Remarks on the mathematics of the free energy principle (especially with regard to path entropy and free action)

Dalton A R Sakthivadivel

16 October 2025

Remarks

These slides can be found later at darsakthi.github.io/talks

Some useful references will be noted on the final slide

We begin from a system of two coupled random variables evolving in time separated by a boundary,

$$X_t \xrightarrow{g} B_t \xrightarrow{h} Y_t$$

assumed to satisfy Itō SDEs

$$\begin{aligned}
\mathrm{d}X_t &= f_1(X_t, B_t, t)\,\mathrm{d}t + D_1(X_t, B_t, t)\,\mathrm{d}W_t^1 \\
\mathrm{d}B_t &= f_2(X_t, B_t, Y_t, t)\,\mathrm{d}t + D_2(X_t, B_t, Y_t, t)\,\mathrm{d}W_t^2 \\
\mathrm{d}Y_t &= f_3(B_t, Y_t, t)\,\mathrm{d}t + D_3(B_t, Y_t, t)\,\mathrm{d}W_t^3
\end{aligned}$$

The precise coupling structure is specific to a given system and defines different classes of dynamics [Friston et al, 2023].

We begin from a system of two coupled random variables evolving in time separated by a boundary,

$$X_t \xrightarrow{g} B_t \xrightarrow{h} Y_t$$

assumed to satisfy Itō SDEs

$$\begin{aligned}
\mathrm{d}X_t &= f_1(X_t, B_t, t)\,\mathrm{d}t + D_1(X_t, B_t, t)\,\mathrm{d}W_t^1 \\
\mathrm{d}B_t &= f_2(X_t, B_t, Y_t, t)\,\mathrm{d}t + D_2(X_t, B_t, Y_t, t)\,\mathrm{d}W_t^2 \\
\mathrm{d}Y_t &= f_3(B_t, Y_t, t)\,\mathrm{d}t + D_3(B_t, Y_t, t)\,\mathrm{d}W_t^3
\end{aligned}$$

The precise coupling structure is specific to a given system and defines different classes of dynamics [Friston et al, 2023].

We begin from a system of two coupled random variables evolving in time separated by a boundary,

$$X_t \xrightarrow{g} B_t \xrightarrow{h} Y_t$$

assumed to satisfy Itō SDEs

$$\begin{aligned}
\mathrm{d}X_t &= f_1(X_t, B_t, t)\,\mathrm{d}t + D_1(X_t, B_t, t)\,\mathrm{d}W_t^1 \\
\mathrm{d}B_t &= f_2(X_t, B_t, Y_t, t)\,\mathrm{d}t + D_2(X_t, B_t, Y_t, t)\,\mathrm{d}W_t^2 \\
\mathrm{d}Y_t &= f_3(B_t, Y_t, t)\,\mathrm{d}t + D_3(B_t, Y_t, t)\,\mathrm{d}W_t^3
\end{aligned}$$

The precise coupling structure is specific to a given system and defines different classes of dynamics [Friston et al, 2023].

Suppose f_1 is not the gradient of a smooth function and X_t has a pullback attractor in the state space

Then under certain regularity assumptions there exists a non-equilibrium steady state density $p^*(x)$

In which case we have the normal form

$$\mathrm{d}X_t = -(Q - \Gamma)\nabla_x \log p^*(X_t)\,\mathrm{d}t + D\,\mathrm{d}W_t$$

with $Q^\top = -Q$, $\Gamma \geqslant 0$, and $2\Gamma = DD^\top$.

Suppose f_1 is not the gradient of a smooth function and X_t has a pullback attractor in the state space

Then under certain regularity assumptions there exists a non-equilibrium steady state density $p^*(x)$

In which case we have the normal form

$$\mathrm{d}X_t = -(Q - \Gamma)\nabla_x \log p^*(X_t)\,\mathrm{d}t + D\,\mathrm{d}W_t$$

with $Q^\top = -Q$, $\Gamma \geqslant 0$, and $2\Gamma = DD^\top$.

Occupation measure up to $t = 36.0$

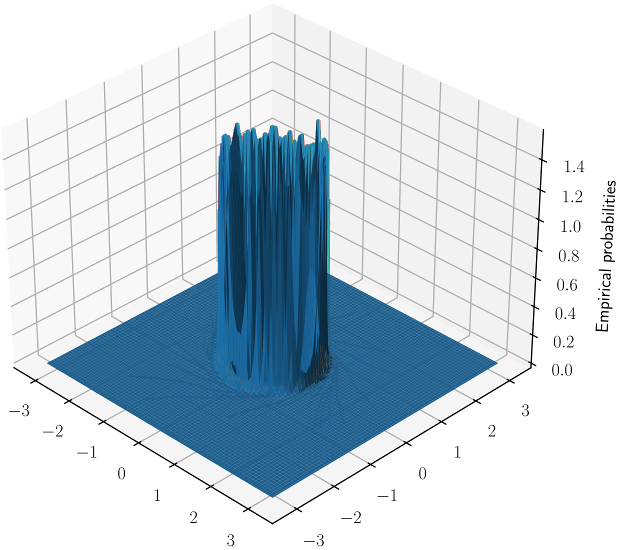

Suppose f_1 is not the gradient of a smooth function and X_t has a pullback attractor in the state space

Then under certain regularity assumptions there exists a non-equilibrium steady state density $p^*(x)$

In which case we have the normal form

$$\mathrm{d}X_t = -(Q - \Gamma)\nabla_x \log p^*(X_t)\,\mathrm{d}t + D\,\mathrm{d}W_t$$

with $Q^\top = -Q$, $\Gamma \geqslant 0$, and $2\Gamma = DD^\top$.

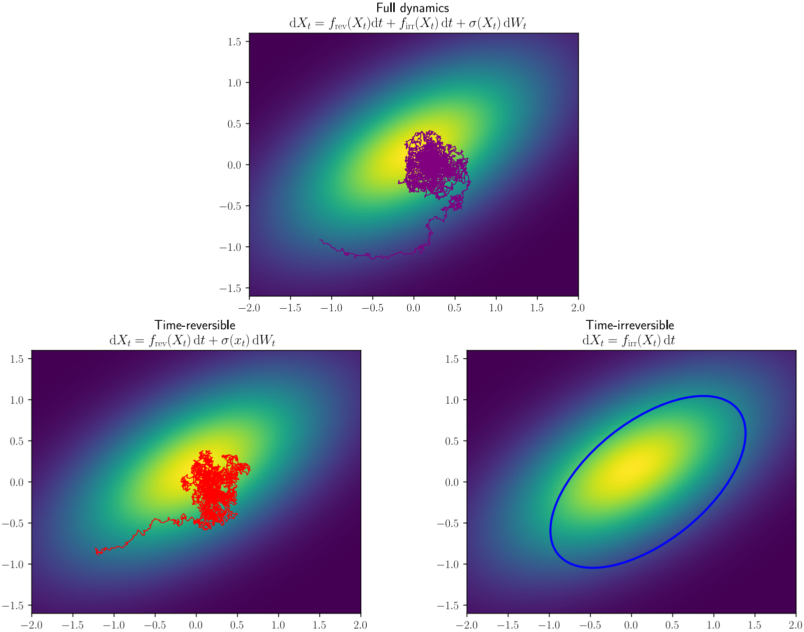
Full dynamics
$\mathrm{d}X_t = f_{\mathrm{rev}}(X_t)\mathrm{d}t + f_{\mathrm{irr}}(X_t)\,\mathrm{d}t + \sigma(X_t)\,\mathrm{d}W_t$
Time-reversible
$\mathrm{d}X_t = f_{\mathrm{rev}}(X_t)\,\mathrm{d}t + \sigma(x_t)\,\mathrm{d}W_t$
Time-irreversible
$\mathrm{d}X_t = f_{\mathrm{irr}}(X_t)\,\mathrm{d}t$

Now stipulate a variational posterior $q(y; x)$ and write

$$F(b, x) := \int q(y; x) \log q(y; x) \, \mathrm{d}y - \int q(y; x) \log p(y, b, x) \, \mathrm{d}y$$

Recall that $p(y, b, x) = p(y \mid b, x) p(b, x)$ and $\log ab = \log a + \log b$

Applying this and conditional independence we have

$$F(b, x) = \mathop{\mathbf{E}}_{q(y;x)} [\log q(y; x)] - \mathop{\mathbf{E}}_{q(y;x)} [\log p(y \mid b)] - \log p(b, x)$$

. . . so that

$$F(b, x) \geqslant - \log p(b, x).$$

Now stipulate a variational posterior $q(y; x)$ and write

$$F(b, x) := \int q(y; x) \log q(y; x) \, \mathrm{d}y - \int q(y; x) \log p(y, b, x) \, \mathrm{d}y$$

Recall that $p(y, b, x) = p(y \mid b, x)p(b, x)$ and $\log ab = \log a + \log b$

Applying this and conditional independence we have

$$F(b, x) = \mathop{\mathbf{E}}_{q(y;x)} [\log q(y; x)] - \mathop{\mathbf{E}}_{q(y;x)} [\log p(y \mid b)] - \log p(b, x)$$

. . . so that

$$F(b, x) \geqslant -\log p(b, x).$$

Now stipulate a variational posterior $q(y;x)$ and write

$$F(b,x) := \int q(y;x) \log q(y;x)\,\mathrm{d}y - \int q(y;x) \log p(y,b,x)\,\mathrm{d}y$$

Recall that $p(y,b,x) = p(y \mid b,x)p(b,x)$ and $\log ab = \log a + \log b$

Applying this and conditional independence we have

$$F(b,x) = \underset{q(y;x)}{\mathbf{E}}[\log q(y;x)] - \underset{q(y;x)}{\mathbf{E}}[\log p(y \mid b)] - \log p(b,x)$$

... so that

$$F(b,x) \geqslant -\log p(b,x).$$

Now stipulate a variational posterior $q(y;x)$ and write

$$F(b,x) := \int q(y;x)\log q(y;x)\,\mathrm{d}y - \int q(y;x)\log p(y,b,x)\,\mathrm{d}y$$

Recall that $p(y,b,x) = p(y \mid b,x)p(b,x)$ and $\log ab = \log a + \log b$

Applying this and conditional independence we have

$$F(b,x) = \underbrace{\boxed{\mathop{\mathbf{E}}_{q(y;x)}[\log q(y;x)] - \mathop{\mathbf{E}}_{q(y;x)}[\log p(y \mid b)]}}_{\geqslant 0\ldots} - \log p(b,x)$$

... so that

$$F(b,x) \geqslant -\log p(b,x).$$

Now stipulate a variational posterior $q(y;x)$ and write

$$F(b,x) := \int q(y;x)\log q(y;x)\,\mathrm{d}y - \int q(y;x)\log p(y,b,x)\,\mathrm{d}y$$

Recall that $p(y,b,x) = p(y \mid b,x)p(b,x)$ and $\log ab = \log a + \log b$

Applying this and conditional independence we have

$$F(b,x) = \underbrace{\boxed{\mathop{\mathbf{E}}_{q(y;x)}[\log q(y;x)] - \mathop{\mathbf{E}}_{q(y;x)}[\log p(y \mid b)]}}_{\geqslant 0\ldots} - \log p(b,x)$$

... so that

$$F(b,x) \geqslant -\log p(b,x).$$

Suppose there exists an x such that $q(y; x)$ equals $p(y \mid b)$ almost surely. Denote that x as x^*

Then $\mathbf{E}[\log q(y; x^*)] = \mathbf{E}[\log p(y \mid b)]$ and our SDE becomes

$$\mathrm{d}X_t = -(Q - \Gamma)\nabla_x F(b, x^*)\,\mathrm{d}t + D\,\mathrm{d}W_t$$

The difference of expectations is the *KL divergence* between a variational posterior and target distribution; the free energy is a tractable upper bound on model evidence

Implication:

All Markov-blanketted non-equilibrium processes on an attractor in their state space can be written as if implementing Bayesian inference over the likely causes of sensations

Suppose there exists an x such that $q(y; x)$ equals $p(y \mid b)$ almost surely. Denote that x as x^*

Then $\mathbf{E}[\log q(y; x^*)] = \mathbf{E}[\log p(y \mid b)]$ and our SDE becomes

$$\mathrm{d}X_t = -(Q - \Gamma)\nabla_x F(b, x^*)\,\mathrm{d}t + D\,\mathrm{d}W_t$$

The difference of expectations is the *KL divergence* between a variational posterior and target distribution; the free energy is a tractable upper bound on model evidence

Implication:

> All Markov-blanketted non-equilibrium processes on an attractor in their state space can be written as if implementing Bayesian inference over the likely causes of sensations

Suppose there exists an x such that $q(y; x)$ equals $p(y \mid b)$ almost surely. Denote that x as x^*

Then $\mathbf{E}[\log q(y; x^*)] = \mathbf{E}[\log p(y \mid b)]$ and our SDE becomes

$$\mathrm{d}X_t = -(Q - \Gamma)\nabla_x F(b, x^*)\,\mathrm{d}t + D\,\mathrm{d}W_t$$

The difference of expectations is the *KL divergence* between a variational posterior and target distribution; the free energy is a tractable upper bound on model evidence

Implication:

All Markov-blanketted non-equilibrium processes on an attractor in their state space can be written as if implementing Bayesian inference over the likely causes of sensations

Suppose there exists an x such that $q(y; x)$ equals $p(y \mid b)$ almost surely. Denote that x as x^*

Then $\mathbf{E}[\log q(y; x^*)] = \mathbf{E}[\log p(y \mid b)]$ and our SDE becomes

$$\mathrm{d}X_t = -(Q - \Gamma)\nabla_x F(b, x^*)\,\mathrm{d}t + D\,\mathrm{d}W_t$$

The difference of expectations is the *KL divergence* between a variational posterior and target distribution; the free energy is a tractable upper bound on model evidence

Implication:

All Markov-blanketted non-equilibrium processes on an attractor in their state space can be written as if implementing Bayesian inference over the likely causes of sensations

Suppose there exists an x such that $q(y;x)$ equals $p(y \mid b)$ almost surely. Denote that x as x^*

Then $\mathbf{E}[\log q(y;x^*)] = \mathbf{E}[\log p(y \mid b)]$ and our SDE becomes

$$\mathrm{d}X_t = -(Q - \Gamma)\nabla_x F(b, x^*)\,\mathrm{d}t + D\,\mathrm{d}W_t$$

The difference of expectations is the *KL divergence* between a variational posterior and target distribution; the free energy is a tractable upper bound on model evidence

Implication:

> All Markov-blanketted non-equilibrium processes on an attractor in their state space can be written as if implementing Bayesian inference over the likely causes of sensations


Ben Nevis Inn
WEATHER FORECASTING STONE
CONDITION
FORECAST
Stone is Wet
Rain
Stone is Dry
Not Raining
Shadow on Ground
Sunny
White on Top
Snowing
Can't see Stone
Foggy
Swinging Stone
Windy
Stone Jumping Up & Down
Earthquake
Stone Gone
Tornado

Interpretation:

If a system mainly does what one (a modeller) expects it to do, it can only be so surprising

For instance

- Observable stones must be concentrated on stone-like states
- Observable control systems must be concentrated on set points

Interpretation:

If a system mainly does what one (a modeller) expects it to do, it can only be so surprising

For instance

- Observable stones must be concentrated on stone-like states
- Observable control systems must be concentrated on set points

Interpretation:

If a system mainly does what one (a modeller) expects it to do, it can only be so surprising

For instance

- Observable stones must be concentrated on stone-like states
- Observable control systems must be concentrated on set points

So far all we have said is that, *via* the interactions across a shared boundary, coupled random dynamical systems estimate each others statistics

Ultimately: any 'thing' encodes a probability distribution over possible environmental states... because the environment must be conducive to it existing

Question: why bother?

Answer: complex systems are difficult to understand because of their interactions, so replacing couplings with the study of variational free energy is fruitful

So far all we have said is that, *via* the interactions across a shared boundary, coupled random dynamical systems estimate each others statistics

Ultimately: any ‘thing’ encodes a probability distribution over possible environmental states... because the environment must be conducive to it existing

Question: why bother?

Answer: complex systems are difficult to understand because of their interactions, so replacing couplings with the study of variational free energy is fruitful

So far all we have said is that, *via* the interactions across a shared boundary, coupled random dynamical systems estimate each others statistics

Ultimately: any ‘thing’ encodes a probability distribution over possible environmental states. . . because the environment must be conducive to it existing

Question: why bother?

Answer: complex systems are difficult to understand because of their interactions, so replacing couplings with the study of variational free energy is fruitful

So far all we have said is that, *via* the interactions across a shared boundary, coupled random dynamical systems estimate each others statistics

Ultimately: any 'thing' encodes a probability distribution over possible environmental states... because the environment must be conducive to it existing

Question: why bother?

Answer: complex systems are difficult to understand because of their interactions, so replacing couplings with the study of variational free energy is fruitful

Let us now imagine being modellers in the system's environment. . . and undertaking the task of trying to infer the probabilities of different states the observed system can occupy.

Let us now imagine being modellers in the system's environment... and undertaking the task of trying to infer the probabilities of different states the observed system can occupy.

Viewing 'finding x^*' as a constrained variation of entropy, one can write

$$\arg\max_p \int p(x \mid b) \log p(x \mid b) \, \mathrm{d}x - \lambda \Big(\mathbf{E} \left[\left\| x - u^{-1}(\hat{y}_b) \right\|^2_{L^2} \mid b \right] - S_\star^{-1} \Big)$$

for $p^*(x \mid b)$.

The solution is

$$p^*(x \mid b) = e^{-\lambda \left(x - u^{-1}(\hat{y}_b)\right)^\top S_\star^{-1} \left(x - u^{-1}(\hat{y}_b)\right)}$$

In words: the optimal distribution is such that the conditional expectation maps (by u) to the conditional expectation of the environment given the blanket state, and the conditional variance maps (by $\star$) to the inverse precision over the environment

This recovers the Laplace approximation of $p(x \mid b)$

Viewing 'finding x^*' as a constrained variation of entropy, one can write

$$\arg\max_p \int p(x \mid b) \log p(x \mid b)\,\mathrm{d}x - \lambda\Big(\mathbf{E}\left[\left\| x - u^{-1}(\hat{y}_b) \right\|^2_{L^2} \mid b \right] - S_\star^{-1} \Big)$$

for $p^*(x \mid b)$.

The solution is

$$p^*(x \mid b) = e^{-\lambda\left(x - u^{-1}(\hat{y}_b)\right)^\top S_\star^{-1}\left(x - u^{-1}(\hat{y}_b)\right)}$$

In words: the optimal distribution is such that the conditional expectation maps (by u) to the conditional expectation of the environment given the blanket state, and the conditional variance maps (by $\star$) to the inverse precision over the environment

This recovers the Laplace approximation of $p(x \mid b)$

Viewing 'finding x^*' as a constrained variation of entropy, one can write

$$\arg\max_{p} \int p(x \mid b) \log p(x \mid b) \, \mathrm{d}x - \lambda \Big(\mathbf{E}\left[\left\| x - u^{-1}(\hat{y}_b) \right\|^2_{L^2} \mid b \right] - S_\star^{-1} \Big)$$

for $p^*(x \mid b)$.

The solution is

$$p^*(x \mid b) = e^{-\lambda \left(x - u^{-1}(\hat{y}_b) \right)^\top S_\star^{-1} \left(x - u^{-1}(\hat{y}_b) \right)}$$

In words: the optimal distribution is such that the conditional expectation maps (by u) to the conditional expectation of the environment given the blanket state, and the conditional variance maps (by $\star$) to the inverse precision over the environment

This recovers the Laplace approximation of $p(x \mid b)$

Viewing 'finding x^*' as a constrained variation of entropy, one can write

$$\arg\max_p \int p(x \mid b) \log p(x \mid b) \, \mathrm{d}x - \lambda \Big(\mathbf{E} \left[\left\| x - u^{-1}(\hat{y}_b) \right\|_{L^2}^2 \mid b \right] - S_\star^{-1} \Big)$$

for $p^*(x \mid b)$.

The solution is

$$p^*(x \mid b) = e^{-\lambda \left(x - u^{-1}(\hat{y}_b)\right)^\top S_\star^{-1} \left(x - u^{-1}(\hat{y}_b)\right)}$$

In words: the optimal distribution is such that the conditional expectation maps (by u) to the conditional expectation of the environment given the blanket state, and the conditional variance maps (by $\star$) to the inverse precision over the environment

This recovers the Laplace approximation of $p(x \mid b)$

Viewing 'finding x^*' as a constrained variation of entropy, one can write

$$\arg\max_p \int p(x \mid b) \log p(x \mid b)\,\mathrm{d}x - \lambda\Big(\mathbf{E}\left[\left\|x - u^{-1}(\hat{y}_b)\right\|^2_{L^2} \mid b\right] - S_\star^{-1}\Big)$$

for $p^*(x \mid b)$.

The solution is

$$p^*(x \mid b) = e^{-\lambda\left(x-u^{-1}(\hat{y}_b)\right)^\top S_\star^{-1}\left(x-u^{-1}(\hat{y}_b)\right)}$$

In words: the optimal distribution is such that the conditional expectation maps (by u) to the conditional expectation of the environment given the blanket state, and the conditional variance maps (by $\star$) to the inverse precision over the environment

This recovers the Laplace approximation of $p(x \mid b)$

This tells a somehow dual story: given our knowledge of the environment and the coupling, assuming the free energy principle, we can infer the steady state density over system states by maximising entropy

Or more actively: the system maintains a regime of states by constraining itself, and these constraints are equivalent to certain environmental compatibility conditions

However . . . a small *caveat* exists.

Maximum entropy is only naïvely applicable at equilibrium!

This tells a somehow dual story: given our knowledge of the environment and the coupling, assuming the free energy principle, we can infer the steady state density over system states by maximising entropy

Or more actively: the system maintains a regime of states by constraining itself, and these constraints are equivalent to certain environmental compatibility conditions

However ... a small *caveat* exists.

Maximum entropy is only naïvely applicable at equilibrium!

This tells a somehow dual story: given our knowledge of the environment and the coupling, assuming the free energy principle, we can infer the steady state density over system states by maximising entropy

Or more actively: the system maintains a regime of states by constraining itself, and these constraints are equivalent to certain environmental compatibility conditions

However ... a small *caveat* exists.

Maximum entropy is only naïvely applicable at equilibrium!

This tells a somehow dual story: given our knowledge of the environment and the coupling, assuming the free energy principle, we can infer the steady state density over system states by maximising entropy

Or more actively: the system maintains a regime of states by constraining itself, and these constraints are equivalent to certain environmental compatibility conditions

However . . . a small *caveat* exists.

Maximum entropy is only naïvely applicable at equilibrium!

This tells a somehow dual story: given our knowledge of the environment and the coupling, assuming the free energy principle, we can infer the steady state density over system states by maximising entropy

Or more actively: the system maintains a regime of states by constraining itself, and these constraints are equivalent to certain environmental compatibility conditions

However . . . a small *caveat* exists.

Maximum entropy is only naïvely applicable at equilibrium!

This tells a somehow dual story: given our knowledge of the environment and the coupling, assuming the free energy principle, we can infer the steady state density over system states by maximising entropy

Or more actively: the system maintains a regime of states by constraining itself, and these constraints are equivalent to certain environmental compatibility conditions

However ... a small *caveat* exists.

Maximum entropy is only naïvely applicable at equilibrium!

In general an enduring question is: what about time-varying systems?

In fact the original formulation of the free energy principle concerns generalised filtering performed by the expected evolution of a random dynamical system

We will investigate this formulation and extend the previous comparison theorem

In general an enduring question is: what about time-varying systems?

In fact the original formulation of the free energy principle concerns generalised filtering performed by the expected evolution of a random dynamical system

We will investigate this formulation and extend the previous comparison theorem

In general an enduring question is: what about time-varying systems?

In fact the original formulation of the free energy principle concerns generalised filtering performed by the expected evolution of a random dynamical system

We will investigate this formulation and extend the previous comparison theorem

In general an enduring question is: what about time-varying systems?

In fact the original formulation of the free energy principle concerns generalised filtering performed by the expected evolution of a random dynamical system

We will investigate this formulation and extend the previous comparison theorem

Let us pass to 'generalised coordinates of motion' now

i.e., a tuple $\tilde{x} = [x, x', x'', \dots]$ containing instantaneous derivatives of the solution to a (stochastic) ordinary differential equation

Denote $\tilde{x}_t$ as a generalised state at time t

Our SDE can be written as

$$\mathrm{d}\tilde{X}_t = \mathbf{D}f(\tilde{X}_t) + D\,\mathrm{d}\tilde{W}_t$$

(*e.g.* simultaneous random evolutions of all derivatives)

Generalised coordinates admit a mapping to solutions of the original SDE by reconstructing any path from its derivatives*

Let us pass to 'generalised coordinates of motion' now

i.e., a tuple $\tilde{x} = [x, x', x'', \dots]$ containing instantaneous derivatives of the solution to a (stochastic) ordinary differential equation

Denote $\tilde{x}_t$ as a generalised state at time t

Our SDE can be written as

$$\mathrm{d}\tilde{X}_t = \mathbf{D}f(\tilde{X}_t) + D\,\mathrm{d}\tilde{W}_t$$

(*e.g.* simultaneous random evolutions of all derivatives)

Generalised coordinates admit a mapping to solutions of the original SDE by reconstructing any path from its derivatives*

Let us pass to 'generalised coordinates of motion' now

i.e., a tuple $\tilde{x} = [x, x', x'', \ldots]$ containing instantaneous derivatives of the solution to a (stochastic) ordinary differential equation

Denote $\tilde{x}_t$ as a generalised state at time t

Our SDE can be written as

$$\mathrm{d}\tilde{X}_t = \mathbf{D}f(\tilde{X}_t) + D\,\mathrm{d}\tilde{W}_t$$

(*e.g.* simultaneous random evolutions of all derivatives)

Generalised coordinates admit a mapping to solutions of the original SDE by reconstructing any path from its derivatives*

Let us pass to 'generalised coordinates of motion' now

i.e., a tuple $\tilde{x} = [x, x', x'', \ldots]$ containing instantaneous derivatives of the solution to a (stochastic) ordinary differential equation

Denote $\tilde{x}_t$ as a generalised state at time t

Our SDE can be written as

$$\mathrm{d}\tilde{X}_t = \mathbf{D}f(\tilde{X}_t) + D\,\mathrm{d}\tilde{W}_t$$

(*e.g.* simultaneous random evolutions of all derivatives)

Generalised coordinates admit a mapping to solutions of the original SDE by reconstructing any path from its derivatives*

Let us pass to 'generalised coordinates of motion' now

i.e., a tuple $\tilde{x} = [x, x', x'', \ldots]$ containing instantaneous derivatives of the solution to a (stochastic) ordinary differential equation

Denote $\tilde{x}_t$ as a generalised state at time t

Our SDE can be written as

$$\mathrm{d}\tilde{X}_t = \mathbf{D}f(\tilde{X}_t) + D\,\mathrm{d}\tilde{W}_t$$

(*e.g.* simultaneous random evolutions of all derivatives)

Generalised coordinates admit a mapping to solutions of the original SDE by reconstructing any path from its derivatives*

Separately, general theory tells us the following:

$$-\log p(X_\tau) = \int_0^\tau \left\|\partial_t X_t - f(X_t)\right\|^2 \mathrm{d}t + o(D)$$

(Freidlin–Wentzell; Onsager–Machlup)

Denote γ_t as a path up until time t. Observation: our normal form is intact ... because this fact implies

$$\mathrm{d}\gamma_t = -\nabla_\gamma \log p(\gamma_t) + D\,\mathrm{d}W_t$$

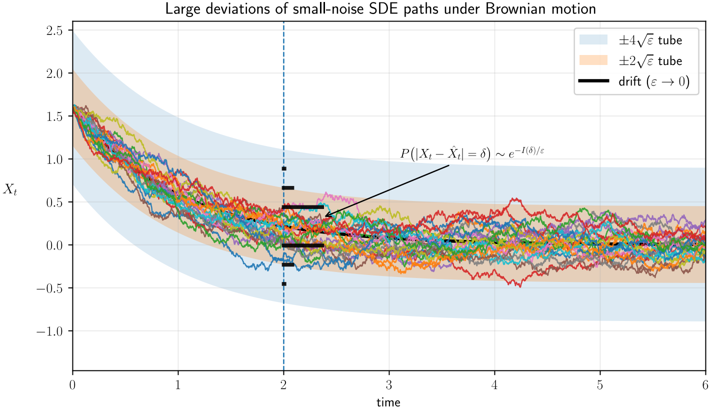

Large deviations of small-noise SDE paths under Brownian motion
$\pm 4\sqrt{\varepsilon}$ tube
$\pm 2\sqrt{\varepsilon}$ tube
drift ($\varepsilon \to 0$)
$P(|X_t - \hat{X}_t| = \delta) \sim e^{-I(\delta)/\varepsilon}$
X_t
time
2.5
2.0
1.5
1.0
0.5
0.0
−0.5
−1.0
0
1
2
3
4
5
6

Separately, general theory tells us the following:

$$-\log p(X_\tau) = \int_0^\tau \left\|\partial_t X_t - f(X_t)\right\|^2 \mathrm{d}t + o(D)$$

(Freidlin–Wentzell; Onsager–Machlup)

Denote γ_t as a path up until time t. Observation: our normal form is intact ... because this fact implies

$$\mathrm{d}\gamma_t = -\nabla_\gamma \log p(\gamma_t) + D\,\mathrm{d}W_t$$

Separately, general theory tells us the following:

$$-\log p(X_\tau) = \int_0^\tau \left\|\partial_t X_t - f(X_t)\right\|^2 \mathrm{d}t + o(D)$$

(Freidlin–Wentzell; Onsager–Machlup)

Denote γ_t as a path up until time t. Observation: our normal form is intact ... because this fact implies

$$\mathrm{d}\gamma_t = -\nabla_\gamma \log p(\gamma_t) + D\,\mathrm{d}W_t$$

The idea of generalised filtering is: we want to infer the evolving causes of evolving sensations

Given a model m and generalised sensor data $\tilde{s}$ one wants to maximise the accumulated log-evidence

$$\varepsilon = \int_0^\tau \log p(\tilde{s}_t \mid m)$$

It is almost never possible to tractably evaluate ε, so instead we study the *free action*

$$\mathscr{F} := \int_0^\tau \mathop{\mathbf{E}}_q[\log q(\vartheta_t)] - \mathop{\mathbf{E}}_q[\log p(\vartheta_t \mid \tilde{s}_t, m)] - \log p(\tilde{s}_t \mid m) \, \mathrm{d}t \,.$$

Question: why is this satisfactory?

The idea of generalised filtering is: we want to infer the evolving causes of evolving sensations

Given a model m and generalised sensor data $\tilde{s}$ one wants to maximise the accumulated log-evidence

$$\varepsilon = \int_0^\tau \log p(\tilde{s}_t \mid m)$$

It is almost never possible to tractably evaluate ε, so instead we study the *free action*

$$\mathscr{F} := \int_0^\tau \mathop{\mathbf{E}}_q[\log q(\vartheta_t)] - \mathop{\mathbf{E}}_q[\log p(\vartheta_t \mid \tilde{s}_t, m)] - \log p(\tilde{s}_t \mid m)\,\mathrm{d}t\,.$$

Question: why is this satisfactory?

The idea of generalised filtering is: we want to infer the evolving causes of evolving sensations

Given a model m and generalised sensor data $\tilde{s}$ one wants to maximise the accumulated log-evidence

$$\varepsilon = \int_0^\tau \log p(\tilde{s}_t \mid m)$$

It is almost never possible to tractably evaluate ε, so instead we study the *free action*

$$\mathscr{F} := \int_0^\tau \mathop{\mathbf{E}}_q[\log q(\vartheta_t)] - \mathop{\mathbf{E}}_q[\log p(\vartheta_t \mid \tilde{s}_t, m)] - \log p(\tilde{s}_t \mid m)\,\mathrm{d}t\,.$$

Question: why is this satisfactory?

The idea of generalised filtering is: we want to infer the evolving causes of evolving sensations

Given a model m and generalised sensor data $\tilde{s}$ one wants to maximise the accumulated log-evidence

$$\varepsilon = \int_0^\tau \log p(\tilde{s}_t \mid m)$$

It is almost never possible to tractably evaluate ε, so instead we study the *free action*

$$\mathscr{F} := \int_0^\tau \underset{q}{\mathbf{E}}[\log q(\vartheta_t)] - \underset{q}{\mathbf{E}}[\log p(\vartheta_t \mid \tilde{s}_t, m)] - \log p(\tilde{s}_t \mid m)\,\mathrm{d}t\,.$$

Question: why is this satisfactory?

Answer: this is satisfactory because

$$\int_0^\tau \mathop{\mathbf{E}}_q[\log q(\vartheta_t)] - \mathop{\mathbf{E}}_q[\log p(\vartheta_t \mid \tilde{s}_t, m)]\,\mathrm{d}t \geqslant 0$$

so that

$$\int_0^\tau \mathop{\mathbf{E}}_q[\log q(\vartheta_t)] - \mathop{\mathbf{E}}_q[\log p(\vartheta_t \mid \tilde{s}_t, m)]\,\mathrm{d}t - \int_0^\tau \log p(\tilde{s}_t \mid m)\,\mathrm{d}t \geqslant -\int_0^\tau \log p(\tilde{s}_t \mid m)\,\mathrm{d}t\,.$$

Conclusion: minimising free action maximises accumulated log-evidence.

Answer: this is satisfactory because

$$\int_0^\tau \mathop{\mathbf{E}}_q[\log q(\vartheta_t)] - \mathop{\mathbf{E}}_q[\log p(\vartheta_t \mid \tilde{s}_t, m)]\,\mathrm{d}t \geqslant 0$$

so that

$$\begin{aligned}\int_0^\tau \mathop{\mathbf{E}}_q[\log q(\vartheta_t)] - \mathop{\mathbf{E}}_q[\log p(\vartheta_t \mid \tilde{s}_t, m)]\,\mathrm{d}t - \int_0^\tau \log p(\tilde{s}_t \mid m)\,\mathrm{d}t \\ \geqslant -\int_0^\tau \log p(\tilde{s}_t \mid m)\,\mathrm{d}t\,.\end{aligned}$$

Conclusion: minimising free action maximises accumulated log-evidence.

Now suppose we have two coupled random processes with non-stationary statistics

We will write the system as γ, the environment as ξ, and the boundary as β

We know the evolution of each minimises its log-path-probability

Introducing the free action

$$\mathscr{F}(\beta_t, \gamma_t) = \int_0^\tau D_{\mathrm{KL}}\big(q(\xi_t; \gamma_{t,\beta}) \| \, p(\xi_t \mid \beta_t)\big) - \log p(\beta_t, \gamma_t) \, \mathrm{d}t$$

it follows that for γ_t^* we have

$$\mathrm{d}\gamma_t = \mathbf{D}\mathscr{F}(\beta_t, \gamma_t^*) + D \, \mathrm{d}W_t$$

Now suppose we have two coupled random processes with <u>non-stationary statistics</u>

We will write the system as γ, the environment as ξ, and the boundary as β

We know the evolution of each minimises its log-path-probability

Introducing the free action

$$\mathscr{F}(\beta_t, \gamma_t) = \int_0^\tau D_{\mathrm{KL}}\big(q(\xi_t; \gamma_{t,\beta}) \| \, p(\xi_t \mid \beta_t)\big) - \log p(\beta_t, \gamma_t) \, \mathrm{d}t$$

it follows that for γ_t^* we have

$$\mathrm{d}\gamma_t = \mathbf{D}\mathscr{F}(\beta_t, \gamma_t^*) + D \, \mathrm{d}W_t$$

Now suppose we have two coupled random processes with non-stationary statistics

We will write the system as γ, the environment as ξ, and the boundary as β

We know the evolution of each minimises its log-path-probability

Introducing the free action

$$\mathscr{F}(\beta_t, \gamma_t) = \int_0^\tau D_{\mathrm{KL}}\big(q(\xi_t; \gamma_{t,\beta}) \| \, p(\xi_t \mid \beta_t)\big) - \log p(\beta_t, \gamma_t) \,\mathrm{d}t$$

it follows that for γ_t^* we have

$$\mathrm{d}\gamma_t = \mathbf{D}\mathscr{F}(\beta_t, \gamma_t^*) + D \,\mathrm{d}W_t$$

Now suppose we have two coupled random processes with non-stationary statistics

We will write the system as γ, the environment as ξ, and the boundary as β

We know the evolution of each minimises its log-path-probability

Introducing the free action

$$\mathscr{F}(\beta_t, \gamma_t) = \int_0^\tau D_{\mathrm{KL}}\big(q(\xi_t; \gamma_{t,\beta}) \,\|\, p(\xi_t \mid \beta_t)\big) - \log p(\beta_t, \gamma_t) \,\mathrm{d}t$$

it follows that for γ_t^* we have

$$\mathrm{d}\gamma_t = \mathbf{D}\mathscr{F}(\beta_t, \gamma_t^*) + D \,\mathrm{d}W_t$$

Maximising path entropy gives us the same picture as what we had before

Namely: viewing 'finding γ_t^*' as a constrained variation of path entropy, one can write

$$\arg\max_p \int p(\gamma_t) \log p(\gamma_t) \, \mathrm{d}\gamma - \lambda \Big(\mathbf{E}\left[\left\| \gamma_t - u^{-1}(\hat{\xi}_{t,\beta}) \right\|_{L^2}^2 \mid \beta \right] - S_\star^{-1} \Big)$$

for $p^*(\gamma \mid \beta)$.

The solution is

$$p^*(\gamma) = e^{-\lambda \left(\gamma - u^{-1}(\hat{\xi}_\beta)\right)^\top S_\star^{-1} \left(\gamma - u^{-1}(\hat{\xi}_\beta)\right)}$$

Maximising path entropy gives us the same picture as what we had before

Namely: viewing 'finding γ_t^*' as a constrained variation of path entropy, one can write

$$\arg\max_p \int p(\gamma_t) \log p(\gamma_t)\,\mathrm{d}\gamma - \lambda\Big(\mathbf{E}\left[\left\|\gamma_t - u^{-1}(\hat{\xi}_{t,\beta})\right\|_{L^2}^2 \mid \beta \right] - S_\star^{-1} \Big)$$

for $p^*(\gamma \mid \beta)$.

The solution is

$$p^*(\gamma) = e^{-\lambda\left(\gamma - u^{-1}(\hat{\xi}_\beta)\right)^\top S_\star^{-1}\left(\gamma - u^{-1}(\hat{\xi}_\beta)\right)}$$

Maximising path entropy gives us the same picture as what we had before

Namely: viewing 'finding γ_t^*' as a constrained variation of path entropy, one can write

$$\arg\max_p \int p(\gamma_t) \log p(\gamma_t) \, \mathrm{d}\gamma - \lambda\Big(\mathbf{E}\left[\left\| \gamma_t - u^{-1}(\hat{\xi}_{t,\beta}) \right\|_{L^2}^2 \mid \beta \right] - S_\star^{-1} \Big)$$

for $p^*(\gamma \mid \beta)$.

The solution is

$$p^*(\gamma) = e^{-\lambda\left(\gamma - u^{-1}(\hat{\xi}_\beta)\right)^\top S_\star^{-1}\left(\gamma - u^{-1}(\hat{\xi}_\beta)\right)}$$

In words: the optimal distribution is such that the expected evolution tracks (by u) to the expected evolution of the environment given the blanket state, and the conditional variance maps (by $\star$) to the inverse precision over generalised states of the environment

This recovers the Laplace approximation of $p(\gamma)$

This tells a somehow dual story: given our knowledge of the system's environment and the coupling, assuming the free energy principle over paths, we can infer the likely evolutions of the system by maximising path entropy

Or more actively: the system maintains a regime of evolutions by constraining itself, and these constraints are equivalent to certain environmental compatibility conditions

In words: the optimal distribution is such that the expected evolution tracks (by u) to the expected evolution of the environment given the blanket state, and the conditional variance maps (by $\star$) to the inverse precision over generalised states of the environment

This recovers the Laplace approximation of $p(\gamma)$

This tells a somehow dual story: given our knowledge of the system's environment and the coupling, assuming the free energy principle over paths, we can infer the likely evolutions of the system by maximising path entropy

Or more actively: the system maintains a regime of evolutions by constraining itself, and these constraints are equivalent to certain environmental compatibility conditions

In words: the optimal distribution is such that the expected evolution tracks (by u) to the expected evolution of the environment given the blanket state, and the conditional variance maps (by $\star$) to the inverse precision over generalised states of the environment

This recovers the Laplace approximation of $p(\gamma)$

This tells a somehow dual story: given our knowledge of the system's environment and the coupling, assuming the free energy principle over paths, we can infer the likely evolutions of the system by maximising path entropy

Or more actively: the system maintains a regime of evolutions by constraining itself, and these constraints are equivalent to certain environmental compatibility conditions

In words: the optimal distribution is such that the expected evolution tracks (by u) to the expected evolution of the environment given the blanket state, and the conditional variance maps (by $\star$) to the inverse precision over generalised states of the environment

This recovers the Laplace approximation of $p(\gamma)$

This tells a somehow dual story: given our knowledge of the system's environment and the coupling, assuming the free energy principle over paths, we can infer the likely evolutions of the system by maximising path entropy

Or more actively: the system maintains a regime of evolutions by constraining itself, and these constraints are equivalent to certain environmental compatibility conditions

In words: the optimal distribution is such that the expected evolution tracks (by u) to the expected evolution of the environment given the blanket state, and the conditional variance maps (by $\star$) to the inverse precision over generalised states of the environment

This recovers the Laplace approximation of $p(\gamma)$

This tells a somehow dual story: given our knowledge of the system's environment and the coupling, assuming the free energy principle over paths, we can infer the likely evolutions of the system by maximising path entropy

Or more actively: the system maintains a regime of evolutions by constraining itself, and these constraints are equivalent to certain environmental compatibility conditions

In words: the optimal distribution is such that the expected evolution tracks (by u) to the expected evolution of the environment given the blanket state, and the conditional variance maps (by $\star$) to the inverse precision over generalised states of the environment

This recovers the Laplace approximation of $p(\gamma)$

This tells a somehow dual story: given our knowledge of the system's environment and the coupling, assuming the free energy principle over paths, we can infer the likely evolutions of the system by maximising path entropy

Or more actively: the system maintains a regime of evolutions by constraining itself, and these constraints are equivalent to certain environmental compatibility conditions

In words: the optimal distribution is such that the expected evolution tracks (by u) to the expected evolution of the environment given the blanket state, and the conditional variance maps (by $\star$) to the inverse precision over generalised states of the environment

This recovers the Laplace approximation of $p(\gamma)$

This tells a somehow dual story: given our knowledge of the system's environment and the coupling, assuming the free energy principle over paths, we can infer the likely evolutions of the system by maximising path entropy

Or more actively: the system maintains a regime of evolutions by constraining itself, and these constraints are equivalent to certain environmental compatibility conditions

Question: why bother?

Answer: contemporary theories of stochastic thermodynamics offer arguments that the maintenance of non-equilibrium evolution is a property of inference about an environment (see *e.g.* England's theory of dissipative adaptation). Generalised fluctuation relations are then derived from

$$p(X_t) = e^{-\int_0^\tau \left\|\partial_t X_t - f(X_t)\right\|^2 \mathrm{d}t + o(D)}$$

In a cylindrical neighbourhood of the expected trajectory, $p(X_t)$ is equivalent to

$$p(X_t) = e^{-\mathscr{F}(\beta_t, \gamma_t)}.$$

This means the free energy principle offers (i) an automatic and pleasing account of *why* non-equilibrium systems do inference, and (ii) a method to prove conceptual arguments that fluctuations relations out of equilibrium are obtained through inference.

Question: why bother?

Answer: contemporary theories of stochastic thermodynamics offer arguments that the maintenance of non-equilibrium evolution is a property of inference about an environment (see *e.g.* England's theory of dissipative adaptation). Generalised fluctuation relations are then derived from

$$p(X_t) = e^{-\int_0^\tau \left\| \partial_t X_t - f(X_t) \right\|^2 \mathrm{d}t + o(D)}$$

In a cylindrical neighbourhood of the expected trajectory, $p(X_t)$ is equivalent to

$$p(X_t) = e^{-\mathscr{F}(\beta_t, \gamma_t)}.$$

This means the free energy principle offers (i) an automatic and pleasing account of *why* non-equilibrium systems do inference, and (ii) a method to prove conceptual arguments that fluctuations relations out of equilibrium are obtained through inference.

Question: why bother?

Answer: contemporary theories of stochastic thermodynamics offer arguments that the maintenance of non-equilibrium evolution is a property of inference about an environment (see *e.g.* England's theory of dissipative adaptation). Generalised fluctuation relations are then derived from

$$p(X_t) = e^{-\int_0^\tau \left\| \partial_t X_t - f(X_t) \right\|^2 \mathrm{d}t + o(D)}$$

In a cylindrical neighbourhood of the expected trajectory, $p(X_t)$ is equivalent to

$$p(X_t) = e^{-\mathscr{F}(\beta_t, \gamma_t)}.$$

This means the free energy principle offers (i) an automatic and pleasing account of *why* non-equilibrium systems do inference, and (ii) a method to prove conceptual arguments that fluctuations relations out of equilibrium are obtained through inference.

I offer the following takeaway.

It becomes prudent to ask not what non-equilibrium statistical mechanics can do for the free energy principle, but what the free energy principle can do for non-equilibrium statistical mechanics

I offer the following takeaway.

It becomes prudent to ask not what non-equilibrium statistical mechanics can do for the free energy principle, but what the free energy principle can do for non-equilibrium statistical mechanics

D A R Sakthivadivel, *Towards a geometry and analysis for Bayesian mechanics*, arXiv:2204.11900.

D A R Sakthivadivel, *Regarding flows under the free energy principle*, Phys Life Rev **42** (2022), 25–28.

D A R Sakthivadivel, *A worked example of the Bayesian mechanics of classical objects*. Active Inference: Third International Workshop, CCIS vol. 1721, 2022, pp. 298–318.

K J Friston, L Da Costa, D A R Sakthivadivel, C Heins, G A Pavliotis, M Ramstead, T Parr, *Path integrals, particular kinds, and strange things*, Phys Life Rev **47** (2023), 35–62.

K J Friston, M Ramstead, D A R Sakthivadivel, *A framework for the use of generative modelling in non-equilibrium statistical mechanics*, arXiv:2406.11630. To appear in *Proc Roy Soc A*.