$
- This encoding affects the development of the system, such that the identity of the system is contingent on those preferences being expressed (i.e., $\hat{\mu}_b = k$)
- This morphological constraint is an ontological potential, and it leads to a system occupying states that parameterise beliefs about external states.

Morphological constraints

We could also call these morphological constraints. Suppose there exists some method for encoding and enforcing preferences about states.

- The encoding becomes a particular functional form for σ, and the enforcing is a penalty $(\mu_b - k)^2$
- This encoding affects the development of the system, such that the identity of the system is contingent on those preferences being expressed (i.e., $\hat{\mu}_b = k$)
- This morphological constraint is an ontological potential, and it leads to a system occupying states that parameterise beliefs about external states.

Example: cell differentiation

In the example of cell differentiation, certain cell structures are characteristic of neurones that are not characteristic of epithelial cells.

Example: cell differentiation

In the example of cell differentiation, certain cell structures are characteristic of neurones that are not characteristic of epithelial cells. (For example, long axons and dendrites.)

Example: cell differentiation

In the example of cell differentiation, certain cell structures are characteristic of neurones that are not characteristic of epithelial cells. (For example, long axons and dendrites.)

These unique states are definitional of a neurone. The phenotype of the cell is the expression of a *genetic* constraint on states.

Example: cell differentiation

In the example of cell differentiation, certain cell structures are characteristic of neurones that are not characteristic of epithelial cells. (For example, long axons and dendrites.)

These unique states are definitional of a neurone. The phenotype of the cell is the expression of a *genetic* constraint on states.

This state also parameterises beliefs about neighbouring cells, since neural cells form communicative networks.

Example: cell differentiation

In the example of cell differentiation, certain cell structures are characteristic of neurones that are not characteristic of epithelial cells. (For example, long axons and dendrites.)

These unique states are definitional of a neurone. The phenotype of the cell is the expression of a *genetic* constraint on states.

This state also parameterises beliefs about neighbouring cells, since neural cells form communicative networks.

Hence, inference over the self (what are my morphological constraints k?)

Example: cell differentiation

In the example of cell differentiation, certain cell structures are characteristic of neurones that are not characteristic of epithelial cells. (For example, long axons and dendrites.)

These unique states are definitional of a neurone. The phenotype of the cell is the expression of a *genetic* constraint on states.

This state also parameterises beliefs about neighbouring cells, since neural cells form communicative networks.

Hence, inference over the self (what are my morphological constraints k?) is dual to

Example: cell differentiation

In the example of cell differentiation, certain cell structures are characteristic of neurones that are not characteristic of epithelial cells. (For example, long axons and dendrites.)

These unique states are definitional of a neurone. The phenotype of the cell is the expression of a *genetic* constraint on states.

This state also parameterises beliefs about neighbouring cells, since neural cells form communicative networks.

Hence, inference over the self (what are my morphological constraints k?) is dual to inference over others (is there signalling happening around me? Do I have an opportunity to form a synapse?)

Example: cell differentiation

In the example of cell differentiation, certain cell structures are characteristic of neurones that are not characteristic of epithelial cells. (For example, long axons and dendrites.)

These unique states are definitional of a neurone. The phenotype of the cell is the expression of a *genetic* constraint on states.

This state also parameterises beliefs about neighbouring cells, since neural cells form communicative networks.

Hence, inference over the self (what are my morphological constraints k?) is dual to inference over others (is there signalling happening around me? Do I have an opportunity to form a synapse?)

And both lead to self-organisation

Example II: plant germination

What makes a seed turn into a stem with leaves?

Example II: plant germination

What makes a seed turn into a stem with leaves? (Disclaimer: the author is not a plant biologist.)

Example II: plant germination

What makes a seed turn into a stem with leaves? (Disclaimer: the author is not a plant biologist.)

Beliefs about the state of the environment: sun exposure, temperature, etc, are parameterised by the state of sprouting.

Example II: plant germination

What makes a seed turn into a stem with leaves? (Disclaimer: the author is not a plant biologist.)

Beliefs about the state of the environment: sun exposure, temperature, etc, are parameterised by the state of sprouting. An example of a false belief, based on limited data or an imperfect model: sprouting just prior to a cold snap.

Example II: plant germination

What makes a seed turn into a stem with leaves? (Disclaimer: the author is not a plant biologist.)

Beliefs about the state of the environment: sun exposure, temperature, etc, are parameterised by the state of sprouting. An example of a false belief, based on limited data or an imperfect model: sprouting just prior to a cold snap.

A morphological constraint exists that plant systems grow leaves to feed off of sunlight.

Example II: plant germination

What makes a seed turn into a stem with leaves? (Disclaimer: the author is not a plant biologist.)

Beliefs about the state of the environment: sun exposure, temperature, etc, are parameterised by the state of sprouting. An example of a false belief, based on limited data or an imperfect model: sprouting just prior to a cold snap.

A morphological constraint exists that plant systems grow leaves to feed off of sunlight.

Hence, inference over the self (is it time to transition to photosynthesis, and thus, to become a plant?)

Example II: plant germination

What makes a seed turn into a stem with leaves? (Disclaimer: the author is not a plant biologist.)

Beliefs about the state of the environment: sun exposure, temperature, etc, are parameterised by the state of sprouting. An example of a false belief, based on limited data or an imperfect model: sprouting just prior to a cold snap.

A morphological constraint exists that plant systems grow leaves to feed off of sunlight.

Hence, inference over the self (is it time to transition to photosynthesis, and thus, to become a plant?) is dual to

Example II: plant germination

What makes a seed turn into a stem with leaves? (Disclaimer: the author is not a plant biologist.)

Beliefs about the state of the environment: sun exposure, temperature, etc, are parameterised by the state of sprouting. An example of a false belief, based on limited data or an imperfect model: sprouting just prior to a cold snap.

A morphological constraint exists that plant systems grow leaves to feed off of sunlight.

Hence, inference over the self (is it time to transition to photosynthesis, and thus, to become a plant?) is dual to inference over others (how much sun must be available, given how warm the soil is?)

Conclusion: mind everywhere

In both examples, systems with no brains appeared to perform inference

Using max ent, we can even say that diffusing particles perform inference over k

Where is the line between cognitive and non-cognitive complexity?

Conclusion: mind everywhere

In both examples, systems with no brains appeared to perform inference

Using max ent, we can even say that diffusing particles perform inference over k

Where is the line between cognitive and non-cognitive complexity? Under the FEP and its mirror image, max ent, no solid line can exist

Gauge theory, intuitively

Gauge forces are closely related to the intuitive idea of a physical force: a free particle goes through space-time on flat paths, whilst a particle undergoing a force curves.

One of the ways we can model the FEP is as a Helmholtz decomposition, which is exactly a decomposition into flat paths and curved paths

Helmholtz decomposition

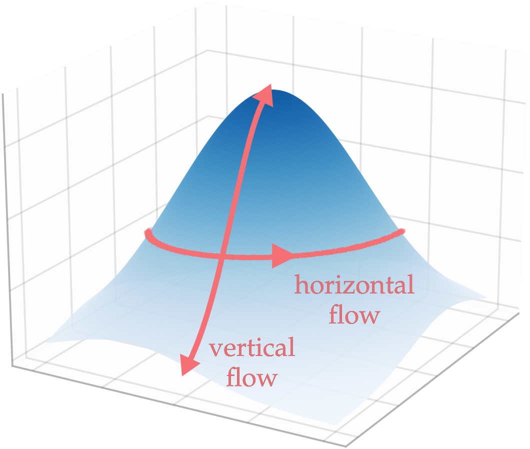

Adapted from *On Bayesian Mechanics*. Credit to Brennan Klein.

Bundles of tricks

Is self-organisation just a realisation of a system following a path of least action (a geodesic) under a gauge force?

- Molecules are not intelligent, but at the ensemble level, 'do inference' against some set of constraints
- Gene expression in developing cells flows on Waddington's landscape
- Proteins fold by travelling down funnels[3] in an energy landscape
- Natural selection flows species towards phenotypic optima under perturbation, over long time-scales

[3] Ken A Dill and Justin L MacCallum. The protein-folding problem, 50 years on. *Science*. 2012.

Bundles of tricks

Is self-organisation just a realisation of a system following a path of least action (a geodesic) under a gauge force?

Examples:

- Molecules are not intelligent, but at the ensemble level, 'do inference' against some set of constraints
- Gene expression in developing cells flows on Waddington's landscape
- Proteins fold by travelling down funnels[3] in an energy landscape
- Natural selection flows species towards phenotypic optima under perturbation, over long time-scales

[3] Ken A Dill and Justin L MacCallum. The protein-folding problem, 50 years on. *Science*. 2012.

Bundles of tricks

Is self-organisation just a realisation of a system following a path of least action (a geodesic) under a gauge force?

Examples:

- Molecules are not intelligent, but at the ensemble level, 'do inference' against some set of constraints
- Gene expression in developing cells flows on Waddington's landscape
- Proteins fold by travelling down funnels[3] in an energy landscape
- Natural selection flows species towards phenotypic optima under perturbation, over long time-scales

[3] Ken A Dill and Justin L MacCallum. The protein-folding problem, 50 years on. *Science*. 2012.

Bundles of tricks

Is self-organisation just a realisation of a system following a path of least action (a geodesic) under a gauge force?

Examples:

- Molecules are not intelligent, but at the ensemble level, ‘do inference’ against some set of constraints
- Gene expression in developing cells flows on Waddington’s landscape
- Proteins fold by travelling down funnels[3] in an energy landscape
- Natural selection flows species towards phenotypic optima under perturbation, over long time-scales

[3] Ken A Dill and Justin L MacCallum. The protein-folding problem, 50 years on. *Science*. 2012.

Bundles of tricks

Is self-organisation just a realisation of a system following a path of least action (a geodesic) under a gauge force?

Examples:

- Molecules are not intelligent, but at the ensemble level, 'do inference' against some set of constraints
- Gene expression in developing cells flows on Waddington's landscape
- Proteins fold by travelling down funnels[3] in an energy landscape
- Natural selection flows species towards phenotypic optima under perturbation, over long time-scales

[3] Ken A Dill and Justin L MacCallum. The protein-folding problem, 50 years on. *Science*. 2012.

Bundles of tricks

Is self-organisation just a realisation of a system following a path of least action (a geodesic) under a gauge force?

Examples:

- Molecules are not intelligent, but at the ensemble level, 'do inference' against some set of constraints
- Gene expression in developing cells flows on Waddington's landscape
- Proteins fold by travelling down funnels[3] in an energy landscape
- Natural selection flows species towards phenotypic optima under perturbation, over long time-scales

[3] Ken A Dill and Justin L MacCallum. The protein-folding problem, 50 years on. *Science*. 2012.

Bundles of tricks

This is called a 'bag of tricks' view of complexity: lots of unintelligent things are intelligent together, or in expectation, or at a higher scale;

Bundles of tricks

This is called a ‘bag of tricks’ view of complexity: lots of unintelligent things are intelligent together, or in expectation, or at a higher scale; conversely, complexity arises from some use-case applying downwards constraints to a bunch of simpler things

Bundles of tricks

This is called a ‘bag of tricks’ view of complexity: lots of unintelligent things are intelligent together, or in expectation, or at a higher scale; conversely, complexity arises from some use-case applying downwards constraints to a bunch of simpler things

In this view, complexity develops from simple tools

Bundles of tricks

This is called a ‘bag of tricks’ view of complexity: lots of unintelligent things are intelligent together, or in expectation, or at a higher scale; conversely, complexity arises from some use-case applying downwards constraints to a bunch of simpler things

In this view, complexity develops from simple tools

Does the fact that minds emerge everywhere come from the fact that most macroscopic or condensed matter systems are actually multiscale systems,

Bundles of tricks

This is called a ‘bag of tricks’ view of complexity: lots of unintelligent things are intelligent together, or in expectation, or at a higher scale; conversely, complexity arises from some use-case applying downwards constraints to a bunch of simpler things

In this view, complexity develops from simple tools

Does the fact that minds emerge everywhere come from the fact that most macroscopic or condensed matter systems are actually multiscale systems, whose lower-level dynamics follow geodesics shaped by some organising principle

Bundles of tricks

This is called a 'bag of tricks' view of complexity: lots of unintelligent things are intelligent together, or in expectation, or at a higher scale; conversely, complexity arises from some use-case applying downwards constraints to a bunch of simpler things

In this view, complexity develops from simple tools

Does the fact that minds emerge everywhere come from the fact that most macroscopic or condensed matter systems are actually multiscale systems, whose lower-level dynamics follow geodesics shaped by some organising principle (the existence of an ensemble constraint or control parameter, for instance)?

Bundles of tricks

Is a gauge force for ensemble-level self-organisation a way to speak about downward causation?[4]

[4] Jessica C Flack. Coarse-graining as a downward causation mechanism. *Phil Trans R Soc A*. 2017.

Bundles of tricks

Is a gauge force for ensemble-level self-organisation a way to speak about downward causation?[4]

Gauge forces are often referred to as illusory forces, since there appears to be a force pulling things, which is actually just a curvature (here, of a system towards a mode, k)

[4] Jessica C Flack. Coarse-graining as a downward causation mechanism. *Phil Trans R Soc A*. 2017.

Bundles of tricks

Is a gauge force for ensemble-level self-organisation a way to speak about downward causation?[4]

Gauge forces are often referred to as illusory forces, since there appears to be a force pulling things, which is actually just a curvature (here, of a system towards a mode, k)

Is the illusory or mystical component of things like natural selection, cognition, etc, theoretically explainable by gauge theory?

[4] Jessica C Flack. Coarse-graining as a downward causation mechanism. *Phil Trans R Soc A*. 2017.

Gauge bosons

What carries gauge forces? Particles called gauge bosons

If the tendency to do inference is a force driving the system towards a mode, is there a gauge boson for self-organisation?

Gauge bosons

What carries gauge forces? Particles called gauge bosons

If the tendency to do inference is a force driving the system towards a mode, is there a gauge boson for self-organisation?

Conjecture: we can model the actions a controlled system takes as a topological excitation, of the sorts that appear in gauge theories

Gauge bosons

What carries gauge forces? Particles called gauge bosons

If the tendency to do inference is a force driving the system towards a mode, is there a gauge boson for self-organisation?

Conjecture: we can model the actions a controlled system takes as a topological excitation, of the sorts that appear in gauge theories

Preliminary evidence: neural action potentials can be modelled[5] as topological excitations...

[5]Hillel Ori et al. Observation of topological action potentials in engineered tissues. 2022.

Gauge bosons

What carries gauge forces? Particles called gauge bosons

If the tendency to do inference is a force driving the system towards a mode, is there a gauge boson for self-organisation?

Conjecture: we can model the actions a controlled system takes as a topological excitation, of the sorts that appear in gauge theories

Preliminary evidence: neural action potentials can be modelled[5] as topological excitations. . .
. . . which are how we model certain particles, called instantons, in quantum gauge field theory

[5] Hillel Ori et al. Observation of topological action potentials in engineered tissues. 2022.

Some further questions

Some questions surrounding morphological constraints:

- How are morphological constraints generated?
  - As stored target morphologies? If so, where?
  - From a 'bundle of tricks?' If so, how could we test this empirically?
- How are Markov blankets generated and maintained?
  - What is the physical signature of individuation in complex systems?
  - How can they be engineered by an experimenter?
- Can we find evidence of MCs and MBs in nature?
  - What empirical methods will reveal the presence of morphological constraints or Markov blankets?

Some further questions

Some questions surrounding morphological constraints:

- How are morphological constraints generated?
  - As stored target morphologies? If so, where?
  - From a ‘bundle of tricks?’ If so, how could we test this empirically?
- How are Markov blankets generated and maintained?
  - What is the physical signature of individuation in complex systems?
  - How can they be engineered by an experimenter?
- Can we find evidence of MCs and MBs in nature?
  - What empirical methods will reveal the presence of morphological constraints or Markov blankets?

Some further questions

Some questions surrounding morphological constraints:

- How are morphological constraints generated?
  - As stored target morphologies? If so, where?
  - From a 'bundle of tricks?' If so, how could we test this empirically?
- How are Markov blankets generated and maintained?
  - What is the physical signature of individuation in complex systems?
  - How can they be engineered by an experimenter?
- Can we find evidence of MCs and MBs in nature?
  - What empirical methods will reveal the presence of morphological constraints or Markov blankets?

Some further questions

Some questions surrounding morphological constraints:

- How are morphological constraints generated?
  - As stored target morphologies? If so, where?
  - From a 'bundle of tricks?' If so, how could we test this empirically?
- How are Markov blankets generated and maintained?
  - What is the physical signature of individuation in complex systems?
  - How can they be engineered by an experimenter?
- Can we find evidence of MCs and MBs in nature?
  - What empirical methods will reveal the presence of morphological constraints or Markov blankets?

Writing self-organisation under the FEP as a problem of constrained maximum entropy, statements about

- Degrees of cognition
- Potential ways of viewing morphogenesis
- The physics of morphogenesis

can be made

Writing self-organisation under the FEP as a problem of constrained maximum entropy, statements about

- Degrees of cognition
- Potential ways of viewing morphogenesis
- The physics of morphogenesis

can be made

Writing self-organisation under the FEP as a problem of constrained maximum entropy, statements about

- Degrees of cognition
- Potential ways of viewing morphogenesis
- The physics of morphogenesis

can be made

Writing self-organisation under the FEP as a problem of constrained maximum entropy, statements about

- Degrees of cognition
- Potential ways of viewing morphogenesis
- The physics of morphogenesis

can be made

Writing self-organisation under the FEP as a problem of constrained maximum entropy, statements about

- Degrees of cognition
- Potential ways of viewing morphogenesis
- The physics of morphogenesis

can be made

Many opportunities to evaluate these potential connections empirically