CORE
Metadata, citation and similar papers at core.ac.uk
Provided by Edinburgh Research Explorer

THE UNIVERSITY of EDINBURGH

Edinburgh Research Explorer

Borel Kernels and their Approximation, Categorically

Citation for published version:
Dahlqvist, F, Silva, A, Danos, V & Garnier, I 2018, 'Borel Kernels and their Approximation, Categorically' Electronic Notes in Theoretical Computer Science, vol. 341, pp. 91-119. DOI: 10.1016/j.entcs.2018.11.006

Digital Object Identifier (DOI):
10.1016/j.entcs.2018.11.006

Link:
Link to publication record in Edinburgh Research Explorer

Document Version:
Publisher's PDF, also known as Version of record

Published In:
Electronic Notes in Theoretical Computer Science

General rights
Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy
The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

OPEN ACCESS

Download date: 05. Apr. 2019

Available online at www.sciencedirect.com

ScienceDirect

Electronic Notes in Theoretical Computer Science 341 (2018) 91–119

Electronic Notes in Theoretical Computer Science

www.elsevier.com/locate/entcs

Borel Kernels and their Approximation, Categorically

Fredrik Dahlqvist[1,2] Alexandra Silva[2]

Department of Computer Science
University College London

Vincent Danos

CNRS, ENS-PSL, INRIA

Ilias Garnier[3]

ENS-PSL

Abstract

This paper introduces a categorical framework to study the exact and approximate semantics of probabilistic programs. We construct a dagger symmetric monoidal category of Borel kernels where the dagger-structure is given by Bayesian inversion. We show functorial bridges between this category and categories of Banach lattices which formalize the move from kernel-based semantics to predicate transformer (backward) or state transformer (forward) semantics. These bridges are related by natural transformations, and we show in particular that the Radon-Nikodym and Riesz representation theorems - two pillars of probability theory - define natural transformations.
With the mathematical infrastructure in place, we present a generic and endogenous approach to approximating kernels on standard Borel spaces which exploits the involutive structure of our category of kernels. The approximation can be formulated in several equivalent ways by using the functorial bridges and natural transformations described above. Finally, we show that for sensible discretization schemes, every Borel kernel can be approximated by kernels on finite spaces, and that these approximations converge for a natural choice of topology.
We illustrate the theory by showing that our approximation scheme can be used in practice as an approximate Bayesian inference algorithm and as an approximation scheme for programs in the probabilistic network specification language ProbNetKAT.

Keywords: Probabilistic programming, probabilistic semantics, Markov process, Bayesian inference, approximation

1 Introduction

Finding a good category in which to study probabilistic programs is a subject of active research [22,17,6,21]. In this paper we present a dagger symmetric monoidal

[1] Email: f.dahlqvist@ucl.ac.uk

[2] This work was partially supported by ERC grant ProfoundNet.

[3] This work was partially supported by ANR project REPAS.

https://doi.org/10.1016/j.entcs.2018.11.006

1571-0661/© 2018 The Author(s). Published by Elsevier B.V.
This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

category of kernels whose dagger-structure is given by Bayesian inversion. The advantages of this new category are two-fold.

Firstly, the most important new construct introduced by probabilistic programming, *viz. Bayesian inversion*, is interpreted completely straightforwardly by the †-operation which is native to our category. In particular we never leave the world of kernels and we therefore do not require any normalization construct. Consider for example the following simple Bayesian inference problem in Anglican ([24])

```
(defquery example
  (let [x (sample (normal 0 1))]
    (observe (normal x 1) 0.5)
    (> x 1)))
```

The semantics of this program is built easily and compositionally (by functoriality of †) in our category:

- The second line builds a Borel space equipped with a normally distributed probability measure – an object $(\mathbb{R}, \mu)$ of our category.
- The `(normal x 1)` instruction builds a Borel kernel – a morphism $f : (\mathbb{R}, \mu) \to (\mathbb{R}, \nu)$ in our category.
- The `observe` statement builds the Bayesian inverse of the kernel – the morphism $f^\dagger : (\mathbb{R}, \nu) \to (\mathbb{R}, \mu)$ in our †-category.
- Finally, the kernel $f^\dagger$ is evaluated, i.e. the denotation of the program above is $f^\dagger(0.5)(]1, \infty[)$. Since $f^\dagger$ is only defined μ-a.s., we understand the formal expression $f^\dagger(0.5)(]1, \infty[)$ as the evaluation of *some* representative of $f^\dagger$ at (0.5) and $(]1, \infty[)$. The choice of representative is μ-a.s. irrelevant, which justifies the notation. Note also that Anglican cleanly separates the modelling process from the inference process. The snippet of code above belongs to the modelling level, whereas the choice of a representative of $f^\dagger$ – that is to say of an inference algorithm – belongs to the inference level, which is specified separately. We only claim to provide a semantics at the modelling level.

Secondly, since Bayesian inference problems are in general very hard to compute (although the one given above has an analytical solution), it makes sense to seek approximate solutions, i.e. *approximate denotations* to probabilistic programs. As we will show, our category of kernels comes equipped with a generic and endogenous approximating scheme which relies on its involutive structure and on the structure of standard Borel spaces. Moreover, this approximation scheme converges for any choice of kernel for a natural choice of topology.

Main contributions and structure of the paper.

Sections 2-6 align in a chain of logical dependencies culminating in our Convergence of Approximations Theorem (Theorem 6.3) which relies on the framework developed in §2, the link with operators developed in §3-4 which allows us to access a natural operator topology, and the presentation of our approximation scheme in §5. More precisely, our main contributions can be summed up section by section as

follows:

- We build a category **Krn** of Borel kernels (§2) and we show how two kernels which agree almost everywhere can be identified under a categorical quotient operation. This technical construction is what allows us to define *Bayesian inversion* as an involutive functor, denoted †. This is a key technical improvement on [6] where the †-structure [4] was hinted at but was not functorial. We show that **Krn** is a dagger symmetric monoidal category.
- We introduce the category $\mathbf{BL}_\sigma$ of Banach lattices and σ-order continuous positive operators as well as the Köthe dual functor $(-)^\sigma : \mathbf{BL}_\sigma^{\mathrm{op}} \to \mathbf{BL}_\sigma$ (§3). These will play a central role in studying convergence of our approximation schemes.
- We provide the first [5] categorical understanding of the Radon-Nikodym and the Riesz representation theorems. These arise as natural transformations between functors relating kernels and Banach lattices (§4).
- We show how the †-structure of **Krn** can be exploited to approximate kernels by averaging (§5). Due to an important structural feature of **Krn** (Th. 2.1) every kernel in **Krn** can be approximated by *finite* kernels.
- We describe a natural class of approximations schemes where the sequence of approximating kernels converges to the kernel to be approximated. The notion of convergence is given naturally by moving to $\mathbf{BL}_\sigma$ and considering convergence in the Strong Operator Topology (§6).
- We show how Bayesian inference can be performed approximately by showing that the †-operation commutes with taking approximations. This provides the basis for an approximate Bayesian inference algorithm which we apply to a simple Anglican program. A second application, as an approximation scheme for the denotation of ProbNetKAT programs [13,20] is given in the Appendix.

All the proofs can be found in the Appendix.

Related work.

Quasi-Borel sets have recently been proposed as a semantic framework for higher-order probabilistic programs in [22]. The main differences with our approach are: (i) unlike [22,21] we never leave the realm of kernels, and in particular we never need to worry about normalization. This makes the interpretation of `observe` statements, i.e. of Bayesian inversion, simpler and more natural. However, (ii) unlike the quasi-Borel sets of [22], our category is not Cartesian closed. We can therefore not give a semantics to all higher-order programs. This shortcoming is partly mitigated by the fact that the category of Polish spaces, on which our category ultimately rests, does have access to many function spaces, in particular all the spaces of functions whose domain is locally compact. We can thus in principle provide a semantics to higher-order programs, provided that λ-abstraction is restricted to locally compact

[4] Suggested to us by Chris Heunen.
[5] To the best of our knowledge.

spaces like the reals and the integers, although this won't be investigated in this paper.

The approximation of probabilistic kernels has been a topic of investigation in theoretical computer science for nearly twenty years (see e.g. [10,8,9,4]), and for much longer in the mathematical literature (e.g. [5]). Our results build on the formalism developed in [4] with the following differences: (i) we can approximate kernels, their associated *stochastic operator* (backward predicate transformer), or their associated *Markov operator* (forward state transformer) with equivalent ease, and move freely across the three formalisms. (ii) Given a kernel $f : X \rightarrowtriangle Y$, we can define its approximation $f' : X' \to Y'$ along any quotients X' of X and Y' of Y as in [4], but we can also 'internalize' the approximation as a kernel $f^* : X \to Y$ of the original type. Morally f' and f^* are the same approximation, but the second approximant, being of the same type as the original kernel, can be compared with it. In particular it becomes possible to study the convergence of ever finer approximations, which we do in Section 6. Finally, (iii) we opt to work with Banach lattices rather than the normed cones of [18,4] because it allows us to formulate the operator side of the theory very naturally, and it connects to a large body of classic mathematical results ([2,25]) which have been used in the semantics of probabilistic programs as far back as Kozen's seminal [16].

2 A category of Borel kernels

In [6] the first three authors presented a category of Borel kernels similar in spirit to the construction of this section, but with a major shortcoming. As we will shortly see, our category **Krn** of Borel kernels can be equipped with an involutive functor – a *dagger* operation † in the terminology of [19] – which captures the notion of *Bayesian inversion* and is absolutely crucial to everything that follows. In [6] this operation had merely been identified as a map, i.e. not even as a functor. In this section we show that Bayesian inversion defines a †-structure on a more sophisticated – but measure-theoretically very natural – category of kernels.

2.1 Standard Borel spaces and the Giry monad

A *standard Borel space* – or *SB space* for short – is a measurable space $(X, \mathcal{S})$ for which there exists a Polish topology $\mathcal{T}$ on X whose Borel sets are the elements of $\mathcal{S}$, i.e. such that $\mathcal{S} = \sigma(\mathcal{T})$ (see e.g. [15] for an overview). Let us write **SB** for the category of standard Borel spaces and measurable maps. One key structural feature of **SB** is the following:

Theorem 2.1 *Every* **SB** *object is a limit of a countable co-directed diagram of finite spaces.*

Proof. This is a consequence of the Isomorphism Theorem (Theorem 15.6 of [15]): two SB spaces are isomorphic iff they have the same cardinality. Uncountable SB spaces are thus all isomorphic to the Cantor space $2^{\mathbb{N}}$ which is the limit of the countable co-directed diagram $(2^n)_{n\in\mathbb{N}}$ with the connecting morphisms $p_{n+1,n}$:

$2^{n+1} \to 2^n$ truncating binary words of length $n+1$ at length n. Similarly all SB-spaces of cardinality $\aleph_0$ are isomorphic to the one-point compactification of $\mathbb{N}$, which is the limit of the countable co-directed diagram $(n)_{n\in\mathbb{N}}$ with the connecting morphisms $p_{n+1,n} : n+1 \to n, i \mapsto \min(i,n)$. The case of finite SB spaces is trivial.□

The *Giry monad* was originally defined in two variants [14]:

- As an endofunctor $\mathsf{G}_{\mathbf{Pol}}$ of **Pol**, the category of *Polish spaces*, one sets $\mathsf{G}_{\mathbf{Pol}}(X,\mathcal{T})$ to be the space of Borel probability measures over X together with the weak topology. This space is Polish [15, Th 17.23], and the Portmanteau Theorem [15, Th 17.20]) gives multiple characterizations of the weak topology.
- As an endofunctor $\mathsf{G}_{\mathbf{Meas}}$ of **Meas**, the category of *measurable spaces*: $\mathsf{G}_{\mathbf{Meas}}(X,\mathcal{S})$ is the set of probability measures on X together with the initial σ-algebra for the maps $ev_A : \mathsf{G}_{\mathbf{Meas}}(X,\mathcal{S}) \to \mathbb{R}, \mu \mapsto \mu(A), A \in \mathcal{S}$.

In both cases the Giry monad is defined on an arrow $f : X \to Y$ as the map f_* which sends a measure μ on X to the *pushforward measure* $f_*\mu$ on Y, defined as $G(f)(\mu)(B) = f_*\mu(B) := \mu(f^{-1}(B))$ for B a measurable subset of Y. We want to define the Giry monad on the category **SB** of standard Borel spaces (and measurable maps), and the two versions of the Giry monad described above offer us natural ways to do this: given an SB space $(X,\sigma(\mathcal{T}))$ we can either compute $\mathsf{G}_{\mathbf{Pol}}(X,\mathcal{T})$ and take the associated standard Borel space, or directly compute $\mathsf{G}_{\mathbf{Meas}}(X,\sigma(\mathcal{T}))$. Fortunately, the two methods agree.

Theorem 2.2 ([15], Th 17.24) *Let* $\mathsf{B} : \mathbf{Pol} \to \mathbf{SB}$ *denote the functor sending a Polish space* $(X,\mathcal{T})$ *to its associated SB-space* $(X,\sigma(\mathcal{T}))$ *and leaving morphisms unchanged, then*

$$\mathsf{G}_{\mathbf{Meas}} \circ \mathsf{B} = \mathsf{B} \circ \mathsf{G}_{\mathbf{Pol}}.$$

We define the *Giry monad* on SB spaces to be the endofunctor $\mathsf{G} : \mathbf{SB} \to \mathbf{SB}$ defined by either of the two equivalent constructions above. The monadic data of G is given at each SB space X by the unit $\delta_X : X \to \mathsf{G}X, x \mapsto \delta_x$, the Dirac δ measure at x, and the multiplication $m_X : \mathsf{G}^2X \to \mathsf{G}X, \mathbb{P} \mapsto \lambda A. \int_{\mathsf{G}X} ev_A d\mathbb{P}$. We refer the reader to [14] for proofs that δ_X and m_X are measurable.

2.2 The construction of **Krn**

Let us denote by $\mathbf{SB}_{\mathsf{G}}$ the Kleisli category associated with the Giry monad (G,δ,m). We denote Kleisli arrows, i.e. Markov kernels, by $X \rightarrowtriangle Y$, and we call such an arrow *deterministic* if it can be factorized as an ordinary measurable function followed by the unit δ. Kleisli composition is denoted by $\bullet$. The category $* \downarrow \mathbf{SB}_{\mathsf{G}}$ has arrows $* \rightarrowtriangle X$ as objects, where $*$ is the one point SB space (the terminal object in **SB**). An arrow from $\mu : * \rightarrowtriangle X$ to $\nu : * \rightarrowtriangle Y$ is a $\mathbf{SB}_{\mathsf{G}}$ arrow $f : X \rightarrowtriangle Y$ such that $\nu = f \bullet \mu$, i.e. such that $\nu(A) = \int_X f(x)(A)d\mu$ for any measurable subset A of Y. This situation will be denoted in short by $f : (X,\mu) \rightarrowtriangle (Y,\nu)$, and we will call a pair (X,μ) a *measured SB space*.

We construct a quotient of $* \downarrow \mathbf{SB}_{\mathsf{G}}$, such that two $* \downarrow \mathbf{SB}_{\mathsf{G}}$ arrows are identified

if they disagree on a null set w.r.t. the measure on their domain. For $g, g' : (X, \mu) \rightarrowtriangle (Y, \nu)$, we define $N(g, g') = \{x \in X \mid g(x) \neq g'(x)\}$.

Lemma 2.3 *$N(g, g')$ is a measurable set.*

Proof. By Dinkyn's π-λ theorem, two finite measures are equal if and only if they agree on a π-system generating the σ-algebra. Any standard Borel space admits such a countable π-system (any countable basis for a Polish topology generating the σ-algebra). Let $\{B_n\}_{n\in\mathbb{N}}$ be such a π-system. Then, for all $x \in X$, $g(x) \neq g'(x) \Leftrightarrow \exists n. g(x)(B_n) \neq g'(x)(B_n)$. Hence,

$$\begin{aligned} N(g, g') &= \cup_n \{x \in X \mid g(x)(B_n) \neq g'(x)(B_n)\} \\ &= \cup_n \{x \in X \mid ev_{B_n}(g(x)) \neq ev_{B_n}(g'(x))\} \\ &= \cup_n (ev_{B_n} \circ g - ev_{B_n} \circ g')^{-1}(\mathbb{R} \setminus \{0\}) \end{aligned}$$

By definition of the measurable structure of $\mathsf{G}(Y)$, $ev_{B_n} \circ g - ev_{B_n} \circ g'$ is measurable, hence $N(g, g')$ is also measurable. $\square$

We now define a relation $\sim$ on $\mathrm{Hom}((X, \mu), (Y, \nu))$ by saying that for any two arrows $g, g' : (X, \mu) \rightarrowtriangle (Y, \nu)$, $g \sim g'$ if $\mu(N(g, g')) = 0$. This clearly defines an equivalence relation on $\mathrm{Hom}((X, \mu), (Y, \nu))$. In order to perform the quotient of the category $* \downarrow \mathbf{SB}_{\mathsf{G}}$ modulo $\sim$, we need to check that it is compatible with composition.

Proposition 2.4 *If $g \sim g'$, then $h \bullet g \bullet f \sim h \bullet g' \bullet f$.*

Proof. We first show that if $g \sim g'$, then $h \bullet g \sim h \bullet g'$. Clearly, for any space V and any deterministic function $u : Y \to V$, $N(u \circ g, u \circ g') \subseteq N(g, g')$. By definition of the Kleisli category, $h \bullet g = m_Z \circ \mathsf{G}(h) \circ g$ and similarly for $h \bullet g'$. Taking $u = m_Z \circ \mathsf{G}(h)$, we obtain that $\mu(N(h \bullet g, h \bullet g')) \leq \mu(N(g, g'))$.

It is now enough to show that $\lambda(N(g \bullet f, g' \bullet f)) = 0$. Let us reason contraposi-

tively. We have:

$$
\begin{array}{lll}
 & \lambda(N(g \bullet f, g' \bullet f)) & > 0 \\
\Leftrightarrow & \int_{w \in W} 1_{N(g \bullet f, g' \bullet f)}(w)\, d\lambda & > 0 \\
\rightarrow & \int_{w \in W} \sum_{n \in \mathbb{N}} 1_{(g \bullet f)(w)(B_n) \neq (g' \bullet f)(w)(B_n)}\, d\lambda & > 0 \\
\overset{\exists n}{\rightarrow} & \int_{w \in W} 1_{(g \bullet f)(w)(B_n) \neq (g' \bullet f)(w)(B_n)}\, d\lambda & > 0 \\
\rightarrow & \int_{w \in W} |(g \bullet f)(w)(B_n) - (g' \bullet f)(w)(B_n)|\, d\lambda & > 0 \\
\Leftrightarrow & \int_{w \in W} \int_{x \in X} |g(x)(B_n) - g'(x)(B_n)|\, df(w)\, d\lambda & > 0 \\
\Leftrightarrow & \int_{x \in X} |g(x)(B_n) - g'(x)(B_n)|\, d\mu & > 0 \\
\overset{\exists X^+ \subseteq X}{\rightarrow} & \int_{x \in X^+} g(x)(B_n) - g'(x)(B_n)\, d\mu & > 0 \\
\rightarrow & \int_{x \in X^+} 1_{g(-)(B_n) > g'(-)(B_n)}(x)\, d\mu & > 0 \\
\rightarrow & \int_{x \in X^+} 1_{N(g,g')}(x)\, d\mu & > 0
\end{array}
$$

The last line implies $\mu(N(g, g')) > 0$, a contradiction. □

Definition 2.5 Let **Krn** be the category obtained by quotienting $* \downarrow \mathbf{SB}_{\mathsf{G}}$ hom-sets with $\sim$.

The following Theorem is of great practical use and generalizes the well-known result for deterministic arrows.

Theorem 2.6 (Change of Variables in Krn) *Let $f : (X, \mu) \rightarrowtriangle (Y, \nu)$ be a* **Krn**-*morphism. For any measurable function $\phi : Y \to \mathbb{R}$, if ϕ is ν-integrable, then $\phi \bullet f(x) = \int_Y \phi\, df(x)$ is μ-integrable and*

$$\int_Y \phi\ d\nu = \int_X \phi \bullet f\ d\mu$$

Proof. If ϕ is ν-integrable, there exists a monotone sequence $\{\phi_n\}$ of simple functions such that $\phi_n \uparrow \phi$ and $\int_Y \phi_n d\nu \to \int_Y \phi d\nu < \infty$. By definition each $\phi_n = \sum_{i=0}^{k} \alpha_i 1_{B_i}$, and by unravelling the definition we have

$$\int_Y 1_{B_i} d\nu = \nu(B_i) = \int_X f(x)(B_i) d\mu = \int_X \int_Y 1_{B_i} df(x) d\mu = \int_X (1_{B_i} \bullet f) d\mu$$

From which it follows that

$$\int_Y \phi_n d\nu = \int_X \int_Y \sum_{i=0}^{k} \alpha_i 1_{B_i} df(x) d\mu = \int_X (\phi_n \bullet f) d\mu$$

and the result follows from the Monotone Convergence Theorem (MCT). □

2.3 The symmetric monoidal structure of **Krn**

The symmetric monoidal structure of **Krn** is defined on a pair of objects $(X, \mu), (Y, \nu)$ by the Cartesian product and the product of measures, i.e. $(X, \mu) \otimes (Y, \nu) = (X \times Y, \mu \otimes \nu)$. On pairs of morphisms $f : (X, \mu) \rightarrowtriangle (Y, \nu)$ and $f' : (X', \mu') \rightarrowtriangle (Y', \nu')$ it is defined by $(f \otimes f')(x, x') := f(x) \otimes f'(x')$. The unitors, associator and braiding transformations are given by the obvious bijections.

2.4 The dagger structure of **Krn**

Krn has an extremely powerful inversion principle:

Theorem 2.7 (Measure Disintegration Theorem, [15], 17.35) *Let $f : (X, \mu) \rightarrowtriangle (Y, \nu)$ be a* deterministic **Krn**-*morphism, there exists a unique morphism $f^\dagger_\mu : (Y, \nu) \rightarrowtriangle (X, \mu)$ such that*

$$f \bullet f^\dagger_\mu = \mathrm{id}_{(Y,\nu)}. \tag{1}$$

The kernel $f^\dagger_\mu$ is called the *disintegration of μ along f*. As our notation suggests, the disintegration depends fundamentally on the measure μ over the domain, however we will omit this subscript when there is no ambiguity. The following lemma relates disintegrations to conditional expectations.

Lemma 2.8 ([7]) *Let $f : (X, \mu) \to (Y, \nu)$ be a deterministic* **Krn**-*morphism, let $\sigma(f)$ be the σ-algebra generated by f, and let $\phi : X \to \mathbb{R}$ be measurable, then μ-a.e.*

$$\phi \bullet f^\dagger \bullet f = \mathbb{E}\left[\phi \mid \sigma(f)\right]$$

We can extend the definition of $(-)^\dagger$ to *any* **Krn**-morphism $f : (X, \mu) \rightarrowtriangle (Y, \nu)$ in a functorial way, although $f^\dagger$ will not in general be a right inverse to f. The construction of $f^\dagger$ is detailed in [6], but let us briefly recall how it works. The category **SB** has products which are built in the same way as in **Meas** via the product of σ-algebras[6]. Given any kernel $f : (X, \mu) \rightarrowtriangle (Y, \nu)$, we can canonically construct a probability measure γ_f on the product $X \times Y$ of SB-space by defining it on the rectangles of $X \times Y$ as

$$\gamma_f(A \times B) = \int_{x \in X} 1_A(x) \cdot f(x)(B)\ d\mu. \tag{2}$$

Equivalently, $\gamma_f = (\delta_X \otimes f) \bullet \Delta_X \bullet \mu$, where $\Delta_X : X \to X \times X$ is the diagonal map. Letting $\pi_X : X \times Y \to X$ and $\pi_Y : X \times Y \to Y$ be the canonical projections, we observe that $\mathsf{G}\pi_X(\gamma_f) = \mu$ and $\mathsf{G}\pi_Y(\gamma_f) = \nu$: in other words, γ_f is a coupling of μ and ν. The disintegration of γ_f along π_Y is a kernel $\pi^\dagger_Y : (Y, \nu) \to (X \times Y, \gamma_f)$. Finally we define:

$$f^\dagger = \pi_X \bullet \pi^\dagger_Y. \tag{3}$$

[6] Unlike the category **Krn**, which does *not* have products.

The following **Krn** diagram sums up the situation:

$$\begin{tikzcd}
(X,\mu) \arrow[rr, bend left, "f"] \arrow[r, bend left, "\pi_X^\dagger"] & (X \times Y, \gamma_f) \arrow[l, bend left, "\pi_X"] \arrow[r, bend left, "\pi_Y"] & (Y,\nu) \arrow[l, bend left, "\pi_Y^\dagger"] \arrow[ll, bend left, "f^\dagger"]
\end{tikzcd}$$

where $\pi_X^\dagger$ is explicitly given by $(\delta_X \otimes f) \bullet \Delta_X$. The following property characterizes the action of $(-)^\dagger$ on **Krn**-morphisms:

Theorem 2.9 *For all $f : (X,\mu) \rightarrowtriangle (Y,\nu)$, $f^\dagger : (Y,\nu) \rightarrowtriangle (X,\mu)$ is the unique* **Krn** *morphism satisfying for all measurable sets $A \subseteq X$, $B \subseteq Y$ the following equation:*

$$\int_{x \in X} 1_A(x) \cdot f(x)(B)\ d\mu = \int_{y \in Y} f^\dagger(y)(A) \cdot 1_B(y)\ d\nu \tag{4}$$

Proof. It follows by definition of $f^\dagger$ and from the disintegration theorem that

$$\int_{y \in Y} f^\dagger(y)(A) \cdot 1_B(y)\ d\nu = \gamma_f(A \times B), \tag{5}$$

from which Eq. 4 follows easily. It remains to prove that this uniquely characterizes $f^\dagger$. Let us reason contrapositively. Assume there exists $g : (Y,\nu) \rightarrowtriangle (X,\mu)$ verifying for all A, B measurable $\int_{y \in Y} g(y)(A) \cdot 1_B(y) = \gamma_f(A \times B)$ as in Eq. 5 and such that $\nu(N(f^\dagger, g)) > 0$ (assuming we take some representative of $f^\dagger$). Let $\{A_n\}_{n \in \mathbb{N}}$ be a countable π-system generating the σ-algebra of X. It is enough to test equality of measures on X on this π-system. Therefore, $N(f^\dagger, g) = \cup_n \{y \mid f^\dagger(y)(A_n) \neq g(y)(A_n)\}$. Since $\nu(N(f^\dagger, g)) > 0$, there must exist a $k \in \mathbb{N}$ such that $\nu(\{y \mid f^\dagger(y)(A_k) \neq g(y)(A_k)\}) > 0$. Therefore, $N_k^+ = \{y \mid f^\dagger(y)(A_k) > g(y)(A_k)\}$ must also have positive measure for ν. But then, $\int_{y \in Y} g(y)(A_k) \cdot 1_{N_k^+}(y) \neq \gamma_f(A_k \times N_k^+)$, a contradiction. □

In view of Eq. (4), we will call $f^\dagger$ the *Bayesian inversion of f*, and refer to $(-)^\dagger$ as the *Bayesian inversion operation* on **Krn**. It will be crucial throughout the rest of this paper. It is important to see that $f^\dagger$ absolutely depends on the choice of μ and not only on f seen as a function. We can now improve on [6] and show that $(-)^\dagger$ is indeed a †-operation in the strict categorical meaning of the term.

Theorem 2.10 **Krn** *is a dagger symmetric monoidal category, with $(-)^\dagger$ given by Bayesian inversion.*

Proof. Let us first show that $(-)^\dagger$ is a functor $\mathbf{Krn} \to \mathbf{Krn}^{\mathrm{op}}$, i.e. that $\mathrm{id}^\dagger_{(X,\mu)} = \mathrm{id}_{(X,\mu)}$ and that for any $f : (X,\mu) \rightarrowtriangle (Y,\nu)$ and $g : (Y,\nu) \rightarrowtriangle (Z,\rho)$ we have $(g \bullet f)^\dagger = f^\dagger \circ g^\dagger$.

Let (X,μ) be an object of **Krn** and $\mathrm{id}_{X,\mu}$ the corresponding identity. By Th. 2.9, it is enough to prove, for all A, A' measurable subsets of X, that

$$\int_{x \in X} 1_A(x) \cdot id_{X,\mu}(x)(A')\ d\mu = \int_{x \in X} id_{X,\mu}(x)(A) \cdot 1_{A'}(x)\ d\mu.$$

We have:

$$\int_{x\in X} 1_A(x)\cdot id_{X,\mu}(x)(A')\ d\mu = \int_{x\in X} 1_A(x)\cdot 1_{A'}\ d\mu = \mu(A\cap A')$$

The same calculation on the right hand side of the first equation yields trivially the same result. Hence the equality is verified.

Now, on to compatibility w.r.t. composition. In sight of Th. 2.9, it is enough to show that for all $A \subseteq X$, $C \subseteq Z$,

$$\int_{x\in X} (g\bullet f)(x)(C)\cdot 1_A(x)\ d\mu = \int_{z\in Z} 1_C(z)\cdot (f^\dagger \bullet g^\dagger)(z)(A)\ d\rho$$

In the following, for X a measurable space, we denote by $SF(X)$ the set of simple functions over X (finite linear combinations of indicator functions of measurable sets). We will use repeatedly the monotone convergence theorem (MCT). The left hand side of the above equation can be re-written as:

$$\begin{aligned}
&\textstyle\int_{x\in X}\left(\int_{y\in Y} g(y)(C)\ df(x)\right)\cdot 1_A(x)\ d\mu\\
\overset{(1)}{=}\ &\textstyle\int_{x\in X}\left(\int_{y\in Y} \lim_{n\to\infty} g_n(y)\ df(x)\right)\cdot 1_A(x)\ d\mu\\
\overset{(2)}{=}\ &\textstyle\int_{x\in X}\lim_{n\to\infty}\left(\int_{y\in Y} g_n(y)\ df(x)\right)\cdot 1_A(x)\ d\mu
\end{aligned}$$

where (1) is because $g_n \uparrow g(-)(C), g_n \in SF(Y)$ and (2) by monotone convergence. Note that the n-indexed family $x \mapsto \int_{y\in Y} g_n(y)\ df(x)$ is pointwise increasing. Therefore,

$$\begin{aligned}
&\overset{(*)}{=}\ \textstyle\lim_{n\to\infty}\int_{x\in X}\left(\int_{y\in Y} g_n(y)\ df(x)\right)\cdot 1_A(x)\ d\mu\\
&\overset{(1)}{=}\ \textstyle\lim_{n\to\infty}\int_{x\in X}\left(\sum_{i=1}^{k_n}\alpha_i^n f(x)(C_i^n)\right)\cdot 1_A(x)\ d\mu\\
&=\ \textstyle\lim_{n\to\infty}\sum_{i=1}^{k_n}\alpha_i^n\int_{x\in X} f(x)(C_i^n)\cdot 1_A(x)\ d\mu\\
&\overset{(2)}{=}\ \textstyle\lim_{n\to\infty}\sum_{i=1}^{k_n}\alpha_i^n\int_{y\in Y} 1_{C_i^n}(y)\cdot f^\dagger(y)(A)\ d\nu\\
&\overset{(*)}{=}\ \textstyle\int_{y\in Y} g(y)(C)\cdot f^\dagger(y)(A)\ d\nu\\
&\overset{(3)}{=}\ \textstyle\int_{y\in Y} g(y)(C)\cdot \lim_n f_n(y)\ d\nu\\
&\overset{(*)}{=}\ \textstyle\lim_n\sum_{i=1}^{k_n}\beta_i^n\int_{y\in Y} g(y)(C)\cdot 1_{D_i^n}(y)\ d\nu\\
&\overset{(2)}{=}\ \textstyle\lim_n\sum_{i=1}^{k_n}\beta_i^n\int_{z\in Z} 1_C(z)\cdot g^\dagger(z)(D_i^n)\ d\rho\\
&=\ \textstyle\lim_n\int_{z\in Z} 1_C(z)\cdot\int_{y\in Y}\sum_{i=1}^{k_n}\beta_i^n 1_{D_i^n}(y)\ dg^\dagger(z)\ d\rho\\
&\overset{(*)}{=}\ \textstyle\int_{z\in Z} 1_C(z)\cdot\int_{y\in Y} f^\dagger(y)(C)\ dg^\dagger(z)\ d\rho\\
&=\ \textstyle\int_{z\in Z} 1_C(z)\cdot (f^\dagger\bullet g^\dagger)(z)(A)\ d\rho
\end{aligned}$$

where $(*)$ is by monotone convergence, (1) is because $g_n \in SF(Y)$, (2) is by Th. 2.9 and (3) is because $f_n \uparrow f^\dagger(-)(A), f_n \in SF(Y)$. We have proved the sought identity.

Finally let us show that $(-)^\dagger$ is involutive, i.e. that for any $f : (X, \mu) \rightarrowtriangle (Y, \nu)$, $(f^\dagger)^\dagger = f$. This follows easily by two applications of Th. 2.9): we have

$$\begin{aligned}\int_{x\in X} 1_A(x) \cdot (f^\dagger)^\dagger(x)(B)\ d\mu &= \int_{y\in Y} f^\dagger(y)(A) \cdot 1_B(y)\ d\nu \\ &= \int_{x\in X} 1_A(x) \cdot f(x)(B)\ d\mu;\end{aligned}$$

and since adjoints are unique, $f = (f^\dagger)^\dagger$.

The fact that $(f \otimes g)^\dagger = f^\dagger \otimes g^\dagger$ follows immediately from the definitions and the property of disintegrations given by Th. 2.9. The fact that the associator, unitors and braiding transformations are unitary follows immediately from the fact that they are deterministic isomorphisms and Th. 2.7. □

3 Banach lattices

It is well-known that kernels can alternatively be seen as predicate – i.e. real-valued function –transformers, or as state – i.e. probability measure – transformers. The latter perspective was adopted by Kozen in [16] to describe the denotational semantics of probabilistic programs (without conditioning). We shall see in this section and the next, that the predicate and state transformer perspectives are dual to one another in the category of *Banach lattices*, a framework incidentally also used in [16]. For an introduction to the theory of Banach lattices we refer the reader to e.g. [2,25].

An *ordered real vector space* V is a real vector space together with a partial order $\leq$ which is compatible with the linear structure in the sense that for all $u, v, w \in V, \lambda \in \mathbb{R}^+$

$$u \leq v \Rightarrow u + w \leq v + w \qquad \text{and} \qquad u \leq v \Rightarrow \lambda u \leq \lambda v$$

An ordered vector space $(V, \leq)$ is called a *Riesz space* if the poset structure forms a lattice. A vector v in a Riesz space $(V, \leq)$ is called *positive* if $0 \leq v$, and its *absolute value* $|v|$ is defined as $|v| = v \vee (-v)$. A Riesz space $(V, \leq)$ is *σ-order complete* if every non-empty countable subset of V which is order bounded has a supremum.

A *normed Riesz space* is a Riesz space $(V, \leq)$ equipped with a *lattice norm*, i.e. a map $\|\cdot\| : V \to \mathbb{R}$ such that:

$$|v| \leq |w| \ \text{ implies } \ \|v\| \leq \|w\|\,. \tag{6}$$

A normed Riesz space is called a *Banach lattice* if it is (norm-) complete, i.e. if every Cauchy sequence (for the norm $\|\cdot\|$) has a limit in V.

Example 3.1 For each measured space (X, μ) – and in particular **Krn**-objects – and each $1 \leq p \leq \infty$, the space $\mathrm{L}_p(X, \mu)$ is a Riesz space with the pointwise order. When it is equipped with the usual L_p-norm, it is a Banach lattice. This fact is often referred to as the *Riesz-Fischer theorem* (see [2, Th 13.5]). We will say that

$p, q \in \mathbb{N} \cup \{\infty\}$ are *Hölder conjugate* if either of the following conditions hold: (i) $1 < p, q < \infty$ and $\frac{1}{p} + \frac{1}{q} = 1$, or (ii) $p = 1$ and $q = \infty$, or (iii) $p = \infty$ and $q = 1$.

Theorem 3.2 (Lemma 16.1 and Theorem 16.2 of [25]) *Every Banach lattice is σ-order complete.*

There are two very natural modes of 'convergence' in a Banach lattice: *order convergence* and *norm convergence*. The latter is well-known, the former less so. An order bounded sequence $\{v_n\}_{n\in\mathbb{N}}$ in a σ-complete Riesz space (e.g. a Banach lattice) *converges in order to* v if either of the following equivalent conditions holds:

$$v = \liminf_n v_n := \bigvee_n \bigwedge_{n\leq m} v_m, \quad v = \limsup_n v_n := \bigwedge_n \bigvee_{n\leq m} v_m.$$

For a monotone increasing sequence v_n, this definition simplifies to $v = \bigvee_n v_n$, which is often written $v_n \uparrow v$.

In a general σ-complete Riesz space, order and norm convergence are disjoint concepts, i.e. neither implies the other (see [25, Ex. 15.2] for two counter-examples). However if a sequence converges both in order and in norm then the limits are the same (see [25, Th. 15.4]). Moreover, for *monotone* sequences norm convergence implies order convergence:

Proposition 3.3 ([25] Theorem 15.3) *If $\{v_n\}_{n\in\mathbb{N}}$ is an increasing sequence in a normed Riesz space and if v_n converges to v in norm (notation $v_n \to v$), then $v_n \uparrow v$.*

In a Banach lattice we have the following stronger property.

Proposition 3.4 (Lemma 16.1 and Theorem 16.2 of [25]) *If $\{v_n\}_{n\in\mathbb{N}}$ is a sequence of positive vectors in a Banach lattice such that $\sup_n \|v_n\|$ converges, then $\bigvee_n v_n$ exists and $\|\bigvee_n v_n\| = \bigvee_n \|v_n\|$.*

It can also happen that order convergence implies norm convergence. A lattice norm on a Riesz space is called *σ-order continuous* if $v_n \downarrow 0$ (v_n is a decreasing sequence whose infimum is 0) implies $\|v_n\| \downarrow 0$.

Example 3.5 For $1 \leq p < \infty$, the L_p-norm is *σ-order continuous*, and thus order convergence and norm convergence coincide. However, for $p = \infty$ this is not the case as the following simple example shows. Consider the sequence of essentially bounded functions $v_n = 1_{[n,+\infty[}$: it is decreasing for the order on $\mathrm{L}_\infty(\mathbb{R}, \lambda)$ with the constant function 0 as its infimum, i.e. $v_n \downarrow 0$. However $\|v_n\| = 1$ for all n.

Many types of morphisms between Banach lattices are considered in the literature but most are at least *linear* and *positive*, that is to say they send positive vectors to positive vectors. From now on, we will assume that all morphisms are positive (linear) operators. Other than that, we will only mention two additional properties, corresponding to the two modes of convergence which we have examined. The first notion is very well-known: a linear operator $T : V \to W$ between normed

vector spaces is called *norm-bounded* if there exists $C \in \mathbb{R}$ such that $\|Tv\| \leq C\,\|v\|$ for every $v \in V$. The following result is familiar:

Theorem 3.6 *An operator $T : V \to W$ between normed vector spaces is norm-bounded iff it is continuous.*

Thus norm-bounded operators preserve norm-convergence. The corresponding order-convergence concept is defined as follows: an operator $T : V \to W$ between σ-order complete Riesz spaces is said to be *σ-order continuous* if whenever $v_n \uparrow v$, $Tv = \bigvee Tv_n$. It follows that we can consider two types of dual spaces on a Banach lattice V: on the one hand we can consider the *norm-dual*:

$$V^* = \{f : V \to \mathbb{R} \mid f \text{ is norm-continuous}\}$$

and the *σ-order-dual*:

$$V^\sigma = \{f : V \to \mathbb{R} \mid f \text{ is } \sigma\text{-order continuous}\}$$

The latter is sometimes known as the *Köthe dual* of V (see [11,25]). The two types of duals coincide for a large class of Banach spaces of interest to us.

Theorem 3.7 *If a Banach lattice V admits a strictly positive linear functional and has a σ-order-continuous norm, then $V^* = V^\sigma$.*

Example 3.8 The result above can directly be applied to our running example: given a measured space (X, μ) and an integer $1 \leq p < \infty$, the Lebesgue integral provides a strictly positive functional on $\mathrm{L}_p(X, \mu)$, and we already know from Example 3.5 that $\mathrm{L}_p(X, \mu)$ has a σ-order-continuous norm. It follows that

$$\mathrm{L}_p(X, \mu)^* = \mathrm{L}_p(X, \mu)^\sigma$$

Moreover, it is well-known that if (p, q) are Hölder conjugate and $1 < p, q < \infty$, then $\mathrm{L}_p(X, \mu)^* = \mathrm{L}_q(X, \mu)$, and thus $\mathrm{L}_p(X, \mu)^\sigma = \mathrm{L}_q(X, \mu)$. It is also known that $\mathrm{L}_1(X, \mu)^* = \mathrm{L}_\infty(X, \mu)$, and thus $\mathrm{L}_1(X, \mu)^\sigma = \mathrm{L}_\infty(X, \mu)$.

However Theorem 3.7 does *not* hold for $\mathrm{L}_\infty(X, \mu)$ since the L_∞-norm is not σ-order continuous, as was shown in Example 3.5. It is well-known that $\mathrm{L}_\infty(X, \mu)^* \neq \mathrm{L}_1(X, \mu)$, and in fact $\mathrm{L}_\infty(X, \mu)^*$ can be concretely described as the Banach lattice $\mathrm{ba}(X, \mu)$ of *charges* (i.e. finitely additive finite signed measures) which are absolutely continuous w.r.t, μ on X (see [12, IV.8.16]). However, as is shown in e.g. [25,4]

$$\mathrm{L}_\infty(X, \mu)^\sigma = \mathrm{L}_1(X, \mu) \tag{7}$$

As Examples 3.5 and 3.8 show, the $(-)^\sigma$ operation brings a lot of symmetry to the relationship between L_p-spaces since $\mathrm{L}_p(X, \mu)^\sigma = \mathrm{L}_q(X, \mu)$ for *any* Hölder conjugate pair $1 \leq p \leq \infty$. For this reason we will consider the category $\mathbf{BL}_\sigma$ whose objects are Banach lattices and whose morphisms are σ-order continuous positive operators. Note that the Köthe dual of a Banach lattice is a Banach lattice, and it easily follows that $(-)^\sigma$ in fact defines a contravariant functor $\mathbf{BL}_\sigma^{\mathrm{op}} \to \mathbf{BL}_\sigma$ which

acts on morphisms by pre-composition. As we will now see, $\mathbf{BL}_\sigma$ is the category in which predicate and state transformers are most naturally defined.

4 From Borel kernels to Banach lattices

The aim of this section is two-fold. First, we establish functorial bridges between Borel kernels (i.e. **Krn**) and operators on Banach lattices (i.e. $\mathbf{BL}_\sigma$). This will allow us to describe the convergence of a sequence of approximating kernels in terms of operator topology in §6. Second, we show that the functors between **Krn** and $\mathbf{BL}_\sigma$ also provide an interesting structural insight into some of the most important results in classical probability theory, in particular we show that the Radon-Nikodym and Riesz representation theorems can be described as natural transformations.

The functors S_p and T_p.

For $1 \leq p \leq \infty$, the operation which associates to a **Krn**-object (X, μ) the space $\mathrm{L}_p(X, \mu)$ can be thought of as either a contravariant or a covariant functor. We define the functors $\mathsf{S}_p : \mathbf{Krn} \to \mathbf{BL}_\sigma^{\mathrm{op}}, 1 \leq p \leq \infty$ as expected on objects, and on **Krn**-morphisms $f : X \rightarrowtriangle Y$ via the well-known 'predicate transformer' perspective:

$$\mathsf{S}_p(f) : \mathrm{L}_p(Y, \nu) \to \mathrm{L}_p(X, \mu), \phi \mapsto \lambda x. \int_Y \phi \, df(x) = \phi \bullet f$$

This defines a functor (see [6]). We define the covariant functors $\mathsf{T}_p : \mathbf{Krn} \to \mathbf{BL}_\sigma, 1 \leq p \leq \infty$ as $\mathsf{T}_p = \mathsf{S}_p \circ (-)^\dagger$.

The functor $\mathcal{M}^{\ll\cdot}$.

An *ideal* of a Riesz space V is a sub-vector space $U \subseteq V$ with the property that if $|u| \leq |v|$ and $v \in U$ then $u \in U$. An ideal U is called a *band* when for every subset $D \subseteq U$ if $\bigvee D$ exists in V, then it also belongs to U. Every band in a Banach lattice is itself a Banach lattice. Of particular importance is the band B_v generated by a singleton $\{v\}$, which can be described explicitly as

$$B_v = \{w \in V \mid (|w| \wedge n\,|v|) \uparrow |w|\}$$

Example 4.1 Let X be an SB-space and $\mathsf{ca}(X)$ denote the set of measures of bounded variation on X. It can be shown ([2, Th 10.56]) that $\mathsf{ca}(X)$ is a Banach lattice. The linear structure on $\mathsf{ca}(X)$ is as expected, the Riesz space structure is given by

$$(\mu \vee \nu)(A) = \sup\{\mu(B) + \nu(A \setminus B) \mid B \text{ measurable }, B \subseteq A\}$$

and the dual definition for the meet operation. The norm is given by the total variation i.e.

$$\|\mu\| = \sup \left\{ \sum_i^n |\mu(A_i)| \,\middle|\, \{A_1, \ldots, A_n\} \text{ a meas. partition of } X \right\}$$

Given $\mu \in \mathsf{ca}(X)$, the band B_μ generated by μ is just the set of measures of bounded variation which are absolutely continuous w.r.t. μ. In particular B_μ is a Banach lattice.

We can now define the functor $\mathcal{M}^{\ll\cdot} : \mathbf{Krn} \to \mathbf{BL}_\sigma$ by:

$$\begin{cases} \mathcal{M}^{\ll\cdot}(X,\mu) := B_\mu \\ \mathcal{M}^{\ll\cdot} f : \mathcal{M}^{\ll\cdot}(X,\mu) \to \mathcal{M}^{\ll\cdot}(Y,\nu), \rho \mapsto f \bullet \rho \end{cases}$$

We will usually write $\mathcal{M}^{\ll\cdot}(X,\mu)$ as $\mathcal{M}^{\ll\mu}(X)$.

Proposition 4.2 *Let $f : (X,\mu) \to (Y,\nu)$ be a* **Krn** *arrow. Let ρ be a finite measure on X such that $\rho \ll \mu$. Then $f \bullet \rho \ll \nu$, and thus $\mathcal{M}^{\ll\cdot}$ defines a functor.*

Proof. Let $B \subseteq Y$ be a measurable set. By definition, we have $(f \bullet \rho)(B) = \int_X ev_B \circ f \; d\rho$ where we recall that $ev_B : \mathsf{G}(X) \to \mathbb{R}_+$ is the evaluation morphism. Let $\{f_n^B\}_{n\in\mathbb{N}}$ be an increasing chain of simple functions converging pointwise to $ev_B \circ f$ such that for each n, $f_n^B = \sum_{i=1}^{k_n} \alpha_i^n 1_{A_i^n}$ with $\alpha_i^n \geq 0$. By the MCT,

$$(f \bullet \rho)(B) = \lim_n \int_X f_n^B \; d\rho = \lim_n \sum_{i=1}^{k_n} \alpha_i^n \rho(A_i^n).$$

Similarly,

$$\nu(B) = (f \bullet \mu)(B) = \lim_n \int_X f_n^B \; d\mu = \lim_n \sum_{i=1}^{k_n} \alpha_i^n \mu(A_i^n).$$

Notice that since the integral is linear and the sequence $\{f_n^B\}_n$ is increasing, the sequences $\{\int_X f_n^B \; d\rho\}_n$ and $\{\int_X f_n^B \; d\mu\}_n$ are also increasing. Assume $\nu(B) = 0$. Then for all n, $\int_X f_n^B \; d\mu = 0$. We deduce that for all n, for all $1 \leq i \leq k_n$, either $\alpha_i^n = 0$ or $\mu(A_i^n) = 0$. Using that $\rho \ll \mu$, we deduce that for all $1 \leq i \leq k_n$, either $\alpha_i^n = 0$ or $\rho(A_i^n) = 0$, from which we conclude that for all n, $\int_X f_n^B \; d\rho = 0$ and finally, $(f \bullet \rho)(B) = 0$. Hence, $f \bullet \rho \ll \nu$. □

Radon-Nikodym is natural.

We now present a first pair of natural transformations which will establish a natural isomorphism between the functors T_1 and $\mathcal{M}^{\ll\cdot}$. First, we define the *Radon-Nikodym transformation* $\mathsf{rn} : \mathcal{M}^{\ll\cdot} \to \mathsf{T}_1$ at each **Krn**-object (X,μ) by the map

$$\mathsf{rn}_{(X,\mu)} : \mathcal{M}^{\ll\mu}(X) \to \mathrm{L}_1(X,\mu), \quad \mathsf{rn}_{(X,\mu)}(\rho) = \frac{d\rho}{d\mu}$$

where $d\rho/d\mu$ is of course the Radon-Nikodym derivative of ρ w.r.t. μ. The fact that this transformation defines a positive operator between Banach lattices is simply a restatement of the usual Radon-Nikodym theorem [12, III.10.7.], combined with the well-known linearity property of the Radon-Nikodym derivative. To see that it is also σ-order-continuous, consider a monotone sequence $\mu_n \uparrow \mu$ converging in order

to μ in $\mathcal{M}^{\ll\nu}(X)$. This means that for any measurable set A of X, $\lim_{n\to\infty}\mu_n(A) = \mu(A)$. Since $(d\mu_n/d\nu)_{n\in\mathbb{N}}$ is bounded in L_1-norm the function $g = \bigvee_n d\mu_n/d\nu$ exists and is simply the pointwise limit $g(x) = \lim_{n\to\infty} d\mu_n/d\nu(x)$. It now follows from the monotone convergence theorem (MCT) that

$$\int_A g d\nu = \int_A \lim_{n\to\infty}\frac{d\mu_n}{d\nu}d\nu = \lim_{n\to\infty}\int_A \frac{d\mu_n}{d\nu}d\nu = \lim_{n\to\infty}\mu_n(A) = \mu(A)$$

in other words, $g = d\mu/d\nu$ and rn is well-defined. That rn is also natural has – to our knowledge – never been published.

Theorem 4.3 *The Radon-Nikodym transformation is natural.*

Proof. We start by proving the following Lemma

Lemma 4.4 *For any $f : (X,\mu) \to (Y,\nu)$, $\phi \in \mathrm{L}_1(Y,\nu)$, and $B_X \subseteq X$ measurable*

$$\int_{B_X}\left(\int_Y \phi df(x)\right)\, d\mu = \int_Y \phi(y) f^\dagger(y)(B_X)\, d\nu$$

Proof. We start by showing the equation on characteristic functions. If B_Y is measurable in Y, we have

$$\begin{aligned}\int_{B_X}\left(\int_Y 1_{B_Y} df(x)\right)\, d\mu &= \int_{B_X} f(x)(B_Y)\, d\mu \\ &= \int_Y 1_{B_Y}(y) f^\dagger(y)(B_X)\, d\nu \qquad \text{Eq. (4)}\end{aligned}$$

Since ϕ is measurable and integrable, there exists a sequence $\phi_n \uparrow \phi$ of simple functions such that $\lim_n \int_Y \phi_n\, d\nu < \infty$, and the results follows by the linearity of integration and the MCT. □

We can now prove the naturality of rn. Let $f : (X,\mu) \to (Y,\nu)$ be a **Krn**-morphism; we have on the one hand

$$\begin{aligned}\mathsf{rn}_{(Y,\mu)} \circ \mathcal{M}^{\ll\cdot}(f)(\rho)(y) &= \mathsf{rn}_{(Y,\nu)}(\int_X f(x)(-)\, d\rho)(y) \\ &= \frac{d\int_X f(x)(-)d\rho}{d\nu}(y) \qquad (*)\end{aligned}$$

and on the other

$$\begin{aligned}\mathsf{T}_1(f^\dagger) \circ \mathsf{rn}_{(X,\mu)}(\rho)(y) &= \mathsf{T}_p(f^\dagger)\left(\frac{d\rho}{d\mu}\right)(y) \\ &= \int_X \frac{d\rho}{d\mu}\, df^\dagger(y) \qquad (**)\end{aligned}$$

To show the equality of these two maps in $\mathrm{L}_1(Y,\nu)$ it is enough to show that they are equal ν-a.e. To see this, we show that $(**)$ satisfies the condition to be the

Radon-Nikodym derivative $(*)$. Let B_Y be a measurable subset of Y. We have from the well-known property of Radon-Nikodym derivatives:

$$\int_{B_Y} \frac{d\int_X f(x)(-)\ d\rho}{d\nu} d\nu = \int_{x\in X} f(x)(B_Y)\ d\rho$$

Moreover, we have

$$\begin{aligned}\int_{B_Y}\int_X \frac{d\rho}{d\mu}\ df^\dagger(y)\ d\nu &\overset{(1)}{=} \int_{x\in X} \frac{d\rho}{d\mu}(x) f(x)(B_Y)\ d\mu \\ &\overset{(2)}{=} \int_{x\in X} f(x)(B)\ d\rho\end{aligned}$$

where (1) is by Lemma 4.4 and (2) is a well-known property of Radon-Nikodym derivatives. □

Secondly, we define the *Measure Representation transformation* $\mathsf{mr} : \mathsf{T}_1 \to \mathcal{M}^{\ll\cdot}$ at each **Krn**-object (X, μ) by the map $\mathsf{mr}_{(X,\mu)} : \mathsf{T}_1(X, \mu) \to \mathcal{M}^{\ll\cdot}(X, \mu)$ defined as

$$\mathsf{mr}_{(X,\mu)}(f)(B_X) = \int_{B_X} f d\mu$$

This is a very well-known construction in measure theory, and the fact that $\mathsf{mr}_{(X,\mu)}$ is a σ-order continuous operator between Banach lattices is immediate from the linearity of integrals and the MCT.

Theorem 4.5 *The Measure Representation transformation is natural.*

Proof. We start with the following elementary lemma.

Lemma 4.6 *If* $\psi, \phi \in \mathrm{L}_1(X, \mu)$ *then*

$$\int_X \psi\phi\ d\mu = \int_X \psi\ d(\mathsf{mr}_{(M,\mu)}\phi)$$

Proof. The proof of naturality now follows easily: it is enough to show the equality in the case where $\psi = 1_{B_X}$ for a measurable subset B_X of X, and the result then extends to all measurable functions by linearity of integrals and the MCT. We have

$$\begin{aligned}\int_X 1_{B_X}\phi\ d\mu = \int_{B_X} \phi d\mu &:= \mathsf{mr}_{(M,\mu)}(\phi)(B_X) \\ &= \int_X 1_{B_X}\ d(\mathsf{mr}_{(M,\mu)}(\phi))\end{aligned}$$

□

To show naturality we now let $f : (X, \mu) \to (Y, \nu)$ be a **Krn**-morphism, $\phi \in \mathrm{L}_1(X, \mu)$

and B_Y measurable in Y

$$\begin{aligned}
&\mathsf{mr}_{(Y,\nu)}\mathsf{T}_1(f^\dagger)(\phi)(B_X)\\
&= \mathsf{mr}_{(Y,\nu)}(\phi \bullet f^\dagger)(B_Y)\\
&= \int_{B_Y} \phi \bullet f^\dagger \; d\nu\\
&= \int_{B_Y}\int_X \phi df^\dagger(y) \; d\nu\\
&= \int_X f(x)(B_Y)\phi(x) \; d\mu && \text{Lemma 4.4}\\
&= \int_X f(x)(B_Y) \; d(\,\mathrm{d}(\mathsf{mr}_{(X,\mu)}(\phi)) && \text{Lemma 4.6}\\
&= \mathcal{M}^{\ll\cdot} f \circ \mathsf{mr}_{(X,\mu)}(\phi)(B_Y)
\end{aligned}$$

□

Riesz representations are natural.

We now present a second pair of natural transformations which will establish a natural isomorphism between $(-)^\sigma \circ \mathsf{S}_\infty$ and $\mathcal{M}^{\ll\cdot}$. First, we define the *Riesz Representation* transformation $\mathsf{rr} : (-)^\sigma \circ \mathsf{S}_\infty \to \mathcal{M}^{\ll\cdot}$ at each **Krn**-object (X,μ) by the map $\mathsf{rr}_{(X,\mu)} : (-)^\sigma \circ \mathsf{S}_\infty(X,\mu) \to \mathcal{M}^{\ll\cdot}(X,\mu)$ defined as

$$\mathsf{rr}_{(X,\mu)}(F)(B_X) = F(1_{B_X})$$

This construction is key to a whole collection of results in functional analysis commonly known as Riesz Representation Theorems (see [2] Chapter 14 for an overview). One can readily check that the Riesz Representation transformation is well-defined: $\mathsf{rr}_{(X,\mu)}(F)(\emptyset) = F(0) = 0$ and the σ-additivity of $\mathsf{rr}_{(X,\mu)}(F)$ follows from the σ-order-continuity of F. To see that $\mathsf{rr}_{(X,\mu)}(F) \ll \mu$, assume that $\mu(B_X) = 0$, then clearly $1_{B_X} = 0$ μ-a.e., i.e. $1_{B_X} = 0$ in $\mathrm{L}_\infty(X,\mu)$, and thus $F(1_{B_X}) = 0$.

Theorem 4.7 *The Riesz Representation transformation is natural.*

Proof. Again, we start with a simple but helpful Lemma.

Lemma 4.8 *Let $F \in (\mathsf{S}_\infty(X,\mu)^\sigma$ and $\phi \in \mathsf{S}_\infty(X,\mu)$, then*

$$F(\phi) = \int_X \phi \; d(\mathsf{rr}_{(X,\mu)}(F))$$

Proof. Starting with characteristic functions, let $\phi = 1_B$ for some measurable subset B of X. We then have

$$F(1_B) := \mathsf{rr}_{(X,\mu)}(F)(B) = \int_X 1_B \; d(\mathsf{rr}_{(X,\mu)}(F))$$

We can then extend the result to simple functions by linearity and to all functions in $\mathrm{L}_\infty(X,\mu)$ by the MCT. □

To show naturality we now let $f : (X,\mu) \rightarrowtriangle (Y,\nu)$ be a **Krn**-morphism, $F \in (\mathsf{S}_\infty(X,\mu))^\sigma$ and B_Y measurable in Y. We have

$$\begin{aligned}
\mathcal{M}^{\ll\cdot} f \circ \mathsf{rr}_{(X,\mu)}(F)(B_Y) &= \int_X f(x)(B_Y)\ d(\mathsf{rr}_{(X,\mu)}(F)) \\
&= F(f(\cdot)(B_Y)) && \text{Lemma 4.8} \\
&= F(\int_X 1_{B_Y} df(\cdot)) \\
&= F(1_{B_Y} \bullet f) \\
&= \mathsf{rr}_{(Y,\nu)}(F(- \bullet f))(B_Y) \\
&= \mathsf{rr}_{(Y,\nu)} \circ (\mathsf{T}_1 f)^\sigma(F)(B_Y)
\end{aligned}$$

□

Finally, we define the *Functional Representation transformation* fr at each **Krn**-object (X,μ) by the map $\mathsf{fr}_{(X,\mu)} : \mathcal{M}^{\ll\cdot}(X,\mu) \to (-)^\sigma \circ \mathsf{S}_\infty(X,\mu)$ by

$$\mathsf{fr}_{(X,\mu)}(\mu)(\phi) = \int_X \phi d\mu$$

This construction is also completely standard in measure theory, although it has never to our knowledge been seen as a natural transformation.

Theorem 4.9 *The Functional Representation transformation is well-defined, i.e.* $\mathsf{fr}_{(X,\mu)}$ *is a σ-order continuous positive operator, and is natural.*

Proof. We start by showing that fr is well defined. The linearity of $\mathsf{fr}_{(X,\mu)}$ is easily checked on simple functions and extended by the CMT. Positivity is also immediate. For the σ-order continuity, let $\mu_m \uparrow \mu$, $\phi \in \mathrm{L}_\infty(X,\mu)$, and $\phi_n \uparrow \phi$ be a monotone approximation of ϕ by simple functions. We need to show that

$$\lim_{m\to\infty} \int_X \phi\ d\mu_m = \int_X \phi\ d\mu$$

For note first that the doubly indexed series $\int_X \phi_n d\mu_m$ is monotonically increasing in m, since the μ_m are monotonically increasing. Note also that the differences

$$d_{mn} := \int_X \phi_n\ d\mu_{m+1} - \int_X \phi_n\ d\mu_m$$

are monotonically increasing in n. Indeed we have

$$\begin{aligned}
&\left(\int_X \phi_{n+1}\ d\mu_{m+1} - \int_X \phi_{n+1}\ d\mu_m\right) - \left(\int_X \phi_n d\mu_{m+1} - \int_X \phi_n\ d\mu_m\right) \\
&= \int_X (\phi_{n+1} - \phi_n)\ d\mu_{m+1} - \int_X (\phi_{n+1} - \phi_n)\ d\mu_m > 0
\end{aligned}$$

since the sequences ϕ_n and μ_m are monotonically increasing. Since d_{mn} is monotonically increasing in n we can apply the CMT to d_{mn} seen as a function of m w.r.t. the counting measure, i.e.

$$\lim_{n\to\infty} \sum_{m}^{\infty} d_{mn} = \sum_{m}^{\infty} \lim_{n\to\infty} d_{mn}$$

which is to say, by taking partial sums

$$\begin{aligned}
\lim_{n\to\infty} \lim_{m\to\infty} \sum_{k=1}^{m} d_{kn} &= \lim_{n\to\infty} \lim_{m\to\infty} \int_X \phi_n \, d\mu_m \\
&= \lim_{m\to\infty} \sum_{k=1}^{m} \lim_{n\to\infty} d_{kn} \\
&= \lim_{m\to\infty} \lim_{n\to\infty} \int_X \phi_n \, d\mu_m
\end{aligned}$$

which concludes the proof that fr is well-defined.

We now prove naturality. Let $f : (X,\mu) \rightarrowtriangle (Y,\nu)$ be a **Krn**-morphism, $\rho \in \mathcal{M}^{\ll\mu}(X)$ and $\phi \in \mathsf{S}_\infty(Y,\nu)$ we then have

$$\begin{aligned}
\mathsf{fr}_{(Y,\nu)} \circ \mathcal{M}^{\ll\cdot} f(\rho)(\phi) &= \int_Y \phi \, d(\mathcal{M}^{\ll\cdot} f(\rho)(\phi)) \\
&= \int_Y (\phi \bullet f) \, d\rho && \text{Theorem 2.6} \\
&= \mathsf{S}_\infty(f) \left(\int_Y (-) \, d\rho \right) (\phi) \\
&= \mathsf{S}_\infty(f) \circ \mathsf{fr}_{(X,\mu)}(\rho)(\phi)
\end{aligned}$$

□

Natural Isomorphisms

We have now defined the following four natural transformations:

$$\mathsf{T}_1 \underset{\mathsf{rn}}{\overset{\mathsf{mr}}{\rightleftarrows}} \mathcal{M}^{\ll\cdot} \underset{\mathsf{rr}}{\overset{\mathsf{fr}}{\rightleftarrows}} (\mathsf{S}_\infty)^\sigma$$

In fact, both pairs form natural isomorphisms which can be restricted to arbitrary Hölder conjugate pairs (p,q).

Theorem 4.10 rn *and* mr *are inverse of one another, in particular there exists a* natural *isomorphism between* $\mathcal{M}^{\ll\mu}(X)$ *and* $\mathrm{L}_1(X,\mu)$.

Proof. The fact that $\mathsf{rn}_{(X,\mu)}$ and $\mathsf{mr}_{(X,\mu)}$ are inverse of each other is just a restatement of the two well-known equalities for Radon-Nikodym derivatives:

$$\frac{d \int_- \phi \, d\mu}{d\mu} = \phi \quad \text{and} \quad \int_{B_X} \frac{d\rho}{d\mu} \, d\mu = \rho(B_X)$$

□

Theorem 4.11 rr *and* fr *are inverse of each other, in particular there exists a* natural *isomorphism between* $\mathcal{M}^{\ll\mu}(X)$ *and* $(\mathrm{L}_\infty(X,\mu))^\sigma$.

Proof. Let (X,μ) be a **Krn**-object, let $F \in (\mathrm{L}_\infty(X,\mu))^\sigma$ and let $\phi \in \mathrm{L}_\infty(X,\mu)$. We have

$$\mathsf{fr}_{(X,\mu)} \circ \mathsf{rr}_{(X,\mu)}(F)(\phi) = \int_X \phi \; d(\mathsf{rr}_{(X,\mu)}(F)) = F(\phi)$$

where the last equality follows from Lemma 4.8. Similarly, we have

$$\mathsf{rr}_{(X,\mu)} \circ \mathsf{fr}_{(X,\mu)}(\rho)(B_X) = \mathsf{fr}_{(X,\mu)}(\rho)(1_{B_X}) = \int_X 1_{B_X} d\rho = \rho(B_X)$$

□

We can now conclude that the isomorphism proved in Theorem 6 of [6] is in fact natural.

Corollary 4.12 *There exists a* natural *isomorphism between* $\mathsf{T}_1 := \mathsf{S}_1 \circ (-)^\dagger$ *and* $(-)^\sigma \circ \mathsf{S}_\infty$.

We can in fact restrict this result to *any* Hölder conjugate pair (p,q):

Theorem 4.13 *For* $1 \leq p \leq \infty$ *with Hölder conjugate* q*, the natural transformation* $\mathsf{rn} \circ \mathsf{rr}$ *restricts to a natural transformation* $(-)^\sigma \circ \mathsf{S}_q \to \mathsf{T}_p$.

Proof. The case $p = 1$ has been treated already, for the case of $1 < p < \infty$, see for example the proof of Theorem 4.4.1 of [3]. Finally for the case of $p = \infty$, see Proposition 3.3 of [4]. □

The correspondence between the various categories and functors discussed in this section are summarized by

$$\begin{tikzcd}
 & \mathbf{BL}_\sigma & \\
\mathbf{BL}_\sigma^{\mathrm{op}} \arrow[ur, "(-)^\sigma"] \arrow[rr, shift right, "\mathsf{rr}"'] & & \mathbf{Krn}^{\mathrm{op}} \arrow[ul, "\mathsf{S}_1"'] \arrow[ll, shift right, "\mathsf{fr}"'] \arrow[ll, shift left, "\mathsf{mr}"] \arrow[rr, shift right, "\mathsf{rn}"'] \arrow[rr, shift left, "\mathsf{rn}"] \\
 & \mathbf{Krn} \arrow[ul, "\mathsf{S}_\infty"] \arrow[ur, "(-)^\dagger"'] \arrow[uu, "\mathcal{M}^{\ll\cdot}"] &
\end{tikzcd} \tag{8}$$

5 Approximations

In this section we develop a scheme for approximating kernels which follows naturally from the †-structure of **Krn**. Consider $f : (X,\mu) \rightarrowtriangle (Y,\nu)$ and a pair of

deterministic maps $p : (X, \mu) \to (X', p_*\mu)$ and $q : (Y, \nu) \to (Y', q_*\nu)$ (typically these maps coarsen the spaces X and Y).

$$\begin{tikzcd}
(X,\mu) \arrow[r, shift left, "f"] \arrow[r, shift right, "f^{p,q}"'] \arrow[d, bend right, "p"'] & (Y,\nu) \arrow[d, bend right, "q"'] \\
(X',p_*\mu) \arrow[r, "f_{p,q}"'] \arrow[u, bend right, "p^\dagger_\mu"'] & (Y',q_*\nu) \arrow[u, bend right, "q^\dagger_\nu"']
\end{tikzcd} \tag{9}$$

The †-structure of **Krn** allows us to define the new kernels

$$f^{p,q} := q^\dagger_\nu \bullet q \bullet f \bullet p^\dagger_\mu \bullet p \qquad : X \rightarrowtriangle Y \tag{10}$$

$$f_{p,q} := q \bullet f \bullet p^\dagger_\mu \qquad : X' \rightarrowtriangle Y' \tag{11}$$

The supscript notation is meant to indicate that the approximation lives 'upstairs' in Diagram (9) and conversely for the subscripts. Intuitively, $f_{p,q}$ and $f^{p,q}$ take the average of f over the fibres given by p, q according to μ and ν (see §7 for concrete calculations). The advantage of (11) is that we can approximate a kernel on a huge space by a kernel on a, say, finite one. The advantage of (10) is that although it is more complicated, it is morally equivalent and has the same type as f, which means that we can compare it to f.

A very simple consequence of our definition is that Bayesian inversion commutes with approximations. We shall use this in §7 to perform *approximate Bayesian inference*.

Theorem 5.1 *Let $f : (X, \mu) \rightarrowtriangle (Y, \nu)$, let $p : X \to X'$ and $q : Y \to Y'n$ be a pair of deterministic maps, then*

$$(f^\dagger)^{q,p} = (f^{p,q})^\dagger \text{ and } (f^\dagger)_{q,p} = (f_{p,q})^\dagger$$

In practice we will often consider endo-kernels $f : X \rightarrowtriangle X$ with a single coarsening map $p : X \to X'$ to a *finite space*. In this case (10) simplifies greatly.

Proposition 5.2 *Under the situation described above*

$$f^p := p^\dagger_\nu \bullet p \bullet f \bullet p^\dagger_\mu \bullet p = f \bullet p^\dagger_\mu \bullet p \tag{12}$$

In the case covered by Proposition 5.2, the interpretation of f^p is very natural: for each $x \in X$ the measure $f(x)$ is approximated by its average over the fibre to which x belongs, conditioned on being in the fibre. For fibres with strictly positive μ-probability, this is simply

$$f^p(x)(A) = \frac{\int_{y \in p^{-1}(p(x))} f(y)(A)\ d\mu}{\mu(p^{-1}(p(x))}$$

However (12) also covers the case of μ-null fibres. Note also that in the case where $f^p = f$, the map p corresponds to what is known as a strong functional bisimulation for f.

Approximating is non-expansive.

It is well-known that conditional expectations are non-expansive and we know from Lemma 2.8 that pre-composing by $p_\mu^\dagger \bullet p$ as in (12) amounts to conditioning. The following lemma is an easy consequence.

Lemma 5.3 *Let $f : (X, \mu) \rightarrowtriangle (Y, \nu)$ and $q : X \to X'$ be a deterministic quotient, then for all $1 \leq p \leq \infty$ and $\phi \in \mathrm{L}_p(Y, \nu)$*

$$\|\mathsf{S}_p f^q(\phi)\|_p \leq \|\mathsf{S}_p f(\phi)\|_p$$

Compositionality of approximations.

In the case where we wish to approximate a composite kernel $g \bullet f$, it might be convenient, for modularity reasons, to approximate f and g separately. This does not entail any loss of information provided the quotient maps are hemi-bisimulations, in the following sense. Let $p : X \to X', q : Y \to Y', r : Z \to Z'$ be deterministic quotients and let $f : (X, \mu) \rightarrowtriangle (Y, \nu), g : (Y, \nu) \rightarrowtriangle (Z, \rho)$ be composable kernels. We say that q is a left hemi-bisimulation for f if $f = q^\dagger \bullet q \bullet f$, and conversely that it is a right hemi-bisimulation for g if $g = g \bullet q^\dagger \bullet q$ holds. In either case, one can verify using Theorems 2.7 and 5.1 that approximation commutes with composition, i.e. that $(g \bullet f)^{p,r} = g^{q,r} \bullet f^{p,q}$.

Discretization schemes.

We will use (11) and (12) to build sequences of arbitrarily good approximations of kernels. For this we introduce the following terminology.

Definition 5.4 We define a *discretization scheme for an SB-space X* to be a countable co-directed diagram (ccd) of finite spaces for which X is a cone (not necessarily a limit).

If $(X_i)_{i \in I}$ is a discretization scheme of X and $p_i : X \to X_i$ are the maps making X a cone, then it follows from the definition that if $i < j$, $\sigma(p_i) \subseteq \sigma(p_j)$ where $\sigma(p_i)$ is the σ-algebra generated by p_i. For each $i \in I$ the finite quotient p_i defines a *measurable partition* of X whose disjoint components $p_i^{-1}(\{k\}), k \in X_i$ we will call *cells*. By Theorem 2.1 every SB-space has a discretization scheme for which it is not just a cone but a limit. In practice we will work with discretization schemes linearly ordered by $\mathbb{N}$. In this case the sequence $(X, \sigma(p_n))_{n \in \mathbb{N}}$ defines what probabilists call a *filtration* and we will denote the approximation f^{p_n} given by (12) simply by f^n.

6 Convergence

We now turn to the question of convergence of approximations. There appears to be little literature on the subject of the convergence of approximations of Markov

kernels. One rare reference is [5]. Via the functor S_p defined above in Sections 3 and 4 we can seek a topology in terms of the *operators associated to a sequence of kernels*. Indeed, following [5], we will prove convergence results for the *Strong Operator Topology* (SOT).

Definition 6.1 We will say that a sequence of kernels $f^n : X \rightarrowtriangle Y$ *converges to* $f : X \rightarrowtriangle Y$ *in strong operator topology*, and write $f^n \longrightarrow_{\mathrm{s}} f$, if $\mathsf{S}_1 f^n$ converges to $\mathsf{S}_1 f$ in the strong operator topology, i.e. if

$$\lim_{n\to\infty} \|\mathsf{S}_1 f^n(\phi) - \mathsf{S}_1 f(\phi)\|_1 = 0$$

Proving convergence.

The following key lemma is a consequence of Lévy's upward convergence Theorem ([23, Th. 14.2]) .

Lemma 6.2 *Let $f : (X, \mu) \rightarrowtriangle (Y, \nu)$ be a* **Krn**-*morphism and let $p_n : X \to X_n, n \in \mathbb{N}$ be a discretization scheme such that for $\mathcal{B}_X$ the Borel σ-algebra of X we have $\mathcal{B}_X = \sigma\left(\bigcup_n \sigma(p_n)\right)$, and let $A \subseteq Y$ be measurable, then for $f^n := f \bullet p_n^\dagger \bullet p_n$*

$$\lim_{n\to\infty} f^n(x)(A) = f(x)(A)$$

for μ-almost every $x \in X$. Moreover,

$$\lim_{n\to\infty} \|\mathsf{S}_1 f^n(1_A) - \mathsf{S}_1 f(1_A)\|_1 = 0$$

Proof. The map $f(-)(A) : X \to \mathbb{R}$ defines a random variable, and the discretization scheme defines a filtration $\sigma(p_n) \subseteq \sigma(p_{n+1})$ whose union is $\mathcal{B}_X$. Following Lemma 2.8 and Proposition 5.2 we have

$$f^n(x)(A) := f(x)(A) \bullet p_n^\dagger \bullet p_n = \mathbb{E}\left[f(x)(A) \mid \sigma(p_n)\right]$$

We thus have a sequence $f^n(-)(A)$ of random variables $X \to \mathbb{R}$ which is adapted to the filtration $\sigma(p_n), n \in \mathbb{N}$ by construction. We can now compute for any $m < n$

$$\begin{aligned}
&\mathbb{E}\left[f^n(x)(A) \mid \sigma(p_m)\right] \\
=& f(x)(A) \bullet p_n^\dagger \bullet p_n \bullet p_m^\dagger \bullet p_m \\
\overset{(1)}{=}& f(x)(A) \bullet p_n^\dagger \bullet p_n \bullet (p_{nm} \bullet p_n)^\dagger \bullet p_m \\
\overset{(2)}{=}& f(x)(A) \bullet p_n^\dagger \bullet p_n \bullet p_n^\dagger \bullet p_{nm}^\dagger \bullet p_m \\
\overset{(3)}{=}& f(x)(A) \bullet p_m^\dagger \bullet p_m = f^m(x)(A)
\end{aligned}$$

where (1) is by definition (5.4), (2) is by Thm (2.10) and (3) is by Theorem (2.7). We have thus shown that $f^n(-)(A)$ is a martingale for the filtration generated by the discretization scheme, and the result now follows from Lévy's upward convergence Theorem ([23, Th. 14.2]) since $f(x)(A) = \mathbb{E}\left[f(x)(A) \mid \sigma\left(\bigcup_n \sigma(p_n)\right)\right]$. □

Theorem 6.3 (Convergence of Approximations Theorem) *Under the conditions of Lemma 6.2, for μ-almost every $x \in X$*

$$\lim_{n\to\infty} f^n(x)(A) = f(x)(A)$$

for all Borel subsets A. Moreover,

$$\lim_{n\to\infty} \|\mathsf{S}_1 f^n(\phi) - \mathsf{S}_1 f(\phi)\|_1 = 0$$

for any $\phi \in \mathrm{L}_1(X,\nu)$. In other words $f^n \longrightarrow_{\mathrm{s}} f$.

Proof. Let $(B_n)_{n\in\mathbb{N}}$ be a countable basis for the Borel σ-algebra of X, which we assume w.l.o.g. is closed under finite unions and intersections. It follows from Lemma 6.2 that for each B_n, $\lim_k f^k(x)(B_n) = f(x)(B_n)$ for all $x \in X \setminus N_n$ where $\mu(N_n) = 0$. It follows that for every $x \in X \setminus \bigcup_i N_i$

$$\lim_{k\to\infty} f^k(x)(B_n) = f(x)(B_n)$$

for all basic Borel sets B_n, and $\mu(\bigcup_i N_i) = 0$. Now we use the $\pi - \lambda$-lemma with $(B_n)_{n\in\mathbb{N}}$ as our π-system. We define

$$\mathcal{L} := \{C \mid f^n(x)(C) \to f(x)(C) \text{ for all } x \in X \setminus \cup_i N_i\}$$

and show that it is a λ-system. Clearly each $B_n \in \mathcal{L}$. Suppose $C \in \mathcal{L}$, it is then immediate that $C^c \in \mathcal{L}$. Now consider a sequence $C_i \in \mathcal{L}$ with $C_i \subseteq C_{i+1}$, and let $C_\infty := \cup_{i=1}^\infty C_i$. We want to show that

$$\begin{aligned}
\lim_n f^n(x)(C_\infty) &= \lim_n \lim_m f^n(x)(C_m) \\
&\overset{(*)}{=} \lim_m \lim_n f^n(x)(C_m) \\
&= \lim_m f(x)(C_m) \\
&= f(x)(C_\infty)
\end{aligned}$$

where $(*)$ is the only step we need to justify. To show the iterated limits can be switched, note first that since

$$\begin{aligned}
|f^n(x)(C_m) - f(x)(C_\infty)| &= |f^n(x)(C_m) - f(x)(C_m) + f(x)(C_m) - f(x)(C_\infty)| \\
&\leq |f^n(x)(C_m) - f(x)(C_m)| + |f(x)(C_m) - f(x)(C_\infty)|
\end{aligned}$$

since the two terms converge separately, for any $\epsilon > 0$ we can find $N > 0$ s. th. for all $m, n \geq N$, $|f^n(x)(C_m) - f(x)(C_\infty)| < \epsilon/2 + \epsilon/2 = \epsilon$. Thus $\lim_{(m,n)\to\infty} f^n(x)(C_m) = f(x)(C_\infty)$.

Now note also that for all $m_0 \in \mathbb{N}$ the sequence $f^n(x)(C_{m_0})$ converges to $f(x)(C_{m_0})$ (by definition of $\mathcal{L}$), and it is not hard to see that $f(x)(C_m)$ converges to $f(x)(C_\infty)$. Conversely for all $n_0 \in \mathbb{N}$, the sequence $f^{n_0}(x)(C_m)$ converges to

$f^{n_0}(C_\infty)$ (by virtue of $f^{n_0}(x)$ being a measure). For $\epsilon/2 > 0$ we can find $N > 0$ such that for all $m, n > N$, $|f^n(x)(C_m) - f(x)(C_\infty)| < \epsilon/2$. We can also find $M > 0$ such that for all $m > M$, $|f^n(x)(C_m) - f^n(x)(C_\infty)| < \epsilon/2$. By taking the maximum of N and M it is clear that for all m, n above this maximum

$$\begin{aligned}&|f^n(x)(C_\infty) - f(x)(C_\infty)|\\ &\leq |f^n(x)(C_\infty) - f^n(x)(C_m)| + |f^n(x)(C_m) - f(x)(C_\infty)| < \epsilon\end{aligned}$$

We have thus shown that

$$\begin{aligned}\lim_{m\to\infty}\lim_{n\to\infty} f^n(x)(C_m) &= \lim_{(m,n)\to\infty} f^n(x)(C_m)\\ &= f(x)(C_\infty) = \lim_{n\to\infty}\lim_{m\to\infty} f^n(x)(C_m).\end{aligned}$$

Thus $\mathcal{L}$ is a λ-system, and it follows from the $\pi - \lambda$-lemma that $\sigma((B_n)_{n\in\mathbb{N}}) \subseteq \mathcal{L}$ which concludes the proof of pointwise almost everywhere convergence.

For the proof of L_1-convergence we start by showing that

$$\lim_{n\to\infty} \|\mathsf{S}_p f^n(1_A) - \mathsf{S}_p(1_A)\|_1 = 0 \tag{13}$$

for any Borel subset A. For this we use exactly the same reasoning as above. The only difference is that we need to check that

$$\lim_{n\to\infty} \int_X |f^n(x)(C_\infty) - f(x)(C_\infty)|\, d\mu = 0$$

For this we use the fact that we have just shown $f^n(x)(C_\infty) \to f(x)(C_\infty)$ pointwise almost everywhere, and that $|f^n(x)(C_\infty) - f(x)(C_\infty)| \leq 1$ with 1 μ-integrable. It follows by dominated convergence that

$$\begin{aligned}&\lim_{n\to\infty} \int_X |f^n(x)(C_\infty) - f(x)(C_\infty)|\, d\mu\\ &= \int_X \lim_{n\to\infty} |f^n(x)(C_\infty) - f(x)(C_\infty)|\, d\mu = 0\end{aligned}$$

which concludes the proof of (13). To extend the result to simple functions and then to arbitrary functions $\phi \in \mathrm{L}_1(X, \nu)$ is routine. □

Note that operators of the shape $\mathsf{S}_p f^n$ obtained from a discretization scheme are *finite rank operators*. Thus, we, in fact, also obtained a theorem to approximate stochastic operators by *stochastic operators of finite rank* for the SOT topology. In general, we cannot hope for convergence in the stronger norm topology since the identity operator – which is stochastic – is a limit of operators of finite rank in the norm topology iff the space is finite dimensional.

Note also that the various relationships established in Section 4 allow us to move from an approximation of a kernel to an approximation of the corresponding Markov operator. Since a discretization scheme making $f^n \longrightarrow_{\mathrm{s}} f$ will also make

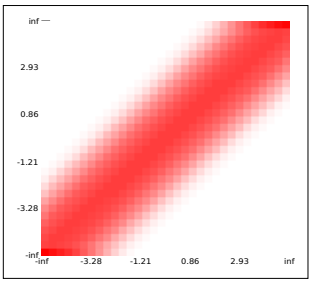

Fig. 1. Log-likelihood of $f^{5,3}$

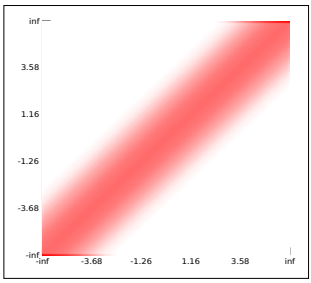

Fig. 2. Log-likelihood of $f^{6,10}$

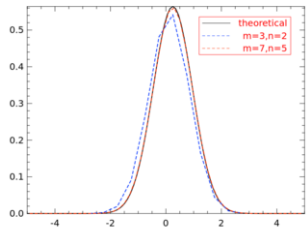

Fig. 3. Approximate posteriors

$(f^\dagger)^n \longrightarrow_{\mathrm{s}} f^\dagger$, it follows from Theorem 4.10 that we get a finite rank approximation of the Markov operator $\mathcal{M}^{\ll\cdot}(f)$.

7 Application: Approximate Bayesian Inference

Consider again the inference problem from the introduction. There one needed to invert $f(x) = \mathcal{N}(x, 1)$ with prior $\mu = \mathcal{N}(0, 1)$. We can use Theorem 5.1 to see how our approximate Bayesian inverse compares to the exact solution which in this simple case is known to be $f^\dagger_\mu(0.5) = \mathcal{N}(1/4, 1/2)$. To do this, we use a doubly indexed discretization scheme:

$$q_{mn} : \mathbb{R} \to 2 \times m \times n + 2$$

defining a window of width $2m$ centred at 0 divided in $2mn$ equal intervals; with the remaining intervals $(-\infty, -m]$ and (m, ∞) each sent to a point (hence the $+2$ above).

Since all classes induced by q_{mn} have positive μ-mass, approximants can be computed simply as:

$$f^{m,n}([k])([l]) = \mu[k]^{-1} \int_{x \in [k]} \mathcal{N}(x, 1)([l])\ d\mu$$

where $[k]$, $[l]$ range over classes of q_{mn}. The corresponding stochastic matrices are shown in Fig. 1 and 2 for $m, n = 5, 3$ and $6, 10$ respectively.

Since these approximants are finite, their Bayesian inverse can be computed directly by Bayes theorem (i.e. taking the adjoint of the stochastic matrices):

$$f^{m,n\dagger}([l])([k]) = \frac{\mu[k] \cdot f^{m,n}([k])([l])}{\nu[l]} \tag{14}$$

with $\nu = f_*(\mu)$. Commutation of inversion and approximation guarantees that the $f^{m,n\dagger}$ converge to $f^\dagger$.

Indeed, Fig. 3 shows the the Lebesgue density of $f^{m,n\dagger}(0.5)$ for $m, n = 3, 2$ (in dashed blue) and $7, 5$ (dashed red). The latter approximant is already hardly distinguishable from the exact solution (solid black).

It must be emphasized that this example is meant only as an illustration and does not constitute a universal solution to the irreducibly hard (not even computable

in general [1]) problem of performing Bayesian inversion. Also, not all quotients are equally convenient: what makes the approach computationally tractable is that the fibres are easily described and the measure conveniently evaluated on such fibres.

8 Conclusion

We have presented a framework for the exact *and* approximate semantics of first-order probabilistic programming. The semantics can be read off either in terms of kernels between measured spaces, or in terms of operators between L_p spaces. Either forms come with related involutive structures: Bayesian inversion for (measured) kernels between Standard Borel spaces, and Köthe duality for positive linear and σ-continuous operators between Banach lattices. Functorial relations between both forms can themselves be related by way of natural isomorphisms. Our main result is the convergence of general systems of finite approximants in terms of the strong operator topology (the SOT theorem). Thus, in principle, one can compute arbitrarily good approximations of the semantics of a probabilistic program of interest for any given (measurable) query. Future work may allow one to derive stronger notions of convergences given additional Lipschitz control on kernels, or to develop approximation schemes that are adapted to the measured kernel of interest. More ambitiously perhaps, one could investigate whether MCMC sampling schemes commonly used to perform approximate Bayesian inference in the context of probabilistic programming could be seen as randomized approximations of the type considered in this paper.

References

[1] N. L. Ackerman, C. E. Freer, and D. M. Roy. Noncomputable conditional distributions. In *LICS 2011*, pages 107–116, 2011.

[2] C. Aliprantis and K. Border. *Infinite dimensional analysis*, volume 32006. Springer, 1999.

[3] V. I. Bogachev. *Measure Theory I*. Springer, 2006.

[4] P. Chaput, V. Danos, P. Panangaden, and G. Plotkin. Approximating Markov Processes by averaging. *Journal of the ACM*, 61(1), January 2014.

[5] K. Choo-Whan. Approximation theorems for Markov operators. *Probability Theory and Related Fields*, 21(3):207–214, 1972.

[6] F. Clerc, V. Danos, F. Dahlqvist, and I. Garnier. Pointless learning. In *FoSSaCS*, pages 355–369. Springer, 2017.

[7] F. Dahlqvist, V. Danos, I. Garnier, and O. Kammar. Bayesian Inversion by Omega-Complete Cone Duality (Invited Paper). In *CONCUR 2016*, volume 59 of *LIPIcs*, pages 1:1–1:15. Schloss Dagstuhl, 2016.

[8] V. Danos, J. Desharnais, and P. Panangaden. Conditional expectation and the approximation of labelled Markov processes. In *International Conference on Concurrency Theory*, pages 477–491. Springer, 2003.

[9] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Metrics for labelled Markov processes. *TCS*, 318(3):323–354, 2004.

[10] J. Desharnais, R. Jagadeesan, V. Gupta, and P. Panangaden. Approximating labeled Markov processes. In *LiCS 2000*, pages 95–106. IEEE, 2000.

[11] J. Dieudonné. Sur les espaces de Köthe. *Jour. dAnalyse Math.*, 1(1):81–115, 1951.

[12] N. Dunford, J. T. Schwartz, W. G Bade, and R. G. Bartle. *Linear operators I.* Wiley-interscience New York, 1971.

[13] N. Foster, D. Kozen, K. Mamouras, M. Reitblatt, and A. Silva. Probabilistic netkat. In *ESOP 2016*, pages 282–309. Springer, 2016.

[14] 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.

[15] A. S. Kechris. *Classical descriptive set theory*, volume 156 of *Graduate Text in Mathematics*. Springer, 1995.

[16] D. Kozen. Semantics of probabilistic programs. *Journal of computer and system sciences*, 22(3):328–350, 1981.

[17] D. Kozen. Kolmogorov extension, martingale convergence, and compositionality of processes. In *LiCS 2016*, pages 692–699, 2016.

[18] P. 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.

[19] P. Selinger. Dagger compact closed categories and completely positive maps. *Electronic Notes in Theoretical computer science*, 170:139–163, 2007.

[20] S. Smolka, P. Kumar, N. Foster, D. Kozen, and A. Silva. Cantor meets Scott: Domain-theoretic foundations for probabilistic network programming. In *POPL 2017*, 2017.

[21] S. Staton. Commutative semantics for probabilistic programming. In *European Symposium on Programming*, pages 855–879. Springer, 2017.

[22] S. Staton, H. Yang, F. Wood, C. Heunen, and O. Kammar. Semantics for probabilistic programming: higher-order functions, continuous distributions, and soft constraints. In *LICS, 2016*, pages 525–534, 2016.

[23] D. Williams. *Probability with martingales.* Cambridge University Press, 1991.

[24] F. Wood, Jan W. van de Meent, and V. Mansinghka. A new approach to probabilistic programming inference. In *International conference on Artificial Intelligence and Statistics*, pages 1024–1032, 2014.

[25] A. C. Zaanen. *Introduction to operator theory in Riesz spaces.* Springer, 2012.