Polynomial Life

and the Structure of Adaptive Systems

Toby St Clere Smithe

Topos Institute
&
University of Oxford

12 July 2021

Cyber Kittens Growing Up

- Last year, I sketched a categorical account of certain cybernetic systems, based on the idea that such systems should *perceive* and *act*, *i.e.*, perform approximate Bayesian inference.

- I showed that approximate inference problems collect into categories of *statistical games*.

- But, that presentation had a number of problems — particularly:
  - unsatisfactory notions of 'action' and 'dynamics';
  - too much generality in the wrong places;
  - no good way to talk about interacting systems;
  - plus, some technical issues!

- Since then, I've been trying to iron out these wrinkles …

This Talk

- I was inspired by David Spivak's advocacy of polynomial functors for interacting systems.
  - We could summarize what I've been doing as "trying to breathe life into polynomials".

- So today, I'll give a progress report on that work. I will:
  - present a simplification of statistical games;
  - briefly re-introduce polynomial functors;
  - define a new category of dynamical systems (Markov processes) over polynomials;
  - sketch *approximate* and *active inference doctrines*,
    the latter via a new category of *statistical games over polynomials*;
  - extend active inference to systems with *goals*;
  - and describe how this can model *homeostasis* and *morphogenesis*.

But first: let me apologize for the terseness of the abstract I submitted!
You will get a much clearer version shortly...

Overview

Recap: Bidirectional Structure of Bayesian Inference (1/2)

- We work in a **Markov category** of *stochastic channels*: inputs give uncertain outputs.
- Given a 'prediction' channel $c : X \to \mathcal{P}\, Y$, the corresponding 'update' channel has a 'state-dependent' type $c^\dagger : \mathcal{P}\, X \times Y \to \mathcal{P}\, X$.
- Such pairs of a forwards map with a 'dependent' backwards map are **lenses**.
- Theorem: the Bayesian inversion of a composite channel is the "lens composite" of its components [1].
- This seems to match (*e.g.*) what we see in the brain (see right).

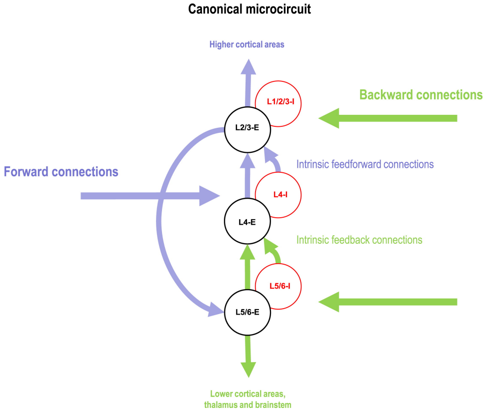

Figure: Bastos et al. [2]

Recap: Bidirectional Structure of Bayesian Inference (2/2)

Such *Bayesian lenses* form a category **BayesLens** of morphisms between (pairs of) spaces. They compose like this:

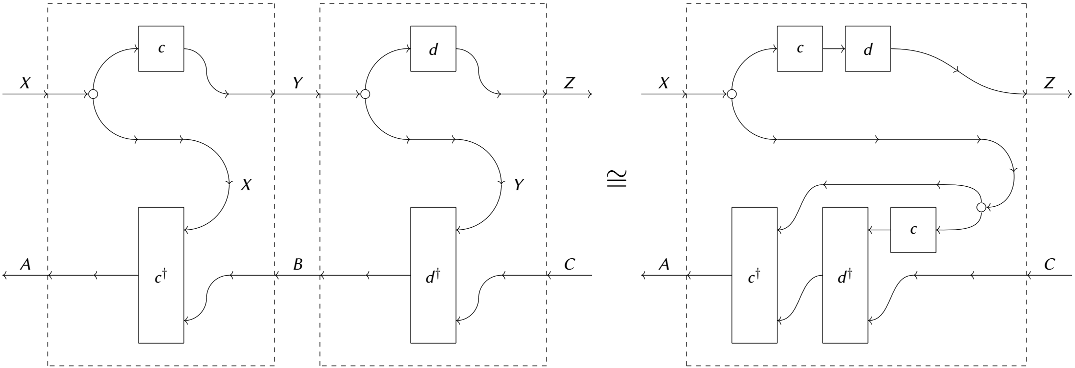

This captures the structure of inverting a hierarchical or causal model.
(Example: brain's visual predictions at cinema...)

Simpler Statistical Games for Approximate Inference (1/2)

- However, given a stochastic channel $c : X \to \mathcal{P}\, Y$ and a prior $\pi : \mathcal{P}\, X$, computing the inversion $c^{\dagger}_{\pi} : Y \to \mathcal{P}\, X$ is often computationally hard: we usually need to approximate it.
- This gives us a lot of freedom. Often, one approximation scheme might be 'better' than another, and we should like to quantify this.
- And, often, the fitness of our approximation depends on how it interacts with the world: the prior we choose, and the dataset we have.
- So the approximation is typically *context-dependent* and *parameterized*.
    - We can capture all this in a category of **statistical games**.

Simpler Statistical Games for Approximate Inference (2/2)

The objects of **SGame** are the objects (X, A) of **BayesLens**.

Then a **statistical game** is a morphism $(X, A) \to (Y, B)$ in **SGame**:
a lens $(X, A) \nrightarrow (Y, B)$ paired with a contextual *loss function* $\operatorname{ctx}\big((X, A), (Y, B)\big) \to \mathbb{R}$.

$\operatorname{ctx}\big((X, A), (Y, B)\big)$ is the set of contexts for lenses $(X, A) \nrightarrow (Y, B)$:

- Everything needed to "close off" the lens.
- $\operatorname{ctx}\big((X, A), (Y, B)\big) = \mathbf{BayesLens}\big((1, 1), (X, A)\big) \times \mathbf{BayesLens}\big((Y, B), (1, 1)\big)$
- That is: a prior $\pi : \mathcal{D}X$ on X and a *continuation* channel $Y \to \mathcal{D}B$.

Composition of statistical games is lens composition paired with the sum of the 'local' fitnesses.
Identities are identity lenses (which just pass on information) with 0 fitness.
And there is a monoidal product inherited from the underlying Markov category.

Free Energy Games

Definition

A (simple) **free energy game** is a simple statistical game $(Z, Z) \to (X, X)$ for some space X with loss function $\phi : \text{ctx}\big((Z, Z), (X, X)\big) \to \mathbb{R}$ given by

$$\phi(\pi, k) = \mathop{\mathbb{E}}_{x\sim k\bullet c\bullet\pi}\left[\mathop{\mathbb{E}}_{z\sim c'_\pi(x)}[-\log p_c(x|z)] + D_{KL}(c'_\pi(x), \pi)\right] = \mathop{\mathbb{E}}_{z\sim c'_\pi\bullet k\bullet c\bullet\pi}\left[\mathcal{F}(x)\right]$$

where $(c, c') : (Z, Z) \nrightarrow (X, X)$ constitutes the lens part of the game, where $p_c : X \times Z \to \mathbb{R}_+$ is a density function for c, and where $\mathcal{F}(x)$ is called the **free energy**.

Nota bene:

- We have $\mathcal{F}(x) = D_{KL}\left[c'_\pi(x), c^\dagger_\pi(x)\right] - \log p_{c\bullet\pi}(x) \geq D_{KL}\left[c'_\pi(x), c^\dagger_\pi(x)\right]$, so free energy games are approximate Bayesian inference games.
- And, in our examples, they are closed under composition.

Overview

Polynomial Morphology (1/2)

Polynomials can represent interfaces;
their morphisms patterns of interconnection:

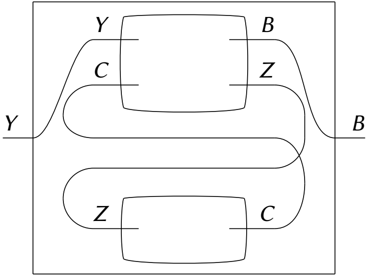

$$BZy^{YC} \otimes Cy^{Z} \to By^{Y}$$

But that's not all!

Polynomial Morphology (1/2)

Polynomials can represent interfaces;
their morphisms patterns of interconnection:

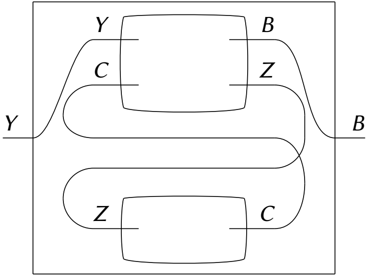

$$BZy^{YC} \otimes Cy^{Z} \to By^{Y}$$

But that’s not all!

Polynomials can describe much richer kinds of interconnection:

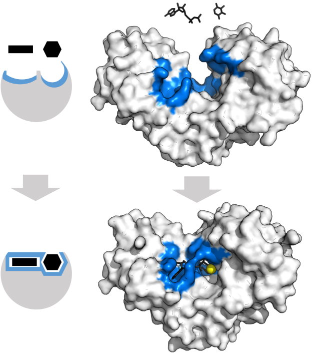

Even systems that change their shape!
How does this work?

Polynomial Morphology (2/2)

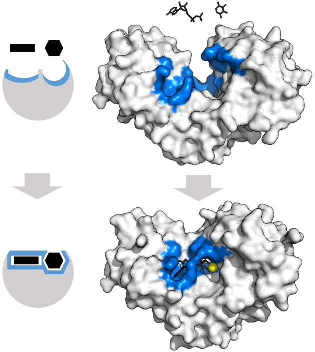

- A general polynomial p looks like: $Ay^X + By^Y + Cy^Z + \cdots$
  - "a sum of representable copresheaves"
- So, that means: $\sum_{i:p(1)} p[i]$
  - For example: $p(1) = A + B + C$; $p[i]_{i \in A} = X$, etc
  - So $p : \sum_{i:p(1)} p[i] \to p(1)$ is a bundle over $p(1)$!
- I like to think of each polynomial as a **phenotype**:
  - The base type $p(1)$ is **configurations** or **morphology**;
  - Each $p[i]$ is the type of **immanent signals** or 'sensorium'.
  - Think eyelid, or "Markov blanket"!
- Then poly. morphisms model interactions as on the left!
  - A 'forwards' map on configurations, and
    a 'backwards' map on sensoria.
- And we can 'nest' polynomials:
  systems within systems...

Now, let's try to animate these gadgets.

Markov Processes and Generalized Coalgebras

Definition (Open Markov Process $(\Theta, \vartheta^o, \vartheta^u)$ on p with time $\mathbb{T}$)

- a state space $\Theta : \mathcal{E}$; and an 'output' map $\vartheta^o : \mathbb{T} \times \Theta \to p(1)$; and
- an 'update' map $\vartheta^u : \sum_{t:\mathbb{T}} \sum_{s:\mathbb{S}} p[\vartheta^o(t,s)] \to \mathcal{P}\, S$

such that, for any section σ of p, the **closure** $\vartheta^\sigma : \mathbb{T} \times \Theta \to \mathcal{P}\, S$ given by

$$\sum_{t:\mathbb{T}} \Theta \xrightarrow{\vartheta^o(t)^* \sigma} \sum_{t:\mathbb{T}} \sum_{s:S} p[\vartheta^o(t,s)] \xrightarrow{\vartheta^u} \mathcal{P}\, \Theta$$

induces an object in the functor category $\mathbf{Cat}\big(\mathbf{B}\mathbb{T}, \mathcal{K}\ell(\mathcal{P})\big)$.

These things:

- form an indexed category $\mathbf{MrkProc}^{\mathbb{T}}_{\mathcal{P}} : \mathbf{Poly}_{\mathcal{E}} \to \mathbf{Cat}$;
- correspond to $p\,\mathcal{P}$-coalgebras when $\mathbb{T} = \mathbb{N}$, and regular Markov processes when $p = y$;
- ... and satisfy various other nice results (for a different talk!)

Overview

Approximate Inference Doctrines

We can use such indexed categories of dynamical systems over polynomials to construct something like a dynamical analogue of our lenses:

Proposition (Hierarchical Bidirectional D-systems)

Suppose $\mathsf{D} : \mathbf{Poly}_{\mathcal{E}} \to \mathbf{Cat}$. There is a category $\mathbf{HiBi}(\mathsf{D})$:

- Objects are pairs (X, A) of objects in $\mathcal{E}$;
- morphisms are functors $\mathsf{D}(Xy^A) \to \mathsf{D}(Yy^B)$;
- composition is just composition of functors.

Approximate Inference Doctrines

We can use such indexed categories of dynamical systems over polynomials to construct something like a dynamical analogue of our lenses:

Proposition (Hierarchical Bidirectional D-systems)

Suppose $\mathrm{D} : \mathbf{Poly}_{\mathcal{E}} \to \mathbf{Cat}$. There is a category $\mathbf{HiBi}(\mathrm{D})$:

- Objects are pairs (X, A) of objects in $\mathcal{E}$;
- morphisms are functors $\mathrm{D}(Xy^A) \to \mathrm{D}(Yy^B)$;
- composition is just composition of functors.

And then we can use this as 'dynamical semantics' for approximate inference:

Definition (Approximate Inference Doctrine)

Let $\mathcal{G}$ be a subcategory of $\mathbf{SGame}_{\mathcal{K}\ell(\mathcal{P})}$. An **approximate inference doctrine** in $\mathcal{G}$ is a lax monoidal functor from $\mathcal{G}$ to $\mathbf{HiBi}\left(\mathbf{MrkProc}^{\mathbb{T}}_{\mathcal{P}}\right)|_{\mathcal{G}}$.

The Laplace Doctrine, now well-typed (1/2)

And with these tools, we begin to have a satisfactory compositional formalism for Friston's free energy framework (amongst other approximate inference schemes):

Theorem (Laplace Doctrine)

Let $\mathcal{G}$ denote the subcategory of free energy games generated by channels with independent Gaussian noise. Then the "Laplace approximation" induces an approximate inference doctrine in $\mathcal{G}$ with time $\mathbb{N}$, $\text{Laplace} : \mathcal{G} \to \mathbf{HiBi}\left(\mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{N}}\right)|_{\mathcal{G}}$. (Illustration on next side!)

(Briefly: "Laplace" means, assume everything is Gaussian, with posteriors of small variance.)
And, of course, different assumptions give different doctrines!

The Laplace Doctrine, now well-typed (2/2)

This is the wiring diagram resulting from the Laplace doctrine, given a game $(X, X) \to (Y, Y)$ and a system emitting priors π on X:

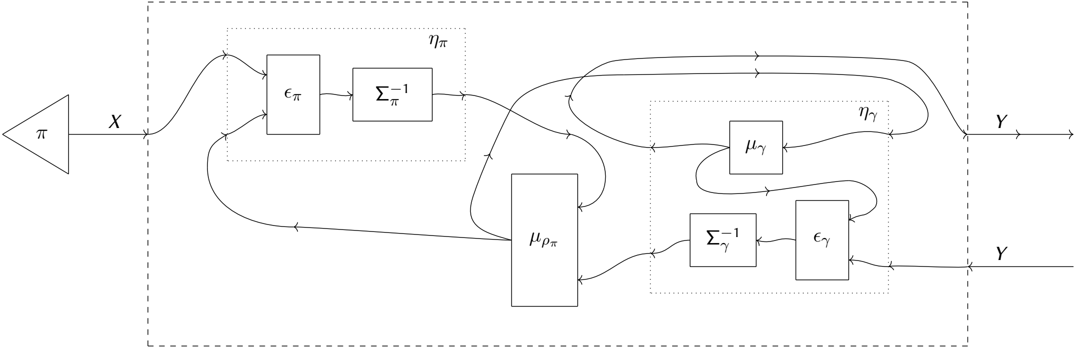

But still, we don't have good notions of *action* and *interaction* here!

Internal Models and Polynomial Statistical Games

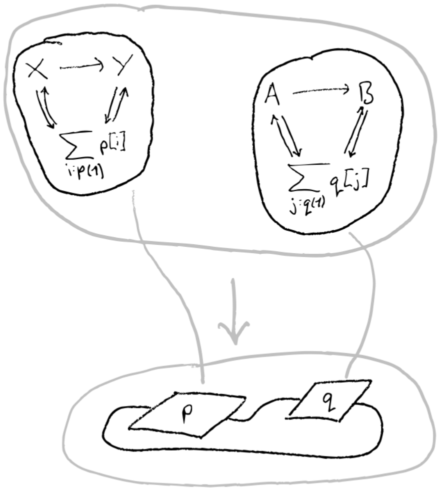

- We can 'index' the category of statistical games by polynomials: like a slice category.
- That is, for each polynomial *p*, we obtain a category of "statistical games over *p*".
- This way, we can model (inter)action:
  - Each object over *p* is an "internal model of *p*-sensations".
  - By sampling from the predicted configurations, we get something like 'action'.
    - (What more is action than a change in configuration?)
- We also get a recipe for describing the generative model of a composite ('multi-agent') system, in terms of the component systems' models.

Polynomial Statistical Games, More Formally

Definition (Simple Statistical Games With Interface $X : \mathcal{E}$)

Define category $\mathbf{IntGame}_{\mathcal{P}}(X)$:

- Objects are simple statistical games with codomain X; that is, points of $\sum_{A:\mathcal{E}} \mathbf{SimpSGame}_{\mathcal{K}\ell(\mathcal{P})}(A, X)$.
- Morphisms $(\gamma, \rho, \phi) \to (\delta, \sigma, \chi)$ are deterministic functions $f : A \to B$—that is, points of $\mathcal{E}(A, B)$—such that $\gamma = \delta \circ f$.

Proposition

There is a polynomially indexed category of statistical games $\mathbf{PSGame}_{\mathcal{P}} : \mathbf{Poly}_{\mathcal{E}} \to \mathbf{Cat}$, defined on objects p as $\mathbf{IntGame}_{\mathcal{P}}\left(\sum_{i:p(1)} p[i]\right)$.

(The action on morphisms of polynomials and the resulting proof are slightly intricate...)

Active Inference Doctrines (1/2)

To breathe life into these games, we define *active inference doctrines.*
Formally, these organize themselves into indexed functors:

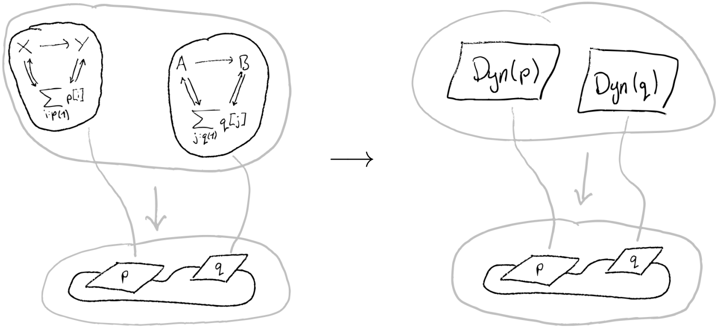

The resulting dynamical systems **perceive** their sensoria, and change their configurations in **action**, to bring their beliefs into alignment with 'reality', and improve their chances of survival, or ability to find abstractions, or their fulfilment, ... or *whatever.*

Active Inference Doctrines (2/2)

We can write this down more formally as follows:

Definition

An **active inference doctrine** is an indexed monoidal functor from (a sub-indexed category of) $\mathbf{PSGame}_{\mathcal{P}}$ to $\mathbf{MrkProc}^{\mathbb{T}}_{\mathcal{P}}$.

And then we have the following result:

Proposition

The Laplace approximate inference doctrine extends to an active inference doctrine.

Overview

Systems with Volition

Recall that a system's performance at a statistical game is contextual:
it depends on prior beliefs (inc. its model structure!), and on its environment.

To improve its performance a system can:

1. Update its beliefs (perception);
2. update its model structure or parameters (part of its beliefs);
   - (we will return to this later)
3. change its shape (a kind of action);
4. couple with part of the world, and change that shape (another kind of action).

By equipping it with 'strong' initial beliefs, we can give it preferences for particular world states.
And the system will attempt to achieve those states: *i.e.*, display **volition**.

Games with Goals, Formally

Definition (Games with Goals on Interface $X : \mathcal{E}$)

Define category $\mathbf{IntGame}_{\mathcal{P}}(X)_*$:

- Objects are dependent pairs in $\sum_{A:\mathcal{E}} 1/\mathcal{K}\ell(\mathcal{P}) \times \mathbf{SimpSGame}_{\mathcal{K}\ell(\mathcal{P})}(A, X)$.
- Morphisms $(\pi, (\gamma, \rho, \phi)) \to (\pi', (\delta, \sigma, \chi))$ are deterministic functions $f : A \to B$ such that $\gamma = \delta \bullet f$ and $\pi = \pi' \bullet f$.

Proposition

There is an indexed category of games with goals $\mathbf{PSGame}_{\mathcal{P}*} : \mathbf{Poly}_{\mathcal{E}} \to \mathbf{Cat}$, defined on polynomials p as $\mathbf{IntGame}_{\mathcal{P}} \left(\sum_{i:p(1)} p[i]\right)_*$.

Definition

An **active inference doctrine with goals** is an indexed monoidal functor from (a sub-indexed category of) $\mathbf{PSGame}_{\mathcal{P}*}$ to $\mathbf{MrkProc}^{\mathbb{T}}_{\mathcal{P}}$.

Homeostasis and Morphogenesis

We sketch a couple of basic examples: (Full workings to follow soon!)

Homeostasis

- Suppose: system's sensorium includes a key parameter, such as the ambient temperature.
- Suppose: by adjusting configuration, the system can move around to sample the parameter.
- Suppose: the 'prior' encodes a high-precision distribution centred on the acceptable range of this parameter.
- Then: by minimizing free energy, the system will attempt to configure itself to remain within the acceptable parameter range.

Homeostasis and Morphogenesis

We sketch a couple of basic examples: (Full workings to follow soon!)

Homeostasis

- Suppose: system's sensorium includes a key parameter, such as the ambient temperature.
- Suppose: by adjusting configuration, the system can move around to sample the parameter.
- Suppose: the 'prior' encodes a high-precision distribution centred on the acceptable range of this parameter.
- Then: by minimizing free energy, the system will attempt to configure itself to remain within the acceptable parameter range.

Morphogenesis (very briefly...)

- Suppose: multiple 'homeostasis' systems, each sensing some signalling molecule.
- Suppose: poly map encodes pattern of signal molecule concentrations.
- Suppose: 'priors' encode "target local configurations".
- Then: free-energy minimization induces systems to arrange themselves to obtain global target pattern.

Overview

Autopoiesis

- An active inference doctrine assigns dynamical systems to statistical games.
- Often, it seems like we should be able to go the other way:
  that is, take an open dynamical system to some "canonical game" that it seems to realize.
  - (Particularly when the dynamical system seems *autopoietic*!)
- I expect this mapping to be adjoint to the doctrine, perhaps as in the triangle below:

$$\begin{array}{ccc} \int \mathbf{PSGame}_{\mathcal{P}} & \underset{\text{Doctrine}}{\overset{\text{CanonicalGame}}{\rightleftarrows}} & \int \mathbf{MrkProc}^{\mathbb{T}}_{\mathcal{P}} \\ & \searrow \quad \swarrow & \\ & \mathbf{Poly}_{\mathcal{E}} & \end{array}$$

- Something like this should work for Ornstein-Uhlenbeck processes.
  But I haven't proved it yet.

Nested Systems and Mutually Coinductive Types

$$
\begin{array}{ccc}
 & \vdots & \\
F & \longrightarrow & C \\
\downarrow & & \downarrow \\
E & \longrightarrow & B \\
\downarrow & & \downarrow \\
J & \longrightarrow & I
\end{array}
$$

- Each polynomial corresponds to a ‘bridge’ diagram $J \leftarrow E \rightarrow B \rightarrow I$.
- Our definition of $\mathbf{Poly}_{\mathcal{E}}$ corresponds to the case when $J = I = 1$.
- We think of this category as considering all systems in the universe together, on equal footing.
- But systems are often ‘nested’ within each other:
  - consider bacteria in your gut microbiome; or
  - creatures on the Earth within the solar system.
- In these cases, we should expect the ‘inner’ and ‘outer’ morphologies to be compatible, which means considering *squares*, not just bridges.
- We then get another indexed category, $\mathbf{Poly}_{\mathcal{E}}(-) : \mathbf{Poly}_{\mathcal{E}}(1) \to \mathbf{Cat}$, and this should in turn be (co)recursive, to model iterated nesting.
  - In turn, this gives a notion of “hierarchical internal model”.
- But what is such a (‘mutually’) coinductive type, precisely?

Connections with ‘Strathclyde’ Cybernetics

Semantics of Parameterization

- At the heart of the MSP group’s account of cybernetics are *parameterized* maps [3].
- These crop up for me, too: for instance, in learning to predict or navigate.
- Indexing over polynomials similarly “adds a dimension” to the algebra of interconnection.
- But there’s also a difference: between indefinitely extended processes, and “process fragments” ...

Nested Systems and Parameters

- The **Para** construction can also be iterated.
- Is that shape the same as my “nested polynomials”?

‘Playing’ Open (Economic) Games

- Open games are like “fragments of a Markov decision process”.
- Active inference systems can optimally solve such problems.
  - (up to a finite time horizon)
- I expect an elegant connection between Bayesian open games and statistical games with goals.

And More...

- More approximate and active inference doctrines! In particular:
  - We can define statistical games over dynamical processes; doctrines here should correspond to 'filtering'.
  - Some of the hierarchical models here seem to demand iterated polynomial nesting...
- Learning the parameters of the model is like learning the 'context' of a dependent type theory, or the base of a slice topos.
  - How far can we push this? (I think, quite far!)
- We haven't made especial use of the functoriality of the polynomial indexing, yet.
  - Can we build compositional models of active multi-agent systems, like corporations?
  - What is the internal model of such a composite system? What are its goals?
- Then, can we use any answers to these questions to study consensus?
  - From a 'fuzzy' or 'Bayesian' angle?
  - Can we apply cohomological tools?

References

[1] Toby St. Clere Smithe. “Bayesian Updates Compose Optically”. In: (05/31/2020). arXiv: 2006.01631v1 [math.CT].

[2] A. M. Bastos et al. “Canonical microcircuits for predictive coding”. In: *Neuron* 76.4 (11/2012), pp. 695–711. DOI: 10.1016/j.neuron.2012.10.038.

[3] Matteo Capucci et al. “Towards foundations of categorical cybernetics”. In: (05/13/2021). arXiv: 2105.06332 [math.CT].