



HAL
open science

Pointless learning (long version)

Florence Clerc, Vincent Danos, Fredrik Dahlqvist, Ilias Garnier

► **To cite this version:**

Florence Clerc, Vincent Danos, Fredrik Dahlqvist, Ilias Garnier. Pointless learning (long version). Proceedings of FoSSaCS 2017, 2017, Uppsala, Sweden. hal-01974692

HAL Id: hal-01974692

<https://hal.science/hal-01974692>

Submitted on 8 Jan 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Pointless learning (long version)

Florence Clerc, Vincent Danos, Fredrik Dahlqvist, Ilias Garnier

► **To cite this version:**

Florence Clerc, Vincent Danos, Fredrik Dahlqvist, Ilias Garnier. Pointless learning (long version). Accepted to the 20th International Conference on Foundations of Software Science and Computation .. 2017. <hal-01429663v2>

HAL Id: hal-01429663

<https://hal.archives-ouvertes.fr/hal-01429663v2>

Submitted on 23 Jan 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Pointless Learning (long version)

Florence Clerc^{*1}, Vincent Danos^{2,4}, Fredrik Dahlqvist^{**3}, and Ilias Garnier^{☒***4}

¹ McGill University

² ENS Paris/CNRS

³ UCL

⁴ University of Edinburgh

☒ ilias.gar@gmail.com

Abstract. Bayesian inversion is at the heart of probabilistic programming and more generally machine learning. Understanding inversion is made difficult by the pointful (kernel-centric) point of view usually taken in the literature. We develop a pointless (kernel-free) approach to inversion. While doing so, we revisit some foundational objects of probability theory, unravel their category-theoretical underpinnings and show how pointless Bayesian inversion sits naturally at the centre of this construction.

1 Introduction

The soaring success of Bayesian machine learning has yet to be matched with a proper foundational understanding of the techniques at play. These statistical models are fundamentally programs that manipulate probability distributions. Therefore, the semantics of programming languages can and should inform the semantics of machine learning. This point of view, upheld by the proponents of *probabilistic programming*, has given rise to a growing body of work on matters ranging from the computability of disintegrations [1] to operational and denotational semantics of probabilistic programming languages [12]. These past approaches have all relied on a pointful, *kernel-centric* view of the key operation in Bayesian learning, namely *Bayesian inversion*. In this paper, we show that a *pointless*, operator-based approach to Bayesian inversion is both more general, simpler and offers a more structured view of Bayesian machine learning.

Let us recall the underpinnings of Bayesian inversion in the finite case. Bayesian statistical inference is a method for updating subjective probabilities on an unknown random process as observations are collected. In a finite setting, this update mechanism is captured by Bayes' law:

$$P(d) \cdot P(h \mid d) = P(d \mid h) \cdot P(h) \tag{1}$$

* Research supported by NSERC, Canada.

** Research partially supported by the ERC project *ProFoundNet - Probabilistic Foundations for Networks* (grant number 679127).

*** Research supported by the ERC project *RULE* (grant number 320823).

On the right-hand side, the *likelihood* $P(d | h)$ encodes a parameter-dependent probability over data, weighted by the *prior* $P(h)$ which corresponds to our current belief on which parameters best fit the law underlying the unknown random process. The left-hand side of Eq.1 involves the *marginal likelihood* $P(d)$, which is the probability of observing the data d under the current subjective probability, and the *posterior* $P(h | d)$ which tells us how well the occurrence of d is explained by the parameter h . More operationally, the posterior tells us how we should revise our prior as a function of the observed data d . In a typical Bayesian setup, the prior and likelihood are given and the marginal likelihood can be computed from the two first ingredients. The only unknown is the *posterior* $P(h | d)$. Eq. 1 allows one to compute the posterior from the two first ingredients—whenever $P(d) > 0$! This formulation emphasises the fundamental symmetry between likelihood and posterior, and hopefully makes clear why the process of computing the posterior is called *Bayesian inversion*. The key observation is that both the likelihood and posterior can be seen as matrices, and Eq. 1 encodes nothing more than a relation of adjunction between these matrices seen as (finite-dimensional) operators. This simple change of point of view, where one thinks no longer directly in terms of kernels (which transform probability measures forward), but in terms of their semantics as operators (which transform real-valued observables backward) generalises well and gives us a much more comprehensive account of Bayesian learning as adjunction. If one thinks of observables as extended predicates, this change of point of view is nothing but a predicate transformer semantics of kernels: a well-established idea planted in the domain of probabilistic semantics by Kozen in the 80s [10]. The object of this paper is to develop in this setting a *pointless* approach to Bayesian inversion.

Our contributions are as follows. In Sec. 3, we recall how Bayesian inversion is formulated using the language of kernels, following the seminal work of [5] and our own preliminary elaboration of the ideas developed in the current paper [6]. The adequate setting is a category of *typed* kernels, i.e. measure-preserving kernels between probability spaces. We observe that Bayesian inversion fits somewhat awkwardly in this pointful setting. Drawing from domain-theoretic ideas [11], we develop in Sec. 4 a categorical theory of ordered Banach cones, including an adjunction theorem for L_p^+/L_q^+ cones taken from Ref. [4]. In Sec. 5, we define a functorial operator interpretation of kernels in the category of Banach cones and prove that pointful Bayesian inversion corresponds through this functorial bridge to adjunction, expanding our recent result [6] to arbitrary L_p^+/L_q^+ cones. Unlike the pointful case, the pointless, adjunction-based approach works with arbitrary measurable spaces. Finally, in Sec. 6 we extract from the pointful and pointless approaches what we consider to be the essence of Bayesian inversion: a correspondence between *couplings* and linear operators. In this new light, adjunction (and therefore Bayesian inversion) is nothing more than a permutation of coordinates. We conclude with a sketch of some directions for future research where one could most profit of the superior agility and extension of the pointless approach.

2 Preliminaries

We refer the reader to e.g. [2] for the concepts of measure theory and functional analysis used in this paper. For convenience, some basic definitions are recalled here.

A *measurable space* (X, Σ) is given by a set X together with a σ -algebra of subsets of X denoted by Σ . Where unambiguous, we will omit the σ -algebra and denote a measurable space by its underlying set. We will also consider the measurable spaces generated from *Polish* (completely metrisable and separable) topological spaces, called *standard Borel spaces* [9]. A measurable function $f : (X, \Sigma) \rightarrow (Y, A)$ is a function $f : X \rightarrow Y$ such that for all $B \in A$, $f^{-1}(B) \in \Sigma$. The category of measurable spaces and measurable functions will be denoted by **Mes**. For B a measurable set, we denote by $\mathbb{1}_B$ the indicator function of that set. A *finite measure* μ over a measurable space (X, Σ) is a σ -additive function $\mu : \Sigma \rightarrow [0, \infty)$ that verifies $\mu(X) < \infty$. Whenever $\mu(X) = 1$, μ is a *probability measure*. A pair (X, μ) with X a measurable space and μ a probability measure on X is called a *probability space*. A measurable set B will be qualified of μ -*null* if $\mu(B) = 0$.

The Giry endofunctor, denoted by $\mathbf{G} : \mathbf{Mes} \rightarrow \mathbf{Mes}$, maps each measurable space X to the space $\mathbf{G}(X)$ of probability measures over X . The measurable structure of $\mathbf{G}(X)$ is given by the initial σ -algebra for the family $\{ev_B\}_B$ of evaluation maps $ev_B(\mu) = \mu(B)$, where B ranges over measurable sets in X . The action of \mathbf{G} on arrows is given by the pushforward (or image measure): for $f : X \rightarrow Y$ measurable, we have $\mathbf{G}(f) : \mathbf{G}(X) \rightarrow \mathbf{G}(Y)$ given by $\mathbf{G}(f)(\mu) = \mu \circ f^{-1}$. This functor admits the familiar monad structure (\mathbf{G}, m, δ) where $m : \mathbf{G}^2 \Rightarrow \mathbf{G}$ and $\delta : Id \Rightarrow \mathbf{G}$ are natural transformations with components at X defined by $m_X(P)(B) = \int_{\mathbf{G}(X)} ev_B dP$ and $\delta_X(x)(B) = \delta_x(B)$. It is well-known that when restricted to standard Borel spaces, the Giry functor admits the same monad structure. See [7] for more details on this construction. The Kleisli category of the Giry monad, corresponding to Lawvere's category of probabilistic maps, will be denoted by \mathcal{Kl} . The objects of \mathcal{Kl} correspond to those of **Mes** and arrows from X to Y correspond to so-called *kernels* $f : X \rightarrow \mathbf{G}(Y)$. Kleisli arrows will be denoted by $f : X \rightarrow Y$. For $f : X \rightarrow Y, g : Y \rightarrow Z$, the Kleisli composition is defined as usual by $g \circ' f = m_Z \circ \mathbf{G}(g) \circ f$. We distinguish deterministic Kleisli maps as those that can be factored as a measurable function followed by δ and denote these arrows $f : X \rightarrow_{\delta} Y$. We write 1 for the one element measurable space (which is the terminal object in **Mes**). Clearly the Homset $\mathcal{Kl}(1, Y)$ is in bijection with the set of probabilities over Y . This justifies the following slight abuse of notation: if $\mu \in \mathbf{G}(X)$ is a probability and $f : X \rightarrow Y$ is a kernel, the pushforward of μ through f will be denoted $f \circ' \mu$. Observe that for $f : X \rightarrow Y$ an usual measurable map, $\mathbf{G}(f)(\mu) = (\delta_Y \circ f) \circ' \mu$, so the pushforward through a kernel extends the earlier definition.

Consider the full subcategory of \mathcal{Kl} restricted to finite spaces. In that setting, any kernel $f : X \rightarrow Y$ can be presented as a positive, real-valued matrix that we denote $T(f) = \{f(x)(y)\}_{x,y}$ with X rows, Y columns and where all rows sum to 1 (aka a stochastic matrix). Matrix multiplication corresponds to Kleisli

composition: taking f, g as above, one has $T(g \circ' f) = T(f)T(g)$ (hence, this representation of kernels as matrices is *contravariant*). Such matrices act on vectors of dimension Y (*observables on Y*) and map them to observables on X : for $v \in \mathbb{R}^Y$, $T(f)v$ corresponds to the expectation of v according to f . This is the basis for the “operator interpretation” of kernels which we will extend to **Mes** below.

3 Bayesian inversion in a category of typed kernels

We introduce the category **Krn** of typed kernels and recall the statement of Bayesian inversion in this setting.

3.1 Definition of Krn

Our starting point is the under category $1 \downarrow \mathcal{Kl}$, where 1 is the one-element measurable space. Objects of $1 \downarrow \mathcal{Kl}$ are Kleisli arrows $\mu : 1 \rightarrow X$, i.e. probability spaces (X, μ) with $\mu \in \mathbf{G}(X)$; while typed kernels from (X, μ) to (Y, ν) are Kleisli arrows $f : X \rightarrow Y$ such that $f \circ' \mu = \nu$. We will call these arrows “kernels” for short. For a deterministic map $f_\delta : X \rightarrow_\delta Y$ (factoring as $f_\delta = \delta_Y \circ f$), the constraint boils down to $\nu = \mathbf{G}(f)(\mu)$. In other words, the subcategory of $1 \downarrow \mathcal{Kl}$ consisting of deterministic maps is isomorphic to the usual category of probability spaces and measure-preserving maps. We define **Krn** to be the subcategory of $1 \downarrow \mathcal{Kl}$ restricted to standard Borel spaces.

3.2 Bayesian inversion in the finite subcategory of Krn

We translate the presentation of Bayesian inversion of Sec. 1 in the language of **Krn**. We are given finite spaces of data D and parameters H and it is assumed that there exists an unknown probability on D , called the “truth” and denoted τ in the following, that we wish to learn. The likelihood corresponds to a \mathcal{Kl} arrow $f : H \rightarrow D$. The prior is a probability $\mu \in \mathbf{G}(H)$ while the marginal likelihood $\nu \in \mathbf{G}(D)$ is obtained as $\nu = f \circ' \mu$. Thus the entire situation is captured by a **Krn** arrow $f : (H, \mu) \rightarrow (D, \nu)$. If our prior was perfect, we would have $\nu = \tau$ but of course (by assumption) this is not the case! The only access we have to the truth is through an infinite, independent family $\{d_n\}_{n \in \mathbb{N}}$ of random elements in D each distributed according to τ . The Bayesian update is the process of using this sequence of data (sometimes called *evidence*) to iteratively revise our prior. In this language, Bayes’s law reads as follows:

$$\nu(d) \cdot f^\dagger(d)(h) = f(h)(d) \cdot \mu(h) \tag{2}$$

where $f^\dagger : (D, \nu) \rightarrow (H, \mu)$ denotes the sought posterior map, to be computed in function of μ and f . Observe that both the left and right hand side of Eq. 2 define the same *joint probability* $\gamma \in \mathbf{G}(H \times D)$ given by $\gamma(h, d) = f(h)(d) \cdot \mu(h) = \nu(d) \cdot f^\dagger(d)(h)$. Denoting π_H, π_D the left and right projections from $H \times D$, one easily verifies that $\mathbf{G}(\pi_H) = \mu$ and $\mathbf{G}(\pi_D) = \nu$. In other terms, γ is a *coupling* of μ and ν . We draw the attention of the reader to the following points.

- As hinted before, $f^\dagger(d)$ is uniquely defined only when $\nu(d) > 0$. Conversely, f^\dagger does not depend on f on μ -null sets. These hurdles will be circumvented by considering equivalence classes of kernels up to null sets. This is the object of Sec. 3.3.
- Sec. 2 introduces a correspondence between (finite) kernels and Markov or stochastic matrices. This raises the following question: what is Bayesian inversion seen through that lens? The answer is *adjunction*. As we show in Sec. 5, this *pointless* point of view generalises to arbitrary measurable spaces and is better behaved than the *pointful* one.

We now proceed to the generalisation of this situation to the case of standard Borel spaces, i.e. to that of **Krn**.

3.3 Bayesian inversion in **Krn**

Bayesian inversion in **Krn** relies crucially on the construction of an (almost sure) bijection between the **Krn** Homset $\mathbf{Krn}(X, \mu; Y, \nu)$ and the set of couplings $\mathbf{\Gamma}(X, \mu; Y, \nu)$ of μ and ν (to be defined next).

Couplings and kernels. To any pair of objects $(X, \mu), (Y, \nu)$, one can associate the space of *couplings* of μ and ν , i.e. the set of all probabilities $\gamma \in \mathbf{G}(X \times Y)$ such that $\mathbf{G}(\pi_X)(\gamma) = \mu$ and $\mathbf{G}(\pi_Y)(\gamma) = \nu$. We denote this set of couplings $\mathbf{\Gamma}(X, \mu; Y, \nu)$. It is a standard Borel space, as the set of couplings of two measures is a closed convex subset in $\mathbf{G}(X \times Y)$ for any choice of a Polish topology for X, Y . In order to construct a mapping from couplings to **Krn** arrows, we will need the *disintegration theorem*, which requires us to introduce some terminology. In the following, we denote $N(f, f') = \{x \mid f(x) \neq f'(x)\}$.

Lemma 1. *For all $f, f' : (X, \mu) \rightarrow (Y, \nu)$, $N(f, f')$ is measurable.*

Proof. We work with standard Borel spaces, hence two measures ρ, ρ' on Y are equal if and only if they coincide on a countable generating π -system $\{B_n\}_{n \in \mathbb{N}}$ of the σ -algebra of Y (this follows from the Carathéodory extension theorem [14]). Dually, if $f(x) \neq f'(x)$ then there must exist an n such that $f(x)(B_n) \neq f'(x)(B_n)$. Therefore, $N(f, f') = \cup_n \{x \mid f(x)(B_n) \neq f'(x)(B_n)\}$. Each set $C_n = \{x \mid f(x)(B_n) \neq f'(x)(B_n)\}$ can be written as

$$C_n = \{x \mid (ev_{B_n} \circ f)(x) \neq (ev_{B_n} \circ f')(x)\}$$

which is measurable. A countable union of measurable sets is measurable, hence so is $N(f, f')$. \square

Note that in more general measurable spaces, $N(f, f')$ is not necessarily a measurable set, as those spaces are not always countably generated. We can now introduce the disintegration theorem.

Theorem 1 (Disintegration ([8], Thm. 5.4)). *For all deterministic \mathbf{Krn} arrow $f : (X, \mu) \rightarrow_{\delta} (Y, \nu)$, there exists $f^{\dagger} : (Y, \nu) \rightarrow (X, \mu)$ such that $f \circ' f^{\dagger} = id_{(Y, \nu)}$ and such that for all $h : (Y, \nu) \rightarrow (X, \mu)$ verifying $f \circ' h = id_{(Y, \nu)}$, $\nu(N(f^{\dagger}, h)) = 0$. In short, we say that f^{\dagger} is the ν -almost surely unique kernel verifying $f \circ' f^{\dagger} = id_{(Y, \nu)}$.*

Disintegrations correspond to *regular conditional probabilities* (see e.g. [8]). The deterministic map $f : X \rightarrow Y$ along which the disintegration of μ is computed acts through its fibers as a parameterised family of subsets on each of which μ is conditioned, resulting in a measurable family of conditional probabilities parameterised by Y . Note that the characteristic property of disintegrations can be equivalently stated as the fact that $f^{\dagger}(y)$ is ν -almost surely supported by $f^{-1}(y)$.

Example 1. In the finite case, disintegration is simply the formula for conditional probabilities. Given X, Y finite and $f : (X, \mu) \rightarrow_{\delta} (Y, \nu)$, for $y \in Y$ s.t. $\nu(y) = \mu(f^{-1}(y)) > 0$, it holds that $f^{\dagger}(y)(x) = \frac{\mu(x)}{\nu(y)}$. However, when $\nu(y) = 0$, the disintegration theorem does not constrain the value of $f^{\dagger}(y)$.

Disintegration establishes a bijective (up to null sets) correspondence between couplings and kernels. Let us make this formal.

Definition 1. *For fixed $(X, \mu), (Y, \nu)$, we define on $\mathbf{Krn}(X, \mu; Y, \nu) \sim$ as the smallest equivalence relation such that $f \sim f'$ if $\mu(N(f, f')) = 0$. We denote $\mathbf{Krn}(X, \mu; Y, \nu)/\mu$ the set of \sim -equivalence classes of $\mathbf{Krn}(X, \mu; Y, \nu)$.*

Any \mathbf{Krn} arrow $f : (X, \mu) \rightarrow (Y, \nu)$ induces a measure on $X \times Y$, defined on measurable rectangles $B_X \times B_Y$ as:

$$I_{X, \mu}^{Y, \nu}(f)(B_X \times B_Y) = \int_{x \in B_X} f(x)(B_Y) d\mu. \quad (3)$$

Lemma 2. $I_{X, \mu}^{Y, \nu}$ is a **Set** injection from $\mathbf{Krn}(X, \mu; Y, \nu)/\mu$ to $\Gamma(X, \mu; Y, \nu)$.

Proof. One easily verifies that $I_{X, \mu}^{Y, \nu}(f)$ is a coupling of μ and ν by evaluating Eq. 3 for respectively $B_X = X$ and $B_Y = Y$. Let us prove injectivity. For $f \not\sim f' : (X, \mu) \rightarrow (Y, \nu)$, let $N(f, f')$ be as in Sec. 3.1. Y is standard Borel, hence its σ -algebra is generated by a countable π -system $\{B_n\}_{n \in \mathbb{N}}$ and it is enough to test measures for equality on this family. Therefore, $N(f, f') = \cup_{n \in \mathbb{N}} \{x \mid f(x)(B_n) \neq f'(x)(B_n)\}$. Since $\mu(N(f, f')) > 0$, we can construct measurable sets A in X and B in Y s.t. $I_{X, \mu}^{Y, \nu}(f)(A \times B) \neq I_{X, \mu}^{Y, \nu}(f')(A \times B)$, from which we conclude that $I_{X, \mu}^{Y, \nu}$ is injective. \square

The second part of the bijection between couplings and quotiented \mathbf{Krn} arrows relies crucially on disintegration.

Lemma 3. *There is a **Set** injection $D_{X, \mu}^{Y, \nu} : \Gamma(X, \mu; Y, \nu) \rightarrow \mathbf{Krn}(X, \mu; Y, \nu)/\mu$. Moreover, $D_{X, \mu}^{Y, \nu}$ and $I_{X, \mu}^{Y, \nu}$ are inverse of one another.*

Proof. Any coupling $\gamma \in \mathbf{\Gamma}(X, \mu; Y, \nu)$ induces two (equivalence classes of) **Krn** arrows by disintegrating along the projections, namely $\pi_X^\dagger : (X, \mu) \rightarrow (X \times Y, \gamma)$ and $\pi_Y^\dagger : (Y, \nu) \rightarrow (X \times Y, \gamma)$. Postcomposing with the adequate projections, we get from γ an equivalence class of kernels $\mathbf{G}(\pi_Y) \circ \pi_X^\dagger : (X, \mu) \rightarrow (Y, \nu)$. We set $D_{X,\mu}^{Y,\nu}(\gamma) = \mathbf{G}(\pi_Y) \circ \pi_X^\dagger$. Let $\gamma \in \mathbf{\Gamma}(X, \mu; Y, \nu)$ be a coupling. We have:

$$\begin{aligned} I_{X,\mu}^{Y,\nu} \circ D_{X,\mu}^{Y,\nu}(\gamma)(B_X \times B_Y) &= \int_{x \in B_X} D_{X,\mu}^{Y,\nu}(\gamma)(x)(B_Y) d\mu \\ &= \int_{x \in B_X} (\mathbf{G}(\pi_Y) \circ \pi_X^\dagger(x))(B_X \times B_Y) d\mu \\ &= \int_{x \in B_X} \pi_X^\dagger(x)(X \times B_Y) d\mu \\ &\stackrel{(1)}{=} \int_{x \in X} \pi_X^\dagger(x)(B_X \times B_Y) d\mu \stackrel{(2)}{=} \gamma \end{aligned}$$

where (1) follows from the fact that $\pi_X^\dagger(x)$ is concentrated on $\{x\} \times Y$ and (2) follows from the characteristic property of disintegrations. We deduce that $D_{X,\mu}^{Y,\nu}$ must be injective.

Conversely, let $f : (X, \mu) \rightarrow (Y, \nu)$ be a kernel and let us write $\delta_X \otimes f = x \mapsto \delta_X(x) \otimes f(x)$, where \otimes denotes the product of probabilities. Clearly, $\delta_X \otimes f : (X, \mu) \rightarrow (X \times Y, \gamma)$. Trivially, $(\delta_X \circ \pi_X) \circ' (\delta_X \otimes f) = id_{(X,\mu)}$. We deduce that $\pi_X^\dagger \sim (\delta_X \otimes f)$, hence that $D_{X,\mu}^{Y,\nu} \circ I_{X,\mu}^{Y,\nu}(f) = \mathbf{G}(\pi_Y) \circ \pi_X^\dagger \sim f$. We conclude that $I_{X,\mu}^{Y,\nu}$ and $D_{X,\mu}^{Y,\nu}$ are inverse to each other. \square

*Bayesian inversion in **Krn**.* Bayesian inversion corresponds to the composition of the bijections we just defined with the pushforward along the permutation map $\sigma : X \times Y \rightarrow Y \times X$.

Theorem 2 (Bayesian inversion). *Let $-^\dagger$ be defined as $f^\dagger = D_{Y,\mu}^{X,\mu} \circ \mathbf{G}(\sigma) \circ I_{X,\mu}^{Y,\nu}$. The map $-^\dagger : \mathbf{Krn}(X, \mu; Y, \nu)/\mu \rightarrow \mathbf{Krn}(Y, \nu; X, \mu)/\nu$ is a bijection.*

Proof. By Lemma 2 and Lemma 3. \square

This section would be incomplete if we didn't address *learning* in its relation to Bayesian inversion. It is known that in good cases,⁵ Bayesian inversion will make the sequence of marginal likelihoods converge to the truth in some appropriate topology. However, issues of convergence are not the subject of this paper and will not be discussed further.

3.4 Pointfulness is harmful

Let us take a critical look at the approach to Bayesian inversion developed so far. The fact that $-^\dagger$ is by construction \sim -invariant and yields \sim -equivalence classes of **Krn** arrows suggests that $-^\dagger$ would be better typed on a hypothetical

⁵ E.g. H, D finite and μ putting strictly positive measure on $f^{-1}(\tau)$

quotient of \mathbf{Krn} by \sim . This mismatch between the behaviour of $-^\dagger$ and its actual type already arises in the finite case where Bayes' rule yields kernels only defined up to a null set (see discussion after Eq. 2), and is an inevitable consequence of the pointful point of view: kernels should respect the measures endogenous to their domain. Constructing the quotient of \mathbf{Krn} w.r.t. \sim would require proving that this equivalence relation is compatible with the composition of \mathbf{Krn} . However, carrying out this approach successfully seems non-trivial:⁶ our past attempts are riddled with obstructions stemming from accumulation of negligible sets—the very technical hurdles that make the theory of disintegration of measures so unintuitive in the first place, while moreover relying on standard Borel assumptions.

This improper typing obscures the categorical structure of Bayesian inversion. In the next sections, we leave the inhospitable world of kernels and relocate the theory of Bayesian inversion in a category of Banach cones and linear maps where these problems vanish, and the structure we seek for becomes manifest.

4 Banach cones

Following [11] and [4], we introduce a category of Banach cones and ω -continuous linear maps, with the intent of interpreting Markov kernels as linear operators between well-chosen function spaces. In the subcategory corresponding to these function spaces, we develop a powerful adjunction theorem that will be used in Sec. 5 to implement pointless Bayesian inversion.

4.1 The category \mathbf{Ban}

A *Banach cone*, informally, corresponds to a normed convex cone of a Banach space which is ω -complete with respect to a particular order. Let us introduce these cones progressively.

Definition 2. *A normed, convex cone $(C, +, \cdot, 0, \|\cdot\|_C)$ of a normed vector space $(V, +, \cdot, 0, \|\cdot\|_V)$ is a subset $C \subseteq V$ that is closed under addition, convex combinations and multiplications by non-negative scalars, endowed with the restriction of the ambient norm, which must be monotone w.r.t. the partial order $u \leq_C v \Leftrightarrow \exists w \in C. u + w = v$.*

We require our Banach cones to be ω -complete with respect to this order, and to be subsets of Banach spaces.

Definition 3 (Banach cones). *A normed convex cone C is ω -complete if for all chain (i.e. \leq_C -increasing countable family) $\{u_n\}_{n \in \mathbb{N}}$ of bounded norm, the least upper bound $\bigvee_n u_n$ exists and $\|\bigvee_n u_n\|_C = \bigvee_n \|u_n\|_C$. A Banach cone is an ω -complete normed cone of a Banach space.*

⁶ Without additional assumptions the quotient is not compatible with pre-composition, differently to what we mistakenly stated in ([6], Lemma 3).

Norm convergence and order convergence are related by the following result.

Lemma 4 ([4], **Lemma 2.12**). *Let $\{u_n\}_{n \in \mathbb{N}}$ be a chain of bounded norm in a Banach cone. Then $\lim_{i \rightarrow \infty} \|\bigvee_n u_n - u_i\| = 0$.*

A prime example of Banach cones is given by the *positive cones* associated to classical L_p spaces of real-valued functions (see e.g. [2] for a definition of those spaces). In details: for (X, μ) a measure space with μ finite and $p \in [1, \infty]$, the set of elements $f \in L_p(X, \mu)$ which are non-negative μ -a.e. is closed under addition, multiplication by non-negative scalars and under linear combinations with non-negative coefficients. Equipped with the restriction of the norm of $L_p(X, \mu)$, this subset forms a normed convex cone that we denote $L_p^+(X, \mu)$. The partial order associated to these L_p^+ cones can be defined explicitly: for $f, g \in L_p^+(X, \mu)$, we write that $f \leq g$ if $f(x) \leq g(x)$ μ -a.e. One easily checks that this coincides with the definitional partial order.

Proposition 1 (ω -completeness of L_p^+ cones, [4]). *For all X measurable, $\mu \in \mathbf{G}(X)$ and $p \in [1, \infty]$, $L_p^+(X, \mu)$ is a Banach cone.*

This result is a direct consequence of the definition of suprema in $L_p^+(X, \mu)$. We are going to construct a category of all Banach cones and we thus have to specify what a morphism between such cones is. We consider only linear maps which are Scott-continuous, which in this case⁷ boils down to commuting with supremas of increasing chains.

Definition 4. *Let C, C' be Banach cones and $A : C \rightarrow C'$ be a linear map. A is ω -continuous if for every chain $\{f_n\}_{n \in \mathbb{N}}$ such that $\bigvee_n f_n$ exists, $A(\bigvee_n f_n) = \bigvee_n A(f_n)$.*

The following example should help make ω -continuity less mysterious. Observe that for $Y = 1$ (the singleton set), all Banach cones $L_p^+(Y, \mu)$ (for μ nonzero, otherwise $L_p^+(Y, \mu) \cong \{0\}$) are isomorphic to $\mathbb{R}_{\geq 0}$ – therefore, $\mathbb{R}_{\geq 0}$ is a *bona fide* Banach cone.

Example 2. There exists a familiar linear map from $L_p^+(X, \mu)$ to $\mathbb{R}_{\geq 0}$, namely the Lebesgue integral $\int : L_p^+(X, \mu) \rightarrow \mathbb{R}_{\geq 0}$, taking $u \in L_p^+(X, \mu)$ to $\int_X u \, d\mu$. In this case, ω -continuity of the integral is simply the monotone convergence theorem!

Unless stated otherwise, all maps in the remainder of this section are ω -continuous. The property of ω -continuity is closed under composition and the identity function is trivially ω -continuous. This takes us to the following definition.

⁷ These cones have the “countable sup property” [2]. Therefore, all directed sets admit a countable subset having the same least upper bound, and we can restrict our attention to chains.

Definition 5 (Categories of Banach cones and of L_p^+ cones). *The category \mathbf{Ban} has Banach cones as objects and ω -continuous linear maps as morphisms. We distinguish the full subcategory \mathbf{L} having as objects all L_p^+ -spaces (ranging over all $p \in [1, \infty]$). Further, \mathbf{L} admits a family of full subcategories $\{\mathbf{Lp}\}_{p \in [1, \infty]}$, each having as objects L_p^+ spaces (for fixed p).*

\mathbf{Ban} is itself a full subcategory of the category $\omega\mathbf{CC}$ of ω -complete normed cones and ω -continuous maps, as defined in [4]. Let us denote by $\mathbf{Ban}(C, C')$ the set of ω -continuous linear maps from C to C' . Denoting $\|\cdot\|_C$ the norm of C , we recall that the operator norm of a linear map $A : C \rightarrow C'$ is given by $\|A\|_{op} = \inf \{K \geq 0 \mid \forall u \in C, \|Au\|_{C'} \leq K \|u\|_C\}$. A partial order on $\mathbf{Ban}(C, C')$ is given by $A \leq B$ iff for all $u \in C$, $A(u) \leq_{C'} B(u)$. Selinger proved in [11] that ω -continuous linear maps between ω -complete cones have automatically bounded norm (i.e. they are continuous in the usual sense), therefore we can and will abstain from asking continuity explicitly. The following result is a cone-theoretic counterpart to the well-known fact that the vector space of bounded linear operators between two Banach spaces forms a Banach space for the operator norm.

Proposition 2. *For all Banach cones C, C' , the cone of ω -continuous linear maps $\mathbf{Ban}(C, C')$ is a Banach cone for the operator norm and the pointwise order.*

Proof. It remains to prove that $\mathbf{Ban}(C, C')$ is ω -complete. Let us check that the pointwise order corresponds to the definitional cone order. Assume $A \leq B$ pointwise. We need to prove that $B - A \in \mathbf{Ban}(C, C')$, which amounts to prove that $B - A$ is ω -continuous. Let $\{u_n\}_{n \in \mathbb{N}}$ be a chain s.t. $\bigvee_n u_n$ exists. By ω -continuity of A, B , $(B - A)(\bigvee_n u_n) = \bigvee_n B(u_n) - \bigvee_n A(u_n)$. It is enough to prove that $\|\bigvee_n B(u_n) - \bigvee_n A(u_n) - (\bigvee_n B(u_n) - A(u_n))\| = 0$. Using Lemma 4, it suffices to prove that

$$\lim_{k \rightarrow \infty} \left\| \bigvee_n B(u_n) - \bigvee_n A(u_n) - (B(u_k) - A(u_k)) \right\| = 0.$$

Notice that $\bigvee_n B(u_n) = B(\bigvee_n u_n)$ implies $\lim_k \|\bigvee_n B(u_n) - B(u_k)\| = 0$ (using Lemma 4) and similarly for A . An application of the triangle inequality allows us to conclude.

Let us prove that $\mathbf{Ban}(C, C')$ is ω -complete. Let $\{A_n\}_{n \in \mathbb{N}}$ be a chain in the pointwise order, s.t. $\{\|A_n\|\}_n$ is bounded. Therefore, there exists $K \geq 0$ s.t. for all u , $\|A_n(u)\|_{C'} \leq K \|u\|$. By ω -completeness of C' , $\bigvee_n A_n(u)$ exists. For all $u \in C$, we set $A(u) = \bigvee_n A_n(u)$. Linearity is trivial. Since the norms are ω -continuous, $\|A(u)\| = \bigvee_n \|A_n(u)\|$ and therefore, A is precisely of norm $\bigvee_n \|A_n\|$. \square

4.2 Duality in Banach cones

We use a powerful Banach cone duality result initially proved in the supplementary material to [4]. We say that a pair (p, q) with $p, q \in [1, \infty]$ is *Hölder*

conjugate if $\frac{1}{p} + \frac{1}{q} = 1$. For any Banach cone C , its dual C^* is by definition the Banach cone of ω -continuous linear functionals, i.e. the cone $C^* = \mathbf{Ban}(C, \mathbb{R}_{\geq 0})$. This operation defines a contravariant endofunctor $-^* : \mathbf{Ban} \rightarrow \mathbf{Ban}^{op}$ mapping each cone C to C^* and each map of cone $A : C \rightarrow C'$ to the map $A^* : C'^* \rightarrow C^*$ defined by $A^*(\varphi) = \varphi \circ A$, for $\varphi \in C'^*$. For Hölder conjugate (p, q) , we have the following extension to the classical isomorphism of L_p^+ spaces.

Theorem 3 (L_p^+ cone duality ([4])). *There is a Banach cone isomorphism $\varepsilon_p : L_p^+(X, \mu) \cong L_q^{+,*}(X, \mu)$.*

We won't reproduce the proof of this theorem here, which can be found in the supplementary material to [4]. Suffice it to say it is a Riesz duality type argument which relies entirely on the Radon-Nikodym theorem. Note that Theorem 3 implies in particular that $L_\infty^{+,*}(X, \mu) \cong L_1^+(X, \mu)$, which classically fails in the usual setting of L_p Banach spaces. It is instructive to study how ω -continuity wards off a classical counter-example to duality in the general Banach case.

Example 3 (Taken from [11]). Let μ be a probability measure on \mathbb{N} with full support. We consider the cone $\ell_\infty^+ = L_\infty^+(\mathbb{N}, \mu)$ of bounded sequences of real numbers. Let \mathcal{U} be a non-principal ultrafilter on \mathbb{N} (i.e. an ultrafilter on the partial order of subsets of \mathbb{N} without a least element). We define the function $\lim_{\mathcal{U}} : \ell_\infty^+ \rightarrow \mathbb{R}$ as $\lim_{\mathcal{U}}(\{x_n\}_{n \in \mathbb{N}}) = \sup\{y \mid \{n \mid x_n \geq y\} \in \mathcal{U}\}$. This function is linear and bounded. However, consider the chain $\{u^k \in \ell_\infty^+\}_{k \in \mathbb{N}}$ with $u_n^k = 1$ for all $n \leq k$ and $u_n^k = 0$ for all $n > k$. The supremum of this chain is the constant 1 sequence. On the other hand, we have $\lim_{\mathcal{U}}(u^k) = 0$ for all k , whereas $\lim_{\mathcal{U}}(\bigvee_k u^k) = 1$. Therefore, $\lim_{\mathcal{U}}(u^k)$ is not ω -continuous—i.e., $\lim_{\mathcal{U}} \notin \ell_\infty^{+,*}$.

It is useful to have a concrete representation of the isomorphism stated in Theorem 3. This theorem implies that for all $u \in L_p^+(X, \mu)$, there exists a unique ω -continuous linear functional $\varepsilon(u) \in L_q^{+,*}(X, \mu)$ —which must therefore correspond to $\varepsilon(u)(v) = \int_X uv \, d\mu$. The pairing between L_p^+ and L_q^+ cones that we introduce below corresponds to the evaluation of such a functional against some argument.

Definition 6 (Pairing). *For Hölder conjugate (p, q) , the pairing is the map $\langle \cdot, \cdot \rangle_X : L_p^+(X, \mu) \times L_q^+(X, \mu) \rightarrow \mathbb{R}_{\geq 0}$ defined by $\langle u, u' \rangle = \int uu' \, d\mu$.*

The pairing is bilinear, continuous and ω -continuous in each argument (consequences of the corresponding properties of the Lebesgue integral). We can now state the adjunction theorem.

4.3 Adjunctions between conjugate L_p^+ cones

It is instructive to look at Theorem 3 under a slightly more general light. Observe that $L_p^+(X, \mu)$ is isomorphic to $\mathbf{Ban}(\mathbb{R}_{\geq 0}, L_p^+(X, \mu))$: indeed, any map A in this function space is entirely constrained by linearity by its value at 1. Therefore, Theorem 3 really states a Banach cone isomorphism between




Fig. 1. Categories of cones




Fig. 2. Kernels, AMKs and MOs

$\mathbf{Ban}(\mathbb{R}_{\geq 0}, L_p^+(X, \mu))$ and $\mathbf{Ban}(L_q^+(X, \mu), \mathbb{R}_{\geq 0})$. This isomorphism generalises to the case where $\mathbb{R}_{\geq 0}$ is replaced by an arbitrary conjugate pair of cones $L_p^+(Y, \mu), L_q^+(Y, \nu)$ (i.e. s.t. (p, q) are Hölder conjugate).

Theorem 4 (L_p^+/L_q^+ adjunction). For (p, q) Hölder conjugate and for all $A : L_p^+(X, \mu) \rightarrow L_p^+(Y, \nu)$, $A^* : L_q^+(Y, \nu) \rightarrow L_q^+(X, \mu)$ is unique such that

$$\forall u \in L_p^+(X, \mu), v \in L_q^+(Y, \nu), \langle v, A(u) \rangle_Y = \langle A^*(v), u \rangle_X. \quad (4)$$

The essence of the previous theorem is neatly captured as follows.

Corollary 1. For all Hölder conjugate (p, q) , the duality functor $-^* : \mathbf{Ban} \rightarrow \mathbf{Ban}^{op}$ restricts to an equivalence of categories $-^* : \mathbf{Lp} \rightarrow \mathbf{Lq}^{op}$.

Fig. 1 recapitulates the categories of Banach cones mentioned in this section along their relationships.

5 Pointless Bayesian inversion

Krn arrows can be represented as linear maps between function spaces. This bridge allows one to manipulate Markov kernels both from the measure-theoretic side and from the functional-analytic side. Concretely, this linear interpretation of kernels is presented as a family of functors from **Krn** to **L**, the subcategory of **Ban** restricted to L_p^+ cones and ω -continuous linear maps. We show that pointful Bayesian inversion, whenever it is defined, coincides with adjunction.

5.1 Representing Krn arrows as AMKs

More precisely, kernels are associated to *abstract Markov kernels* (AMKs for short), which are a generalisation of stochastic matrices. Below, we denote by $\mathbb{1}_X$ the constant function equal to 1 on the space X . Since all measures we consider are finite, $\mathbb{1}_X \in L_p^+(X, \mu)$ for all $p \in [1, \infty]$.

Definition 7 (Abstract Markov kernels). An \mathbf{Lp} morphism $A : L_p^+(Y, \nu) \rightarrow L_p^+(X, \mu)$ is an AMK if $A(\mathbb{1}_Y) = \mathbb{1}_X$ and $\|A\| = 1$. Clearly, AMKs are closed under composition and the identity operator is trivially an AMK. \mathbf{AMK}_p is the subcategory of \mathbf{Lp} having the same objects and where morphisms are restricted to AMKs.

Example 4. Let us look at the particular case where X and Y are finite discrete spaces and μ, ν finite measures with full support. Then $L_p^+(X, \mu) \cong \mathbb{R}^X$ and similarly for $L_p^+(Y, \nu)$. Therefore, A corresponds to an $Y \times X$ matrix. The constraint that $A(\mathbb{1}_Y) = \mathbb{1}_X$ amounts to asking that the rows of A sum to 1, i.e. that A is a stochastic matrix.

The adjoint of an AMK is in general *not* an AMK. In the finite case, this reflects the fact that the transpose of a stochastic matrix is not necessarily stochastic. Adjoints of AMKs are called *Markov operators* (MOs for short). Whereas AMKs pull back observables, an MO pushes densities forward. In the following, we make use of the fact that for $p \leq q$, any $u \in L_q^+(X, \mu)$ belongs to $L_p^+(X, \mu)$.

Definition 8 (Markov operators). An arrow $A : L_p^+(X, \mu) \rightarrow L_p^+(Y, \nu)$ is an MO if for all $u \in L_p^+(X, \mu)$, $\|A(u)\|_1 = \|u\|_1$ and $\|A\| = 1$. \mathbf{MO}_p is the subcategory of \mathbf{Lp} having the same objects and where morphisms are restricted to MOs.

Notice that we require an MO to be norm preserving for the L_1^+ norm. This is a mass preservation constraint in disguise. Adjunction maps AMKs to MOs and conversely.

Proposition 3. The equivalence of categories $-^* : \mathbf{Lp} \rightarrow \mathbf{Lq}^{op}$ restricts to an equivalence of categories $-^* : \mathbf{AMK}_p \rightarrow \mathbf{MO}_q^{op}$.

Proof. It is enough to prove that an operator A is an AMK if and only if A^* is an MO. Let $A : L_p^+(Y, \nu) \rightarrow L_p^+(X, \mu)$ be an AMK. For all $u \in L_q^+(X, \mu)$, $\|u\|_1 = \langle A(\mathbb{1}_Y), u \rangle_X = \langle \mathbb{1}_Y, A^*(u) \rangle_Y = \|A^*(u)\|_1$. Conversely, let $A : L_p^+(X, \mu) \rightarrow L_p^+(Y, \nu)$ be an MO. For all $u \in L_p^+(X, \mu)$, $\langle A^*(\mathbb{1}_Y), u \rangle = \langle \mathbb{1}_Y, A(u) \rangle = \|A(u)\|_1 = \|u\|_1$, therefore $A^*(\mathbb{1}_Y)$ must be equal to $\mathbb{1}_X$. \square

We now introduce a family of contravariant functors $\mathbb{T}_p : \mathbf{Krn}^{op} \rightarrow \mathbf{AMK}_p$. On objects, we set $\mathbb{T}_p(X, \mu) = L_p^+(X, \mu)$. For $f : (X, \mu) \rightarrow (Y, \nu)$ a \mathbf{Krn} arrow, and for $v \in \mathbb{T}_p(Y, \nu) = L_p^+(Y, \nu)$, we define $\mathbb{T}_p(f)(v)(x) = \int_Y v df(x)$. The following theorem generalises the interpretation of kernels as stochastic matrices given in Sec. 2.

Theorem 5. \mathbb{T}_p is a functor from \mathbf{Krn}^{op} to \mathbf{AMK}_p .

Proof. We consider \mathbf{Krn} arrows $f : (X, \mu) \rightarrow (Y, \nu), g : (Y, \nu) \rightarrow (Z, \rho)$. Let us proceed stepwise. (i) We first consider the case $p \in [1, \infty)$. We show that

$\forall v \in L_p^+(Y, \nu), \mathbb{T}_p(f)(v) \in L_p^+(X, \mu)$. We have:

$$\begin{aligned} \int_{x \in X} (\mathbb{T}_p(f)(v)(x))^p d\mu &= \int_{x \in X} \left(\int_Y v df(x) \right)^p d\mu \\ &\leq \int_{x \in X} \int_Y v^p df(x) d\mu \quad (1) \\ &= \int_Y v^p d(f \circ' \mu) \quad (2) \\ &= \int_Y v^p d\nu < \infty \end{aligned}$$

where (1) follows from Jensen's inequality and (2) by the monotone convergence theorem – see ([7], Theorem 1, d)) for more details. Therefore $\|\mathbb{T}_p(f)(v)\|_p \leq \|v\|_p$. Observe that $\mathbb{T}_p(f)(\mathbb{1}_Y)(x) = \int_Y \mathbb{1}_Y df(x) = 1$, therefore $\mathbb{T}_p(f)(\mathbb{1}_Y) = \mathbb{1}_X$. This implies that $\|\mathbb{T}_p(f)(\mathbb{1}_Y)\|_p = \|\mathbb{1}_X\|_p = 1$. Therefore, $\mathbb{T}_p(f)$ has operator norm equal to 1 and it is ω -continuous by the monotone convergence theorem. We conclude that $\mathbb{T}_p(f)$ is an AMK. In the case $p = \infty$, given $v \in L_\infty^+(Y, \nu)$, we have by definition:

$$\begin{aligned} \|\mathbb{T}_\infty(f)(v)\|_\infty &= \inf \{ C \mid \mu \{ x \mid \mathbb{T}_p(f)(v)(x) > C \} = 0 \} \\ &= \inf \left\{ C \mid \mu \left\{ x \mid \int_Y v df(x) > C \right\} = 0 \right\} \\ &\leq \|v\|_\infty. \end{aligned}$$

The bound is reached by taking $v = \mathbb{1}_Y$, therefore $\|\mathbb{T}_p(f)\| = 1$.

(ii) We now turn to the property of \mathbb{T}_p of being a functor. Let $id' : (X, \mu) \rightarrow (X, \mu)$ be the identity at some object (X, μ) , i.e. the identity function postcomposed with the monadic unit δ ; let also be $u \in L_p^+(X, \mu)$. We have trivially $\mathbb{T}_p(id')(u)(x) = \int_X u id'(x) = u(x)$. Finally, we must prove that \mathbb{T}_p commutes with composition: we must prove $\mathbb{T}_p(g \circ' f) = \mathbb{T}_p(f)\mathbb{T}_p(g)$. For all $w \in L_p^+(Z, \rho)$, we have:

$$\begin{aligned} \mathbb{T}_p(g \circ' f)(w)(x) &= \int_Z w d(g \circ' f)(x) \\ &= \int_{y \in Y} \left(\int_Z w dg(y) \right) df(x) \quad (1) \\ &= \int_{y \in Y} \mathbb{T}_p(g)(w) df(x) \\ &= \mathbb{T}_p(f)\mathbb{T}_p(g)(w)(x) \end{aligned}$$

where (1) is an application of ([7], Theorem 1, d)). □

The relationship between AMKs and MOs is summed up in Fig. 2. Notice that \mathbf{AMK}_p and \mathbf{MO}_p are subcategories of \mathbf{Lp} which are not full.

5.2 Bayesian inversion in \mathbf{Krn}

Recall that Theorem 2 gives Bayesian inversion as a bijection

$$-\dagger : \mathbf{Krn}(X, \mu; Y, \nu)/\mu \cong \mathbf{Krn}(Y, \nu; X, \mu)/\nu.$$

\mathbb{T}_p is \sim -invariant, which allows us to apply it to \sim -equivalence classes of arrows.

Lemma 5. *Let $f, f' : (X, \mu) \rightarrow (Y, \nu)$ be such that $f \sim f'$. Then for all $p \in [1, \infty]$, $\mathbb{T}_p(f) = \mathbb{T}_p(f')$.*

Proof. Since $\mu(\{x \mid f(x) \neq f'(x)\}) = 0$, we have for all function $g : \mathbf{G}(Y) \rightarrow [0, \infty]$ that $\mu(\{x \mid g \circ f(x) \neq g \circ f'(x)\}) = 0$. Taking $g = ev_v(\lambda) = \int_Y v \, d\lambda$, the sought property follows. \square

The following theorem states that pointful Bayesian inversion implements adjunction.

Theorem 6. *For all \mathbf{Krn} arrow $f : (X, \mu) \rightarrow (Y, \nu)$ and all Hölder conjugate (p, q) , $\mathbb{T}_p(f^\dagger) = \mathbb{T}_q(f)^*$.*

Proof. It is enough to prove that for all $u \in L_p^+(X, \mu)$, $v \in L_q^+(Y, \nu)$, we have $\langle \mathbb{T}_p(f^\dagger)(u), v \rangle_Y = \langle u, \mathbb{T}_q(f)(v) \rangle_X$. We compute:

$$\begin{aligned}
\langle \mathbb{T}_p(f^\dagger)(u), v \rangle_Y &= \int_{y \in Y} v(y) \int_{x \in X} u(x) \, df^\dagger(y) \, d\nu \\
&= \int_{y \in Y} \int_{(x, -) \in X \times Y} u(x)v(y) \, d\pi_Y^\dagger(y) \, d\nu \quad (*) \\
&= \int_{(x, y) \in X \times Y} u(x)v(y) \, dI_{X, \mu}^{Y, \nu}(f) \\
&= \int_{x \in X} \int_{(-, y) \in X \times Y} u(x)v(y) \, d\pi_X^\dagger(x) \, d\mu \\
&= \int_{x \in X} u(x) \int_{y \in Y} v(y) \, df(x) \, d\mu \quad (*) \\
&= \langle u, \mathbb{T}_q(f)(v) \rangle_X
\end{aligned}$$

This string of equations follows from the definition of $-^\dagger$ (Theorem 2). At the equations marked $(*)$ we used the characteristic property of disintegrations to move u (resp. v) in (resp. out of) the integral (see Theorem 1). \square

This proves that Bayesian inversion is really just adjunction. However, performing Bayesian inversion in \mathbf{Krn} relies on standard Borel assumptions, while adjunction does not! Most importantly, Bayesian inversion in $\omega\mathbf{CC}$ is better structured, as it corresponds to a categorical duality.

6 Pointless Bayesian inversion through couplings

Under standard Borel assumptions, Bayesian models can be given equivalently either in terms of \mathbf{Krn} arrows or more classically in terms of *joint probabilities* (i.e. couplings). The latter appears crucially in the definition of the pointful inverse, as demonstrated in Theorem 2. However, pointless Bayesian inversion seems *prima facie* to do away with these objects entirely. We conclude this work by shedding some light on the status of couplings w.r.t. pointless inversion:

we show that the bijection $I_{X,\mu}^{Y,\nu} : \mathbf{Krn}(X, \mu; Y, \nu) / \mu \rightarrow \mathbf{\Gamma}(X, \mu; Y, \nu)$ defined in Sec. 3.3 generalises, for X, Y arbitrary measurable spaces, to a bijection from the Homset $\mathbf{MO}_\infty(L_\infty^+(X, \mu); L_\infty^+(Y, \nu))$ to the set of couplings $\mathbf{\Gamma}(X, \mu; Y, \nu)$. In this setting, we prove that Bayesian inversion amounts to permuting the coordinates of the coupling. Our first ingredient is a map from couplings to ω -continuous linear operators. The key observation is the following.

Proposition 4. *For all $p \in [1, \infty]$, any coupling $\gamma \in \mathbf{\Gamma}(X, \mu; Y, \nu)$ induces an ω -continuous linear operator $\mathbf{K}_p(\gamma) \in \mathbf{AMK}_p(X, \mu; Y, \nu)$ defined for $u \in L_p^+(X, \mu)$ and $v \in L_q^+(Y, \nu)$ (using $\varepsilon_p : L_p^+(Y, \nu) \cong L_q^{+,*}(Y, \nu)$ for (p, q) Hölder conjugate) as $\mathbf{K}_p(\gamma)(u)(v) = \int_{(x,y) \in X \times Y} u(x)v(y) d\gamma$.*

Proof. Linearity is trivial. Let us prove that the integral converges. Any function $u \in L_p^+(X, \mu)$ extends to a function $\hat{u} \in L_p^+(X \times Y, \gamma)$ defined as $\hat{u}(x, y) = u(y)$. Indeed, one trivially has $\int_{X \times Y} \hat{u}^p d\gamma = \int_Y u^p d\nu$ by inserting the relevant projection and applying a change of variables. The operation $\hat{\cdot} : L_p^+(X, \mu) \rightarrow L_p^+(X \times Y, \gamma)$ is easily seen to be linear, ω -continuous and norm-preserving (and similarly from $L_q^+(Y, \nu)$ to $L_q^+(X \times Y, \gamma)$), since its action is only to precompose with a projection. The case $p = \infty$ is treated similarly. Therefore, we have the equation $\mathbf{K}_p(\gamma)(u) = \langle \hat{u}, \hat{\cdot} \rangle_{X \times Y} = \varepsilon_p(\hat{u})(\hat{\cdot})$. This proves that $\mathbf{K}_p(\gamma)$ is linear and ω -continuous (hence continuous).

Observe that $\mathbf{K}_p(\gamma)(\mathbb{1}_X) \in L_q^{+,*}(Y, \nu)$ verifies $\mathbf{K}_p(\gamma)(\mathbb{1}_X)(v) = \int_Y v d\nu$. Clearly, the functional $v \mapsto \int_Y v d\nu$ corresponds through the ε_p isomorphism to the element $\mathbb{1}_Y \in L_p^+(Y, \nu)$. For all $u \in L_p^+(X, \mu)$, we have

$$\begin{aligned} \|\mathbf{K}_p(\gamma)(u)\|_p &= \sup \left\{ \|\mathbf{K}_p(\gamma)(u)(v)\| \mid v \in L_q^+(Y, \nu), \|v\|_q = 1 \right\} \\ &= \sup \left\{ \int_{X \times Y} u(x)v(y) d\gamma \mid v \in L_q^+(Y, \nu), \|v\|_q = 1 \right\}. \end{aligned}$$

But by Hölder's inequality, $\int_{X \times Y} u(x)v(y) d\gamma \leq \|u\|_p \|v\|_q = \|u\|_p$ therefore, $\|\mathbf{K}_p(\gamma)\| \leq 1$. Taking $u = \mathbb{1}_X$, we conclude that $\|\mathbf{K}_p(\gamma)\| = 1$. Therefore, $\mathbf{K}_p(\gamma) \in \mathbf{AMK}_p(X, \mu; Y, \nu)$. \square

Dually to Prop. 4, any MO gives rise to a probability measure (but not necessarily a coupling!). For $A : L_p^+(X, \mu) \rightarrow L_p^+(Y, \nu)$ and $B_X \times B_Y$ a basic measurable rectangle in $X \times Y$, we define:

$$\mathbf{C}_p(A)(B_X \times B_Y) = \int_Y \mathbb{1}_{B_Y} A(\mathbb{1}_{B_X}) d\nu. \quad (5)$$

Lemma 6. *For all MO $A : L_p^+(X, \mu) \rightarrow L_p^+(Y, \nu)$, $\mathbf{C}_p(A) \in \mathbf{G}(X \times Y)$.*

Proof. Let us first prove that $\mathbf{C}_p(A)$ is a finite measure. Clearly, $\mathbf{C}_p(A)(\emptyset \times \emptyset) = 0$. Since A is a Markov operator, $\mathbf{C}_p(A)(X \times Y) = 1$. Finite additivity of $\mathbf{C}_p(A)$ is a consequence of linearity. Therefore, $\mathbf{C}_p(A)$ is a finitely additive measure on the algebra generated by basic measurable rectangles. Note that rectangles form a *semialgebra* in the sense of ([3], Def. 1.2.13). ω -continuity of A and of

the Lebesgue integral implies that $C_p(A)$ is σ -additivity on this semialgebra. By ([3], Prop. 1.3.10), this implies σ -additivity of $C_p(A)$ on the algebra generated by the rectangles. Then, the Carathéodory extension theorem [14] implies the existence of a unique (by finiteness) extension of the function defined in Eq. 5 to a probability measure on $X \times Y$. \square

It is not obvious what a necessary and sufficient condition should be for $C_p(A)$ to give rise to a coupling. However, we have the following reasonable sufficient condition.

Proposition 5. *For all MO $A : L_\infty^+(X, \mu) \rightarrow L_\infty^+(Y, \nu)$, $C_\infty(A) \in \mathbf{\Gamma}(X, \mu; Y, \nu)$.*

Proof. Let us prove that $C_\infty(A)$, which is an element of $\mathbf{G}(Y \times X)$ by Lemma 6, has the right marginals. Projecting on X , we have

$$\begin{aligned} C_\infty(A)(B_X \times Y) &= \int_Y \mathbb{1}_Y A(\mathbb{1}_{B_X}) d\nu \\ &= \int_X A^*(\mathbb{1}_Y) \mathbb{1}_{B_X} d\mu \quad (*_1) \\ &= \int_X \mathbb{1}_{B_X} d\mu = \mu(B_X), \quad (*_2) \end{aligned}$$

where we used the adjunction theorem (Theorem 4) at $(*_1)$ and the fact that the adjoint of an MO is an AMK (Prop. 3) at $(*_2)$. Projecting on Y , we get using that A^* is an **AMK**₁ arrow:

$$\begin{aligned} C_\infty(A)(X \times B_Y) &= \int_Y \mathbb{1}_{B_Y} A(\mathbb{1}_X) d\nu \\ &= \int_X A^*(\mathbb{1}_{B_Y}) d\mu = \|A^*(\mathbb{1}_{B_Y})\|_1 \leq \|\mathbb{1}_{B_Y}\|_1 = \nu(B_Y). \end{aligned}$$

Performing the same computation with $B_Y^c = Y \setminus B_Y$ instead of B_Y yields that $C_\infty(A)(B_Y^c \times X) \leq \nu(B_Y^c)$. Using these two inequations together with $\nu(B_Y) + \nu(B_Y^c) = 1$ and the fact that $C_\infty(A)$ is a probability measure allows to conclude that $C_\infty(A)(B_Y) = \nu(B_Y)$. \square

C and **K** are the counterparts of respectively *I* and *D* in Sec. 3.3, with kernels replaced by respectively MOs and AMKs. However, no quotient is needed to obtain the following result, which states that pointless Bayesian inversion (i.e. adjunction) coincides in the world of couplings to the operation which permutes the coordinates (namely the isomorphism $\mathbf{G}(\sigma) : \mathbf{G}(X \times Y) \rightarrow \mathbf{G}(Y \times X)$).

Proposition 6. *For all MO $A : L_\infty^+(X, \mu) \rightarrow L_\infty^+(Y, \nu)$, $A^* = \mathbf{K}_1 \circ \mathbf{G}(\sigma) \circ C_\infty(A)$.*

Proof. It is enough to check the adjointness relation. A monotone convergence argument shows that for all $v \in L_1^+(Y, \nu)$ and $u \in L_\infty^+(X, \mu)$, $\int_{X \times Y} v u dC_\infty(A) = \int_Y v A(u) d\nu$. Therefore,

$$\begin{aligned} \langle \mathbf{K}_1 \circ \mathbf{G}(\sigma) \circ C_\infty(A)(v), u \rangle_X &= \int_{(x,y) \in X \times Y} u(x)v(y) d(\mathbf{G}(\sigma) \circ C_\infty(A)) \\ &= \int_Y v A(u) d\nu = \langle v, A(u) \rangle_Y. \end{aligned}$$

□

In order to close the circle, we prove that couplings are indeed in bijections with \mathbf{MO}_∞ arrows (and hence, by duality, \mathbf{AMK}_1 arrows).

Theorem 7. *For all X, Y measurable and $\mu \in \mathbf{G}(X), \nu \in \mathbf{G}(Y)$, \mathbf{C}_∞ defines a bijection $\mathbf{MO}_\infty(L_\infty^+(X, \mu); L_\infty^+(Y, \nu)) \cong \mathbf{\Gamma}(X, \mu; Y, \nu)$.*

Proof. The fact that $-^*$ is an equivalence of categories implies that $\mathbf{K}_1 \circ \mathbf{G}(\sigma) \circ \mathbf{C}_\infty$ is bijective as a map of Homsets. Hence, \mathbf{C}_∞ must be injective. Let us prove that it is surjective. For all $\gamma \in \mathbf{\Gamma}(X, \mu; Y, \nu)$, we have:

$$\begin{aligned} \mathbf{C}_\infty((\mathbf{K}_1(\mathbf{G}(\sigma)(\gamma)))^*)(B_X \times B_Y) &= \int_{X \times Y} \mathbb{1}_{B_Y} \cdot (\mathbf{K}_1(\mathbf{G}(\sigma)(\gamma)))^*(\mathbb{1}_{B_X}) d\nu \\ &= \int_{X \times Y} \mathbf{K}_1(\mathbf{G}(\sigma)(\gamma))(\mathbb{1}_{B_Y}) \cdot \mathbb{1}_{B_X} d\mu \\ &= \mathbf{G}(\sigma)(\gamma)(B_Y \times B_X) = \gamma(B_X \times B_Y) \end{aligned}$$

which concludes the proof. □

The correspondence between adjunction and the permutation of coupling coordinates together with this last result show that couplings are really at the heart of the semantics of Bayesian inversion.

7 Conclusion

Pointless Bayesian inversion has several qualities which its pointful counterpart lacks: it does not rely on Polish assumptions on the underlying space, it is better typed (as it boils down to an equivalence of categories between abstract Markov kernels and Markov operators) and it admits a trivial and elegant computational interpretation in terms of couplings (as well as the structure of a self-duality on the category of couplings sketched above).

This pointless categorical approach to Bayesian inversion opens the way for exciting new research. First, one yearns to reinterpret previous constructions performed in a kernel-centric way in this new light, such as [12]. Also, the connection between our categories of operators and couplings hints at connections with the Kantorovich distance [13]. For instance, one could study issues of convergence of learning using the weak topology on the space of couplings, which suggests possibly fruitful connections with information geometry.

But chiefly, our more structured framework allows one to reason on the interactions between the approximation of Markov processes by averaging [4] and Bayesian inversion. For instance, we can now ask whether some properties of the Bayesian learning procedure are *profinite*, i.e. entirely characterised by considering the finite approximants (one thinks of issues of convergence of learning, for instance). More generally, we posit that pointless inversion is the right tool to investigate *approximate learning*.

References

1. Nathanael Leedom Ackerman, Cameron E. Freer, and Daniel M. Roy. Noncomputable conditional distributions. In *Proceedings of the 26th Annual IEEE Symposium on Logic in Computer Science, LICS 2011, June 21-24, 2011, Toronto, Ontario, Canada*, pages 107–116, 2011.
2. Charalambos Aliprantis and Kim Border. *Infinite dimensional analysis*, volume 32006. Springer, 1999.
3. Vladimir I Bogachev. *Measure Theory I*. Springer, 2006.
4. Philippe Chaput, Vincent Danos, Prakash Panangaden, and Gordon Plotkin. Approximating Markov Processes by averaging. *Journal of the ACM*, 61(1), January 2014. 45 pages.
5. Jared Culbertson and Kirk Sturtz. A categorical foundation for Bayesian probability. *Applied Categorical Structures*, pages 1–16, 2012.
6. Fredrik Dahlqvist, Vincent Danos, Ilias Garnier, and Ohad Kammar. Bayesian Inversion by Omega-Complete Cone Duality (Invited Paper). In Josée Desharnais and Radha Jagadeesan, editors, *27th International Conference on Concurrency Theory (CONCUR 2016)*, volume 59 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 1:1–1:15, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
7. M. Giry. A categorical approach to probability theory. In *Categorical Aspects of Topology and Analysis*, number 915 in *Lecture Notes In Math.*, pages 68–85. Springer-Verlag, 1981.
8. Olav Kallenberg. *Foundations of Modern Probability*. Springer, 1997.
9. A. S. Kechris. *Classical descriptive set theory*, volume 156 of *Graduate Text in Mathematics*. Springer, 1995.
10. Dexter Kozen. A probabilistic pdl. In *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing, STOC '83*, pages 291–297, New York, NY, USA, 1983. ACM.
11. Peter Selinger. Towards a semantics for higher-order quantum computation. In *Proceedings of the 2nd International Workshop on Quantum Programming Languages, TUCS General Publication*, volume 33, pages 127–143, 2004.
12. Sam Staton, Hongseok Yang, Chris Heunen, Ohad Kammar, and Frank Wood. Semantics for probabilistic programming: higher-order functions, continuous distributions, and soft constraints. *CoRR*, abs/1601.04943, 2016.
13. Cedric Villani. *Optimal transport, Old and New*. Grundlehren der mathematischen Wissenschaften. Springer, 2006.
14. David Williams. *Probability with Martingales*. Cambridge University Press, 1991.