BAYESIAN GAMES, DIEGETICALLY

1. Reverse differential structure

Let **Markov** be the category of Markov kernels on measurable spaces, i.e. the Kleisli category of the Giry monad $D : \mathbf{Msbl} \to \mathbf{Msbl}$. Let **Cvx** be the category of D-**Alg**, aka convex spaces.

We recall there is a 'universal parametric map' (monadic strength):

$$\sigma : D(X \to Y) \longrightarrow (X \to DY) \tag{1.1}$$

given by

$$\sigma(\mathrm{d}f) = \lambda x \,.\, \mathrm{ev}_*(\mathrm{d}f \otimes \delta_x) = \lambda x \,.\, \lambda A \,.\, \int_{X \to Y} \mathbf{1}_A(f(x))\, \mathrm{d}f. \tag{1.2}$$

On the other hand, given a stochastic function $f : X \to DY$ we can consider it as an X-indexed family of probability spaces $(Y, f(x))_{x \in X}$. Taking their cartesian product in the category of probability spaces we obtain a probability space $(Y^X, \prod_{x \in X} f(x))$. This defines a map:

$$K : (X \to DY) \longrightarrow D(X \to Y) \tag{1.3}$$

related to σ in the following fashion:

Lemma 1.1. *Let X and Y be measurable spaces, and suppose Y is such that every singleton is measurable. Then $K : (X \to DY) \to D(X \to Y)$ is a section of $\sigma : D(X \to Y) \to (X \to DY)$.*

Proof.

$$\begin{aligned} \sigma(K(\mathrm{d}y)) &= \sigma(\textstyle\bigotimes_{x \in X} \mathrm{d}y(x)) \\ &= \lambda \bar{x} \,.\, \lambda A \,.\, (\textstyle\bigotimes_{x \in X} \mathrm{d}y(x) \otimes \delta_{\bar{x}})(\mathrm{ev}^{-1} A) \\ &= \lambda \bar{x} \,.\, \lambda A \,.\, \textstyle\prod_{x \in X} \mathrm{d}y(x)(\underbrace{\pi_{\bar{x}}\{f \mid f(\bar{x}) \in A\}}_{Y \text{ if } x \neq \bar{x},\ A \text{ otherwise}}) \\ &= \lambda \bar{x} \,.\, \lambda A \,.\, \mathrm{d}y(\bar{x})(A) \\ &= \mathrm{d}y. \end{aligned} \tag{1.4}$$

□

We define the following functor, which we could call the *dependent type theory of stochastic types*:

Received by the editors October 28, 2022.

$$\begin{array}{ccc} \mathbf{Markov}^{\mathsf{op}} & \xrightarrow{F_D} & \mathbb{C}\mathbf{at} \\ X & & \mathbf{Cvx}^{DX} \\ {\scriptstyle f}\downarrow & \longmapsto & \uparrow{\scriptstyle f^*} \\ Y & & \mathbf{Cvx}^{DY} \end{array} \tag{1.5}$$

Remark 1.1.1. Given a family $p : X \to \mathbf{Markov}$, one can freely extend it along $\eta_X : X \to DX$. In fact one can first imagine to convert p to a display map $\sum_{x\in X} p(x) \stackrel{\pi_X}{\to} X$. Then notice $DX = \sum_{x\in X} 1$ in $\mathbf{Cvx} = D$-$\mathbf{Alg}$, so that π_X can be realized as an actual map there:

$$\sum_{x\in X} Dp(x) \xrightarrow{\pi_X} DX \tag{1.6}$$

which, converted back to a family $Dp : X \to \mathbf{Markov}$, gives

$$Dp(\mathrm{d}x) = \pi_X^{-1}(\mathrm{d}x) = \{\textstyle\sum_{x\in X} \mathrm{d}x(x) y_x \mid \forall x \in X, y_x \in p(x)\}. \tag{1.7}$$

This is better considered in $\mathbf{Cvx}$ though, since one would expect maps out of any such 'interpolated' fiber to be determined by maps out of the original fibers.

Definition 1.2. A D-lens $\begin{pmatrix} X' \\ X \end{pmatrix} \leftrightarrows \begin{pmatrix} Y' \\ Y \end{pmatrix}$ is a pair of morphisms:

$$\begin{gathered} f : X \to Y, \\ \mathrm{d}x \in DX \vdash f^\sharp_{\mathrm{d}x} : Y'(f(\mathrm{d}x)) \to X'(\mathrm{d}x) \end{gathered} \tag{1.8}$$

with $f : \mathbf{Markov}$ and $f^\sharp : \mathbf{Cvx}^{DX}$ (i.e. it's componentwise convex linear).

Remark 1.2.1. There's also another approach to Markov/Bayes lenses, through the point of view of optics:

$$\mathbf{Optic}_{\mathbf{Kl}(D),\mathbf{Alg}(D)}\left(\begin{pmatrix} X' \\ X \end{pmatrix}, \begin{pmatrix} Y' \\ Y \end{pmatrix}\right) = \int^{M:\mathbf{Markov}} \mathbf{Markov}(X, M \times Y) \times \mathbf{Cvx}(M \times Y', X') \tag{1.9}$$

At least in the constant case (when $X'(\mathrm{d}x) = X'(\mathrm{d}x')$ for all $\mathrm{d}x, \mathrm{d}x' \in DX$) D-lenses are also Markov optics [**hedges2022value**] of a particular kind. Markov optics are more general though, since they allow non-trivial couplings between the output of f and that of $f^\sharp$.

Moreover, one can promote the optics to the dependent case by deploying the definitions in [**braithwaite2021fibre**; **capucci2022seeing**]. In $\mathbf{DOptic}_{\mathbf{Markov}}$ one has objects dependent pairs $\begin{pmatrix} X' \\ X \end{pmatrix}$ of a measurable space X and a family of convex sets $X' : DX \to \mathbf{Cvx}$ as before. A morphism is given by a span $X \stackrel{p}{\leftarrow} Z \stackrel{q}{\to} Y$, and a family of convex morphisms:

$$\mathrm{d}x \in DX \vdash f^\sharp_{\mathrm{d}x} : \sum_{p(z)=\mathrm{d}x} Y'(q(z)) \to X'(\mathrm{d}x) \tag{1.10}$$

A span $X \overset{p}{\leftarrow} Z \overset{q}{\to} Y$ can be seen as a dependent type p plus an actual forward part q: given an $x : X$, a type $Z(x)$ (the residual) and family of maps $q_x : Z(x) \to Y$. Then the optic can be written as:

$$\begin{gathered} f : (x : X) \times Z(x) \to Y, \\ \mathrm{d}x \in DX \vdash f^\sharp_{\mathrm{d}x} : (z : DZ(\mathrm{d}x)) \times Y'(q_{\mathrm{d}x}(z)) \to X'(\mathrm{d}x) \end{gathered} \tag{1.11}$$

where DZ is the convex extension of $p : Z \to X$ as described in Remark 1.1.1. It is evident that one recovers D-lenses as the special case when $p = 1_X$, hence residual space is trivial. In that case, $q_x : 1 \to Y$ picks a single distribution over Y and the extra z argument in $f^\sharp$ is spurious, recovering Definition 1.2.

For simplicity, we stick to D-lenses here.

Definition 1.3. There is a monoidal structure on D-lenses, given by

$$\begin{pmatrix} X' \\ X \end{pmatrix} \otimes \begin{pmatrix} Y' \\ Y \end{pmatrix} = \begin{pmatrix} X' \otimes Y' \\ X \times Y \end{pmatrix} \tag{1.12}$$

where, in turn, $X' \otimes Y'$ is defined by the lax monoidal structure of F_D so given, for any $\mathrm{d}xy \in D(X \times Y)$:

$$(X' \otimes Y')(\mathrm{d}xy) = X'(\pi_X(\mathrm{d}xy)) \times Y'(\pi_Y(\mathrm{d}xy)). \tag{1.13}$$

We introduce now the following functor, the *support functor*:

$$\begin{array}{ccc} \mathbf{Markov} & \xrightarrow{\quad \mathrm{supp} \quad} & \mathbf{DChart}(D) \\ X & & \begin{pmatrix} D\,\mathrm{supp}_X \\ X \end{pmatrix} \\ {\scriptstyle f}\big\downarrow & \longmapsto & {\scriptstyle f}\big\downarrow\big\downarrow{\scriptstyle \mathrm{supp}\, f} \\ Y & & \begin{pmatrix} D\,\mathrm{supp}_Y \\ Y \end{pmatrix} \end{array} \tag{1.14}$$

For a given distribution $\mathrm{d}x \in DX$, the space $\mathrm{supp}_X\, \mathrm{d}x$ is supposed to be the subspace of X where $\mathrm{d}x$ has strictly positive density.

Remark 1.3.1. We handwave its existence here, which is debatable for arbitrary probability measures over arbitrary measurable spaces, but can be guaranteed if we restrict either the measures or the spaces or both. For instance we can restrict spaces to be topological spaces equipped with the Borel measurable structure, or the probability measures to be atomic. More attention to this problem is devolved in [**braithwaite2022**], where the type families supp_X are first defined.

In any case, $\mathrm{supp}\, f$ has type

$$\mathrm{d}x \in DX \vdash \mathrm{supp}\, f_{\mathrm{d}x} : D\,\mathrm{supp}_X\, \mathrm{d}x \to D\,\mathrm{supp}_Y\, f(\mathrm{d}x) \tag{1.15}$$

and its defined by convex extension of the co/restriction of f.

Lemma 1.4. $\operatorname{supp} f$ *is well-defined.*

Proof. Notice first the following property of supports:

$$\forall p \in (0,1), \quad \operatorname{supp}_X(p\pi + (1-p)\sigma) = \operatorname{supp}_X \pi \cup \operatorname{supp}_X \sigma \tag{1.16}$$

From this we deduce that

$$\operatorname{supp}_Y f(\mathrm{d}x) = \bigcup \{\operatorname{supp}_Y f(x) \mid x \in \operatorname{supp}_X \mathrm{d}x\}. \tag{1.17}$$

Now $\operatorname{supp} f$ is a map of free algebras hence can be defined as a Kleisli map $\operatorname{supp}_X \mathrm{d}x \to D \operatorname{supp}_Y f(\mathrm{d}x)$. The argument above proves that given a point in $\operatorname{supp}_X \mathrm{d}x$, the distribution we get on Y is actually supported on $\operatorname{supp}_Y f(\mathrm{d}x)$, meaning it's a well-defined element of $D \operatorname{supp}_Y f(\mathrm{d}x)$. □

Moreover, supp admits a lax monoidal structure:

$$\begin{pmatrix} \operatorname{supp}_\Theta \otimes \operatorname{supp}_\Omega \\ \Theta \times \Omega \end{pmatrix} \overset{\otimes}{\Longrightarrow} \begin{pmatrix} \operatorname{supp}_{\Theta\times\Omega} \\ \Theta \times \Omega \end{pmatrix} \tag{1.18}$$

where $\otimes$ is the product of measures. There also is an oplax monoidal structure, given by taking marginals, of which $\otimes$ is a section but not an inverse.

Apparently very similar to this functor, but actually very different, is the *Bayesian inversion functor*, as defined in [**braithwaite2022**]:

$$\begin{array}{ccc} \mathbf{Markov} & \xrightarrow{(-)^\dagger} & \mathbf{DLens}(D) \\ X & & \begin{pmatrix} D\operatorname{supp}_X \\ X \end{pmatrix} \\ f \big\downarrow & \longmapsto & f \big\downarrow \big\uparrow f^\dagger \\ Y & & \begin{pmatrix} D\operatorname{supp}_Y \\ Y \end{pmatrix} \end{array} \tag{1.19}$$

The family of convex maps $f^\dagger$ have type

$$\mathrm{d}x \in DX \vdash f^\dagger_{\mathrm{d}x} : D\operatorname{supp}_Y(f(\mathrm{d}x)) \to D\operatorname{supp}_X \mathrm{d}x \tag{1.20}$$

and are defined by Bayesian inversion: given $\mathrm{d}y \in \operatorname{supp}_Y(f(\mathrm{d}x))$, $f^\dagger_{\mathrm{d}x}(\mathrm{d}y)$ is the distribution on X (which happens to be also supported on $\operatorname{supp}_X(\mathrm{d}x)$) obtained by conditioning $\mathrm{d}x$ with the observation $\mathrm{d}y$ (normalization omitted):

$$f^\dagger(\mathrm{d}y)(x) = \mathrm{d}x(x)\mathrm{d}y(f(x)). \tag{1.21}$$

We are then ready to define our functor of valuations, or 'reverse-mode differentiation functor':

$$\mathbf{Markov} \xrightarrow{\mathrm{supp}^*} \mathbf{DLens}(D)$$

$$\begin{array}{ccc} X & & \begin{pmatrix} \mathrm{supp}^*_X \\ X \end{pmatrix} \\ {\scriptstyle f}\downarrow & \longmapsto & {\scriptstyle f}\downarrow\uparrow{\scriptstyle \mathrm{supp}^* f} \\ Y & & \begin{pmatrix} \mathrm{supp}^*_Y \\ Y \end{pmatrix} \end{array} \tag{1.22}$$

where for a given $\mathrm{d}x \in DX$, $\mathrm{supp}^*_X(\mathrm{d}x) = \mathrm{supp}_X \mathrm{d}x \to D\mathbb{R}$ is the set of Kleisli maps from the support of $\mathrm{d}x$ to $\mathbb{R}$, or equivalently, the set of convex linear maps $D \,\mathrm{supp}_X \mathrm{d}x \to D\mathbb{R}$.

Then $\mathrm{supp}^* f$ is defined as $D\mathbb{R}^{\mathrm{supp}\, f}$, precomposing a given $u \in \mathrm{supp}^*_X$ with $\mathrm{supp}\, f$.

We can equip supp^* with a lax monoidal structure:

$$\begin{pmatrix} \mathrm{supp}^*_{\Theta\times\Omega} \\ \Theta\times\Omega \end{pmatrix} \overset{\mathbf{n}_{\Theta,\Omega}}{\underset{}{\Longrightarrow}} \begin{pmatrix} \mathrm{supp}^*_{\Theta} \times \mathrm{supp}^*_{\Omega} \\ \Theta\times\Omega \end{pmatrix} \tag{1.23}$$

with $\mathbf{n}_{\Theta,\Omega}$ defined as

$$\begin{aligned} \mathrm{d}\theta\omega \in D(\Theta\times\Omega) \vdash \mathbf{n}_{\Theta,\Omega} : D\mathbb{R}^{\mathrm{supp}_{\Theta\times\Omega}(\mathrm{d}\theta\omega)} &\longrightarrow D\mathbb{R}^{\mathrm{supp}_\Theta \pi_\Theta(\mathrm{d}\theta\omega)} \times D\mathbb{R}^{\mathrm{supp}_\Omega \pi_\Omega(\mathrm{d}\theta\omega)} \\ u &\longmapsto \langle \lambda\theta \,.\, u(\pi^\dagger_\Theta(\mathrm{d}\theta\omega, \theta)),\ \lambda\omega \,.\, u(\pi^\dagger_\Omega(\mathrm{d}\theta\omega, \omega))\rangle \end{aligned} \tag{1.24}$$

We remark this is well-defined: $\pi^\dagger_\Theta(\mathrm{d}\theta\omega, \theta)$ (and similarly $\pi^\dagger_\Omega(\mathrm{d}\theta\omega, \omega)$) is a well-defined expression since $\theta \in \mathrm{supp}_\Theta \pi_\Theta(\mathrm{d}\theta\omega)$ by definition. Moreover, all the operations involved are convex linear (precomposition and Bayesian inversion).

Intuitively, this takes a stochastic function u defined on $\Theta\times\Omega$ and turns it into a pair of stochastic functions defined, respectively, on Θ and Ω. These are defined by using the given argument, say $\theta \in \Theta$, to condition the joint probability $\mathrm{d}\theta\omega$ into a new probability on $\Theta \times \Omega$. All the shenanigans with supports is to make sure this conditioning is well-defined.

2. The structure of a Bayesian game

In Bayesian games, players employ *mixed strategies*, i.e. their canonical set of strategies in a game with states X and moves Y is $D(X \to Y)$. Their play map $X \to DY$ is extracted from a given mixed strategy by the strength σ defined previously (1.1). Observe that $X \to DY$ is instead the type of *behavioural strategies*, so σ is the operation that turns a mixed strategy into a behavioural strategy, and losing information.[1].

In general we assume players have a parametrization $p : \Omega \to (X \to DY)$ of the space of plays, which corresponds to a parametric map $\Omega \times X \to DY$. Observe how every such

[1]This is true unless $X = 1$, a commonplace occurence in game theory, since most of the time a game has a single initial state. In a game of perfect information therefore, behavioural and mixed strategies do not differ. When the game has imperfect information instead, not all DY is accessible, and which parts are inaccesible depends on whether we consider DY to be a set of behavioural or mixed strategies!

p factors through σ:

$$\begin{array}{ccc} \Omega & & \\ {\scriptstyle \exists!\, p'} \downarrow & \searrow^{p} & \\ D(X \to Y) & \xrightarrow[\sigma]{} & (X \to DY) \end{array} \tag{2.1}$$

In fact, by Lemma 1.1 we know there is a section K of σ, so that one can obtain $p' = p \mathbin{⨾} K$. Notice, however, that p' is not unique (since a behavioural strategy can descend from many different mixed strategies), otherwise **Markov** would be cartesian closed.

Players will thus have a space Ω of strategies and a map p into mixed strategies $D(X \to Y)$. On the backward part, they receive a payoff function, which is something of type $\Omega \to \mathbb{R}$. Of course in Bayesian games the payoff function of the game depends on an additional parameter Θ, the type. However this is not of concern, as we can treat Θ like any other strategy space: indeed it can be considered Nature's strategy space. This is one of the first advantages that the diegetic and open games formalism provide.

Thanks to this, a diegetic Bayesian open game has a very similar structure to that of a deterministic diegetic open game, as described in [**capucci2022diegetic**]. The major difference lies in the way strategies are broadcasted to various players, which leads to the fact that once an arena has been composed and passed through **Para**(supp*), then instead of applying the Nashator a single time, we separately apply it for each player.

This gives the chance to specify the amount of information on everyone's strategy each player has, including information about the type Θ. The latter is usually where this discussion is mostly relevant, as much is encoded in the distribution over Θ, and different players will have different knowledge of this.

This information asymmetry leads to the various zoology of Bayesian equilibria. Complete knowledge of Θ leads to ex post equilibria, 'personal' knowledge to ex interim, and no knowledge at all to ex ante. Indeed, notice that ex ante means knowing the distribution, ex interim means knowing it conditioned on the player's own actual type, and ex post conditioned on everyone's actual type.

A selection lens for a Bayesian game looks as follows:

$$\begin{pmatrix} D\,\mathrm{supp}_\Omega \\ \Omega \end{pmatrix} \overset{\varepsilon}{\Lleftarrow} \begin{pmatrix} \mathrm{supp}^*_\Omega \\ \Omega \end{pmatrix} \tag{2.2}$$

The backward part looks as follows:

$$\mathrm{d}\omega \in D\Omega \vdash \varepsilon_{\mathrm{d}\omega} : D\mathbb{R}^{\mathrm{supp}_\Omega \mathrm{d}\omega} \to D\,\mathrm{supp}_\Omega\, \mathrm{d}\omega \tag{2.3}$$

Notice such a map could factor through the algebra map of $\mathbb{R}$, $\mathbb{E} : D\mathbb{R} \to \mathbb{R}$, that yields a map $\binom{\mathbb{E}}{\Omega} : \binom{\mathbb{R}^{\mathrm{supp}_\Omega}}{\Omega} \leftrightarrows \binom{D\mathbb{R}^{\mathrm{supp}_\Omega}}{\Omega}$. If ε factors like this, it means the selection depends on expected utility. Then the remaining part is a convex map

$$\mathrm{d}\omega \in \bar{\varepsilon}_{\mathrm{d}\omega} : D\Omega \vdash \mathbb{R}^{\mathrm{supp}_\Omega \mathrm{d}\omega} \to D\,\mathrm{supp}_\Omega\, \mathrm{d}\omega \tag{2.4}$$

Convexity for this map means that, given $u, v : \mathrm{supp}_\Omega \mathrm{d}\omega \to \mathbb{R}$ and $p \in [0, 1]$:

$$\bar{\varepsilon}_{\mathrm{d}\omega}(pu + (1-p)v) = p\bar{\varepsilon}_{\mathrm{d}\omega}(u) + (1-p)\bar{\varepsilon}_{\mathrm{d}\omega}(v) \tag{2.5}$$

Since $u = \sum_{\omega \in \mathrm{supp}_\Omega \mathrm{d}\omega} u(\omega)\delta_\omega$, calling $\bar{u} = u / \int u \mathrm{d}\omega$, we get

$$\bar{\varepsilon}_{\mathrm{d}\omega}(\bar{u}) = \int \bar{u}(\omega)\bar{\varepsilon}_{\mathrm{d}\omega}(\delta_\omega) \tag{2.6}$$

Hence $\bar{\varepsilon}_{\delta\omega}$ is substantially determined by its values on the 'points' of $\mathrm{supp}_\Omega \mathrm{d}\omega$. A very sensible choice is $\bar{\varepsilon}_{\mathrm{d}\omega}(\delta_\omega) = \delta_\omega$, which leads to updating a strategy using Luce's rule and then to adopt so called *Boltzmann rationality*.