

Higher order programming with probabilistic effects: A model of stochastic memoization and name generation

Younesse Kaddar Sam Staton
Department of Computer Science, University of Oxford, UK

This abstract is about probabilistic programming, which is a method of programming statistical models and machine learning models. By combining higher-order functions, we can specify increasingly complex models. This abstract focuses on stochastic memoization, a higher-order method that is simple and useful in practice, but semantically elusive, particularly regarding dataflow transformations.

Deterministic versus stochastic memoization. Deterministic memoization, for a function f , is the process of storing the result when f is applied to an argument, so that we can reuse it later when the same call is made [Mic68]. Since we only need to compute results once, this leads to a speed-up of the program, but it does not change its semantics. However, in the presence of probabilistic effects, memoizing is no longer just an optimization technique. It does change the semantics [Roy+08; Sta21; Woo+09], enabling us to define infinite random sequences, which are of paramount importance in probability and statistics.

Typed stochastic memoization with monads. Suppose we have a monad for probabilistic effects, Prob . *Stochastic memoization* is a higher-order function with probabilistic effects, of type $\text{mem} :: (a \rightarrow \text{Prob } b) \rightarrow \text{Prob } (a \rightarrow b)$. For example, suppose we start with a function which randomly picks a number in $[0, 1]$ every time it is called, $\text{const } (\text{uniform } 0 \ 1) :: \mathbb{R} \rightarrow \text{Prob } \mathbb{R}$. Then the memoized function randomly assigns a number to every input number, $\text{mem } (\text{const } (\text{uniform } 0 \ 1)) :: \text{Prob } (\mathbb{R} \rightarrow \mathbb{R})$. This is sometimes called white noise. White noise is an important first example; we give a more substantial motivating example based on random graphs in Section 1. In practice, memoization is crucial in Church [GS14] and WebPPL [GS14], and in this typed form with monads in our library LazyPPL [Sta+].

Data flow properties. Stochastic memoization is easy to implement using state. (When a is enumerable, in LazyPPL we can also use laziness and tries, following [Hin].). Unlike a fully stateful language, however, stochastic memoization is still compatible with commutativity / data flow program transformations:

$$x \leftarrow t ; y \leftarrow u \quad = \quad y \leftarrow u ; x \leftarrow t \quad \text{where } x \notin \text{freevars}(u), y \notin \text{freevars}(t)$$

These transformations are very useful in program optimization and inference algorithms. On the foundational side, data flow is a fundamental concept that corresponds to monoidal categories. The challenge addressed in this work is to validate these transformations.

Challenges for probability theory. On the semantic side, these transformations are not trivial. Informally, stochastic memoization appears related to Kolmogorov's extension theorem, which relates probability measures on infinite product spaces to probability measures on the projections from the product. Arguably, then, stochastic memoization validates dataflow transformations because it is not *intrinsically* stateful, rather, it is linked to a fundamental part of pure probability theory. However, traditional probability theory does not actually support higher-order functions [Aum61], and higher-order probability is a burgeoning field [CJ19; Fri20; Heu+17; Koc11; Ste21]. Existing models of higher-order probability do not support stochastic memoization, and this is the contribution of our work: a first model of stochastic memoization with a non-enumerable type and validating the dataflow transformations.

1 Stochastic Memoization

Stochastic memoization is, on the practical side, a convenient way to implement infinite sequences of random variables, point processes, clustering and, more generally, nonparametric Bayesian models in probabilistic programming [Roy+08;

Submitted to:
HOPE 2022

[Woo+09]. On the theoretical side, stochastic memoization may become of utmost importance to obtain general representation theorems for exchangeable data types, similar to de Finetti's theorem [Fin37] (for exchangeable sequences) or the Aldous-Hoover theorem [Ald81; Hoo79] (for exchangeable arrays). Consider the following example.

Exchangeable random graphs. Random arrays $(X_{i,j})_{i,j}$ of Boolean random variables (where i and j range over countable collections of potentially related objects) are a useful data analysis tool to model relational data, which are observations of binary relationships between collections of objects, e.g. graph social networks, world wide web, biochemical pathways, etc. When we have a single (countable) collection of objects that we wish to compare pairwise (i.e. i and j range over the same countable index), the array $(X_{i,j})_{i,j}$ can be seen as the adjacency matrix a random graph. Such random adjacency matrices $(X_{i,j})_{i,j}$ are said to be *exchangeable* when their joint distribution is invariant under relabeling the nodes of the corresponding graph. If the random graph is assumed to be simple (i.e. $(X_{i,j})_{i,j}$ is symmetric and has zero diagonal), the celebrated Aldous-Hoover theorem states that every corresponding exchangeable adjacency matrix can be parametrized by a (random) measurable function $G: [0, 1]^2 \rightarrow [0, 1]$ called a *graphon*. If u_i, u_j are iid random variables corresponding to two nodes i, j in the random simple graph, $G(u_i, u_j)$ gives the probability that there is an edge between i and j . As a result, Bayesian models modeling exchangeable simple graphs are entirely determined (in distribution) by a prior on the space of graphons.

Consider a typeclass `RandomGraph`, corresponding to an abstract data type whose interface allows us to draw vertices at random (with `newVertex`) from a random graph with a countable set $V = \mathbf{v}$ of vertices, and inspect the presence of edges between vertices (with `isEdge`). In the spirit of the Aldous-Hoover representation theorem, for every graphon $G = \text{graphon}: [0, 1]^2 \rightarrow [0, 1]$, we can write an implementation for this interface in our library LazyPPL [Sta+]:

```
class RandomGraph v where
  type Graph v
  newVertex :: Prob v
  newGraph :: Prob (Graph v)
  isEdge :: Graph v → (v, v) → Bool

instance Graphon => RandomGraph ℝ where
  type Graph Double = (ℝ, ℝ) → Bool
  newVertex = uniform

  -- return a randomly sampled function '(ℝ, ℝ) → Bool'
  newGraph = mem $ bernoulli . graphon

class Graphon where
  graphon :: (ℝ, ℝ) → ℝ

  isEdge g (x, y) = g (x, y)
```

assuming we have a memoization function `mem :: (a → Prob b) → Prob (a → b)`. The idea is that, once an edge between x and y has been sampled (with probability `graphon (x, y)`), its presence (or absence) remains unchanged in the rest of the program, hence the need to memoize the result. Staton showed that such graphon-based implementations are the only ones up to contextual equivalence, provided that they satisfy the dataflow property:

Theorem 1.1 (Staton's Aldous-Hoover representation [Sta20]). *An implementation for the interface `RandomGraph` satisfies the dataflow property (i.e. program lines can be reordered (commutativity) and discarded (affineness) as long as the dataflow is preserved) iff it is observationally equivalent to a graphon implementation.*

The proof involves building a categorical model based on the *Rado graph* (the countable graph that embeds every at most countable graph), paving the way for a very general approach to tackle similar representation problems.

2 Semantics

Stochastic functions are interpreted as probabilistic kernels $f: X \rightarrow PY$, where P is a probability monad, in a suitable cartesian closed category (to model higher-order functions), e.g. the category of quasi-Borel spaces ([Heu+17]). Stochastic memoization (simply referred to as 'memoization' henceforth) is then internally expressed as a morphism $\text{mem}_{X,Y}: (PY)^X \rightarrow P(Y^X)$ converting a probabilistic kernel (which associates, for every given input in X , a random output in Y) into a random function $X \rightarrow Y$ (randomly choosing all the outputs for all the possible inputs at once). If f is written as a lambda-abstraction $\lambda x. u: X \rightarrow PY$, $\text{mem}_{X,Y}(f)$ will be denoted by $\lambda_{\mathfrak{N}} x. u$.

$\begin{array}{l} f \leftarrow \lambda_{\mathbb{A}} x. u \\ f \ n \end{array}$	$\underline{\text{one sample}}$	$u[n/x]$	$\underline{\text{several samples}}$	$\begin{array}{l} f \leftarrow \lambda_{\mathbb{A}} x. u \\ v_1 \leftarrow f \ n \\ w \leftarrow f \ m \\ v_2 \leftarrow f \ n \\ \text{return } (v_1, w, v_2) \end{array}$	$\begin{array}{l} v \leftarrow u[n/x] \\ w \leftarrow u[m/x] \\ \text{return } (v, w, v) \end{array}$
--	---------------------------------	----------	--------------------------------------	---	---

For finite types X , memoization is a matter of simply sampling a value of $f(x)$ for all inhabitants of $x \in X$, and returning the assignment as a finite mapping. For countable types, we can sometimes successfully take advantage of the host language's laziness to circumvent the issue of manipulating potentially infinite structures (this is done in LazyPPL). But what about uncountable types X ? (We cannot just range over an uncountable space for all $x \in X$ anymore.) The category of quasi-Borel spaces is unfortunately known not to support memoization in general [Sta21]. The first implementation of 'memoize' that comes to mind uses state to store previously seen values. But usually, using memory compromises the dataflow property (the state monad is not commutative). However, we conjecture that 'memoize' is a special kind of stateful operation that do preserve the dataflow property:

Proposition 2.1. *Stochastic memoization still admits the dataflow property (even if one resorts to hidden state).*

To give a semantics to stochastic memoization for Boolean-valued functions (to start small), the idea is that we may wish to have the following interface:

```
new_atom ::  $\mathbb{A}$  -- Atoms (randomly generated fresh names)
new_function ::  $\mathbb{F}$  -- Function labels : type to be thought of as  $\mathbb{A} \rightarrow \text{Bool}$ 
(@) :: ( $\mathbb{F}, \mathbb{A}$ )  $\rightarrow$  Bool -- Application operator making every function memoized: type of a bipartite graph
```

where every function from a set of atoms \mathbb{A} (each one of which can be randomly generated) to **Bool** would be defunctionalized and viewed as an inhabitant of a type \mathbb{F} (which would then be thought of as $\mathbb{A} \rightarrow \text{Bool}$). Applying a function to an argument and memoizing the result would then be made possible by an 'apply' operator $(@) :: \mathbb{F} \times \mathbb{A} \rightarrow \text{Bool}$. But requiring that the results be memoized is precisely saying that $(@) :: \mathbb{F} \times \mathbb{A} \rightarrow \text{Bool}$ ought to be seen as the 'edge' relation of a bipartite graph (bigraph, for short) with set \mathbb{F} of left nodes and \mathbb{A} of right nodes – whose edges are such that their presence (or absence) remain unchanged after being sampled, exactly like `isEdge` in `RandomGraph`. Analogously to the Rado topos setting, this suggests, on the denotational semantics side, that we are looking for a topos where a random countable bigraph would play the role of the Rado graph in the Rado topos.

3 Toy language for stochastic memoization and name generation

One way to prove [proposition 2.1](#) is to exhibit a denotational model, which we attempt in this section, in a restricted setting. We consider a small simply typed language to shed light on three features that we model semantically:

- **name generation:** we can generate fresh names – referred to as *atomic* names or *atoms*, in the sense of Pitt's nominal set theory [Pit13] – with constructs such as `let x = fresh() in \dots` .
- **basic probabilistic effects:** for illustrative purposes, the only distribution we consider, as a first step, is the Bernoulli distribution with bias $p = 1/2$. Constructs like `let b = flip() in \dots` amount to flipping a fair coin and storing its result in a variable b .
- **stochastic memoization:** if a stochastic function f – memoized with the new $\lambda_{\mathbb{A}}$ operator – is called twice on the same argument, it should return the same result.

We have the following base types: `bool` (booleans), \mathbb{A} (atomic names), and \mathbb{F} (intended for memoized functions $\mathbb{A} \rightarrow \text{bool}$). For the sake of simplicity, we do not have arbitrary function types.

$$A, B ::= \text{bool} \mid \mathbb{A} \mid \mathbb{F} \mid A \times B$$

In fine-grained call-by-value fashion [Lev06], there are two kinds of judgments: typed values, and typed computations.

Values:

$$\frac{}{\Gamma, x : A \Vdash x : A} \quad \frac{\Gamma \Vdash v : A \quad \Gamma \Vdash w : B}{\Gamma \Vdash (v, w) : A \times B} \quad \frac{}{\Gamma \Vdash \text{true} : \text{bool}} \quad \frac{}{\Gamma \Vdash \text{false} : \text{bool}}$$

Computations:

$$\frac{\Gamma \Vdash v : A}{\Gamma \Vdash \text{return}(v) : A} \quad \frac{\Gamma \Vdash u : A \quad \Gamma, x : A \Vdash t : B}{\Gamma \Vdash \text{let val } x \leftarrow u \text{ in } t : B}$$

Matching:

$$\frac{\Gamma \Vdash v : \text{bool} \quad \Gamma \Vdash u : A \quad \Gamma \Vdash t : A}{\Gamma \Vdash \text{if } v \text{ then } u \text{ else } t : A} \quad \frac{\Gamma \Vdash v : A \times B \quad \Gamma, x : A, y : B \Vdash t : C}{\Gamma \Vdash \text{match } v \text{ as } (x, y) \text{ in } t : C}$$

Language-specific commands:

$$\frac{}{\Gamma \Vdash \text{flip}() : \text{bool}} \quad \frac{}{\Gamma \Vdash \text{fresh}() : \mathbb{A}} \quad \frac{\Gamma \Vdash v : \mathbb{A} \quad \Gamma \Vdash w : \mathbb{A}}{\Gamma \Vdash (v = w) : \text{bool}}$$

$$\frac{\Gamma \Vdash v : \mathbb{F} \quad \Gamma \Vdash w : \mathbb{A}}{\Gamma \Vdash (v @ w) : \text{bool}} \quad \frac{\Gamma, x : \mathbb{A} \Vdash u : \text{bool}}{\Gamma \Vdash \lambda_{\mathbb{A}} x. u : \mathbb{F}}$$

We work in the (cartesian closed) category of covariant presheaves on the category \mathbf{BiGrph}_{emb} of finite bipartite graphs (henceforth called *bigraphs*) and embeddings (that do not add or remove edges). For a bigraph g , we denote by g_L (resp. g_R) and E^g its set of left (resp. right) nodes and its edge relation.

The denotation of basic types is given by:

$$\llbracket \mathbb{F} \rrbracket = \mathbf{BiGrph}_{emb}(\circ, -) \quad \llbracket \mathbb{A} \rrbracket = \mathbf{BiGrph}_{emb}(\bullet, -)$$

where \circ and \bullet are the one-vertex left and right graphs respectively. The denotation of the type of booleans is the constant presheaf $2 \cong 1 + 1$, as usual. For a bigraph g and a presheaf $X = \llbracket \mathcal{X} \rrbracket$, $X(g)$ is thought of as the set of generative models/programs of type \mathcal{X} that may use the bigraph g , in the following sense: probabilistic function (that we want to memoize) and atom labels are stored as left and right nodes respectively. The presence (resp. absence) of an edge between a given left and right node memoizes the fact that a probabilistic call of the corresponding function on the corresponding atom has resulted in true (resp. false). For every embedding $\iota : g \hookrightarrow g'$, the function $X\iota : X(g) \rightarrow X(g')$ models substitution in the programs in $X(g)$ according to ι .

4 Probabilistic local state monad

Inspired from Plotkin and Power's local state monad [PP02] (which was defined on the covariant presheaf category $[\mathbf{Inj}, \mathbf{Set}]$, where \mathbf{Inj} is the category of finite sets and injections), we model probabilistic and name generation effects by a new monad that we name 'probabilistic local state monad'. In the following, $X, Y, Z : \mathbf{BiGrph}_{emb} \rightarrow \mathbf{Set}$ denote presheaves, $g = (g_L, g_R, E^g)$, $g', h, h' \in \mathbf{BiGrph}_{emb}$ bigraphs, and $\iota, \iota' : g \hookrightarrow g'$ bigraph embeddings. We will omit subscripts when they are clear from the context.

Definition 4.1 (Probabilistic local state monad). For all covariant presheaf $X : \mathbf{BiGrph}_{emb} \rightarrow \mathbf{Set}$ and bigraph $g \in \mathbf{BiGrph}_{emb}$:

$$T(X)(g) := \left(P_{\mathbb{F}} \int^{g \hookrightarrow h} X(h) \times [0, 1]^{(h-g)_L} \right)^{[0, 1]^{g_L}}$$

where $P_{\mathbb{F}}$ is the finite distribution monad.

The assignment T is similar to the read-only local state monad, except that any fresh node can be initialized. Every $\lambda \in [0, 1]^{g_L}$ is thought of as the probability of the corresponding function/left node being true on a new fresh atom. We will refer to such a λ as a *state of biases*. The coend takes care of garbage collection.

Notation 4.2. Equivalence classes in $\int^{g \hookrightarrow h} X(h) \times [0, 1]^{(h-g)_L}$ are written $[x_h, \lambda^h]_g$. We use Dirac's bra-ket notation $[x_h, \lambda^h]_g \rangle_h$ to denote a formal column vector of equivalence classes ranging over a finite set of h 's. As such, a formal convex sum $\sum_i p_i [x_{h_i}, \lambda^{h_i}]_g \in P_{\mathbb{F}} \int^{g \hookrightarrow h} X(h) \times [0, 1]^{(h-g)_L}$ will be concisely denoted by $\langle \vec{p} \mid [x_h, \lambda^h]_g \rangle_h$.

Definition 4.3 (Action of $T(X)$ on morphisms).

$$T(X)(g \xrightarrow{\iota} g') = \begin{cases} \left(P_{\mathbb{F}} \int^{g \hookrightarrow h} X(h) \times [0, 1]^{(h-g)_L} \right)^{[0, 1]^{g_L}} \longrightarrow \left(P_{\mathbb{F}} \int^{g' \hookrightarrow h'} X(h') \times [0, 1]^{(h'-g')_L} \right)^{[0, 1]^{g'_L}} \\ \vartheta \longmapsto \lambda' \mapsto \text{let } \vartheta(\lambda' \iota_L) = \langle \vec{p} \mid [x_h, \lambda^h]_g \rangle_h \text{ in } \langle \vec{p} \mid [X(h \hookrightarrow h \amalg_g g')(x_h), \lambda^h]_{g'} \rangle_{h'} \end{cases}$$

Theorem 4.4. *The construction T is functorial and can be endowed with the structure of a monad.*

We now use this monad to give a denotational semantics to our language.

5 Categorical semantics

In our language, the denotational interpretation of values, computations (return and let binding), and matching (elimination of bool's and product types) is standard. We interpret computation judgements $\Gamma \Vdash t : A$ as morphisms $[\Gamma] \rightarrow T([A])$, by induction on the structure of typing derivations. The context Γ is built of bool's, \mathbb{A} and \mathbb{F} and products. Therefore, $[\Gamma]$ is isomorphic to an object of the form $2^k \times \mathbf{BiGrph}_{emb}(\circ, -)^\ell \times \mathbf{BiGrph}_{emb}(\bullet, -)^m$.

Definition 5.1. For every bigraph g , we denote by R_g (resp. L_g) the set of bigraphs $h \in g/\mathbf{BiGrph}_{emb}$ having one more right (resp. left) node than g , and that are the same otherwise.

$$\begin{aligned} R_g &:= \{ h \in \mathbf{BiGrph}_{emb} \mid h_L = g_L, g_R \subseteq h_R \text{ and } |h_R| = |g_R| + 1 \} \\ L_g &:= \{ h \in \mathbf{BiGrph}_{emb} \mid h_R = g_R, g_L \subseteq h_L \text{ and } |h_L| = |g_L| + 1 \} \end{aligned}$$

Denotation of $\Gamma \Vdash \text{fresh}() : \mathbb{A}$

The map $[\text{fresh}]_g : [\Gamma](g) \rightarrow T([\mathbb{A}])(g)$ randomly chooses connections to each left node according to the state of biases, and makes a fresh right node with those connections.

$$[\text{fresh}]_g : \begin{cases} 2^k \times \mathbf{BiGrph}_{emb}(\circ, g)^\ell \times \mathbf{BiGrph}_{emb}(\bullet, g)^m \longrightarrow P_{\mathbb{F}}(g_R + 2g_L)^{[0, 1]^{g_L}} \\ -, -, - \mapsto \lambda \mapsto \left\langle \frac{1}{Z} \prod_{f \in g_L} \lambda(f)^{E^h(f, a_h(\bullet))} (1 - \lambda(f))^{1 - E^h(f, a_h(\bullet))} \mid \left[\underbrace{\bullet}_{\cong (h-g)_R} \xrightarrow{a_h} h, ! \right]_g \right\rangle_{h \in R_g} \end{cases}$$

where Z is a normalization constant.

Denotation of $\Gamma \Vdash \lambda_{\mathbb{A}} x. u : \mathbb{F}$

As $\lambda_{\mathbb{A}}$ -abstractions are formed based on computation judgements of the form $\Gamma, x : \mathbb{A} \Vdash u : \text{bool}$, we first note that

$$T([\text{bool}])g \cong P_{\mathbb{F}}(2)^{[0, 1]^{g_L}} \cong [0, 1]^{[0, 1]^{g_L}}$$

Also, we can decompose the extra variable x in the environment $\Gamma, x : \mathbb{A}$, the denotation of which is of the form $[\Gamma, x : \mathbb{A}](g) = 2^k \times \mathbf{BiGrph}_{emb}(\circ, g)^\ell \times \mathbf{BiGrph}_{emb}(\bullet, g)^m \times \mathbf{BiGrph}_{emb}(\bullet, g)$ for a bigraph $g \in \mathbf{BiGrph}_{emb}$.

Now, the extra part x is a right node, and its valuation will either be a node already in the graph described in the rest of the environment, or a new one with particular edges to the rest of the environment. The argument u can test (if it wants) what kind of node x is, before returning a probability.

As a result, for $\llbracket u \rrbracket_g : 2^k \times \mathbf{BiGrph}_{emb}(\circ, g)^\ell \times \mathbf{BiGrph}_{emb}(\bullet, g)^m \times \mathbf{BiGrph}_{emb}(\bullet, g) \rightarrow [0, 1]^{[0, 1]^{g_L}}$, the denotation $\llbracket u \rrbracket$ gives us the edge probability of the left node that we need to generate, both to the existing right nodes, and to any future right nodes (which needs to be remembered). This can be formalized into a natural transformation $\llbracket \lambda_{\mathfrak{P}} x. u \rrbracket : \llbracket \Gamma \rrbracket \rightarrow T(\llbracket \mathbb{F} \rrbracket)$.

$$\llbracket \lambda_{\mathfrak{P}} x. u \rrbracket_g : \begin{cases} 2^k \times \mathbf{BiGrph}_{emb}(\circ, g)^\ell \times \mathbf{BiGrph}_{emb}(\bullet, g)^m \rightarrow P_{\mathfrak{f}}(g_L + 2^{g_R} \times [0, 1]^{[0, 1]^{g_L}} \\ b^k, (\circ \xrightarrow{\kappa_i} g)_i, \mapsto \lambda \mapsto \left\langle \frac{1}{Z} \prod_{a \in g_R} p_a^{E^h(f_h(\circ), a)} (1 - p_a)^{1 - E^h(f_h(\circ), a)} \left| \left[\underbrace{\circ \xrightarrow{f_h} h}_{\cong (h-g)_L}, - \mapsto \tilde{p} \right]_g \right\rangle_{h \in L_g} \end{cases}$$

where Z is a normalization constant, and

- for every $a \in g_R$, $p_a := \llbracket u \rrbracket_g(b^k, (\circ \xrightarrow{\kappa_i} g)_i, (\bullet \xrightarrow{\tau_j} g)_j, \bullet \xrightarrow{a} g, \lambda)$
- $\tilde{p} := \llbracket u \rrbracket_g(b^k, (\circ \xrightarrow{\kappa_i} g \xrightarrow{\iota_1} g + \bullet)_i, (\bullet \xrightarrow{\tau_j} g \xrightarrow{\iota_1} g + \bullet)_j, \bullet \xrightarrow{\iota_2} g + \bullet, \lambda)$ where ι_1, ι_2 are the coprojections.

Finally, we show that the monad T is commutative and affine, which yields an abstract model of probability (as defined in [Koc11]) and implies that the language enjoys the dataflow property:

Theorem 5.2. *The probabilistic local state monad T is strong commutative and affine.*

References

- [Ald81] David J. Aldous. “Representations for Partially Exchangeable Arrays of Random Variables”. In: *Journal of Multivariate Analysis* 11.4 (Dec. 1, 1981), pages 581–598. ISSN: 0047-259X. DOI: [10.1016/0047-259x\(81\)90099-3](https://doi.org/10.1016/0047-259x(81)90099-3). URL: <http://www.sciencedirect.com/science/article/pii/0047259X81900993> (visited on 11/25/2020).
- [Aum61] R. Aumann. “Borel Structures for Function Spaces”. In: (1961). DOI: [10.1215/IJM/1255631584](https://doi.org/10.1215/IJM/1255631584).
- [CJ19] Kenta Cho and Bart Jacobs. “Disintegration and Bayesian Inversion via String Diagrams”. In: *Mathematical Structures in Computer Science* 29.7 (Aug. 2019), pages 938–971. ISSN: 0960-1295, 1469-8072. DOI: [10.1017/S0960129518000488](https://doi.org/10.1017/S0960129518000488). arXiv: [1709.00322](https://arxiv.org/abs/1709.00322). URL: <http://arxiv.org/abs/1709.00322> (visited on 01/25/2021).
- [Fin37] Bruno De Finetti. “La prévision : ses lois logiques, ses sources subjectives”. In: *Annales de l’institut Henri Poincaré* 7 (1937), page 69. URL: http://www.numdam.org/item/?id=AIHP_1937__7_1_1_0.
- [Fri20] Tobias Fritz. “A Synthetic Approach to Markov Kernels, Conditional Independence and Theorems on Sufficient Statistics”. In: *Advances in Mathematics* 370 (Aug. 2020), page 107239. ISSN: 00018708. DOI: [10.1016/j.aim.2020.107239](https://doi.org/10.1016/j.aim.2020.107239). arXiv: [1908.07021](https://arxiv.org/abs/1908.07021). URL: <http://arxiv.org/abs/1908.07021> (visited on 01/31/2021).
- [GS14] Noah D Goodman and Andreas Stuhlmüller. *The Design and Implementation of Probabilistic Programming Languages*. 2014. URL: <http://dippl.org/> (visited on 08/17/2021).
- [Heu+17] Chris Heunen et al. “A Convenient Category for Higher-Order Probability Theory”. Apr. 18, 2017. arXiv: [1701.02547](https://arxiv.org/abs/1701.02547) [cs, math]. URL: <http://arxiv.org/abs/1701.02547> (visited on 02/11/2020).
- [Hin] Ralf Hinze. *Generalizing Generalized Tries*.

- [Hoo79] Douglas N. Hoover. *Relations on Probability Spaces and Arrays of Random Variables*. Institute for Advanced Study, Princeton, 1979. URL: <http://www.stat.berkeley.edu/~aldous/Research/hover.pdf>.
- [Koc11] Anders Kock. “Commutative Monads as a Theory of Distributions”. Aug. 30, 2011. arXiv: [1108.5952](https://arxiv.org/abs/1108.5952) [math]. URL: <http://arxiv.org/abs/1108.5952> (visited on 02/07/2021).
- [Lev06] Paul Blain Levy. “Call-by-Push-Value: Decomposing Call-by-Value and Call-by-Name”. In: *Higher-Order and Symbolic Computation* 19.4 (Dec. 2006), pages 377–414. ISSN: 1388-3690, 1573-0557. DOI: [10/fwb7vh](https://doi.org/10.1007/s10990-006-0480-6). URL: <http://link.springer.com/10.1007/s10990-006-0480-6> (visited on 11/12/2021).
- [Mic68] Donald Michie. ““Memo” Functions and Machine Learning”. In: (1968), page 4.
- [Pit13] A. M. Pitts. *Nominal Sets: Names and Symmetry in Computer Science*. Cambridge Tracts in Theoretical Computer Science 57. Cambridge ; New York: Cambridge University Press, 2013. 276 pages. ISBN: 978-1-107-01778-8.
- [PP02] Gordon Plotkin and John Power. “Notions of Computation Determine Monads”. In: *Foundations of Software Science and Computation Structures*. Edited by Mogens Nielsen and Uffe Engberg. Redacted by Gerhard Goos, Juris Hartmanis and Jan van Leeuwen. Volume 2303. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pages 342–356. ISBN: 978-3-540-43366-8 978-3-540-45931-6. DOI: [10.1007/3-540-45931-6_24](https://doi.org/10.1007/3-540-45931-6_24). URL: http://link.springer.com/10.1007/3-540-45931-6_24 (visited on 05/24/2021).
- [Roy+08] D. Roy et al. “A Stochastic Programming Perspective on Nonparametric Bayes”. In: 2008. URL: <https://www.semanticscholar.org/paper/A-stochastic-programming-perspective-on-Bayes-Roy-Mansinghka/946ed04f705a2a28539a8af1f5e9ccdedd7fab2> (visited on 01/21/2022).
- [Sta20] Sam Staton. “Categorical Models of Probability with Symmetries”. Categorical Probability and Statistics. June 2020. URL: http://perimeterinstitute.ca/personal/tfritz/2019/cps_workshop/slides/staton.pdf.
- [Sta21] Sam Staton. “Some Formal Structures in Probability”. 2021. URL: <http://www.cs.ox.ac.uk/people/samuel.staton/2021fscd.pdf>.
- [Sta+] Sam Staton et al. *LazyPPL*. URL: <https://lazyppp.bitbucket.io/> (visited on 04/04/2022).
- [Ste21] Dario Maximilian Stein. “Structural Foundations for Probabilistic Programming Languages”. University of Oxford, 2021. 221 pages.
- [Woo+09] Frank Wood et al. “A Stochastic Memoizer for Sequence Data”. In: *Proceedings of the 26th Annual International Conference on Machine Learning - ICML '09*. The 26th Annual International Conference. Montreal, Quebec, Canada: ACM Press, 2009, pages 1–8. ISBN: 978-1-60558-516-1. DOI: [10/fg8z4q](https://doi.org/10.1145/1553374.1553518). URL: <http://portal.acm.org/citation.cfm?doid=1553374.1553518> (visited on 01/21/2022).