Luke Ong's background and current research interests

Semantics of programming languages

Verification

Probabilistic programming in statistical ML (from mid-2017)

1. Semantics of probabilistic programming languages (PPL)
   - modelling PPCF / Prob. Idealised Algol in game semantics and ω-quasi Borel spaces, aiming at full abstraction
2. S-finite measures and metaprogramming for PPL pragmatics
   - classical theorems (e.g. Radon-Nikodym / Lebesgue decomposition, disintegration) for s-finite measures / kernels, and applications to metaprogramming
3. Martingales and static analysis of probabilistic programs
   - super / martingales for analysing liveness / invariance of (higher-order functional) PPL

S-finite measures for probabilistic (meta)programming

The semantic basis of probabilistic programming is s-finite measure theory: fully-fledged PPL computes s-finite kernels. (Staton ESOP17)

DEF. Let μ be a measure on measurable space (X, Σ_X).

- μ is σ-finite if $X = \biguplus_{i \in \omega} X_i$ with each $X_i \in \Sigma_X$ and $\mu(X_i) < \infty$.
- μ is s-finite if $\mu = \sum_{i \in \omega} \mu_i$, and each $\mu_i(X) < \infty$.

σ-finite $\subset$ s-finite

Standard results for infinite measures assume σ-finite measures; e.g.

1. Radon-Nikodym Theorem
2. Lebesgue Decomposition Theorem
3. Disintegration Theorem

Matthijs Vákár and I have extended the above to s-finite measures.

Characterising s-finite measures

Intuition: "bad ∞" is ∞ concentrated at a point.

- σ-finiteness only admits "good ∞"
- s-finiteness can admit "bad ∞", but only countably many.

Examples (infinite measures)

1. The Lebesgue measure, Leb, is σ-finite.
2. The ∞-measure on the point 1 is s-finite, but not σ-finite.
3. $\#_S$ on uncountable standard Borel space S is not s-finite.

DEF. $U \in \Sigma_X$ is an ∞-set w.r.t. measure μ if (i) $\mu(U) = \infty$, and (ii) for all $V \in \Sigma_U$, $\mu(V) = 0$ or ∞.

IDEA: Presence of ∞-sets distinguishes s-finite from σ-finite measures.

Theorem. Let μ be an s-finite measure on X. Then there exists $U \in \Sigma_X$ such that μ is σ-finite on $X \setminus U$, and U is an ∞-set or a null-set.

Definition: disintegration of a measure

Given measure spaces (X, μ) and (Y, ν), and measurable $T : X \to Y$.

DEF. A (T, ν)-disintegration of μ is a family $\{\mu_y\}_{y \in Y}$ of measures on X and a ν-null set $\mathsf{N} \in \Sigma_Y$ s.t.

i. Regularity: $(y, U) \mapsto \mu_y(U)$ is a kernel from Y to X;
ii. Concentration: μ_y concentrates on $\{T = y\}$, for all $y \in Y \setminus \mathsf{N}$
iii. Weighted average: $\mu(V) = \int_Y \nu(\mathrm{d}y)\, \mu_y(V)$, for all $V \in \Sigma_X$.

Conditional distribution: μ_y "is" $\mu(- \mid T = y)$.

The standard Disintegration Theorem for σ-finite measures asserts something weaker: disintegration $\mu_y(U)$ is a measure for every fixed y.

Maharam's 1950 Problem: Theorem (Back et al. 2015). *If CH holds, $\mu_-(-)$ cannot be a kernel.*

A disintegration theorem for s-finite measures

Given measure spaces (X, μ) and (Y, ν), and measurable $T : X \to Y$.

Existence

Assume

(i) μ and ν are s-finite

(ii) $T_* \mu \ll \nu$

(iii) for all ν-∞-sets U, $T^{-1}(U)$ is a μ-∞-set or a μ-null-set.

Then there exists a (T, ν)-disintegration of μ, $\{\mu_y\}_{y \in Y}$, which is an s-finite kernel.

Uniqueness

If ν is s-finite, then the (T, ν)-disintegration of μ (*qua* s-finite kernel) is unique up to ν-∞-equivalence.

Problem

Exact Bayesian inf. by symbolic disintegration (Shan & Ramsey POPL17)

Conjecture. Let ρ be an s-finite measure on $X \times Y$ and μ be s-finite measure on X, satisfying condition (C). Then there exists an s-finite kernel $k : X \rightsquigarrow Y$ such that $\rho = \mu \otimes k$. Further the kernel is unique up to μ-∞-equivalence.

Desiderata:

1. Higher order & definability. Take $\mathcal{L}$ an idealised higher-order PPL; e.g. core Hakaru$^{\rightarrow}$(?). Extend ρ and μ to $\mathcal{L}$-definable measures; prove: k is $\mathcal{L}$-definable (Staton ESOP17).

2. Constructiveness / relativised computability. Show that k is $\mathcal{L}$-definable via partial evaluation (type-directed / continuation-based); prove correctness via synthetic measure theory.

3. Compositionality / "parametricity law". Replace ρ and μ by s-finite kernels (appropriately typed).