

# Polynomial Life: the Structure of Adaptive Systems

Toby St. Clere Smithe

Topos Institute

toby@topos.institute

We extend our earlier work on the compositional structure of cybernetic systems in order to account for the embodiment of such systems. All their interactions proceed through their bodies' boundaries: sensations impinge on their surfaces, and actions correspond to changes in their configurations. We formalize this morphological perspective using polynomial functors. The 'internal universes' of systems are shown to constitute an indexed category of statistical games over polynomials; their dynamics form an indexed category of behaviours. We characterize *active inference doctrines* as indexed functors between such categories, resolving a number of open problems in our earlier work, and pointing to a formalization of the *free energy principle* as adjoint to such doctrines. We illustrate our framework through fundamental examples from biology, including homeostasis, morphogenesis, and autopoiesis, and suggest a formal connection between spatial navigation and the process of proof.

## 1 Introduction

In a submission to ACT 2020 [15], we presented some first steps towards a theory of categorical cybernetics, motivated by concerns about what gives physical systems life. We explained that perception and action could both be described as processes of Bayesian inference: on the one hand, adjusting beliefs about the world on the basis of observational evidence; on the other, adjusting the world itself in order better to match beliefs. In each case, the system must instantiate a number of structures: a choice of 'prior' belief about the state of the world; a mechanism to generate predictions about sense-data on the basis of that belief, called a 'stochastic channel'; and a (typically approximate) Bayesian inversion of that channel, by which to update those beliefs, in light of sensory observations.

The pairing of a prior with a stochastic channel corresponds to what is called in the informal literature a *generative model*, and it is common to suppose that these models are hierarchical: that is, that the stochastic channel factors as some composite; one imagines that each factor corresponds to predictions at some level of detail, cascading for example down from high-level abstractions to individual photoreceptors. In our earlier submission, we formalized this compositional structure using the bidirectional 'lens' pattern<sup>1</sup> — since predictions and inversions are oppositely directed — and characterized a number of approximate inference processes using a novel category of *statistical games*, whose best responses correspond to optimal inferences. A cybernetic system was then defined as a 'dynamical realisation' of such a game.

This formalism left some things to be desired: our notion of dynamical realisation was ill-defined, and the notion of 'action' was overly abstract. In this submission, we resolve these issues, substantially simplifying our presentation along the way. We explain that various recipes for performing approximate inference, corresponding to our earlier informal notion of dynamical realisation, form functorial *approximate inference doctrines*, between appropriate categories of statistical games and dynamical systems. Then, to formalize a satisfactory notion of action, we note that any active system has a boundary defining

---

<sup>1</sup>See [14] for a pedagogical presentation of the fundamental structures.

its morphology, and that it acts by changing the shape of this boundary; in order to *act on* another system, it couples part of this boundary to that other system, so as to change the composite shape.

To formalize the shapes of systems and their interactions, we adopt polynomial functors: each polynomial will encode the ‘phenotype’ (possible shapes or configurations) of a system, and the sensorium possible in each configuration. To give such systems life, we construct categories of statistical games and dynamical behaviours indexed by polynomials. An *active inference doctrine* is then an indexed functor between such categories. This framework enables a number of possibilities: we can construct a generative model for a corporation on the basis of models for its employees; we give compositional descriptions of fundamental processes of life such as homeostasis and morphogenesis, and point towards an account of autopoiesis; and we sketch the process by which living systems internalize the structures of their environments, and navigate accordingly, noting that such navigation in abstract spaces corresponds precisely to the process of proof.

The work presented here is work in progress, and owing to constraints of space and time, it has not been possible to elaborate everything that we would have liked; this means that we defer proofs of the main results to subsequent elaborations of this extended abstract. Two such elaborations are our series of papers on compositional active inference, beginning with [16], and our work on open dynamical systems with polynomial interfaces [18]. Notwithstanding these limitations, we believe the results here go some of the way to answering the open questions, of elegance and interaction, sketched at the end of our earlier submission. We see this work as making baby steps towards a theory of *embodied* cybernetics.

**Acknowledgements** We thank the members of the Topos Institute for stimulating discussions, and the Foundational Questions Institute and Topos Institute for financial support.

## 2 Simpler Statistical Games

We begin by sketching a refinement of the statistical games formalism developed in [15]. For much more detail, we refer the reader to [16]. First, we recall the bidirectional structure of Bayesian inversion.

### 2.1 Bayesian Lenses

**Definition 2.1** ([11, Def. 3.3]). The category  $\mathbf{GrLens}_F$  of *Grothendieck lenses* for a pseudofunctor  $F : \mathcal{C}^{\text{op}} \rightarrow \mathbf{Cat}$  is the total category of the Grothendieck construction for the pointwise opposite of  $F$ .

**Proposition 2.2** ( $\mathbf{GrLens}_F$  is a category). The objects  $(\mathbf{GrLens}_F)_0$  of  $\mathbf{GrLens}_F$  are (dependent) pairs  $(C, X)$  with  $C : \mathcal{C}$  and  $X : F(C)$ , and its hom-sets  $\mathbf{GrLens}_F((C, X), (C', X'))$  are dependent sums

$$\mathbf{GrLens}_F((C, X), (C', X')) = \sum_{f : \mathcal{C}(C, C')} F(C)(F(f)(X'), X)$$

so that a morphism  $(C, X) \rightarrowtail (C', X')$  is a pair  $(f, f^\dagger)$  of  $f : \mathcal{C}(C, C')$  and  $f^\dagger : F(C)(F(f)(X'), X)$ . We call such pairs *Grothendieck lenses for  $F$*  or  *$F$ -lenses*.

*Proof sketch.* The identity Grothendieck lens on  $(C, X)$  is  $\text{id}_{(C, X)} = (\text{id}_C, \text{id}_X)$ . Sequential composition is as follows. Given  $(f, f^\dagger) : (C, X) \rightarrowtail (C', X')$  and  $(g, g^\dagger) : (C', X') \rightarrowtail (D, Y)$ , their composite  $(g, g^\dagger) \circ (f, f^\dagger)$  is defined to be the lens  $(g \circ f, f^\dagger \circ F(f)(g^\dagger)) : (C, X) \rightarrowtail (D, Y)$ . Associativity and unitality of composition follow from the functoriality of  $F$ .  $\square$

**Definition 2.3.** Suppose  $F(C)_0 = \mathcal{C}_0$ , with  $F : \mathcal{C}^{\text{op}} \rightarrow \mathbf{Cat}$  a pseudofunctor. Define  $\mathbf{SimpGrLens}_F$  to be the full subcategory of  $\mathbf{GrLens}_F$  whose objects are duplicate pairs  $(C, C)$  of objects  $C$  in  $\mathcal{C}$ . We call  $\mathbf{SimpGrLens}_F$  the category of *simple*  $F$ -lenses. More generally, any lens between such duplicate pairs will be called a *simple lens*. Since duplicating the objects in the pairs  $(C, C)$  is redundant, we will write the objects simply as  $C$ .

**Definition 2.4.** Let  $(\mathcal{C}, \otimes, I)$  be a monoidal category enriched in a Cartesian closed category  $\mathbf{V}$ . Define the  $\mathcal{C}$ -state-indexed category  $\text{Stat} : \mathcal{C}^{\text{op}} \rightarrow \mathbf{V-Cat}$  as follows.

$$\text{Stat} : \mathcal{C}^{\text{op}} \rightarrow \mathbf{V-Cat}$$

$$X \mapsto \text{Stat}(X) := \left( \begin{array}{lll} \text{Stat}(X)_0 & := & \mathcal{C}_0 \\ \text{Stat}(X)(A, B) & := & \mathbf{V}(\mathcal{C}(I, X), \mathcal{C}(A, B)) \\ \text{id}_A : \text{Stat}(X)(A, A) & := & \begin{cases} \text{id}_A : \mathcal{C}(I, X) \rightarrow \mathcal{C}(A, A) \\ \rho \mapsto \text{id}_A \end{cases} \end{array} \right) \quad (1)$$

$$f : \mathcal{C}(Y, X) \mapsto \left( \begin{array}{lll} \text{Stat}(f) : \text{Stat}(X) & \rightarrow & \text{Stat}(Y) \\ \text{Stat}(X)_0 & = & \text{Stat}(Y)_0 \\ \mathbf{V}(\mathcal{C}(I, X), \mathcal{C}(A, B)) & \rightarrow & \mathbf{V}(\mathcal{C}(I, Y), \mathcal{C}(A, B)) \\ \alpha & \mapsto & f^* \alpha : (\sigma : \mathcal{C}(I, Y)) \mapsto (\alpha(f \bullet \sigma) : \mathcal{C}(A, B)) \end{array} \right)$$

Composition in each fibre  $\text{Stat}(X)$  is as in  $\mathcal{C}$ . Explicitly, indicating morphisms  $\mathcal{C}(I, X) \rightarrow \mathcal{C}(A, B)$  in  $\text{Stat}(X)$  by  $A \xrightarrow{X} B$ , and given  $\alpha : A \xrightarrow{X} B$  and  $\beta : B \xrightarrow{X} C$ , their composite is  $\beta \circ \alpha : A \xrightarrow{X} C := \rho \mapsto \beta(\rho) \bullet \alpha(\rho)$ , where here we indicate composition in  $\mathcal{C}$  by  $\bullet$  and composition in the fibres  $\text{Stat}(X)$  by  $\circ$ . Given  $f : Y \rightarrow X$  in  $\mathcal{C}$ , the induced functor  $\text{Stat}(f) : \text{Stat}(X) \rightarrow \text{Stat}(Y)$  acts by pullback.

**Example 2.5.** Typically, a choice of  $\mathcal{C}$  above will be the Kleisli category  $\mathcal{Kl}(\mathcal{P})$  of a probability monad, such as the Giry monad  $\mathcal{G} : \mathbf{Meas} \rightarrow \mathbf{Meas}$  on measurable spaces; and  $\mathbf{V}$  will either be  $\mathbf{Set}$  or (better) some ‘nice’ category of measurable spaces. However, the category  $\mathbf{Meas}$  of general measurable spaces is not Cartesian closed, as there is no general way to make the evaluation maps  $\mathbf{Meas}(X, Y) \times X \rightarrow Y$  measurable, meaning that if we take  $\mathcal{C} = \mathcal{Kl}(\mathcal{G})$  above, then we are forced to take  $\mathbf{V} = \mathbf{Set}$ . In turn, this makes the inversion maps  $c^\dagger : \mathcal{Kl}(\mathcal{G})(1, X) \rightarrow \mathcal{Kl}(\mathcal{G})(Y, X)$  into mere functions. We can salvage measurability by working instead with  $\mathcal{Kl}(\mathcal{Q})$ , where  $\mathcal{Q} : \mathbf{QBS} \rightarrow \mathbf{QBS}$  is the analogue of the Giry monad for quasi-Borel spaces [10]. The category  $\mathbf{QBS}$  is indeed Cartesian closed, and  $\mathcal{Kl}(\mathcal{Q})$  is enriched in  $\mathbf{QBS}$ , so that we can instantiate  $\text{Stat}$  there, and the corresponding inversion maps are accordingly measurable.

We define the category of Bayesian lenses in  $\mathcal{C}$  to be the category of  $\text{Stat}$ -lenses.

**Definition 2.6.** The category  $\mathbf{BayesLens}_{\mathcal{C}}$  of Bayesian lenses in  $\mathcal{C}$  is the category  $\mathbf{GrLens}_{\text{Stat}}$  of Grothendieck lenses for the functor  $\text{Stat}$ . A *Bayesian lens* is a morphism in  $\mathbf{BayesLens}_{\mathcal{C}}$ . Where the category  $\mathcal{C}$  is evident from the context, we will just write  $\mathbf{BayesLens}$ .

Unpacking this definition, we find that the objects of  $\mathbf{BayesLens}_{\mathcal{C}}$  are pairs  $(X, A)$  of objects of  $\mathcal{C}$ . Morphisms (that is, Bayesian lenses)  $(X, A) \rightarrow (Y, B)$  are pairs  $(c, c^\dagger)$  of a channel  $c : X \rightarrow Y$  and a “generalized Bayesian inversion”  $c^\dagger : B \xrightarrow{X} A$ ; that is, elements of the hom objects

$$\begin{aligned} \mathbf{BayesLens}_{\mathcal{C}}((X, A), (Y, B)) &:= \mathbf{GrLens}_{\text{Stat}}((X, A), (Y, B)) \\ &\cong \mathcal{C}(X, Y) \times \mathbf{V}(\mathcal{C}(I, X), \mathcal{C}(B, A)). \end{aligned}$$

The identity Bayesian lens on  $(X, A)$  is  $(\text{id}_X, \text{id}_A)$ , where by abuse of notation  $\text{id}_A : \mathcal{C}(I, Y) \rightarrow \mathcal{C}(A, A)$  is the constant map  $\text{id}_A$  defined in equation (1) that takes any state on  $Y$  to the identity on  $A$ .

The sequential composite  $(d, d^\dagger) \circ (c, c^\dagger)$  of  $(c, c^\dagger) : (X, A) \leftrightarrow (Y, B)$  and  $(d, d^\dagger) : (Y, B) \leftrightarrow (Z, C)$  is the Bayesian lens  $((d \bullet c), (c^\dagger \circ c^* d^\dagger)) : (X, A) \leftrightarrow (Z, C)$  where  $(c^\dagger \circ c^* d^\dagger) : C \xrightarrow{X} A$  takes a state  $\pi : I \rightarrow X$  to the channel  $c_\pi^\dagger \bullet d_{c \bullet \pi}^\dagger : C \rightarrow A$ .

**Definition 2.7.** Given a Bayesian lens  $(c, c') : (X, A) \leftrightarrow (Y, B)$ , we will call  $c$  its *forwards* or *prediction* channel and  $c'$  its *backwards* or *update* channel (even though  $c'$  is really a family of channels).

**Definition 2.8** (Bayesian inversion). We say that a channel  $c : X \rightarrow Y$  admits *Bayesian inversion* with respect to  $\pi : I \rightarrow X$  if there exists a channel  $c_\pi^\dagger : Y \rightarrow X$ , called the *Bayesian inversion* of  $c$  with respect to  $\pi$ , satisfying the following equation [6, eq. 5] in the graphical calculus of  $\mathcal{C}$ :

$$(2)$$

We say that  $c$  admits Bayesian inversion *tout court* if  $c$  admits Bayesian inversion with respect to all states  $\pi : I \rightarrow X$  such that  $c \bullet \pi$  has non-empty support. We say that a category  $\mathcal{C}$  admits Bayesian inversion if all its morphisms admit Bayesian inversion *tout court*.

**Definition 2.9.** We call the pairing  $(\pi, c)$  of a state  $\pi : I \rightarrow X$  with a channel  $c : X \rightarrow Y$  a *generative model*  $X \rightarrow Y$ . It induces a joint distribution  $\omega_{(\pi, c)} := (\text{id}_X \otimes c) \bullet \text{y}_X \bullet \pi : I \rightarrow X \otimes Y$ .

**Definition 2.10.** Let  $(c, c^\dagger) : (X, X) \leftrightarrow (Y, Y)$  be a Bayesian lens. We say that  $(c, c^\dagger)$  is *exact* if  $c$  admits Bayesian inversion and, for each  $\pi : I \rightarrow X$  such that  $c \bullet \pi$  has non-empty support,  $c$  and  $c_\pi^\dagger$  together satisfy equation (2). Bayesian lenses that are not exact are said to be *approximate*.

**Definition 2.11** (Almost-equality). Given a state  $\pi : I \rightarrow X$ , we say that two parallel channels  $c, d : X \rightarrow Y$  are  $\pi$ -almost-equal, denoted  $c \stackrel{\pi}{\sim} d$ , if the joint distributions of the two generative models  $(\pi, c)$  and  $(\pi, d)$  are equal; that is, if  $(\text{id} \otimes c) \bullet \text{y} \bullet \pi = (\text{id} \otimes d) \bullet \text{y} \bullet \pi$ .

**Theorem 2.12** ([14]). Let  $(c, c^\dagger)$  and  $(d, d^\dagger)$  be sequentially composable exact Bayesian lenses. Then the contravariant component of the composite lens  $(d, d^\dagger) \circ (c, c^\dagger) = (d \bullet c, c^\dagger \circ c^* d^\dagger)$  is, up to  $d \bullet c \bullet \pi$ -almost-equality, the Bayesian inversion of  $d \bullet c$  with respect to any state  $\pi$  on the domain of  $c$  such that  $c \bullet \pi$  has non-empty support.

## 2.2 Statistical Games

The performance of a statistical or cybernetic system depends upon its interaction with its environment, and the prior beliefs that it started with. We will therefore define a statistical game to be a Bayesian lens paired with a fitness function measuring performance in context; and for this we need a notion of context. For this notion, we take inspiration from compositional game theory [2].

**Definition 2.13.** A *context* for a Bayesian lens is an element of the profunctor  $\overline{\mathbf{BayesLens}}_{\mathcal{C}}$  defined by

$$\begin{array}{ccc} \mathbf{BayesLens}_{\mathcal{C}} & \times & \mathbf{BayesLens}_{\mathcal{C}}^{\text{op}} \rightarrow \mathbf{V} \\ - & \times & = \mapsto \mathbf{BayesLens}_{\mathcal{C}}((I, I), -) \times \mathbf{BayesLens}_{\mathcal{C}}(=, (I, I)) \end{array}$$

where  $I$  is the monoidal unit in  $\mathcal{C}$ ,  $\mathbf{BayesLens}_{\mathcal{C}}((I, I), -) : \mathbf{BayesLens}_{\mathcal{C}} \rightarrow \mathbf{V}$  is the representable copresheaf on  $(I, I)$ , and  $\mathbf{BayesLens}_{\mathcal{C}}(=, (I, I)) : \mathbf{BayesLens}_{\mathcal{C}}^{\text{op}} \rightarrow \mathbf{V}$  is the representable presheaf. Hence a context for a lens  $(X, A) \rightarrow (Y, B)$  is an element of

$$\overline{\mathbf{BayesLens}}_{\mathcal{C}}((X, A), (Y, B)) := \mathbf{BayesLens}_{\mathcal{C}}((I, I), (X, A)) \times \mathbf{BayesLens}_{\mathcal{C}}((Y, B), (I, I)).$$

**Proposition 2.14.** When  $I$  is terminal in  $\mathcal{C}$  and the base  $\mathbf{V}$  of enrichment of  $\mathcal{C}$  is well pointed (as when  $\mathbf{V} = \mathbf{Set}$ ), we have  $\overline{\mathbf{BayesLens}}_{\mathcal{C}}((X, A), (Y, B)) \cong \mathcal{C}(I, X) \times \mathbf{V}(\mathcal{C}(I, B), \mathcal{C}(I, Y))$ .

*Proof.* A straightforward calculation which we omit: use that  $X \xrightarrow{c} Y \xrightarrow{!} I = X \xrightarrow{!} I$  for any  $c : X \rightarrow Y$ .  $\square$

**Proposition 2.15.** Given a context  $(\pi, k) : \overline{\mathbf{BayesLens}}_{\mathcal{C}}((X, A), (Z, C))$  for a composite lens  $(X, A) \xrightarrow{f} (Y, B) \xrightarrow{g} (Z, C)$ , we obtain contexts for the factors:

$$\begin{aligned} \overline{\mathbf{BayesLens}}_{\mathcal{C}}((X, A), g)(\pi, k) &= (\pi, k \circ g) : \overline{\mathbf{BayesLens}}_{\mathcal{C}}((X, A), (Y, B)), \\ \overline{\mathbf{BayesLens}}_{\mathcal{C}}(f, (Z, C))(\pi, k) &= (f \circ \pi, k) : \overline{\mathbf{BayesLens}}_{\mathcal{C}}((Y, B), (Z, C)). \end{aligned}$$

We are now in a position to define the category of statistical games over  $\mathcal{C}$ .

**Proposition 2.16.** Let  $\mathcal{C}$  be a  $\mathbf{V}$ -category admitting Bayesian inversion and let  $(R, +, 0)$  be a monoid in  $\mathbf{V}$ . Then there is a category  $\mathbf{SGame}_{[R]}^{\mathcal{C}}$  whose objects are the objects of  $\mathbf{BayesLens}_{\mathcal{C}}$  and whose morphisms  $(X, A) \rightarrow (Y, B)$  are *statistical games*: pairs  $(f, \phi)$  of a lens  $f : \mathbf{BayesLens}_{\mathcal{C}}((X, A), (Y, B))$  and a *fitness function*  $\phi : \overline{\mathbf{BayesLens}}_{\mathcal{C}}((X, A), (Y, B)) \rightarrow R$ . When  $R$  is the monoid of reals  $\mathbb{R}$ , then we just denote the category by  $\mathbf{SGame}_{\mathcal{C}}$ .

*Proof.* Suppose given statistical games  $(f, \phi) : (X, A) \rightarrow (Y, B)$  and  $(g, \psi) : (Y, B) \rightarrow (Z, C)$ . We seek a composite game  $(g, \psi) \circ (f, \phi) := (gf, \psi\phi) : (X, A) \rightarrow (Z, C)$ . We have  $gf = g \circ f$  by lens composition. Proposition 2.15 gives us a family of functions  $\text{localCtx}$  with signature

$$\begin{aligned} &\overline{\mathbf{BayesLens}}_{\mathcal{C}}((X, A), (Z, C)) \times \mathbf{BayesLens}_{\mathcal{C}}((X, A), (Y, B)) \times \mathbf{BayesLens}_{\mathcal{C}}((Y, B), (Z, C)) \\ &\rightarrow \overline{\mathbf{BayesLens}}_{\mathcal{C}}((X, A), (Y, B)) \times \overline{\mathbf{BayesLens}}_{\mathcal{C}}((Y, B), (Z, C)). \end{aligned}$$

We therefore identify the composite fitness function  $\psi\phi$  as

$$\psi\phi := + \circ (\phi, \psi) \circ \text{localCtx}(-, f, g)$$

where  $+$  :  $R \times R \rightarrow R$  is the monoid operation. The identity game  $(X, A) \rightarrow (X, A)$  is given by  $(\text{id}, 0)$ , the pairing of the identity lens on  $(X, A)$  with the unit  $0$  of the monoid  $R$ . Associativity and unitality are immediate from lens composition and the monoid laws.  $\square$

**Definition 2.17.** We will write  $\mathbf{SimpSGame}_{\mathcal{C}} \hookrightarrow \mathbf{SGame}_{\mathcal{C}}$  for the full subcategory of  $\mathbf{SGame}_{\mathcal{C}}$  defined on simple Bayesian lenses  $(X, X) \rightarrow (Y, Y)$ . As in the case of simple lenses (Definition 2.3), we will eschew redundancy by writing the objects  $(X, X)$  simply as  $X$ .

We now present some key examples of statistical games.

**Example 2.18** (Bayesian inference). Let  $D : \mathcal{C}(I, X) \times \mathcal{C}(I, X) \rightarrow \mathbb{R}$  be a measure of divergence between states on  $X$ . Then a (simple) *D*-Bayesian inference game is a statistical game  $(X, X) \rightarrow (Y, Y)$  with fitness function  $\phi : \overline{\text{BayesLens}}_{\mathcal{C}}((X, X), (Y, Y)) \rightarrow \mathbb{R}$  given by  $\phi(\pi, k) = \mathbb{E}_{y \sim k \bullet c \bullet \pi} \left[ D(c'_\pi(y), c_\pi^\dagger(y)) \right]$ , where  $(c, c')$  constitutes the lens part of the game and  $c_\pi^\dagger$  is the exact inversion of  $c$  with respect to  $\pi$ .

Note that we say that  $D$  is a “measure of divergence between states on  $X$ ”. By this we mean any function of the given type with the semantical interpretation that it acts like a distance measure between states. But this is not to say that  $D$  is a metric or even pseudometric. One usually requires that  $D(\pi, \pi') = 0 \iff \pi = \pi'$ , but typical choices do not also satisfy symmetry nor subadditivity. An important such typical choice is the *relative entropy* or *Kullback-Leibler divergence*, denoted  $D_{KL}$ .

**Definition 2.19.** The *Kullback-Leibler divergence*  $D_{KL} : \mathcal{C}(I, X) \times \mathcal{C}(I, X) \rightarrow \mathbb{R}$  is defined by

$$D_{KL}(\alpha, \beta) := \mathbb{E}_{x \sim \alpha} [\log p(x)] - \mathbb{E}_{x \sim \alpha} [\log q(x)]$$

where  $p$  and  $q$  are density functions corresponding to the states  $\alpha$  and  $\beta$ .

Typically, computing  $D_{KL}(c'_\pi(x), c_\pi^\dagger(x))$  is computationally difficult, so one resorts to optimizing an upper bound; a prominent choice is the *free energy* [8].

**Definition 2.20** (*D*-free energy). Let  $(\pi, c)$  be a generative model with  $c : X \rightarrow Y$ . Let  $p_c : Y \times X \rightarrow \mathbb{R}_+$  and  $p_\pi : X \rightarrow \mathbb{R}_+$  be density functions corresponding to  $c$  and  $\pi$ . Let  $p_{c \bullet \pi} : Y \rightarrow \mathbb{R}_+$  be a density function for the composite  $c \bullet \pi$ . Let  $c'_\pi$  be a channel  $Y \rightarrow X$  that we take to be an approximation of the Bayesian inversion of  $c$  with respect to  $\pi$  and that admits a density function  $q : X \times Y \rightarrow \mathbb{R}_+$ . Finally, let  $D : \mathcal{C}(I, X) \times \mathcal{C}(I, X) \rightarrow \mathbb{R}$  be a measure of divergence between states on  $X$ . Then the *D*-free energy of  $c'_\pi$  with respect to the generative model given an observation  $y : Y$  is the quantity

$$\mathcal{F}_D(c'_\pi, c, \pi, y) := \mathbb{E}_{x \sim c'_\pi(y)} [-\log p_c(y|x)] + D(c'_\pi(y), \pi). \quad (3)$$

We will elide the dependence on the model when it is clear from the context, writing only  $\mathcal{F}_D(y)$ .

The *D*-free energy is an upper bound on  $D$  when  $D$  is the relative entropy  $D_{KL}$ .

**Proposition 2.21** (Evidence upper bound). The  $D_{KL}$ -free energy satisfies the following equality:

$$\mathcal{F}_{D_{KL}}(y) = D_{KL}[c'_\pi(y), c_\pi^\dagger(y)] - \log p_{c \bullet \pi}(y) = \mathbb{E}_{x \sim c'_\pi(y)} \left[ \log \frac{q(x|y)}{p_c(y|x) \cdot p_\pi(x)} \right]$$

Since  $\log p_{c \bullet \pi}(y)$  is always negative, the free energy is an upper bound on  $D_{KL}[c'_\pi(y), c_\pi^\dagger(y)]$ , where  $c_\pi^\dagger$  is the exact Bayesian inversion of the channel  $c$  with respect to the prior  $\pi$ . Similarly, the free energy is an upper bound on the negative log-likelihood  $-\log p_{c \bullet \pi}(y)$ . Thinking of this latter quantity as a measure of the “model evidence” gives us the alternative name *evidence upper bound* for the  $D_{KL}$ -free energy.

Any system that performs (approximate) Bayesian inversion can thus be seen as minimizing some free energy. The *free energy principle* says that all it means to be an adaptive system is to embody a process of approximate inference in this way. We therefore define free energy games:

**Example 2.22** (Free energy game). Let  $D : \mathcal{C}(I, X) \times \mathcal{C}(I, X) \rightarrow \mathbb{R}$  be a measure of divergence between states on  $X$ . Then a simple *D*-free energy game is a simple statistical game  $(X, X) \rightarrow (Y, Y)$  with fitness function  $\phi : \overline{\text{BayesLens}}_{\mathcal{C}}((X, X), (Y, Y)) \rightarrow \mathbb{R}$  given by  $\phi(\pi, k) = \mathbb{E}_{y \sim k \bullet c \bullet \pi} [\mathcal{F}_D(c'_\pi, c, \pi, y)]$  where  $(c, c') : (X, X) \rightarrow (Y, Y)$  constitutes the lens part of the game.

**Remark 2.23.** It is also often of interest to consider *parameterized* channels, for which we can use the **Para** construction [13]. This acts by adjoining an object of parameters to the domain of the category at hand (such as the category of Bayesian lenses), and tensoring parameters of composite morphisms. Formally, this corresponds to a generalization of the indexed category of state-dependent morphisms, and forms the subject of another paper in these proceedings [4]. The parameters might represent the ‘weights’ of a neural network, or encode some structure about the possible predictions. Lacking the space to do justice to this structure here, we nonetheless leave it at that, and refer the reader to our paper [16] for more details.

### 3 Systems Within Interfaces; Worlds Within Worlds

In this section, we develop the structures required for extending the formalism of statistical games to embodied systems.

#### 3.1 Polynomials for Embodiment and Interaction

Each system in our universe inhabits some interface or boundary. It receives signals from its environments through this boundary, and can act by changing its shape (and, as we will see later, its position). As a system changes its shape, the set of possible immanent signals might change accordingly: consider a hedgehog rolling itself into a ball, thereby protecting its soft underbelly from harm (amongst other immanent signals). A system may also change its shape by coupling itself to some other system, such as when we pick up chalk to work through a problem. And shapes can be abstract: we change our ‘shapes’ when we enter an online video conference, or move within a virtual reality. We describe all of these interactions formally using polynomial functors, drawing on the work of [12].

**Definition 3.1.** Let  $\mathcal{E}$  be a locally Cartesian closed category, and denote by  $y^A$  the representable copresheaf  $y^A := \mathcal{E}(A, -) : \mathcal{E} \rightarrow \mathcal{E}$ . A *polynomial functor*  $p$  is a coproduct of representable functors, written  $p := \sum_{i:p(1)} y^{p_i}$ , where  $p(1) : \mathcal{E}$  is the indexing object. The category of polynomial functors in  $\mathcal{E}$  is the full subcategory  $\mathbf{Poly}_{\mathcal{E}} \hookrightarrow [\mathcal{E}, \mathcal{E}]$  of the  $\mathcal{E}$ -copresheaf category spanned by coproducts of representables. A morphism of polynomials is therefore a natural transformation.

**Remark 3.2.** Every polynomial functor  $P : \mathcal{E} \rightarrow \mathcal{E}$  corresponds to a bundle  $p : E \rightarrow B$  in  $\mathcal{E}$ , for which  $B = P(1)$  and for each  $i : P(1)$ , the fibre  $p_i$  is  $P(i)$ . We will henceforth elide the distinction between a copresheaf  $P$  and its corresponding bundle  $p$ , writing  $p(1) := B$  and  $p[i] := p_i$ , where  $E = \sum_i p[i]$ . A natural transformation  $f : p \rightarrow q$  between copresheaves therefore corresponds to a map of bundles. In the case of polynomials, by the Yoneda lemma, this map is given by a ‘forwards’ map  $f_1 : p(1) \rightarrow q(1)$  and a family of ‘backwards’ maps  $f^\# : q[f_1(-)] \rightarrow p[-]$  indexed by  $p(1)$ , as in the left diagram below. Given  $f : p \rightarrow q$  and  $g : q \rightarrow r$ , their composite  $g \circ f : p \rightarrow r$  is as in the right diagram below.

$$\begin{array}{ccc}
 E & \xleftarrow{f^\#} f^* F & \longrightarrow F \\
 p \downarrow & \downarrow & \downarrow q \\
 B & \xlongequal{\quad} B & \xrightarrow{f_1} C
 \end{array}
 \qquad
 \begin{array}{ccc}
 E & \xleftarrow{(gf)^\#} f^* g^* G & \longrightarrow G \\
 p \downarrow & \downarrow & \downarrow r \\
 B & \xlongequal{\quad} B & \xrightarrow{g_1 \circ f_1} D
 \end{array}$$

Here,  $(gf)^\#$  is given by the  $p(1)$ -indexed family of composite maps  $r[g_1(f_1(-))] \xrightarrow{f^* g^\#} q[f_1(-)] \xrightarrow{f^\#} p[-]$ .

In our morphological semantics, we will call a polynomial  $p$  a *phenotype*, its base type  $p(1)$  its *morphology* and the total space  $\sum_i p[i]$  its *sensorium*. We will call elements of the morphology *shapes* or *configurations*, and elements of the sensorium *immanent signals*.

**Proposition 3.3** ([12]). There is a monoidal structure  $(\mathbf{Poly}_{\mathcal{E}}, \otimes, y)$  that we interpret as “putting systems in parallel”. Given  $p : \sum_i p[i] \rightarrow p(1)$  and  $q : \sum_j q[j] \rightarrow q(1)$ , we have  $p \otimes q = \sum_i \sum_j p[i] \times q[j] \rightarrow p(1) \times q(1)$ .  $y : 1 \rightarrow 1$  is then clearly unital.  $\square$

**Proposition 3.4** ([12]). The monoidal structure  $(\mathbf{Poly}_{\mathcal{E}}, \otimes, y)$  is closed, with corresponding internal hom denoted  $[-, -]$ .

We interpret morphisms  $(f_1, f^\#)$  of polynomials as encoding interaction patterns; in particular, such morphisms encode how composite systems act as unities. For example, a morphism  $f : p \otimes q \rightarrow r$  specifies how the systems  $p$  and  $q$  come together to form a system  $r$ : the map  $f_1$  encodes how  $r$ -configurations are constructed from configurations of  $p$  and  $q$ ; and the map  $f^\#$  encodes how immanent signals on  $p$  and  $q$  result from signals on  $r$  or from the interaction of  $p$  and  $q$ . For intuition, consider two people engaging in a handshake, or an enzyme acting on a protein to form a complex. The internal hom  $[o, p]$  encodes all the possible ways that an  $o$ -phenotype system can “plug into” a  $p$ -phenotype system.

**Remark 3.5.** In the literature on active inference and the free energy principle, there is much debate about the concept and role of ‘Markov blankets’, which are sometimes conceived to represent the boundary of an adaptive system. We believe that the algebra of polynomials will help to make such thinking precise, where it departs from the established usage of Markov blankets in Bayesian networks.

### 3.2 Dynamical Systems on Polynomial Interfaces

Although some dynamical systems can be modelled within  $\mathbf{Poly}_{\mathcal{E}}$  itself, we prefer to take a fibrational perspective, following the idea that polynomials represent the boundaries of systems, separating ‘internal’ states from ‘external’. We will therefore adopt a pattern of indexing categories by polynomials: in the case of dynamics, the fibre over a polynomial will be a category of possible internal systems, whose projection forgets the internal structure. We can only sketch the relevant structures here, and so refer the reader to our subsequent paper [18] where we give a general account of open dynamical systems as coalgebras for polynomial functors.

**Definition 3.6.** Let  $\mathcal{P} : \mathcal{E} \rightarrow \mathcal{E}$  be a probability monad on the category  $\mathcal{E}$ , and let  $p : \mathbf{Poly}_{\mathcal{E}}$  be a polynomial in  $\mathcal{E}$ . Let  $(\mathbb{T}, +, 0)$  be a monoid in  $\mathcal{E}$ , representing time. Then an **open Markov process** on the interface  $p$  with time  $\mathbb{T}$  consists in a triple  $\vartheta := (S, \vartheta^o, \vartheta^u)$  of a **state space**  $S : \mathcal{E}$  and two morphisms  $\vartheta^o : \mathbb{T} \times S \rightarrow p(1)$  and  $\vartheta^u : \sum_{t:\mathbb{T}} \sum_{s:S} p[\vartheta^o(t, s)] \rightarrow \mathcal{P}S$ , such that for any section  $\sigma : p(1) \rightarrow \sum_{i:p(1)} p[i]$  of  $p$ , the maps  $\vartheta^\sigma : \mathbb{T} \times S \rightarrow \mathcal{P}S$  given by

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

constitute an object in the functor category  $\mathbf{Cat}(\mathbf{BT}, \mathcal{Kl}(\mathcal{P}))$ , where  $\mathbf{BT}$  is the delooping of  $\mathbb{T}$  and  $\mathcal{Kl}(\mathcal{P})$  is the Kleisli category of  $\mathcal{P}$ . We call  $\vartheta^\sigma$  the **closure** of  $\vartheta$  by  $\sigma$ .

**Proposition 3.7.** Open Markov processes over  $p$  with time  $\mathbb{T}$  form a category, denoted  $\mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{T}}(p)$ . Its morphisms are defined as follows. Let  $\vartheta := (X, \vartheta^o, \vartheta^u)$  and  $\psi := (Y, \psi^o, \psi^u)$  be two Markov processes over  $p$ . A morphism  $f : \vartheta \rightarrow \psi$  consists in a morphism  $f : X \rightarrow Y$  such that, for any time  $t : \mathbb{T}$  and global section  $\sigma : p(1) \rightarrow \sum_{i:p(1)} p[i]$  of  $p$ , we have a natural transformation  $\vartheta^\sigma \rightarrow \psi^\sigma$  between the closures. The identity morphism  $\text{id}_\vartheta$  on the open Markov process  $\vartheta$  is given by the identity morphism  $\text{id}_X$  on its state space  $X$ . Composition of morphisms of open Markov processes is given by composition of the morphisms of the state spaces.



**Proposition 3.8.**  $\mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{T}}$  extends to an indexed category,  $\mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{T}} : \mathbf{Poly}_{\mathcal{E}} \rightarrow \mathbf{Cat}$ . Suppose  $\varphi : p \rightarrow q$  is a morphism of polynomials. We define the functor  $\mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{T}}(\varphi) : \mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{T}}(p) \rightarrow \mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{T}}(q)$  as follows. Suppose  $(X, \vartheta^o, \vartheta^u)$  is an object (open Markov process) in  $\mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{T}}(p)$ . Then  $\mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{T}}(\varphi)(X, \vartheta^o, \vartheta^u)$  is defined as the triple  $(X, \varphi_1 \circ \vartheta^o, \vartheta^u \circ \vartheta^{o*} \varphi^\#) : \mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{T}}(q)$ . On morphisms,  $\mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{T}}(\varphi)(f) : \mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{T}}(\varphi)(X, \vartheta^o, \vartheta^u) \rightarrow \mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{T}}(\varphi)(Y, \psi^o, \psi^u)$  is given by the same underlying map  $f : X \rightarrow Y$  of state spaces.

**Remark 3.9** (Closed Markov chains and Markov processes). A closed *Markov chain* is given by a map  $X \rightarrow \mathcal{P}X$ , where  $\mathcal{P} : \mathcal{E} \rightarrow \mathcal{E}$  is a probability monad on  $\mathcal{E}$ ; this is equivalently an object in  $\mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{T}}(y)$ , and (again equivalently) an object in  $\mathbf{Cat}(\mathbf{BN}, \mathcal{Hl}(\mathcal{P}))$ . With more general time  $\mathbb{T}$ , one obtains closed *Markov processes*: objects in  $\mathbf{Cat}(\mathbf{BT}, \mathcal{Hl}(\mathcal{P}))$ . More explicitly, a closed Markov process is a time-indexed family of Markov kernels; that is, a morphism  $\vartheta : \mathbb{T} \times X \rightarrow \mathcal{P}X$  such that, for all times  $s, t : \mathbb{T}$ ,  $\vartheta_{s+t} = \vartheta_s \bullet \vartheta_t$  as a morphism in  $\mathcal{Hl}(\mathcal{P})$ . Note that composition  $\bullet$  in  $\mathcal{Hl}(\mathcal{P})$  is given by the Chapman-Kolmogorov equation, so this means that

$$\vartheta_{s+t}(y|x) = \int_{x':X} \vartheta_s(y|x') \vartheta_t(dx'|x).$$

We will need to convert our fibration of Markov processes into an ordinary category, in order to supply ‘dynamical semantics’ for statistical games. We can do so, with objects being appropriately typed pairs and morphisms representing appropriately ‘bidirectional’ dynamical systems, by using the following structure.

**Proposition 3.10.** Let  $D : \mathbf{Poly}_{\mathcal{E}} \rightarrow \mathbf{Cat}$  be an indexed category over polynomials. Then there is a category of hierarchical bidirectional  $D$ -systems, denoted  $\mathbf{HiBi}(D)$ , and defined as follows. The objects of  $\mathbf{HiBi}(D)$  are pairs  $(X, A)$  of objects in  $\mathcal{E}$ . Morphisms  $(X, A) \rightarrow (Y, B)$  are functors  $D(Xy^A) \rightarrow D(Yy^B)$ . The composition rule is just composition of the corresponding functors, and identity morphisms  $\mathrm{id}_{(X,A)} : (X, A) \rightarrow (X, A)$  are given by the identity functor  $\mathrm{id}_{D(Xy^A)} : D(Xy^A) \rightarrow D(Xy^A)$ .

*Proof.* Immediate from the associativity and unitality of composition of functors.  $\square$

In particular, we can take  $D = \mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{T}}$  to obtain “hierarchical bidirectional Markov systems”.


### 3.3 Nested Systems and Dependent Polynomials

The foregoing formalism suffices to describe systems’ shapes, and behaviours of those shapes that depend on their sensoria. But in our world, a system has a *position* as well as a shape! Indeed, one might want to consider systems nested within systems, such that the outer systems constitute the ‘universes’ of the inner systems; in this way, inner shapes may depend on outer shapes, and inner sensoria on outer sensoria.<sup>2</sup> We can model this situation polynomially.

Recall that an object in  $\mathbf{Poly}_{\mathcal{E}}$  corresponds to a bundle  $E \rightarrow B$ , equivalently a diagram  $1 \leftarrow E \xrightarrow{p} B \rightarrow 1$ , and note that the unit polynomial  $y$  corresponds to a bundle  $1 \rightarrow 1$ . We can then think of  $\mathbf{Poly}_{\mathcal{E}}$  as the category of “polynomials in one variable”, or “polynomials over  $y$ ”. This presents a natural generalization, to polynomials in many variables, corresponding to diagrams  $J \leftarrow E \rightarrow B \rightarrow I$ ; these diagrams form the objects of a category  $\mathbf{Poly}_{\mathcal{E}}(J, I)$ . When  $J$  is a (polynomial) bundle  $\beta$  over  $I$ , then we can take the subcategory of  $\mathbf{Poly}_{\mathcal{E}}(J, I)$  whose objects are commuting squares and whose morphisms are prisms as follows; the commutativity ensures that inner and outer sensoria are compatible.

<sup>2</sup>We might even consider the outer shapes explicitly as positions in some world-space, and the outer sensorium as determined by possible paths between positions, in agreement with the perspective of [12] on polynomials.

**Proposition 3.11.** There is an indexed category of **nested polynomials** which by abuse of notation we will call  $\mathbf{Poly}_{\mathcal{E}}(-) : \mathbf{Poly}_{\mathcal{E}} \rightarrow \mathbf{Cat}$ . Given  $\beta : J \rightarrow I$ , the category  $\mathbf{Poly}_{\mathcal{E}}(\beta)$  has commuting squares as on the left below as objects and prisms as on the right as morphisms. Its action on polynomial morphisms  $\beta \rightarrow \gamma$  is given by composition.



**Remark 3.12.** This construction can be repeatedly iterated, modelling systems within systems within systems. We leave the consideration of the structure of this iteration to future work, though we expect it to have an opetopic shape equivalent to that obtained by iterating the **Para** construction (cf. Remark 2.23).

**Observation 3.13.** Our polynomially indexed categories of dynamical behaviours and statistical games (Prop. 4.9) generalize to the case of nested polynomials, giving a doubly-indexed structure. For more information on the former case, we refer the reader to our preprint [17].

## 4 Theories of Approximate and Active Inference

We now start to bring together the structures of the previous sections, in order to breathe life into polynomials. We begin by sketching *approximate inference doctrines*, which characterize dynamical systems that optimize their performance at statistical games, without reference to morphology. In this paper, we do not concentrate on the detailed structure of these doctrines, leaving their exposition and comparison to future work, in which we will also be interested in morphisms between doctrines.

### 4.1 Approximate Inference Doctrines

An approximate inference doctrine will be a monoidal functor from a category of statistical games into an appropriate category of dynamical systems, taking games to systems that ‘play’ those games, typically by implementing an optimization process. In the free-energy literature (e.g., [3]), these systems have a hierarchical structure in which the realization of a game has access to the dynamics realizing the prior, mirroring the context-dependence of the games themselves, and it is for this reason that we use the category of hierarchical bidirectional systems defined above.

**Proposition 4.1.** Let  $\mathcal{P} : \mathcal{E} \rightarrow \mathcal{E}$  be a probability monad on  $\mathcal{E}$ , and let  $\mathcal{C}$  be a category that admits Bayesian inversion, equipped with a functor  $F : \mathcal{C} \rightarrow \mathcal{Kl}(\mathcal{P})$ . Since  $\mathbf{SGame}_{\mathcal{E}}$  is a Grothendieck fibration, it is equipped with a canonical projection functor  $\mathbf{SGame}_{\mathcal{E}} \rightarrow \mathcal{C}$ . Suppose that  $\mathcal{G}$  is a subcategory of  $\mathbf{SGame}_{\mathcal{E}}$ . Then there is a functor  $F_{\mathcal{G}} : \mathcal{G} \hookrightarrow \mathbf{SGame}_{\mathcal{E}} \rightarrow \mathcal{C} \xrightarrow{F} \mathcal{Kl}(\mathcal{P})$  whose image  $\text{im } F_{\mathcal{G}}$  in  $\mathcal{Kl}(\mathcal{P})$  consists of the forwards maps of  $\mathcal{G}$  realized as stochastic channels in  $\mathcal{Kl}(\mathcal{P})$ .

**Definition 4.2.** Suppose  $p : \mathbf{Poly}_{\mathcal{E}}$  and that  $\mathcal{C}$  is a subcategory of  $\mathcal{Kl}(\mathcal{P})$ . Denote by  $\mathbf{MrkProc}^{\mathbb{T}}_{\mathcal{P}}(p)|_{\mathcal{C}}$  the subcategory of  $\mathbf{MrkProc}^{\mathbb{T}}_{\mathcal{P}}(p)$  whose objects  $(\Theta, \theta^o, \theta^u)$  satisfy the condition that, for all  $t : \mathbb{T}$ , the following composite belongs to  $\mathcal{C}$ :

$$\sum_{t:\mathbb{T}} \sum_{x:\Theta} p[\theta^o(t, x)] \xrightarrow{\theta^u(t)} \mathcal{P} \Theta \xrightarrow{\mathcal{P} \theta^o(t)} \mathcal{P} p(1).$$

Note that  $\mathbf{MrkProc}^{\mathbb{T}}_{\mathcal{P}}(-)|_{\mathcal{C}}$  does not in general define an indexed category.

**Definition 4.3.** Suppose  $\mathcal{C}$  is a subcategory of  $\mathcal{Kl}(\mathcal{P})$ . Denote by  $\mathbf{HiBi}(\mathbf{MrkProc}^{\mathbb{T}}_{\mathcal{P}})|_{\mathcal{C}}$  the restriction of  $\mathbf{HiBi}(\mathbf{MrkProc}^{\mathbb{T}}_{\mathcal{P}})$  defined as follows. The objects of  $\mathbf{HiBi}(\mathbf{MrkProc}^{\mathbb{T}}_{\mathcal{P}})|_{\mathcal{C}}$  are the objects of  $\mathbf{HiBi}(\mathbf{MrkProc}^{\mathbb{T}}_{\mathcal{P}})$ , but morphisms  $(X, A) \rightarrow (Y, B)$  are now functors  $\mathbf{MrkProc}^{\mathbb{T}}_{\mathcal{P}}|_{\mathcal{C}}(Xy^A) \rightarrow \mathbf{MrkProc}^{\mathbb{T}}_{\mathcal{P}}|_{\mathcal{C}}(Yy^B)$ . Composition and identities are defined as for  $\mathbf{HiBi}(\mathbf{MrkProc}^{\mathbb{T}}_{\mathcal{P}})$ .

We are now in a position to define approximate inference doctrines.

**Definition 4.4.** Let  $\mathcal{P} : \mathcal{E} \rightarrow \mathcal{E}$  be a probability monad on  $\mathcal{E}$  and let  $\mathcal{G}$  be a subcategory of  $\mathbf{SGame}_{\mathcal{Kl}(\mathcal{P})}$ . An **approximate inference doctrine** in  $\mathcal{G}$  with time  $\mathbb{T}$  is a functor from  $\mathcal{G}$  to  $\mathbf{HiBi}(\mathbf{MrkProc}^{\mathbb{T}}_{\mathcal{P}})|_{\mathcal{C}}$  where  $\mathcal{C}$  is defined as in Proposition 4.1.

Many informal approximate inference schemes—including Markov chain Monte Carlo, variational Bayes, expectation-maximization, particle filtering—give rise to approximate inference doctrines; functoriality typically follows from Theorem 2.12. Here we note one explicitly, for later reference; a detailed presentation will appear in an upcoming paper.

**Definition 4.5.** We say that  $f : X \rightarrowtail Y$  in  $\mathcal{Kl}(\mathcal{P})$  is **Gaussian** if, for any  $x : X$ , the state  $f(x) : \mathcal{P}Y$  is Gaussian. Note that Gaussian morphisms are not themselves closed under composition, though we can consider the subcategory of morphisms generated by Gaussian morphisms.

**Lemma 4.6** (Laplace approximation). Suppose:

- (a) **Gauss** is the subcategory of channels generated by Gaussian morphisms between finite-dimensional Euclidean spaces in  $\mathcal{Kl}(\mathcal{P})$ ;
- (b)  $(\gamma, \rho, \phi) : X \rightarrow Y$  is a simple  $D_{KL}$ -free energy game with Gaussian channels;
- (c) for all priors  $\pi : \mathcal{P}X$ , the statistical parameters of  $\rho_{\pi} : Y \rightarrow \mathcal{P}X$  are denoted by  $(\mu_{\rho_{\pi}}, \Sigma_{\rho_{\pi}}) : Y \rightarrow \mathbb{R}^{|X|} \times \mathbb{R}^{|X| \times |X|}$ , where  $|X|$  is the dimension of  $X$ ; and
- (d) for all  $y : Y$ , the eigenvalues of  $\Sigma_{\rho_{\pi}}(y)$  are small.

Then the loss function  $\phi : \mathbf{BayesLens}_{\mathbf{Gauss}}((X, X), (Y, Y)) \rightarrow \mathbb{R}$  can be approximated by

$$\phi(\pi, k) = \mathbb{E}_{y \sim k \bullet \gamma \bullet \pi} [\mathcal{F}(y)] \approx \mathbb{E}_{y \sim k \bullet \gamma \bullet \pi} [\mathcal{F}^L(y)]$$

where

$$\begin{aligned} \mathcal{F}^L(y) &= E_{(\pi, \gamma)}(\mu_{\rho_{\pi}}(y), y) - S_X[\rho_{\pi}(y)] \\ &= -\log p_{\gamma}(y | \mu_{\rho_{\pi}}(y)) - \log p_{\pi}(\mu_{\rho_{\pi}}(y)) - S_X[\rho_{\pi}(y)] \end{aligned} \quad (4)$$

and where  $S_X[\rho_{\pi}(y)] = \mathbb{E}_{x \sim \rho_{\pi}(y)}[-\log p_{\rho_{\pi}}(x|y)]$  is the Shannon entropy of  $\rho_{\pi}(y)$ , and  $p_{\gamma} : Y \times X \rightarrow [0, 1]$ ,  $p_{\pi} : X \rightarrow [0, 1]$ , and  $p_{\rho_{\pi}} : X \times Y \rightarrow [0, 1]$  are density functions for  $\gamma$ ,  $\pi$ , and  $\rho_{\pi}$  respectively. The approximation is valid when  $\Sigma_{\rho_{\pi}}$  satisfies

$$\Sigma_{\rho_{\pi}}(y) = (\partial_x^2 E_{(\pi, \gamma)})(\mu_{\rho_{\pi}}(y), y)^{-1}. \quad (5)$$

**Theorem 4.7.** Suppose:

- (a) **Gauss** is the subcategory of channels generated by Gaussian morphisms between finite-dimensional Euclidean spaces in  $\mathcal{KL}(\mathcal{P})$ ;
- (b)  $\mathbf{SimpSGame}_{\mathbf{Gauss}} \hookrightarrow \mathbf{SGame}_{\mathcal{KL}(\mathcal{P})}$  is the subcategory of simple statistical games over  $\mathcal{KL}(\mathcal{P})$  with channels in **Gauss**; and
- (c)  $\mathcal{G}$  is the subcategory of  $D_{KL}$ -Bayesian inference games in  $\mathbf{SimpSGame}_{\mathbf{Gauss}}$ .

Then the discrete-time free-energy principle under the Laplace approximation induces an approximate inference doctrine in  $\mathcal{G}$  with time  $\mathbb{N}$ ,  $\text{Laplace} : \mathcal{G} \rightarrow \mathbf{HiBi}(\mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{T}}) |_{\text{imF}_{\mathcal{G}}}$ .

## 4.2 Statistical Games over Polynomials

Approximate inference doctrines of the foregoing type do not supply a satisfactory model of active systems. One piece of structure is still missing, with which we can describe action and interaction faithfully: an indexed category of statistical games over polynomials. In order to construct this, we first define categories of “games on interfaces”: this is simpler than slicing the category of statistical games, as we do not require games between games.

**Definition 4.8.** Let  $\mathcal{P} : \mathcal{E} \rightarrow \mathcal{E}$  be a probability monad on  $\mathcal{E}$ . Let  $X : \mathcal{E}$  be an object in  $\mathcal{E}$ . Define a category of **simple statistical games on the interface**  $X$ , denoted by  $\mathbf{IntGame}_{\mathcal{P}}(X)$ , as follows. Its objects are simple statistical games with codomain  $X$ ; that is, points of  $\sum_{A:\mathcal{E}} \mathbf{SimpSGame}_{\mathcal{KL}(\mathcal{P})}(A, X)$ . Let  $(\gamma, \rho, \phi) : A \rightarrow X$  and  $(\delta, \sigma, \chi) : B \rightarrow X$  be two such simple statistical games. Then a morphism  $(\gamma, \rho, \phi) \rightarrow (\delta, \sigma, \chi)$  is a deterministic function  $f : A \rightarrow B$ —that is, a point of  $\mathcal{E}(A, B)$ —such that  $\gamma = \delta \circ f$ . Unitality and associativity follow immediately from those properties in  $\mathcal{E} / \mathcal{P} X$ .

We then use this to construct games over polynomials. The intuition here is that ‘inside’ a system with a polynomial phenotype is a statistical model of the system’s sensorium. This involves an object representing the space of possible causes of observations, and a simple statistical game from this object onto the sensorium; by its nature, this model induces predictions about the system’s configurations, as well as about the immanent signals. By sampling from these predicted configurations, the system can act; by observing its actual configuration and the corresponding immanent signals, it can update its internal beliefs, and any parameters of the model. Later, we will equip this process with (random) dynamics, thereby giving the systems life.

**Proposition 4.9.** Let  $\mathcal{P} : \mathcal{E} \rightarrow \mathcal{E}$  be a probability monad on a locally Cartesian closed category  $\mathcal{E}$ . There is a **polynomially indexed category of statistical games**  $\mathbf{PSGame}_{\mathcal{P}} : \mathbf{Poly}_{\mathcal{E}} \rightarrow \mathbf{Cat}$ , defined on objects  $p$  as  $\mathbf{IntGame}_{\mathcal{P}}(\sum_{i:p(1)} p[i])$ .

**Example 4.10.** To understand the action of  $\mathbf{PSGame}_{\mathcal{P}}(\varphi)$  on statistical games, it may help to consider the example of a corporation. Such a system is composed of a number of active systems, which instantiate statistical games and interact according to some pattern formalized by the polynomial map  $\varphi$ . Given such a collection of games,  $\mathbf{PSGame}_{\mathcal{P}}(\varphi)$  tells us how to construct a game for the corporation as a whole: in particular, we obtain a stochastic channel generating predictions for the (exposed) sensorium of the corporation, and an inversion updating the constituent systems’ beliefs accordingly.

## 4.3 Active Inference

We are now ready to define active inference doctrines; given all the foregoing structure, this proves relatively simple.

**Definition 4.11.** Let  $\mathcal{P} : \mathcal{E} \rightarrow \mathcal{E}$  be a probability monad, and let  $\mathbb{T}$  be a time monoid. An **active inference doctrine** is a monoidal indexed functor from  $\mathbf{PSGame}_{\mathcal{P}}$  to  $\mathbf{MrkProc}_{\mathcal{P}}^{\mathbb{T}}$ .

**Proposition 4.12.** The Laplace doctrine lifts from approximate to active inference. The functors on each fibre are as before on games (here, objects), and on morphisms between games they are given merely by lifting the corresponding maps to maps between state spaces. One then checks that morphisms of polynomials correspond to natural transformations between these functors.

**Remark 4.13.** It would be desirable to incorporate the compositional structure of the games themselves, rather than treat them opaquely as objects. This suggests a double-categorical structure the investigation of which we leave to future work. Similarly, we do not here elaborate the extension of these indexed categories to the dependent-polynomial case.

## 5 Polynomial Life and Embodied Cognition

Finally, we sketch how a number of classic biological processes can be modelled as processes of active inference over polynomials. The key insight is that, by fixing the prior of an ‘active’ free-energy game to encode high-precision (low-variance) beliefs about the external state, we can induce the system to prefer acting (*i.e.*, reifying those beliefs) over perceiving (*i.e.*, updating the beliefs to match perceptions). In doing so, one can induce *volition* or *goal-directedness* in the system. A key feature of these examples is that they demonstrate ‘embodied’ cognition, in which a system’s form and interactions become part of its cognitive apparatus.

**Remark 5.1.** Of course, one must be careful not to choose a prior with excessively high precision (such as a Dirac delta distribution), as this would cause the system to forego any belief-updating, thereby rendering its actions independent of the ‘actual’ external state.

**Example 5.2.** Suppose that the system’s sensorium includes a key parameter such as ambient temperature or blood pH. Suppose that by adjusting its configuration, the system can move around in order to sample this parameter. And suppose that the prior encodes a high-precision distribution centred on the acceptable range of this parameter. Then it is straightforward to show that the system, by minimizing the free energy, will attempt to configure itself so as to remain within the acceptable parameter range. We can consider this as a simple model of **homeostasis**.

**Example 5.3.** We can extend the previous example to a system with multiple (polynomial) components, each equipped with a “homeostasis game”, in order to model **morphogenesis**. Suppose the environmental parameter in the sensorium is the local concentration of some signalling molecule, and suppose the polynomial morphism forming the composite system encodes the pattern of signalling molecule concentrations in the neighbourhood of each system, as a result of their mutual configurations. Suppose then that the target state encoded in the prior of each system corresponds to the system being positioned in a particular way relative to the systems around it, as represented by the signal concentrations. Free-energy minimization then induces the systems to arrange themselves in order to obtain the target pattern.

**Remark 5.4.** The foregoing examples begin to point towards a compositional theory of *autopoiesis*: here, one might expect the target state to encode the proposition “maintain my morphology”, which appears self-referential. The most elegant way of encoding this proposition in the prior is not immediately clear, although a number of possibilities present themselves (such as avoiding some undesirable configuration representing dissolution). We expect a satisfactory answer to this to be related to “Bayesian mechanics” (see §6).

**Remark 5.5.** It has been shown informally that, given a finite time horizon Markov decision problem, active inference can recover the Bellman-optimal policy traditionally obtained by backward induction [7]. Ongoing work by the present author is directed at formalizing the structure of this relationship. In particular, the result rests on encoding directly into the prior the expectation of the loss function given a policy and a goal, which strikes us as a large amount of information to push into an unstructured distribution over numbers.

**Example 5.6.** The examples need not be restricted to simple biological cases. For instance, we can model spatial navigation quite generally: we can use parameterized statistical games to encode uncertainty about the structure of the ‘external space’ (for instance: which points or neighbourhoods are connected to which, and by which paths). By setting a high-precision prior at some location, the system will attempt to reach that location, learning the spatial structure along the way; reducing the precision of the prior causes the system to prefer “mere exploration”. One can attach sense-data to each location using the natural polynomial bundle structure. Moreover, the ‘external space’ need not be a simple topological space: it may be something more structured. For instance, categories and sites can themselves be modelled polynomially. One can think of “taking an action” as precisely analogous to “following a morphism”: thus, in a topos-theoretic setting, one can consider the structure of the ‘external space’ to be a type-theoretic context, and positions in the world to be objects in the corresponding topos. One could then encode in the prior a target proposition, and free-energy minimization would cause the system then to explore the ‘space’ (learning its structure), and seek a path to the target. But such a path is precisely a proof! There is increasing evidence that the neural mechanisms underlying spatial and abstract navigation are the same [1], and a line of thinking as sketched here may supply a mathematical justification.

## 6 Future Directions

Besides expanding the examples above in detail, there are many future directions to pursue. Our last example points towards a ‘well-typed’ theory of cognition, finding type-theoretic analogues of cognitive processes (such as action, planning, or navigation). By formalizing the connection between polynomial statistical games and Markov decision processes, we hope finally to relate our ‘statistical’ account of cybernetics with the account emerging from research in compositional game theory. In particular, we believe that the hierarchical/nested structure of our polynomial systems is structurally similar to that of (parameterized) players in open games. Along similar lines, we expect a connection between statistical games and ‘learners’ [13] implementing backprop.

In a more physical direction, there is a controversy in the informal literature about whether one should expect *any* system with a boundary (and hence any system on a polynomial) to admit a canonical statistical-game description; the typical suggestion is that such a description should obtain at non-equilibrium steady state, through a manipulation of the corresponding Fokker-Planck equation; this is the notion of “Bayesian mechanics” [9]. Our results suggest that such a canonical description should form a left adjoint to some active inference doctrine; this is a matter of ongoing research by the author.

Working topos-theoretically points further in a metaphysical direction: a Bayesian perspective lends itself to subjectivism, but considering the “internal universe” of a navigating system to be a topos in some context also points to a subjective realism. It seems likely then that composite systems need not in general agree about their observations. We should therefore expect to find evidence of contextuality and disagreement in multi-agent systems, and to investigate this using cohomological tools (*e.g.*, [5]).

## References

- [1] Timothy EJ Behrens, Timothy H Muller, James CR Whittington, Shirley Mark, Alon B Baram, Kimberly L Stachenfeld & Zeb Kurth-Nelson (2018): *What is a cognitive map? Organizing knowledge for flexible behavior*. *Neuron* 100(2), pp. 490–509, doi:10.1016/j.neuron.2018.10.002.
- [2] Joe Bolt, Jules Hedges & Philipp Zahn (2019): *Bayesian open games*. *arXiv preprint*. Available at <http://arxiv.org/abs/1910.03656>.
- [3] Christopher L Buckley, Chang Sub Kim, Simon McGregor & Anil K Seth (2017): *The free energy principle for action and perception: A mathematical review*. *Journal of Mathematical Psychology* 81, pp. 55–79, doi:10.1016/j.jmp.2017.09.004.
- [4] Matteo Capucci, Bruno Gavranović, Jules Hedges & Eigil Fjeldgren Rischel (2021): *Towards foundations of categorical cybernetics*. this volume of EPTCS, Open Publishing Association. Available at <https://arxiv.org/abs/2105.06332>.
- [5] Giovanni Carù (2017): *On the Cohomology of Contextuality*. In: *Proceedings of the 13th International Conference on Quantum Physics and Logic*, 236, pp. 21–39, doi:10.4204/EPTCS.236.2.
- [6] Kenta Cho & Bart Jacobs (2017): *Disintegration and Bayesian Inversion via String Diagrams*. *Mathematical Structures in Computer Science* 29, pp. 938–971, doi:10.1017/S0960129518000488.
- [7] Lancelot Da Costa, Noor Sajid, Thomas Parr, Karl Friston & Ryan Smith (2020): *The Relationship between Dynamic Programming and Active Inference: The Discrete, Finite-horizon Case*. *arXiv preprint*. Available at <https://arxiv.org/abs/2009.08111>.
- [8] K. Friston, J. Mattout, N. Trujillo-Barreto, J. Ashburner & W. Penny (2007): *Variational free energy and the Laplace approximation*. *Neuroimage* 34(1), pp. 220–234, doi:10.1016/j.neuroimage.2006.08.035.
- [9] Karl Friston (2019): *A free energy principle for a particular physics*. *arXiv preprint*. Available at <http://arxiv.org/abs/1906.10184>.
- [10] Chris Heunen, Ohad Kammar, Sam Staton & Hongseok Yang (2017): *A Convenient Category for Higher-Order Probability Theory*. In: *32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, IEEE, doi:10.1109/lics.2017.8005137.
- [11] David I. Spivak (2019): *Generalized Lens Categories via functors  $\mathcal{C}^{op} \rightarrow \mathbf{Cat}$* . *arXiv preprint*. Available at <http://arxiv.org/abs/1908.02202>.
- [12] David I. Spivak (2020): *Poly: An abundant categorical setting for mode-dependent dynamics*. *arXiv preprint (presented at the 3rd Annual International Applied Category Theory Conference)*. Available at <https://arxiv.org/abs/2005.01894>.
- [13] David I. Spivak (2021): *Learners’ languages*. this volume of EPTCS, Open Publishing Association. Available at <https://arxiv.org/abs/2103.01189>.
- [14] Toby St. Clere Smithe (2020): *Bayesian Updates Compose Optically*. *arXiv preprint*. Available at <https://arxiv.org/abs/2006.01631>.
- [15] Toby St. Clere Smithe (2020): *Cyber Kittens, or Some First Steps Towards Categorical Cybernetics*. In: *Proceedings 3rd Annual International Applied Category Theory Conference 2020 (ACT 2020)*, 333, pp. 108–124, doi:10.4204/EPTCS.333.8.

- [16] Toby St. Clere Smithe (2021): *Compositional Active Inference I: Bayesian Lenses. Statistical Games*. *arXiv preprint*. Available at <https://arxiv.org/abs/2109.04461>.
- [17] Toby St. Clere Smithe (2021): *Some Notions of (Open) Dynamical System on Polynomial Interfaces*. *arXiv preprint*. Available at <https://arxiv.org/abs/2108.11137>.
- [18] Toby St. Clere Smithe (2022): *Open dynamical systems as coalgebras for polynomial functors, with application to predictive processing*. In: *Proceedings 5th Annual International Applied Category Theory Conference 2021 (forthcoming)*. Available at <https://arxiv.org/abs/2206.03868>.