SPPL: Probabilistic Programming with Fast Exact Symbolic Inference

Feras A. Saad

Massachusetts Institute of Technology
Cambridge, MA, USA

Martin C. Rinard

Massachusetts Institute of Technology
Cambridge, MA, USA

Vikash K. Mansinghka

Massachusetts Institute of Technology
Cambridge, MA, USA

Abstract

We present the Sum-Product Probabilistic Language (SPPL), a new probabilistic programming language that automatically delivers exact solutions to a broad range of probabilistic inference queries. SPPL translates probabilistic programs into *sum-product expressions*, a new symbolic representation and associated semantic domain that extends standard sum-product networks to support mixed-type distributions, numeric transformations, logical formulas, and pointwise and set-valued constraints. We formalize SPPL via a novel translation strategy from probabilistic programs to sum-product expressions and give sound exact algorithms for conditioning on and computing probabilities of events. SPPL imposes a collection of restrictions on probabilistic programs to ensure they can be translated into sum-product expressions, which allow the system to leverage new techniques for improving the scalability of translation and inference by automatically exploiting probabilistic structure. We implement a prototype of SPPL with a modular architecture and evaluate it on benchmarks the system targets, showing that it obtains up to 3500x speedups over state-of-the-art symbolic systems on tasks such as verifying the fairness of decision tree classifiers, smoothing hidden Markov models, conditioning transformed random variables, and computing rare event probabilities.

CCS Concepts: • **Mathematics of computing** → *Probabilistic representations; Probabilistic inference problems*; • **Software and its engineering** → *Formal language definitions*.

Keywords: probabilistic programming, symbolic execution

ACM Reference Format:

Feras A. Saad, Martin C. Rinard, and Vikash K. Mansinghka. 2021. SPPL: Probabilistic Programming with Fast Exact Symbolic Inference. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI '21)*, June 20–25, 2021, Virtual, Canada. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3453483.3454078>

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike International 4.0 License.

PLDI '21, June 20–25, 2021, Virtual, Canada

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8391-2/21/06.

<https://doi.org/10.1145/3453483.3454078>

1 Introduction

Reasoning under uncertainty is a well-established theme across diverse fields including robotics, cognitive science, natural language processing, algorithmic fairness, amongst many others [14, 21, 31, 60]. A common approach for modeling uncertainty is to use probabilistic programming languages (PPLs [29]) to both represent complex probability distributions and perform probabilistic inference within the language. There is growing recognition of the utility of PPLs for solving challenging tasks that involve probabilistic reasoning in various application domains [8, 27, 34, 35].

Probabilistic inference is central to reasoning about uncertainty and is a central concern for both PPL implementors and users. Several PPLs leverage approximate inference techniques [28, 59, 67], which have been used effectively in a variety of settings [11, 17, 55]. Drawbacks of approximate inference, however, include a lack of accuracy and/or soundness guarantees [18, 38]; difficulties with programs that combine continuous, discrete, or mixed-type distributions [11, 68]; challenges assessing the quality of iterative solvers [9]; and the substantial expertise needed to write custom inference programs that deliver acceptable performance [17, 39]. To address the shortcomings of approximate inference, several PPLs instead use exact symbolic techniques [6, 10, 23, 43, 69]. These languages can typically express a large class of models, using general computer algebra to solve queries. However, the generality of the symbolic computations causes them to sometimes fail, even on problems with tractable solutions.

Our Work We introduce the Sum-Product Probabilistic Language (SPPL), a system that occupies a new point in the expressiveness vs. performance trade-off space for exact symbolic inference. A key idea in SPPL is to incorporate certain modeling restrictions that avoid the need for general computer algebra, instead using a new, specialized class of “sum-product” symbolic expressions to exactly represent probability distributions specified by SPPL programs. These new symbolic expressions extend and generalize sum-product networks [47], which are computational graphs that have received widespread attention for their clear probabilistic semantics and tractable properties for exact inference—see [64] for a comprehensive and curated literature review. These sum-product expressions are used to automatically obtain exact solutions to probabilistic inference queries about SPPL programs, which are fast and scalable in tractable regimes.

Figure 1. SPPL system overview. Probabilistic programs are translated into symbolic sum-product expressions that represent the joint distribution over all program variables and are used to deliver exact solutions to probabilistic inference queries.

System Overview Fig. 1 shows an overview of our approach. Given a probabilistic program written in SPPL (Lst. 2) a translator (Lst. 3) produces a sum-product expression that represents the prior distribution over all program variables. Given this expression and a query specified by the user, the SPPL inference engine returns an exact answer, where:

1. **simulate**(Vars) returns random samples of program variables from their joint probability distribution;
2. **prob**(Event) returns the probability of an event, which is a predicate on program variables;
3. **condition**(Event) returns a new sum-product expression for the posterior distribution over program variables, given that the specified event is true.

A key aspect of the system design in Fig. 1 is modularity: modeling, conditioning, and querying are factored into distinct stages that reflect the essential components of a Bayesian workflow. Moreover, the dashed back-edge in Fig. 1 indicates that the new sum-product expression returned by **condition** can be reused to interactively invoke additional queries on the posterior distribution. This closure property enables substantial runtime gains across multiple datasets and queries.

Trade-offs SPPL imposes restrictions on probabilistic programs that specifically rule out the following constructs: (i) unbounded loops; (ii) multivariate numeric transformations; and (iii) arbitrary prior distributions on continuous parameters. As a result, SPPL is not designed to express model classes such as regression with a prior on real coefficients; neural networks; support-vector machines; spatial Poisson processes; urn processes; and hidden Markov models with unknown transition matrices. The aforesaid model classes cannot be represented as sum-product expressions, and most of them do not have tractable algorithms for exact inference.

We impose these restrictions to ensure that valid SPPL programs can always be translated into finite sum-product expressions, as opposed to general symbolic algebra expressions. The resulting sum-product expressions delivered by SPPL have a number of characteristics that make them a particularly useful translation target for probabilistic programs:

- **Completeness and Decomposability:** By satisfying important completeness and decomposability conditions from the literature [47, Defs. 4,5], sum-product expressions are guaranteed to represent normalized probability distributions.
- **Efficient Factorization:** By specifying multivariate probability distributions compositionally in terms of sums and products of simpler distributions, sum-product expressions can be simplified by algebraic “factorization” (Fig. 3d, Fig. 6a).
- **Efficient Deduplication:** When an SPPL program specifies a generative model with conditional independence structure, the translated sum-product expression typically contains identical subexpressions that can be “deduplicated” into a single logical node in memory (Fig. 3d, Fig. 6b).
- **Efficient Caching:** Inference algorithms for sum-product expressions proceed from root to leaves to root, allowing intermediate results to be cached and reused at deduplicated internal subexpressions in a depth-first graph traversal.
- **Closure Under Conditioning:** Sum-product expressions are closed under probabilistic conditioning (Thm. 4.1), which allows them to be reused across multiple datasets and inference queries about the same probabilistic program.
- **Linear-Time Exact Inferences:** For a well-defined class of common queries, inference scales linearly in the expression size (Thm. 4.3); when SPPL delivers a “small” expression after factorization and deduplication, inference is also fast.

It is well-known that a very large class of tractable models can be cast as sum-product networks [47, Thm. 2]. SPPL automatically constructs these representations from generative probabilistic programs that use standard constructs such as arrays, if/else branches, for-loops, and numeric and logical operators. To enable this translation, SPPL introduces new sum-product expressions and inference algorithms that extend standard sum-product networks by supporting (many-to-one) univariate transformations, mixed-type base measures, and pointwise and set-valued constraints. These constructs make SPPL expressive enough to solve prominent inference tasks in the PPL literature [2, 36, 46, 68] for which standard sum-product networks have not been previously

used. Example model classes include most finite discrete models, latent variable models with discrete hidden states and arbitrary observed states, and decision trees over discrete and continuous variables. Taken together, these characteristics make SPPL particularly effective for fast and scalable inference on tractable problems, with low variance runtime and complete, usable answers to users. Our experimental evaluation (Sec. 6) indicates that SPPL delivers these benefits on the problems it is designed to solve, whereas more general and expressive techniques in previous solvers [2, 4, 23] typically exhibit orders of magnitude worse performance on these problems, runtime has higher variance, and/or results may be unusable, i.e., with unsimplified symbolic integrals.

Key contributions We identify the following contributions:

- **New semantic domain for sum-product expressions** (Sec. 3) that extends sum-product networks [47] by including mixed-type distributions, numeric transforms, logical formulas, and events with pointwise and set-valued constraints.
- **Provably sound exact symbolic inference algorithms** (Sec. 4) based on a proof that sum-product expressions are closed under conditioning on any event that can be specified in the domain. We use these algorithms to build an efficient and multi-stage inference architecture that separates model translation, conditioning, and querying into distinct stages, enabling interactive workflows and computation reuse.
- **The Sum-Product Probabilistic Language** (Sec. 5), a PPL built on a novel translation semantics from generative code to sum-product expressions, which are used to deliver exact inferences to queries. We present optimization techniques to improve scalability of translation and inference by exploiting conditional independences and repeated structure.
- **Empirical measurements of efficacy** (Sec. 6) on inference tasks from the literature that SPPL targets, which show that it delivers substantial improvements over existing baselines, including up to 3500x speedup over state-of-the-art fairness verifiers [2, 4] and symbolic integration [23], as well as many orders of magnitude speedup over sampling-based inference [40] for computing rare event probabilities.

2 Overview

We next describe two examples that illustrate the programming style in SPPL and queries supported by the system.

2.1 Indian GPA Problem

The Indian GPA problem is a canonical example that has been widely considered in the probabilistic programming literature [44, 45, 48, 57, 68] for its use of a “mixed-type” random variable that takes both continuous and discrete values, depending on the random branch taken by the program.

Specifying the Prior Fig. 2a shows the generative process for three variables (Nationality, Perfect and GPA) of a student. The student’s nationality is either India or USA with

equal probability (line 1). Students from India (line 2) have a 10% probability of a perfect 10 GPA (lines 3–4), otherwise the GPA is uniform over $[0, 10]$ (line 5). Students from USA (line 6) have a 15% probability of a perfect 4 GPA (lines 6–7), otherwise the GPA is uniform over $[0, 4]$ (line 8).

Prior Sum-Product Expression The graph in Fig. 2d shows the translated sum-product expression for this program, which represents a sampler for the distribution over program variables as follows: (i) if a node is a sum (+), visit a random child with probability equal to the weight of the edge pointing to the child; (ii) if a node is a product ($\times$), visit each child exactly once, in no specific order; (iii) if a node is a leaf, sample a value from the distribution at the leaf and assign it to the variable at the leaf. Similarly, the graph encodes the joint distribution of the variables by treating (i) each sum node as a probabilistic mixture; (ii) each product node as a tuple of independent variables; and (iii) each leaf node as a primitive random variable. Thus, the prior distribution is:

$$\begin{aligned} \Pr[\text{Nationality} = n, \text{Perfect} = p, \text{GPA} \leq g] & \quad (1) \\ = & 0.5 [\delta_{\text{India}}(n) \cdot (0.1[(\delta_{\text{True}}(p) \cdot 1[10 \leq g])] \\ & + 0.9[(\delta_{\text{False}}(p) \cdot (g/10 \cdot 1[0 \leq g < 10] + 1[10 \leq g]))]] \\ & + 0.5 [\delta_{\text{USA}}(n) \cdot (0.15[(\delta_{\text{True}}(p) \cdot 1[4 \leq g])] \\ & + 0.85[(\delta_{\text{False}}(p) \cdot (g/4 \cdot 1[0 \leq g < 4] + 1[4 \leq g]))]]]. \end{aligned}$$

Fig. 2b shows SPPL queries for the prior marginal distributions of the three variables, plotted in Fig. 2e. The two jumps in the cumulative distribution function (CDF¹) of GPA at 4 and 10 correspond to the atoms that occur when Perfect is true. The piecewise linear behavior on $[0, 4]$ and $[4, 10]$ follows from the conditional uniform distributions of GPA.

Conditioning the Program Fig. 2f shows an example of the **condition** query, which specifies an event e on which to constrain executions of the program. An event is a predicate on (possibly transformed) program variables that can be used for both **condition** (Fig. 2f) and **prob** (Fig. 2c). SPPL is the first system with inference algorithms for sum-product expressions that handle predicates of this form. Given e , the object of inference is the full posterior distribution:

$$\begin{aligned} \Pr[\text{Nationality} = n, \text{Perfect} = p, \text{GPA} \leq g \mid e] & \quad (2) \\ := & \Pr[\text{Nationality} = n, \text{Perfect} = p, \text{GPA} \leq g, e] / \Pr[e]. \end{aligned}$$

Posterior Sum-Product Expression Given the prior expression (Fig. 2d) and conditioning event e (Fig. 2f), SPPL produces a new expression (Fig. 2g) that specifies a distribution which is precisely equal to Eq. (2). From Thm. 4.1, conditioning an SPPL program on any event that can be specified in the language results in a posterior distribution that also admits an exact sum-product expression. Conditioning on e performs several transformations on the prior expression, which are:

¹ For a real-valued random variable X , the cumulative distribution function $F : \text{Real} \rightarrow [0, 1]$ is given by $F(r) := \Pr[X \leq r]$.

Figure 2. Analyzing the Indian GPA problem in SPPL.

1. Eliminating the subtree rooted at the parent of leaf δ_{10} , which is inconsistent with the conditioning event.
2. Rescaling the distribution $U(0, 10)$ at the leaf node in the India subtree to $U(8, 10)$.
3. Rescaling the distribution $U(0, 4)$ at the leaf node in the USA subtree to $U(3, 4)$.
4. Reweighting the branch probabilities of the sum node in the USA subtree from $[\cdot 15, \cdot 85]$ to $[\cdot 41, \cdot 59]$, where $\cdot 41 = \cdot 15 / (\cdot 15 + \cdot 2125)$ is the posterior probability of $(\text{Perfect} = 1, \text{GPA} = 4)$ given the condition e .

5. Reweighting the branch probabilities at the root from $[\cdot 5, \cdot 5]$ to $[\cdot 33, \cdot 67]$ (same rules as in the previous item).

Fig. 2g shows the posterior expression obtained by applying these transformations. Using this expression, the right-hand side of Eq. (2), which is the object of inference, is

$$\begin{aligned}
 & \Pr[\text{Nationality} = n, \text{Perfect} = p, \text{GPA} \leq g \mid e] \\
 &= \cdot 33 [\delta_{\text{India}}(n) \cdot \delta_{\text{False}}(p) \cdot (\frac{g-8}{2} \cdot \mathbf{1}[8 \leq g < 10] + \mathbf{1}[10 \leq g])] \\
 &+ \cdot 67 [\delta_{\text{USA}}(n) \cdot (\cdot 41 [(\delta_{\text{True}}(p) \cdot \mathbf{1}[4 \leq g])] \\
 &+ \cdot 59 [(\delta_{\text{False}}(p) \cdot (\frac{g}{4} \cdot \mathbf{1}[0 \leq g < 4] + \mathbf{1}[4 \leq g]))]]].
 \end{aligned} \tag{3}$$

(Floats are shown to two decimal places.) We can now run the **prob** queries in Fig. 2b on the conditioned program to plot the posterior marginal distributions, which are shown in Fig. 2h. The example in Fig. 2 illustrates a typical modular workflow in SPPL (Fig. 1), where modeling (Fig. 2a), conditioning (Fig. 2f) and querying (Figs. 2b–2c) are separated into distinct and reusable stages that together express the main components of Bayesian modeling and inference.

2.2 Scalable Inference in a Hierarchical HMM

The next example shows how to perform efficient smoothing in a hierarchical hidden Markov model (HMM [42]) and illustrates the optimization techniques used by the SPPL translator (Sec. 5.1), which exploit conditional independence to ensure that the size of the sum-product expression grows linearly in the number of timesteps. The code box in Fig. 3a shows a hierarchical hidden Markov model with Bernoulli hidden states Z_t and Normal–Poisson observations (X_t, Y_t) . The separated variable indicates whether the mean values of X_t and Y_t at $Z_t = 0$ and $Z_t = 1$ are well-separated. The `p_transition` vector specifies that the current state Z_t switches from the previous state Z_{t-1} with 20% probability. This example leverages the SPPL **array**, **for**, and **switch-cases** statements, where the latter is a macro that expands to **if-else** statements (as in, e.g., the C language):

```
switch x cases (x' in values) {C}
  ↪ desugar
  ↪ if (x in values[0]) {C[x'/values[0]]}
  elif ...
  elif (x in values[n-1]) {C[x'/values[n-1]]},
```

where n is the length of *values* and $C[x/E]$ indicates syntactic replacement of x with expression E in command C .

The top and middle plots in Fig. 3b show a realization of X and Y that result from simulating the process for 100 time steps. The blue and orange regions along the x-axes indicate whether the true hidden state Z is 0 or 1, respectively (these “ground-truth” values of Z are not observed but need to be inferred from X and Y). The bottom plot in Fig. 3b shows the exact posterior marginal probabilities $\Pr[Z_t = 1 \mid x_{0:99}, y_{0:99}]$ for each $t = 0, \dots, 99$ as inferred by SPPL (an inference known as “smoothing”). These probabilities track the true hidden state, i.e., the posterior probabilities that $Z_t = 1$ are low in the blue and high in the orange regions.

Fig. 3c shows a “naive” sum-product expression for the distribution of all program variables up to the first two time steps. This expression is a sum-of-products, where the products in the second level are an enumeration of all possible realizations of program variables, so that the number of terms scales exponentially in the number of time steps. Fig. 3d shows the expression constructed by SPPL, which is (conceptually) based on factoring and sharing common terms in the two level sum-of-products in Fig. 3c. These factorizations

and deduplications exploit conditional independences and repeated structure in the program (Sec. 5.1), which here delivers a expression whose size scales linearly in the number of time points. SPPL can also efficiently solve variants of smoothing, e.g., computing posterior marginals $\Pr[Z_t \mid x_{0:t}, y_{0:t}]$ (filtering) or the posterior joint $\Pr[Z_{0:t} \mid x_{0:t}, y_{0:t}]$ for any t .

3 A Core Calculus for Sum-Product Expressions

This section presents a semantic domain of sum-product expressions that generalizes sum-product networks [47] and enables precise reasoning about them. This domain will be used to (i) establish the closure of sum-product expressions under conditioning on events expressible in the calculus (Thm. 4.1); (ii) describe sound algorithms for exact Bayesian inference in our system (Appx. D); and (iii) describe a procedure for translating a probabilistic program into a sum-product expression in the core language (Sec. 5). Lst. 1 shows denotations of the key syntactic elements (Lst. 9 in Appx. A) in the calculus, which includes real and nominal outcomes (Lst. 1a); real transforms (Lst. 1b); predicates with pointwise and set-valued constraints (Lst. 1c); primitive distributions (Lst. 1e); and multivariate distributions specified compositionally as sums and products of primitive distributions (Lst. 1f).

Basic Outcomes Random variables in the calculus take values in the Outcome := Real + String domain. The symbol + here indicates a sum (disjoint-union) data type, whose elements are formed by the injection operation, e.g., $\downarrow_{\text{Outcome}}^{\text{Real}} r$ for $r \in \text{Real}$. This domain is used to model mixed-type random variables, such as X in the following SPPL program:

```
Z ~ normal(0, 1)
if (Z <= 0): X ~ "negative" # string
elif (0 < Z < 4): X ~ 2*exp(Z) # continuous real
elif (4 <= Z): X ~ atomic(4) # discrete real
```

The Outcomes domain denotes a subset of Outcome as defined by the valuation function $\mathbb{V}$ (Lst. 1a). For example, $((b_1 r_1) (r_2 b_2))$ specifies a (open, closed, or clopen) real interval and $\{s_1 \dots s_m\}^b$ is a set of strings, where $b = \#t$ indicates the complement (meta-variables such as m indicate an arbitrary but finite number of repetitions of a domain variable or subexpression). The operations *union*, *intersection*, and *complement* operate on Outcomes in the usual way (while preserving certain semantic invariants, see Appx. B)

A Sigma Algebra of Outcomes To speak precisely about random variables and measures on Outcome, we define a sigma-algebra $\mathcal{B}(\text{Outcome}) \subset \mathcal{P}(\text{Outcome})$ as follows:

1. Let τ_{Real} be the usual topology on Real.
2. Let τ_{String} be the discrete topology on String.
3. Let $\tau_{\text{Outcome}} := \tau_{\text{Real}} \uplus \tau_{\text{String}}$ be the disjoint-union topology on Outcome, where U is open iff $\{r \mid (\downarrow_{\text{Outcome}}^{\text{Real}} r) \in U\}$ is open in Real and $\{s \mid (\downarrow_{\text{Outcome}}^{\text{String}} s) \in U\}$ is open in String.
4. Let $\mathcal{B}(\text{Outcome})$ be the Borel sigma-algebra of τ_{Outcome} .

```

1 p_transition = [.2, .8]
2 mu_x = [[5, 7], [5, 15]]
3 mu_y = [[5, 8], [3, 8]]
4
5 n_step = 100
6 Z = array(n_step)
7 X = array(n_step)
8 Y = array(n_step)
9
10 separated ~ bernoulli(p=.4)
11 switch separated cases (s in [0,1]):
12   Z[0] ~ bernoulli(p=.5)
13   switch Z[0] cases (z in [0, 1]):
14     X[0] ~ normal(mu_x[s][z], 1)
15     Y[0] ~ poisson(mu_y[s][z])
16   for t in range(1, n_step):
17     switch Z[t-1] cases (z in [0, 1]):
18       Z[t] ~ bernoulli(p_transition[z])
19       switch Z[t] cases (z in [0, 1]):
20         X[t] ~ normal(mu_x[s][z], 1)
21         Y[t] ~ poisson(mu_y[s][z])

```

(a) Probabilistic Program

(b) Observed Data X, Y and Inferred Hidden States Z

(c) Naive Sum-Product Expression (Scales Exponentially)

(d) Optimized Sum-Product Expression (Scales Linearly)

Figure 3. Fast smoothing in a hierarchical hidden Markov model using SPPL by constructing an efficient sum-product expression that exploits conditional independences in the generative process. Optimization techniques are discussed in Sec. 5.1.

Remark 3.1. As measures on Real are defined by their values on open intervals and measures on String on singletons, we can speak of probability measures on $\mathcal{B}(\text{Outcome})$ as mappings from Outcomes to $[0, 1]$.

Real Transformations Lst. 1b shows real transformations that can be applied to variables in the calculus. The Identity Transform, written $\text{Id}(x)$, is a terminal subexpression of any Transform t and contains a single variable name that specifies the “dimension” over which t is defined. The list of all transforms is in Appx. C.1. The key operation involving

transforms is computing the preimage of Outcomes v under t using $\text{preimg} : \text{Transform} \rightarrow \text{Outcomes} \rightarrow \text{Outcomes}$ which satisfies the following properties:

$$(\downarrow_{\text{Outcome}}^{\text{Real}} r) \in \mathbb{V}[\text{preimg } t \ v] \iff \mathbb{T}[t](r) \in \mathbb{V}[v]$$

$$(\downarrow_{\text{Outcome}}^{\text{String}} s) \in \mathbb{V}[\text{preimg } t \ v] \iff (t \in \text{Identity}) \wedge (s \in \mathbb{V}[v]).$$

Appx. C.2 presents a symbolic solver that implements preimg for each Transform, which is leveraged to enable exact probabilistic inferences on transformed variables in SPPL. Fig. 4 and Appx. C.3 show example inferences with transforms.

<p>Outcome := Real + String $\mathbb{V} : \text{Outcomes} \rightarrow \mathcal{P}(\text{Outcome})$ $\mathbb{V}[\emptyset] := \emptyset$ [Empty] $\mathbb{V}[\{s_1 \dots s_m\}^b] := \text{if } b \text{ then } \bigcup_{i=1}^m \{(\downarrow_{\text{Outcome}}^{\text{String}} s_i)\}$ [FiniteStr] $\quad \text{else } \{(\downarrow_{\text{Outcome}}^{\text{Real}} s) \mid \forall i.s \neq s_i\}$ $\mathbb{V}[\{r_1 \dots r_m\}] := \bigcup_{i=1}^m \{(\downarrow_{\text{Outcome}}^{\text{Real}} r_i)\}$ [FiniteReal] $\mathbb{V}[(\langle b_1 r_1 \rangle (r_2 b_2))] := \{(\downarrow_{\text{Outcome}}^{\text{Real}} r) \mid r_1 <_{b_1} r <_{b_2} r_2\}$ [Interval] $\quad \text{where } <_{\#t} := <; <_{\#f} := <; r_1 < r_2$ $\mathbb{V}[v_1 \sqcup \dots \sqcup v_m] := \bigcup_{i=1}^m \mathbb{V}[v_i]$ [Union]</p> <p>(a) Outcomes</p>	<p>$\mathbb{D} : \text{Distribution} \rightarrow \text{Outcomes} \rightarrow [0, 1]$ $\mathbb{D}[\text{DistS}((s_i w_i)_{i=1}^m)] v :=$ [DistStr] $\text{match } (\text{intersection } v \{s_1 \dots s_m\}^{\#f})$ $\triangleright \emptyset \mid \{r'_1 \dots r'_m\} \mid (\langle b_1 r_1 \rangle (r_2 b_2)) \Rightarrow 0$ $\triangleright v_1 \sqcup \dots \sqcup v_m \Rightarrow \sum_{i=1}^m \mathbb{D}[\text{DistS}((s_i w_i)_{i=1}^m)] v_i$ $\triangleright \{s'_1 \dots s'_k\}^b \Rightarrow \text{let } w \text{ be } \sum_{i=1}^m (w_i \text{ if } s_i \in \{s'_j\}_{j=1}^k \text{ else } 0)$ $\quad \text{in if } \bar{b} \text{ then } w \text{ else } 1 - w$ $\mathbb{D}[\text{DistR}(F r_1 r_2)] v := \text{match } (\text{intersection } ((\#f r_1) (r_2 \#f)) v)$ [DistReal] $\triangleright \emptyset \mid \{r'_1 \dots r'_m\} \mid \{s'_1 \dots s'_k\}^b \Rightarrow 0$ $\triangleright v_1 \sqcup \dots \sqcup v_m \Rightarrow \sum_{i=1}^m \mathbb{D}[\text{DistR}(F r_1 r_2)] v_i$ $\triangleright (\langle b'_1 r'_1 \rangle (r'_2 b'_2)) \Rightarrow \frac{F(r'_2) - F(r'_1)}{F(r_2) - F(r_1)}$ $\mathbb{D}[\text{DistI}(F r_1 r_2)] v := \text{match } (\text{intersection } ((\#f r_1) (r_2 \#f)) v)$ [DistInt] $\triangleright \emptyset \mid \{s'_1 \dots s'_k\}^b \Rightarrow 0$ $\triangleright v_1 \sqcup \dots \sqcup v_m \Rightarrow \sum_{i=1}^m \mathbb{D}[\text{DistI}(F r_1 r_2)] v_i$ $\triangleright \{r'_1 \dots r'_m\} \Rightarrow \frac{\sum_{i=1}^m [\text{if } (r'_i = \lfloor r'_i \rfloor) \wedge (r_1 \leq r'_i \leq r_2) \text{ then } F(r'_i) - F(r'_i - 1) \text{ else } 0]}{F(\lfloor r_2 \rfloor) - F(\lceil r_1 \rceil - 1)}$ $\triangleright (\langle b'_1 r'_1 \rangle (r'_2 b'_2)) \Rightarrow \text{let } \tilde{r}_1 \text{ be } \lfloor r'_1 \rfloor - 1 \mid (r'_1 = \lfloor r'_1 \rfloor) \wedge \bar{b}'_1$ $\quad \text{in let } \tilde{r}_2 \text{ be } \lfloor r'_2 \rfloor - 1 \mid (r'_2 = \lfloor r'_2 \rfloor) \wedge \bar{b}'_2$ $\quad \text{in } \frac{F(\tilde{r}_2) - F(\tilde{r}_1)}{F(\lfloor r_2 \rfloor) - F(\lceil r_1 \rceil - 1)}$</p>
<p>$\mathbb{T} : \text{Transform} \rightarrow \text{Real} \rightarrow \text{Real}$ $\mathbb{T}[\text{Id}(x)] := \lambda r'.r'; \mathbb{T}[\text{Reciprocal}(t)] := \lambda r'.1/(\mathbb{T}[t](r'));$ $\mathbb{T}[\text{Abs}(t)] := \lambda r'. \mathbb{T}[t](r') ; \mathbb{T}[\text{Root}(t n)] := \lambda r'.\sqrt[n]{\mathbb{T}[t](r')};$ $\mathbb{T}[\text{Poly}(t r_0 \dots r_m)] := \lambda r'.\sum_{i=0}^m r_i (\mathbb{T}[t](r'))^i; \dots$</p> <p>(b) Transformations (Lst. 17 in Appx. C.1)</p>	<p>$\mathbb{E} : \text{Event} \rightarrow \text{Var} \rightarrow \text{Outcomes}$ $\mathbb{E}[(t \text{ in } v)] x := \text{if } (\text{vars } t) = \{x\} \text{ then } (\text{preimg } t v) \text{ else } \emptyset$ [Contains] $\mathbb{E}[e_1 \sqcap \dots \sqcap e_m] x := (\text{intersection } \mathbb{E}[e_1] x \dots \mathbb{E}[e_m] x)$ [Conjunction] $\mathbb{E}[e_1 \sqcup \dots \sqcup e_m] x := (\text{union } \mathbb{E}[e_1] x \dots \mathbb{E}[e_m] x)$ [Disjunction]</p>
<p>(c) Events</p>	<p>(e) Primitive Distributions</p>
<p>$\mathbb{P}_0[S] : \text{SPE} \rightarrow \text{Event} \rightarrow \text{Natural} \times [0, \infty)$ $\mathbb{P}_0[\text{Leaf}(x d \sigma)] (\text{Id}(x) \text{ in } \{rs\}) := \text{match } d$ [Leaf] $\triangleright \text{DistR}(F r_1 r_2) \Rightarrow \text{match } rs$ $\triangleright r \Rightarrow (1, 1 \mid r_1 \leq r \leq r_2 \mid F'(r)/[F(r_2) - F(r_1)])$ $\triangleright s \Rightarrow (1, 0)$ $\triangleright \text{else} \Rightarrow \text{let } w \text{ be } \mathbb{P}[\text{Leaf}(x d \sigma)] (\text{Id}(x) \text{ in } \{rs\}) \text{ in } (1 \mid w = 0, w)$ $\mathbb{P}_0[(S_1 w_1) \oplus \dots \oplus (S_m w_m)] \cap_{i=1}^\ell (\text{Id}(x_i) \text{ in } \{rs_i\}) :=$ $\text{let}_{1 \leq i \leq m} (d_i, p_i) \text{ be } \mathbb{P}_0[S_i] \left(\cap_{i=1}^\ell (\text{Id}(x_i) \text{ in } \{rs_i\}) \right)$ [Sum] $\text{in if } \forall_{1 \leq i \leq m}. p_i = 0 \text{ then } (1, 0)$ $\quad \text{else let } d^* \text{ be } \min\{d_i \mid 1 \leq i \leq m, 0 < p_i\}$ $\quad \text{in } (d^*, \sum_{i=1}^m 1[d_i = d^*] w_i p_i)$ $\mathbb{P}_0[S_1 \otimes \dots \otimes S_m] \cap_{i=1}^\ell (\text{Id}(x_i) \text{ in } \{rs_i\}) :=$ [Product] $\text{let}_{1 \leq i \leq m} (d_i, p_i) \text{ be match } \{x_1, \dots, x_m\} \cap (\text{scope } S_i)$ $\triangleright \{n_1, \dots, n_k\} \Rightarrow \mathbb{P}_0[S_i] \cap_{i=1}^k (\text{Id}(x_{n_i}) \text{ in } \{rs_i\})$ $\triangleright \{\} \Rightarrow (0, 1)$ $\text{in } (\sum_{i=1}^n d_i, \prod_{i=1}^m p_i)$</p> <p>(d) Sum-Product Expressions (Density Semantics)</p>	<p>$\mathbb{P} : \text{SPE} \rightarrow \text{Event} \rightarrow [0, 1]$ $\mathbb{P}[\text{Leaf}(x d \sigma)] e := \mathbb{D}[d] (\mathbb{E}[(\text{subsen } v e \sigma)] x)$ [Leaf] $\mathbb{P}[(S_1 w_1) \oplus \dots \oplus (S_m w_m)] e := \text{let } Z \text{ be } \sum_{i=1}^m w_i$ [Sum] $\quad \text{in } \sum_{i=1}^m (\mathbb{P}[S_i] e) w_i / Z$ $\mathbb{P}[S_1 \otimes \dots \otimes S_m] e := \text{match } (\text{dnf } e)$ [Product] $\triangleright (t \text{ in } v) \Rightarrow \text{let } n \text{ be } \min\{1 \leq i \leq m \mid (\text{vars } e) \subset (\text{scope } S_i)\}$ $\quad \text{in } \mathbb{P}[S_n] e$ $\triangleright (e_1 \sqcap \dots \sqcap e_\ell) \Rightarrow$ $\quad \prod_{1 \leq i \leq m} \left[\text{match } \{1 \leq j \leq \ell \mid (\text{vars } e_j) \subset (\text{scope } S_i)\} \right.$ $\quad \left. \triangleright \{n_1, \dots, n_k\} \Rightarrow \mathbb{P}[S_i] (e_{n_1} \sqcap \dots \sqcap e_{n_k}) \right]$ $\triangleright (e_1 \sqcup \dots \sqcup e_\ell) \Rightarrow \sum_{J \subset [\ell]} \left[(-1)^{ J -1} \mathbb{P}[S_1 \otimes \dots \otimes S_m] (\cap_{i \in J} e_i) \right]$</p> <p>(f) Sum-Product Expressions (Distribution Semantics)</p>

Listing 1. Syntax and semantics of a core calculus for sum-product expressions and related domains.

Events Lst. 1c shows the Event domain, which specifies predicates on variables. The valuation $\mathbb{E}[e] : \text{Var} \rightarrow \text{Outcomes}$ of an Event takes a variable x and returns the set $v \in \text{Outcomes}$ of elements that satisfy the predicate along dimension x , leveraging the properties of *preimg*. This domain specifies measurable sets of an n -dimensional distribution on variables $\{x_1, \dots, x_n\}$ as follows: let $\sigma_{\text{gen}}(\{A_1, A_2, \dots\})$ be the sigma-algebra generated by $A_1, A_2, \dots$, and define $\mathcal{B}^n(\text{Outcome}) := \sigma_{\text{gen}}(\{\prod_{i=1}^n U_i \mid \forall_{1 \leq i \leq n}. U_i \in \mathcal{B}(\text{Outcomes})\})$. In other words, $\mathcal{B}^n(\text{Outcome})$ is the n -fold product sigma-algebra generated by open rectangles of Outcomes. Any $e \in \text{Event}$ specifies a measurable set U in $\mathcal{B}^n(\text{Outcome})$, whose i th coordinate $U_i = \mathbb{E}[e] x_i$ if $x_i \in \text{vars } e$, and $U_i = \text{Outcomes}$ otherwise. Any Transform in e is solved and any Var that does not appear in e is marginalized, as in the next example.

Example 3.2. Let $\{X, Y, Z\}$ be elements of Var. Then

$$e := \text{Reciprocal}(\text{Id}(X)) \text{ in } ((\#f 1) (2 \#f))$$

corresponds to the $\mathcal{B}^3(\text{Outcome})$ -measurable set

$$\{(\downarrow_{\text{Outcome}}^{\text{Real}} r) \mid 1/2 \leq r \leq 1\} \times \text{Outcomes} \times \text{Outcomes}.$$

As in Remark 3.1, we may speak about probability distributions on $\mathcal{B}^n(\text{Outcome})$ as mappings from Event to $[0, 1]$.

Primitive Distributions Lst. 1e presents the primitive distributions out of which multivariate distributions are constructed. The CDF domain contains cumulative distribution functions F , whose quantile function is denoted F^{-1} and derivative F' . CDF is in 1-1 correspondence with all distributions and random variables on Real [7, Thms 12.4, 14.1]. The Distribution domain specifies continuous real, atomic

Figure 4. Inference on a stochastic many-to-one transformation of a real random variable in SPPL.

real (on the integers) and nominal distributions. The denotation $\mathbb{D} \llbracket d \rrbracket$ of a Distribution is a distribution on Outcomes (recall Remark 3.1). For example, $\text{DistR}(F r_1 r_2)$ is the restriction of F to a positive measure interval $[r_1, r_2]$. The distributions specified by DistR and DistI can be simulated using a variant of the integral probability transform (Prop. A.1 in Appx. A), which also defines their sampling semantics.

Sum-Product Expressions Lsts. 1d and 1f show the probability density and distribution semantics of the SPE domain, respectively, whose elements are probability measures. The following conditions specify well-definedness for SPE:

- (C1) $\forall \text{Leaf}(x \ d \ \sigma). x \in \sigma$ and $\sigma(x) = \text{Id}(x)$.
- (C2) $\forall \text{Leaf}(x \ d \ \sigma)$. If $\text{dom}(\sigma) = \{x, x_1, \dots, x_m\}$ for some $m > 0$ then $\forall_{1 \leq t \leq m}. (\text{vars } \sigma(x_t)) \subset \{x, x_1, \dots, x_{t-1}\}$.
- (C3) $\forall (S_1 \otimes \dots \otimes S_m). \forall i \neq j. (\text{scope } S_i) \cap (\text{scope } S_j) = \emptyset$.
- (C4) $\forall (S_1 w_1) \oplus \dots \oplus (S_m w_m). \forall i. (\text{scope } S_i) = (\text{scope } S_1)$.
- (C5) $\forall (S_1 w_1) \oplus \dots \oplus (S_m w_m). w_1 + \dots + w_m > 0$.

For Leaf , (C1) ensures that σ maps the leaf variable x to the Identity Transform and (C2) ensures there are no cyclic dependencies or undefined variables in Environment σ . Condition (C3) ensures the scopes of all children of a Product are disjoint and (C4) ensures the scopes of all children of a Sum are identical, which together ensure completeness and decomposability from sum-product networks [47, Defs. 4, 5].

In Lst. 1f, the denotation $\mathbb{P} \llbracket S \rrbracket$ of $S \in \text{SPE}$ is a map from $e \in \text{Event}$ to its probability under the n -dimensional distribution defined by S , where $n := |\text{scope } S|$ is the number of variables in S . A terminal node $\text{Leaf}(x \ d \ \sigma)$ is comprised of a Var x , Distribution d , and Environment σ that maps other variables to a Transform of x , e.g., $Z \mapsto \text{Poly}(\text{Root}(\text{Id}(X) \ 2))$ [11, 5].

When assessing the probability of e at a Leaf, *subsenv* (Lst. 13 in Appx. A) rewrites e as an Event e' on one variable x , so that the probability of Outcomes that satisfy e is exactly $\mathbb{D} \llbracket d \rrbracket (\mathbb{E} \llbracket e' \rrbracket x)$. The *scope* function (Lst. 12 in Appx. A) returns the list of variables in S . For a Sum, the probability of

e is a weighted average of the probabilities under each subexpression. For a Product, the semantics are defined in terms of (*dnf* e) (Lst. 15 in Appx. B), leveraging inclusion-exclusion.

In Lst. 1d, the denotation $\mathbb{P}_0 \llbracket S \rrbracket$ defines the density semantics of SPE, used for measure zero events such as $\{X = 3, Y = \pi, Z = \text{"foo"}\}$ under a mixed-type base measure. These semantics, which define the density as a pair, adapt “lexicographic likelihood-weighting”, an approximate inference algorithm for discrete-continuous Bayes Nets [68], to exact inference using “lexicographic enumeration” for SPE.

4 Conditioning Sum-Product Expressions

We next present the main theoretical result for exact inference on probability distributions defined by an expression $S \in \text{SPE}$ and describe the inference algorithm for conditioning on an Event (Lst. 1c) in the core calculus, which includes transformations and predicates with set-valued constraints.

Theorem 4.1 (Closure under conditioning). *Let $S \in \text{SPE}$ and $e \in \text{Event}$ be given, where $\mathbb{P} \llbracket S \rrbracket e > 0$. There exists an algorithm which, given S and e , returns $S' \in \text{SPE}$ such that, for all $e' \in \text{Event}$, the probability of e' according to S' is equal to the conditional probability of e' given e according to S , i.e.,*

$$\mathbb{P} \llbracket S' \rrbracket e' \equiv \mathbb{P} \llbracket S \rrbracket (e' \mid e) := \frac{\mathbb{P} \llbracket S \rrbracket (e \cap e')}{\mathbb{P} \llbracket S \rrbracket e}. \quad (5)$$

Thm. 4.1 is a structural conjugacy property [20] for the family of probability distributions defined by the SPE domain, where both the prior and posterior are identified by elements of SPE. We establish Eq. (5) constructively, by describing a new algorithm *condition* : $\text{SPE} \rightarrow \text{Event} \rightarrow \text{SPE}$ that satisfies

$$\mathbb{P} \llbracket (\text{condition } S \ e) \rrbracket e' = \mathbb{P} \llbracket S \rrbracket (e' \mid e) \quad (6)$$

for all $e, e' \in \text{Event}$ with $\mathbb{P} \llbracket S \rrbracket e > 0$. Refer to Appx. D for the proof. Fig. 5 shows a conceptual example of how *condition*

Figure 5. Conditioning a Product S on an Event e that defines a union of hyperrectangles in Real^3 . The inference algorithm partitions the region into a disjoint union, in this case converting two overlapping regions into five disjoint regions. The result is a Sum-of-Product, where each child is the restriction of S to one of the disjoint hyperrectangles.

works, where the prior distribution is a Product S and the conditioned distribution is a Sum-of-Product S' . Fig. 4 shows an example of the closure property when the expression has transformed variables (details in Appx. C).

Remark 4.2. Thm. 4.1 refers to a positive probability Event e . As with sum-product networks, SPE is also closed under conditioning on a Conjunction of possibly measure zero equality constraints on non-transformed variables, which appear in many PPL interfaces [17, 41, 51]. Appx. D.3 presents the condition_0 algorithm for inference on such events, leveraging the generalized mixed-type density semantics in Lst. 1d.

The next result, Thm. 4.3, states a sufficient requirement for inference using $(\text{condition } S e)$ to scale linearly in the size of S , which holds for both zero and positive measure events.

Theorem 4.3. *The runtime of $(\text{condition } S e)$ scales linearly in the number of nodes in the graph representing S whenever e is a single Conjunction $(t_1 \text{ in } v_1) \sqcap \dots \sqcap (t_m \text{ in } v_m)$ of Containment constraints on non-transformed variables.*

5 Translating Probabilistic Programs to Sum-Product Expressions

We next present a probabilistic language called SPPL and show how to translate each program in the language to an element $S \in \text{SPE}$ that symbolically represents (via $\mathbb{P} \llbracket S \rrbracket$) the probability distribution specified by the program. As in Fig. 1, S can then be used to answer queries about an Event e :

simulate: Samples from the distribution defined by $\mathbb{P} \llbracket S \rrbracket$;
prob: Computes the probability of e , using $\mathbb{P} \llbracket S \rrbracket e$ (Lst. 1f);
condition: Conditions on e , using condition (Eq. (6)).

Lst. 2 shows the source syntax of SPPL, which contains standard constructs of an imperative language such as **array** data structures, **if-else** statements, and bounded **for** loops. The **switch-case** macro is defined in Eq. (4). Random variables are defined using “sample” ($\sim$) and $\text{condition}(E)$ can be used to restrict executions to those for which $E \in \text{Expr}$ evaluates to $\#t$ as part of the prior definition. Lst. 3 defines a relation $\langle C, S \rangle \rightarrow_{\text{SPE}} S'$, which translates a “current” $S \in \text{SPE}$ and $C \in \text{Command}$ into $S' \in \text{SPE}$, where the initial

$x \in \text{Var}; y \in \text{ArrayVar}; n \in \text{Natural}; b \in \text{Boolean}; r \in \text{Real}; s \in \text{String};$
 $\circ_{\text{arith}} \in \{+, -, *, /, **\}; \quad \circ_{\text{bool}} \in \{\text{and}, \text{or}\}; \quad \circ_{\text{neg}} \in \{\text{not}\};$
 $\circ_{\text{rel}} \in \{<, <=, >, >=, ==, \text{in}\}; \quad D \in \{\text{normal}, \text{poisson}, \text{choice}, \dots\};$
 $E \in \text{Expr} := x \mid n \mid b \mid r \mid s \mid y[E] \mid D(E^*) \mid (E_1, \dots, E_m)$
 $\mid E_1 \circ_{\text{arith}} E_2 \mid \circ_{\text{neg}} E \mid E_1 \circ_{\text{bool}} E_2 \mid E_1 \circ_{\text{rel}} E_2$
 $C \in \text{Command} := x = E \mid y[E_1] = E_2 \mid x \sim E \mid y[E_1] \sim E_2 \mid y = \text{array}(E)$
 $\mid \text{skip} \mid C_1; C_2 \mid \text{if } E \{C_1\} \text{ else } \{C_2\} \mid \text{condition}(E)$
 $\mid \text{for } x \text{ in range}(E_1, E_2) \{C\} \mid \text{switch } x_1 \text{ cases } (x_2 \text{ in } E) \{C\}$

Listing 2. Source syntax of SPPL.

(SAMPLE)

$$\frac{E \Downarrow d; \quad \text{where } x \notin \text{scope } S}{\langle x \sim E, S \rangle \rightarrow_{\text{SPE}} S \otimes (x \text{ d } \{x \mapsto \text{Id}(x)\})}$$

 (TRANSFORM-LEAF)

$$\frac{E \Downarrow t; \quad \text{where vars } t \in \text{dom}(\sigma), x \notin \text{dom}(\sigma)}{\langle x = E, \text{Leaf}(x' \text{ d } \sigma) \rangle \rightarrow_{\text{SPE}} \text{Leaf}(x' \text{ d } (\sigma \cup \{x \mapsto t\}))}$$

 (TRANSFORM-SUM)

$$\frac{E \Downarrow t, \forall 1 \leq i \leq m. \langle x = E, S_i \rangle \rightarrow_{\text{SPE}} S'_i}{\langle x = E, \oplus_{i=1}^m (S_i \text{ w}_i) \rangle \rightarrow_{\text{SPE}} \oplus_{i=1}^m (S'_i \text{ w}_i)}$$

 (TRANSFORM-PROD)

$$\frac{E \Downarrow t, \langle x = E, S_j \rangle \rightarrow_{\text{SPE}} S'_j; \quad \text{where } j := \min\{i \mid (\text{vars } E) \in \text{scope } S_i\} > 0}{\langle x = E, \otimes_{i=1}^m S_i \rangle \rightarrow_{\text{SPE}} \otimes_{i=1, i \neq j}^m S_i \otimes S'_j}$$

 (IFELSE)

$$\frac{E \Downarrow e, \langle C_1, \text{condition } S e \rangle \rightarrow_{\text{SPE}} S_1, \langle C_2, \text{condition } S (\text{negate } e) \rangle \rightarrow_{\text{SPE}} S_2}{\langle \text{if } E \{C_1\} \text{ else } \{C_2\}, S \rangle \rightarrow_{\text{SPE}} (S_1 \mathbb{P} \llbracket S \rrbracket e) \oplus (S_2 (1 - \mathbb{P} \llbracket S \rrbracket e))}$$

 (FOR-REPEAT)

$$\frac{E_1 \Downarrow n_1, E_2 \Downarrow n_2; \quad \text{where } n_1 < n_2}{\langle \text{for } x \text{ in range}(E_1, E_2) \{C\}, S \rangle \rightarrow_{\text{SPE}} \langle C[x/n_1]; \text{for } x \text{ in range}(n_1 + 1, E_2) \{C\}, S \rangle}$$

Listing 3. Example rules for translating an SPPL command C (Lst. 2) to an element of SPE (Lst. 1f).

step operates on an “empty” S . (Lst. 8 in Appx. E defines a semantics-preserving inverse of $\rightarrow_{\text{SPE}}$). The $\Downarrow$ relation evaluates $E \in \text{Expr}$ to other domains in the core calculus using straightforward rules. We briefly describe key rules of $\rightarrow_{\text{SPE}}$:

(TRANSFORM-LEAF) updates the environment σ at each Leaf.
 (TRANSFORM-SUM) delegates to all subexpressions.
 (TRANSFORM-PROD) delegates to the subexpression whose scope contains the transformed variable.
 (FOR-REPEAT) unrolls a **for** loop into a Command sequence.
 (IFELSE) returns a Sum with two subexpressions, where the **if** branch is conditioned on the test Event and the **else** branch is conditioned on the negation of the test Event. This translation step involves running inference (using condition , Eq. (6)) on the current $S \in \text{SPE}$ translated so far.

The rule for $\text{condition}(E)$ (not shown) calls $(\text{condition } S e)$ (Eq. 6) or $(\text{condition}_0 S e)$ (Remark 4.2), where $E \Downarrow e$. To ensure SPPL programs translate to well-defined element of SPE, per (C1)–(C5), each program must satisfy these restrictions:

(R1) Variables x in $x \sim E$ (SAMPLE) and $x = E$ (TRANSFORM-LEAF) must be fresh (ensures conditions (C1), (C2) and (C3)).
 (R2) The branches in an **if-else** statement must define identical variables (ensures conditions (C4) and (C5)).

(a) Invalid program (translates to an infinite-sized SPE)

```

mu ~ beta(a=4, b=3, scale=7)
num_loops ~ poisson(mu)      # invalid (real integral)
for i in range(0, num_loops): # invalid (infinite series)
  [... commands ... ]

```

(b) Valid program (translates to a finite-sized SPE)

```

mu ~ beta(a=4, b=3, scale=7)
# binspace partitions [0,7] into 10 intervals
switch (mu) cases (m in binspace(0, 7, n=10)):
  num_loops ~ poisson(m.mean()) # discretization
condition (num_loops < 50)      # truncation
switch num_loops cases (n in range(50)):
  for i in range(0, n):
    [... commands ... ]

```

Listing 4. Examples of valid and invalid SPPL programs.

(R3) Derived random variables are obtained via (many-to-one) *univariate* transformations (Lst. 1b).

(R4) Parameters of distributions D or **range** must be either constants or random variables with finite support.

(R3) is required since the distribution of a multivariate transform (e.g., $Z = X/Y^2$) is typically intractable and does not factor into Sum and Product expressions. (R4) is required to ensure a finite-size SPE: distributional parameters with infinite support require integrals (uncountable support) or infinite series (countable support), which are not in SPE. Lst. 4 shows an example of using **switch** and **condition** to work around these restrictions by discretization and truncation.

5.1 Building Compact Sum-Product Expressions

As discrete Bayesian networks can be encoded as SPPL programs, it is possible to write programs where exact inference is NP-Hard [16], which corresponds to an element of SPE that is exponentially large. It is well known that the complexity of exact inference in Bayesian networks is worst case exponential in the treewidth, which is the only structural restriction that can ensure tractability [12]. As computing treewidth is NP-Complete [3], for fundamental theoretical reasons we cannot generally check conditions needed for even simple SPPL programs, such as those that only use if/else statements on binary variables, to translate into a “small” expression.

However, many models of interest contain (conditional) independence relationship [33] that induce a compact factorization of the model into tractable subparts, as in, e.g., Sec. 2.2. SPPL uses several optimization techniques to improve scalability of translation (Lst. 3) and inference (Eq. (6)) by automatically exploiting independences and repeated structure, when they exist, to build compact sum-product expressions.

Factorization Using standard algebraic manipulations, a sum-product expression can be made smaller without changing its semantics (Lst. 1f) by “factoring out” common terms (Fig. 6a), provided that the new expression satisfies (C1)–(C5). Factorization plays a key role in the (IfElse) rule of $\rightarrow_{\text{SPE}}$: since all statements before the **if-else** are shared by the

Figure 6. Exploiting independences and repeated structure during translation of SPPL programs to build compact sum-product expressions. Blue subtrees are identical components.

Table 1. Measurements of SPE graph size with and without the factorization and deduplication optimizations in Fig. 6.

Benchmark	No. of Nodes in Translated SPE		Data Compression Ratio (unopt/opt)
	Unoptimized	Optimized	
Hiring [2]	33	27	1.2x
Alarm [46]	58	45	1.3x
Grass [46]	130	59	2.2x
Noisy OR [46]	783	132	4.1x
Clinical Trial [46]	43761	4131	10.6x
Heart Disease [56]	1041235	6257	166.4x
Hierarchical HMM (Sec. 2.2)	29273397577908185	1787	16381308101795x

bodies of the **if** and **else** branches, statements outside the branch that are independent of statements inside the branch often produce subexpressions that can be factored out.

Deduplication When a sum-product expression contains duplicate subexpressions that cannot be factored out without violating the definedness conditions, we instead resolve duplicates into a single physical representative. Fig. 6b shows an example where the left and right components of the original expression contain an identical subexpression S (in blue), but factorization would lead to an invalid sum-product expression. The optimizer represents the computation graph of this expression using a single data structure S shared by the left and right subtrees (see also Figs. 3c–3d).

Memoization While deduplication reduces memory overhead, memoization is used to reduce runtime overhead. Consider either SPE in Fig. 6b: calling **condition** on the Sum root will dispatch the query to the left and right subexpressions (Lst. 6b). We cache the results of $(\text{condition } S \ e)$ or $\mathbb{P}[\llbracket S \rrbracket e]$ when S is visited in the left subtree to avoid recomputing the result when S is visited again in the right subtree via a depth-first traversal. Memoization delivers large runtime gains not only for solving queries but also for detecting duplicates returned by **condition** in the (IfElse) translation step.

Measurements Table 1 shows measurements of performance gains delivered by the factorization and deduplication optimizations on seven benchmarks. Compression ratios range between 1.2x to $1.64 \times 10^{13}x$ and are highest in the presence of independence or repeated structure. The deduplication and memoization optimizations together enable fast detection of duplicate subtrees by comparing logical memory addresses of internal nodes in $O(1)$ time, instead of computing hash functions that require an expensive subtree traversal.

Table 2. Runtime measurements and speedup for 15 fairness verification tasks using SPPL, FairSquare [2], and VeriFair [4].

Decision Program	Population Model	Lines of Code	Fairness Judgment	Wall-Clock Runtime (seconds)			SPPL Speedup Factor	
				FairSquare	VeriFair	SPPL	vs. FairSquare	vs. VeriFair
DT ₄	Independent	15	Unfair	1.4	16.0	0.01	140x	1600x
	Bayes Net. 1	25	Unfair	2.5	1.27	0.03	83x	42x
	Bayes Net. 2	29	Unfair	6.2	0.91	0.03	206x	30x
DT ₁₄	Independent	32	Fair	2.7	105	0.03	90x	3500x
	Bayes Net. 1	46	Fair	15.5	152	0.07	221x	2171x
	Bayes Net. 2	50	Fair	70.1	151	0.08	876x	1887x
DT ₁₆	Independent	36	Fair	4.1	13.6	0.03	136x	453x
	Bayes Net. 1	49	Unfair	12.3	1.58	0.08	153x	19x
	Bayes Net. 2	53	Unfair	30.3	2.02	0.08	378x	25x
DT ₁₆ ^α	Independent	62	Fair	5.1	2.01	0.06	85x	33x
	Bayes Net. 1	58	Fair	15.4	21.6	0.12	128x	180x
	Bayes Net. 2	45	Fair	53.8	24.5	0.12	448x	204x
DT ₄₄	Independent	93	Fair	15.6	23.1	0.05	312x	462x
	Bayes Net. 1	109	Unfair	264.1	19.8	0.09	2934x	220x
	Bayes Net. 2	113	Unfair	t/o	20.1	0.09	—	223x

6 Evaluation

We implemented a prototype of SPPL² and evaluated its performance on benchmark problems from the literature. Sec. 6.1 compares the runtime of verifying fairness properties of decision trees using SPPL to FairSquare [2] and VeriFair [4], two state-of-the-art fairness verification tools. Sec. 6.2 compares the runtime of conditioning and querying probabilistic programs using SPPL to PSI [23], a state-of-the-art tool for exact symbolic inference. Sec. 6.3 compares the runtime of computing exact rare event probabilities in SPPL to sampling-based estimation in BLOG [40]. Experiments were run on Intel i7-8665U 1.9GHz CPU with 16GB RAM.

6.1 Fairness Benchmarks

Characterizing the fairness of classification algorithms is a growing application area in machine learning [21]. Recently, Albarghouthi et al. [2] precisely cast the problem of verifying the fairness of a classifier in terms of computing ratios of conditional probabilities in a probabilistic program that specifies the data generating and classification processes. Briefly, if (i) D is a decision program that classifies whether applicant A should be hired; (ii) H is a population program that generates random applicants; and (iii) ϕ_m (resp. ϕ_q) is a predicate on A that is true if the applicant is a minority (resp. qualified), then D is ϵ -fair on H (where $\epsilon > 0$) if

$$\frac{\Pr_{A \sim H} [D(A) \mid \phi_m(A) \wedge \phi_q(A)]}{\Pr_{A \sim H} [D(A) \mid \neg \phi_m(A) \wedge \phi_q(A)]} > 1 - \epsilon, \quad (7)$$

i.e., the probability of hiring a qualified minority applicant is ϵ -close to that of hiring a qualified non-minority applicant.

In this evaluation, we compare the runtime needed by SPPL to obtain a fairness judgment (Eq. (7)) for machine-learned decision and population programs against the FairSquare [2] and VeriFair [4] solvers. We evaluate performance on the decision tree benchmarks from Albarghouthi et al. [2, Sec. 6.1],

which are one-third of the full benchmark set. SPPL cannot solve the neural network and support-vector machine benchmarks, as they contain multivariate transforms which do not have exact tractable solutions and are ruled out by the SPPL restriction (R3). FairSquare and VeriFair can express these benchmarks as they have approximate inference.

Table 2 shows the results. The first column shows the decision making program (DT _{n} means “decision tree” with n conditionals); the second column shows the population model used to generate data; the third column shows the lines of code (in SPPL); and the fourth column shows the result of the fairness analysis (FairSquare, VeriFair, and SPPL produce the same judgment on all fifteen benchmarks). The remaining columns show the runtime and speedup factors. We note that SPPL, VeriFair, and FairSquare are all implemented in Python, which allows for a fair comparison. The measurements indicate that SPPL consistently obtains probability estimates in milliseconds, whereas the two baselines can each require over 100 seconds. The SPPL speedup factors are up to 3500x (vs. VeriFair) and 2934x (vs. FairSquare). We further observe that the runtimes in FairSquare and VeriFair vary significantly. For example, VeriFair uses rejection sampling to estimate Eq. (7) with a stopping rule to determine when the estimate is close enough, leading to unpredictable runtime (e.g., >100 seconds for DT₁₄ but <1 second for DT₄, Bayes Net. 2). FairSquare, which uses symbolic volume computation and hyperrectangle sampling to approximate Eq. (7), is faster than VeriFair in some cases (e.g., DT₁₄), but times out in others (DT₄₄, Bayes Net. 2). In contrast, SPPL, computes exact probabilities for Eq. (7) and its runtime does not vary significantly across the various benchmark problems. The performance–expressiveness trade-off here is that SPPL computes exact probabilities and is substantially faster on the decision tree problems that it can express. FairSquare and VeriFair compute approximate probabilities that enable them to express more fairness problems, at the cost of a higher and less predictable runtime on the decision trees.

²Available in supplement and online at <https://github.com/probcomp/sppl>.

Figure 7. Comparison of multi-stage and single-stage inference workflows. In SPPL, modeling, observing data, and querying are separated into distinct stages, enabling substantial efficiency gains from computation reuse across multiple datasets and/or queries, as opposed to single-stage workflows such as in PSI which combine all these tasks into one large symbolic computation (daggers/colors used in Table 4).

6.2 Comparison to Symbolic Integration

We next compare SPPL to PSI [23], a state-of-the-art symbolic inference engine, on benchmark problems that include discrete, continuous, and transformed random variables. PSI can express more inference problems than SPPL, as it uses general computer algebra without having restrictions (R3) and (R4) in SPPL. As a result, SPPL can solve 14/21 benchmarks listed in [23, Table 1]. We first discuss key architecture novelties in SPPL that contribute to its performance gains.

Workflow Comparison In SPPL, the multi-stage modeling and inference workflow (Fig. 7a) involves three steps that reflect the key elements of a Bayesian inference problem:

- (S1) Translating the model program into a prior SP S .
- (S2) Conditioning S on data to obtain a posterior SP S' .
- (S3) Querying S' , using, e.g., `prob` or `simulate`.

An advantage of this multi-stage workflow is that multiple tasks can be run at a given stage without rerunning previous stages. For example, multiple datasets can be observed in (S2) without translating the prior expression in (S1) once per dataset; and, similarly, multiple queries can be run in (S3) without conditioning on data in (S2) once per query. In contrast, PSI adopts a single-stage workflow (Fig. 7b), where a single program contains the prior distribution over variables, “observe” (i.e., “condition”) statements for conditioning on a dataset, and a “return” statement for the query. PSI converts the program into a symbolic expression for the distribution over the return value: if this expression is “complete” (i.e., no unevaluated symbolic integrals) it can be used to obtain interpretable answers (e.g., for plotting or tabulating); otherwise, the result is “partial” and is too complex to be used for practical purposes. A consequence of the single-stage workflow in a system like PSI is that the entire solution is recomputed from scratch on a per-dataset or per-query basis.

Table 3. Distribution of end-to-end inference runtime for four benchmarks from Table 4 using PSI [23] and SPPL.

Benchmark	Mean/Std Runtime (sec/sec)	
	PSI	SPPL
Digit Recognition	26.5/1.3	15.9/0.5
Markov Switching	22.5/3.8	0.1/0.0
Student Interviews	539/663	7.8/0.2
Clinical Trial	107.3/153.2	12.7/0.3

Runtime Comparison Table 4 compares the runtime of SPPL and PSI on seven benchmark problems: Digit Recognition [23]; TrueSkill [36]; Clinical Trial [23]; Gamma transforms (described below); Student Interviews [36] (two variants); and Markov Switching (two variants, from Sec. 2.2). The second column shows the distributions in each benchmark, which include continuous, discrete, and transformed variables. The third column shows the number of datasets on which to condition the program. The next three columns show the time needed to translate the program (stage (S1)), condition the program on a dataset (stage (S2)), and query the posterior (stage (S3))—entries in the latter two columns are written as $n \times t$, where n is the number of datasets and t the average time per dataset. For PSI: (i) modeling and observing data are a single stage, shown in the merged gray cell; and (ii) querying the posterior times out whenever the system returns a result with unsimplified integrals ($\times$). The last column shows the overall runtime for solving all n tasks.

For benchmarks that both systems solve completely, SPPL realizes speedups between 3x (Digit Recognition) to 3600x (Markov Switching₃). In addition, the measurements show the advantage of our multi-stage workflow; for example, in TrueSkill, which uses a Poisson–Binomial distribution, SPPL translation (3.4 seconds) is more expensive than both conditioning on data (0.7 seconds) and querying (0.1 seconds), which highlights the benefit of amortizing the translation cost over several datasets or queries. In PSI, solving TrueSkill takes 2×41.6 seconds, but the solution contains unsimplified integrals and is thus unusable. The Markov Switching and Student Interviews benchmarks show that PSI may not perform well in the presence of many discrete random variables.

The Gamma Transform benchmark tests the robustness of many-to-one transforms of random variables (Lst. 1b), where $X \sim \text{Gamma}(3, 1)$; $Y = 1/\exp X^2$ if $X < 1$ and $Y = 1/\ln X$ otherwise; and $Z = -Y^3 + Y^2 + 6Y$. Each of the $n = 5$ datasets specifies a different constraint $\phi(Z)$ and a query about the posterior $Y | \phi(Z)$, which needs to compute and integrate out $X | \phi(Z)$. PSI reports that there is an error in its answer for all five datasets, whereas SPPL, using the symbolic transform solver from Appx. C.2, solves all five problems effectively.

Table 3 compares the runtime variance using SPPL and PSI for four of the benchmarks in Table 4, repeating one query over 10 datasets. In all benchmarks, the SPPL variance is lower than that of PSI, with a maximum standard deviation $\sigma = 0.5$ sec. In contrast, the spread of PSI runtime is high

Table 4. Runtime comparison of PSI [23] and SPPL.

Benchmark	Distribution	Datasets $\diamond$	System	Wall-Clock Runtime of Inference Stages			Overall Time	Legend
				Translation $\dagger$	Conditioning $\ddagger$	Querying $\star$		
Digit Recognition	$C \times B^{784}$	10	SPPL	6.9 sec	10×7.7 sec	$10 \times (<0.01)$ sec	84 sec	A: Atomic B: Bernoulli Be: Beta Bi: Binomial C: Categorical N: Normal G: Gamma P: Poisson T: Transform U: Uniform
			PSI	10×24.3 sec		$10 \times (<0.01)$ sec	244 sec	
TrueSkill	$P \times B^2$	2	SPPL	3.4 sec	2×0.7 sec	2×0.1 sec	4.9 sec	$\diamond$ Number of distinct datasets on which to condition the benchmark program. $\dagger, \ddagger$ Runtime of first two phases in Fig. 7; PSI implements these phases in a single computation. $\star$ Runtime of final phase in Fig. 7; same query used for all datasets of a given benchmark program.
			PSI	2×41.60 sec		∞	∞	
Clinical Trial	$B \times U^3 \times B^{50} \times B^{50}$	10	SPPL	9.5 sec	10×2.2 sec	$10 \times (<0.01)$ sec	31 sec	h/m High-Memory o/m Out-of-Memory i/e Integration Error ∞ Unsimplified Symbolic Integrals $\emptyset$ No Value
			PSI	10×107.3 sec		$10 \times (<0.01)$ sec	1073 sec	
Gamma Transforms	$G \times T \times (T+T)$	5	SPPL	0.02 sec	5×0.52 sec	5×0.03 sec	2.8 sec	
			PSI	5×0.68 sec; i/e		∞	∞	
Student Interviews ₂	$P \times B^2 \times B^4 \times (A+Be)^2$	10	SPPL	4.0 sec	10×0.7 sec	10×0.2 sec	13.5 sec	
			PSI	10×540 sec; h/m (35GB)		∞	∞	
Student Interviews ₁₀	$P \times B^{10} \times B^{20} \times (A+Be)^{10}$	10	SPPL	24.6 sec	10×3.9 sec	10×1.2 sec	75 sec	
			PSI	o/m (64GB+)		∞	∞	
Markov Switching ₃	$B \times B^3 \times N^3 \times P^3$	10	SPPL	0.05 sec	$10 \times (<0.01)$ sec	$10 \times (<0.01)$ sec	0.5 sec	
			PSI	10×182.9 sec		$10 \times (<0.01)$ sec	1829 sec	
Markov Switching ₁₀₀	$B \times B^{100} \times N^{100} \times P^{100}$	10	SPPL	4.1 sec	10×6.5 sec	10×0.5 sec	74 sec	
			PSI	o/m (64GB+)		∞	∞	

for Student Interviews ($\sigma = 540$ sec, range 64–1890 sec) and Clinical Trial ($\sigma = 153$ sec, range 2.75–470 sec). In PSI, the symbolic analyses are sensitive to the numeric values in the dataset, leading to unpredictable runtime across different datasets, even for a fixed query pattern. In SPPL, the runtime depends only on the query pattern not the observed data and therefore behaves predictably across different datasets.

As with the fairness benchmarks in Sec. 6.1, PSI trades off expressiveness with efficacy on tractable problems, and our measurements show that its runtime and memory do not scale well or are unpredictable on benchmarks that SPPL solves very efficiently. Moreover, the evaluations show that PSI can return unusable inference results to the user and that it needs to recompute entire symbolic solutions from scratch for each new dataset or query, whereas SPPL is less expressive than PSI but carries neither of these limitations.

6.3 Comparison to Sampling-Based Estimates

We next compare the runtime and accuracy of estimating probabilities of rare events in a canonical Bayesian network [33] using SPPL and BLOG [40]. As discussed by Koller and Friedman [33, Sec 12.13], rare events are the rule, not the exception, in many applications, as the probability of a predicate $\phi(X)$ decreases exponentially with the number of observed variables in X . Small estimation errors can magnify substantially when, e.g., taking ratios of probabilities.

In Fig. 8, each subplot shows the runtime and probability estimates for a low-probability predicate ϕ . In BLOG, the rejection sampler estimates the probability of ϕ by computing the fraction of times it holds in a size n i.i.d. random sample from the prior. The horizontal red line shows the “ground truth” probability. The x marker shows the runtime needed by SPPL to (exactly) compute the probability and the dots show the estimates from BLOG with increasing runtime (i.e., more samples n). SPPL consistently returns an exact answer in less than 2ms. The accuracy of BLOG estimates improve as the runtime increases: by the strong law of large numbers, these estimates converge to the true value, but

the fluctuations for any single run can be large (the standard error decays as $1/\sqrt{n}$). Each “jump” corresponds to a new sample $X^{(i)}$ that satisfies $\phi(X^{(i)})$, which increases the estimate. Without ground truth, it is hard to predict how much computation is needed for BLOG to obtain accurate results: estimates for predicates with $\log p = -12.73$ and $\log p = -17.32$ did not converge within the allotted time, while those for $\log p = -14.48$ converged after 180 seconds.

7 Related Work

SPPL is distinguished by being the first system to deliver exact symbolic inference by translating probabilistic programs to sum-product expressions, which extend and generalize sum-product networks. We briefly discuss related approaches.

Symbolic Integration Several systems deliver exact inference by translating a probabilistic program and observed dataset into a symbolic expression whose solution is the answer to the query [6, 10, 23, 43, 69]. Our approach to exact inference, which uses sum-product expressions instead of general computer algebra, enables effective performance on a range of models and queries, primarily at the expense of the expressiveness of the language on continuous priors. The state-of-the-art solver, PSI [23], can effectively solve many inference problems that SPPL cannot express due to restrictions (R1)–(R4), including higher-order programs [24]. However, comparisons on benchmarks that SPPL targets (Sec. 6.2) find PSI has less scalable and higher variance runtime, and can return partial results with unsimplified symbolic integrals. In contrast, SPPL exploits conditional independences, when they exist, to improve scalability (Sec. 5.1) and delivers complete, usable answers to users. Moreover, SPPL’s multi-stage workflow (Fig. 7) allows expensive computations such as translation and conditioning to be amortized over multiple datasets or queries, whereas PSI recomputes the symbolic solution from scratch each time. Hakaru [43] is a symbolic solver that delivers exact inference in a multi-stage workflow based on program transformations, and can disintegrate against a variety of base measures [44]. This paper compares

Figure 8. Runtime comparison for computing probabilities using exact inference in SPPL and rejection sampling in BLOG.

against PSI because the reference Hakaru implementation crashes or delivers incorrect or partial results on several benchmark problems [23, Table 1], and, as mentioned by the developers, does not support constructs such as arrays needed to support dozens or hundreds of observations.

Symbolic Execution and Volume Computation: Previous work has addressed the problem of computing the probability of a predicate by integrating a distribution defined by a program [2, 25, 55, 61]. For example, Geldenhuys et al. [25] present a probabilistic symbolic execution technique that uses model counting to compute path probabilities, assuming that all program variables are discrete and uniformly distributed. While SPPL can model a variety of distributions, due to restriction (R3) it only supports predicates that specify rectangular regions, whereas several of the aforementioned systems can (approximately) handle non-rectangular regions. More specifically, predicates in SPPL may include combinations of nonlinear transforms, each of a single variable, which are solved into linear expressions that specify unions of disjoint hyperrectangles (Appx. C.2). Table 2 shows that SPPL delivers substantial speedup on the hyperrectangular regions specified by the important class of decision trees, which are widely used in interpretable machine learning applications.

Sum-Product Networks: The SPFlow library [41] is an object-oriented “graphical model toolkit” in Python for constructing and querying sum-product networks. SPPL leverages a new and more general sum-product representation (Lst. 1) and solves probability and conditioning queries that are not supported by SPFlow (Thm. 4.1), which include mixed random variables, numeric transforms, and logical predicates with set-valued constraints. In addition, we introduce a novel translation strategy (Sec. 5) that allows users to specify models as generative code in a PPL (using e.g., variables, arrays, arithmetic and logical expressions, loops, branches) without having to manually manipulate low-level data structures. “Factored sum-product networks” [58] have been used as intermediate representations for converting a probabilistic program and any functional interpreter into a system of equations whose solution is the marginal probability of the program’s return value. These algorithms handle recursive procedures and leverage dynamic programming, but only apply to discrete variables, cannot handle transforms, and require solving fixed-points. Moreover, they have not been quantitatively evaluated on PPL benchmark problems.

Weighted Model Counting/Integration: A common approach to probabilistic inference is using algorithmic reductions from probabilistic programs to weighted-model counting (WMC) or integration (WMI) via knowledge compilation [5, 15, 19, 22, 66]. For example, Symbo [70] leverages WMI for exact inference in hybrid domains, using sentinel decision diagrams as the representation and the PSI solver to symbolically integrate over continuous variables. Dice [30] leverages WMC for scaling exact inference in discrete probabilistic programs and uses binary decision diagram representations that automatically exploit program structure to factorize inference. The representations in Dice enable substantial computation reuse for querying and/or conditioning, such as computing “all-marginal” probabilities by reusing the same compiled representation multiple times. SPPL also leverages factorization and computation reuse, but uses a different representation based on sum-product expressions that handle additional computations such as numeric transforms and continuous and mixed-type random variables.

Probabilistic Program Synthesis: Existing PPL synthesis systems for tabular data [13, 53] produce programs in languages that are subsets of SPPL, which enable automatic synthesis of full SPPL programs from data. SPPL can also unify and extend custom PPL query engines used in these systems for tasks such as similarity search and dependence detection [49, 50, 52]. It may also be fruitful to use structure discovery methods for time series [1, 54] or relational data [32] to synthesize SPPL programs for these domains.

8 Conclusion

We have presented SPPL, a new system that automatically delivers exact answers to a range of probabilistic inference queries. A key insight in SPPL is to impose restrictions on probabilistic programs that enable them to be translated to sum-product expressions, which are highly effective representations for inference. Our evaluation highlights the efficacy of SPPL on inference tasks in the literature and underscores the importance of key design decisions, including the multi-stage inference workflow and techniques used to build compact expressions by exploiting probabilistic structure. In addition to its efficacy as a standalone language, we further anticipate that SPPL could be useful as an embedded domain-specific language within more expressive PPLs, combining the benefits of exact and approximate inference.

References

- [1] Amr Ahmed and Eric Xing. 2008. Dynamic Non-Parametric Mixture Models and The Recurrent Chinese Restaurant Process: With Applications to Evolutionary Clustering. In *SDM 2008: Proc. 2008 SIAM Int. Conf. Data Min.* SIAM, Philadelphia, PA, USA, 219–230. <https://doi.org/10.1137/1.9781611972788.20>
- [2] Aws Albarghouthi, Loris D'Antoni, Samuel Drews, and Aditya V. Nori. 2017. FairSquare: Probabilistic Verification of Program Fairness. *Proc. ACM Program. Lang.* 1, OOPSLA, Article 80 (Oct. 2017), 30 pages. <https://doi.org/10.1145/3133904>
- [3] Stefan Arnborg, Derek G Corneil, and Andrzej Proskurowski. 1987. Complexity of Finding Embeddings in a k-Tree. *SIAM J. Alg. Disc. Meth.* 8, 2 (April 1987), 277–284. <https://doi.org/10.1137/0608024>
- [4] Osbert Bastani, Xin Zhang, and Armando Solar-Lezama. 2019. Probabilistic Verification of Fairness Properties via Concentration. *Proc. ACM Program. Lang.* 3, OOPSLA, Article 118 (Oct. 2019), 27 pages. <https://doi.org/10.1145/3360544>
- [5] Vaishak Belle, Andrea Passerini, and Guy Van den Broeck. 2015. Probabilistic Inference in Hybrid Domains by Weighted Model Integration. In *IJCAI 2015: Proc. 24th Int. Jt. Conf. Artif. Intell.* International Joint Conferences on Artificial Intelligence, Cham, 2770–2776.
- [6] Sooraj Bhat, Johannes Borgström, Andrew D. Gordon, and Claudio Russo. 2013. Deriving Probability Density Functions from Probabilistic Functional Programs. In *TACAS 2013: Proc. 19th Int. Conf. Tools Alg. Constr. Anal. Syst.* Springer, Berlin, 508–522. https://doi.org/10.1007/978-3-642-36742-7_35
- [7] Patrick Billingsley. 1986. *Probability and Measure* (2nd ed.). John Wiley & Sons, New York.
- [8] Andrew D. Bolton, Martin Haesemeyer, Josua Jordi, Ulrich Schaechtle, Feras A. Saad, Vikash K. Mansinghka, Joshua B. Tenenbaum, and Florian Engert. 2019. Elements of a Stochastic 3D Prediction Engine in Larval Zebrafish Prey Capture. *eLife* 8 (Nov. 2019), e51975. <https://doi.org/10.7554/eLife.51975>
- [9] Stephen P. Brooks and Andrew Gelman. 1998. General Methods for Monitoring Convergence of Iterative Simulations. *J. Comput. Graph. Stat.* 7, 4 (1998), 434–455. <https://doi.org/10.1080/10618600.1998.10474787>
- [10] Jacques Carette and Chung-chieh Shan. 2016. Simplifying Probabilistic Programs Using Computer Algebra. In *PADL 2016: Proc. 18th Int. Sympos. Pract. Asp. Declar. Lang.* Springer, Cham, 135–152. https://doi.org/10.1007/978-3-319-28228-2_9
- [11] Bob Carpenter, Andrew Gelman, Matthew Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. Stan: A Probabilistic Programming Language. *J. Stat. Soft.* 76, 1 (Jan. 2017), 1–32. <https://doi.org/10.18637/jss.v076.i01>
- [12] Venkat Chandrasekaran, Nathan Srebro, and Prahladh Harsha. 2008. Complexity of Inference in Graphical Models. In *UAI 2008: Proc. 24th Annu. Conf. Uncertain. Artif. Intell.* AUAI Press, Arlington, VA, USA, 70–78.
- [13] Sarah Chasins and Phitchaya M. Phothilimthana. 2017. Data-driven Synthesis of Full Probabilistic Programs. In *CAV 2017: Proc. 29th Int. Conf. Comput. Aided Verif.* Springer, Cham, 279–304. https://doi.org/10.1007/978-3-319-63387-9_14
- [14] Nick Chater, Joshua B. Tenenbaum, and Alan Yuille. 2006. Probabilistic Models of Cognition: Conceptual Foundations. *Trends Cogn. Sci.* 10, 7 (July 2006), 287–291. <https://doi.org/10.1016/j.tics.2006.05.007>
- [15] Mark Chavira and Adnan Darwiche. 2008. On Probabilistic Inference by Weighted Model Counting. *Artif. Intell.* 172, 6 (April 2008), 772–799. <https://doi.org/10.1016/j.artint.2007.11.002>
- [16] Gregory F. Cooper. 1990. The Computational Complexity of Probabilistic Inference Using Bayesian Belief Networks. *Artif. Intell.* 42, 2 (March 1990), 393–405. [https://doi.org/10.1016/0004-3702\(90\)90060-D](https://doi.org/10.1016/0004-3702(90)90060-D)
- [17] Marco F. Cusumano-Towner, Feras A. Saad, Alexander K. Lew, and Vikash K. Mansinghka. 2019. Gen: A General-Purpose Probabilistic Programming System with Programmable Inference. In *PLDI 2019: Proc. 40th ACM SIGPLAN Conf. Program. Lang. Des. Impl.* ACM, New York, NY, USA, 221–236. <https://doi.org/10.1145/3314221.3314642>
- [18] Paul Dagum and Michael Luby. 1993. Approximating Probabilistic Inference in Bayesian Belief Networks is NP-Hard. *Artif. Intell.* 60, 1 (March 1993), 141–153. [https://doi.org/10.1016/0004-3702\(93\)90036-B](https://doi.org/10.1016/0004-3702(93)90036-B)
- [19] Adnan Darwiche and Pierre Marquis. 2002. A Knowledge Compilation Map. *J. Artif. Intell. Res.* 17 (2002), 229–264. <https://doi.org/10.1613/jair.989>
- [20] Persi Diaconis and Donald Ylvisaker. 1979. Conjugate Priors for Exponential Families. *Ann. Statist.* 7, 2 (March 1979), 269–281. <https://doi.org/10.1214/aos/1176344611>
- [21] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through Awareness. In *ITCS 2012: Proc. 3rd Innov. Theor. Comput. Sci. Conf.* ACM, New York, NY, USA, 214–226. <https://doi.org/10.1145/2090236.2090255>
- [22] Daan Fierens, Guy Van den Broeck, Ingo Thon, Bernd Gutmann, and Luc De Raedt. 2011. Inference in Probabilistic Logic Programs using Weighted CNF's. In *UAI 2011: Proc. 27th Annu. Conf. Uncertain. Artif. Intell.* AUAI Press, Arlington, VA, USA, 211–220.
- [23] Timon Gehr, Sasa Misailovic, and Martin Vechev. 2016. PSI: Exact Symbolic Inference for Probabilistic Programs. In *CAV 2016: Proc. 28th Int. Conf. Comput. Aided Verif.* Springer, Cham, 62–83. https://doi.org/10.1007/978-3-319-41528-4_4
- [24] Timon Gehr, Samuel Steffen, and Martin Vechev. 2020. λPSI: Exact Inference for Higher-Order Probabilistic Programs. In *PLDI 2020: Proc. 41st ACM SIGPLAN Conf. Program. Lang. Des. Impl.* ACM, New York, NY, USA, 883–897. <https://doi.org/10.1145/3385412.3386006>
- [25] Jaco Geldenhuys, Matthew B. Dwyer, and Willem Visser. 2012. Probabilistic Symbolic Execution. In *ISSTA 2012: Proc. 2012 Int. Sympos. Soft. Test. Anal.* ACM, New York, NY, USA, 166–176. <https://doi.org/10.1145/2338965.2336773>
- [26] Robert Gens and Pedro Domingos. 2013. Learning the Structure of Sum-Product Networks. In *ICML 2013: Proc. 30th Int. Conf. Mach. Learn.* PMLR, 873–880.
- [27] Zoubin Ghahramani. 2015. Probabilistic Machine Learning and Artificial Intelligence. *Nature* 521, 7553 (May 2015), 452–459. <https://doi.org/10.1038/nature14541>
- [28] Noah Goodman, Vikash Mansinghka, Daniel M. Roy, Keith Bonawitz, and Joshua B. Tenenbaum. 2008. Church: A Language For Generative Models. In *UAI 2008: Proc. 24th Annu. Conf. Uncertain. Artif. Intell.* AUAI Press, Arlington, VA, USA, 220–229.
- [29] Andrew D. Gordon, Thomas A. Henzinger, Aditya V. Nori, and Sri-ram K. Rajamani. 2014. Probabilistic Programming. In *FOSE 2014: Future Soft. Eng. Proc.* ACM, New York, NY, USA, 167–181. <https://doi.org/10.1145/2593882.2593900>
- [30] Steven Holtzen, Guy Van den Broeck, and Todd Millstein. 2020. Scaling Exact Inference for Discrete Probabilistic Programs. *Proc. ACM Program. Lang.* 4, OOPSLA, Article 140 (Oct. 2020), 31 pages. <https://doi.org/10.1145/3133904>
- [31] Frederick Jelinek. 1997. *Statistical Methods for Speech Recognition*. MIT Press, Cambridge, MA, USA.
- [32] Charles Kemp, Joshua B. Tenenbaum, Thomas L. Griffiths, Takeshi Yamada, and Naonori Ueda. 2006. Learning Systems of Concepts with An Infinite Relational Model. In *AAAI 2006: Proc. 21st AAAI Conf. Artif. Intell.* AAAI Press, Palo Alto, CA, USA, 381–388.
- [33] Daphne Koller and Nir Friedman. 2009. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, Cambridge, MA, USA.
- [34] Christopher Krapu and Mark Borsuk. 2019. Probabilistic Programming: A Review for Environmental Modellers. *Environ. Model. Softw.* 114 (April 2019), 40–48. <https://doi.org/10.1016/j.envsoft.2019.01.014>
- [35] Tejas D. Kulkarni, Pushmeet Kohli, Joshua B. Tenenbaum, and Vikash Mansinghka. 2015. Picture: A Probabilistic Programming Language for Scene Perception. In *CVPR 2015: Proc. 2015 IEEE Conf. Comput. Vis.*

- Pattern Recognit.* IEEE Press, Piscataway, NJ, USA, 4390–4399.
- [36] Jacob Laurel and Sasa Misailovic. 2020. Continualization of Probabilistic Programs with Correction. In *ESOP 2020: Proc. 29th Eur. Sympos. Program. Springer*, Cham, 366–393. https://doi.org/10.1007/978-3-030-44914-8_14
 - [37] Sang-Woo Lee, Christopher Watkins, and Byoung-Tak Zhang. 2014. Non-Parametric Bayesian Sum-Product Networks. In *LTPM 2014: Workshop Learn. Tractable Probab. Model.*
 - [38] Alexander K. Lew, Marco F. Cusumano-Towner, Benjamin Sherman, Michael Carbin, and Vikash K. Mansinghka. 2020. Trace Types and Denotational Semantics for Sound Programmable Inference in Probabilistic Languages. *Proc. ACM Program. Lang.* 4, POPL, Article 19 (Jan. 2020), 32 pages. <https://doi.org/10.1145/3371087>
 - [39] Vikash K. Mansinghka, Ulrich Schaechtle, Shivam Handa, Alexey Radul, Yutian Chen, and Martin Rinard. 2018. Probabilistic Programming with Programmable Inference. In *PLDI 2018: Proc. 39th ACM SIGPLAN Conf. Program. Lang. Des. Impl.* ACM, New York, NY, USA, 603–616. <https://doi.org/10.1145/3192366.3192409>
 - [40] Brian Milch, Bhaskara Marthi, Stuart Russell, David Sontag, Daniel L. Ong, and Andrey Kolobov. 2005. BLOG: Probabilistic Models with Unknown Objects. In *IJCAI 2005: Proc. 19th Int. Jt. Conf. Artif. Intell.* International Joint Conferences on Artificial Intelligence, Cham, 1352–1359.
 - [41] Alejandro Molina et al. 2020. SPFlow: An Easy and Extensible Library for Deep Probabilistic Learning using Sum-Product Networks. (Jan. 2020). arXiv:1901.03704
 - [42] Kevin Murphy and Mark A. Paskin. 2002. Linear-Time Inference in Hierarchical HMMs. In *NIPS 2001: Adv. Neural Inf. Proc. Syst. 14*. MIT Press, Cambridge, MA, USA, 833–840.
 - [43] Praveen Narayanan, Jacques Carette, Wren Romano, Chung-chieh Shan, and Robert Zinkov. 2016. Probabilistic Inference by Program Transformation in Hakaru (System Description). In *FLOPS 2016: Proc. 13th Int. Sympos. Func. Log. Program.* Springer, Cham, 62–79. https://doi.org/10.1007/978-3-319-29604-3_5
 - [44] Praveen Narayanan and Chung-chieh Shan. 2020. Symbolic Disintegration with a Variety of Base Measures. *ACM Trans. Program. Lang. Syst.* 42, 2, Article 9 (May 2020), 60 pages. <https://doi.org/10.1145/3374208>
 - [45] Davide Nitti, Tinne De Laet, and Luc De Raedt. 2016. Probabilistic Logic Programming for Hybrid Relational Domains. *Mach. Learn.* 103 (April 2016), 407–449. <https://doi.org/10.1007/s10994-016-5558-8>
 - [46] Aditya Nori, Chung-Kil Hur, Sriram Rajamani, and Selva Samuel. 2014. R2: An Efficient MCMC Sampler for Probabilistic Programs. In *AAAI 2014: Proc. 28th AAAI Conf. Artif. Intell.* AAAI Press, 2476–2482.
 - [47] Hoifung Poon and Pedro Domingos. 2011. Sum-Product Networks: A New Deep Architecture. In *UAI 2011: Proc. 27th Annu. Conf. Uncertain. Artif. Intell.* AUAI Press, Arlington, VA, USA, 337–346.
 - [48] Fabrizio Riguzzi. 2018. *Foundations of Probabilistic Logic Programming: Languages, Semantics, Inference and Learning*. River Publishers, Delft, The Netherlands.
 - [49] Feras Saad, Leonardo Casarsa, and Vikash Mansinghka. 2017. Probabilistic Search For Structured Data Via Probabilistic Programming and Nonparametric Bayes. (April 2017). arXiv:1704.01087
 - [50] Feras Saad and Vikash Mansinghka. 2016. Probabilistic Data Analysis with Probabilistic Programming. (Aug. 2016). arXiv:1608.05347
 - [51] Feras Saad and Vikash Mansinghka. 2016. A Probabilistic Programming Approach to Probabilistic Data Analysis. In *NIPS 2016: Adv. Neural Inf. Proc. Syst. 29*. Curran Associates, Inc., Red Hook, NY, USA, 2011–2019.
 - [52] Feras Saad and Vikash Mansinghka. 2017. Detecting Dependencies in Sparse, Multivariate Databases using Probabilistic Programming and Non-parametric Bayes. In *AISTATS 2017: Proc. 20th Int. Conf. Artif. Intell. Stat.* PMLR, 632–641.
 - [53] Feras A. Saad, Marco F. Cusumano-Towner, Ulrich Schaechtle, Martin C. Rinard, and Vikash K. Mansinghka. 2019. Bayesian Synthesis of Probabilistic Programs for Automatic Data Modeling. *Proc. ACM Program. Lang.* 3, POPL, Article 36 (Jan. 2019), 32 pages. <https://doi.org/10.1145/3290350>
 - [54] Feras A. Saad and Vikash K. Mansinghka. 2018. Temporally-reweighted Chinese Restaurant Process Mixtures for Clustering, Imputing, and Forecasting Multivariate Time Series. In *AISTATS 2018: Proc. 21st Int. Conf. Artif. Intell. Stat.* PMLR, 755–764.
 - [55] Sriram Sankaranarayanan, Aleksandar Chakarov, and Sumit Gulwani. 2013. Static Analysis for Probabilistic Programs: Inferring Whole Program Properties from Finitely Many Paths. In *PLDI 2013: Proc. 34th ACM SIGPLAN Conf. Program. Lang. Des. Impl.* ACM, New York, NY, USA, 447–458. <https://doi.org/10.1145/2491956.2462179>
 - [56] David J. Spiegelhalter, A. Philip Dawid, Steffen L. Lauritzen, and Robert G. Cowell. 1993. Bayesian Analysis in Expert Systems. *Statist. Sci.* 8, 3 (Aug. 1993), 219–247. <https://doi.org/10.1214/ss/1177010888>
 - [57] Siddharth Srivastava, Nicholas Hay, Yi Wu, and Stuart Russell. 2017. The Extended Semantics for Probabilistic Programming Languages. In *PPS 2017: Workshop Probab. Program. Semant.*
 - [58] Andreas Stuhlmüller and Noah Goodman. 2012. A Dynamic Programming Algorithm for Inference in Recursive Probabilistic Programs. (Sept. 2012). arXiv:1206.3555
 - [59] Andrew Thomas. 1994. BUGS: A Static Modelling Package. *RTA/BCS Modul. Lang. Newsl.* 2 (1994), 36–38.
 - [60] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. 2005. *Probabilistic Robotics*. MIT Press, Cambridge, MA, USA.
 - [61] Neil Toronto, Jay McCarthy, and David Van Horn. 2015. Running Probabilistic Programs Backwards. In *ESOP 2015: Proc. 24th Eur. Sympos. Program.* Springer, Berlin, 53–79. https://doi.org/10.1007/978-3-662-46669-8_3
 - [62] Martin Trapp, Robert Peharz, Hong Ge, Franz Pernkopf, and Zoubin Ghahramani. 2019. Bayesian Learning of Sum-Product Networks. In *NeurIPS 2019: Adv. Neural Inf. Proc. Syst. 32*. Curran Associates, Inc., Red Hook, NY, USA, 6347–6358.
 - [63] F. Turbak, D. Gifford, and M.A. Sheldon. 2008. *Design Concepts in Programming Languages*. MIT Press, Cambridge, MA, USA.
 - [64] Antonio Vergari. 2021. *A Systematic View of the Literature on Probabilistic Circuits, Computational Graphs Encoding Tractable Probability Distributions*. Retrieved Mar 7, 2021 from <https://arranger1044.github.io/probabilistic-circuits/>
 - [65] Antonio Vergari, Alejandro Molina, Robert Peharz, Zoubin Ghahramani, Kristian Kersting, and Isabel Valera. 2019. Automatic Bayesian Density Analysis. In *AAAI 2019: Proc. 23rd AAAI Conf. Artif. Intell.* AAAI Press, Palo Alto, CA, USA, 5207–5215.
 - [66] Jonas Vlasselaer, Guy Van den Broeck, Angelika Kimmig, Wannes Meert, and Luc De Raedt. 2015. Anytime Inference in Probabilistic Logic Programs with Tp-Compilation. In *IJCAI 2015: Proc. 24th Int. Jt. Conf. Artif. Intell.* International Joint Conferences on Artificial Intelligence, Cham, 1852–1858.
 - [67] David Wingate and Theophane Weber. 2013. Automated Variational Inference in Probabilistic Programming. (Jan. 2013). arXiv:1301.1299
 - [68] Yi Wu, Siddharth Srivastava, Nicholas Hay, Simon Du, and Stuart Russell. 2018. Discrete-Continuous Mixtures in Probabilistic Programming: Generalized Semantics and Inference Algorithms. In *ICML 2018: Proc. 35th Int. Conf. Mach. Learn.* PMLR, 5343–5352.
 - [69] Jieyuan Zhang and Jingling Xue. 2019. Incremental Precision-Preserving Symbolic Inference for Probabilistic Programs. In *PLDI 2019: Proc. 40th ACM SIGPLAN Conf. Program. Lang. Des. Impl.* ACM, New York, NY, USA, 237–252. <https://doi.org/10.1145/3314221.3314623>
 - [70] Pedro Zuidberg Dos Martires, Anton Dries, and Luc De Raedt. 2019. Exact and Approximate Weighted Model Integration with Probability Density Functions Using Knowledge Compilation. In *AAAI 2019: Proc. 23rd AAAI Conf. Artif. Intell.* AAAI Press, Palo Alto, CA, USA, 7825–7833.