

PMAF: An Algebraic Framework for Static Analysis of Probabilistic Programs

Di Wang
Carnegie Mellon University

Jan Hoffmann
Carnegie Mellon University

Thomas Reps
University of Wisconsin
GrammaTech, Inc.

Abstract

Automatically establishing that a probabilistic program satisfies some property φ is a challenging problem. While a sampling-based approach—which involves running the program repeatedly—can *suggest* that φ holds, to establish that the program *satisfies* φ , analysis techniques must be used. Despite recent successes, probabilistic static analyses are still more difficult to design and implement than their deterministic counterparts. This paper presents a framework, called *PMAF*, for designing, implementing, and proving the correctness of static analyses of probabilistic programs with challenging features such as recursion, unstructured control-flow, divergence, nondeterminism, and continuous distributions. *PMAF* introduces *pre-Markov algebras* to factor out common parts of different analyses. To perform *interprocedural analysis* and to create *procedure summaries*, *PMAF* extends ideas from non-probabilistic interprocedural dataflow analysis to the probabilistic setting. One novelty is that *PMAF* is based on a semantics formulated in terms of a control-flow *hypergraph* for each procedure, rather than a standard control-flow graph. To evaluate its effectiveness, *PMAF* has been used to reformulate and implement existing *intraprocedural* analyses for Bayesian-inference and the Markov decision problem, by creating corresponding *interprocedural* analyses. Additionally, *PMAF* has been used to implement a new interprocedural *linear expectation-invariant analysis*. Experiments with benchmark programs for the three analyses demonstrate that the approach is practical.

Keywords Program analysis, probabilistic programming, expectation invariant, pre-Markov algebra

ACM Reference Format:

Di Wang, Jan Hoffmann, and Thomas Reps. 2018. PMAF: An Algebraic Framework for Static Analysis of Probabilistic Programs. In

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PLDI '18, June 18–22, 2018, Philadelphia, PA, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5698-5/18/06...\$15.00

<https://doi.org/10.1145/3192366.3192408>

PLDI '18: PLDI '18: ACM SIGPLAN Conference on Programming Language Design and Implementation, June 18–22, 2018, Philadelphia, PA, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3192366.3192408>

1 Introduction

Probabilistic programming is becoming increasingly popular because it provides a rich framework for implementing randomized algorithms [5], cryptography protocols [6], cognitive [39] models, and machine learning [37] algorithms. Static analysis of probabilistic programs has received a lot of attention [10, 11, 13–18, 25, 29–31, 33, 36, 48–50, 72, 79]. Unfortunately, analyses of probabilistic programs have usually been standalone developments, and it is not immediately clear how different techniques relate.

This paper presents a framework, which we call *PMAF* (for Pre-Markov Algebra Framework), for designing, implementing, and proving the correctness of static analyses of probabilistic programs. We show how several analyses that may appear to be quite different, can be formulated—and generalized—using *PMAF*. Examples include Bayesian inference [18, 30, 31], Markov decision problem with rewards [75], and probabilistic-invariant generation [13, 15, 49]

New constructs in probabilistic programs are of two kinds, to express *data randomness* (e.g., sampling) and *control-flow randomness* (e.g., probabilistic choice). To express both features, we introduce a new algebraic structure, called a *pre-Markov algebra*, which is equipped with operations corresponding to control-flow actions in probabilistic programs: *sequencing*, *conditional-choice*, *probabilistic-choice*, and *nondeterministic-choice*. *PMAF* is based on a new fixed-point semantics that models challenging features such as divergence, unstructured control-flow, nondeterminism, and continuous distributions. To establish correctness, we introduce *probabilistic abstractions* between two pre-Markov algebras that represent the concrete and abstract semantics.

Our work shows how, with suitable extensions, a blending of ideas from prior work on (i) static analysis of single-procedure probabilistic programs, and (ii) interprocedural dataflow analysis of standard (non-probabilistic) programs can be used to create a framework for interprocedural analysis of multi-procedure probabilistic programs. In particular,

- The semantics on which *PMAF* is based is an interpretation of the control-flow graphs (CFGs) for a program's

procedures. One insight is to treat each CFG as a *hyper-graph* rather than a standard graph.

- The abstract semantics is formulated so that the analyzer can obtain *procedure summaries*.

Hyper-graphs contain *hyper-edges*, each of which consists of one source node and possibly several destination nodes. Conditional-branching, probabilistic-branching, and nondeterministic-branching statements are represented by hyper-edges. In ordinary CFGs, nodes can also have several successors; however, the operator applied at a confluence point q when analyzing a CFG is join (\sqcup), and the paths leading up to q are analyzed *independently*. For reasons discussed in §2.3, PMAF is based on a backward analysis, so the confluence points represent the program’s branch points (i.e., for if-statements and while-loops). If the CFG is treated as a graph, join would be applied at each branch-node, and the subpaths from each successor would be analyzed independently. In contrast, when the CFG is treated as a hyper-graph, the operator applied at a probabilistic-choice node with probability p is $\lambda a. \lambda b. a_p \oplus b$ —where $_p \oplus$ is not join, but an operator that weights the two successor paths by p and $1 - p$. For instance, in Fig. 2(b), the hyper-edge $\langle v_0, \{v_1, v_5\} \rangle$ generates the inequality $\mathcal{A}[v_0] \sqsupseteq \mathcal{A}[v_1]_{0.75} \oplus \mathcal{A}[v_5]$, for some analysis \mathcal{A} . This approach allows the (hyper-)subpaths from the successors to be analyzed *jointly*.

To perform interprocedural analyses of probabilistic programs, we adopt a common practice from interprocedural analysis of standard non-probabilistic programs: the abstract domain is a *two-vocabulary* domain (each value represents an abstraction of a state transformer) rather than a *one-vocabulary* domain (each value represents an abstraction of a state). In the algebraic approach, an element in the algebra represents a two-vocabulary transformer. Elements can be “multiplied” by the algebra’s formal multiplication operator, which is typically interpreted as (an abstraction of) the reversal of transformer composition. The transformer obtained for the set of hyper-paths from the entry of procedure P to the exit of P is the summary for P .

In the case of loops and recursive procedures, PMAF uses widening to ensure convergence. Here our approach is slightly non-standard: we found that for some instantiations of the framework, we could improve precision by using different widening operators for loops controlled by conditional, probabilistic, and nondeterministic branches.

The main advantage of PMAF is that instead of starting from scratch to create a new analysis, you only need to instantiate PMAF with the implementation of a new pre-Markov algebra. To establish soundness, you have to establish some well-defined algebraic properties, and can then rely on the soundness proof of the framework. To implement your analysis, you can rely on PMAF to perform sound interprocedural analysis, with respect to the abstraction that you provided. The PMAF implementation supplies common parts of different static analyses of probabilistic programs, e.g., efficient

iteration strategies with widenings and interprocedural summarization. Moreover, any improvements made to the PMAF implementation immediately translate into improvements to *all of its instantiations*.

To evaluate PMAF, we created a prototype implementation, and reformulated two existing intraprocedural probabilistic-program analyses—the Bayesian-inference algorithm proposed by Claret et al. [18], and Markov decision problem with rewards [75]—to fit into PMAF: Reformulation involved changing from the one-vocabulary abstract domains proposed in the original papers to appropriate two-vocabulary abstract domains. We also developed a new program analysis: *linear expectation-invariant analysis* (LEIA). Linear expectation-invariants are equalities involving expected values of linear expressions over program variables.

A related approach to static analysis of probabilistic programs is *probabilistic abstract interpretation* (PAI) [25, 67–69], which lifts standard program analysis to the probabilistic setting. PAI is both general and elegant, but the more concrete approach developed in our work on PMAF has a couple of advantages. First, PMAF is algebraic and provides a simple and well-defined interface for implementing new abstractions. We provide an actual implementation of PMAF that can be easily instantiated to specific abstract domains. Second, PMAF is based on a different semantic foundation, which follows the standard interpretation of non-deterministic probabilistic programs in domain theory [27, 46, 47, 64, 65, 82].

The concrete semantics of PAI isolates probabilistic choices from the non-probabilistic part of the semantics by interpreting programs as distributions $P : \Omega \rightarrow (D \rightarrow D)$, where Ω is a probability space and $D \rightarrow D$ is the space of non-probabilistic transformers. As a result, the PAI interpretation of the following non-deterministic program is that with probability $\frac{1}{2}$, we have a program that non-deterministically returns 1 or 2; with probability $\frac{1}{4}$, we have a program that returns 1; and with probability $\frac{1}{4}$, a program that returns 2.

```

if ★ then if prob( $\frac{1}{2}$ ) then return 1 else return 2
           else if prob( $\frac{1}{2}$ ) then return 1 else return 2 fi

```

In contrast, the semantics used in PMAF resolves non-determinism on the outside, and thus the semantics of the program is that it returns 1 with probability $\frac{1}{2}$ and 2 with $\frac{1}{2}$. As a result, one can conclude that the expected return value r is 1.5. However, PAI—and every static analysis based on PAI—can only conclude $1.25 \leq r \leq 1.75$.

- Contributions.** Our work makes five main contributions:
- We present a new denotational semantics for probabilistic programs, which is capable of expressing several nontrivial features of probabilistic-programming languages.
 - We develop PMAF, an algebraic framework for static analyses of probabilistic programs. PMAF provides a novel approach to analyzing probabilistic programs with non-deterministic choice (§4.1 and §4.2) and general recursion

(§4.3) (as well as continuous sampling and unstructured control-flow).

- We show that two previous intraprocedural probabilistic program analyses can be reformulated to fit into PMAF, thereby creating new interprocedural analyses for such previous work.
- We develop a new program analysis, linear expectation-invariant analysis, by means of a suitable instantiation of PMAF. This analysis is more general than previous approaches to finding expectation invariants.
- We report on experiments with PMAF that show that the framework is easy to instantiate and works effectively. The experiments also show that linear expectation-invariant analysis can derive nontrivial invariants.

2 Overview

In this Section, we familiarize the reader with probabilistic programming, and briefly introduce two different static analyses of probabilistic programs: Bayesian inference and linear expectation invariant analysis. We then informally explain the main ideas behind our algebraic framework for analyzing probabilistic programs and show how it generalizes the aforementioned analyses.

2.1 Probabilistic Programming

Probabilistic programs contain two sources of randomness: (i) *data randomness*, i.e., the ability to draw random values from distributions, and (ii) *control-flow randomness*, i.e., the ability to branch probabilistically. A variety of probabilistic programming languages and systems has been proposed [12, 38, 54, 61, 63, 74]. In this paper, our prototypical language is imperative.

We use the Boolean program in Fig. 1a to illustrate data randomness. In the program, b_1 and b_2 are two Boolean-valued variables. The *sampling statement* $x \sim \text{Dist}(\theta)$ draws a value from a distribution Dist with a vector of parameters θ , and assigns it to the variable x , e.g., $b_1 \sim \text{Bernoulli}(0.5)$ assigns to b_1 a random value drawn from a Bernoulli distribution with mean 0.5. Intuitively, the program tosses two fair Boolean-valued coins repeatedly, until one coin is *true*.

We introduce control-flow randomness through the arithmetic program in Fig. 1b. In the program, x , y , and z are real-valued variables. As in the previous example, we have sampling statements, and $\text{Uniform}(l, r)$ represents a uniform distribution on the interval (l, r) . The *probabilistic choice* $\text{prob}(p)$ returns true with probability p and false with probability $1 - p$. Moreover, the program also exhibits *nondeterminism*, as the symbol \star stands for a *nondeterministic choice* that can behave like standard nondeterminism, as well as an arbitrary probabilistic choice [60, §6.6]. Intuitively, the program describes two players x and y playing a round-based game that ends with probability $\frac{1}{4}$ after each round. In each round, either player x or player y gains some reward that is uniformly distributed on $[0, 2]$.

<pre> $b_1 \sim \text{Bernoulli}(0.5);$ $b_2 \sim \text{Bernoulli}(0.5);$ while ($\neg b_1 \wedge \neg b_2$) do $b_1 \sim \text{Bernoulli}(0.5);$ $b_2 \sim \text{Bernoulli}(0.5)$ od </pre> <p style="text-align: center;">(a)</p>	<pre> while $\text{prob}(\frac{3}{4})$ do $z \sim \text{Uniform}(0, 2);$ if \star then $x := x + z$ else $y := y + z$ fi od </pre> <p style="text-align: center;">(b)</p>
---	---

Figure 1. (a) Boolean probabilistic program; (b) Arithmetic probabilistic program

2.2 Two Static Analyses

Bayesian inference (BI). Probabilistic programs can be seen as descriptions of probability distributions [12, 38, 63]. For a Boolean probabilistic program, such as the one in Fig. 1a, *Bayesian-inference analysis* [18] calculates the distribution over variable valuations at the end of the program, conditioned on the program terminating. The inferred probability distribution is called the *posterior probability distribution*. The program in Fig. 1a specifies the posterior distribution over the variables (b_1, b_2) given by: $\mathbb{P}[b_1 = \text{false}, b_2 = \text{false}] = 0$, and $\mathbb{P}[b_1 = \text{false}, b_2 = \text{true}] = \mathbb{P}[b_1 = \text{true}, b_2 = \text{false}] = \mathbb{P}[b_1 = \text{true}, b_2 = \text{true}] = \frac{1}{3}$. This distribution also indicates that the program terminates *almost surely*, i.e., the probability that the program terminates is 1.¹

Linear expectation invariant analysis (LEIA). Loop invariants are crucial to verification of imperative programs [28, 34, 43]. Although loop invariants for traditional programs are usually Boolean-valued expressions over program variables, real-valued invariants are needed to prove the correctness of probabilistic loops [55, 60]. Such *expectation invariants* are usually defined as random variables—specified as expressions over program variables—with some desirable properties [13, 14, 49]. In this paper, we work with a more general kind of expectation invariant, defined as follows:

Definition 2.1. For a program P , $\mathbb{E}[\mathcal{E}_2] \bowtie \mathcal{E}_1$ is called an *expectation invariant* if \mathcal{E}_1 and \mathcal{E}_2 are real-valued expressions over P 's program variables, \bowtie is one of $\{=, <, >, \leq, \geq\}$, and the following property holds: For any initial valuation of the program variables, the expected value of \mathcal{E}_2 in the final valuation (i.e., after the execution of P) is related to the value of \mathcal{E}_1 in the initial valuation by \bowtie .

We typically use variables with primes in \mathcal{E}_2 to denote the values in the final valuation. For example, for the program in Fig. 1b, $\mathbb{E}[x' + y'] = x + y + 3$, $\mathbb{E}[z'] = \frac{1}{4}z + \frac{3}{4}$, $\mathbb{E}[x'] \leq x + 3$, $\mathbb{E}[x'] \geq x$, $\mathbb{E}[y'] \leq y + 3$, and $\mathbb{E}[y'] \geq y$ are several linear expectation invariants, and our analysis can derive all of these automatically! The expectation invariant $\mathbb{E}[x' + y'] = x + y + 3$ indicates that the expected value of the total reward that the two players would gain is exactly 3.

¹In general, we work with with *subprobability distributions*, where the probabilities add up to strictly less than 1. In the case of a program that diverges with probability $p > 0$, the posterior distribution is a subprobability distribution in which the probabilities of the states sum up to $1 - p$.




Figure 2. (a) Control-flow hyper-graph of the program in Fig. 1a. (b) Control-flow hyper-graph of the program in Fig. 1b.

2.3 The Algebraic Framework

This section explains the main ideas behind PMAF, which is general enough to encode the two analyses from §2.2.

Data Randomness vs. Control-Flow Randomness. Our first principle is to make an *explicit separation between data randomness and control-flow randomness*. This distinction is intended to make the framework more flexible for analysis designers by providing multiple ways to translate the constructs of their specific probabilistic programming language into the constructs of PMAF. Analysis designers may find it useful to use the control-flow-randomness construct directly (e.g., “if **prob**(0.3) . . .”), rather than simulating control-flow randomness by data randomness (e.g., “ $p \sim \text{Uniform}(0, 1)$; if ($p < 0.3$) . . .”). For program analysis, such a simulation can lead to suboptimal results if the constructs used in the simulation require properties to be tracked that are outside the class of properties that a particular analysis’s abstract domain is capable of tracking. For example, if an analysis domain only keeps track of expectations, then analysis of “ $p \sim \text{Uniform}(0, 1)$ ” only indicates that $\mathbb{E}[p] = 0.5$, which does not provide enough information to establish that $\mathbb{P}[p < 0.3] = 0.3$ in the then-branch of “if ($p < 0.3$) . . .”. In contrast, when “**prob**(0.3) . . .” is analyzed in the fragment with the explicit control-flow-randomness construct (“if **prob**(0.3) . . .”) the analyzer can directly assign the probabilities 0.3 and 0.7 to the outgoing branches, and use those probabilities to compute appropriate expectations in the respective branches.

We achieve the separation between data randomness and control-flow randomness by capturing the different types of randomness in the graphs that we use for representing programs. In contrast to traditional program analyses, which usually work on control-flow graphs (CFGs), we use *control-flow hyper-graphs* to model probabilistic programs. Hyper-graphs are directed graphs, each edge of which (i) has one source and possibly multiple destinations, and (ii) has an associated *control-flow action*—either *sequencing*, *conditional-choice*, *probabilistic-choice*, or *nondeterministic-choice*. A traditional CFG represents a collection of execution paths, while in probabilistic programs, paths are no longer independent, and the program specifies probability distributions over the paths. It is natural to treat a collection of paths as a whole and define distributions over the collections. These kinds of collections can be precisely formalized as *hyper-paths* made up of *hyper-edges* in hyper-graphs.

Fig. 2 shows the control-flow hyper-graphs of the two programs in Fig. 1. Every edge has an associated action, e.g., the control-flow actions $\text{cond}[\neg b_1 \wedge \neg b_2]$, $\text{prob}[\frac{3}{4}]$, and ndet are conditional-choice, probabilistic-choice, and nondeterministic-choice actions. Data actions, like $x := x + z$ and $b_1 \sim \text{Bernoulli}(0.5)$, also perform a trivial control-flow action to transfer control to their one destination node.

Just as the control-flow graph of a procedure typically has a single entry node and a single exit node, a procedure’s control-flow hyper-graph also has a single entry node and a single exit node. In Fig. 2a, the entry and exit nodes are v_0 and v_3 , respectively; in Fig. 2b, the entry and exit nodes are v_0 and v_5 , respectively.

Backward Analysis. Traditional static analyses assign to a CFG node v either backward assertions—about the computations that can lead up to v —or forward assertions—about the computations that can continue from v [21, 23]. Backward assertions are computed via a forward analysis (in the same direction as CFG edges); forward assertions are computed via a backward analysis (counter to the flow of CFG edges).

Because we work with hyper-graphs rather than CFGs, from the perspective of a node v , there is a difference in how things “look” in the backward and forward direction: hyper-edges fan *out* in the forward direction. Hyper-edges can have two destination nodes, but only one source node.

The second principle of the framework is essentially dictated by this structural asymmetry: the framework *supports backward analyses that compute a particular kind of forward assertion*. In particular, the property to be computed for a node v in the control-flow hyper-graph for procedure P is (an abstraction of) a transformer that summarizes the transformation carried out by the hyper-graph fragment that extends from v to the exit node of P . It is possible to reason in the forward direction—i.e., about computations that lead up to v —but one would have to “break” hyper-paths into paths and “relocate” probabilities, which is more complicated than reasoning in the backward direction. The framework interprets an edge as a property transformer that computes properties of the edge’s source node as a function of properties of the edge’s destination node(s) and the edge’s associated action. These property transformers propagate information in a *hypergraph-leaf-to-hypergraph-root* manner, which is natural in hyper-graph problems. For example, standard formulations of interprocedural dataflow analysis [52, 57, 71, 80]

can be viewed as hyper-graph analyses, and propagation is performed in the leaf-to-root direction there as well.

Recall the Boolean program in Fig. 1a. Suppose that we want to perform BI to analyze $\mathbb{P}[b_1 = \text{true}, b_2 = \text{true}]$ in the posterior distribution. The property to be computed for a node will be a mapping from variable valuations to probabilities, where the probability reflects the chance that a given state will cause the program to terminate in the post-state ($b_1 = \text{true}, b_2 = \text{true}$). For example, the property that we would hope to compute for node v_1 is the function $\lambda(b_1, b_2).[b_1 \wedge b_2] + [\neg b_1 \wedge \neg b_2] \cdot \frac{1}{3}$, where $[\varphi]$ is an *Iverson bracket*, which evaluates to 1 if φ is true, and 0 otherwise.

Two-Vocabulary Program Properties. In the example of BI above, we observe that the property transformation discussed above is not suitable for *interprocedural* analysis. Suppose that (i) we want analysis results to tell us something about $\mathbb{P}[b_1 = \text{true}, b_2 = \text{true}]$ in the posterior distribution of the main procedure, but (ii) to obtain the answer, the analysis must also analyze a call to some other procedure Q . In the procedure main, the analysis is driven by the posterior-probability query $\mathbb{P}[b_1 = \text{true}, b_2 = \text{true}]$; in general, however, Q will need to be analyzed with respect to some other posterior probability (obtained from the distribution of valuations at the point in main just after the call to Q). One might try to solve this issue by analyzing each procedure multiple times with different posterior probabilities. However, in an infinite state space, this approach is no longer feasible.

Following common practice in interprocedural static analysis of traditional programs, the third principle of the framework is to work with *two-vocabulary program properties*. The property sketched in the BI example above is actually *one-vocabulary*, i.e., the property assigned to a control-flow node only involves the state at that node. In contrast, a two-vocabulary property at node v (in the control-flow hyper-graph for procedure P) should describe the state transformation carried out by the hyper-graph fragment that extends from v to the exit node of P .

For instance, LEIA assigns to each control-flow node a conjunction of expectation invariants, which relate the state at the node to the state at the exit node; consequently, LEIA deals with two-vocabulary properties. In §5, we show that we can reformulate BI to manipulate two-vocabulary properties. As in interprocedural dataflow analysis [22, 80], procedure summaries are used to interpret procedure calls.

Separation of Concerns. Our fourth principle—which is common to most analysis frameworks—is *separation of concerns*, by which we mean

Provide a declarative interface for a client to specify the program properties to be tracked by a desired analysis, but leave it to the framework to furnish the analysis implementation by which the analysis is carried out.

We achieve this goal by adopting (and adapting) ideas from previous work on *algebraic program analysis* [32, 76, 81].

Algebraic program analysis is based on the following idea:

Any static analysis method performs reasoning in some space of program properties and property transformers; such property transformers should obey algebraic laws.

For instance, the data action **skip**, which does nothing, can be interpreted as the *identity* element in an algebra of program-property transformers.

Concretely, our fourth principle has three aspects:

1. For our intended domain of probabilistic programs, identify an appropriate set of algebraic laws that hold for useful sets of property transformers.
2. Define a specific algebra \mathcal{A} for a program-analysis problem by defining a specific set of property transformers that obey the laws identified in item 1. Give translations from data actions and control-flow actions to such property transformers. (When such a translation is applied to a specific program, it sets up an equation system to be solved over \mathcal{A} .)
3. Develop a generic analysis algorithm that solves an equation system over any algebra that satisfies the laws identified in item 1.

Items 1 and 3 are tasks for us, the framework designers; they are the subjects of §3 and §4. Item 2 is a task for a client of the framework: examples are given in §5.

A client of the framework must furnish an *interpretation*—which consists of a *semantic algebra* and a *semantic function*—and a program. The semantic algebra consists of a *universe*, which defines the space of possible program-property transformers, and sequencing, conditional-choice, probabilistic-choice, and nondeterministic-choice operators, corresponding to control-flow actions. The semantic function is a mapping from data actions to the universe. (An interpretation is also called a *domain*.)

To address Item 3, our prototype implementation follows the standard *iterative* paradigm of static analysis [21, 51]: We first transform the control-flow hyper-graph into a system of inequalities, and then use a chaotic-iteration algorithm to compute a solution to it (e.g., [9]), which repeatedly applies the interpretation until a fixed point is reached (possibly using widening to ensure convergence). For example, the control-flow hyper-graph in Fig. 2b can be transformed into the system shown in Fig. 3, where $S(v) \in \mathcal{M}$ are elements in the semantic algebra; \sqsubseteq is the approximation order on \mathcal{M} ; $[\cdot]$ is the semantic function, which maps data actions to \mathcal{M} ; and $\underline{\cdot}$ is the transformer associated with the exit node.

The soundness of the analysis (with respect to a concrete semantics) is proved by (i) establishing an approximation relation between the concrete domain and the abstract domain; (ii) showing that the abstract semantic function approximates the concrete one; and (iii) showing that the abstract operators (sequencing, conditional-choice, probabilistic-choice, and nondeterministic-choice) approximate the concrete ones.

$$\begin{array}{ll}
S(v_0) \sqsupseteq \text{prob}[\frac{3}{4}](S(v_1), S(v_5)) & S(v_3) \sqsupseteq \text{seq}[x := x + z](S(v_0)) \\
S(v_1) \sqsupseteq \text{seq}[z \sim \text{Uniform}(0, 2)](S(v_2)) & S(v_4) \sqsupseteq \text{seq}[y := y + z](S(v_0)) \\
S(v_2) \sqsupseteq \text{ndet}(S(v_3), S(v_4)) & S(v_5) \sqsupseteq \perp
\end{array}$$

Figure 3. The system of inequalities corresponding to Fig. 2b

For BI, we instantiate our framework to give lower bounds on posterior distributions, using with an interpretation in which state transformers are probability matrices (see §5.1). For LEIA, we design an interpretation using a Cartesian product of polyhedra (see §5.3). Once the functions of the interpretations are implemented, and a program is translated into the appropriate hyper-graph, the framework handles the rest of the work, namely, solving the equation system.

3 Probabilistic Programs

In this Section, we first review the concepts of *hyper-graphs* [35] and introduce a probabilistic-program model based on them. Then we briefly sketch a new denotational semantics for our hyper-graph based imperative program model.

3.1 A Hyper-Graph Model of Probabilistic Programs

Definition 3.1 (Hyper-graphs). A *hyper-graph* H is a quadruple $\langle V, E, v^{\text{entry}}, v^{\text{exit}} \rangle$, where V is a finite set of nodes, E is a set of hyper-edges, $v^{\text{entry}} \in V$ is a distinguished *entry node*, and $v^{\text{exit}} \in V$ is a distinguished *exit node*. A *hyper-edge* is an ordered pair $\langle x, Y \rangle$, where $x \in V$ is a node and $Y \subseteq V$ is an ordered, non-empty set of nodes. For a hyper-edge $e = \langle x, Y \rangle$ in E , we use $\text{src}(e)$ to denote x and $\text{Dst}(e)$ to denote Y . Following the terminology from graphs, we say that e is an *outgoing* edge of x and an *incoming* edge of each of the nodes $y \in Y$. We assume that v^{entry} has no incoming edges, and v^{exit} has no outgoing edges.

Definition 3.2 (Probabilistic programs). A *probabilistic program* contains a finite set of procedures $\{H_i\}_{1 \leq i \leq n}$, where each procedure $H_i = \langle V_i, E_i, v_i^{\text{entry}}, v_i^{\text{exit}} \rangle$ is a control-flow hyper-graph in which each node except v_i^{exit} has exactly one outgoing hyper-edge. We assume that the nodes of each procedure are pairwise disjoint. To assign meanings to probabilistic programs modulo *data actions* \mathcal{A} and *logical conditions* \mathcal{L} , we associate with each hyper-edge $e \in E = \bigcup_{1 \leq i \leq n} E_i$ a control-flow action $\text{Ctrl}(e)$, where Ctrl is

$$\begin{array}{l}
\text{Ctrl} ::= \text{seq}[\text{act}] \text{ where } \text{act} \in \mathcal{A} \mid \text{call}[i] \text{ where } 1 \leq i \leq n \\
\quad \mid \text{cond}[\varphi] \text{ where } \varphi \in \mathcal{L} \mid \text{prob}[p] \text{ where } 0 \leq p \leq 1 \\
\quad \mid \text{ndet}
\end{array}$$

where the number of destination nodes $|\text{Dst}(e)|$ of a hyper-edge e is 1 if $\text{Ctrl}(e)$ is $\text{seq}[\text{act}]$ or $\text{call}[i]$, and 2 otherwise.

Fig. 2 shows two examples of hyper-graph-based probabilistic programs. See Fig. 4 for data actions \mathcal{A} and logical conditions \mathcal{L} that would be used for an arithmetic program like the one shown in Fig. 1b.

$$\begin{array}{l}
\mathcal{A} ::= x := e \mid x \sim \mathcal{D} \mid \text{skip} \mid \text{observe}(\varphi) \\
\varphi \in \mathcal{L} ::= \text{true} \mid \text{false} \mid e \bowtie u, \text{ where } \bowtie \in \{=, \leq, \geq\} \mid \neg \varphi \\
e, u \in \text{Exp} ::= x \mid c, \text{ where } c \in \mathbb{R} \mid e \bullet u, \text{ where } \bullet \in \{+, -, \times, /\} \\
x \in \text{Var} ::= x \mid y \mid z \mid \dots \\
\mathcal{D} \in \text{Dist} ::= \text{Uniform}(e, u) \mid \text{Gaussian}(e, u) \mid \dots
\end{array}$$

Figure 4. Examples of data actions and logical conditions

3.2 Background from Measure Theory

To define denotational semantics for probabilistic programs modulo data actions \mathcal{A} and logical conditions \mathcal{L} , we review some standard definitions from measure theory [7, 73].

A *measurable space* is a pair $\langle X, \Sigma \rangle$ where X is a non-empty set called the *sample space*, and Σ is a σ -*algebra* over X (i.e. a set of subsets of X which contains \emptyset and is closed under complement and countable union). A *measurable function* from a measurable space $\langle X_1, \Sigma_1 \rangle$ to another measurable space $\langle X_2, \Sigma_2 \rangle$ is a mapping $f : X_1 \rightarrow X_2$ such that for all $A \in \Sigma_2$, $f^{-1}(A) \in \Sigma_1$. The measurable functions from a measurable space $\langle X, \Sigma \rangle$ to the Borel space $\mathcal{B}(\mathbb{R}_{\geq 0})$ on nonnegative real numbers (the smallest σ -algebra containing all open intervals) is called Σ -*measurable*.

A *measure* μ on a measurable space $\langle X, \Sigma \rangle$ is a function from Σ to $[0, \infty]$ such that: (i) $\mu(\emptyset) = 0$, and (ii) for all pairwise-disjoint countable sequences of sets $A_1, A_2, \dots \in \Sigma$ (i.e., $A_i \cap A_j = \emptyset$ for all $i \neq j$) we have $\sum_{i=1}^{\infty} \mu(A_i) = \mu(\bigcup_{i=1}^{\infty} A_i)$. The measure μ is called a (*sub-probability*) *distribution* if $\mu(X) \leq 1$. A *measure space* is a triple $\mathcal{M} = \langle X, \Sigma, \mu \rangle$ where μ is a measure on the measurable space $\langle X, \Sigma \rangle$. The *integral* of a Σ -measurable function f over the measurable space $\mathcal{M} = \langle X, \Sigma, \mu \rangle$ can be defined following Lebesgue's theory and denoted either by $\int f d\mu$ or $\int f(x)\mu(dx)$. The *Dirac measure* $\delta(x)$ is defined as $\lambda A.[x \in A]$.

A (*sub-probability*) *kernel* from a measurable space $\langle X_1, \Sigma_1 \rangle$ to a measurable space $\langle X_2, \Sigma_2 \rangle$ is a function $\kappa : X_1 \times \Sigma_2 \rightarrow [0, 1]$ such that: (i) for each A in Σ_2 , the function $\lambda x.\kappa(x, A)$ is Σ_1 -measurable, and (ii) for each x in X_1 , the function $\kappa_x \stackrel{\text{def}}{=} \lambda A.\kappa(x, A)$ is a distribution on $\langle X_2, \Sigma_2 \rangle$. We write the integral of a measurable function $f : \Sigma_2 \rightarrow [0, \infty]$ with respect to the distribution in (ii) as $\int f(y)\kappa(x, dy)$.

3.3 A Denotational Semantics

The next step is to define semantics based on the control-flow hyper-graphs. We use a denotational approach because it abstracts away how a program is evaluated and concentrates only on the effect of the program. This property makes it suitable as a starting point for static analysis, which is aimed at reasoning about program properties.

We develop a new semantics for probabilistic programming by combining Borgström et al.'s distribution-based semantics using the concept of kernels from measure theory [8] and existing results on domain-theoretic probabilistic nondeterminism [27, 46, 47, 64, 65, 82]. This semantics can

describe several nontrivial constructs, including continuous sampling, nondeterministic choice, and recursion.

Three components are used to define the semantics:

- A *measurable space* $\mathcal{P} = \langle \Omega, \mathcal{F} \rangle$ over program states (e.g., finite mappings from program variables to values).
- A mapping from data act actions to *kernels* $\widehat{\text{act}} : \Omega \times \mathcal{F} \rightarrow \mathbb{R}$. The intuition to keep in mind is that $\widehat{\text{act}}(\omega, F)$ is the probability that the action, starting in state $\omega \in \Omega$, halts in a state that satisfies $F \in \mathcal{F}$ [56].²
- A mapping from logical conditions φ to measurable functions $\widehat{\varphi} : \Omega \rightarrow \{true, false\}$.

Example 3.3. For an arithmetic program with a finite set Var of program variables, Ω is defined as $\text{Var} \rightarrow \mathbb{R}$ and \mathcal{P} as the Borel space on Ω . Fig. 5 shows interpretation of the data actions and logical conditions in Fig. 4, where $e(\omega)$ evaluates expression e in state ω , $(x \mapsto v)\omega$ updates x in ω with v , and $\mu_{\mathcal{D}} : \mathcal{B}(\mathbb{R}) \rightarrow [0, 1]$ is the measure corresponding to the distribution \mathcal{D} on reals. Note that the action $\text{observe}(\varphi)$ performs conditioning on states that satisfy φ .

While distributions can be seen as one-vocabulary specifications, kernels are indeed two-vocabulary transformers over program states. When there is no nondeterminism, we can assign a kernel to every control-flow node. We can also define control-flow actions on kernels. A sequence of actions $\text{act}_1; \text{act}_2$ with kernels κ_1 and κ_2 , respectively, is modeled by their *composition*, denoted by $\kappa_1 \otimes \kappa_2$, which yields a new kernel defined as follows:³

$$\kappa_1 \otimes \kappa_2 \stackrel{\text{def}}{=} \lambda(x, A). \int \kappa_1(x, dy) \kappa_2(y, A). \quad (1)$$

The conditional-choice $\kappa_1 \diamond_{\varphi} \kappa_2$ is defined as a new kernel $\lambda(x, A). [\widehat{\varphi}(x)] \cdot \kappa_1(x, A) + [\neg \widehat{\varphi}(x)] \cdot \kappa_2(x, A)$. The probabilistic-choice $\kappa_1 \oplus_p \kappa_2$ is defined as a new kernel $\lambda(x, A). p \cdot \kappa_1(x, A) + (1 - p) \cdot \kappa_2(x, A)$.

Example 3.4. Consider the following program that models a variation on a geometric distribution.

```

n := 0;
while prob(0.9) do
  n := n + 1;  if n ≥ 10 then break else continue fi
od

```

Fig. 6 shows its control-flow hyper-graph. The assignment $n := n + 1$ is interpreted as a kernel $n := n + 1 = \lambda(\omega, F). [(n \mapsto \omega(n) + 1)\omega \in F]$. The comparison $n \geq 10$ is interpreted as a measurable function $\widehat{n \geq 10} = \lambda\omega. (\omega(n) \geq 10)$. Let K stand for 0.3486784401; the semantics assigned to node v_0 is

² As explained by Kozen [56], for finite or countable Ω , $\widehat{\text{act}}$ has a representation as a Markov transition matrix $M_{\widehat{\text{act}}}$, in which each entry $M_{\widehat{\text{act}}}(\omega, \omega')$ for a pair of states $\langle \omega, \omega' \rangle$ gives the probability that ω transitions to ω' under action act .

³ For finite or countable Ω , and the matrix representation described in footnote 2, the integral in Eqn. (1) degenerates to matrix multiplication [56].

$$\begin{array}{ll}
\widehat{x := e} = \lambda(\omega, F). [(x \mapsto e(\omega))\omega \in F] & \widehat{\text{true}} = \lambda\omega. true \\
\widehat{x \sim \mathcal{D}} = \lambda(\omega, F). \mu_{\mathcal{D}}(\{v \mid (x \mapsto v)\omega \in F\}) & \widehat{\text{false}} = \lambda\omega. false \\
\widehat{\text{observe}(\varphi)} = \lambda(\omega, F). \widehat{\varphi}(\omega) \cdot [\omega \in F] & \widehat{e \bowtie u} = \lambda\omega. [e(\omega) \bowtie u(\omega)] \\
\widehat{\text{skip}} = \lambda(\omega, F). [\omega \in F] & \widehat{\neg \varphi} = \lambda\omega. \neg \widehat{\varphi}(\omega)
\end{array}$$

Figure 5. Interpretation of actions and conditions

$$\lambda(\omega, F). \sum_{k=0}^9 (0.1 \times 0.9^k) \cdot [(n \mapsto k)\omega \in F] + K \cdot [(n \mapsto 10)\omega \in F].$$

When nondeterminism comes into the picture, we need to associate each control-flow node with a *collection* of kernels. In other words, we need to consider *powerdomains* [41] of kernels. We adopt Tix et al.'s constructions of probabilistic powerdomains [82], and extend them to work on kernels instead of distributions. We denote the set of feasible collections of kernels by $\mathbb{P}\Omega$, and the composition, conditional-choice, probabilistic-choice, and nondeterministic-choice operators on that by $\overline{\otimes}$, $\overline{\diamond}_{\varphi}$, $\overline{\oplus}_p$, and $\overline{\sqcup}$. $\mathbb{P}\Omega$ is also equipped with a partial order \sqsubseteq .

We reformulated distributions and kernels in a domain-theoretic way to adopt existing studies on powerdomains. We discuss the details of the construction of $\mathbb{P}\Omega$ in a companion paper [83]; the focus of this paper is static analysis and we will keep the domain-theoretic terminology to a minimum.

We adopted Hoare powerdomains and Smyth powerdomains [1, §6.2] over kernels. Kernels are ordered pointwise, i.e., $\kappa_1 \leq \kappa_2$ if and only if for all ω and F , $\kappa_1(\omega, F) \leq \kappa_2(\omega, F)$. The zero kernel $\lambda(\omega, F). 0$ is the bottom element of this order. Intuitively, the Hoare powerdomain is used for *partial* correctness, in which the order is set inclusion on the *lower* closures of the elements—because each downward-closed set contains a kernel that represents nontermination (i.e., the zero kernel), terminating and nonterminating executions cannot be distinguished. The Smyth powerdomain is used for *total* correctness: the order is reverse inclusion on the *upper* closures of the elements—nontermination is interpreted as the worst output, and the kernel that represents nontermination does not occur in an upward-closed set that represents the semantics of a terminating computation.

Given a probabilistic program $P = \{H_i\}_{1 \leq i \leq n}$, where $H_i = \langle V_i, E_i, v_i^{\text{entry}}, v_i^{\text{exit}} \rangle$, we want to define the semantics of each node v as a set of kernels that represent the effects from v to the exit node of the procedure that contains v . Let $\mathcal{S}(v) \in \mathbb{P}\Omega$ be the semantics assigned to the node v ; the following local properties should hold:




Figure 6. Control-flow hyper-graph of the program in Ex. 3.4

- if $e = \langle v, \{u_1, \dots, u_k\} \rangle \in E$, $\mathcal{S}\langle v \rangle = \widehat{Ctrl}(e)(\mathcal{S}\langle u_1 \rangle, \dots, \mathcal{S}\langle u_k \rangle)$, and
- otherwise, $\mathcal{S}\langle v \rangle = \perp_{\mathbb{P}}$.

The function \widehat{act} for the different kinds of control-flow actions is defined as follows:

$\widehat{seq[act]}(S_1) \stackrel{\text{def}}{=} \widehat{act} \otimes S_1$	$\widehat{cond[\varphi]}(S_1, S_2) \stackrel{\text{def}}{=} S_1 \varphi \diamond S_2$
$\widehat{call[i]}(S_1) \stackrel{\text{def}}{=} \mathcal{S}\langle v_i^{\text{entry}} \rangle \otimes S_1$	$\widehat{prob[p]}(S_1, S_2) \stackrel{\text{def}}{=} S_1 p \oplus S_2$
	$\widehat{ndet}(S_1, S_2) \stackrel{\text{def}}{=} S_1 \sqcup S_2$

Lemma 3.5. *The function F_P defined as*

$$\lambda \mathcal{S}. \lambda v. \begin{cases} \widehat{Ctrl}(e)(\mathcal{S}\langle u_1 \rangle, \dots, \mathcal{S}\langle u_k \rangle) & e = \langle v, \{u_1, \dots, u_k\} \rangle \in E \\ \perp_{\mathbb{P}} & \text{otherwise} \end{cases}$$

is ω -continuous on $\langle V \rightarrow \mathbb{P}\Omega, \sqsubseteq \rangle$, which is an ω -cpo with the least element $\lambda v. \perp_{\mathbb{P}}$.

A dot over an operator denotes its application pointwise. By Kleene's fixed-point theorem, we have

Theorem 3.6. $\text{lfp}_{\lambda v. \perp_{\mathbb{P}}}^{\dot{}} F_P$ exists for all prob. programs P .

Thus, the semantics of a node v is defined as $(\text{lfp}_{\lambda v. \perp_{\mathbb{P}}}^{\dot{}} F_P)(v)$.

4 Analysis Framework

To aid in creating abstractions of probabilistic programs, we first identify, in §4.1, some algebraic properties that underlie the mechanisms used in the semantics from §3.3. This algebra will aid our later definitions of abstractions in §4.2. We then discuss interprocedural analysis (in §4.3) and widening (§4.4).

4.1 An Algebraic Characterization of Fixpoint Semantics

In the denotational semantics, the concrete semantics is obtained by composing $\widehat{Ctrl}(e)$ operations along hyper-paths. Hence in the algebraic framework, the semantics of probabilistic programs is denoted by an *interpretation*, which consists of two parts: (i) a *semantic algebra*, which defines a set of possible program meanings, and which is equipped with sequencing, conditional-choice, probabilistic-choice, and nondeterministic-choice operators to compose these meanings, and (ii) a *semantic function*, which assigns a meaning to each basic program action.

The semantic algebras that we use—and the lattices used for abstract interpretation—are *pre-Markov algebras*:

Definition 4.1 (Pre- ω -continuous functions). A function $f : X \rightarrow Y$ between two ω -cpo's X and Y is *pre- ω -continuous* if it is monotone, and for every ω -chain $C \subseteq X$, $f(\text{sup}(C)) \leq \text{sup}(f(C))$.

Definition 4.2 (Pre-Markov algebras). A *pre-Markov algebra* (PMA) over a set of logical conditions \mathcal{L} is an 8-tuple $\mathcal{M} = \langle M, \sqsubseteq, \otimes, \varphi \diamond, p \oplus, \sqcup, \perp, \underline{1} \rangle$, where $\langle M, \sqsubseteq, \perp \rangle$ forms an ω -cpo with a least element \perp ; $\langle M, \otimes, \underline{1} \rangle$ forms a monoid (i.e., \otimes is an associative binary operator with $\underline{1}$ as its identity element); $\varphi \diamond$ is a binary operator parametrized by φ which is

a condition in \mathcal{L} ; $p \oplus$ is a binary operator parametrized by $p \in [0, 1]$; \sqcup is a binary operator that is idempotent, commutative, and associative; $\otimes, p \oplus, \varphi \diamond$, and \sqcup are pre- ω -continuous; and the following properties hold:

$$\begin{aligned} a \varphi \diamond b &\sqsubseteq a \sqcup b & a p \oplus b &\sqsubseteq a \sqcup b \\ a &\sqsubseteq a \varphi \diamond a & a &\sqsubseteq a p \oplus a \\ a &\sqsubseteq a_{\text{true}} \diamond b & a &\sqsubseteq a \perp \oplus b \\ a \varphi \diamond b &= b_{\neg\varphi} \diamond a & a p \oplus b &= b_{1-p} \oplus a \\ a \varphi \diamond (b \psi \diamond c) &= (a \varphi' \diamond b) \psi' \diamond c \\ &\text{where } \varphi = \varphi' \wedge \psi', \varphi \vee \psi = \psi' \\ a p \oplus (b q \oplus c) &= (a p' \oplus b) q' \oplus c \\ &\text{where } p = p'q', (1-p)(1-q) = 1-q' \end{aligned}$$

The precedence of the operators is that \otimes binds tightest, followed by $\varphi \diamond, p \oplus$, and \sqcup .

Remark 4.3. *These algebraic laws are not needed to prove soundness of the framework (stated in Thm. 4.8). These laws helped us when designing the abstract domains. Exploiting these algebraic laws to design better algorithms is an interesting direction for future work.*

Lemma 4.4. *The denotational semantics in §3.3 is a PMA $\mathcal{C} = \langle \mathbb{P}\Omega, \sqsubseteq_{\mathbb{P}}, \otimes, \varphi \diamond, p \oplus, \sqcup, \perp_{\mathbb{P}}, \underline{1}_{\mathbb{P}} \rangle$ (which we call the concrete domain for our framework).*

As is standard in abstract interpretation, the order on the algebra should represent an approximation order: $a \sqsubseteq b$ iff a is approximated by b (i.e., if a represents a more precise property than b).

Definition 4.5 (Interpretations). An interpretation is a pair $\mathcal{I} = \langle \mathcal{M}, [\cdot] \rangle$, where \mathcal{M} is a pre-Markov algebra, and $[\cdot] : \mathcal{A} \rightarrow \mathcal{M}$, where \mathcal{A} is the set of data actions for probabilistic programs. We call \mathcal{M} the *semantic algebra* of the interpretation and $[\cdot]$ the *semantic function*.

Given a probabilistic program P and an interpretation $\mathcal{I} = \langle \mathcal{M}, [\cdot] \rangle$, we define $\mathcal{I}[P]$ to be the interpretation of the probabilistic program. $\mathcal{I}[P]$ is then defined as the *least prefixed point* (i.e., the least ρ such that $f(\rho) \leq \rho$ for a function f) of the function F_P^{\sharp} , which is defined as

$$\lambda \mathcal{S}^{\sharp}. \lambda v. \begin{cases} \widehat{Ctrl}(e)^{\sharp}(\mathcal{S}^{\sharp}(u_1), \dots, \mathcal{S}^{\sharp}(u_k)) & e = \langle v, \{u_1, \dots, u_k\} \rangle \in E \\ \perp & \text{otherwise} \end{cases}$$

where

$\widehat{seq[act]}^{\sharp}(a_1) \stackrel{\text{def}}{=} [\text{act}] \otimes a_1$	$\widehat{cond[\varphi]}^{\sharp}(a_1, a_2) \stackrel{\text{def}}{=} a_1 \varphi \diamond a_2$
$\widehat{call[i]}^{\sharp}(a_1) \stackrel{\text{def}}{=} \mathcal{S}^{\sharp}(v_i^{\text{entry}}) \otimes a_1$	$\widehat{prob[p]}^{\sharp}(a_1, a_2) \stackrel{\text{def}}{=} a_1 p \oplus a_2$
	$\widehat{ndet}^{\sharp}(a_1, a_2) \stackrel{\text{def}}{=} a_1 \sqcup a_2$

We generalize Kleene's fixed-point theorem to prove the existence of the least prefixed point of F_P^{\sharp} .

Theorem 4.6 (Generalized Kleene's fixed point theorem). *Suppose $\langle X, \leq \rangle$ is an ω -cpo with a least element \perp , and let $f : X \rightarrow X$ be a pre- ω -continuous function. Then f has a*

least prefixed point, which is the supremum of the ascending Kleene chain of f , denoted by $\text{lpp}_{\perp}^{\leq} f$.

We use the least prefixed point of $F_P^{\#}$ to define the interpretation of a probabilistic program P as $\mathcal{I}[P] = \text{lpp}_{\lambda v. \perp}^{\leq} F_P^{\#}$. The interpretation of a control-flow node v is then defined as $\mathcal{I}[v] = \mathcal{I}[P](v)$.

4.2 Abstractions of Probabilistic Programs

Given two PMAs C and \mathcal{A} , a *probabilistic abstraction* is defined as follows:

Definition 4.7 (Probabilistic abstractions). A *probabilistic over-abstraction* (or *under-abstraction*, resp.) from a PMA C to a PMA \mathcal{A} is a concretization mapping, $\gamma : \mathcal{A} \rightarrow C$, such that

- γ is monotone, i.e., for all $Q_1, Q_2 \in \mathcal{A}$, $Q_1 \sqsubseteq_{\mathcal{A}} Q_2$ implies that $\gamma(Q_1) \sqsubseteq_C \gamma(Q_2)$,
- $\perp_C \sqsubseteq_C \gamma(\perp_{\mathcal{A}})$ (or $\gamma(\perp_{\mathcal{A}}) \sqsubseteq_C \perp_C$, resp.),
- $\perp_C \sqsubseteq_C \gamma(\perp_{\mathcal{A}})$ (or $\gamma(\perp_{\mathcal{A}}) \sqsubseteq_C \perp_C$, resp.),
- for all $Q_1, Q_2 \in \mathcal{A}$, $\gamma(Q_1) \otimes_C \gamma(Q_2) \sqsubseteq_C \gamma(Q_1 \otimes_{\mathcal{A}} Q_2)$ (or $\gamma(Q_1 \otimes_{\mathcal{A}} Q_2) \sqsubseteq_C \gamma(Q_1) \otimes_C \gamma(Q_2)$, resp.),
- for all $Q_1, Q_2 \in \mathcal{A}$, $\gamma(Q_1) \diamond_C \gamma(Q_2) \sqsubseteq_C \gamma(Q_1 \diamond_{\mathcal{A}} Q_2)$ (or $\gamma(Q_1 \diamond_{\mathcal{A}} Q_2) \sqsubseteq_C \gamma(Q_1) \diamond_C \gamma(Q_2)$, resp.),
- for all $Q_1, Q_2 \in \mathcal{A}$, $\gamma(Q_1) \oplus_C \gamma(Q_2) \sqsubseteq_C \gamma(Q_1 \oplus_{\mathcal{A}} Q_2)$ (or $\gamma(Q_1 \oplus_{\mathcal{A}} Q_2) \sqsubseteq_C \gamma(Q_1) \oplus_C \gamma(Q_2)$, resp.), and
- for all $Q_1, Q_2 \in \mathcal{A}$, $\gamma(Q_1) \sqcup_C \gamma(Q_2) \sqsubseteq_C \gamma(Q_1 \sqcup_{\mathcal{A}} Q_2)$ (or $\gamma(Q_1 \sqcup_{\mathcal{A}} Q_2) \sqsubseteq_C \gamma(Q_1) \sqcup_C \gamma(Q_2)$, resp.).

A probabilistic abstraction leads to a sound analyses:

Theorem 4.8. Let \mathcal{C} and \mathcal{A} be interpretations over PMAs C and \mathcal{A} ; let γ be a probabilistic over-abstraction (or under-abstraction, resp.) from C to \mathcal{A} ; and let P be an arbitrary probabilistic program. If for all basic actions act , $\llbracket \text{act} \rrbracket^{\mathcal{C}} \sqsubseteq_C \gamma(\llbracket \text{act} \rrbracket^{\mathcal{A}})$ (or $\gamma(\llbracket \text{act} \rrbracket^{\mathcal{A}}) \sqsubseteq_C \llbracket \text{act} \rrbracket^{\mathcal{C}}$, resp.), then we have $\mathcal{C}[P] \sqsubseteq_C \gamma(\mathcal{A}[P])$ (or $\gamma(\mathcal{A}[P]) \sqsubseteq_C \mathcal{C}[P]$, resp.).

4.3 Interprocedural Analysis Algorithm

We are given a probabilistic program P and an interpretation $\mathcal{A} = \langle \mathcal{A}, \llbracket \cdot \rrbracket^{\mathcal{A}} \rangle$, where $\mathcal{A} = \langle M_{\mathcal{A}}, \sqsubseteq_{\mathcal{A}}, \otimes_{\mathcal{A}}, \diamond_{\mathcal{A}}, \oplus_{\mathcal{A}}, \sqcup_{\mathcal{A}}, \perp_{\mathcal{A}}, \perp_{\mathcal{A}} \rangle$ is a PMA and $\llbracket \cdot \rrbracket^{\mathcal{A}}$ is a semantic function. The goal is to compute (an overapproximation of) $\mathcal{A}[P] = \text{lpp}_{\lambda v. \perp}^{\leq} F_P^{\#}$. An equivalent way to define $\mathcal{A}[P]$ is to specify it as the least solution to a system of inequalities on $\{\mathcal{A}[v] \mid v \in V\}$ (where $e \in E$ in each case):

	e	$\text{Ctrl}(e)$
$\mathcal{A}[v] \sqsupseteq_{\mathcal{A}} \llbracket \text{act} \rrbracket^{\mathcal{A}} \otimes_{\mathcal{A}} \mathcal{A}[u_1]$	$\langle v, \{u_1\} \rangle$	$\text{seq}[\text{act}]$
$\mathcal{A}[v] \sqsupseteq_{\mathcal{A}} \mathcal{A}[u_1] \diamond_{\mathcal{A}} \mathcal{A}[u_2]$	$\langle v, \{u_1, u_2\} \rangle$	$\text{cond}[\varphi]$
$\mathcal{A}[v] \sqsupseteq_{\mathcal{A}} \mathcal{A}[u_1] \oplus_{\mathcal{A}} \mathcal{A}[u_2]$	$\langle v, \{u_1, u_2\} \rangle$	$\text{prob}[p]$
$\mathcal{A}[v] \sqsupseteq_{\mathcal{A}} \mathcal{A}[u_1] \sqcup_{\mathcal{A}} \mathcal{A}[u_2]$	$\langle v, \{u_1, u_2\} \rangle$	ndet
$\mathcal{A}[v] \sqsupseteq_{\mathcal{A}} \mathcal{A}[v_i^{\text{entry}}] \otimes_{\mathcal{A}} \mathcal{A}[u_1]$	$\langle v, \{u_1\} \rangle$	$\text{call}[i]$
$\mathcal{A}[v] \sqsupseteq_{\mathcal{A}} \perp_{\mathcal{A}}$	if $v = v_i^{\text{exit}}$	

Note that in line 5 a call is treated as a hyper-edge with the action $\lambda(\text{entry}, \text{succ}).\text{entry} \otimes_{\mathcal{A}} \text{succ}$. There is no explicit

return edge to match a call (as in many multi-procedure program representations, e.g., [77]); instead, each exit node is initialized with the constant $\perp_{\mathcal{A}}$ (line 6).

We mainly use known techniques from previous work on interprocedural dataflow analysis, with some adaptations to our setting, which uses hyper-graphs instead of ordinary graphs (i.e., CFGs).⁴ The analysis direction is backward, and the algorithm is similar to methods for computing summary edges in demand interprocedural-dataflow-analysis algorithms ([44, Fig. 4], [78, Fig. 10]). The algorithm uses a standard chaotic-iteration strategy (except that propagation is performed along hyper-edges instead of edges); it uses a fair iteration strategy for selecting the next edge to consider.

4.4 Widening

Widening is a general technique in static analysis to ensure and speed up convergence [20, 22]. To choose the nodes at which widening is to be applied, we treat the hyper-graph as a graph—i.e., each hyper-edge (including calls) contributes one or two ordinary edges. More precisely, we construct a *dependence graph* $G(H) = \langle N, A \rangle$ from hyper-graph $H = \{\langle V_i, E_i, v_i^{\text{entry}}, v_i^{\text{exit}} \rangle\}_{1 \leq i \leq n}$ by defining $N \stackrel{\text{def}}{=} \bigcup_{1 \leq i \leq n} V_i$, and

$$A \stackrel{\text{def}}{=} \{ \langle u, v \rangle \mid \exists e \in E. (v = \text{src}(e) \wedge u \in \text{Dst}(e)) \} \cup \{ \langle v_i^{\text{entry}}, v \rangle \mid \exists e \in E. (v = \text{src}(e) \wedge \text{Ctrl}(e) = \text{call}[i]) \}. \quad (2)$$

We then compute a set W of widening points for $G(H)$ via the algorithm of Bourdoncle [9, Fig. 4]. Because of the second set-former in Eqn. (2), W contains widening points that cut each cycle caused by recursion.

While traditional programs exhibit only one sort of choice operator, probabilistic programs can have three different kinds of choice operators, and hence loops can exhibit three different kinds of behavior. We found that if we used the same widening operator for all widening nodes, there could be a substantial loss in precision. Thus, we equip the framework with three separate widening operators: ∇_c , ∇_p , and ∇_n . Let $v \in W$ be the source of edge $e \in E$. Then the inequalities become

	e	$\text{Ctrl}(e)$
$\mathcal{A}[v] \sqsupseteq_{\mathcal{A}} \mathcal{A}[v] \nabla_n (\llbracket \text{act} \rrbracket^{\mathcal{A}} \otimes_{\mathcal{A}} \mathcal{A}[u_1])$	$\langle v, \{u_1\} \rangle$	$\text{seq}[\text{act}]$
$\mathcal{A}[v] \sqsupseteq_{\mathcal{A}} \mathcal{A}[v] \nabla_c (\mathcal{A}[u_1] \diamond_{\mathcal{A}} \mathcal{A}[u_2])$	$\langle v, \{u_1, u_2\} \rangle$	$\text{cond}[\varphi]$
$\mathcal{A}[v] \sqsupseteq_{\mathcal{A}} \mathcal{A}[v] \nabla_p (\mathcal{A}[u_1] \oplus_{\mathcal{A}} \mathcal{A}[u_2])$	$\langle v, \{u_1, u_2\} \rangle$	$\text{prob}[p]$
$\mathcal{A}[v] \sqsupseteq_{\mathcal{A}} \mathcal{A}[v] \nabla_n (\mathcal{A}[u_1] \sqcup_{\mathcal{A}} \mathcal{A}[u_2])$	$\langle v, \{u_1, u_2\} \rangle$	ndet
$\mathcal{A}[v] \sqsupseteq_{\mathcal{A}} \mathcal{A}[v] \nabla_n (\mathcal{A}[v_i^{\text{entry}}] \otimes_{\mathcal{A}} \mathcal{A}[u_1])$	$\langle v, \{u_1\} \rangle$	$\text{call}[i]$

Observation 4.9. Recall from Defn. 3.2 that in a probabilistic program each non-exit node has exactly one outgoing

⁴As mentioned in §2.3, standard formulations of interprocedural dataflow analysis [52, 57, 71, 80] can be viewed as hyper-graph analyses. In that setting, one deals with hyper-graphs with constituent control-flow graphs. With PMAF, because each procedure is represented as a hyper-graph, one has hyper-graphs of constituent hyper-graphs. Fortunately, each procedure's hyper-graph is a *single-entry/single-exit* hyper-graph, so the basic ideas and algorithms from standard interprocedural dataflow analysis carry over to PMAF.

hyper-edge. In each right-hand side above, the second argument to the widening operator re-evaluates the action of the (one outgoing) hyper-edge. Consequently, during an analysis, we have the invariant that whenever a widening operation $a \nabla b$ is performed, the property $a \sqsubseteq_{\mathcal{A}} b$ holds.

The safety properties for the three widening operators are adaptations of the standard stabilization condition: For every pair of ascending chains $\{a_k\}_{k \in \mathbb{N}}$ and $\{b_k\}_{k \in \mathbb{N}}$,

- the chain $\{c_k\}_{k \in \mathbb{N}}$ defined by $c_0 = a_0 \diamond_{\varphi} b_0$ and $c_{k+1} = c_k \nabla_c (a_{k+1} \diamond_{\varphi} b_{k+1})$ is eventually stable;
- the chain $\{c_k\}_{k \in \mathbb{N}}$ defined by $c_0 = a_0 \oplus_p b_0$ and $c_{k+1} = c_k \nabla_p (a_{k+1} \oplus_p b_{k+1})$ is eventually stable; and
- the chain $\{c_k\}_{k \in \mathbb{N}}$ defined by $c_0 = a_0 \cup_{\mathcal{A}} b_0$ and $c_{k+1} = c_k \nabla_n (a_{k+1} \cup_{\mathcal{A}} b_{k+1})$ is eventually stable.

5 Instantiations

In this Section, we instantiate the framework to derive three important analyses: Bayesian inference (BI) (§5.1), computing rewards in Markov decision processes (§5.2), and linear expectation-invariant analysis (LEIA) (§5.3).

5.1 Bayesian Inference

Claret et al. [18] proposed a technique to perform Bayesian inference on Boolean programs using dataflow analysis. They use a forward analysis to compute the posterior distribution of a single-procedure, well-structured, probabilistic program. Their analysis is similar to an intraprocedural dataflow analysis: they use discrete joint-probability distributions as dataflow facts, merge these facts at join points, and compute fixpoints in the presence of loops. Let Var be the set of program variables; the set of program states is $\Omega = \text{Var} \rightarrow \mathbb{B}$. Note that Ω is isomorphic to $\mathbb{B}^{|\text{Var}|}$, and consequently, a distribution can be represented by a vector of length $2^{|\text{Var}|}$ of reals in $\mathbb{R}_{[0,1]}$. (Their implementation uses Algebraic Decision Diagrams [2] to represent distributions compactly.)

The algorithm by Claret et al. is defined inductively on the structure of programs [18, Alg. 2]—for example, the output distribution of $x \sim \text{Bernoulli}(r)$ from an input distribution μ , denoted by $\text{POST}(\mu, x \sim \text{Bernoulli}(r))$, is computed as $\lambda \sigma'. (r \cdot \sum_{\{\sigma | \sigma' = \sigma[x \leftarrow \text{true}]\}} \mu(\sigma) + (1-r) \cdot \sum_{\{\sigma | \sigma' = \sigma[x \leftarrow \text{false}]\}} \mu(\sigma))$.

We have used PMAF to extend their work in two dimensions, creating (i) an *interprocedural* version of Bayesian inference with (ii) *nondeterminism*. Because of nondeterminism, for a given input state the posterior distribution is not unique; consequently, our goal is to compute procedure summaries that gives *lower bounds* on posterior distributions.

To reformulate the domain in the two-vocabulary setting needed for computing procedure summaries, we introduce Var' , primed versions of the variables in Var . Var and Var' denote the variables in the pre-state and post-state of a state transformer. A distribution transformer (and therefore a procedure summary) is a matrix of size $2^{|\text{Var}'|} \times 2^{|\text{Var}|}$ of reals in $\mathbb{R}_{[0,1]}$. We define a PMA $\mathcal{B} = \langle M_{\mathcal{B}}, \sqsubseteq_{\mathcal{B}}$

, $\otimes_{\mathcal{B}}, \diamond_{\mathcal{B}}, \oplus_{\mathcal{B}}, \cup_{\mathcal{B}}, \perp_{\mathcal{B}}, \underline{1}_{\mathcal{B}} \rangle$ as follows:

$M_{\mathcal{B}} \stackrel{\text{def}}{=} 2^{ \text{Var} } \times 2^{ \text{Var}' } \rightarrow \mathbb{R}_{[0,1]}$	
$a \sqsubseteq_{\mathcal{B}} b \stackrel{\text{def}}{=} a \leq b$	$a \cup_{\mathcal{B}} b \stackrel{\text{def}}{=} \min(a, b)$
$a \otimes_{\mathcal{B}} b \stackrel{\text{def}}{=} a \times b$	$\perp_{\mathcal{B}} \stackrel{\text{def}}{=} \lambda(s, t). 0$
$a \oplus_{\mathcal{B}} b \stackrel{\text{def}}{=} p \cdot a + (1-p) \cdot b$	$\underline{1}_{\mathcal{B}} \stackrel{\text{def}}{=} \lambda(s, t). [s = t]$
$a \diamond_{\mathcal{B}} b \stackrel{\text{def}}{=} \lambda(s, t). \text{if } \widehat{\varphi}(s) \text{ then } a(s, t) \text{ else } b(s, t)$	

The use of pointwise min in the definition of $a \cup_{\mathcal{B}} b$ causes the analysis to compute procedure summaries that provide lower bounds on the posterior distributions.

Theorem 5.1. \mathcal{B} is a PMA.

Let $\mathcal{B} = \langle \mathcal{B}, [\cdot]_{\mathcal{B}} \rangle$ be the interpretation for Bayesian inference. We define the semantic function as $\llbracket x := \mathcal{E} \rrbracket^{\mathcal{C}} = \lambda(s, A). [s[x \leftarrow \mathcal{E}(s)] \in A]$ and $\llbracket x := \mathcal{E} \rrbracket^{\mathcal{B}} = \lambda(s, t). [s[x \leftarrow \mathcal{E}(s)] = t]$, as well as $\llbracket x \sim \text{Bernoulli}(p) \rrbracket^{\mathcal{C}} = \lambda(s, A). p \cdot [s[x \leftarrow \text{true}] \in A] + (1-p) \cdot [s[x \leftarrow \text{false}] \in A]$ and $\llbracket x \sim \text{Bernoulli}(p) \rrbracket^{\mathcal{B}} = \lambda(s, t). p \cdot [s[x \leftarrow \text{true}] = t] + (1-p) \cdot [s[x \leftarrow \text{false}] = t]$.

We define the concretization mapping $\gamma_{\mathcal{B}} : M_{\mathcal{B}} \rightarrow \mathbb{P}\Omega$ as $\gamma_{\mathcal{B}}(a) = \langle\langle \{\kappa \mid \forall s, s'. \kappa(s, \{s'\}) \geq a(s, s') \} \rangle\rangle$ where $\langle\langle C \rangle\rangle$ denotes the smallest element in $\mathbb{P}\Omega$ such that contains C .

Theorem 5.2. $\gamma_{\mathcal{B}}$ is a prob. under-abstraction from \mathcal{C} to \mathcal{B} .

We do not define widening operators for BI, because $\gamma_{\mathcal{B}}$ is an under-abstraction and our algorithm starts from the bottom element in the abstract domain, the intermediate result at any iteration is a sound answer.

5.2 Markov Decision Process with Rewards

Analyses of finite-state Markov decision processes were originally developed in the fields of operational research and finance mathematics [75]. Originally, Markov decision processes were defined as finite-state machines with actions that exhibit probabilistic transitions. In this paper, we use a slightly different formalization, using hyper-graphs.

Definition 5.3 (Markov decision process). A *Markov decision process* (MDP) is a hyper-graph $H = \langle V, E, v^{\text{entry}}, v^{\text{exit}} \rangle$, where every node except v^{exit} has exactly one outgoing hyper-edge; each hyper-edge with just a single destination has an associated *reward*, $\text{seq}[\text{reward}(r)]$, where r is a positive real number; and each hyper-edge with two destinations has either $\text{prob}[p]$, where $0 \leq p \leq 1$, or ndet . Note that MDPs are a specialization of single-procedure probabilistic programs without conditional-choice.

We can also treat the hyper-graph as a graph: each hyper-edge contributes one or two graph edges. A path through the graph has a *reward*, which is the sum of the rewards that label the edges of the path. (Edges from hyper-edges with the actions $\text{prob}[p]$ or ndet are considered to have reward 0.) The analysis problem that we wish to solve is to determine,

for each node v , the greatest expected reward that one can gain by executing the program from v .

It is natural to extend MDPs with procedure calls and multiple procedures, to obtain *recursive Markov decision processes*. The set of program states is defined to be the set of nonnegative real numbers: $\Omega = [0, \infty]$. To address the maximum-expected-reward problem for a recursive Markov decision process, we define a PMA $\mathcal{R} = \langle M_{\mathcal{R}}, \sqsubseteq_{\mathcal{R}}, \otimes_{\mathcal{R}}, \varphi \diamond_{\mathcal{R}}, p \oplus_{\mathcal{R}}, \cup_{\mathcal{R}}, \perp_{\mathcal{R}}, \underline{\perp}_{\mathcal{R}} \rangle$ as follows:

$M_{\mathcal{R}} \stackrel{\text{def}}{=} [0, \infty]$	$\varphi \diamond_{\mathcal{R}} \stackrel{\text{def}}{=} \max$	$\perp_{\mathcal{R}} \stackrel{\text{def}}{=} 0$
$\sqsubseteq_{\mathcal{R}} \stackrel{\text{def}}{=} \leq$	$a p \oplus_{\mathcal{R}} b \stackrel{\text{def}}{=} p \cdot a + (1 - p) \cdot b$	$\underline{\perp}_{\mathcal{R}} \stackrel{\text{def}}{=} 0$
$\otimes_{\mathcal{R}} \stackrel{\text{def}}{=} +$	$\cup_{\mathcal{R}} \stackrel{\text{def}}{=} \max$	

Theorem 5.4. \mathcal{R} is a PMA.

Let $\mathcal{R} = \langle \mathcal{R}, [\cdot]^{\mathcal{R}} \rangle$ be the interpretation for a Markov decision process with rewards. We define the semantic function as $[\text{reward}(r)]^{\mathcal{C}} = \lambda(s, A). [s + r \in A]$ and $[\text{reward}(r)]^{\mathcal{R}} = r$.

We define the concretization mapping $\gamma_{\mathcal{R}} : M_{\mathcal{R}} \rightarrow \mathbb{P}[0, \infty]$ as follows: $\gamma_{\mathcal{R}}(a) = \langle \{ \kappa \mid \forall s. \int y \cdot \kappa(s, dy) \leq s + a \} \rangle$.

Theorem 5.5. $\gamma_{\mathcal{R}}$ is a prob. over-approximation from \mathcal{C} to \mathcal{R} .

We use a trivial widening in this analysis: if after some fixed number of iterations the analysis does not converge, it returns ∞ as the result.

5.3 Linear Expectation-Invariant Analysis

Several examples of expectation invariants obtained via linear expectation-invariant analysis (LEIA) were given in §2.2. This section gives details of the abstract domain for LEIA.

We make use of an existing abstract domain, namely, the domain of *convex polyhedra* [24]. Elements of the polyhedral domain are defined by linear-inequality and linear-equality constraints among program variables. For LEIA, we use two-vocabulary polyhedra over nonnegative program variables. Let $x = (x_1, \dots, x_n)^T$ be a column vector of nonnegative program variables and $x' = (x'_1, \dots, x'_n)^T$ be a column vector of the “primed” versions of corresponding program variables. A polyhedron $P \subseteq \mathbb{R}_{\geq 0}^{2n}$ captures linear-inequality constraints among x and x' , which can be interpreted as a relation between pre-state and post-state variable valuations.

A polyhedron $P = \{ (x'^T, x^T)^T \in \mathbb{R}_{\geq 0}^{2n} \mid A'x' + Ax \leq b \wedge D'x' + Dx = e \}$, can be encoded as the intersection of a finite number of closed half spaces and a finite number of subspaces, where A', A, D', D are matrices and b, e are vectors. The associated *constraint set* is defined as $C_P = \{ A'x' + Ax \leq b, D'x' + Dx = e \}$. Let \mathcal{P} be the set of polyhedra; \mathcal{P} is equipped with meet, join, renaming, forgetting, and comparison operations.

LEIA uses *expectation polyhedra*. They are actually the same as polyhedra, except that the two vocabularies are $x = (x_1, \dots, x_n)^T$ and $\mathbb{E}[x'] = (\mathbb{E}[x'_1], \dots, \mathbb{E}[x'_n])^T$. An expectation polyhedron represents a constraint set of the form

$$\{ A' \mathbb{E}[x'] + Ax \leq b, D' \mathbb{E}[x'] + Dx = e \}. \quad (3)$$

Because of the linearity of the expectation operator \mathbb{E} , an equivalent way to express Eqn. (3) is as follows:

$$\{ \mathbb{E}[A'x'] + Ax \leq b, \mathbb{E}[D'x'] + Dx = e \}.$$

Let \mathcal{EP} be the set of expectation polyhedra. \mathcal{EP} is equipped with the same set of operations as \mathcal{P} .

We define the state space to be $\Omega = \mathbb{R}_{\geq 0}^n$. We then define a PMA \mathcal{I} with a universe $M_{\mathcal{I}} \stackrel{\text{def}}{=} \mathcal{P} \times \mathcal{EP}$. An element $(P, EP) \in \mathcal{I}$ consists of (i) a set of standard constraints $P \in \mathcal{P}$, and (ii) a set of expectation constraints $EP \in \mathcal{EP}$, such that $\mathbf{0} \sqcup P[\mathbb{E}[x']/x'] \sqsupseteq EP$ holds, where $\mathbf{0} \stackrel{\text{def}}{=} \bigwedge_{i=1}^n (\mathbb{E}[x'_i] = 0)$. The latter property means that, if necessary, we can always “rebuild” a pessimistic \mathcal{EP} component from the \mathcal{P} component as $\mathbf{0} \sqcup P[\mathbb{E}[x']/x']$.⁵

We define the concretization mapping $\gamma_{\mathcal{I}}$ as follows:

$$\gamma_{\mathcal{I}}(P, EP) = \langle \left\{ \kappa \mid \forall s. \kappa \left(s, \left\{ s' \mid \begin{bmatrix} s' \\ s \end{bmatrix} \models \neg P \right\} \right) = 0 \wedge \left[\int_s s' \kappa(s, ds') \right] \models EP \right\} \rangle.$$

Comparison. The comparison operation on ordinary polyhedra can be defined as standard set inclusion. For expectation polyhedra, taking into account subprobability distributions, we define $EP_1 \sqsubseteq EP_2$ to be $\mathbf{0} \sqcup EP_1 \subseteq \mathbf{0} \sqcup EP_2$, so that any element inside or below EP_1 should also be inside or below EP_2 . Consequently, we define $(P_1, EP_1) \sqsubseteq_{\mathcal{I}} (P_2, EP_2) \stackrel{\text{def}}{=} P_1 \subseteq P_2 \wedge \mathbf{0} \sqcup EP_1 \subseteq \mathbf{0} \sqcup EP_2$.

Composition. For ordinary polyhedra, the composition of P_1 and P_2 can be defined as

$$(\exists x''. C_{P_1}[x''/x'] \wedge C_{P_2}[x''/x]) \Rightarrow C_{P_1 \otimes P_2},$$

where we introduce an intermediate vocabulary $x'' = (x''_1, \dots, x''_n)^T$, and use it to connect P_1 and P_2 . Consequently, we define $P_1 \otimes P_2$ to be $\exists x''. C_{P_1}[x''/x'] \wedge C_{P_2}[x''/x]$. Operationally, composition involves first introducing a new vocabulary; renaming the variables properly; performing a meet, and finally forgetting the intermediate vocabulary.

Somewhat surprisingly, because of the *tower property* in probability theory, exactly the same steps can be used to compose expectation polyhedra. Informally, the tower property means that $\mathbb{E}[X] = \mathbb{E}[\mathbb{E}[X \mid Y]]$, where X and Y are two random variables, and $\mathbb{E}[X \mid Y]$ is a conditional expectation. For instance, suppose that EP_1 and EP_2 are defined by the constraint sets $\{ \mathbb{E}(x') = x + 2 \}$ and $\{ \mathbb{E}(x') = 7x \}$, respectively. Following the renaming recipe above, we have $\mathbb{E}(x'') = x + 2$ and $\mathbb{E}(x' \mid x'') = 7x''$. By the tower property, we have $\mathbb{E}(x') = \mathbb{E}(\mathbb{E}(x' \mid x'')) = \mathbb{E}(7x'') = 7\mathbb{E}(x'') = 7x + 14$. Operationally, the tower property allows us to compose linear expectation invariants, and eliminate the intermediate

⁵The intuition is that P represents a convex overapproximation to some desired set of points; the expected value has to lie somewhere inside $\tilde{\mathbf{0}} \sqcup P$, where “ $\tilde{\mathbf{0}} \dots$ ” is needed to account for subprobability distributions. For instance, for a nonnegative interval $[lo, hi]$, we must have *expected* $\in ([0, 0] \sqcup [lo, hi])$; i.e., $0 \leq \text{expected} \leq hi$.

vocabulary x'' . Consequently, we define

$$(P_1, EP_1) \otimes_I (P_2, EP_2) \stackrel{\text{def}}{=} (P_1 \otimes P_2, EP_1 \otimes EP_2).$$

Conditional-choice. For the ordinary-polyhedron component, a conditional-choice $\varphi \diamond$ is performed by first meeting each operand with the logical constraint φ , and then joining the results. However, for the expectation-polyhedron component, conditioning can split the probability space in almost arbitrary ways. Consequently, the constraints on post-state expectations as a function of pre-state valuations are not necessarily true after conditioning. Thus, we define

$$(P_1, EP_1) \varphi \diamond_I (P_2, EP_2) \stackrel{\text{def}}{=} \mathbf{let} P = (\{\varphi\} \sqcap P_1) \sqcup (\{\neg\varphi\} \sqcap P_2) \\ \mathbf{in} (P, (EP_1 \sqcup EP_2) \sqcap (\mathbf{0} \sqcup P[\mathbb{E}[x']]/x'))).$$

The \sqcap in the second component is performed to maintain the invariant that $\mathbf{0} \sqcup P[\mathbb{E}[x']]/x' \sqsupseteq$ the second component.

Probabilistic-choice. For the ordinary-polyhedron component, we merely join the components of the two operands. For the expectation-polyhedron component, we introduce two more vocabularies and have

$$(\exists x'', x'''. C_{EP_1}[x''/\mathbb{E}[x']] \wedge C_{EP_2}[x'''/\mathbb{E}[x']] \wedge \\ \bigwedge_{i=1}^n \mathbb{E}[x'_i] = p \cdot x''_i + (1-p) \cdot x'''_i) \Rightarrow C_{EP_1 \oplus EP_2}.$$

Consequently, we define $EP_1 \oplus EP_2$ to be

$$\exists x'', x''' \cdot \left(C_{EP_1}[x''/\mathbb{E}[x']] \wedge C_{EP_2}[x'''/\mathbb{E}[x']] \right. \\ \left. \wedge \bigwedge_{i=1}^n \mathbb{E}[x'_i] = p \cdot x''_i + (1-p) \cdot x'''_i \right),$$

and $(P_1, EP_1) \oplus_I (P_2, EP_2) \stackrel{\text{def}}{=} (P_1 \sqcup P_2, EP_1 \oplus EP_2)$.

Nondeterministic-choice. The nondeterministic-choice operations on both ordinary polyhedra and expectation polyhedra can be defined as join. Hence, we define $(P_1, EP_1) \sqcup_I (P_2, EP_2) \stackrel{\text{def}}{=} (P_1 \sqcup P_2, EP_1 \sqcup EP_2)$.

Bottom and Unit Element. We define $\perp_I \stackrel{\text{def}}{=} (\text{false}, \mathbf{0})$, and $\underline{1}_I \stackrel{\text{def}}{=} (\{x'_i = x_i \mid 1 \leq i \leq n\}, \{\mathbb{E}[x'_i] = x_i \mid 1 \leq i \leq n\})$.

Semantic Function. Some examples of the semantic mapping $\llbracket \cdot \rrbracket^{\mathcal{S}}$ are as follows, where $\min(\mathcal{D})$ and $\max(\mathcal{D})$ represents the interval of the support of a distribution \mathcal{D} , while $\text{mean}(\mathcal{D})$ stands for its average.

$$\llbracket x_i := \mathcal{E} \rrbracket^{\mathcal{S}} \stackrel{\text{def}}{=} \left(\{x'_i = \mathcal{E}(x)\} \cup \{x'_j = x_j \mid j \neq i\}, \right. \\ \left. \{\mathbb{E}[x'_i] = \mathcal{E}(x)\} \cup \{\mathbb{E}[x'_j] = x_j \mid j \neq i\} \right) \\ \llbracket x_i \sim \mathcal{D} \rrbracket^{\mathcal{S}} \stackrel{\text{def}}{=} \left(\{\min(\mathcal{D}) \leq x'_i \leq \max(\mathcal{D})\} \cup \{x'_j = x_j \mid j \neq i\}, \right. \\ \left. \{\mathbb{E}[x'_i] = \text{mean}(\mathcal{D})\} \cup \{\mathbb{E}[x'_j] = x_j \mid j \neq i\} \right) \\ \llbracket \text{skip} \rrbracket^{\mathcal{S}} \stackrel{\text{def}}{=} \underline{1}_I$$

Note we assume all expressions in the program are linear. For nonlinear arithmetic programs, one can adopt some linearization techniques [32, 62].

Theorem 5.6. γ_I is a prob. over-abstraction from C to I .

Widening. Let ∇ be the standard widening operator on ordinary polyhedra [42]. Recall from Obs. 4.9 that whenever a widening operation $a \nabla b$ is performed, the property $a \sqsubseteq_{\mathcal{A}} b$ holds. There is a subtle issue with expectation invariants when dealing with conditional or nondeterministic loops.

Observation 5.7. In a conventional program, if you have a loop “while B do S od,” and I is a loop-invariant, then $I \wedge \neg B$ (which implies I) holds on exiting the loop. In contrast, for a conditional or nondeterministic loop in a probabilistic program, a loop-invariant that holds at the beginning and end of the loop body does not necessarily hold on exiting the loop.

Example 5.8. Consider the following program:

```
while  $\neg(x = y)$  do
  if  $\text{prob}(\frac{1}{2})$  then  $x := x + 1$  else  $y := y + 1$  fi od
```

For the loop body, we can derive an expectation invariant $\mathbb{E}[x' - y'] = x - y$; however, for the entire loop this property does not hold: at the end of the loop $x = y$ must hold, and hence $\mathbb{E}[x' - y']$ should be equal to 0.

Because of this issue, we use a pessimistic widening operator for conditional-choice and nondeterministic-choice: the widening operator forgets the expectation invariants and rebuilds them from standard invariants.

$$(P_1, EP_1) \nabla_c (P_2, EP_2) \stackrel{\text{def}}{=} (P_1 \nabla P_2, \mathbf{0} \sqcup P_2[\mathbb{E}[x']]/x') \\ (P_1, EP_1) \nabla_n (P_2, EP_2) \stackrel{\text{def}}{=} (P_1 \nabla P_2, \mathbf{0} \sqcup P_2[\mathbb{E}[x']]/x')$$

We do not have a good method for $(P_1, EP_1) \nabla_p (P_2, EP_2)$. We found that the following approach loses precision:

$$\mathbf{let} P = (P_1 \nabla P_2) \mathbf{in} (P, (EP_1 \nabla EP_2) \sqcap (\mathbf{0} \sqcup P[\mathbb{E}[x']]/x'))$$

In our experiments, we use $(P_1, EP_1) \nabla_p (P_2, EP_2) \stackrel{\text{def}}{=} (P_1 \nabla P_2, EP_2)$, which does no extrapolation in the \mathcal{EP} component.

6 Evaluation

In this Section, we first describe the implementation of PMAF, and the three instantiations introduced in §5. Then, we evaluate the effectiveness and performance of the three analyses.

6.1 Implementation

PMAF is implemented in OCaml; the core framework consists of about 400 lines of code. The framework is implemented as a functor parametrized by a module representing a PMA, with some extra functions, such as widening and printing. This organization allows any analysis that can be formulated in PMAF to be implemented as a plugin. Also, the core framework relies on control-flow hyper-graphs, and provides users the flexibility to employ it with any front end. We use OCamlGraph [19] as the implementation of fixed-point computation and Bourdoncle’s algorithm.

The plugin for Bayesian inference is about 400 lines of code, including a lexer and a parser for the imperative language that we use in the examples of this paper. We use Lacaml [70] to manipulate matrices. The plugins for the Markov decision problem with rewards and linear expectation-invariant analysis are about 200 lines and 500 lines, respectively. We use APRON [45] for polyhedron operations. Most of the code in the plugins is to implement the PMA structure of the analysis domain.

Program	Expectation invariants	#loc	rec?	#call	time
2d-walk	$\mathbb{E}[x'] = x, \mathbb{E}[y'] = y, \mathbb{E}[dist'] = dist, \mathbb{E}[count'] \leq count + 1, \mathbb{E}[count'] \geq count$	47	n	0	0.24
aggregate-rv	$\mathbb{E}[2x' - i'] = 2x - i, \mathbb{E}[x'] \leq x + \frac{1}{2}, \mathbb{E}[x'] \geq x$	11	n	0	0.06
biased-coin	$\mathbb{E}[x'] \leq x + \frac{1}{2}, \mathbb{E}[x'] \geq x - \frac{1}{2}$	25	n	0	0.06
binom-update ($p=\frac{1}{4}$)	$\mathbb{E}[4x' - n'] = 4x - n, \mathbb{E}[x'] \leq x + \frac{1}{4}, \mathbb{E}[x'] \geq x$	14	n	0	0.06
coupon5	$\mathbb{E}[count' - i'] = count - i$ (1st), $\mathbb{E}[4count' - 5i'] = 4count - 5i$ (2nd), $\mathbb{E}[3count' - 5i'] = 3count - 5i$ (3rd), $\mathbb{E}[2count' - 5i'] = 2count - 5i$ (4th), $\mathbb{E}[count' - 5i'] = count - 5i$ (5th)	58	n	0	0.07
dist	$\mathbb{E}[x'] = x, \mathbb{E}[y'] = y, \mathbb{E}[z'] = \frac{1}{2}x + \frac{1}{2}y$	5	n	0	0.05
eg	$\mathbb{E}[x' + y'] = x + y + 3, \mathbb{E}[z'] = \frac{1}{4}z + \frac{3}{4}, \mathbb{E}[x'] \leq x + 3, \mathbb{E}[x'] \geq x$	8	n	0	0.89
eg-tail	$\mathbb{E}[z'] \geq \frac{1}{4}z, \mathbb{E}[x'] \geq x, \mathbb{E}[y'] \geq y, \mathbb{E}[x' + y'] \geq x + y + \frac{3}{4}$	11	t	1	0.13
hare-turtle	$\mathbb{E}[2h' - 5t'] = 2h - 5t, \mathbb{E}[h'] \leq h + \frac{5}{2}, \mathbb{E}[h'] \geq h$	15	n	0	0.06
hawk-dove	$\mathbb{E}[p1b' - count'] = p1b - count, \mathbb{E}[p2b' - count'] = p2b - count, \mathbb{E}[p1b'] \leq p1b + 1, \mathbb{E}[p1b'] \geq p1b$	29	n	0	0.08
mot-ex	$\mathbb{E}[2x' - y'] = 2x - y, \mathbb{E}[4x' - 3count'] = 4x - 3count, \mathbb{E}[x'] \leq x + \frac{3}{4}, \mathbb{E}[x'] \geq x$	16	n	0	0.06
recursive	$\mathbb{E}[x'] = x + 9$	13	r	2	0.37
uniform-dist	$\mathbb{E}[n'] \leq 2n, \mathbb{E}[n'] \geq n, \mathbb{E}[g'] \leq 2g + \frac{1}{2}, \mathbb{E}[g'] \geq g$	14	n	0	0.06

Table 1. Linear expectation-invariant analysis

Because of the numerical reasoning required when analyzing probabilistic programs, we need to be concerned about finite numerical precision in our implementations of the instantiations (although they are sound on a theoretical machine operating on reals). In our implementation, we use the fact that ascending chains of floating numbers always converge in a finite number of steps. The user could use the technique proposed by Darulova et al. [26] to obtain a sound guarantee on numerical precision.

6.2 Experiments

Evaluation Platform. Our experiments were performed on a machine with an Intel Core i5 2.4 GHz processor and 8GB of RAM under Mac OS X 10.13.4.

Bayesian Inference and Markov Decision Problem with Rewards.

We tested our framework on Bayesian inference and Markov decision problem with rewards on handcrafted examples. The results of the evaluation of the two analyses are described in Tab. 2.

The tables contains the number of lines; whether the program is non-recursive, tail-recursive, or recursive; the number of procedure calls; and the time taken by the implementation (measured by running each program 5 times and computing the 20% trimmed mean).

Our framework computed the same answer (modulo floating-point round-off errors) as PReMo [84], a tool for probabilistic recursive models. We did not compare with

probabilistic abstract interpretation [25] because its semantic foundation is substantially different from that of our framework—as we mentioned in §1, the order for resolving probabilistic behavior and nondeterministic behavior is different.

The analysis time of Bayesian inference grows exponentially with respect to the number of program variables.⁶ The time cost comes from the explicit matrix representation of domain elements. One could use Algebraic Decision Diagrams [2] as a compact representation to improve the efficiency.

The analyzer for the Markov decision problem with rewards works quickly and obtains some interesting results. *quicksort7* is a model of a randomized quicksort algorithm on an array of size 7 (obtained from [84]), and our analysis results are consistent with the worst-case expected number of comparisons being $\Theta(n \log n)$.⁷ *binary10* is a model of randomized binary search algorithm on an array of size 10, and our analysis results are consistent with the worst-case expected number of comparisons being $\Theta(\log n)$.

Linear Expectation-Invariant Analysis. We performed a more thorough evaluation of linear expectation-invariant analysis. We collected several examples from the literature on probabilistic invariant generation [14, 49], and handcrafted some new examples to demonstrate particular capabilities of our domain, e.g., analysis of recursive programs. For the examples obtained from the loop-invariant-generation benchmark, we extracted the loop body as our test programs. Also, we performed a positive-negative decomposition to make sure all program variables are nonnegative. That is, we represented each variable x as $x^+ - x^-$ where $x^+, x^- \geq 0$, and

Program	#loc	rec?	#call	time
compare	17	n	0	2.22
dice	12	n	0	0.02
eg1	10	n	0	0.02
eg1-tail	16	t	2	0.02
eg2	10	n	0	0.02
eg2-tail	16	t	2	0.01
recursive	14	r	1	0.01
binary10	184	n	90	0.03
loop	10	n	0	0.03
quicksort7	109	n	42	0.03
recursive	13	t	1	0.03
student	43	t	8	0.03

Table 2. Top: Bayesian inference. Bottom: Markov decision problem with rewards. (Time is in seconds.)

⁶One should not assume that exponential growth makes the analysis useless; after all, predicate-abstraction domains [40] also grow exponentially: the universe of assignments to a set of Boolean variables grows exponentially in the number of variables. Finding useful coarser abstractions for Bayesian inference—by analogy with the techniques of Ball et al. [3] for predicate abstraction—might be an interesting direction for future work.

⁷The analysis computes *worst-case expected number* because the underlying semantics resolves nondeterminism first and probabilistic-choice second, and thus the analysis computes $\max_{\text{nondet. resolution}} \mathbb{E}[\text{\#comparisons under resolution}]$.

replaced every operation on variables with appropriate operations on the decomposed variables.

The results of the evaluation are shown in Tab. 1, which lists the expectation invariants obtained, and the time taken by the implementation. In general, the analysis runs quickly—all the examples are processed in less than one second. The analysis time mainly depends on the number of program variables and the size of the control-flow hyper-graph.

As shown in Tab. 1, our analysis can derive nontrivial expectation invariants, e.g., relations among different program variables such as $\mathbb{E}[x' + y'] = x + y + 3$, $\mathbb{E}[2x' - y'] = 2x - y$. In most cases, our results are at least as precise as those in [14, 49]. Exceptions are *biased-coin* and *uniform-dist*, collected from [49], where their invariant-generation algorithm uses a template-based approach and the form of expectations can be more complicated, e.g., $[P_1] \cdot \mathcal{E}_1 + [P_2] \cdot \mathcal{E}_2$ where P_1, P_2 are linear assertions and $\mathcal{E}_1, \mathcal{E}_2$ are linear expressions. Nevertheless, our analysis is fully automated and applicable to general programs, while [49] requires interactive proofs for nested loops, and [14] works only for single loops.

7 Related Work

Static Analysis for Standard Programs. Our framework is an extension of interprocedural dataflow analysis [52, 57, 71, 80] to probabilistic programs, but it does not support some language features that standard dataflow analysis has been used to address, e.g., calls through function pointers.

Compared to the Galois connections that are ordinarily used in abstract interpretation [21, 23], our definition of probabilistic abstractions is based on just a concretization function, so PMAF does not have the full power of standard abstract-interpretation machinery.

Static Analysis for Probabilistic Programs. Most closely related to our work is probabilistic abstract interpretation [25, 67–69], which is discussed in the introduction. There is a long line of research on manual reasoning techniques for probabilistic programs [33, 48, 56, 59, 72]. The main difference to this work is that we focus on the design and implementation of automatic techniques that rely on computing fixed points.

Other work focuses on specialized automatic analyses for specific properties. Claret et al. [18] proposed a dataflow analysis for Bayesian inference on Boolean programs that we reformulate in PMAF to lift it to the interprocedural level. There are different techniques for automatically proving termination, such as probabilistic pushdown automata [10, 11] and martingales and stochastic invariants [16, 17]. Martingales for automatic analysis of probabilistic programs have been pioneered by Chakarov et al. [13]. Compared with existing techniques for probabilistic invariant generation [4, 13, 14, 17], the expectation-invariant analysis proposed in §5.3 is designed as a two-vocabulary domain utilizing the well-studied polyhedral abstract domain.

Semantics for Probabilistic Programs. There is a long tradition of using probability kernels to define the semantics of probabilistic programs. Kernels were used by Kozen [56] to give a semantics for Probabilistic Propositional Dynamic Logic (PPDL), a probabilistic generalization of PDL. Kozen considers well-structured programs with sequencing and conditional-choice, but without non-deterministic choice. He does not consider reasoning methods that use abstract interpretation of his PPDL semantics. There are a list of domain-theoretic studies on probabilistic nondeterminism [27, 46, 47, 64, 65, 82], which develop powerdomain constructions over probability distributions, but do not consider powerdomains over kernels. Borgström et al. [8] have used kernels to define the operational semantics of a probabilistic lambda calculus. The main novelty of our denotational semantics in §3.3 is that it is defined for control-flow hyper-graphs, based on kernels.

Other Analyses Based on Hyper-Graphs. Hyper-graph-based analyses go back to the join-over-all-hyper-path-valuations of Knuth [53]. Other analyses based on hyper-graphs includes Möncke and Wilhelm’s [66] framework for finding join-over-all-hyper-path-valuations for *partially ordered* abstract domains. In the hyper-paths in this paper, we use binary hyper-edges to model calls, as well as conditional, probabilistic, and nondeterministic choice. For *acyclic* hyper-graphs, Eisner has considered semirings for computing expectations and variances of random variables [58]. He works with a discrete sample space: all hyper-paths in a given hyper-graph, and the value of a random variable for a given hyper-path is built up as the sum of the values contributed by each hyper-edge. In our work, we consider *cyclic* hyper-graphs, and the nature of the computation that a hyper-path represents is more complex than that considered by Eisner.

Acknowledgments

This article is based on research supported, in part, by a gift from Rajiv and Ritu Batra; by AFRL under DARPA MUSE award FA8750-14-2-0270, DARPA STAC award FA8750-15-C-0082, and DARPA award FA8750-16-2-0274; and by the UW-Madison Office of the Vice Chancellor for Research and Graduate Education with funding from the Wisconsin Alumni Research Foundation. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors, and do not necessarily reflect the views of the sponsoring agencies.

We thank the nurses from the Univ. of Pittsburgh Medical Center for their professional medical care, which was instrumental in finishing the submission before the PLDI deadline.

References

- [1] S. Abramsky and A. Jung. 1994. Domain Theory. In *Handbook of Logic in Computer Science*. Oxford University Press Oxford, UK.
- [2] R. I. Bahar, E. A. Frohm, C. M. Gaona, G. D. Hachtel, E. Macii, A. Pardo, and F. Somenzi. 1997. Algebraic Decision Diagrams and their Applications. *Formal Methods in System Design* 10 (April 1997). Issue 2.
- [3] T. Ball, A. Podelski, and S. K. Rajamani. 2001. Boolean and Cartesian Abstraction for Model Checking C Programs. In *Tools and Algs. for the Construct. and Anal. of Syst. (TACAS'01)*.
- [4] G. Barthe, T. Espitau, L. M. Ferrer Fioriti, and J. Hsu. 2016. Synthesizing Probabilistic Invariants via Doob's Decomposition. In *Computer Aided Verif. (CAV'16)*.
- [5] G. Barthe, T. Espitau, M. Gaboardi, B. Grégoire, J. Hsu, and P.-Y. Strub. 2016. A Program Logic for Probabilistic Programs. Available at justinh.su/files/papers/ellora.pdf.
- [6] G. Barthe, B. Grégoire, and S. Zanella Béguelin. 2009. Formal Certification of Code-based Cryptographic Proofs. In *Princ. of Prog. Lang. (POPL'09)*.
- [7] P. Billingsley. 2012. *Probability and Measure*. John Wiley & Sons, Inc.
- [8] J. Borgström, U. D. Lago, A. D. Gordon, and M. Szymczak. 2016. A Lambda-Calculus Foundation for Universal Probabilistic Programming. In *Int. Conf. on Functional Programming (ICFP'16)*.
- [9] F. Bourdoncle. 1993. Efficient Chaotic Iteration Strategies With Widening. In *Formal Methods in Prog. and Their Applications*.
- [10] T. Brázdil, S. Kiefer, and A. Kučera. 2014. Efficient Analysis of Probabilistic Programs with an Unbounded Counter. *J. ACM* 61 (November 2014). Issue 6.
- [11] T. Brázdil, S. Kiefer, A. Kučera, and I. H. Vařeková. 2015. Runtime Analysis of Probabilistic Programs with Unbounded Recursion. *J. Comput. Syst. Sci.* 81 (February 2015). Issue 1.
- [12] B. Carpenter, A. Gelman, M. D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, and A. Riddell. 2017. Stan: A Probabilistic Programming Language. *J. Statistical Softw.* 76 (2017). Issue 1.
- [13] A. Chakarov and S. Sankaranarayanan. 2013. Probabilistic Program Analysis with Martingales. In *Computer Aided Verif. (CAV'13)*.
- [14] A. Chakarov and S. Sankaranarayanan. 2014. Expectation Invariants for Probabilistic Program Loops as Fixed Points. In *Static Analysis Symp. (SAS'14)*.
- [15] K. Chatterjee, H. Fu, and A. K. Goharshady. 2016. Termination Analysis of Probabilistic Programs Through Positivstellensatz's. In *Computer Aided Verif. (CAV'16)*.
- [16] K. Chatterjee, H. Fu, P. Novotný, and R. Hasheminezhad. 2016. Algorithmic Analysis of Qualitative and Quantitative Termination Problems for Affine Probabilistic Programs. In *Princ. of Prog. Lang. (POPL'16)*.
- [17] K. Chatterjee, P. Novotný, and Đ. Žikelić. 2017. Stochastic Invariants for Probabilistic Termination. In *Princ. of Prog. Lang. (POPL'17)*.
- [18] G. Claret, S. K. Rajamani, A. V. Nori, A. D. Gordon, and J. Borgström. 2013. Bayesian Inference using Data Flow Analysis. In *Found. of Softw. Eng. (FSE'13)*.
- [19] S. Conchon, J.-C. Filliâtre, and J. Signoles. 2007. Designing a Generic Graph Library Using ML Functors. In *Trends in Functional Programming*.
- [20] P. Cousot. 1981. Semantic Foundations of Program Analysis. In *Program Flow Analysis: Theory and Applications*. Prentice-Hall.
- [21] P. Cousot and R. Cousot. 1977. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *Princ. of Prog. Lang. (POPL'77)*.
- [22] P. Cousot and R. Cousot. 1978. Static Determination of Dynamic Properties of Recursive Procedures. In *Formal Descriptions of Programming Concepts, (IFIP WG 2.2, St. Andrews, Canada, August 1977)*. North-Holland.
- [23] P. Cousot and R. Cousot. 1979. Systematic Design of Program Analysis Frameworks. In *Princ. of Prog. Lang. (POPL'79)*.
- [24] P. Cousot and N. Halbwachs. 1978. Automatic Discovery of Linear Constraints Among Variables of a Program. In *Princ. of Prog. Lang. (POPL'78)*.
- [25] P. Cousot and M. Monerau. 2012. Probabilistic Abstract Interpretation. In *European Symp. on Programming (ESOP'12)*.
- [26] E. Darulova and V. Kuncak. 2014. Sound Compilation of Reals. In *Princ. of Prog. Lang. (POPL'14)*.
- [27] J. I. den Hartog and E. P. de Vink. 1999. Mixing Up Nondeterminism and Probability: a preliminary report. *Electr. Notes Theor. Comp. Sci.* 22 (1999).
- [28] E. W. Dijkstra. 1997. *A Discipline of Programming*. Prentice Hall PTR Upper Saddle River.
- [29] K. Etessami, D. Wojtczak, and M. Yannakakis. 2008. Recursive Stochastic Games with Positive Rewards. In *Int. Colloq. on Automata, Langs., and Programming (ICALP'08)*.
- [30] K. Etessami and M. Yannakakis. 2005. Recursive Markov Chains, Stochastic Grammars, and Monotone Systems of Nonlinear Equations. In *Symp. on Theor. Aspects of Comp. Sci. (STACS'05)*.
- [31] K. Etessami and M. Yannakakis. 2015. Recursive Markov Decision Processes and Recursive Stochastic Games. *J. ACM* 62 (May 2015). Issue 2.
- [32] A. Farzan and Z. Kincaid. 2015. Compositional Recurrence Analysis. In *Formal Methods in Computer-Aided Design (FMCAD'15)*.
- [33] L. M. Ferrer Fioriti and H. Hermanns. 2015. Probabilistic Termination: Soundness, Completeness, and Compositionality. In *Princ. of Prog. Lang. (POPL'15)*.
- [34] R. W. Floyd. 1967. Assigning Meanings to Programs. In *Proc. AMS Symposium in Appl. Math.*
- [35] G. Gallo, G. Longo, S. Pallottino, and S. Nguyen. 1993. Directed Hypergraphs and Applications. *Disc. Appl. Math.* 42 (April 1993). Issue 2.
- [36] T. Gehr, S. Misailovic, and M. Vechev. 2016. PSI: Exact Symbolic Inference for Probabilistic Programs. In *Computer Aided Verif. (CAV'16)*.
- [37] Z. Ghahramani. 2015. Probabilistic machine learning and artificial intelligence. *Nature* (2015).
- [38] N. D. Goodman, V. K. Mansinghka, D. M. Roy, and J. B. Tenenbaum. 2008. Church: a language for generative models. In *Uncertainty in Artif. Intelligence*.
- [39] A. D. Gordon, T. A. Henzinger, A. V. Nori, and S. K. Rajamani. 2014. Probabilistic Programming. In *Future of Softw. Eng. (FOSE'14)*.
- [40] S. Graf and H. Saidi. 1997. Construction of Abstract State Graphs with PVS. In *Computer Aided Verif. (CAV'97)*.
- [41] C. A. Gunter, P. D. Mosses, and D. S. Scott. 1989. *Semantic Domains and Denotational Semantics*. Technical Report. University of Pennsylvania Department of Computer and Information Science.
- [42] N. Halbwachs. 1979. *Détermination automatique de relations linéaires vérifiées par les variables d'un programme*. Ph.D. Dissertation. Univ. of Grenoble.
- [43] C. A. R. Hoare. 1969. An Axiomatic Basis for Computer Programming. *Commun. ACM* 12 (October 1969). Issue 10.
- [44] S. Horwitz, T. Reps, and M. Sagiv. 1995. Demand Interprocedural Dataflow Analysis. In *Found. of Softw. Eng. (FSE'95)*.
- [45] B. Jeannot and A. Miné. 2009. Apron: A Library of Numerical Abstract Domains for Static Analysis. In *Computer Aided Verif. (CAV'09)*.
- [46] C. Jones. 1989. *Probabilistic Non-determinism*. Ph.D. Dissertation. University of Edinburgh Edinburgh.
- [47] C. Jones and G. Plotkin. 1989. A Probabilistic Powerdomain of Evaluations. In *Logic in Computer Science (LICS'89)*.
- [48] B. L. Kaminski, J.-P. Katoen, C. Matheja, and F. Olmedo. 2016. Weakest Precondition Reasoning for Expected Run—Times of Probabilistic Programs. In *European Symp. on Programming (ESOP'16)*.

- [49] J.-P. Katoen, A. K. McIver, L. A. Meinicke, and C. C. Morgan. 2010. Linear-Invariant Generation for Probabilistic Programs: Automated Support for Proof-Based Methods. In *Static Analysis Symp. (SAS'10)*.
- [50] M. Kattenbelt, M. Kwiatkowska, G. Norman, and D. Parker. 2009. Abstraction Refinement for Probabilistic Software. In *Verif., Model Checking, and Abs. Interp. (VMCAI'09)*.
- [51] G. A. Kildall. 1973. A Unified Approach to Global Program Optimization. In *Princ. of Prog. Lang. (POPL'73)*.
- [52] J. Knoop and B. Steffen. 1992. The Interprocedural Coincidence Theorem. In *Comp. Construct. (CC'92)*.
- [53] D. E. Knuth. 1977. A Generalization of Dijkstra's Algorithm. *Inf. Proc. Lett.* 6 (February 1977). Issue 1.
- [54] S. Kok, M. Sumner, M. Richardson, P. Singla, H. Poon, D. Lowd, J. Wang, and P. Domingos. 2007. *The Alchemy System for Statistical Relational AI*. Technical Report. University of Washington.
- [55] D. Kozen. 1981. Semantics of Probabilistic Programs. *J. Comput. Syst. Sci.* 22 (June 1981). Issue 3.
- [56] D. Kozen. 1985. A Probabilistic PDL. *J. Comput. Syst. Sci.* 30 (April 1985). Issue 2.
- [57] A. Lal, T. Reps, and G. Balakrishnan. 2005. Extended Weighted Push-down Systems. In *Computer Aided Verif. (CAV'05)*.
- [58] Z. Li and J. Eisner. 2009. First- and Second-Order Expectation Semirings with Applications to Minimum-Risk Training on Translation Forests. In *Conference on Empirical Methods in Natural Language Processing (EMNLP'09)*.
- [59] A. K. McIver and C. C. Morgan. 2001. Partial correctness for probabilistic demonic programs. *Theor. Comp. Sci.* 266 (September 2001). Issue 1.
- [60] A. K. McIver and C. C. Morgan. 2005. *Abstraction, Refinement and Proof for Probabilistic Systems*. Springer Science+Business Media, Inc.
- [61] B. Milch, B. Marthi, S. Russell, D. Sontag, D. L. Ong, and A. Kolobov. 2005. BLOG: Probabilistic Models with Unknown Objects. In *Int. Joint Conf. on Artif. Intelligence (IJCAI'05)*.
- [62] A. Miné. 2006. Symbolic Methods to Enhance the Precision of Numerical Abstract Domains. In *Verif., Model Checking, and Abs. Interp. (VMCAI'06)*.
- [63] T. Minka, J. M. Winn, J. P. Guiver, S. Webster, Y. Zaykov, B. Yangel, A. Spengler, and J. Bronskill. 2014. Infer.NET 2.6. Microsoft Research Cambridge. research.microsoft.com/infernet.
- [64] M. Mislove. 2000. Nondeterminism and Probabilistic Choice: Obeying the Laws. In *Concurrency Theory*.
- [65] M. Mislove, J. Ouaknine, and J. Worrell. 2004. Axioms for Probability and Nondeterminism. *Electr. Notes Theor. Comp. Sci.* 96 (June 2004).
- [66] U. Möncke and R. Wilhelm. 1991. Grammar Flow Analysis. In *Attribute Grammars, Applications and Systems, (Int. Summer School SAGA)*.
- [67] D. Monniaux. 2000. Abstract Interpretation of Probabilistic Semantics. In *Static Analysis Symp. (SAS'00)*.
- [68] D. Monniaux. 2001. Backwards Abstract Interpretation of Probabilistic Programs. In *European Symp. on Programming (ESOP'01)*.
- [69] D. Monniaux. 2003. Abstract Interpretation of Programs as Markov Decision Processes. In *Static Analysis Symp. (SAS'03)*.
- [70] M. Mottl. 2017. Lacaml - Linear Algebra for OCaml. Available at github.com/mmmottl/lacaml.
- [71] M. Müller-Olm and H. Seidl. 2004. Precise Interprocedural Analysis through Linear Algebra. In *Princ. of Prog. Lang. (POPL'04)*.
- [72] F. Olmedo, B. L. Kaminski, J.-P. Katoen, and C. Matheja. 2016. Reasoning about Recursive Probabilistic Programs. In *Logic in Computer Science (LICS'16)*.
- [73] P. Panangaden. 1999. The Category of Markov Kernels. *Electr. Notes Theor. Comp. Sci.* 22 (1999).
- [74] A. Pfeffer. 2005. *The Design and Implementation of IBAL: A General-Purpose Probabilistic Language*. Technical Report. Harvard Computer Science Group.
- [75] M. L. Puterman. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc.
- [76] G. Ramalingam. 1996. *Bounded Incremental Computation*. Springer-Verlag.
- [77] T. Reps, S. Horwitz, and M. Sagiv. 1995. Precise Interprocedural Dataflow Analysis via Graph Reachability. In *Princ. of Prog. Lang. (POPL'95)*.
- [78] M. Sagiv, T. Reps, and S. Horwitz. 1996. Precise Interprocedural Dataflow Analysis with Applications to Constant Propagation. *Theor. Comp. Sci.* 167 (1996). Issue 1.
- [79] S. Sankaranarayanan, A. Chakarov, and S. Gulwani. 2013. Static Analysis for Probabilistic Programs: Inferring Whole Program Properties from Finitely Many Paths. In *Prog. Lang. Design and Impl. (PLDI'13)*.
- [80] M. Sharir and A. Pnueli. 1981. Two Approaches to Interprocedural Data Flow Analysis. In *Program Flow Analysis: Theory and Applications*. Prentice-Hall.
- [81] R. E. Tarjan. 1981. A Unified Approach to Path Problems. *J. ACM* 28 (July 1981). Issue 3.
- [82] R. Tix, K. Keimel, and G. Plotkin. 2009. Semantic Domains for Combining Probability and Non-Determinism. *Electr. Notes Theor. Comp. Sci.* 222 (February 2009).
- [83] D. Wang, J. Hoffmann, and T. Reps. 2018. A Denotational Semantics for Nondeterminism in Probabilistic Programs. Available at www.cs.cmu.edu/~diw3/papers/WangHR18.pdf.
- [84] D. Wojtczak and K. Etessami. 2017. PReMo - Probabilistic Recursive Models analyzer. Available at groups.inf.ed.ac.uk/premo/.