

# Coinductive Trees for Exact Inference of Probabilistic Programs

ALEXANDER BAGNALL, Ohio University, United States

GORDON STEWART, Ohio University, United States

ANINDYA BANERJEE, IMDEA Software Institute, Spain

We propose a coinductive variant of Knuth-Yao trees (variously, Discrete Distribution Generating trees) as an intermediate representation supporting exact inference for probabilistic programs which may contain loops with probabilistic termination conditions. We provide a prototype implementation of a probabilistic programming language, Zar, written in Haskell, exemplifying both Knuth-Yao trees as well as exact inference on the trees.

## 1 INTRODUCTION

Probabilistic programmers define probabilistic models by writing conventional imperative programs extended with primitives for random sampling and conditioning. Inference is the problem of computing an explicit representation of the probability distribution implicitly specified by a probabilistic program in order to support queries such as “What is the probability of event  $e$ ?”, where  $e$  is a predicate on program states or a possible return value. One can infer a probabilistic program’s posterior distribution either approximately (e.g., by sampling) or exactly using symbolic methods (e.g., compilation to binary decision diagrams, or BDDs). When computationally viable, exact inference is preferable to approximate techniques because it is deterministic, trustworthy, and does not propagate errors to subsequent analyses.

Symbolic methods based on decision diagrams like BDDs have seen most currency within the probabilistic model checking community (cf. [Miner and Parker 2004] for a survey). In this work, we revisit the problem of exact inference for probabilistic programs and propose a new intermediate representation based in Knuth-Yao (KY) trees (alternatively, Discrete Distribution Generating trees [Knuth and Yao 1976]). In contrast to existing recent work such as that of Holtzen *et al.* [Holtzen *et al.* 2019] who compile only loop-free programs to BDDs, our KY tree representation enables exact inference even of programs with almost-surely terminating loops. We have a prototype implementation of our technique, Zar<sup>1</sup>, that uses our Knuth-Yao IR to do exact inference of programs in a conventional probabilistic programming language.

## 2 KNUTH-YAO IR FOR ALMOST-SURE TERMINATION

As motivation, consider the probabilistic programs of Listings 1 and 2, the first of which terminates in all executions (it contains no loops) while the second, which simulates a fair coin using a biased one, terminates only with probability 1. If we understood the program of Listing 1 as a random process mapping inputs bits to outputs, we can therefore place an upper bound on the number of bits required for the program to produce an output  $y$  (in this

<sup>1</sup><https://github.com/OUPL/Zar>

Authors’ addresses: Alexander Bagnall, Ohio University, United States, ab667712@ohio.edu; Gordon Stewart, Ohio University, United States, gstewart@ohio.edu; Anindya Banerjee, IMDEA Software Institute, Spain, anindya.banerjee@imdea.org.

```
x <- flip(1/2)
y <- if x then flip(1/2) else flip(1/4)
return y
```

Listing 1. A probabilistic program with no loops

```
x, y <- false, false
while (x = y):
  x, y <- flip(1/3), flip(1/3)
return x
```

Listing 2. Simulating a fair coin with a biased one

```
data TreeF a b = LeafF a | SplitF b b | NeverF
type Cotree a = Fix (TreeF a)
data Tree a = Leaf a
            | Split (Maybe Label) (Tree a) (Tree a)
            | Hole Label
phi :: Tree a -> TreeCoalgebra a
phi = ...
generate :: Tree a -> Cotree a
generate t = (unfold . phi) t (Hole 0)
```

Listing 3. Knuth-Yao trees in Haskell

case, 3). That the program terminates absolutely leads to a natural interpretation as a finite function of boolean variables, which can be compiled to a symbolic formula or BDD.

The program of Listing 2, which simulates a fair coin by flipping the biased coins  $x$  and  $y$  until  $x = y$ , differs fundamentally from that of Listing 1 in that the number of bits required to produce a sample cannot be bounded above by any fixed constant. No finite number of loop iterations suffices to guarantee termination—we can say only that the program *almost surely terminates*, that is, does so with probability 1.

We compile such programs to KY trees, possibly infinite decision diagrams such as the one of Figure 1 in which nodes represent binary decisions and leaves, results. The tree of Figure 1, for example, returns True with probability 1/3, corre-




Fig. 1. Knuth-Yao representation of Bernoulli(1/3)

sponding to the binary expansion of 1/3 as  $0.\overline{01}$ . To indicate the cycle, we use  $T_0$  to label the root of the tree and  $\boxed{0}$  to indicate a labeled “hole”, a tree expansion point that unfolds as  $T_0$ .

KY trees may have arbitrary structure (consider a representation of the Bernoulli distribution with  $p = \pi - 3$ ). To do exact inference, we target the subclass of KY trees that are finitely representable, as sketched in Haskell in Listing 3. The main data type, `Tree a`, defines labeled trees with leaves of type `a` and `Holes`, or corecursive expansion points. We define a mapping from `Trees` to `Cotrees`, the final coalgebra generated by the functor `TreeF`, as the function `generate t`,

<pre> 115 116 e ::= ... 117 c ::= skip 118     x := e 119     x ← e 120     c<sub>1</sub> ; c<sub>2</sub> 121     if e {c<sub>1</sub>} {c<sub>2</sub>} 122     observe e 123     while e c </pre>	<pre> [[skip]](T) ≜ T [[x := e]](T) ≜ bind T (λσ. Leaf σ[x := [[e]]<sub>σ</sub>]) [[x ← e]](T) ≜ ... //Bernoulli(p) and Uniform(... ) [[c<sub>1</sub> ; c<sub>2</sub>]](T) ≜ [[c<sub>2</sub>]]([[c<sub>1</sub>]](T)) [[if e {c<sub>1</sub>} {c<sub>2</sub>}}](T) ≜ bind T (λσ. if [[e]]<sub>σ</sub> then [[c<sub>1</sub>]](Leaf σ) else [[c<sub>2</sub>]](Leaf σ)) [[observe e]](T) ≜ bind T (λσ. if [[e]]<sub>σ</sub> then Leaf σ else Hole 0) [[while e c]](T) ≜ bind T F where F(σ)   ¬[[e]]<sub>σ</sub> ≜ Leaf σ       F(σ)   [[e]]<sub>σ</sub> ≜ Lbl ℓ (bind ([[c]](Leaf σ)) (λσ'. if [[e]]<sub>σ'</sub> then Hole ℓ else Leaf σ')), ℓ fresh </pre>	<pre> 172 173 174 175 176 177 178 179 180 181 182 183 184 185 186 187 188 189 190 191 192 193 194 195 196 197 198 199 200 201 202 203 204 205 206 207 208 209 210 211 212 213 214 215 216 217 218 219 220 221 222 223 224 225 226 227 228 </pre>
(a) Syntax	(b) Semantics	

Fig. 2. Zar syntax and semantics

the anamorphism `unfold` applied to the tree coalgebra produced by `phi t`.

*Semantics.* We map programs in a conventional probabilistic command language (Figure 2a) to KY trees using the semantics of Figure 2b. Commands are interpreted as tree transformers in which input trees  $T$  defining prior distributions over states  $\sigma$  are mapped to output trees defining posterior distributions over results. The function `bind`  $:: \forall a b. \text{Tree } a \rightarrow (a \rightarrow \text{Tree } b) \rightarrow \text{Tree } b$  applies a continuation to the leaves of a tree. When interpreting an observe statement, we bind to the leaves of the tree the continuation that returns `Leaf`  $\sigma$  if the observe condition is satisfied, and `Hole` 0 otherwise, indicating failure (a return to the distinguish root labeled 0). The interpretation of while loops does not take a fixed point, instead generating a finite `Tree` describing the infinite process that repeatedly executes the interpretation of the loop body `[[c]](Leaf  $\sigma$ )`. Sampling a discrete distribution expression  $x \leftarrow e$  (implementation elided) generates a tree that assigns  $x$  the outcomes dictated by  $e$ , with the appropriate probabilities (our prototype currently supports simple discrete distributions like Bernoulli and Uniform).

*Exact Inference.* Figure 3 shows the `Tree` generated by the semantics of Figure 2b for the fair coin program of Listing 2. To perform exact inference, we fix a return predicate  $f : \text{State} \rightarrow \mathbb{R}$  whose expected value we wish to compute, then construct and solve the system of equations induced by mapping  $f$  over the tree. The Figure 3 tree, for example, has leaves corresponding to  $f(\sigma) \triangleq \sigma(x)$  (projection of  $x$ , corresponding to `return x`). The resulting system of linear equations (RHS of Figure 3) computes the weight of each subtree, the probability that  $x = \text{True}$ . The weight of the root  $T_0$  gives the total weight of  $x$ . The system of linear equations generated in this way has a unique solution iff the program interpreted by the tree terminates with probability 1. The inferred distribution is normalized despite the presence of observe statements in the source language because branches inconsistent with observations are made to loop to the root  $T_0$ . An alternative representation sets such branches to 0, yielding an unnormalized distribution.

### 3 RELATED WORK AND DISCUSSION

Both approximate and exact inference are theoretically hard [Roth 1996], but approximate techniques tend to perform better at scale.




Fig. 3. Knuth-Yao tree (left) for the fair coin program of Listing 2 together with the corresponding system of equations (right)

Nonetheless, as [Holtzen et al. 2019] show, exact inference is often successfully applied in practice by exploiting repeating and compositional structure that appears in many problem instances.

Holtzen et al. perform exact inference on finite-domain probabilistic Boolean programs without loops by compiling them to Boolean formulae (represented by BDDs) and applying weighted model-counting (WMC) techniques on the resulting representations. Compilation to BDDs, in contrast to path-based enumeration inference methods (e.g., [Gehr et al. 2016]), exploits duplication and conditional independence in order to minimize the size of the representation, which in turn improves efficiency. Coinductive KY trees, in addition to supporting probabilistic loops, lie in between full path enumeration and compilation to BDDs. Since they are finitely represented, duplication and otherwise unnecessary structure in KY trees can be eliminated to improve performance.

In other closely related work, [Claret et al. 2013] use dataflow analysis techniques to perform exact inference directly on the syntax of finite-domain probabilistic programs, supporting inference on probabilistic loops via computation of fixpoints. Zar supports programs over (possibly infinite) discrete domains, but requires finite support, a restriction our implementation enforces by doing dataflow analysis on Zar loops to ensure the absence of loop-carried dependences (for example, between the values of integer  $i$  in an almost surely terminating loop that increments  $i$  at each step).

### REFERENCES

Guillaume Claret, Sriram K Rajamani, Aditya V Nori, Andrew D Gordon, and Johannes Borgström. 2013. Bayesian inference using data flow analysis. In *Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering*. ACM, 92–102.

229	Timon Gehr, Sasa Misailovic, and Martin Vechev. 2016. PSI: Exact symbolic inference for probabilistic programs. In <i>International Conference on Computer Aided Verification</i> . Springer, 62–83.	286
230		287
231	Steven Holtzen, Todd Millstein, and Guy Van den Broeck. 2019. Symbolic Exact Inference for Discrete Probabilistic Programs. <a href="https://arxiv.org/abs/1904.02079">arXiv:cs.PL/1904.02079</a>	288
232		289
233	Donald Knuth and Andrew Yao. 1976. The complexity of nonuniform random number generation. <i>Algorithm and Complexity, New Directions and Results (1976)</i> , 357–428.	290
234	Andrew Miner and David Parker. 2004. Symbolic representations and analysis of large probabilistic systems. In <i>Validation of Stochastic Systems</i> . Springer, 296–338.	291
235		292
236	Dan Roth. 1996. On the hardness of approximate reasoning. <i>Artificial Intelligence</i> 82, 1-2 (1996), 273–302.	293
237		294
238		295
239		296
240		297
241		298
242		299
243		300
244		301
245		302
246		303
247		304
248		305
249		306
250		307
251		308
252		309
253		310
254		311
255		312
256		313
257		314
258		315
259		316
260		317
261		318
262		319
263		320
264		321
265		322
266		323
267		324
268		325
269		326
270		327
271		328
272		329
273		330
274		331
275		332
276		333
277		334
278		335
279		336
280		337
281		338
282		339
283		340
284		341
285		342