# Guaranteed Bounds for Posterior Inference in Universal Probabilistic Programming

Raven Beutner\* CISPA Helmholtz Center for Information Security Germany C.-H. Luke Ong University of Oxford United Kingdom Fabian Zaiser University of Oxford United Kingdom

# Abstract

We propose a new method to approximate the posterior distribution of probabilistic programs by means of computing guaranteed bounds. The starting point of our work is an interval-based trace semantics for a recursive, higher-order probabilistic programming language with continuous distributions. Taking the form of (super-/subadditive) measures, these lower/upper bounds are non-stochastic and provably correct: using the semantics, we prove that the actual posterior of a given program is sandwiched between the lower and upper bounds (soundness); moreover, the bounds converge to the posterior (completeness). As a practical and sound approximation, we introduce a weight-aware interval type system, which automatically infers interval bounds on not just the return value but also the weight of program executions, simultaneously. We have built a tool implementation, called GuBPI, which automatically computes these posterior lower/upper bounds. Our evaluation on examples from the literature shows that the bounds are useful, and can even be used to recognise wrong outputs from stochastic posterior inference procedures.

CCS Concepts: • Mathematics of computing  $\rightarrow$  Probabilistic inference problems; • Theory of computation  $\rightarrow$  Program analysis; • Software and its engineering  $\rightarrow$  Formal methods.

*Keywords:* probabilistic programming, Bayesian inference, verification, abstract interpretation, operational semantics, interval arithmetic, type system, symbolic execution

## **ACM Reference Format:**

Raven Beutner, C.-H. Luke Ong, and Fabian Zaiser. 2022. Guaranteed Bounds for Posterior Inference in Universal Probabilistic Programming. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI '22), June 13–17, 2022, San Diego, CA, USA.* ACM, New York, NY, USA, 32 pages. https://doi.org/10.1145/3519939.3523721

\*Member of the Saarbrücken Graduate School of Computer Science.

PLDI '22, June 13-17, 2022, San Diego, CA, USA

© 2022 Copyright held by the owner/author(s).

# 1 Introduction

Probabilistic programming is a rapidly developing discipline at the interface of programming and Bayesian statistics [32, 33, 62]. The idea is to express probabilistic models (incorporating the prior distributions) and the observed data as programs, and to use a general-purpose Bayesian inference engine, which acts directly on these programs, to find the posterior distribution given the observations.

Some of the most influential probabilistic programming languages (PPLs) used in practice are *universal* (i.e. the underlying language is Turing-complete); e.g. Church [31], Anglican [61], Gen [18], Pyro [5], and Turing [25]. Using stochastic branching, recursion, and higher-order features, universal PPLs can express arbitrarily complex models. For instance, these language constructs can be used to incorporate probabilistic context free grammars [43], statistical phylogenetics [52], and even physics simulations [3] into probabilistic models. However, expressivity of the PPL comes at the cost of complicating the posterior inference. Consider, for example, the following problem from [41, 42].

**Example 1.1** (Pedestrian). A pedestrian has gotten lost on a long road and only knows that they are a random distance between 0 and 3 km from their home. They repeatedly walk a uniform random distance of at most 1 km in either direction, until they find their home. When they arrive, a step counter tells them that they have traveled a distance of 1.1 km in total. Assuming that the measured distance is normally distributed around the true distance with standard deviation 0.1 km, what is the posterior distribution of the starting point? We can model this with a probabilistic program:

```
let start = 3 \times sample uniform(0, 1) in
letrec walk x = if x \le 0 then 0 else
let step = sample uniform(0, 1) in
step + walk((x + step) \oplus_{0.5} (x - step))
let distance = walk start in
observe distance from Normal(1.1, 0.1);
start
```

Here **sample** uniform(a, b) samples a uniformly distributed value in [a, b],  $\oplus_{0.5}$  is probabilistic branching, and **observe** *M* **from** *D* observes the value of *M* from distribution *D*.

Example 1.1 is a challenging model for inference algorithms in several regards: not only does the program use

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI '22), June 13–17, 2022, San Diego, CA, USA*, https://doi.org/10.1145/3519939.3523721.



**Figure 1.** Histogram of samples from the posterior distribution of Example 1.1 and wrong samples produced by the probabilistic programming system Pyro.

stochastic branching and recursion, but the number of random variables generated is unbounded – it's *nonparametric* [29, 37, 42]. To approximate the posterior distribution of the program, we apply two standard inference algorithms: likelihood-weighted importance sampling (IS), a simple algorithm that works well on low-dimensional models with few observations [49]; and Hamiltonian Monte Carlo (HMC) [21], a successful MCMC algorithm that uses gradient information to efficiently explore the parameter space of highdimensional models. Figure 1 shows the results of the two inference methods as implemented in Anglican [61] (for IS) and Pyro [5] (for HMC): they clearly disagree! But how is the user supposed to know which (if any) of the two results is correct?

Note that *exact* inference methods (i.e. methods that try to compute a closed-form solution of the posterior inference problem using computer algebra and other forms of symbolic computation) such as PSI [26, 27], Hakaru [47], Dice [38], and SPPL [55] are only applicable to non-recursive models, and so they don't work for Example 1.1.

#### 1.1 Guaranteed Bounds

The above example illustrates central problems with both approximate stochastic and exact inference methods. For approximate methods, there are no guarantees for the results they output after a finite amount of time, leading to unclear inference results (as seen in Fig. 1).<sup>1</sup> For exact methods, the symbolic engine may fail to find a closed-form description of the posterior distribution and, more importantly, they are

only applicable to very restricted classes of programs (most notably, non-recursive models).

Instead of computing approximate or exact results, this work is concerned with computing guaranteed bounds on the posterior distribution of a probabilistic program. Concretely, given a probabilistic program *P* and a measurable set  $U \subseteq \mathbb{R}$ (given as an interval), we infer upper and lower bounds on  $\llbracket P \rrbracket(U)$  (formally defined in Section 2), i.e. the posterior probability of *P* on  $U^2$ . Such bounds provide a ground truth to compare approximate inference results with: if the approximate results violate the bounds, the inference algorithm has not converged yet or is even ill-suited to the program in question. Crucially, our method is applicable to arbitrary (and in particular recursive) programs of a universal PPL. For Example 1.1, the bounds computed by our method (which we give in Section 7) are tight enough to separate the IS and HMC output. In this case, our method infers that the results given by HMC are wrong (i.e. violate the guaranteed bounds) whereas the IS results are plausible (i.e. lie within the guaranteed bounds). To the best of our knowledge, no existing methods can provide such definite answers for programs of a universal PPL.

# 1.2 Contributions

The starting point of our work is an interval-based operational semantics [4]. In our semantics, we evaluate a program on interval traces (i.e. sequences of intervals of reals with endpoints between 0 and 1) to approximate the outcomes of sampling, and use interval arithmetic [19] to approximate numerical operations (Section 3). Our semantics is sound in the sense that any (compatible and exhaustive) set of interval traces yields lower and upper bounds on the posterior distribution of a program. These lower/upper bounds are themselves super-/subadditive measures. Moreover, under mild conditions (mostly restrictions on primitive operations), our semantics is also complete, i.e. for any  $\epsilon > 0$  there exists a countable set of interval traces that provides  $\epsilon$ -tight bounds on the posterior. Our proofs hinge on a combination of stochastic symbolic execution and the convergence of Riemann sums, providing a natural correspondence between our interval trace semantics and the theory of (Riemann) integration (Section 4).

Based on our interval trace semantics, we present a practical algorithm to automate the computation of guaranteed bounds. It employs an interval type system (together with constraint-based type inference) that bounds both the value of an expression in a refinement-type fashion *and* the score weight of any evaluation thereof. The (interval) bounds inferred by our type system fit naturally in the domain of our semantics. This enables a sound approximation of the

<sup>&</sup>lt;sup>1</sup>Take MCMC sampling algorithms. Even though the Markov chain will eventually converge to the target distribution, we do not know how long to iterate the chain to ensure convergence [49, 53]. Likewise for variational inference [64]: given a variational family, there is no guarantee that a given value for the KL-divergence (from the approximating to the posterior distribution) is attainable by the minimising distribution.

<sup>&</sup>lt;sup>2</sup>By repeated application of our method on a discretisation of the domain we can compute histogram-like bounds.

behaviour of a program with finitely many interval traces (Section 5).

We implemented our approach in a tool called GuBPI<sup>3</sup> (**Gu**aranteed **B**ounds for **P**osterior Inference), described in Section 6, and evaluate it on a suite of benchmark programs from the literature. We find that the bounds computed by GuBPI are competitive in many cases where the posterior could already be inferred exactly. Moreover, GuBPI's bounds are useful (in the sense that they are precise enough to rule out erroneous approximate results as in Fig. 1, for instance) for recursive models that could not be handled rigorously by any method before (Section 7).

#### 1.3 Scope and Limitations

The contributions of this paper are of both theoretical and practical interest. On the theoretical side, our novel semantics underpins a sound and deterministic method to compute guaranteed bounds on program denotations. As shown by our completeness theorem, this analysis is applicable-in the sense that it computes arbitrarily tight bounds-to a very broad class of programs. On the practical side, our analyser GuBPI implements (an optimised version of) our semantics. As is usual for exact/guaranteed<sup>4</sup> methods, our semantics considers an exponential number of program paths, and partitions each sampled value into a finite number of interval approximations. Consequently, GuBPI generally struggles with high-dimensional models. We believe GuBPI to be most useful for unit-testing of implementations of Bayesian inference algorithms such as Example 1.1, or to compute results on (recursive) programs when non-stochastic, guaranteed bounds are needed.

# 2 Background

# 2.1 Basic Probability Theory and Notation

We assume familiarity with basic probability theory, and refer to [51] for details. Here we just fix the notation. A *measurable space* is a pair  $(\Omega, \Sigma_{\Omega})$  where  $\Omega$  is a set (of outcomes) and  $\Sigma_{\Omega} \subseteq 2^{\Omega}$  is a  $\sigma$ -algebra defining the measurable subsets of  $\Omega$ . A *measure* on  $(\Omega, \Sigma_{\Omega})$  is a function  $\mu : \Sigma_{\Omega} \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ that satisfies  $\mu(\emptyset) = 0$  and is  $\sigma$ -additive. For  $\mathbb{R}^n$ , we write  $\Sigma_{\mathbb{R}^n}$  for the Borel  $\sigma$ -algebra and  $\lambda_n$  for the Lebesgue measure on  $(\mathbb{R}^n, \Sigma_{\mathbb{R}^n})$ . The Lebesgue integral of a measurable function f with respect to a measure  $\mu$  is written  $\int f d\mu$ or  $\int f(x) \mu(dx)$ . Given a predicate  $\psi$  on  $\Omega$ , we define the Iverson brackets  $[\psi] : \Omega \to \mathbb{R}$  by mapping all elements that satisfy  $\psi$  to 1 and all others to 0. For  $A \in \Sigma_{\Omega}$  we define the bounded integral  $\int_A f d\mu := \int f(x) \cdot [x \in A] \mu(dx)$ .

$$\begin{array}{c} \hline ((\lambda x.M)V,\mathbf{s},w) \rightarrow (M[V/x],\mathbf{s},w) & \hline (\operatorname{sample},r\,\mathbf{s},w) \rightarrow (\underline{r},\mathbf{s},w) \\ \hline \hline ((\mu_x^{\varphi}.M)V,\mathbf{s},w) \rightarrow (M[V/x,(\mu_x^{\varphi}.M)/\varphi],\mathbf{s},w) \\ \hline \hline (f(\underline{r}_1,\ldots,\underline{r}_{|f|}),\mathbf{s},w) \rightarrow (f(r_1,\ldots,r_{|f|}),\mathbf{s},w) \\ \hline \hline (\underline{r \leq 0} & r > 0 \\ \hline (\operatorname{if}(\underline{r},N,P),\mathbf{s},w) \rightarrow (N,\mathbf{s},w) & \hline (\operatorname{if}(\underline{r},N,P),\mathbf{s},w) \rightarrow (P,\mathbf{s},w) \\ \hline \hline (\operatorname{score}(\underline{r}),\mathbf{s},w) \rightarrow (\underline{r},\mathbf{s},w\cdot r) & \hline (E[R],\mathbf{s},w) \rightarrow (E[M],\mathbf{s}',w') \end{array}$$

**Figure 2.** Standard (CbV) reduction rules for SPCF  $(\rightarrow)$ .

### 2.2 Statistical PCF (SPCF)

As our probabilistic programming language of study, we use *statistical PCF* (SPCF) [41], a typed variant of [7]. SPCF includes primitive operations which are measurable functions  $f : \mathbb{R}^{|f|} \to \mathbb{R}$ , where  $|f| \ge 0$  denotes the arity of the function. *Values* and *terms* of SPCF are defined as follows:

$$V := x \mid \underline{r} \mid \lambda x.M \mid \mu_x^{\varphi}.M$$
$$M, N, P := V \mid MN \mid \text{if}(M, N, P) \mid \underline{f}(M_1, \dots, M_{|f|})$$
$$\mid \text{sample} \mid \text{score}(M)$$

where *x* and  $\varphi$  are variables, *f* is a primitive operation, and <u>r</u> a constant with  $r \in \mathbb{R}$ . Note that we write  $\mu_x^{\varphi}$ . M instead of  $Y(\lambda \varphi x.M)$  for the fixpoint construct. The branching construct is if (M, N, P), which evaluates to N if  $M \leq 0$  and *P* otherwise. In SPCF, sample draws a random value from the uniform distribution on [0, 1], and score(M) weights the current execution with the value of M. Samples from a different real-valued distribution D can be obtained by applying the inverse of the cumulative distribution function for D to a uniform sample [54]. Most PPLs feature an observe statement instead of manipulating the likelihood weight directly with score, but they are equally expressive [59].<sup>5</sup> As usual, we write let x = M in N for  $(\lambda x.N)M, M; N$ for let  $\_$  = M in N and  $M \oplus_p N$  for if (sample -p, M, N). The type system of our language is as expected, with simple types being generated by  $\alpha, \beta := \mathbf{R} \mid \alpha \to \beta$ . Selected rules are given below:

$$\frac{\Gamma \vdash sample : \mathbf{R}}{\Gamma \vdash score(M) : \mathbf{R}} = \frac{\Gamma \vdash M : \mathbf{R}}{\Gamma \vdash score(M) : \mathbf{R}} \\
\frac{\Gamma, \varphi : \alpha \to \beta, x : \alpha \vdash M : \beta}{\Gamma \vdash \mu_x^{\varphi} \cdot M : \alpha \to \beta} = \frac{\{\Gamma \vdash M_i : \mathbf{R}\}_{i=1}^{|f|}}{\Gamma \vdash \underline{f}(M_1, \dots, M_{|f|}) : \mathbf{R}}$$

<sup>&</sup>lt;sup>3</sup>GuBPI (pronounced "guppy") is available at gubpi-tool.github.io.

<sup>&</sup>lt;sup>4</sup>By "exact/guaranteed methods", we mean inference algorithms that compute deterministic (non-stochastic) results about the mathematical denotation of a program. In particular, they are correct with probability 1, contrary to stochastic methods.

<sup>&</sup>lt;sup>5</sup> In Bayesian terms, an **observe** statement multiplies the likelihood function by the probability (density) of the observation [33] (as we have used in Example 1.1). Scoring makes this explicit by keeping a weight for each program execution [7]. Observing a value v from a distribution D then simply multiplies the current weight by  $pdf_D(v)$  where  $pdf_D$  is the probability density function of D (for continuous distributions) or the probability mass function of D (for discrete distributions).

#### 2.3 Trace Semantics

Following [7], we endow SPCF with a trace-based operational semantics. We evaluate a probabilistic program P on a fixed *trace*  $s = \langle r_1, \ldots, r_n \rangle \in \mathbb{T} := \bigcup_{n \in \mathbb{N}} [0, 1]^n$ , which predetermines the probabilistic choices made during the evaluation. Our semantics therefore operates on configurations of the form (M, s, w) where M is an SPCF term, s is a trace and  $w \in \mathbb{R}_{>0}$  a weight. The call-by-value (CbV) reduction is given by the rules in Fig. 2, where  $E[\cdot]$  denotes a CbV evaluation context. The definition is standard [4, 7, 41]. Given a program  $\vdash P$  : **R**, we call a trace *s* terminating just if  $(P, \mathbf{s}, 1) \rightarrow^* (V, \langle \rangle, w)$  for some value V and weight w, i.e. if the samples drawn are as specified by s, the program P terminates. Note that we require the trace *s* to be completely used up. As P is of type **R** we can assume that V = r for some  $r \in \mathbb{R}$ . Each terminating trace *s* therefore uniquely determines the returned value  $\underline{r}$  where  $r =: \operatorname{val}_{P}(\mathbf{s}) \in \mathbb{R}$ , and the *weight*  $w =: wt_P(s) \in \mathbb{R}_{\geq 0}$ , of the execution. For a nonterminating trace s, val<sub>*P*</sub>(s) is undefined and wt<sub>*P*</sub>(s) := 0.

**Example 2.1.** Consider Example 1.1. On the trace  $s = \langle 0.1, 0.2, 0.4, 0.7, 0.8 \rangle \in [0, 1]^5 \subseteq \mathbb{T}$ , the pedestrian walks 0.2 away from their home (taking the left branch of  $\oplus_{0.5}$  as  $0.4 \le 0.5$ ) and 0.7 towards their home (as 0.8 > 0.5), hence:

 $val_P(s) = 3 \times 0.1 = 0.3, wt_P(s) = pdf_{Normal(1,1,0,1)}(0.9).$ 

In order to do measure theory, we need to turn our set of traces into a measurable space. The trace space  $\mathbb{T}$  is equipped with the  $\sigma$ -algebra  $\Sigma_{\mathbb{T}} := \{\bigcup_{n \in \mathbb{N}} U_n \mid U_n \in \Sigma_{[0,1]^n}\}$  where  $\Sigma_{[0,1]^n}$  is the Borel  $\sigma$ -algebra on  $[0, 1]^n$ . We define a measure  $\mu_{\mathbb{T}}$  by  $\mu_{\mathbb{T}}(U) := \sum_{n \in \mathbb{N}} \lambda_n (U \cap [0, 1]^n)$ , as in [7].

We can now define the semantics of an SPCF program  $\vdash P : \mathbf{R}$  by using the weight and returned value of (executions of *P* determined by) individual traces. Given  $U \in \Sigma_{\mathbb{R}}$ , we need to define the likelihood of *P* evaluating to a value in *U*. To this end, we set  $\operatorname{val}_{P}^{-1}(U) := \{\mathbf{s} \in \mathbb{T} \mid (P, \mathbf{s}, 1) \rightarrow^{*} (\underline{r}, \langle\rangle, w), r \in U\}$ , i.e. the set of traces on which the program *P* reduces to a value in *U*. As shown in [7, Lem. 9],  $\operatorname{val}_{P}^{-1}(U)$  is measurable. Thus, we can define (cf. [7, 41])

$$\llbracket P \rrbracket(U) := \int_{\operatorname{val}_P^{-1}(U)} \operatorname{wt}_P(s) \,\mu_{\mathbb{T}}(\mathrm{d}s).$$

That is, the integral takes all traces s on which P evaluates to a value in U, weighting each s with the weight wt<sub>P</sub>(s) of the corresponding execution. A program P is called **almost** surely terminating (AST) if it terminates with probability 1, i.e.  $\mu_{\mathbb{T}}(val_P^{-1}(\mathbb{R})) = 1$ . This is a necessary assumption for approximate inference algorithms (since they execute the program). See [7] for a more in-depth discussion of this (standard) sampling-style semantics.

*Normalizing constant and integrability.* In Bayesian statistics, one is usually interested in the *normalised* posterior, which is a conditional probability distribution. We obtain the normalised denotation as posterior<sub>*p*</sub> :=  $\frac{\|P\|}{Z_p}$  where

 $Z_P := \llbracket P \rrbracket(\mathbb{R})$  is the *normalising constant*. We call *P integrable* if  $0 < Z_P < \infty$ . The bounds computed in this paper (on the unnormalised denotation  $\llbracket P \rrbracket$ ) allow us to compute bounds on the normalizing constant  $Z_P$ , and thereby also on the normalised denotation. All bounds reported in this paper (in particular in Section 7) refer to the *normalised* denotation.

## **3** Interval Trace Semantics

In order to obtain guaranteed bounds on the distribution denotation  $[\![P]\!]$  (and also on posterior<sub>*P*</sub>) of a program *P*, we present an interval-based semantics. In our semantics, we approximate the outcomes of sample with intervals and handle arithmetic operations by means of interval arithmetic (which is similar to the approach by Beutner and Ong [4] in the context of termination analysis). Our semantics enables us to reason about the denotation of a program *without* considering the uncountable space of traces explicitly.

# 3.1 Interval Arithmetic

For our purposes, an *interval* has the form [a, b] which denotes the set  $\{x \in \mathbb{R} \mid a \le x \le b\}$ , where  $a \in \mathbb{R} \cup \{-\infty\}$ ,  $b \in \mathbb{R} \cup \{\infty\}$ , and  $a \le b$ . For consistency, we write  $[0, \infty]$  instead of the more typical  $[0, \infty)$ . For  $X \subseteq \mathbb{R} \cup \{-\infty, \infty\}$ , we denote by  $\mathbb{I}_X$  the set of intervals with endpoints in X, and simply write  $\mathbb{I}$  for  $\mathbb{I}_{\mathbb{R} \cup \{-\infty,\infty\}}$ . An *n*-dimensional *box* is the Cartesian product of *n* intervals.

We can lift functions on real numbers to intervals as follows: for each  $f : \mathbb{R}^n \to \mathbb{R}$  we define  $f^{\mathbb{I}} : \mathbb{I}^n \to \mathbb{I}$  by

$$f^{\mathbb{I}}([a_1, b_1], \dots, [a_n, b_n]) := [\inf F, \sup F]$$

where  $F := f([a_1, b_1], ..., [a_n, b_n])$ . For common functions like +, -, ×,  $|\cdot|$ , min, max, and monotonically increasing or decreasing functions  $f : \mathbb{R} \to \mathbb{R}$ , their interval-lifted counterparts can easily be computed, from the values of the original function on just the endpoints of the input interval. For example, addition lifts to  $[a_1, b_1] + [a_2, b_2] = [a_1 + a_2, b_1 + b_2]$ ; similarly for multiplication ×<sup>I</sup>.

#### 3.2 Interval Traces and Interval SPCF

In our interval interpretation, probabilistic programs are run on *interval traces*. An *interval trace*,  $\langle I_1, \ldots, I_n \rangle \in \mathbb{T}_{\mathbb{I}} := \bigcup_{n \in \mathbb{N}} (\mathbb{I}_{[0,1]})^n$ , is a finite sequence of intervals  $I_1, \ldots, I_n$ , each with endpoints between 0 and 1. To distinguish ordinary traces  $s \in \mathbb{T}$  from interval traces  $t \in \mathbb{T}_{\mathbb{I}}$ , we call the former *concrete* traces.

We define the *refinement* relation  $\triangleleft$  between concrete and interval traces as follows: for  $s = \langle r_1, \ldots, r_n \rangle \in \mathbb{T}$  and  $t = \langle I_1, \ldots, I_m \rangle \in \mathbb{T}_{\mathbb{I}}$ , we define  $s \triangleleft t$  just if n = m and for all  $i, r_i \in I_i$ . For each interval trace t, we denote by  $|\langle t \rangle := \{s \in \mathbb{T} \mid s \triangleleft t\}$  the set of all refinements of t.

To define a reduction of a term on an interval trace, we extend SPCF with *interval literals* [a, b], which replace the literals *r* but are still considered values of type **R**. In fact, *r* 

$\overline{((\lambda x.M)V, t, w) \rightarrow_{\mathbb{I}} (M[V/x], t, w)}$	$\overline{(\text{sample}, It, w)} \rightarrow_{\mathbb{I}} (\underline{I}, t, w)$
$((\mu_x^{\varphi}.M)V, t, w) \to_{\mathbb{I}} (M[V])$	$/x, (\mu_x^{\varphi}.M)/\varphi], t, w)$
$\frac{b \le 0}{(if([a, b], N, P), t, w)}$	$\rightarrow_{\mathbb{I}} (N, t, w)$
$\frac{a > 0}{(if(\underline{[a,b]}, N, P), t, w)}$	$) \rightarrow_{\mathbb{I}} (P, t, w)$
$(f(\underline{I_1},\ldots,\underline{I_{ f }}),\boldsymbol{t},\boldsymbol{w}) \to_{\mathbb{I}} (f(\underline{I_1},\ldots,\underline{I_{ f }}),\boldsymbol{w})$	$f^{\mathbb{I}}(I_1,\ldots,I_{ f }), t, w)$
$\frac{a \ge 0}{(a + b) + c}$	
$(\text{score}(\underline{[a, b]}), t, w) \to_{\mathbb{I}} (\underline{[a, b]})$ $(R, t, w) \to_{\mathbb{I}} (R, t, w) \to_{\mathbb{I}} (R, t, w) \to_{\mathbb{I}} (R, t, w)$	$\frac{[a, b]}{[a, t', w']}$

**Figure 3.** Interval reduction rules for (interval) SPCF  $(\rightarrow_{\mathbb{I}})$ .

can be read as an abbreviation for [r, r]. We call such terms *interval terms*, and the resulting language *Interval SPCF*.

**Reduction.** The interval-based reduction  $\rightarrow_{\mathbb{I}}$  now operates on configurations (M, t, w) of interval terms M, interval traces  $t \in \mathbb{T}_{\mathbb{I}}$ , and interval weights  $w \in \mathbb{I}_{\mathbb{R}_{\geq 0} \cup \{\infty\}}$ . The redexes and evaluation contexts of SPCF extend naturally to interval terms. The reduction rules are given in Fig. 3.<sup>6</sup>

Given a program  $\vdash P : \mathbf{R}$ , the reduction relation  $\rightarrow_{\mathbb{I}}$  allows us to define the *interval weight* function  $(\mathsf{wt}_P^{\mathbb{I}} : \mathbb{T}_{\mathbb{I}} \rightarrow \mathbb{I}_{\mathbb{R}_{>0} \cup \{\infty\}})$  and *interval value* function  $(\mathsf{val}_P^{\mathbb{I}} : \mathbb{T}_{\mathbb{I}} \rightarrow \mathbb{I})$  by:

$$\mathsf{wt}_P^{\mathbb{I}}(t) \coloneqq \begin{cases} w & \text{if } (P, t, 1) \to_{\mathbb{I}}^* (\underline{I}, \langle \rangle, w) \\ [0, \infty] & \text{otherwise,} \end{cases}$$

$$\mathsf{val}_P^{\mathbb{I}}(t) \coloneqq \begin{cases} I & \text{if } (P, t, 1) \to_{\mathbb{I}}^* (\underline{I}, \langle \rangle, w) \\ [-\infty, \infty] & \text{otherwise.} \end{cases}$$

It is not difficult to prove the following relationship between standard and interval reduction.

**Lemma 3.1.** Let  $\vdash P : \mathbf{R}$  be a program. For any interval trace t and concrete trace  $s \triangleleft t$ , we have  $wt_P(s) \in wt_P^{\mathbb{I}}(t)$  and  $val_P(s) \in val_P^{\mathbb{I}}(t)$  (provided  $val_P(s)$  is defined).

# 3.3 Bounds from Interval Traces

*Lower bounds.* How can we use this interval trace semantics to obtain lower bounds on  $[\![P]\!]$ ? We need a few definitions. Two intervals  $[a_1, b_1], [a_2, b_2] \in \mathbb{I}$  are called *almost disjoint* if  $b_1 \leq a_2$  or  $b_2 \leq a_1$ . Interval traces  $\langle I_1, \ldots, I_m \rangle$  and  $\langle J_1, \ldots, J_n \rangle \in \mathbb{T}_I$  are called *compatible* if there is an index  $i \in \{1, \ldots, \min(m, n)\}$  such that  $I_i$  and  $J_i$  are almost disjoint.

A set of interval traces is called *compatible* if its elements are pairwise compatible. We define the *volume* of an interval trace  $t = \langle [a_1, b_1], ..., [a_n, b_n] \rangle$  as vol $(t) := \prod_{i=1}^n (b_i - a_i)$ .

Let  $\mathcal{T} \subseteq \mathbb{T}_{\mathbb{I}}$  be a countable and compatible set of interval traces. Define the *lower bound* on  $\llbracket P \rrbracket$  by

$$\mathsf{lowerBd}_{P}^{\mathcal{T}}(U) \coloneqq \sum_{t \in \mathcal{T}} \mathsf{vol}(t) \cdot (\min \mathsf{wt}_{P}^{\mathbb{I}}(t)) \cdot [\mathsf{val}_{P}^{\mathbb{I}}(t) \subseteq U]$$

for  $U \in \Sigma_{\mathbb{R}}$ . That is, we sum over each interval trace in  $\mathcal{T}$  whose value is *guaranteed* to be in U, weighted by its volume and the lower bound of its weight interval. Note that, in general, lowerBd $_{P}^{\mathcal{T}}$  is not a measure, but merely a *superadditive measure*.<sup>7</sup>

Upper bounds. For upper bounds, we require the notion of a set of interval traces being *exhaustive*, which is easiest to express in terms of infinite traces. Let  $\mathbb{T}_{\infty} := [0, 1]^{\omega}$  be the set of infinite traces. Every interval trace *t* covers the set of infinite traces with a prefix contained in *t*, i.e. cover(t) := $|t| > \mathbb{T}_{\infty}$  (where the Cartesian product × can be viewed as trace concatenation). A countable set of (finite) interval traces  $\mathcal{T} \subseteq \mathbb{T}_{\mathbb{I}}$  is called *exhaustive* if  $\bigcup_{t \in \mathcal{T}} cover(t)$  covers almost all of  $\mathbb{T}_{\infty}$ , i.e.  $\mu_{\mathbb{T}_{\infty}}(\mathbb{T}_{\infty} \setminus \bigcup_{t \in \mathcal{T}} cover(t)) = 0.^8$  Phrased differently, almost all concrete traces must have a finite prefix that is contained in some interval trace in  $\mathcal{T}$ . Therefore, the analysis in the interval semantics on  $\mathcal{T}$  covers the behaviour on almost all concrete traces (in the original semantics).

**Example 3.1.** (i) The singleton set { $\langle [0, 1], [0, 0.6] \rangle$ } is not exhaustive as, e.g. all infinite traces  $\langle r_1, r_2, ... \rangle$  with  $r_2 > 0.6$  are not covered. (ii) The set { $\langle [0, 0.6] \rangle$ ,  $\langle [0.3, 1] \rangle$ } is exhaustive, but not compatible. (iii) Define  $\mathcal{T}_1 := \{\langle [\frac{1}{2}, 1]^{...n}, [0, \frac{1}{3}] \rangle | n \in \mathbb{N}\}$  and  $\mathcal{T}_2 := \{\langle [\frac{1}{2}, 1]^{...n}, [0, \frac{1}{2}] \rangle | n \in \mathbb{N}\}$  where  $x^{...n}$  denotes *n*-fold repetition of *x*.  $\mathcal{T}_1$  is compatible but not exhaustive. For example, it doesn't cover the set  $[\frac{1}{2}, 1] \times (\frac{1}{3}, \frac{1}{2}) \times \mathbb{T}_{\infty}$ , i.e. all traces  $\langle r_1, r_2, ... \rangle$  where  $r_1 \in [\frac{1}{2}, 1]$  and  $r_2 \in (\frac{1}{3}, \frac{1}{2})$ .  $\mathcal{T}_2$  is compatible and exhaustive (the set of non-covered traces  $(\frac{1}{2}, 1]^{\omega}$  has measure 0).

Let  $\mathcal{T} \subseteq \mathbb{T}_{\mathbb{I}}$  be a countable and exhaustive set of interval traces. Define the *upper bound* on  $[\![P]\!]$  by

upperBd<sup>T</sup><sub>P</sub>(U) := 
$$\sum_{t \in T} \operatorname{vol}(t) \cdot (\sup \operatorname{wt}_{P}^{\mathbb{I}}(t)) \cdot [\operatorname{val}_{P}^{\mathbb{I}}(t) \cap U \neq \emptyset]$$

for  $U \in \Sigma_{\mathbb{R}}$ . That is, we sum over each interval trace in  $\mathcal{T}$  whose value *may* be in *U*, weighted by its volume and the upper bound of its weight interval. Note that upperBd<sup> $\mathcal{T}$ </sup><sub>*p*</sub> is not a measure but only a *subadditive measure*.<sup>9</sup>

<sup>&</sup>lt;sup>6</sup>For conditionals, the interval bound is not always precise enough to decide which branch to take, so the reduction can get stuck if  $a \le 0 < b$ . We could include additional rules to overapproximate the branching behaviour (see Appendix A.4). But the rules given here simplify the presentation and are sufficient to prove soundness and completeness.

<sup>&</sup>lt;sup>7</sup> A superadditive measure  $\mu$  on  $(\Omega, \Sigma_{\Omega})$  is a measure, except that  $\sigma$ -additivity is replaced by  $\sigma$ -superadditivity:  $\mu(\bigcup_{i \in \mathbb{N}} U_i) \ge \sum_{i \in \mathbb{N}} \mu(U_i)$  for a countable, pairwise disjoint family  $(U_i)_{i \in \mathbb{N}} \in \Sigma_{\Omega}$ .

<sup>&</sup>lt;sup>8</sup>The  $\sigma$ -algebra on  $\mathbb{T}_{\infty}$  is defined as the smallest  $\sigma$ -algebra that contains all sets  $U \times \mathbb{T}_{\infty}$  where  $U \in \Sigma_{[0,1]^n}$  for some  $n \in \mathbb{N}$ . The measure  $\mu_{\mathbb{T}_{\infty}}$  is the unique measure with  $\mu_{\mathbb{T}_{\infty}}(U \times \mathbb{T}_{\infty}) = \lambda_n(U)$  when  $U \in \Sigma_{[0,1]^n}$ .

<sup>&</sup>lt;sup>9</sup>A subadditive measure  $\mu$  on  $(\Omega, \Sigma_{\Omega})$  is a measure, except that  $\sigma$ -additivity is replaced by  $\sigma$ -subadditivity:  $\mu(\bigcup_{i \in \mathbb{N}} U_i) \leq \sum_{i \in \mathbb{N}} \mu(U_i)$  for a countable, pairwise disjoint family  $(U_i)_{i \in \mathbb{N}} \in \Sigma_{\Omega}$ .

# 4 Soundness and Completeness

# 4.1 Soundness

We show that the two bounds described above are *sound*, in the following sense.

**Theorem 4.1** (Sound lower bounds). Let  $\mathcal{T}$  be a countable and compatible set of interval traces and  $\vdash P : \mathbf{R}$  a program. Then lower  $\mathsf{Bd}_P^{\mathcal{T}} \leq [\![P]\!]$ .

*Proof.* For any  $U \in \Sigma_{\mathbb{R}}$ , we have:

$$\operatorname{IowerBd}_{P}^{\mathcal{T}}(U) = \sum_{t \in \mathcal{T}} \operatorname{vol}(t) (\min \operatorname{wt}_{P}^{\mathbb{I}}(t)) \left[ \operatorname{val}_{P}^{\mathbb{I}}(t) \subseteq U \right]$$
$$= \sum_{t \in \mathcal{T}} \int_{(t)} (\min \operatorname{wt}_{P}^{\mathbb{I}}(t)) \left[ \operatorname{val}_{P}^{\mathbb{I}}(t) \subseteq U \right] \mathrm{d}s$$
$$\leq \sum_{t \in \mathcal{T}} \int_{(t)} \operatorname{wt}_{P}(s) \left[ \operatorname{val}_{P}(s) \in U \right] \mathrm{d}s \qquad (1)$$

$$= \int_{\bigcup_{t \in \mathcal{T}} (t)} \mathrm{wt}_P(s) \left[ \mathrm{val}_P(s) \in U \right] \mathrm{d}s \qquad (2)$$

$$\leq \int_{\mathbb{T}} \operatorname{wt}_{P}(s) \left[ \operatorname{val}_{P}(s) \in U \right] \mathrm{d}s = \llbracket P \rrbracket(U) \quad (3)$$

where Eq. (1) follows from Lemma 3.1, Eq. (2) from compatibility, and Eq. (3) from  $\bigcup_{t \in \mathcal{T}} (t) \subseteq \mathbb{T}$ .

**Theorem 4.2** (Sound upper bounds). Let  $\mathcal{T}$  be a countable and exhaustive set of interval traces and  $\vdash P : \mathbf{R}$  a program. Then  $[\![P]\!] \leq \text{upperBd}_{P}^{\mathcal{T}}$ .

*Proof sketch.* The formal proof is similar to the soundness proof for the lower bound in Theorem 4.1, but needs an infinite trace semantics [16] for probabilistic programs and is given in Appendix C.1. The idea is that each interval trace *t* summarises all infinite traces starting with (t), i.e. all traces in *cover*(*t*). Exhaustivity ensures that almost all infinite traces are "covered".

#### 4.2 Completeness

The soundness results for upper and lower bounds allow us to derive bounds on the denotation of a program. One would expect that a finer partition of interval traces will yield more precise bounds. In this section, we show that for a program P and an interval  $I \in \mathbb{I}$ , the approximations lowerBd $_{P}^{\mathcal{T}}(I)$  and upperBd $_{P}^{\mathcal{T}}(I)$  can in fact come arbitrarily close to  $[\![P]\!](I)$  for suitable  $\mathcal{T}$ . However, this is only possible under certain assumptions.

Assumption 1: use of sampled values. Interval arithmetic is imprecise if the same value is used more than once: consider, for instance, let s = sample in if (s - s, 0, 1) which deterministically evaluates to 0. However, in interval arithmetic, if x is approximated by an interval [a, b] with a < b, the difference x - x is approximated as [a - b, b - a], which always contains both positive and negative values. So no non-trivial interval trace can separate the two branches.

To avoid this, we could consider a call-by-name semantics (as done in [4]) where sample values can only be used once by definition. However, many of our examples cannot be expressed in the call-by-name setting, so we instead propose a less restrictive criterion to guarantee completeness for call-by-value: we allow sample values to be used more than once, but at most once in the guard of each conditional, at most once in each score expression, and at most once in the return value. While this prohibits terms like the one above, it allows, e.g. let s = sample in if (s, f(s), g(s)). This sufficient condition is formalised in Appendix C.3. Most examples we encountered in the literature satisfy this assumption.

Assumption 2: primitive functions. In addition, we require mild assumptions on the primitive functions, called *boxwise continuity* and *interval separability*.

We need to be able to approximate a program's weight function by step functions in order to obtain tight bounds on its integral. A function  $f : \mathbb{R}^n \to \mathbb{R}$  is **boxwise continuous** if it can be written as the countable union of continuous functions on boxes, i.e. if there is a countable union of pairwise almost disjoint boxes  $B_i$  such that  $\bigcup B_i = \mathbb{R}^n$  and the restriction  $f|_{B_i}$  is continuous for each  $B_i$ .

Furthermore, we need to approximate preimages. Formally, we say that *A* is a *tight subset* of *B* (written  $A \\\in B$ ) if  $A \subseteq B$  and  $B \setminus A$  is a null set. A function  $f : \mathbb{R}^n \to \mathbb{R}$  is called *interval separable* if for every interval  $[a, b] \in \mathbb{I}$ , there is a countable set  $\mathcal{B}$  of boxes in  $\mathbb{R}^n$  that tightly approximates the preimage, i.e.  $\bigcup \mathcal{B} \\\in f^{-1}([a, b])$ . A sufficient condition for checking this is the following. If *f* is boxwise continuous and preimages of points have measure zero, then *f* is already interval separable (cf. Lemma C.4).

We assume the set  $\mathcal{F}$  of primitive functions is closed under composition and each  $f \in \mathcal{F}$  is boxwise continuous and interval separable.

*The completeness theorem.* Using these two assumptions, we can state completeness of our interval semantics.

**Theorem 4.3** (Completeness of interval approximations). Let  $I \in \mathbb{I}$  and  $\vdash P : \mathbf{R}$  be an almost surely terminating program satisfying the two assumptions discussed above. Then, for all  $\epsilon > 0$ , there is a countable set of interval traces  $\mathcal{T} \subseteq \mathbb{T}_{\mathbb{I}}$  that is compatible and exhaustive such that

upperBd<sup>$$\mathcal{T}$$</sup><sub>*p*</sub>(*I*) -  $\epsilon \leq \llbracket P \rrbracket(I) \leq \text{lowerBd}_{p}^{\mathcal{T}}(I) + \epsilon.$ 

*Proof sketch.* We consider each branching path through the program separately. The set of relevant traces for a given path is a preimage of intervals under compositions of interval separable functions, hence can essentially be partitioned into boxes. By boxwise continuity, we can refine this partition such that the weight function is continuous on each box. To approximate the integral, we pass to a refined partition again, essentially computing Riemann sums. The latter converge

to the Riemann integral, which agrees with the Lebesgue integral under our conditions, as desired.  $\hfill \Box$ 

For the lower bound, we can actually derive  $\epsilon\text{-close}$  bounds using only finitely many interval traces:

**Corollary 4.4.** Let  $I \in \mathbb{I}$  and  $\vdash P : \mathbf{R}$  be as in Theorem 4.3. There is a sequence of finite, compatible sets of interval traces  $\mathcal{T}_1, \mathcal{T}_2, \ldots \subseteq \mathbb{T}_{\mathbb{I}}$  s.t.  $\lim_{n\to\infty} \text{lowerBd}_{\mathcal{P}}^{\mathcal{T}_n}(I) = \llbracket P \rrbracket(I)$ .

For the upper bound, a restriction to finite sets  $\mathcal{T}$  of interval traces is, in general, not possible: if the weight function for a program is unbounded, it is also unbounded on some  $t \in \mathcal{T}$ . Then  $\operatorname{wt}_P^{\mathbb{I}}(t)$  is an infinite interval, implying upper  $\operatorname{Bd}_P^{\mathcal{T}}(I) = \infty$  (see Example C.3 for details). Despite the (theoretical) need for countably infinite many interval traces, we can, in many cases, compute finite upper bounds by making use of an interval-based static approximation, formalised as a type system in the next section.

# 5 Weight-aware Interval Type System

To obtain sound bounds on the denotation with only finitely many interval traces, we present an interval-based type system that can derive static bounds on a program. Crucially, our type-system is *weight-aware*: we bound not only the return value of a program but also the weight of an execution. Our analyzer GuBPI uses it for two purposes. First, it allows us to derive upper bounds even for areas of the sample space not covered with interval traces. Second, we can use our analysis to derive a *finite* (and sound) approximation of the infinite number of symbolic execution paths of a program (more details are given in Section 6). Note that the bounds inferred by our system are *interval bounds*, which allow for seamless integration with our interval trace semantics. In this section, we present the interval type system and sketch a constraint-based type inference method.

## 5.1 Interval Types

We define interval types by the following grammar:

$$\sigma := I \mid \sigma \to \mathcal{A} \qquad \mathcal{A} := \begin{cases} \sigma \\ I \end{cases}$$

where  $I \in \mathbb{I}$  is an interval. For readers familiar with refinement types, it is easiest to view the type  $\sigma = I$  as the refinement type  $\{x : \mathbb{R} \mid x \in I\}$ . The definition of the syntactic category  $\mathcal{A}$  by mutual recursion with  $\sigma$  gives a bound on the weight of the execution. We call a type  $\sigma$  weightless and a type  $\mathcal{A}$  weighted. The following examples should give some intuition about the types.

Example 5.1. Consider the example term

 $(\mu_x^{\varphi}.5 \times x \oplus_{0.5} sigm(\varphi x + score sample))(4 \times sample)$ 

where  $sigm : \mathbb{R} \to [0, 1]$  is the sigmoid function. In our type system, this term can be typed with the weighted type  $\left\{ \begin{smallmatrix} [0, 20] \\ [0, 1] \end{smallmatrix} \right\}$ , which indicates that any terminating execution of the term



**Figure 4.** Weight-aware interval type system for SPCF. We abbreviate **1** := [1, 1].

reduces to a value (a number) within [0, 20] and the weight of any such execution lies within [0, 1].

**Example 5.2.** We consider the fixpoint subexpression of the pedestrian example in Example 1.1 which is

 $\mu_x^{\varphi}$ .if $(x, 0, (\lambda step.step + \varphi((x + step) \oplus_{0.5} (x - step)))$ sample).

Using the typing rules (defined below), we can infer the type  $\left\{ \begin{smallmatrix} [a,b] \rightarrow \left\{ \begin{bmatrix} [0,\infty] \\ [1,1] \end{smallmatrix} \right\} \end{smallmatrix} \right\}$  for any *a*, *b*. This type indicates that any terminating execution reduces to a function value (of simple type  $\mathbf{R} \rightarrow \mathbf{R}$ ) with weight within [1, 1]. If this function value is then called on a value within [*a*, *b*], any terminating execution reduces to a value within  $[0,\infty]$  with a weight within [1, 1].

**Subtyping.** The partial order on intervals naturally extends to our type system. For base types  $I_1$  and  $I_2$ , we define  $I_1 \sqsubseteq_{\sigma} I_2$  just if  $I_1 \sqsubseteq I_2$ , where  $\sqsubseteq$  is interval inclusion. We then extend this via:

$$\frac{\sigma_2 \sqsubseteq_{\sigma} \sigma_1 \qquad \mathcal{A}_1 \sqsubseteq_{\mathcal{A}} \mathcal{A}_2}{\sigma_1 \to \mathcal{A}_1 \sqsubseteq_{\sigma} \sigma_2 \to \mathcal{A}_2} \qquad \frac{\sigma_1 \sqsubseteq_{\sigma} \sigma_2 \qquad I_1 \sqsubseteq I_2}{\begin{pmatrix} \sigma_1 \\ I_1 \end{pmatrix} \sqsubseteq_{\mathcal{A}} \begin{pmatrix} \sigma_2 \\ I_2 \end{pmatrix}}$$

Note that in the case of weighted types, the subtyping requires not only that the weightless types be subtype-related  $(\sigma_1 \sqsubseteq_{\sigma} \sigma_2)$  but also that the weight bound be refined  $I_1 \sqsubseteq I_2$ . It is easy to see that both  $\sqsubseteq_{\mathcal{R}}$  and  $\sqsubseteq_{\sigma}$  are partial orders on types with the same underlying base type.

## 5.2 Type System

As for the interval-based semantics, we assume that every primitive operation  $f : \mathbb{R}^n \to \mathbb{R}$  has an overapproximating interval abstraction  $f^{\mathbb{I}} : \mathbb{I}^n \to \mathbb{I}$  (cf. Section 3.1). Interval typing judgments have the form  $\Gamma \vdash M : \mathcal{A}$  where  $\Gamma$  is a typing context mapping variables to types  $\sigma$ . They are given via the rules in Fig. 4. Our system is sound in the following sense (which we here only state for first-order programs).

**Theorem 5.1.** Let  $\vdash P : \mathbf{R}$  be a simply-typed program. If  $\vdash P : \{ \begin{bmatrix} [a,b] \\ [c,d] \end{bmatrix} \text{ and } (P, \mathbf{s}, 1) \rightarrow^* (\underline{r}, \langle \rangle, w) \text{ for some } \mathbf{s} \in \mathbb{T} \text{ and } r, w \in \mathbb{R}, \text{ then } r \in [a,b] \text{ and } w \in [c,d].$ 

Note that the bounds derived by our type system only refer to terminating executions, i.e. they are partial correctness statements. Theorem 5.1 formalises the intuition of an interval type, i.e. every type derivation in our system bounds *both* the returned value (in typical refinement-type fashion [24]) and the weight of this derivation. Our type system also comes with a weak completeness statement: for each term, we can derive some bounds in our system.

**Proposition 5.2.** Let  $\vdash P : \alpha$  be a simply-typed program. There exists a weighted interval type  $\mathcal{A}$  such that  $\vdash P : \mathcal{A}$ .

#### 5.3 Constraint-based Type Inference

In this section, we briefly discuss the automated type *inference* in our system, as needed in our tool GuBPI. For space reasons, we restrict ourselves to an informal overview (see Appendix D for a full account).

Given a program *P*, we can derive the symbolic skeleton of a type derivation (the structure of which is determined by *P*), where each concrete interval is replaced by a placeholder variable. The validity of a typing judgment within this skeleton can then be encoded as constraints. Crucially, as we work in the fixed interval domain and the subtyping structure  $\sqsubseteq_{\mathcal{R}}$ is compositional, they are simple constraints over the placeholder variables in the abstract interval domain. Solving the resulting constraints naïvely might not terminate since the interval abstract domain is not chain-complete. Instead, we approximate the least fixpoint (where the fixpoint denotes a solution to the constraints) using widening, a standard approach to ensure termination of static analysis on domains with infinite chains [13, 14]. This is computationally much cheaper compared to, say, types with general first-order refinements where constraints are typically phrased as constrained Horn clauses (see e.g. [11]). This gain in efficiency is crucial to making our GuBPI tool practical.

# 6 Symbolic Execution and GuBPI

In this section, we describe the overall structure of our tool GuBPI (gubpi-tool.github.io), which builds upon symbolic execution. We also outline how the interval-based semantics can be accelerated for programs containing linear subexpressions.

#### 6.1 Symbolic Execution

The starting point of our analysis is a *symbolic exploration* of the term in question [10, 28, 41]. For space reasons we only give an informal overview of the approach. A detailed and formal discussion can be found in Appendix B.

The idea of symbolic execution is to treat outcomes of sample expressions fully symbolically: each sample evaluates to a fresh variable ( $\alpha_1, \alpha_2, \ldots$ ), called *sample variable*. The result of symbolic execution is thus a symbolic value: a term consisting of sample variables and delayed primitive function applications. We postpone branching decisions and the weighting with score expressions because the value in question is symbolic. During execution, we therefore explore both branches of a conditional and keep track of the (symbolic) conditions on the sample variables that need to hold in the current branch. Similarly, we record the (symbolic) values of score expressions. Formally, our symbolic execution operates on symbolic configurations of the form  $\psi = (\mathcal{M}, n, \Delta, \Xi)$  where  $\mathcal{M}$  is a symbolic term containing sample variables instead of sample outcomes;  $n \in \mathbb{N}$  is a natural number used to obtain fresh sample variables;  $\Delta$  is a list of symbolic constraints of the form  $\mathcal{V} \bowtie r$ , where  $\mathcal{V}$  is a symbolic value,  $r \in \mathbb{R}$  and  $\bowtie \in \{\leq, <, >, \geq\}$ , to keep track of the conditions for the current execution path; and  $\Xi$  is a set of values that records all symbolic values of score expressions encountered along the current path. The symbolic reduction relation  $\rightarrow$  includes the following key rules.

$$\begin{aligned} (\text{sample}, n, \Delta, \Xi) &\rightsquigarrow (\alpha_{n+1}, n+1, \Delta, \Xi) \\ (\text{if}(\mathcal{V}, \mathcal{N}, \mathcal{P})), n, \Delta, \Xi) &\rightsquigarrow (\mathcal{N}, n, \Delta \cup \{\mathcal{V} \leq 0\}, \Xi) \\ (\text{if}(\mathcal{V}, \mathcal{N}, \mathcal{P})), n, \Delta, \Xi) &\rightsquigarrow (\mathcal{P}, n, \Delta \cup \{\mathcal{V} > 0\}, \Xi) \\ (\text{score}(\mathcal{V}), n, \Delta, \Xi) &\rightsquigarrow (\mathcal{V}, n, \Delta \cup \{\mathcal{V} \geq 0\}, \Xi \cup \{\mathcal{V}\}) \end{aligned}$$

That is, we replace sample outcomes with fresh sample variables (first rule), explore both paths of a conditional (second and third rule), and record all score values (fourth rule).

**Example 6.1.** Consider the symbolic execution of Example 1.1 where the first step moves the pedestrian towards their home (taking the right branch of  $\oplus_{0.5}$ ) and the second step moves away from their home (the left branch of  $\oplus_{0.5}$ ). We reach a configuration ( $\mathcal{M}, 5, \Delta, \Xi$ ) where  $\mathcal{M}$  is

score  $\left( pdf_{Normal(1.1,0.1)} \left( \alpha_2 + \alpha_4 + (\mu_x^{\varphi} \cdot \mathcal{N})(3\alpha_1 - \alpha_2 + \alpha_4) \right) \right); 3\alpha_1.$ 

Here  $\alpha_1$  is the initial sample for *start*;  $\alpha_2$ ,  $\alpha_4$  the two samples of *step*; and  $\alpha_3$ ,  $\alpha_5$  the samples involved in the  $\oplus_{0.5}$  operator. The fixpoint  $\mu_x^{\varphi}$ .  $\mathcal{N}$  is already given in Example 5.2,  $\Xi = \emptyset$  and  $\Delta = \{3\alpha_1 > 0, \alpha_3 > \frac{1}{2}, 3\alpha_1 - \alpha_2 > 0, \alpha_5 \le \frac{1}{2}\}$ .

For a symbolic value  $\mathcal{V}$  using sample variables  $\overline{\alpha} = \alpha_1$ , ...,  $\alpha_n$  and  $s \in [0, 1]^n$ , we write  $\mathcal{V}[s/\overline{\alpha}] \in \mathbb{R}$  for the substitution of concrete values in *s* for the sample variables. Call a symbolic configuration of the form  $\Psi = (\mathcal{V}, n, \Delta, \Xi)$  (i.e. a configuration that has reached a symbolic value  $\mathcal{V}$ ) a *symbolic path*. We write *symPaths*( $\psi$ ) for the (countable) set of

Guaranteed Bounds for Posterior Inference in Universal Probabilistic Programming

Algorithm 1 Symbolic Analysis in GuBPI.

1: **Input:** Program  $\vdash P : \mathbf{R}$ , depth limit  $D \in \mathbb{N}$ , and  $I \in \mathbb{I}$ 2:  $\psi_{\text{init}} := (P, 0, \emptyset, \emptyset); S := \{(\psi_{\text{init}}, 0)\}; T := \emptyset$ 3: while  $\exists (\psi, depth) \in S$  do 4: if  $\psi$  has terminated **then**  $T := T \cup \{\psi\}; S := S \setminus \{(\psi, depth)\}$ 5: else if  $\psi$  contains no fixpoints or *depth*  $\leq D$  then 6:  $S := S \setminus \{(\psi, depth)\}$ 7: for  $\psi'$  with  $\psi \rightsquigarrow \psi'$  do 8:  $S := S \cup \{(\psi', depth + 1)\}$ 9: else 10:  $S := (S \setminus \{(\psi, depth)\}) \cup \{(approxFix(\psi), depth)\}$ 11: 12: return  $\left[\sum_{\Psi \in T} \llbracket \Psi \rrbracket_{lb}(I), \sum_{\Psi \in T} \llbracket \Psi \rrbracket_{ub}(I)\right]$ 

symbolic paths reached when evaluating from configuration  $\psi$ . Given a symbolic path  $\Psi = (\mathcal{V}, n, \Delta, \Xi)$  and a set  $U \in \Sigma_{\mathbb{R}}$ , we define the denotation along  $\Psi$ , written  $\llbracket \Psi \rrbracket(U)$ , as

$$\int_{[0,1]^n} \left[ \mathcal{V}[\mathbf{s}/\overline{\alpha}] \in U \right] \prod_{C \bowtie r \in \Delta} \left[ C[\mathbf{s}/\overline{\alpha}] \bowtie r \right] \prod_{\mathcal{W} \in \Xi} \mathcal{W}[\mathbf{s}/\overline{\alpha}] \, \mathrm{d}\mathbf{s},$$

i.e. the integral of the product of the score weights  $\Xi$  over the traces of length *n* where the result value is in *U* and all the constraints  $\Delta$  are satisfied. We can recover the denotation of a program *P* (as defined in Section 2) from all its symbolic paths starting from the configuration (*P*, 0,  $\emptyset$ ,  $\emptyset$ ).

**Theorem 6.1.** Let  $\vdash P : \mathbf{R}$  be a program and  $U \in \Sigma_{\mathbb{R}}$ . Then

$$[P]](U) = \sum_{\Psi \in svmPaths(P,0,\emptyset,\emptyset)} [\![\Psi]\!](U)$$

Analogously to interval SPCF (Section 3), we define *symbolic interval terms* as symbolic terms that may contain intervals (and similarly for symbolic interval values, symbolic interval configurations, and symbolic interval paths).

## 6.2 GuBPI

With symbolic execution at hand, we can outline the structure of our analysis tool GuBPI (sketched in Algorithm 1). GuBPI's analysis begins with symbolic execution of the input term to accumulate a set of symbolic interval paths T. If a symbolic configuration  $\psi$  has exceeded the user-defined depth limit D and still contains a fixpoint, we overapproximate all paths that extend  $\psi$  to ensure a finite set T. We accomplish this by using the interval type system (Section 5) to overapproximate all fixpoint subexpressions, thereby obtaining strongly normalizing terms (in line 11). Formally, given a symbolic configuration  $\psi = (\mathcal{M}, n, \Delta, \Xi)$  we derive a typing judgment for the term  $\mathcal{M}$  in the system from Section 5. Each first-order fixpoint subterm is thus given a (weightless) type of the form  $[a, b] \rightarrow {[c, d] \\ [e, f]}$ . We replace this fixpoint with  $\lambda_{-}(\text{score}([e, f]); [c, d])$ . We denote this operation on configurations by *approxFix*( $\psi$ ) (it extends to higher-order fixpoints as expected). Note that  $approxFix(\psi)$  is a symbolic interval configuration.

**Example 6.2.** Consider the symbolic configuration given in Example 6.1. As in Example 5.2 we infer the type of  $\mu_x^{\varphi}$ .  $\mathcal{N}$  to be  $[-1,4] \rightarrow \{ \begin{bmatrix} 0,\infty \\ 1 \end{bmatrix} \}$ . The function *approxFix* replaces  $\mu_x^{\varphi}$ .  $\mathcal{N}$  with  $\lambda$ \_.score([1,1]);  $[0,\infty]$ . By evaluating the resulting symbolic interval configuration further, we obtain the symbolic interval path  $(3\alpha_1, 5, \Delta, \Xi)$  where  $\Delta$  is as in Example 6.1 and  $\Xi = \{ pdf_{Normal(1.1,0.1)}(\alpha_2 + \alpha_4 + [0,\infty]) \}$ . Note that, in general, the further evaluation of *approxFix*( $\psi$ ) can result in multiple symbolic interval paths.

Afterwards, we're left with a finite set *T* of symbolic interval paths. Due to the presence of intervals, we cannot define a denotation of such paths directly and instead define lower and upper bounds. For a symbolic interval value  $\mathcal{V}$  that contains *no* sample variables, we define  $\lceil \mathcal{V} \rceil \subseteq \mathbb{R}$  as the set of all values that the term can evaluate to by replacing every interval [a, b] with some value  $r \in [a, b]$ . Given a symbolic interval path  $\Psi = (\mathcal{V}, n, \Delta, \Xi)$  and  $U \in \Sigma_{\mathbb{R}}$  we define  $\llbracket \Psi \rrbracket_{lb}(U)$  by considering only those concrete traces that fulfill the constraints in  $\Psi$  for *all* concrete values in the intervals and take the infimum over all scoring expressions:

$$\int \left[ \ulcorner \mathcal{V}[s/\overline{\alpha}] \urcorner \subseteq U \right] \prod_{C \bowtie r \in \Delta} \left[ \forall t \in \ulcorner C[s/\overline{\alpha}] \urcorner . t \bowtie r \right] \prod_{\mathcal{W} \in \Xi} \inf \ulcorner \mathcal{W}[s/\overline{\alpha}] \urcorner \mathrm{d}s.$$

Similarly, we define 
$$[\![\Psi]\!]_{ub}(U)$$
 as

$$\int \left[ \ulcorner \mathcal{V}[s/\overline{\alpha}] \urcorner \cap U \neq \emptyset \right] \prod_{C \bowtie r \in \Delta} \left[ \exists t \in \ulcorner C[s/\overline{\alpha}] \urcorner . t \bowtie r \right] \prod_{\mathcal{W} \in \Xi} \sup \ulcorner \mathcal{W}[s/\overline{\alpha}] \urcorner \mathrm{d}s.$$

We note that, if  $\Psi$  contains no intervals,  $\llbracket \Psi \rrbracket$  is defined and we have  $\llbracket \Psi \rrbracket_{lb} = \llbracket \Psi \rrbracket_{ub} = \llbracket \Psi \rrbracket$ . We can now state the following theorem that formalises the observation that *approxFix*( $\psi$ ) soundly approximates all symbolic paths that result from  $\psi$ .

**Theorem 6.2.** Let  $\psi$  be a symbolic (interval-free) configuration and  $U \in \Sigma_{\mathbb{R}}$ . Define  $A = symPaths(\psi)$  as the (possibly infinite) set of all symbolic paths reached when evaluating  $\psi$ and  $B = symPaths(approxFix(\psi))$  as the (finite) set of symbolic interval paths reached when evaluating approxFix( $\psi$ ). Then

$$\sum_{\Psi \in B} \llbracket \Psi \rrbracket_{lb}(U) \le \sum_{\Psi \in A} \llbracket \Psi \rrbracket(U) \le \sum_{\Psi \in B} \llbracket \Psi \rrbracket_{ub}(U).$$

The correctness of Algorithm 1 is then a direct consequence of Theorems 6.1 and 6.2:

**Corollary 6.3.** Let T be the set of symbolic interval paths computed when at line 12 of Algorithm 1 and  $U \in \Sigma_{\mathbb{R}}$ . Then

$$\sum_{\Psi \in T} \llbracket \Psi \rrbracket_{lb}(U) \le \llbracket P \rrbracket(U) \le \sum_{\Psi \in T} \llbracket \Psi \rrbracket_{ub}(U).$$

What remains is to compute (lower bounds on)  $\llbracket \Psi \rrbracket_{lb}(I)$ and (upper bounds on)  $\llbracket \Psi \rrbracket_{ub}(I)$  for a symbolic interval path  $\Psi \in T$  and  $I \in \mathbb{I}$ . We first present the standard interval trace semantics (Section 6.3) and then a more efficient analysis for the case that  $\Psi$  contains only linear functions (Section 6.4).

### 6.3 Standard Interval Trace Semantics

For any symbolic interval path  $\Psi$ , we can employ the semantics as introduced in Section 3. Instead of applying the

analysis to the entire program, we can restrict to the current path  $\Psi$  (intuitively, by adding a score(0) to all other program paths). The interval traces split the domain of each sample variable in  $\Psi$  into intervals. It is easy to see that for any compatible and exhaustive set of interval traces  $\mathcal{T}$ , we have lowerBd $_{\Psi}^{\mathcal{T}}(U) \leq \llbracket \Psi \rrbracket_{lb}(U)$  and  $\llbracket \Psi \rrbracket_{ub}(U) \leq \text{upperBd}_{\Psi}^{\mathcal{T}}(U)$ for any  $U \in \Sigma_{\mathbb{R}}$  (see Theorem 4.1 and 4.2). Applying the interval-based semantics on the level of symbolic interval paths maintains the attractive features, namely soundness and completeness (relative to the current path) of the semantics. Note that the intervals occurring in  $\Psi$  seamlessly integrate with our interval-based semantics.

#### 6.4 Linear Interval Trace Semantics

In case the score values and the guards of all conditionals are linear, we can improve and speed up the interval-based semantics.

Assume all symbolic interval values appearing in  $\Psi$  are interval linear functions of  $\overline{\alpha}$  (i.e. functions  $\overline{\alpha} \mapsto \mathbf{w}^{\intercal} \overline{\alpha} + {}^{\mathbb{I}}[a, b]$  for some  $\mathbf{w} \in \mathbb{R}^n$  and  $[a, b] \in \mathbb{I}$ ). We assume, for now, that each symbolic value  $\mathcal{W} \in \Xi$  denotes an interval-free linear function (i.e. a function  $\overline{\alpha} \mapsto \mathbf{w}^{\intercal} \overline{\alpha} + r$ ). Fix some interval  $I \in \mathbb{I}$ . We first note that both  $\llbracket \Psi \rrbracket_{ub}(I)$  and  $\llbracket \Psi \rrbracket_{ub}(I)$  are the integral of a polynomial over a convex polytope: define

$$\mathfrak{P}_{lb} \coloneqq \left\{ s \in \mathbb{R}^n \mid \ulcorner \mathcal{V}[s/\overline{\alpha}] \urcorner \subseteq I \land \bigwedge_{C \bowtie r \in \Delta} \forall t \in \ulcorner C[s/\overline{\alpha}] \urcorner .t \bowtie r \right\}$$

which is a polytope.<sup>10</sup> Then  $\llbracket \Psi \rrbracket_{lb}(I)$  is the integral of the polynomial  $\overline{\alpha} \mapsto \prod_{W \in \Xi} W$  over  $\mathfrak{P}_{lb}$ . The definition of  $\mathfrak{P}_{ub}$  (as the region of integration for  $\llbracket \Psi \rrbracket_{ub}(I)$ ) is similar. Such integrals can be computed exactly [2], e.g. with the LattE tool [20]. Unfortunately, our experiments showed that this does not scale to interesting probabilistic programs.

Instead, we derive guaranteed bounds on the denotation by means of iterated volume computations. This has the additional benefit that we can handle non-uniform samples and non-liner expressions in  $\Xi$ . We follow an approach similar to that of the interval-based semantics in Section 4 but do not split/bound individual sample variables and instead directly bound linear functions over the sample variables. Let  $\Xi = \{W_1, \dots, W_k\}$ . We define a **box** (by abuse of language) as an element  $t = \langle [a_1, b_1], \dots, [a_k, b_k] \rangle$ , where  $[a_i, b_i]$  gives a bound on  $\mathcal{W}_i$ .<sup>11</sup> We define  $lb(t) := \prod_{i=1}^k a_i$ and  $ub(t) := \prod_{i=1}^{k} b_i$ . The box *t* naturally defines a subset of  $\mathfrak{P}_{lb}$  given by  $\mathfrak{P}_{lb}^t = \{ s \in \mathfrak{P}_{lb} \mid \bigwedge_{i=1}^k \mathcal{W}_i[s/\overline{\alpha}] \in [a_i, b_i] \}.$ Then  $\mathfrak{P}_{lb}^t$  is again a polytope and we write  $\operatorname{vol}(\mathfrak{P}_{lb}^t)$  for its volume. The definition of  $\mathfrak{P}_{ub}^t$  and  $\operatorname{vol}(\mathfrak{P}_{ub}^t)$  is analogous. As for interval traces, we call two boxes  $t_1$ ,  $t_2$  compatible if the intervals are almost disjoint in at least one position. A set

entry of a box bounds the *i*th score value.

of boxes *B* is *compatible* if its elements are pairwise compatible and *exhaustive* if  $\bigcup_{t \in B} \mathfrak{P}_{lb}^t = \mathfrak{P}_{lb}$  and  $\bigcup_{t \in B} \mathfrak{P}_{ub}^t = \mathfrak{P}_{ub}$ (cf. Section 3.3).

**Proposition 6.4.** Let B be a compatible and exhaustive set of boxes. Then  $\sum_{t \in B} lb(t) \cdot vol\left(\mathfrak{P}_{lb}^{t}\right) \leq \llbracket \Psi \rrbracket_{lb}(I)$  and  $\llbracket \Psi \rrbracket_{ub}(I) \leq \sum_{t \in B} ub(t) \cdot vol\left(\mathfrak{P}_{ub}^{t}\right)$ .

As in the standard interval semantics, a finer partition into boxes yields more precise bounds. While the volume computation involved in Proposition 6.4 is expensive [22], the number of splits on the linear functions is much smaller than that needed in the standard interval-based semantics. Our experiments empirically demonstrate that the direct splitting of linear functions (if applicable) is usually superior to the standard splitting. In GuBPI we compute a set of exhaustive boxes by first computing a lower and upper bound on each  $W_i \in \Xi$  over  $\mathfrak{P}_{lb}$  (or  $\mathfrak{P}_{ub}$ ) by solving a linear program (LP) and splitting the resulting range in evenly sized chunks.

**Beyond uniform samples and linear scores.** We can extend our linear optimization to non-uniform samples and arbitrary symbolic values in  $\Xi$ . We accomplish the former by *combining* the optimised semantics (where we bound linear expressions) with the standard interval-trace semantics (where we bound individual sample variables). For the latter, we can identify linear sub-expressions of the expressions in  $\Xi$ , use boxes to bound each linear sub-expression, and use interval arithmetic to infer bounds on the entire expression from bounds on its linear sub-expressions. More details can be found in Appendix E.1.

**Example 6.3.** Consider the path  $\Psi = (3\alpha_1, 5, \Delta, \Xi)$  derived in Example 6.2. We use 1-dimensional boxes to bound  $\alpha_2 + \alpha_4$  (the single linear sub-expression of the symbolic values in  $\Xi$ ). To obtain a lower bound on  $\llbracket \Psi \rrbracket_{lb}(I)$ , we sum over all boxes  $t = \langle [a_1, b_1] \rangle$  and take the product of vol  $(\mathfrak{P}_{lb}^t)$ with the lower interval bound of pdf<sub>Normal(1.1,0.1)</sub> ( $[a_1, b_1] + [0, \infty]$ ) (evaluated in interval arithmetic). Analogously, for the upper bound we take the product of vol  $(\mathfrak{P}_{ub}^t)$  with the upper interval bound of pdf<sub>Normal(1.1,0.1)</sub> ( $[a_1, b_1] + [0, \infty]$ ).

# 7 Practical Evaluation

We have implemented our semantics in the prototype GuBPI (gubpi-tool.github.io), written in F#. In cases where we apply the linear optimisation of our semantics, we use Vinci [8] to discharge volume computations of convex polytopes. We set out to answer the following questions:

- 1. How does GuBPI perform on instances that could already be solved (e.g. by PSI [26])?
- 2. Is GuBPI able to infer useful bounds on recursive programs that could not be handled rigorously before?

#### 7.1 Probability Estimation

We collected a suite of 18 benchmarks from [56]. Each benchmark consists of a program P and a query  $\phi$  over the variables

<sup>&</sup>lt;sup>10</sup>For example, if *C* denotes the function  $\overline{\alpha} \mapsto \mathbf{w}^{\mathsf{T}} \overline{\alpha} + [a, b]$  we can transform a constraint  $\forall t \in {}^{\mathsf{T}} C[s/\overline{\alpha}]^{\mathsf{T}}$ . *t*  $\leq r$  into the linear constraint  $\mathbf{w}^{\mathsf{T}} \overline{\alpha} + b \leq r$ . <sup>11</sup>Note the similarity to the interval trace semantics. While the *i*th position in an interval trace bounds the value of the *i*th sample variable, the *i*th

**Table 1.** Evaluation on selected benchmarks from [56]. We give the times (in seconds) and bounds computed by [56] and GuBPI. Details on the exact queries (the Q column) can be found in Table 4 in the appendix.

		Tool from [56]			GuBPI
Program	Q	t	Result	t	Result
tug-of-war	Q1	1.29	[0.6126, 0.6227]	0.72	[0.6134, 0.6135]
tug-of-war	Q2	1.09	[0.5973, 0.6266]	0.79	[0.6134, 0.6135]
beauquier-3	Q1	1.15	[0.5000, 0.5261]	22.5	[0.4999, 0.5001]
ex-book-s	Q1	8.48	[0.6633, 0.7234]	6.52	[0.7417, 0.7418]
ex-book-s	Q2*	10.3	[0.3365, 0.3848]	8.01	[0.4137, 0.4138]
ex-cart	Q1	2.41	[0.8980, 1.1573]	67.3	[0.9999, 1.0001]
ex-cart	Q2	2.40	[0.8897, 1.1573]	68.5	[0.9999, 1.0001]
ex-cart	Q3	0.15	[0.0000, 0.1150]	67.4	[0.0000, 0.0001]
ex-ckd-epi-s	Q1*	0.15	[0.5515, 0.5632]	0.86	[0.0003, 0.0004]
ex-ckd-epi-s	Q2*	0.08	[0.3019, 0.3149]	0.84	[0.0003, 0.0004]
ex-fig6	Q1	1.31	[0.1619, 0.7956]	21.2	[0.1899, 0.1903]
ex-fig6	Q2	1.80	[0.2916, 1.0571]	21.4	[0.3705, 0.3720]
ex-fig6	Q3	1.51	[0.4314, 2.0155]	24.7	[0.7438, 0.7668]
ex-fig6	Q4	3.96	[0.4400, 3.0956]	27.4	[0.8682, 0.9666]
ex-fig7	Q1	0.04	[0.9921, 1.0000]	0.18	[0.9980, 0.9981]
example4	Q1	0.02	[0.1910, 0.1966]	0.31	[0.1918, 0.1919]
example5	Q1	0.06	[0.4478, 0.4708]	0.27	[0.4540, 0.4541]
herman-3	Q1	0.47	[0.3750, 0.4091]	124	[0.3749, 0.3751]

of *P*. We bound the probability of the event described by  $\phi$ using the tool from [56] and GuBPI (Table 1). While our tool is generally slower than the one in [56], the completion times are still reasonable. Moreover, in many cases, the bounds returned by GuBPI are tighter than those of [56]. In addition, for benchmarks marked with a  $\star$ , the two pairs of bounds contradict each other.<sup>12</sup> We should also remark that GuBPI cannot handle all benchmarks proposed in [56] because the heavy use of conditionals causes our precise symbolic analysis to suffer from the well-documented path explosion problem [6, 9, 30]. Perhaps unsurprisingly, [56] can handle such examples much better, as one of their core contributions is a stochastic method to reduce the number of paths considered (see Section 8). Also note that [56] is restricted to uniform samples, linear guards and score-free programs, whereas we tackle a much more general problem.

### 7.2 Exact Inference

To evaluate our tool on instances that can be solved exactly, we compared it with PSI [26, 27], a symbolic solver which can, in certain cases, compute a closed-form solution of the posterior. We note that whenever exact inference is possible, exact solutions will always be superior to mere bounds and, due to the overhead of our semantics, will often be

**Table 2.** Probabilistic programs with discrete domains from

 PSI [26]. The times for PSI and GuBPI are given in seconds.

Instance	t <sub>PSI</sub>	<b>t</b> <sub>GuBPI</sub>	Instance	t <sub>PSI</sub>	t <sub>GuBPI</sub>
burglarAlarm	0.06	0.21	coins	0.04	0.18
twoCoins	0.04	0.21	ev-model1	0.04	0.21
grass	0.06	0.37	ev-model2	0.04	0.20
noisyOr	0.14	0.72	murderMystery	0.04	0.19
bertrand	0.04	0.22	coinBiasSmall	0.13	1.92
coinPattern	0.04	0.19	gossip	0.08	0.24



(a) coinBias example from [26]. (b) m The program samples a beta prog prior on the bias of a coin and of tw observes repeated coin flips (26 second





(b) max example from [26]. The program compute the maximum of two i.i.d. normal samples (31 seconds).



(c) Binary Gaussian Mixture Model from [65] (39 seconds). MCMC methods usually find only one mode.

(d) Neal's funnel from [34, 48] (2.8 seconds). HMC misses some probability mass around 0.

**Figure 5.** Guaranteed Bounds computed by GuBPI for a selection of non-recursive models from [26, 27, 48, 65].

found faster. Because of the different output formats (i.e. exact results vs. bounds), a direct comparison between exact methods and GuBPI is challenging. As a consistency check, we collected benchmarks from the PSI repository where the output domain is finite and GuBPI can therefore compute *exact* results (tight bounds). They agree with PSI in all cases, which includes 8 of the 21 benchmarks from [26]. We report the computation times in Table 2.

We then considered examples where GuBPI computes non-tight bounds. For space reasons, we can only include a selection of examples in this paper. The bounds computed by GuBPI and a short description of each example are shown in Fig. 5. We can see that, despite the relatively loose bounds, they are still useful and provide the user with a rough and most importantly, *guaranteed to be correct*—idea of the denotation.

<sup>&</sup>lt;sup>12</sup>A stochastic simulation using 10<sup>6</sup> samples in Anglican [61] yielded results that fall within GuBPI's bounds but violate those computed by [56].



(a) cav-example-7. Program taken from the PSI repository. PSI bounds the depth resulting in a spike at 10, whereas GuBPI can compute bounds on the denotation of the unbounded program (112 seconds).



(d) random-box-walk. The program models the cumulative distance traveled by a biased random walk. If a uniformly sampled step *s* has length less than  $\frac{1}{2}$ , we move *s* to the left, otherwise *s* to the right. The walk stops when it crosses a threshold (167 seconds).



(b) cav-example-5. Program taken from the PSI repository. PSI cannot handle this program due to the unbounded loops (192 seconds).



(e) growing-walk. The program models a geometric random walk where (with increasing distance) the step size of the walk is increased. The cumulative distance is observed from a normal distribution centered at 3 (67 seconds).



(c) add\_uniform\_with\_counter\_large. Program taken from the PSI repository. GuBPI can handle the unbounded loop, whereas PSI unrolls to a fixed depth (21 seconds).



(f) param-estimation-recursive. We sample a uniform prior p and (in each step) travel to the left with probability p and to the right with probability (1 - p). We observe the walk to come to a halt at location 1 (observed from a normal) and wish to find the posterior on p (162 seconds).

Figure 6. Guaranteed bounds computed by GuBPI for a selection of recursive models.

The success of exact solvers such as PSI depends on the underlying symbolic solver (and the optimisations implemented). Consequently, there are instances where the symbolic solver cannot compute a closed-form (integral-free) solution. Conversely, while our method is (theoretically) applicable to a very broad class of programs, there exist programs where the symbolic solver finds solutions but the analysis in GuBPI becomes infeasible due to the large number of interval traces required.

# 7.3 Recursive Models

We also evaluated our tool on complex models that *cannot* be handled by any of the existing methods. For space reasons, we only give an overview of some examples. Unexpectedly, we found recursive models in the PSI repository: there are examples that are created by unrolling loops to a fixed depth. This fixed unrolling changes the posterior of the model. Using our method we can handle those examples *without* bounding the loop. Three such examples are shown in Figs. 6a to 6c. In Fig. 6a, PSI bounds the iterations resulting in a spike at 10 (the unrolling bound). For Fig. 6b, PSI does not provide any solution whereas GuBPI provides useful bounds. For Fig. 6c, PSI bounds the loop to compute results (displayed

in blue) whereas GuBPI computes the green bounds on the unbounded program. It is obvious that the bounds differ significantly, highlighting the impact that unrolling to a fixed depth can have on the denotation. This again strengthens the claim that rigorous methods that can handle unbounded loops/recursion are needed. There also exist unbounded discrete examples where PSI computes results for the bounded version that differ from the denotation of the unbounded program. Figs. 6d to 6f depict further recursive examples (alongside a small description).

Lastly, as a *very* challenging example, we consider the pedestrian example (Example 1.1) again. The bounds computed by GuBPI are given in Fig. 7 together with the two stochastic results from Fig. 1. The bounds are clearly precise enough to rule out the HMC samples. Since this example has infinite expected running time, it is very challenging and GuBPI takes about 1.5h (84 min).<sup>13</sup> In fact, guaranteed bounds are the only method to recognise the wrong samples with certainty (see the next section for statistical methods).

<sup>&</sup>lt;sup>13</sup>While the running time seems high, we note that Pyro HMC took about an hour to generate 10<sup>4</sup> samples and produce the (wrong) histogram. Diagnostic methods like simulation-based calibration took even longer (>30h) and delivered inconclusive results (see Section 7.4 for details).



Figure 7. Bounds for the pedestrian example (Example 1.1).

**Table 3.** Running times of GuBPI and SBC for (Pyro's) HMC. Times are given in seconds (s) and hours (h).

Instance	<b>t</b> <sub>GuBPI</sub>	t <sub>SBC</sub>
Binary GMM (1-dimensional) (Fig. 5c)	39s	1h
Binary GMM (2-dimensional)	4h	1.5h
Pedestrian Example (Fig. 7)	1.5h	>300h

#### 7.4 Comparison with Statistical Validation

A general approach to validating inference algorithms for a generative Bayesian model is simulation-based calibration (SBC) [12, 60]. SBC draws a sample  $\theta$  from the prior distribution of the parameters, generates data y for these parameters, and runs the inference algorithm to produce posterior samples  $\theta_1, \ldots, \theta_L$  given y. If the posterior samples follow the true posterior distribution, the rank statistic of the prior sample  $\theta$  relative to the posterior samples will be uniformly distributed. If the empirical distribution of the rank statistic after many such simulations is non-uniform, this indicates a problem with the inference. While SBC is very general, it is computationally expensive because it performs inference in every simulation. Moreover, as SBC is a stochastic validation approach, any fixed number of samples may fail to diagnose inference errors that only occur on a very low probability region.

We compare the running times of GuBPI and SBC for three examples where Pyro's HMC yields wrong results (Table 3). Running SBC on the pedestrian example (with a reduced sample size and using the parameters recommended in [60]) took 32 hours and was still inconclusive because of strong autocorrelation. Reducing the latter via thinning requires more samples, and would increase the running time to >300 hours. Similarly, GuBPI diagnoses the problem with the mixture model in Fig. 5c in significantly less time than SBC. However, for higher-dimensional versions of this mixture model, SBC clearly outperforms GuBPI. We give a more detailed discussion of SBC for these examples in Appendix F.3.

#### 7.5 Limitations and Future Improvements

The theoretical foundations of our interval-based semantics ensure that GuBPI is applicable to a very broad class of programs (cf. Section 4). In practice, as usual for exact methods, GuBPI does not handle all examples equally well.

Firstly, as we already saw in Section 7.1, the symbolic execution—which forms the entry point of the analysis suffers from path explosion. On some extreme loop/recursionfree programs (such as example-ckd-epi from [56]), our tool cannot compute all (finitely many) symbolic paths in reasonable time, let alone analyse them in our semantics. Extending the approach from [56], to sample representative program paths (in the presence of conditioning), is an interesting future direction that we can combine with the rigorous analysis provided by our interval type system.

Secondly, our interval-based semantics imposes bounds on each sampled variable and thus scales exponentially with the dimension of the model; this is amplified in the case where the optimised semantics (Section 6.4) is not applicable. It would be interesting to explore whether this can be alleviated using different trace splitting techniques.

Lastly, the bounds inferred by our interval type system take the form of a single interval with no information about the exact distribution on that interval. For example, the most precise bound derivable for the term  $\mu_x^{\varphi}.x \oplus [\varphi(x+\text{sample}) \oplus \varphi(x-\text{sample})]$  is  $[a,b] \rightarrow \{ [-\infty,\infty] \\ [1,1] \}$  for any a, b. After unrolling to a fixed depth, the approximation of the paths not terminating within the fixed depth is therefore imprecise. For future work, it would be interesting to improve the bounds in our type system to provide more information about the distribution by means of rigorous approximations of the denotation of the fixpoint in question (i.e. a probabilistic summary of the fixpoint [46, 50, 63]).

# 8 Related Work

Interval trace semantics and Interval SPCF. Our interval trace semantics to compute bounds on the denotation is similar to the semantics introduced by Beutner and Ong [4], who study an interval approximation to obtain *lower* bounds on the termination probability. By contrast, we study the more challenging problem of bounding the program denotation which requires us to track the weight of an execution, and to prove that the denotation approximates a Lebesgue integral, which requires novel proof ideas. Moreover, whereas the termination probability of a program is always upper bounded by 1, here we derive both lower and *upper* bounds.

**Probability estimation.** Sankaranarayanan et al. [56] introduced a static analysis framework to infer bounds on a class of definable events in (*score-free*) probabilistic programs. The idea of their approach is that if we find a finite set  $\mathcal{T}$  of symbolic traces with cumulative probability at least 1 - c, and a given event  $\phi$  occurs with probability at most b on the

traces in  $\mathcal{T}$ , then  $\phi$  occurs with probability at most b + c on the entire program. In the presence of conditioning, the problem becomes vastly more difficult, as the aggregate weight on the unexplored paths can be unbounded, giving  $\infty$  as the only derivable upper bound (and therefore also  $\infty$  as the best upper bound on the normalising constant). In order to infer guaranteed bounds, it is necessary to analyse all paths in the program, which we accomplish via static analysis and in particular our interval type system. The approach from [56] was extended by Albarghouthi et al. [1] to compute the probability of events defined by arbitrary SMT constraints but is restricted to score-free and non-recursive programs. Our interval-based approach, which may be seen as a variant of theirs, is founded on a complete semantics (Theorem 4.3), can handle recursive program with (soft) scoring, and is applicable to a broad class of primitive functions.

Note that we consider programs with *soft* conditioning in which scoring cannot be reduced to volume computation directly.<sup>14</sup> Intuitively, soft conditioning performs a (global) re-weighting of the set of traces, which cannot be captured by (local) volume computations. In our interval trace semantics, we instead track an approximation of the weight along each interval trace.

Exact inference. There are numerous approaches to inferring the exact denotation of a probabilistic program. Holtzen et al. [38] introduced an inference method to efficiently compute the denotation of programs with discrete distributions. By exploiting program structure to factorise inference, their system Dice can perform exact inference on programs with hundreds of thousands of random variables. Gehr et al. [26] introduced PSI, an exact inference system that uses symbolic manipulation and integration. A later extension,  $\lambda PSI$ [27], adds support for higher-order functions and nested inference. The PPL Hakaru [47] supports a variety of inference algorithms on programs with both discrete and continuous distributions. Using program transformation and partial evaluation, Hakaru can perform exact inference via symbolic disintegration [57] on a limited class of programs. Saad et al. [55] introduced SPPL, a system that can compute exact answers to a range of probabilistic inference queries, by translating a restricted class of programs to sum-product expressions, which are highly effective representations for inference.

While exact results are obviously desirable, this kind of inference only works for a restricted family of programs: none of the above exact inference systems allow (unbounded) recursion. Unlike our tool, they are therefore unable to handle, for instance, the challenging Example 1.1 or the programs in Fig. 6. Abstract interpretation. There are various approaches to probabilistic abstract interpretation, so we can only discuss a selection. Monniaux [44, 45] developed an abstract domain for (score-free) probabilistic programs given by a weighted sum of abstract regions. Smith [58] considered truncated normal distributions as an abstract domain and developed analyses restricted to score-free programs with only linear expressions. Extending both approaches to support soft conditioning is non-trivial as it requires the computation of integrals on the abstract regions. In our interval-based semantics, we abstract the concrete traces (by means of interval traces) and not the denotation. This allows us to derive bounds on the weight along the abstracted paths.

Huang et al. [39] discretise the domain of continuous samples into interval cubes and derive posterior approximations on each cube. The resulting approximation converges to the true posterior (similarly to approximate/stochastic methods) but does not provide exact/guaranteed bounds and is not applicable to recursive programs.

**Refinement types.** Our interval type system (Section 5) may be viewed as a type system that refines not just the value of an expression but also its weight [24]. To our knowledge, no existing type refinement system can bound the weight of program executions. Moreover, the seamless integration with our interval trace semantics by design allows for much cheaper type inference, without resorting to an SMT or Horn constraint solver. This is a crucial advantage since a typical GuBPI execution queries the analysis numerous times.

**Stochastic methods.** A general approach to validating inference algorithms for a generative Bayesian model is *simulation-based calibration* (SBC) [12, 60], discussed in Section 7.4. Grosse et al. [36] introduced a method to estimate the log marginal likelihood of a model by constructing stochastic lower/upper bounds. They show that the true value can be sandwiched between these two stochastic bounds with high probability. In closely related work [17, 35], this was applied to measure the accuracy of approximate probabilistic inference algorithms on a specified dataset. By contrast, our bounds are non-stochastic and our method is applicable to arbitrary programs of a universal PPL.

# 9 Conclusion

We have studied the problem of inferring guaranteed bounds on the posterior of programs written in a universal PPL. Our work is based on the interval trace semantics, and our weightaware interval type system gives rise to a tool that can infer useful bounds on the posterior of interesting recursive programs. This is a capability beyond the reach of existing methods, such as exact inference. As a method of Bayesian inference for statistical probabilistic programs, we can view our framework as occupying a useful middle ground between approximate stochastic inference and exact inference.

<sup>&</sup>lt;sup>14</sup>For programs including only hard-conditioning (i.e. scoring is only possible with 0 or 1), the posterior probability of an event  $\varphi$  can be computed by dividing the probability of all traces with weight 1 on which  $\varphi$  holds by the probability of all traces with weight 1.

Guaranteed Bounds for Posterior Inference in Universal Probabilistic Programming

# References

- Aws Albarghouthi, Loris D'Antoni, Samuel Drews, and Aditya V. Nori. 2017. FairSquare: probabilistic verification of program fairness. *Proceedings of the ACM on Programming Languages* 1, OOPSLA (2017). https://doi.org/10.1145/3133904
- [2] Velleda Baldoni, Nicole Berline, Jesus De Loera, Matthias Köppe, and Michèle Vergne. 2011. How to integrate a polynomial over a simplex. *Math. Comp.* 80, 273 (2011). https://doi.org/10.1090/S0025-5718-2010-02378-6
- [3] Atilim Günes Baydin, Lei Shao, Wahid Bhimji, Lukas Alexander Heinrich, Lawrence Meadows, Jialin Liu, Andreas Munk, Saeid Naderiparizi, Bradley Gram-Hansen, Gilles Louppe, Mingfei Ma, Xiaohui Zhao, Philip H. S. Torr, Victor W. Lee, Kyle Cranmer, Prabhat, and Frank Wood. 2019. Etalumis: bringing probabilistic programming to scientific simulators at scale. In *International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2019.* ACM. https://doi.org/10.1145/3295500.3356180
- [4] Raven Beutner and Luke Ong. 2021. On probabilistic termination of functional programs with continuous distributions. In ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2021. ACM. https://doi.org/10.1145/3453483.3454111
- [5] Eli Bingham, Jonathan P. Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul A. Szerlip, Paul Horsfall, and Noah D. Goodman. 2019. Pyro: Deep Universal Probabilistic Programming. *Journal of Machine Learning Research* 20 (2019). http://jmlr.org/papers/v20/18-403.html
- [6] Peter Boonstoppel, Cristian Cadar, and Dawson R. Engler. 2008. RWset: Attacking Path Explosion in Constraint-Based Test Generation. In International Conference on Tools and Algorithms for the Construction and Analysis of Systems, TACAS 2008. Springer. https://doi.org/10. 1007/978-3-540-78800-3\_27
- [7] Johannes Borgström, Ugo Dal Lago, Andrew D. Gordon, and Marcin Szymczak. 2016. A lambda-calculus foundation for universal probabilistic programming. In ACM SIGPLAN International Conference on Functional Programming, ICFP 2016. ACM. https://doi.org/10.1145/ 2951913.2951942
- [8] Benno Büeler, Andreas Enge, and Komei Fukuda. 2000. Exact Volume Computation for Polytopes: A Practical Study. In *Polytopes combinatorics and computation*. DMV Sem., Vol. 29. Birkhäuser, Basel. https://doi.org/10.1007/978-3-0348-8438-9\_6
- [9] Cristian Cadar, Vijay Ganesh, Peter M. Pawlowski, David L. Dill, and Dawson R. Engler. 2008. EXE: Automatically Generating Inputs of Death. ACM Transactions on Information and System Security 12, 2 (2008). https://doi.org/10.1145/1455518.1455522
- [10] Arun Tejasvi Chaganty, Aditya V. Nori, and Sriram K. Rajamani. 2013. Efficiently Sampling Probabilistic Programs via Program Analysis. In International Conference on Artificial Intelligence and Statistics, AISTATS 2013 (JMLR, Vol. 31). http://proceedings.mlr.press/v31/chaganty13a. html
- [11] Adrien Champion, Tomoya Chiba, Naoki Kobayashi, and Ryosuke Sato. 2020. ICE-Based Refinement Type Discovery for Higher-Order Functional Programs. *Journal of Automated Reasoning* 64, 7 (2020). https://doi.org/10.1007/s10817-020-09571-y
- [12] Samantha R Cook, Andrew Gelman, and Donald B Rubin. 2006. Validation of software for Bayesian models using posterior quantiles. *Journal of Computational and Graphical Statistics* 15, 3 (2006). https: //doi.org/10.1198/106186006X136976
- [13] Patrick Cousot and Radhia Cousot. 1976. Static determination of dynamic properties of programs. In *International Symposium on Programming*. Dunod.
- [14] Patrick Cousot and Radhia Cousot. 1977. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In ACM Symposium on Principles of Programming Languages, POPL 1977. ACM. https://doi.org/10.1145/ 512950.512973

- [15] Patrick Cousot and Radhia Cousot. 2002. Modular Static Program Analysis. In International Conference on Compiler Construction, CC 2002 (LNCS, Vol. 2304). Springer. https://doi.org/10.1007/3-540-45937-5\_13
- [16] Ryan Culpepper and Andrew Cobb. 2017. Contextual Equivalence for Probabilistic Programs with Continuous Random Variables and Scoring. In European Symposium on Programming, ESOP 2017 (LNCS, Vol. 10201). Springer. https://doi.org/10.1007/978-3-662-54434-1\_14
- [17] Marco F. Cusumano-Towner and Vikash K. Mansinghka. 2017. AIDE: An algorithm for measuring the accuracy of probabilistic inference algorithms. In Annual Conference on Neural Information Processing Systems, NIPS 2017. https://proceedings.neurips.cc/paper/2017/hash/ acab0116c354964a558e65bdd07ff047-Abstract.html
- [18] Marco F. Cusumano-Towner, Feras A. Saad, Alexander K. Lew, and Vikash K. Mansinghka. 2019. Gen: a general-purpose probabilistic programming system with programmable inference. In ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2019. ACM. https://doi.org/10.1145/3314221.3314642
- [19] Hend Dawood. 2011. Theories of Interval Arithmetic: Mathematical Foundations and Applications. LAP Lambert Academic Publishing.
- [20] Jesús A De Loera, Brandon Dutra, Matthias Koeppe, Stanislav Moreinis, Gregory Pinto, and Jianqiu Wu. 2013. Software for exact integration of polynomials over polyhedra. *Computational Geometry* 46, 3 (2013). https://doi.org/10.1016/j.comgeo.2012.09.001
- [21] Simon Duane, Anthony D Kennedy, Brian J Pendleton, and Duncan Roweth. 1987. Hybrid Monte Carlo. *Physics letters B* 195, 2 (1987). https://doi.org/10.1016/0370-2693(87)91197-X
- [22] Martin E. Dyer and Alan M. Frieze. 1988. On the Complexity of Computing the Volume of a Polyhedron. SIAM J. Comput. 17, 5 (1988). https://doi.org/10.1137/0217060
- [23] Christian Fecht and Helmut Seidl. 1999. A Faster Solver for General Systems of Equations. Science of Computer Programming 35, 2 (1999). https://doi.org/10.1016/S0167-6423(99)00009-X
- [24] Timothy S. Freeman and Frank Pfenning. 1991. Refinement Types for ML. In ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 1991. ACM. https://doi. org/10.1145/113445.113468
- [25] Hong Ge, Kai Xu, and Zoubin Ghahramani. 2018. Turing: A Language for Flexible Probabilistic Inference. In International Conference on Artificial Intelligence and Statistics, AISTATS 2018 (PMLR, Vol. 84). https://proceedings.mlr.press/v84/ge18b.html
- [26] Timon Gehr, Sasa Misailovic, and Martin T. Vechev. 2016. PSI: Exact Symbolic Inference for Probabilistic Programs. In *International Conference on Computer Aided Verification, CAV 2016 (LNCS, Vol. 9779)*. Springer. https://doi.org/10.1007/978-3-319-41528-4\_4
- [27] Timon Gehr, Samuel Steffen, and Martin T. Vechev. 2020. λPSI: exact inference for higher-order probabilistic programs. In ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2020. ACM. https://doi.org/10.1145/3385412.3386006
- [28] Jaco Geldenhuys, Matthew B. Dwyer, and Willem Visser. 2012. Probabilistic symbolic execution. In International Symposium on Software Testing and Analysis, ISSTA 2012. ACM. https://doi.org/10.1145/ 2338965.2336773
- [29] Zoubin Ghahramani. 2013. Bayesian non-parametrics and the probabilistic approach to modelling. *Philosophical Transactions of the Royal Society A* 371 (2013). https://doi.org/10.1098/rsta.2011.0553
- [30] Patrice Godefroid. 2007. Compositional dynamic test generation. In ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2007. ACM. https://doi.org/10.1145/1190216.1190226
- [31] Noah D. Goodman, Vikash K. Mansinghka, Daniel M. Roy, Keith Bonawitz, and Joshua B. Tenenbaum. 2008. Church: a language for generative models. In *Conference on Uncertainty in Artificial Intelligence,* UAI 2008. AUAI Press.
- [32] Noah D. Goodman and Andreas Stuhlmüller. 2014. The Design and Implementation of Probabilistic Programming Languages.

- [33] Andrew D. Gordon, Thomas A. Henzinger, Aditya V. Nori, and Sriram K. Rajamani. 2014. Probabilistic programming. In *Future of Software Engineering, FOSE 2014.* ACM. https://doi.org/10.1145/2593882. 2593900
- [34] Maria I. Gorinova, Dave Moore, and Matthew D. Hoffman. 2020. Automatic Reparameterisation of Probabilistic Programs. In *International Conference on Machine Learning, ICML 2020 (PMLR, Vol. 119)*. http://proceedings.mlr.press/v119/gorinova20a.html
- [35] Roger B. Grosse, Siddharth Ancha, and Daniel M. Roy. 2016. Measuring the reliability of MCMC inference with bidirectional Monte Carlo. In Annual Conference on Neural Information Processing Systems, NIPS 2016. https://proceedings.neurips.cc/paper/2016/hash/ 0e9fa1f3e9e66792401a6972d477dcc3-Abstract.html
- [36] Roger B. Grosse, Zoubin Ghahramani, and Ryan P. Adams. 2015. Sandwiching the marginal likelihood using bidirectional Monte Carlo. *CoRR* abs/1511.02543 (2015). https://doi.org/10.48550/arXiv.1511.02543
- [37] N. L. Hjort, Chris Homes, Peter Muller, and Stephen G. Walker. 2010. Bayesian Nonparametrics. Cambridge University Press. https://doi. org/10.1017/CBO9780511802478
- [38] Steven Holtzen, Guy Van den Broeck, and Todd D. Millstein. 2020. Scaling exact inference for discrete probabilistic programs. *Proceedings* of the ACM on Programming Languages 4, OOPSLA (2020). https: //doi.org/10.1145/3428208
- [39] Zixin Huang, Saikat Dutta, and Sasa Misailovic. 2021. AQUA: Automated Quantized Inference for Probabilistic Programs. In International Symposium on Automated Technology for Verification and Analysis, ATVA 2021 (LNCS, Vol. 12971). Springer. https://doi.org/10.1007/978-3-030-88885-5 16
- [40] Andrew Kenyon-Roberts and C.-H. Luke Ong. 2021. Supermartingales, Ranking Functions and Probabilistic Lambda Calculus. In Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021. IEEE. https://doi.org/10.1109/LICS52264.2021.9470550
- [41] Carol Mak, C.-H. Luke Ong, Hugo Paquet, and Dominik Wagner. 2021. Densities of Almost Surely Terminating Probabilistic Programs are Differentiable Almost Everywhere. In *European Sympo*sium on Programming, ESOP 2021 (LNCS, Vol. 12648). Springer. https: //doi.org/10.1007/978-3-030-72019-3\_16
- [42] Carol Mak, Fabian Zaiser, and Luke Ong. 2021. Nonparametric Hamiltonian Monte Carlo. In *International Conference on Machine Learning*, *ICML 2021 (PMLR, Vol. 139)*. http://proceedings.mlr.press/v139/mak21a. html
- [43] Chris Manning and Hinrich Schütze. 1999. Foundations of Statistical Natural Language Processing. MIT Press. Cambridge, MA.
- [44] David Monniaux. 2000. Abstract Interpretation of Probabilistic Semantics. In International Symposium on Static Analysis, SAS 2000 (LNCS, Vol. 1824). Springer. https://doi.org/10.1007/978-3-540-45099-3\_17
- [45] David Monniaux. 2001. An abstract Monte-Carlo method for the analysis of probabilistic programs. In ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2001. ACM. https://doi. org/10.1145/360204.360211
- [46] Markus Müller-Olm and Helmut Seidl. 2004. Precise interprocedural analysis through linear algebra. In ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2004. ACM. https: //doi.org/10.1145/964001.964029
- [47] Praveen Narayanan, Jacques Carette, Wren Romano, Chung-chieh Shan, and Robert Zinkov. 2016. Probabilistic Inference by Program Transformation in Hakaru (System Description). In *International Symposium on Functional and Logic Programming*, FLOPS 2016 (LNCS, Vol. 9613). Springer. https://doi.org/10.1007/978-3-319-29604-3\_5
- [48] Radford M Neal. 2003. Slice sampling. The annals of statistics 31, 3 (2003). https://doi.org/10.1214/aos/1056562461
- [49] Art B. Owen. 2013. Monte Carlo theory, methods and examples.
- [50] Andreas Podelski, Ina Schaefer, and Silke Wagner. 2005. Summaries for While Programs with Recursion. In *European Symposium on Programming, ESOP 2005 (LNCS, Vol. 3444)*. Springer. https://doi.org/10. 1007/978-3-540-31987-0\_8

- [51] David Pollard. 2002. A User's Guide to Measure-Theoretic Probability. Cambridge University Press. https://doi.org/10.1017/ CBO9780511811555
- [52] Fredrik Ronquist, Jan Kudlicka, Viktor Senderov, Johannes Borgström, Nicolas Lartillot, Daniel Lundén, Lawrence Murray, Thomas B Schön, and David Broman. 2021. Universal probabilistic programming offers a powerful approach to statistical phylogenetics. *Communications biology* 4, 1 (2021). https://doi.org/10.1038/s42003-021-01753-7
- [53] Vivekananda Roy. 2020. Convergence Diagnostics for Markov Chain Monte Carlo. Annual Review of Statistics and Its Application 7 (2020). https://doi.org/10.1146/annurev-statistics-031219-041300
- [54] Reuven Y. Rubinstein and Dirk P. Kroese. 2017. Simulation and the Monte Carlo Method (3rd ed.). Wiley. https://doi.org/10.1002/ 9781118631980
- [55] Feras A. Saad, Martin C. Rinard, and Vikash K. Mansinghka. 2021. SPPL: probabilistic programming with fast exact symbolic inference. In ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2021. ACM. https://doi.org/10.1145/ 3453483.3454078
- [56] Sriram Sankaranarayanan, Aleksandar Chakarov, and Sumit Gulwani. 2013. Static analysis for probabilistic programs: inferring whole program properties from finitely many paths. In ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2013. ACM. https://doi.org/10.1145/2491956.2462179
- [57] Chung-chieh Shan and Norman Ramsey. 2017. Exact Bayesian inference by symbolic disintegration. In ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017. ACM. https: //doi.org/10.1145/3009837.3009852
- [58] Michael J. A. Smith. 2008. Probabilistic Abstract Interpretation of Imperative Programs using Truncated Normal Distributions. *Electronic Notes in Theoretical Computer Science* 220, 3 (2008). https://doi.org/10. 1016/j.entcs.2008.11.018
- [59] Sam Staton. 2017. Commutative Semantics for Probabilistic Programming. In European Symposium on Programming, ESOP 2017 (LNCS, Vol. 10201). Springer. https://doi.org/10.1007/978-3-662-54434-1\_32
- [60] Sean Talts, Michael Betancourta, Daniel Simpson, Aki Vehtari, and Andrew Gelman. 2018. Validating Bayesian Inference Algorithms with Simulation-Based Calibration. arXiv 1804.06788 (2018). https: //doi.org/10.48550/arXiv.1804.06788
- [61] David Tolpin, Jan-Willem van de Meent, and Frank D. Wood. 2015. Probabilistic Programming in Anglican. In European Conference on Machine Learning and Knowledge Discovery in Databases, ECML PKDD 2015 (LNCS, Vol. 9286). Springer. https://doi.org/10.1007/978-3-319-23461-8\_36
- [62] Jan-Willem van de Meent, Brooks Paige, Hongseok Yang, and Frank Wood. 2018. An Introduction to Probabilistic Programming. *CoRR* abs/1809.10756 (2018). https://doi.org/10.48550/arXiv.1809.10756
- [63] Di Wang, Jan Hoffmann, and Thomas W. Reps. 2018. PMAF: an algebraic framework for static analysis of probabilistic programs. In ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2018. ACM. https://doi.org/10.1145/3192366. 3192408
- [64] Cheng Zhang, Judith Bütepage, Hedvig Kjellström, and Stephan Mandt. 2019. Advances in Variational Inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 8 (2019). https://doi.org/10.1109/ TPAMI.2018.2889774
- [65] Yuan Zhou, Bradley J. Gram-Hansen, Tobias Kohn, Tom Rainforth, Hongseok Yang, and Frank Wood. 2019. LF-PPL: A Low-Level First Order Probabilistic Programming Language for Non-Differentiable Models. In International Conference on Artificial Intelligence and Statistics, AISTATS 2019 (PMLR, Vol. 89). http://proceedings.mlr.press/v89/ zhou19b.html

# A Supplementary Material for Section 3

# A.1 Intervals as a lattice

Intervals I form a partially ordered set under interval inclusion ( $\sqsubseteq$ ). We will sometimes need the meet  $\sqcap$  and join  $\sqcup$  operations, corresponding to the greatest lower bound and the least upper bound of two intervals. Note that the meet of two intervals does not exist if the two intervals are disjoint. Concretely, these two operations are given by  $[a, b] \sqcap [c, d] := [\max(a, c), \min(b, d)]$  (if the two intervals overlap) and  $[a, b] \sqcup [c, d] := [\min(a, c), \max(b, d)]$ .

For some applications (e.g. the interval type system), we need the interval domain to be a true lattice. To turn I into a lattice, we add a bottom element  $\perp$  (signifying an empty interval). The definition of the meet  $\sqcap$  and join  $\sqcup$  is extended in the natural way. The meet  $\sqcap$  is extended by defining  $I_1 \sqcap$  $\perp = \perp \sqcap I_2 = \perp$  and  $I_1 \sqcap I_2 = \perp$  if the two intervals  $I_1, I_2 \in I$ are disjoint. The join  $\sqcup$  satisfies  $I \sqcup \bot = \bot \sqcup I = I$ .

### A.2 Lifting Functions to Intervals

For constants  $c \in \mathbb{R}$  (i.e. nullary functions), for common functions like +, -, ×,  $|\cdot|$ , min, max, for monotonically increasing functions  $f_{\nearrow} : \mathbb{R} \to \mathbb{R}$ , and for monotonically decreasing functions  $f_{\searrow} : \mathbb{R} \to \mathbb{R}$ , it is easy to describe the interval-lifted functions  $+^{\mathbb{I}}$ ,  $-^{\mathbb{I}}$ ,  $\times^{\mathbb{I}}$ ,  $|\cdot|^{\mathbb{I}}$ , min<sup> $\mathbb{I}$ </sup>, max<sup> $\mathbb{I}$ </sup>,  $f_{\nearrow}^{\mathbb{I}}$ , and  $f_{\searrow}^{\mathbb{I}}$ :

$$c^{\mathbb{I}} = [c, c]$$

$$-^{\mathbb{I}}[x_1, y_1] = [-y_1, -x_1]$$

$$|[x_1, y_1]|^{\mathbb{I}} = \begin{cases} [0, \max(|x_1|, |y_1|)] \text{ if } x_1 \le 0 \le y_1 \\ [\min(|x_1|, |y_1|)], \max(|x_1|, |y_1|)] \text{ else} \end{cases}$$

$$[x_1, y_1] +^{\mathbb{I}} [x_2, y_2] = [x_1 + x_2, y_1 + y_2]$$

$$[x_1, y_1] -^{\mathbb{I}} [x_2, y_2] = [x_1 - y_2, y_1 - x_2]$$

$$[x_1, y_1] \times^{\mathbb{I}} [x_2, y_2] = [\min(x_1 x_2, x_1 y_2, y_1 x_2, y_1 y_2), \max(x_1 x_2, x_1 y_2, y_1 x_2, y_1 y_2)]$$

$$\min^{\mathbb{I}}([x_1, y_1], [x_2, y_2]) = [\min(x_1, x_2), \min(y_1, y_2)]$$

 $\max^{\mathbb{I}}([x_1, y_1], [x_n, y_n]) = [\max(x_1, x_2), \max(y_1, y_2)]$ 

$$\begin{split} f^{\mathbb{I}}_{\nearrow}([x_1,y_1]) &= [f_{\nearrow}(x_1), f_{\nearrow}(y_1)] \\ f^{\mathbb{I}}_{\searrow}([x_1,y_1]) &= [f_{\searrow}(y_1), f_{\searrow}(x_1)] \end{split}$$

where we write  $f(\pm \infty)$  for  $\lim_{x\to\pm\infty} f(x) \in \mathbb{R}_{\infty}$ , respectively.

#### A.3 Properties of Interval Reduction

We can define a refinement relation  $M \triangleleft M'$  ("*M* refines *M*'") between a standard term *M* and an interval term *M'*, if *M* is obtained from *M'* by replacing every occurrence of [a, b] with some  $r \in [a, b]$ .

**Lemma 3.1.** Let  $\vdash P : \mathbf{R}$  be a program. For any interval trace t and concrete trace  $s \triangleleft t$ , we have  $wt_P(s) \in wt_P^{\mathbb{I}}(t)$  and  $val_P(s) \in val_P^{\mathbb{I}}(t)$  (provided  $val_P(s)$  is defined).

*Proof.* If the interval reduction  $\rightarrow_{\mathbb{I}}$  gets stuck,  $\mathsf{wt}_P^{\mathbb{I}}$  is  $[0, \infty]$ and  $\mathsf{val}_P^{\mathbb{I}}$  is  $[-\infty, \infty]$ , so the claim is certainly true. Otherwise, for each  $(M_{\mathbb{I}}, t, w_{\mathbb{I}}) \rightarrow_{\mathbb{I}} (M'_{\mathbb{I}}, t', w'_{\mathbb{I}})$  reduction step, we can do a reduction step  $(M, s, w) \rightarrow (M', s', w')$  where  $M' \triangleleft M'_{\mathbb{I}}$  and  $w' \in w'_{\mathbb{I}}$ , and  $s' \triangleleft t'$  if  $M \triangleleft M_{\mathbb{I}}$ ,  $w \in w_{\mathbb{I}}$ , and  $s \triangleleft t$ . Since the reduction doesn't get stuck, we end up with a value  $\underline{r} \triangleleft [\underline{a}, \underline{b}]$ , so  $\mathsf{val}_P(s) = r \in [a, b] = \mathsf{val}_P^{\mathbb{I}}(t)$ .

# A.4 Additional Possible Reduction Rules

The interval semantics as presented has the unfortunate property that even a simple program like

requires infinitely many interval traces to achieve a finite upper bound. The reason is that the right branch score(1) is taken if the sampled value is in the open interval (0, 1]. To approximate this using closed intervals [a, b] that our analysis supports, we need infinitely many intervals, e.g. { $[2^{-n-1}, 2^{-n}]$  |  $n \in \mathbb{N}$ }. Adding (half-)open intervals to the semantics would solve this specific problem, but not more general ones, where the guard condition is for example

sample – sample 
$$\leq 0$$
.

In that case, we have to approximate the set  $\{(x, y) \in [0, 1]^2 | x \le y\}$ . For the lower bounds, that is not an issue, but for the upper bounds, we need an infinite number of interval traces again. We would like to use the interval traces  $\langle [0, \frac{1}{2}], [0, \frac{1}{2}] \rangle$  and  $\langle [\frac{1}{2}, 1], [0, 1] \rangle$  to cover this set, but the reduction gets stuck on them because it is not clear which branch should be taken.

To remedy this, we could add the following two rules.

$$\frac{a \leq 0 < b}{(if(\underline{[a,b]}, N, P), t, w) \rightarrow_{\mathbb{I}} (N, t, w \times^{\mathbb{I}} [0,1])}$$
$$\frac{a \leq 0 < b}{(if(\underline{[a,b]}, N, P), t, w) \rightarrow_{\mathbb{I}} (P, t, w \times^{\mathbb{I}} [0,1])}$$

They basically express that if the interval bounds are not precise enough to decide what branch to take, we can take both, but have to allow the weight to be zero because it's not guaranteed that the taken branch can actually happen. This change can only improve the upper bounds, not the lower bounds because the lower bound on each weight is zero if the additional rules are used. Then the definition of upper bound can be modified in the following way:

$$upperBd_{P}^{\mathcal{T}}(U) := \sum_{t \in \mathcal{T}} \sum_{\substack{(P,t,[1,1]) \to 1\\ (\underline{[a,b]}, \langle \rangle, [w_1, w_2])}} vol(t) \cdot w_2 \cdot \left[ [a,b] \cap U \neq \emptyset \right]$$

This is the strategy we use for our implementation and is a natural extension of the existing semantics: it requires very few changes to the soundness and completeness proofs.

A downside of the previous approach is that the bounds are not always very tight: for the term if(...) score(50)



Figure 8. Reduction rules for symbolic execution

else score(100), it returns bounds [0, 150] instead of [50, 100]. To improve this, we could omit the multiplication with [0, 1].

$$\frac{a \le 0 < b}{(if(\underline{[a,b]}, N, P), t, w) \to_{\mathbb{I}} (N, t, w)}$$
$$\frac{a \le 0 < b}{(if(\underline{[a,b]}, N, P), t, w) \to_{\mathbb{I}} (P, t, w)}$$

However, this complicates the equations of our bounds. With this semantics, we have to compute minima and suprema instead of a simple sum:

$$lowerBd_{P}^{\mathcal{T}}(U) := \sum_{t \in \mathcal{T}} \min_{\substack{(P,t,[1,1]) \to 1 \\ (\underline{[a,b]}, \langle \rangle, [w_1, w_2])}} vol(t) \cdot w_1 \cdot \left[ [a, b] \subseteq U \right]$$
$$upperBd_{P}^{\mathcal{T}}(U) := \sum_{t \in \mathcal{T}} \sup_{\substack{(P,t,[1,1]) \to 1 \\ ([a,b], \langle \rangle, [w_1, w_2])}} vol(t) \cdot w_2 \cdot \left[ [a, b] \cap U \neq \emptyset \right]$$

This is harder to implement because the sums cannot be computed incrementally, but many temporary results have to be kept in memory to compute the minima and suprema. Proving soundness and completeness for this would require more substantial changes to the proofs.

# **B** Symbolic Execution

In this section we formally introduce stochastic symbolic execution. We make use of this form of symbolic execution in two separate ways. First, our completeness proof hinges on guarantees provided by the symbolic execution in order to identify a suitable set of interval traces. Second, our tool GuBPI relies on the symbolic execution as a first step in the program analysis, in order to identify relevant paths and independent subexpressions, and to avoid repeated evaluation in a small-step semantics.

*High-level idea.* The overarching idea of symbolic execution is to postpone the evaluation of sample expressions and instead use a *sample variable* to symbolically represent its outcome. As a consequence, branching and scoring steps cannot be executed concretely, so we record them symbolically instead.

Symbolic terms. To postpone concrete sample decisions we introduce sample variables  $\alpha_1, \alpha_2, \ldots$  into our language. We then define symbolic terms and symbolic values by extending interval terms and values by adding two new constructs: every sample variable  $\alpha_j$  is a symbolic value and for every primitive function f and symbolic values  $\mathcal{V}_1, \ldots, \mathcal{V}_{|f|}$ , the symbolic term  $f(\mathcal{V}_1, \ldots, \mathcal{V}_{|f|})$  is a symbolic value, denoting a function application that is postponed until all sample variables are instantiated. We denote symbolic terms by  $\mathcal{M}, \mathcal{N}, \mathcal{P}$  and symbolic values by  $\mathcal{V}, \mathcal{W}$ . Formally we define

$$\mathcal{V} := x \mid \underline{r} \mid \lambda x.\mathcal{M} \mid \mu_x^{\varphi}.\mathcal{M} \mid \alpha_i \mid \underline{f}(\mathcal{V}_1, \dots, \mathcal{V}_{|f|})$$
$$\mathcal{M}, \mathcal{N}, \mathcal{P} := \mathcal{V} \mid \mathcal{M}\mathcal{N} \mid \text{if}(\mathcal{M}, \mathcal{N}, \mathcal{P}) \mid \underline{f}(\mathcal{M}_1, \dots, \mathcal{M}_{|f|})$$

 $| \text{ sample} | \text{ score}(\mathcal{M})$ 

The definition of redex and evaluation context extends naturally (recall that we regard  $\alpha_i$  as a value).

Symbolic execution. A symbolic constraint is a pair ( $\mathcal{V} \bowtie$ r) where  $\mathcal{V}$  is a symbolic value,  $\bowtie \in \{\leq, <, \geq, >\}$  and  $r \in \mathbb{R}$ . A symbolic configuration has the form  $\psi = (\mathcal{M}, n, \Delta, \Xi)$ where  $\mathcal{M}$  is a symbolic term,  $n \in \mathbb{N}$  a natural number (used to obtain fresh sample variables),  $\Delta$  a set of symbolic constraints (which track the symbolic conditions on the current execution path), and  $\Xi$  is a set of symbolic values (which records all symbolic values scored on the current path). When executing symbolically: (1) we evaluate each sample to a fresh sample variable, (2) we postpone function application, (3) for each conditional, we explore both branches (our reduction is nondeterministic) and record the symbolic inequalities that must hold along the current path, and (4) we record the symbolic values that we scored with. We give the reduction rules in Fig. 8.

We call a tuple  $\Psi = (\Psi, n, \Delta, \Xi)$  (a symbolic configuration where the symbolic term is a value) a *symbolic path*. For a symbolic configuration  $\psi$ , we write *symPaths*( $\psi$ ) for the set of symbolic paths reached when evaluating from  $\psi$ . Note that *symPaths*( $\psi$ ) is countable.

Let  $\mathcal{V}$  be a symbolic value of type **R** (no  $\lambda$ -abstraction or fixed point) with sample variables within  $\{\alpha_1, \ldots, \alpha_n\}$ . For a trace  $\mathbf{s} = \langle r_1, \ldots, r_n \rangle \in [0, 1]^n$ , we define  $\mathcal{V}[\mathbf{s}/\overline{\alpha}]$  as the value (in  $\mathbb{R}$ ) obtained by replacing the sample variables in  $\overline{\alpha}$  with  $\mathbf{s}$  and evaluate the postponed primitive function applications. For a symbolic path  $\Psi = (\mathcal{V}, n, \Delta, \Xi)$ , we define  $\llbracket \Psi \rrbracket (U)$  as

$$\int_{[0,1]^n} \left[ \mathcal{V}[\boldsymbol{s}/\overline{\alpha}] \in U \right] \prod_{C \bowtie r \in \Delta} \left[ C[\boldsymbol{s}/\overline{\alpha}] \bowtie r \right] \prod_{\mathcal{W} \in \Xi} \mathcal{W}[\boldsymbol{s}/\overline{\alpha}] \, \mathrm{d}\boldsymbol{s}.$$

**Solution to symbolic constraints.** To simplify notation (and avoid extensive use of Iverson brackets) we introduce notation for the set of traces that satisfy a set of symbolic constraints. Given a set of symbolic constraints  $\Delta$  with sample

Guaranteed Bounds for Posterior Inference in Universal Probabilistic Programming

PLDI '22, June 13-17, 2022, San Diego, CA, USA

variables contained in  $\{\alpha_1, \ldots, \alpha_n\}$  we define

$$\operatorname{Sat}_{n}(\Delta) := \bigcap_{(\mathcal{V} \bowtie r) \in \Delta} \{ s \in [0,1]^{n} \mid \mathcal{V}[s/\overline{\alpha}] \bowtie r \}$$

as the set of actual values for the sample variables that satisfy all constraints. It follows immediately from the definitions that we can replace the Iverson brackets in  $\llbracket \Psi \rrbracket$  by directly restricting the integral to the traces in  $\operatorname{Sat}_n(\Delta)$ .

**Lemma B.1.** For any symbolic path  $(\mathcal{V}, n, \Delta, \Xi)$  and any  $U \in \Sigma_{\mathbb{R}}$  we have

$$\llbracket (\mathcal{V}, n, \Delta, \Xi) \rrbracket (U) = \int_{\mathsf{Sat}_n(\Delta)} [\mathcal{V}[\mathbf{s}/\overline{\alpha}] \in U] \prod_{\mathcal{W} \in \Xi} \mathcal{W}[\mathbf{s}/\overline{\alpha}] \, \mathrm{d}\mathbf{s}.$$

**Correctness of symbolic execution.** We can now establish a correspondence between symbolic execution and the ordinary reduction. If we wish to symbolically analyse a term *P*, we consider the (symbolic) reductions starting from  $(P, 0, \emptyset, \emptyset)$ , resulting in the symbolic paths *symPaths* $(P, 0, \emptyset, \emptyset)$ .

**Lemma B.2.** Let  $\vdash P : \mathbf{R}$  and suppose we have  $(\mathcal{V}, n, \Delta, \Xi) \in$ symPaths $(P, 0, \emptyset, \emptyset)$ , where P is interpreted as a symbolic term. Then for any  $\mathbf{s} \in Sat_n(\Delta)$ , we have

$$(P, \mathbf{s}, 1) \to^* (\mathcal{V}[\mathbf{s}/\overline{\alpha}], \langle\rangle, \prod_{\mathcal{W} \in \Xi} \mathcal{W}[\mathbf{s}/\overline{\alpha}])$$

*Proof.* A similar proof can be found in [41, Theorem 1].

**Lemma B.3.** Let  $\vdash P : \mathbf{R}$  and suppose  $(P, \mathbf{s}, 1) \rightarrow^* (\underline{r}, \langle \rangle, w)$ for some  $r \in \mathbb{R}$ . Then there exists a unique  $(\mathcal{V}, n, \Delta, \Xi) \in$ symPaths $(P, 0, \emptyset, \emptyset)$  such that  $\mathbf{s} \in Sat_n(\Delta)$ . For this unique symbolic path we have  $w = \prod_{W \in \Xi} \mathcal{W}[\mathbf{s}/\overline{\alpha}]$  and  $r = \mathcal{V}[\mathbf{s}/\overline{\alpha}]$ . *Proof sketch.* Choose the same branches in the  $\rightsquigarrow$ -reduction of *P* as in its  $\rightarrow$ -reduction. Then it is straightforward to see that this correspondence holds at every symbolic reduction step: if  $(P, 0, \emptyset, \emptyset) \sim^* (\mathcal{P}', n, \Delta, \Xi)$  then the corresponding  $\rightarrow$ -reduction steps yield  $(P, \mathbf{ss}', 1) \rightarrow^* (P', \mathbf{s}', w)$  where  $\mathbf{s}$  has length n, P' is  $\mathcal{P}'[\mathbf{s}/\overline{\alpha}]$  (after evaluating delayed primitve function applications),  $\Delta$  records the guards  $C[\mathbf{s}/\overline{\alpha}] \leq 0$  or  $C[\mathbf{s}/\overline{\alpha}] > 0$  that need to hold for the trace  $\mathbf{s}$ , and finally, the weight w is given by  $\prod_{W \in \Xi} \mathcal{W}[\mathbf{s}/\overline{\alpha}]$  at any point.  $\Box$ 

**Theorem 6.1.** Let  $\vdash P : \mathbf{R}$  be a program and  $U \in \Sigma_{\mathbb{R}}$ . Then

$$\llbracket P \rrbracket(U) = \sum_{\Psi \in symPaths(P,0,\emptyset,\emptyset)} \llbracket \Psi \rrbracket(U).$$

Proof.

$$\llbracket P \rrbracket(U) = \sum_{n \in \mathbb{N}} \int_{[0,1]^n} [\operatorname{val}_P(s) \in U] \operatorname{wt}_P(s) \, \mathrm{d}s$$
$$= \sum_{n \in \mathbb{N}} \int_{[0,1]^n} \sum_{(\mathcal{V}, n, \Delta, \Xi)} \left( [s \in \operatorname{Sat}_n(\Delta)] [\mathcal{V}[s/\overline{\alpha}] \in U] \prod_{\mathcal{W} \in \Xi} \mathcal{W}[s/\overline{\alpha}] \, \mathrm{d}s \right)$$
$$= \sum_{(\mathcal{V}, n, \Delta, \Xi)} \int_{\operatorname{Sat}_n(\Delta)} [\mathcal{V}[s/\overline{\alpha}] \in U] \prod_{\mathcal{W} \in \Xi} \mathcal{W}[s/\overline{\alpha}] \, \mathrm{d}s$$

$$= \sum_{(\mathcal{V},n,\Delta,\Xi)} \llbracket (\mathcal{V},n,\Delta,\Xi) \rrbracket (U)$$

where the sum ranges over symbolic paths  $(\mathcal{V}, n, \Delta, \Xi) \in symPaths(P, 0, \emptyset, \emptyset)$ . The first equality is by definition, the second one by Lemmas B.2 and B.3, the third by noting that  $\operatorname{Sat}_n(\Delta) \subseteq [0, 1]^n$  and exchanging the infinite sum and integral (which is allowed because everything is nonnegative) and the fourth by Lemma B.1.

# C Supplementary Material for Section 4

# C.1 Infinite Trace Semantics

A convenient alternative to the (finite) trace semantics is using infinite traces  $\mathbb{T}_{\infty} := [0, 1]^{\mathbb{N}}$  with a suitable  $\sigma$ -algebra and measure  $\mu_{\mathbb{T}_{\infty}}$  [16, 40]. The  $\sigma$ -algebra on  $\mathbb{T}_{\infty}$  is defined as the smallest  $\sigma$ -algebra that contains all sets  $U \times \mathbb{T}_{\infty}$  where  $U \in \Sigma_{[0,1]^n}$  for some  $n \in \mathbb{N}$ . The measure  $\mu_{\mathbb{T}_{\infty}}$  is the unique measure with  $\mu_{\mathbb{T}_{\infty}}(U \times \mathbb{T}_{\infty}) = \lambda_n(U)$  for  $U \in \Sigma_{[0,1]^n}$ . We use the symbol  $\boldsymbol{u}$  for an infinite trace in  $\mathbb{T}_{\infty}$ . For a finite trace  $\boldsymbol{s}$ and infinite trace  $\boldsymbol{u}$  we write  $\boldsymbol{s}\boldsymbol{u} \in \mathbb{T}_{\infty}$  for their concatenation. For any infinite trace  $\boldsymbol{u} \in \mathbb{T}_{\infty}$ , there is at most one prefix  $\boldsymbol{s} \in \mathbb{T}$  with wt $_P(\boldsymbol{s}) > 0$  since the reduction is deterministic. We can therefore define wt $_P^{\infty}(\boldsymbol{u}) := \text{wt}_P(\boldsymbol{s})$  and  $\text{val}_P^{\infty}(\boldsymbol{u}) :=$  $\text{val}_P(\boldsymbol{s})$  if such a prefix  $\boldsymbol{s}$  exists, and wt $_P^{\infty}(\boldsymbol{u}) := 0$  and  $\text{val}_P^{\infty}(\boldsymbol{u})$ is undefined otherwise. The infinite trace semantics of a term is then defined as

$$\llbracket P \rrbracket(U) := \int_{(\mathsf{val}_P^{\infty})^{-1}(U)} \mathsf{wt}_P^{\infty}(\boldsymbol{u}) \, \mu_{\mathbb{T}_{\infty}}(\mathrm{d}\boldsymbol{u}).$$

**Lemma C.1.** The finite and infinite trace semantics agree, that is:

$$\int_{(\operatorname{val}_P)^{-1}(U)} \operatorname{wt}_P(\boldsymbol{s}) \, \mu_{\mathbb{T}}(\mathrm{d}\boldsymbol{s}) = \int_{(\operatorname{val}_P^{\infty})^{-1}(U)} \operatorname{wt}_P^{\infty}(\boldsymbol{u}) \, \mu_{\mathbb{T}_{\infty}}(\mathrm{d}\boldsymbol{u}).$$

*Proof.* Observe that  $\operatorname{val}_{P}^{\infty}(su) = \operatorname{val}_{P}(s)$  and  $\operatorname{wt}_{P}^{\infty}(su) = \operatorname{wt}_{P}(s)$  for all  $u \in \mathbb{T}_{\infty}$  if  $\operatorname{wt}_{P}(s) > 0$ . Then we get:

$$\begin{split} &\int_{(\operatorname{val}_P)^{-1}(U)} \operatorname{wt}_P(s) \, \mu_{\mathbb{T}}(\mathrm{d}s) \\ &= \int_{(\operatorname{wt}_P)^{-1}(\mathbb{R}_{>0})} [\operatorname{val}_P(s) \in U] \operatorname{wt}_P(s) \int_{\mathbb{T}_{\infty}} \mu_{\mathbb{T}_{\infty}}(\mathrm{d}u) \mu_{\mathbb{T}}(\mathrm{d}s) \\ &= \int_{(\operatorname{wt}_P)^{-1}(\mathbb{R}_{>0})} \int_{\mathbb{T}_{\infty}} [\operatorname{val}_P^{\infty}(su) \in U] \operatorname{wt}_P^{\infty}(su) \mu_{\mathbb{T}_{\infty}}(\mathrm{d}u) \mu_{\mathbb{T}}(\mathrm{d}s) \\ &= \int_{(\operatorname{wt}_P)^{-1}(\mathbb{R}_{>0}) \times \mathbb{T}_{\infty}} [\operatorname{val}_P^{\infty}(u) \in U] \operatorname{wt}_P^{\infty}(u) \mu_{\mathbb{T}_{\infty}}(\mathrm{d}u) \\ &= \int_{(\operatorname{val}_P^{\infty})^{-1}(U)} \operatorname{wt}_P^{\infty}(u) \mu_{\mathbb{T}_{\infty}}(\mathrm{d}u) \end{split}$$

where we used the fact that the sets  $\{s\} \times \mathbb{T}_{\infty}$  are disjoint for different  $s \in \operatorname{wt}_{P}^{-1}(\mathbb{R}_{>0})$  because otherwise we would be able to find a trace s as a prefix of s' and both having positive weight, which is impossible due to the deterministic reduction. Therefore  $(\operatorname{wt}_{P})^{-1}(\mathbb{R}_{>0}) \times \mathbb{T}_{\infty} = (\operatorname{wt}_{P}^{\infty})^{-1}(\mathbb{R}_{>0})$  and everything works as desired. Note that the second to last equality follows from Fubini's theorem and the fact that the product measure of  $\mu_{\mathbb{T}}$  and  $\mu_{\mathbb{T}_{\infty}}$  is  $\mu_{\mathbb{T}_{\infty}}$  again. П

#### C.2 Exhaustivity and Soundness

**Example C.1** (more examples of exhaustive sets). Here are more examples and counterexamples for exhaustivity.

(i)  $\{\langle\rangle\}$  is an (uninteresting) exhaustive set and only useful for deterministic programs.

(ii)  $\{\langle [2^{-n-1}, 2^{-n}] \rangle \mid n \in \mathbb{N}\}$  is exhaustive because only the trace  $\langle 0 \rangle$  (with measure 0) is not covered.

(iii) Let  $a_n \ge 0$  be a converging series, i.e.  $\sum_{i=1}^n a_i < \infty$ , for example  $a_n = n^{-2}$ . Define

$$\mathcal{T} := \{ \langle [0, e^{-a_1}], \dots, [0, e^{-a_n}], [e^{-a_{n+1}}, 1] \rangle \mid n \in \mathbb{N} \}.$$

This is not an exhaustive set since it doesn't cover any of the traces in

$$[0, e^{-a_1}) \times [0, e^{-a_2}, 1) \times \cdots$$

which has measure  $\prod_{i=1}^{\infty} e^{-a_i} = \exp\left(-\sum_{i=1}^{\infty} a_i\right) > 0.$ 

We also note that exhaustivity can be expressed just in terms of finite traces as well, at the cost of a more complicated definition.

**Lemma C.2.** A set of interval traces  $\mathcal{T}$  is exhaustive if and only if

$$\mu_{\mathbb{T}}\left([0,1]^n \setminus \left(\bigcup_{\langle I_1,\dots,I_m \rangle \in \mathcal{T}, m \le n} I_1 \times \dots \times I_m \times [0,1]^{n-m}\right)\right) \to 0$$
  
as  $n \to \infty$ .

*Proof.* Let  $S := \mathbb{T}_{\infty} \setminus \bigcup_{t \in \mathcal{T}} cover(t)$ . By the definition of exhaustivity,  $\mu_{\mathbb{T}_m}(S) = 0$ . Let

$$S_n = [0,1]^n \setminus \left( \bigcup_{\langle I_1, \dots, I_m \rangle \in \mathcal{T}, m \leq n} I_1 \times \dots \times I_m \times [0,1]^{n-m} \right).$$

It's easy to see that  $S = \bigcap_{n=0}^{\infty} S_n \times \mathbb{T}_{\infty}$  where  $S_n \times \mathbb{T}_{\infty}$  is a decreasing sequence of sets:  $S_1 \times \mathbb{T}_{\infty} \supseteq S_2 \times \mathbb{T}_{\infty} \supseteq \cdots$ . Since measures are continuous from above, we have  $\lim_{n\to\infty} \mu_{\mathbb{T}}(S_n) =$  $\lim_{n\to\infty} \mu_{\mathbb{T}_{\infty}}(S_n \times \mathbb{T}_{\infty}) = \mu_{\mathbb{T}_{\infty}}(S) = 0$ , as desired. 

The following lemma establishes a correspondence between infinite trace semantics and interval trace semantics.

**Lemma C.3.** For any interval trace t and infinite trace  $s_{\infty}$ with a prefix s such that  $s \triangleleft t$ , we have  $wt_p^{\infty}(s_{\infty}) \in wt_p^{\mathbb{I}}(t)$ and  $\operatorname{val}_{p}^{\infty}(\boldsymbol{s}_{\infty}) \in \operatorname{val}_{p}^{\mathbb{I}}(\boldsymbol{t}).$ 

*Proof.* Follows directly from the definition of infinite trace semantics and Lemma 3.1. 

Using the previous results, we can prove soundness of upper bounds.

**Theorem 4.2** (Sound upper bounds). Let  $\mathcal{T}$  be a countable and exhaustive set of interval traces and  $\vdash P : \mathbf{R}$  a program. Then  $\llbracket P \rrbracket \leq upperBd_p^{\mathcal{T}}$ .

*Proof.* For any 
$$U \in \Sigma_{\mathbb{R}}$$
, we have

$$upperBd_{P}^{\mathcal{T}}(U) = \sum_{t \in \mathcal{T}} vol(t)(sup wt_{P}^{\mathbb{I}}(t))[val_{P}^{\mathbb{I}}(t) \cap U \neq \emptyset]$$
$$= \sum_{t \in \mathcal{T}} \int_{(t)} (sup wt_{P}^{\mathbb{I}}(t))[val_{P}^{\mathbb{I}}(t) \cap U \neq \emptyset] ds$$
$$\geq \sum_{t \in \mathcal{T}} \int_{(t)} \int_{\mathbb{T}_{\infty}} wt_{P}^{\infty}(su)[val_{P}^{\infty}(su) \in U] du ds \qquad (4)$$
$$\geq \int_{\bigcup_{t \in \mathcal{T}} (t) \times \mathbb{T}_{\infty}} wt_{P}^{\infty}(u)[val_{P}^{\infty}(u) \in U] ds$$
$$\geq \int_{\mathbb{T}_{\infty}} wt_{P}^{\infty}(u)[val_{P}^{\infty}(u) \in U] du \qquad (5)$$
$$= \|P\|(U) \qquad (6)$$

$$\llbracket P \rrbracket(U) \tag{6}$$

where Eq. (4) follows from Lemma C.3, Eq. (5) from exhaustivity and Eq. (6) from Lemma C.1. П

### C.3 Assumptions for Completeness

Remarks on Assumption 1. We can formally express Assumption 1 from Section 4 about a given program  $\vdash P$  : **R** as follows. For each symbolic path  $\Psi = (\mathcal{V}, n, \Delta, \Xi)$ , we require that  $\mathcal{V}$ , each C with  $C \bowtie 0 \in \Delta$ , and each  $\mathcal{W} \in \Xi$ contain each sample variable  $\alpha_i$  at most once.

**Example C.2.** The pedestrian example (Example 1.1) satisfies Assumption 1 because the symbolic paths have the form  $\Psi = (\Psi, n, \Delta, \Xi)$  with:

$$\begin{split} \mathcal{V} &= 3\alpha_1 \\ n &= 2k+1 \\ \Delta &= \{\alpha_3 - \frac{1}{2} \bowtie 0, \alpha_5 - \frac{1}{2} \bowtie 0, \dots, \alpha_{2k+1} - \frac{1}{2} \bowtie 0\} \\ &\cup \{3\alpha_1 > 0, \\ &3\alpha_1 \pm \alpha_2 > 0, \\ &\dots, \\ &3\alpha_1 \pm \alpha_2 \pm \alpha_4 \dots \pm \alpha_{2k-2} > 0, \\ &3\alpha_1 \pm \alpha_2 \pm \alpha_4 \dots \pm \alpha_{2k} \le 0\} \\ \Xi &= \{ \mathrm{pdf}_{\mathrm{Normal}(1,1,0,1)} (\alpha_2 + \alpha_4 + \dots + \alpha_{2k}) \} \end{split}$$

As we can see, none of the symbolic values contains a sample variable twice, so the assumption is satisfied.

Remarks on Assumption 2. We first prove the sufficient condition for interval separability from Section 4.2.

**Lemma C.4.** If a function  $f : \mathbb{R}^n \to \mathbb{R}$  is boxwise continuous and preimages of points are null sets then f is interval separable.

*Proof.* We decompose  $f^{-1}([a, b]) = f^{-1}((a, b)) \cup f^{-1}(\{a, b\})$ and deal with the former set first. By boxwise continuity,  $f = \bigcup_i f|_{B_i}$  where  $\bigcup_i B_i = \mathbb{R}^n$  and each  $f|_{B_i}$  is continuous on  $B_i$ . To show that the preimage  $f^{-1}((a, b))$  can be tightly approximated by a countable set of boxes, it suffices to show this for each  $(f|_{B_i})^{-1}((a, b))$ . This set is open in  $B_i$  by continuity of  $f|_{B_i}$ , so it can be written as a countable union of boxes (e.g. by taking a box within  $B_i$  around each rational point, which exists because it's an open set). By the assumption, the preimage  $f^{-1}(\{a, b\})$  is a null set. Hence  $f^{-1}([a, b])$ can be approximated by a null set. □

Note that a composition of interval separable functions need *not* be interval separable. This is an incorrect assumption made in the completeness proof of [4]. (To fix their Theorem 3.8, one needs to make the additional assumption that the set of primitive functions be closed under composition.) To see this, let  $f, g : \mathbb{R} \to \mathbb{R}$  be interval separable functions and I an interval. By definition, there are intervals  $B_i$  such that  $\bigcup_i B_i \cup N = f^{-1}(I)$  where N is a null set. Then by interval separability, the preimage  $g^{-1}(\bigcup_i B_i)$  can be tightly approximated by interals  $B'_j$ , but the preimage  $g^{-1}(N)$  need not be a null set. It is also not clear at all whether one can approximate the preimage  $(f \circ g)^{-1}(I)$  tightly using intervals without further restrictions on f and g. For this reason, we require the assumption that the set of primitive functions be closed under composition.

It is not immediately obvious that such a set of functions exists. One example is given by the following. A function  $f : \mathbb{R}^n \to \mathbb{R}$  is called a *submersion* if it is continuously differentiable and its gradient is nonzero everywhere.

**Lemma C.5.** The set  $\mathcal{F}_{subm}$  of submersions is closed under composition and each of its functions is boxwise continuous and interval separable.

*Proof.* Boxwise continuity is obvious given that the functions are even continuously differentiable. For interval separability, we use Lemma C.4. Let  $f : \mathbb{R}^n \to \mathbb{R} \in \mathcal{F}_{subm}$ . Since f is a submersion, the preimage  $f^{-1}(x)$  of any point  $x \in \mathbb{R}$  is an (n-1)-dimensional submanifold of  $\mathbb{R}^n$  by the preimage theorem (a variation of the implicit function theorem). Submanifolds of codimension > 1 have measure zero. (This well-known fact can be shown by writing the submanifold as a countable union of graphs and applying Fubini's theorem to each of them.) Therefore, the lemma applies.

For closure under composition, let  $f : \mathbb{R}^m \to \mathbb{R}$  and  $f_i : \mathbb{R}^{n_i} \to \mathbb{R}$  for  $i \in \{1, ..., m\}$ , all in  $\mathcal{F}_{subm}$ . The composition  $g := f \circ (f_1 \times \cdots \times f_m)$  is clearly  $C^1$  again, so we just have to check the submersion property. By the chain rule, we find that the gradient of the composition

$$\nabla g(x_1,\ldots,x_m) = \begin{pmatrix} \partial_1 f(f_1(x_1),\ldots,f_m(x_m)) \cdot \nabla f_1(x_1) \\ \vdots \\ \partial_m f(f_1(x_1),\ldots,f_m(x_m)) \cdot \nabla f_m(x_m) \end{pmatrix}$$

is nonzero because at least one of the  $\partial_i f$  is nonzero and  $\nabla f_i(x_i) \neq 0$  by assumption. Hence the composition is a submersion again.

Unfortunately, the set of submersions does not contain constant functions. This is a problem because then it is not guaranteed that partially applying a primitive function to a constant is still an admissible primitive function. (For example, this would break Lemma C.8.) Hence we need to assume that *all constant functions be primitive functions*. Luckily, the set  $\mathcal{F}_{subm}$  of submersions can be easily extended to accommodate this.

**Lemma C.6.** Let  $\mathcal{F}_{subm}^*$  be the set of functions  $f : \mathbb{R}^n \to \mathbb{R}$  (for all  $n \in \mathbb{N}$ ) such that whenever the partial derivative  $\partial_i f(x)$  is zero for some  $i \in \{1, ..., n\}$  and  $x \in \mathbb{R}^n$  then f is constant in its *i*-th argument, *i*.e. there is a function  $f^* : \mathbb{R}^{n-1} \to \mathbb{R}$  such that  $f(x_1, ..., x_n) = f^*(x_1, ..., x_{i-1}, x_{i+1}, ..., x_n)$ . This set satisfies all the assumptions about sets of primitive functions: it is closed under composition, contains all constant functions, and all its functions are boxwise continuous and interval separable.

*Proof.* Boxwise continuity is obvious given that the functions are even continuously differentiable. Similarly, it is clear that  $\mathcal{F}^*_{subm}$  contains all constant functions.

For interval separability, let  $f : \mathbb{R}^n \to \mathbb{R} \in \mathcal{F}^*_{subm}$  and  $J \subseteq \{1, \ldots, n\}$  be the set of indices in which f is not constant, and J' its complement. Hence there is a submersion  $f_J : \mathbb{R}^{|J|} \to \mathbb{R}$  such that  $f(x) = f_J(x_J)$  where  $x_J$  stands for the vector of coordinates of x with index in J. The preimage of  $f^{-1}(U) \subseteq \mathbb{R}^{|J|}$  of any set  $U \subseteq \mathbb{R}$  can be tightly approximated by boxes if and only if  $f_J^{-1}(U)$  can because  $f^{-1}(U)$  is a Cartesian product of  $f_J^{-1}(U)$  and  $\mathbb{R}^{|J'|}$ . Since  $f_J$  is interval separable by the previous lemma, this shows that f is as well.

For closure under composition, let  $f : \mathbb{R}^m \to \mathbb{R}$  and  $f_i : \mathbb{R}^{n_i} \to \mathbb{R}$  for  $i \in \{1, ..., m\}$ , all in  $\mathcal{F}^*_{subm}$ . The composition  $g := f \circ (f_1 \times \cdots \times f_m)$  is clearly  $C^1$  again, so we just have to check the property of the partial derivatives. By the chain rule, the partial derivatives of the composition are

 $\partial_i g(x_1, \ldots, x_m) = \partial_j f(f_1(x_1), \ldots, f_m(x_m)) \partial_k f_j(x_{j1}, \ldots, x_{jn_j})$ for some  $j \in \{1, \ldots, m\}$  and  $k \in \{1, \ldots, n_j\}$ , and for all  $x_1 \in \mathbb{R}^{n_1}, \ldots, x_m \in \mathbb{R}^{n_m}$ . So if this partial derivative is zero, there are two cases. First, if  $\partial_j f(f_1(x_1), \ldots, f_m(x_m)) = 0$  then fmust be constant in its *j*-th argument (because  $f \in \mathcal{F}^*_{subm}$ ) and thus *g* is constant in  $x_j$ , and in particular the *i*-th argument (which is an entry of  $x_j$ ). Second, if  $\partial_k f_j(x_{j1}, \ldots, x_{jn_j}) = 0$ then  $f_j$  must be constant in its *k*-th argument (because  $f_j \in \mathcal{F}^*_{subm}$ ) and thus *g* is constant in its corresponding *i*-th argument as well. This proves  $g \in \mathcal{F}^*_{subm}$ , as desired.  $\Box$ 

Note that exp, sinh, arctan, *n*-th roots for *n* odd, and all linear functions are in  $\mathcal{F}^*_{subm}$ . So this is a useful set of primitive functions already. Unfortunately, it does not include

multiplication because the gradient of  $(x, y) \mapsto xy$  is zero at (0, 0). To fix this issue, we need to restrict the domain.<sup>15</sup>

For simplicity, we required primitive functions to be defined on all of  $\mathbb{R}^n$ . Suppose we allow for primitive functions to be defined only on an open subset of  $\mathbb{R}^n$ , and applying them to a value outside their domain is disallowed in SPCF programs. Then we can also include multiplication (on  $\mathbb{R}^2 \setminus \{(0,0)\}$ ), logarithms (on  $(0,\infty)$ ), non-constant univariate polynomials (on the complement of their stationary points), quantile functions of continuous distributions with nonzero density (on (0, 1)), and probability density functions (on the complement of their stationary points). The fact that some points in the domain are missing is inconvenient for functions that can be continuously extended to the these points, but one can work around this in a program by checking for the points that are not in the domain and returning the function values as constants for those cases.

# C.4 Completeness Proof

This section uses the definitions of *boxwise continuity* and *interval separability* from Section 4.2. As discussed there, we assume that for each symbolic path  $\Psi = (\mathcal{V}, n, \Delta, \Xi)$ , we have that  $\mathcal{V}$ , each C with  $C \bowtie 0 \in \Delta$ , and each  $\mathcal{W} \in \Xi$  contains each sample variable  $\alpha_i$  at most once (Assumption 1). We also assume that the primitive functions are boxwise continuous, interval separable and closed under composition (Assumption 2). Furthermore, we say that  $\mathcal{T}$  is a *subdivision* of *t* if  $\mathcal{T}$  is compatible and  $\bigcup_{t' \in \mathcal{T}} (|t'|) = ||t|$ ).

**Theorem 4.3** (Completeness of interval approximations). Let  $I \in \mathbb{I}$  and  $\vdash P : \mathbf{R}$  be an almost surely terminating program satisfying the two assumptions discussed above. Then, for all  $\epsilon > 0$ , there is a countable set of interval traces  $\mathcal{T} \subseteq \mathbb{T}_{\mathbb{I}}$  that is compatible and exhaustive such that

upperBd<sup> $\mathcal{T}$ </sup><sub>*P*</sub>(*I*) -  $\epsilon \leq [\![P]\!](I) \leq \text{lowerBd}_{P}^{\mathcal{T}}(I) + \epsilon.$ 

*Proof.* First, we give a brief outline of how the proof works. The idea is to cover  $\{s \in \mathbb{T} \mid val_P(s) \in I\}$  using boxes (interval traces). We can achieve this using symbolic execution: for a fixed path through the program, the result value is just a composition of primitive functions applied to the samples. Similarly, the weight function is a product of such functions, hence boxwise continuous. By passing to smaller boxes, we can assume that it is continuous on each box. In order to approximate the integral of the weight function, we use Riemann sums (as used in the definition of the Riemann

integral). We partition the domain into smaller and smaller boxes such that the lower bound and the upper bound of the weight function come arbitrarily close (by continuity). Then by properties of the Riemann integral, the bounds arising from the interval traces representing the boxes in this partition converge to the desired integral of the weight function. The details of the proof are as follows.

**Step 1:** approximating the branching inequalities. Let  $\Psi = (\mathcal{V}, n, \Delta, \Xi)$  a symbolic path of *P*. To find a countable set  $\mathcal{T}_{\Psi} \subseteq \mathbb{T}_{\mathbb{I}}$  such that  $\bigcup_{t \in \mathcal{T}_{\Psi}} \|t\| \in \operatorname{Sat}_{n}(\Delta)$ . Note that  $\operatorname{Sat}_{n}(\Delta)$  is a finite intersection of sets of the form  $\{s \in [0, 1]^{n} \mid C[s/\overline{\alpha}] \bowtie 0\}$  where  $\bowtie \in \{\leq, >\}$ . In the  $\leq$  case, we can write this constraint as  $C[s/\overline{\alpha}] \in \bigcup_{n \in \mathbb{N}} [-n, 0]$  and in the > case as  $C[s/\overline{\alpha}] \in \bigcup_{n \in \mathbb{N}} [1/n, n]$ . By applying Lemma C.8 to each of the compact intervals in these unions, we obtain a countable union of boxes that is a tight subset of  $\{s \in [0, 1]^{n} \mid C[s/\overline{\alpha}] \bowtie 0\}$ . Since the intersection of two boxes is a box and since  $\operatorname{Sat}_{n}(\Delta)$  is a finite intersection of such countable unions of boxes, it can be rewritten as a countable union of boxes. This yields  $\mathcal{T}_{\Psi}$ , such that  $\bigcup_{t \in \mathcal{T}_{\Psi}} [t] \in \operatorname{Sat}_{n}(\Delta)$ .

**Step 2:** handling the result value. By applying Lemma C.8 and intersecting the obtained interval traces with  $\mathcal{T}_{\Psi}$ , we obtain a countable set  $\mathcal{T}'_{\Psi,I} \subseteq \mathbb{T}_{\mathbb{I}}$  such that  $\bigcup_{t \in \mathcal{T}'_{\Psi,I}} |t| \cong \{s \in \text{Sat}(\Delta) \mid \mathcal{V}[s/\overline{\alpha}] \in I\}$ . By the same lemma, we find a countable set  $\mathcal{T}'_{\Psi,I^c} \subseteq \mathbb{T}_{\mathbb{I}}$  such that  $\bigcup_{t \in \mathcal{T}'_{\Psi,I^c}} |t| \cong \{s \in \text{Sat}(\Delta) \mid \mathcal{V}[s/\overline{\alpha}] \notin I\}$  because the complement of I can be written as a countable union of intervals. By Lemma C.10, we find subdivisions  $\mathcal{T}_{\Psi,I}$  and  $\mathcal{T}_{\Psi,I^c}$  that even satisfy  $\bigcup_{t \in \mathcal{T}_{\Psi,I}} |t| \cong$ val $_P^{-1}(I) \cap \text{Sat}(\Delta)$  and  $\bigcup_{t \in \mathcal{T}_{\Psi,I^c}} |t| \cong \text{val}_P^{-1}(\mathbb{R} \setminus I) \cap \text{Sat}(\Delta)$ . By Lemma C.7, we can assume that the interval traces  $\mathcal{T}_{\Psi,I}$ are almost disjoint. Because of the almost sure termination assumption, the set of traces  $\bigcup_{(\mathcal{V},n,\Delta,\Xi)} \text{Sat}(\Delta) = \text{val}_P^{-1}(\mathbb{R})$ where the union ranges over all symbolic paths of P has measure 1. As a consequence,  $\bigcup_{\Psi \in symPaths(P,0,0,0)} (\mathcal{T}_{\Psi,I} \cup \mathcal{T}_{\Psi,I^c})$ is a compatible and exhaustive set of interval traces. Now

$$\begin{split} \llbracket P \rrbracket(I) &= \sum_{(\mathcal{V}, n, \Delta, \Xi)} \llbracket (\mathcal{V}, n, \Delta, \Xi) \rrbracket(I) \\ &= \sum_{(\mathcal{V}, n, \Delta, \Xi)} \int_{\operatorname{Sat}_n(\Delta)} [\mathcal{V}[\boldsymbol{s}/\overline{\alpha}] \in I] \prod_{\mathcal{W} \in \Xi} \mathcal{W}[\boldsymbol{s}/\overline{\alpha}] \, \mathrm{d}\boldsymbol{s} \\ &= \sum_{(\mathcal{V}, n, \Delta, \Xi)} \sum_{\boldsymbol{t} \in \mathcal{T}_{(\mathcal{V}, n, \Delta, \Xi), I}} \int_{([t])} \prod_{\mathcal{W} \in \Xi} \mathcal{W}[\boldsymbol{s}/\overline{\alpha}] \, \mathrm{d}\boldsymbol{s} \end{split}$$

where the outer sum ranges over the symbolic paths  $(\mathcal{V}, n, \Delta, \Xi) \in symPaths(P, 0, \emptyset, \emptyset)$ . The first equality holds by Theorem 6.1, the second one by Lemma B.1 and the last one by the construction of  $\mathcal{T}_{(\mathcal{V},n,\Delta,\Xi),I}$ .

Step 3: approximating the weight function. Let

$$\mathcal{T}' := \bigcup_{\Psi \in symPaths(P,0,\emptyset,\emptyset)} \mathcal{T}_{\Psi,I}.$$

For each  $t \in \mathcal{T}'$ , fix some  $\epsilon_t > 0$ , such that  $\sum_{t \in \mathcal{T}'} \epsilon_t = \epsilon$ . This can be achieved, for example, by enumerating  $\mathcal{T}'$  as  $t^{(1)}, t^{(2)}, \ldots$  and choosing  $\epsilon_{t^{(i)}} = 2^{-i}\epsilon$ . By Lemma C.9, we

<sup>&</sup>lt;sup>15</sup>Handling these issues at the level of primitive functions directly (without restricting the domain) seems challenging: even if a function has only one point with zero gradient, e.g. multiplication, its preimage under other primitive functions can become very complicated. We tried to handle this by allowing the primitive functions to be submersions except on a null set given by a union of lower-dimensional manifolds. However, the preimages of such manifolds need not be manifolds again. Hence it seems difficult to come up with a broader class of primitive functions satisfying the assumptions without restricting the domain.

can find for each  $t \in \mathcal{T}'$  a countable set  $S_t$  of interval traces such that:

$$\sum_{t' \in S_t} \operatorname{vol}(t') \sup_{s \in (t')} \prod_{W \in \Xi} W[s/\overline{\alpha}] - \epsilon_t/2$$

$$\leq \int_{(t)} \prod_{W \in \Xi} W[s/\overline{\alpha}] \, \mathrm{d}s$$

$$\leq \sum_{t' \in S_t} \operatorname{vol}(t') \min_{s \in (t')} \prod_{W \in \Xi} W[s/\overline{\alpha}] + \epsilon_t/2$$

Next, choose  $\epsilon_{t'} > 0$  for each t' in such a way that  $\sum_{t' \in S'_t} \epsilon_{t'} < \epsilon_t/2$ . By Lemma C.11, we can find for each t' a finite subdivision  $S'_{t'}$  such that for all  $t'' \in S'_{t'}$ , we have

$$\sup_{s \in (t'')} \operatorname{wt}_P(s) = \sup_{s \in (t'')} \prod_{W \in \Xi} \mathcal{W}[s/\overline{\alpha}] \ge \sup \operatorname{wt}_P^{\mathbb{I}}(t'') - \epsilon_{t'}$$
$$\min_{s \in (t'')} \operatorname{wt}_P(s) = \min_{s \in (t'')} \prod_{W \in \Xi} \mathcal{W}[s/\overline{\alpha}] \le \min \operatorname{wt}_P^{\mathbb{I}}(t'') + \epsilon_{t'}.$$

Multiplying by vol(t'') and summing over all t'', we find together with the previous inequality

$$\begin{split} &\sum_{t'\in\mathcal{S}_{t}}\sum_{t''\in\mathcal{S}_{t'}}\operatorname{vol}(t'')\operatorname{sup}\operatorname{wt}_{P}^{\mathbb{I}}(t'') - \epsilon_{t} \\ &\leq \sum_{t'\in\mathcal{S}_{t}}\sum_{t''\in\mathcal{S}_{t'}'}\operatorname{vol}(t'')(\operatorname{sup}\operatorname{wt}_{P}^{\mathbb{I}}(t'') - \epsilon_{t'}) - \epsilon_{t}/2 \\ &\leq \int_{\left\{t\right\}}\prod_{W\in\Xi}\mathcal{W}[s/\overline{\alpha}]\,\mathrm{d}s \\ &\leq \sum_{t'\in\mathcal{S}_{t}}\sum_{t''\in\mathcal{S}_{t'}'}\operatorname{vol}(t'')(\operatorname{min}\operatorname{wt}_{P}^{\mathbb{I}}(t'') + \epsilon_{t'}) + \epsilon_{t}/2 \\ &\leq \sum_{t'\in\mathcal{S}_{t}}\sum_{t''\in\mathcal{S}_{t'}'}\operatorname{vol}(t'')\min\operatorname{wt}_{P}^{\mathbb{I}}(t'') + \epsilon_{t} \end{split}$$

because  $\sum_{t'' \in S'_{t'}} \operatorname{vol}(t'') \leq 1$  and thus the contribution of all the  $\epsilon_{t'}$  is at most  $\epsilon_t/2$ .

Overall, the desired trace set is given by

$$\mathcal{T} := \bigcup_{\Psi = (\mathcal{V}, n, \Delta, \Xi)} \left( \mathcal{T}_{\Psi, I^c} \cup \bigcup_{t \in \mathcal{T}_{\Psi, I}} \bigcup_{t' \in \mathcal{S}_t} \mathcal{S}'_{t'} \right)$$

is compatible and exhaustive because it is a subdivision of  $\mathcal{T}_{\Psi,I^c}$  and  $\mathcal{T}_{\Psi,I}$ . By construction, we have  $\operatorname{val}_P^{\mathbb{I}}(t) \subseteq I$  for  $t \in \mathcal{T}_{\Psi,I}$  and  $\operatorname{val}_P^{\mathbb{I}}(t) \cap I = \emptyset$  for  $t \in \mathcal{T}_{\Psi,I^c}$ . Hence the  $\mathcal{T}_{\Psi,I^c}$ -summands vanish in the sum for the bounds and we obtain

$$upperBd'_{P}(I) - \epsilon$$

$$= \sum_{t \in \mathcal{T}} vol(t)(\sup wt_{P}^{\mathbb{I}}(t))[val_{P}^{\mathbb{I}}(t) \cap I \neq \emptyset] - \epsilon$$

$$= \sum_{\Psi = (\mathcal{W}, n, \Delta, \Xi)} \sum_{t \in \mathcal{T}_{\Psi, I}} vol(t'') \sup wt_{P}^{\mathbb{I}}(t'') - \epsilon_{t}$$

$$\leq \sum_{\Psi = (\mathcal{W}, n, \Delta, \Xi)} \sum_{t \in \mathcal{T}_{\Psi, I}} \int_{(t)} \prod_{\mathcal{W} \in \Xi} \mathcal{W}[s/\overline{\alpha}] ds$$

$$\leq \sum_{\Psi = (\Psi, n, \Delta, \Xi)} \sum_{t \in \mathcal{T}_{\Psi, I}} \left( \sum_{t' \in \mathcal{S}_t} \sum_{t'' \in \mathcal{S}'_{t'}} \operatorname{vol}(t'') \min \operatorname{wt}_P^{\mathbb{I}}(t'') + \epsilon_t \right)$$
$$= \sum_{t \in \mathcal{T}} \operatorname{vol}(t) (\min \operatorname{wt}_P^{\mathbb{I}}(t)) [\operatorname{val}_P^{\mathbb{I}}(t) \subseteq I] + \epsilon$$
$$= \operatorname{lowerBd}_P^{\mathcal{T}}(I) + \epsilon. \qquad \Box$$

**Lemma C.7.** Given a countable set of interval traces  $\mathcal{T} \subseteq \mathbb{I}^n$ , there is a countable set of interval traces  $\mathcal{T}' \subseteq \mathbb{I}^n$  that is compatible and satisfies  $\bigcup_{t \in \mathcal{T}} \langle t \rangle = \bigcup_{t \in \mathcal{T}'} \langle t \rangle$ .

*Proof.* Let  $A : \mathbb{N} \to \mathcal{T}$  be an enumeration. Define  $A' : \mathbb{N} \to \Sigma_{\mathbb{R}^n}$  by  $m \mapsto \overline{A(m)} \setminus \bigcup_{i=0}^{m-1} \overline{A(i)}$  where  $\overline{S}$  denotes the closure of *S*. Then the collection  $\{A'(m) \mid m \in \mathbb{N}\}$  is pairwise almost disjoint, and each A'(m) can be written as a finite union of boxes, proving the claim.

**Lemma C.8.** Let  $\mathcal{V}$  a symbolic value of ground type containing each sample variable  $\alpha_1, \ldots, \alpha_n$  at most once and [x, y]an interval. Then there is a countable set of pairwise disjoint interval traces  $\mathcal{T} \subset \mathbb{I}^n_{[0,1]}$  such that

$$\bigcup_{t\in\mathcal{T}} \langle t \rangle \Subset \{ s \in [0,1]^n \mid \mathcal{V}[s/\overline{\alpha}] \in [x,y] \}.$$

*Proof.* If  $\mathcal{V}$  is of ground type, then it is simply a composition of primitive functions applied to sample variables and literals. Since the set of primitive functions is closed under composition and since no sample variable occurs twice, this composition is still an interval separable function f of the sample variables. By definition of interval separability, there is a countable set of interval traces  $\mathcal{J}$  such that  $\bigcup_{t \in \mathcal{J}} |t| \in f^{-1}([x, y])$ , as desired.

**Lemma C.9.** Let  $t \in \mathbb{I}^n$  be an interval trace and  $\Xi$  a set of symbolic values with sample variables from  $\overline{\alpha} = \alpha_1, \ldots, \alpha_n$ . Then for any  $\epsilon > 0$ , there is a countable subdivision  $\mathcal{T}$  of t such that

$$\sum_{t'\in\mathcal{T}} \operatorname{vol}(t') \sup_{s\in \langle\!\langle t'\rangle\!\rangle} \prod_{W\in\Xi} \mathcal{W}[s/\overline{\alpha}] - \epsilon$$
  
$$\leq \int_{\langle\!\langle t\rangle\!\rangle} \prod_{W\in\Xi} \mathcal{W}[s/\overline{\alpha}] \,\mathrm{d}s$$
  
$$\leq \sum_{t'\in\mathcal{T}} \operatorname{vol}(t') \min_{s\in \langle\!\langle t'\rangle\!\rangle} \prod_{W\in\Xi} \mathcal{W}[s/\overline{\alpha}] + \epsilon.$$

*Proof.* Values are simply boxwise continuous functions applied to the sample variables. Intersecting the boxes for each  $\mathcal{W} \in \Xi$ , we see that the function

$$f: (t) \to \mathbb{R}, \quad s \mapsto \prod_{W \in \Xi} \mathcal{W}[s/\overline{\alpha}]$$

 $= \llbracket P \rrbracket (I)$ 

is boxwise continuous. We can thus find a countable subdivision  $\mathcal{T}_{cont}$  of t such that f is continuous on each  $t' \in \mathcal{T}_{cont}$ .

Since we can sum over the  $t' \in \mathcal{T}_{cont}$ , it suffices to prove that each integral  $\int_{||t'||} f(s) \, ds$  can be approximated arbitrarily closely. Note that each such integral is finite because a continuous function is bounded on a compact set and the measure of (t') is finite. But then such approximations are given by Riemann sums, i.e. the sums that are used to define the Riemann integral. As a concrete example, one can consider the subdivision  $\mathcal{T}_m$  of t' in m equidistant sections in each dimension (consisting of  $m^n$  parts overall). Then  $\sum_{t'' \in \mathcal{T}_m} \operatorname{vol}(t'') \min_{s \in (t'')} f(s)$  converges to the Riemann integral  $\int_{\mathbb{R}^{t}} f(s) \, \mathrm{d}s$  as  $m \to \infty$  (and similarly for the supremum). Since it is known that the Riemann integral and the Lebesgue integral have the same value for continuous functions on a Cartesian product of compact intervals, the claim follows immediately. П

**Lemma C.10** (Relationship between symbolic execution and interval semantics). Let  $\Psi = (\Psi, n, \Delta, \Xi)$  be a symbolic path of *P* and *t* an interval trace with  $||t|| \subseteq \operatorname{Sat}_n(\Delta)$ . Suppose furthermore that all the symbolic values contain each of the sample variables  $\overline{\alpha} = \alpha_1, \ldots, \alpha_n$  at most once. Then there is a subdivision  $\mathcal{T}$  of *t* such that for all  $t' \in \mathcal{T}$ , the interval semantics for the value is precise:

$$\operatorname{val}_{P}^{\mathbb{I}}(\boldsymbol{t}') = \{\operatorname{val}_{P}(\boldsymbol{s}) \mid \boldsymbol{s} \in (\boldsymbol{t}')\}.$$

For each symbolic score value  $W \in \Xi$ , let  $[W_{t'}^-, W_{t'}^+]$  be its interval approximation, i.e.  $W_{t'}^- = \min_{s \in (t')} W[s/\overline{\alpha}]$  and  $W_{t'}^+ = \sup_{s \in (t')} W[s/\overline{\alpha}]$ . Then

$$\mathsf{wt}_P^{\mathbb{I}}(t') = \left[\prod_{\mathcal{W}\in\Xi} \mathcal{W}_{t'}^-, \prod_{\mathcal{W}\in\Xi} \mathcal{W}_{t'}^+\right].$$

*Proof.* The symbolic value  $\mathcal{V}$  is a composition of primitive functions applied to  $\alpha$ 's. Hence the are boxwise continuous functions of the  $\alpha$ 's. We pick a suitable subdivision  $\mathcal{T}$  such that all these functions are continuous when restricted to any  $t' \in \mathcal{T}$ . For any such function f, we have

$$f_{\mathbb{I}}([x_1, y_1], \dots, [x_m, y_m]) = [\inf F, \sup F]$$

where  $F := f([x_1, y_1] \times \cdots \times [x_m, y_m]) \subseteq \mathbb{R}$ , by definition. Then continuity implies that the image of any box is a compact and path-connected subset of  $\mathbb{R}$ , i.e. an interval. Hence we even have  $f_{\mathbb{I}}([x_1, y_1], \dots, [x_m, y_m]) = F$ , i.e. the image of any box equals its interval approximation, proving the claim about the value semantics.

For the interval semantics of the weight, note that the previous argument applies to every symbolic score value  $W \in \Xi$ , proving that

$$\{\mathcal{W}[s/\overline{\alpha}] \mid s \in (t')\} = [\mathcal{W}_{t'}^{-}, \mathcal{W}_{t'}^{+}].$$

Since the interval semantics multiplies the interval approximation of each score value in interval arithmetic, this implies the claim.  $\hfill \Box$ 

Note that the interval approximation of the weight is imprecise in the following sense:

$$\mathsf{wt}_P^{\mathbb{I}}(t') \neq \left\{ \prod_{\mathcal{W} \in \Xi} \mathcal{W}[s/\overline{\alpha}] \, \middle| \, s \in \langle\!\!\!| t' \rangle\!\!\!\!\!\rangle \right\}$$

As an example, if  $\Xi = \{\alpha_1, 1 - \alpha_1\}$  and  $t'' = \langle [0, 1] \rangle$  then the left-hand side is [0, 1] because each of the weights is approximated by [0, 1], but the right-hand side is [0, 1/4]because the function  $\alpha_1(1 - \alpha_1)$  attains its maximum at 1/4, not 1.

**Lemma C.11.** Let  $\Psi = (\mathcal{V}, n, \Delta, \Xi)$  be a symbolic path of Pand t an interval trace with  $(t) \subseteq \operatorname{Sat}_n(\Delta)$ . Suppose furthermore that all the symbolic values contain each of the sample variables  $\overline{\alpha} = \alpha_1, \ldots, \alpha_n$  at most once. Then for all  $\epsilon > 0$ , there is a subdivision  $\mathcal{T}$  of t such that for all  $t' \in \mathcal{T}$ , we have  $\min_{s \in (t')} \operatorname{wt}_P(s) \leq \min \operatorname{wt}_P^{\mathbb{I}}(t') + \epsilon$  and  $\sup_{s \in (t')} \operatorname{wt}_P(s) \geq$  $\sup \operatorname{wt}_P^{\mathbb{I}}(t') - \epsilon$ .

*Proof.* Since for each  $W \in \Xi$ , the function  $f : (|t|) \to \mathbb{R}$ ,  $s \mapsto W[s/\overline{\alpha}]$  is boxwise continuous (a property of primitive functions), we can find a countable subdivision  $\mathcal{T}'$  of t, such that for all  $t' \in \mathcal{T}'$ , f is continuous on (|t'|). Hence it suffices to prove the statement for each t'.

Since (t') is compact (because it's closed and bounded), *f* attains a maximum  $W < \infty$  on (t') and is even uniformly continuous on *t'*. Hence there is a  $\delta > 0$  such that whenever  $||s - s'|| < \delta$  then  $|f(s) - f(s')| < \epsilon' := \frac{\epsilon}{W^{|\Xi|-1}}$ .

Let  $\mathcal{T}$  be a subdivision where each interval trace  $t \in \mathcal{T}$  has diameter less than  $\delta$ . For  $t \in \mathcal{T}$  and  $\mathcal{W} \in \Xi$ , let  $\mathcal{W}_t^- := \min_{s \in \{\!\!\!\ t \ \!\!\!\}} \mathcal{W}[s/\alpha]$  and  $\mathcal{W}_t^+ := \sup_{s \in \{\!\!\!\ t \ \!\!\}} \mathcal{W}[s/\alpha]$ . By the choice of  $\mathcal{T}$ , we have  $\mathcal{W}_t^+ \leq \mathcal{W}_t^- + \epsilon'$  for  $t \in \mathcal{T}$ . By Lemma C.10, we find that  $\sup \mathsf{wt}_P^{\mathbb{I}}(t) = \prod_{\mathcal{W} \in \Xi} \mathcal{W}_t^+$  and  $\min \mathsf{wt}_P^{\mathbb{I}}(t) = \prod_{\mathcal{W} \in \Xi} \mathcal{W}_t^-$ . As a consequence, we have

$$\sup \mathsf{wt}_P^{\mathbb{I}}(t) - \min \mathsf{wt}_P^{\mathbb{I}}(t) = \prod_{\mathcal{W} \in \Xi} \mathcal{W}_t^+ - \prod_{\mathcal{W} \in \Xi} \mathcal{W}_t^-$$
$$< \prod_{\mathcal{W} \in \Xi} (\mathcal{W}_t^- + \epsilon') - \prod_{\mathcal{W} \in \Xi} \mathcal{W}_t^-$$
$$< \epsilon' \mathcal{W}^{|\Xi|-1} = \epsilon$$

So the interval  $\operatorname{wt}_{P}^{\mathbb{I}}(t)$  has diameter less than  $\epsilon$ . Since the interval  $\{\operatorname{wt}_{P}(s) \mid s \in (t')\}$  is contained in it (by soundness), the claim follows.

**Corollary 4.4.** Let  $I \in \mathbb{I}$  and  $\vdash P : \mathbf{R}$  be as in Theorem 4.3. There is a sequence of finite, compatible sets of interval traces  $\mathcal{T}_1, \mathcal{T}_2, \ldots \subseteq \mathbb{T}_{\mathbb{I}}$  s.t.  $\lim_{n\to\infty} \text{lowerBd}_p^{\mathcal{T}_n}(I) = \llbracket P \rrbracket(I)$ .

*Proof.* By Theorem 4.3, we can find for each  $n \in \mathbb{N}$  a set of interval traces  $\mathcal{T}'_n$  such that lower  $\operatorname{Bd}_P^{\mathcal{T}'_n}(I) > \llbracket P \rrbracket(I) - 1/n$ . Since the lower bound is defined as a sum over  $\mathcal{T}'_n$ , there is a finite subset  $\mathcal{T}_n$  such that lower  $\operatorname{Bd}_P^{\mathcal{T}_n}(I) > \llbracket P \rrbracket(I) - 2/n$ . Since  $\mathcal{T}'_n$  is still compatible, the soundness result yields lower  $\operatorname{Bd}_P^{\mathcal{T}_n}(I) \leq \llbracket P \rrbracket(I)$ , implying the claim.  $\Box$ 

Finitely many interval traces are not enough for complete upper bounds, if the weight function is unbounded. This issue arises even if we can compute the tightest possible bounds on the weight function, as the following program illustrates.

**Example C.3.** Consider the following probabilistic program expressed in pseudocode.

threshold := 1  
while (sample 
$$\leq$$
 threshold) do  
threshold :=  $\frac{threshold}{2}$   
score(2)

The program only requires addition and scalar multiplication. It can even be implemented using call-by-name (CbN) semantics (which allows each sampled value to be used at most once). For example in SPCF we can write

$$P \equiv (\mu_s^{\varphi}, \text{ if } (\text{sample} - s, \text{score}(2); \varphi(s/2), 1)) 1$$

The program *P* has the weight function

$$\mathsf{wt}_P(\langle t_0, \dots, t_n \rangle) = \begin{cases} 2^n & \text{if } t_n > 2^{-n} \land \\ & \forall i \in \{0, \dots, n-1\} : t_i \le 2^{-i} \\ 0 & \text{otherwise.} \end{cases}$$

*P* is integrable because the normalizing constant is

$$Z = \int_{\mathbb{T}} \operatorname{wt}_{P}(\mathbf{s}) \, \mathrm{d}\mathbf{s} = \sum_{n=1}^{\infty} 2^{n} \times (1 - 2^{-n}) \prod_{i=0}^{n-1} 2^{-i}$$
$$= \sum_{n=1}^{\infty} 2^{n} (1 - 2^{-n}) 2^{-n(n-1)/2} < \infty.$$

We claim that *P* requires infinitely many interval traces for the upper bound to converge to the true denotation. Define the sets of traces  $T_n$  for  $n \ge 1$  by

$$T_n = [0, 2^0] \times [0, 2^{-1}] \times \cdots \times [0, 2^{-n+1}] \times (2^{-n}, 1].$$

Suppose we are given an arbitrary finite exhaustive set of interval traces. This set needs to cover all of the  $T_n$ 's, so one interval trace, say t, must cover infinitely many  $T_n$ 's. Since wt<sub>*P*</sub>( $\mathbf{s}$ ) = 2<sup>*n*</sup> for  $\mathbf{s} \in T_n$ , the weight function on t is unbounded. Therefore, the only possible upper bound for *t* is  $\infty$ , even if our semantics could compute the set {wt<sub>P</sub>(s) |  $s \in (|t|)$  exactly. Hence any finite exhaustive interval trace has upper bound  $\infty$ , while the true denotation is finite. As we have seen, this is not because of imprecision of interval analysis, but an inherent problem if the weight function is unbounded. So we cannot hope for complete upper bounds with finitely many interval traces.

#### D Supplementary Material for Section 5

We provide additional proofs and material for Section 5. To have access to named rules, we give the type system in Fig. 9 which agrees with the one in Fig. 4 in everything but the labels.

#### **D.1 Soundness**

To show soundness (Theorem 5.1), we establish a (weightaware) subject reduction property for our type system as follows. For an interval  $[a, b] \in \mathbb{I}$  and  $r \in \mathbb{R}_{>0}$ , we define  $r \cdot [a, b] := [r \cdot a, r \cdot b]$ . To simplify notation, we use a modified transition relation that omits the concrete trace (which is irrelevant in Theorem 5.1). We write  $P \rightarrow_{w} P'$  if  $(P, \mathbf{s}, 1) \rightarrow (P', \mathbf{s}', w)$  for some  $w \in \mathbb{R}$  and  $\mathbf{s}, \mathbf{s}' \in \mathbb{T}$ . Note that we could define  $\rightarrow_w$  as a dedicated reduction system by adapting the rules from  $\rightarrow$  in Fig. 2.

**Lemma D.1** (Substitution). If  $\Gamma$ ;  $\{x_i : \sigma_i\}_{i=1}^n \vdash P : \mathcal{A}$  for distinct variables  $x_i$  and if  $\Gamma \vdash M_i : {\sigma_i \\ 1}$  for all  $i \in \{1, ..., n\}$ then  $\Gamma \vdash P[M_i/x_i]_{i=1}^n : \mathcal{A}$ .

*Proof.* By a standard induction on *M*.

Lemma D.2 (Weighted Subject Reduction). Let P be any program such that  $\vdash P : \begin{cases} \sigma \\ J \end{cases}$  and  $P \rightarrow_w P'$  for some w > 0. Then  $\vdash P': \left\{ \begin{matrix} \sigma \\ \frac{1}{2m} \cdot J \end{matrix} \right\}.$ 

*Proof.* We prove this by induction on the structure of *P*.

**Case** P = sample: then  $P \rightarrow_1 \underline{r}$  for some  $r \in [0, 1]$ . As multiple consecutive application of (Sub) can be replaced by a single one since (as  $\sqsubseteq_{\mathcal{R}}$  is transitive), we can assume w.l.o.g. that the last step in  $\vdash P : \begin{cases} \sigma \\ I \end{cases}$  was:

$$\frac{}{\vdash \text{ sample} : \left\{ \begin{bmatrix} [0, 1] \\ 1 \end{bmatrix} \right\}} \text{ (SAMPLE)}$$
$$\frac{}{\vdash \text{ sample} : \left\{ \begin{matrix} I \\ J \end{matrix} \right\}} \text{ (SUB)}$$

By subtyping we have  $[0, 1] \sqsubseteq I$  and  $1 \sqsubseteq J$ . So  $[r, r] \sqsubseteq I$  and we can type  $\vdash \underline{r} : {I \atop J}$  using (LIT) and (SUB) as required. **Case**  $P = f(\underline{r_1}, \dots, \underline{r_{|f|}})$ : then  $P \rightarrow_1 \underline{f(r_1, \dots, r_{|f|})}$ . W.l.o.g.,

we can assume that the last step in  $\vdash P : \begin{cases} \sigma \\ I \end{cases}$  was

$$\frac{\overline{r_{1}:\left\{\begin{bmatrix}r_{1},r_{1}\end{bmatrix}\right\}}}{\frac{\vdash \underline{r_{1}:\left\{I_{1}\right\}}}{J_{1}}} (SUB)} \xrightarrow{\left\{\vdash \underline{r_{|f|}:\left\{I_{|f|},r_{|f|}\right\}}}{\frac{\vdash \underline{r_{|f|}:\left\{I_{|f|}\right\}}}{J_{|f|}}} (SUB)} (SUB)$$

$$\frac{\vdash \underline{r_{|f|}:\left\{I_{|f|}\right\}}}{\frac{\vdash f(\underline{r_{1}},\ldots,\underline{r_{|f|}}):\left\{f^{I}(I_{1},\ldots,I_{|f|})\right\}}{(X^{T})_{i=1}^{I_{|f|}}J_{i}}} (PRIM)$$

$$\frac{\vdash f(\underline{r_{1}},\ldots,\underline{r_{|f|}}):\left\{I_{|f|}\right\}}{\vdash f(\underline{r_{1}},\ldots,\underline{r_{|f|}}):\left\{I_{|f|}\right\}} (SUB)}$$

So by subtyping  $r_i \in I_i$  and  $1 \in J_i$  for all *i*. By definition of  $f^{\mathbb{I}}$ we thus have  $f(r_1, \ldots, r_{|f|}) \in f^{\mathbb{I}}(I_1, \ldots, I_{|f|})$  and (again by subtyping) we have  $f(r_1, \ldots, r_{|f|}) \in I$ . Similarly, by definition of  $\times^{\mathbb{I}}$  we have  $1 \in (\times^{\mathbb{I}})_{i=1}^{|f|} J_i$  and thus  $1 \in J$ . We can type  $\vdash \frac{f(r_1, \dots, r_{|f|}) : {I \choose J} \text{ using (LIT) and (SUB) as required.}}{\mathbf{Case } P = if(\underline{r}, M, N) \text{ and } r \leq 0: \text{ then } P \to_1 M. \text{ W.l.o.g.,}}$ 

we can assume that the last step in  $\vdash P : \begin{cases} \sigma \\ I \end{cases}$  is

$$\frac{x:\sigma\in\Gamma}{\Gamma+x:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{VAR}\right) \qquad \frac{\Gamma\in\mathcal{M}:\mathcal{R}\to\mathcal{R}\subseteq\mathcal{R}}{\Gamma\in\mathcal{M}:\mathcal{B}}\left(\operatorname{SUB}\right) \qquad \frac{\Gamma;x:\sigma\in\mathcal{M}:\mathcal{R}}{\Gamma+\lambda x.M:\left\{\begin{matrix}\sigma\to\mathcal{R}\\1\end{matrix}\right\}}\left(\operatorname{ABS}\right) \qquad \frac{\Gamma;\varphi:\sigma\to\mathcal{R};x:\sigma\in\mathcal{M}:\mathcal{R}}{\Gamma+\mu_{x}^{\varphi},M:\left\{\begin{matrix}\sigma\to\mathcal{R}\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+\mu_{x}^{\varphi},M:\left\{\begin{matrix}\sigma\to\mathcal{R}\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right) \qquad \frac{\Gamma\in\mathcal{R}:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}{\Gamma+\mu_{x}^{\varphi},M:\left\{\begin{matrix}\sigma\to\mathcal{R}\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+\mu_{x}^{\varphi},M:\left\{\begin{matrix}\sigma\to\mathcal{R}\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right) \qquad \frac{\Gamma\in\mathcal{R}:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right) \qquad \frac{\Gamma\in\mathcal{R}:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right) \qquad \frac{\Gamma\in\mathcal{R}:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right) \qquad \frac{\Gamma\in\mathcal{R}:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fix}\right)}{\Gamma+M:\left\{\begin{matrix}\sigma\\1\end{matrix}\right\}}\left(\operatorname{Fi$$

Figure 9. Weight-aware interval type system for SPCF with typing rule names. The rules agree with those in Fig. 4.

$$\frac{\overline{r}: \left\{ \begin{bmatrix} r, r \\ 1 \end{bmatrix} \right\}}{\underbrace{F: \left\{ \begin{smallmatrix} I \\ J'' \end{smallmatrix} \right\}}_{I''}} (SUB)} + M: \left\{ \begin{smallmatrix} \sigma' \\ J' \end{smallmatrix} \right\} + N: \left\{ \begin{smallmatrix} \sigma' \\ J' \end{smallmatrix} \right\}} (IF)$$

$$\frac{\underbrace{F: \left\{ \begin{smallmatrix} I \\ J'' \end{smallmatrix} \right\}}_{I''} (IF)}{\underbrace{F: \left\{ \begin{smallmatrix} r, M, N \end{smallmatrix} \right\}: \left\{ \begin{smallmatrix} \sigma' \\ J' \times^{\mathbb{I}} J'' \end{smallmatrix} \right\}}_{F: if(\underline{r}, M, N): \left\{ \begin{smallmatrix} \sigma' \\ J' \times^{\mathbb{I}} J'' \end{smallmatrix} \right\}} (SUB)$$

and we have  $1 \in J'', \sigma' \sqsubseteq_{\sigma} \sigma$  and  $J' \times^{\mathbb{I}} J'' \sqsubseteq J$  by subtyping. As  $1 \in J''$ , we get  $J' \sqsubseteq J' \times^{\mathbb{I}} J''$ . Thus we obtain  $J' \sqsubseteq J$  and  $\sigma' \sqsubseteq_{\sigma} \sigma$ , and we can type  $\vdash M : \begin{cases} \sigma \\ J \end{cases}$  using (SUB).

**Case**  $P = if(\underline{r}, M, N)$  and r > 0: then  $P \rightarrow_1 N$ . Analogous to the previous case.

**Case**  $P = \text{score}(\underline{r})$  and  $r \ge 0$ : then  $P \rightarrow_r \underline{r}$ . W.l.o.g., we can assume that the last step in  $\vdash P : \begin{cases} \sigma \\ I \end{cases}$  is



By subtyping we have  $r \in I'$  and even  $r \in I' \sqcap [0, \infty]$  because  $r \ge 0$ . Thus  $r \in I$ , again by subtyping. Similarly,  $1 \in J'$  and by definition of  $\times^{\mathbb{I}}$ , we have  $r \in J' \times^{\mathbb{I}} (I' \sqcap [0, \infty])$ . Thus  $r \in J$  by subtyping. This already implies  $1 \in \frac{1}{r} \cdot J$  and we can thus type  $\vdash \underline{r} : \left\{ \frac{I}{r} \cdot J \right\}$  by using (LIT) and (SUB).

**Case**  $P = (\lambda x.M)V$ : then  $P \rightarrow_1 M[V/x]$ . W.l.o.g., the last step in  $\vdash P : \begin{cases} \sigma \\ J \end{cases}$  is

$$\frac{\left\{x:\tilde{\sigma}'\right\} \vdash M: \left\{\tilde{\tilde{\sigma}}''\right\}}{\vdash \lambda x.M: \left\{\tilde{\sigma}' \rightarrow \left[\tilde{\tilde{\sigma}}''\right]^{*}\right\}} (ABS)} \\
\frac{\vdash \lambda x.M: \left\{\tilde{\sigma}' \rightarrow \left[\tilde{\tilde{\sigma}}''\right]^{*}\right\}}{J'} (SUB)} \\
\frac{\vdash \lambda x.M: \left\{\sigma' \rightarrow \left[\tilde{\sigma}''\right]^{*}\right\}}{J'} \vdash V: \left\{\tilde{\sigma}''\right\}} \\
\frac{\vdash (\lambda x.M)V: \left\{J' \times^{\mathbb{I}} J'' \times^{\mathbb{I}} J'''\right\}}{\vdash (\lambda x.M)V: \left\{\sigma'\right\}} (SUB)$$

By subtyping we get that  $\sigma' \sqsubseteq_{\sigma} \tilde{\sigma}'$ . It is easy to see that we can also type  $\vdash V : \begin{cases} \sigma' \\ 1 \end{cases}$  because V is a value and we get  $\vdash V : \begin{cases} \tilde{\sigma}' \\ 1 \end{cases}$  by (SUB). Using Lemma D.1, we can thus type  $\vdash M[V/x] : \{ \tilde{\sigma}'' \\ \tilde{I}'' \}$ .

We have  $1 \in J'$  by subyptying and as V is a value, it is easy to see that  $1 \in J'''$ . Hence  $J'' \sqsubseteq J' \times^{\mathbb{I}} J'' \times^{\mathbb{I}} J'''$ . Also by subtyping, we find  $\tilde{\sigma}'' \sqsubseteq_{\sigma} \sigma'' \sqsubseteq_{\sigma} \sigma$ ,  $\tilde{J}'' \sqsubseteq J''$ , and  $J' \times^{\mathbb{I}} J'' \times^{\mathbb{I}} J''' \sqsubseteq J$ . This implies  $\tilde{\sigma}'' \sqsubseteq_{\sigma} \sigma$  and  $\tilde{J}'' \sqsubseteq J$ . By subtyping  $\vdash M[V/x] : {\sigma \atop J}$ , as required.

**Case**  $P = (\mu_x^{\varphi}.M)V$ : then  $P \to_1 M[V/x, (\mu_x^{\varphi}.M)/\varphi]$ . Analogous to the previous case for abstractions.

**Case** P = E[P'] **for an evaluation context**  $E \neq [\cdot]$ : then  $P' \rightarrow_r P''$  and  $P \rightarrow_r E[P'']$ . All such cases follow easily by case analysis on *E*. As an example, consider the context  $E = [\cdot]N$ . In this situation, we have P = E[P'] = P'N with  $P' \rightarrow_r P''$ , so  $P \rightarrow_r P''N$ . W.l.o.g., the last step in  $\vdash P : \begin{cases} \sigma \\ J \end{cases}$  is

$$\frac{\vdash P': \left\{ \begin{matrix} \sigma' \to \left[ \begin{matrix} \sigma'' \\ J' \end{matrix} \right\} \\ \hline \end{matrix} \right\} \vdash N: \left\{ \begin{matrix} \sigma'' \\ J''' \end{matrix}}{ \begin{matrix} F'N: \left\{ \begin{matrix} \sigma'' \\ J' \times^{\mathbb{I}} J'' \times^{\mathbb{I}} J''' \end{matrix} \right\} \\ \vdash P'N: \left\{ \begin{matrix} \sigma'' \\ J' \times^{\mathbb{I}} J'' \times^{\mathbb{I}} J''' \end{matrix} \right\}} (SUB)$$

By the inductive assumption for  $P' \rightarrow_r P''$ , we get  $\vdash P'' : \begin{cases} \sigma' \rightarrow \begin{pmatrix} \sigma'' \\ J'' \end{pmatrix} \\ \frac{1}{r} \cdot J' \end{pmatrix}$  and can then type

$$\frac{\vdash P'': \left\{ \begin{matrix} \sigma' \to \left\{ \begin{matrix} \sigma'' \\ \frac{1}{r} \cdot J' \end{matrix} \right\} \\ \vdash P'N: \left\{ \begin{matrix} \sigma'' \\ \frac{1}{r} \cdot J' \end{matrix} \right\} \\ \vdash P'N: \left\{ \begin{matrix} \sigma'' \\ \frac{1}{r} \cdot J' \times^{\mathbb{I}} J'' \times^{\mathbb{I}} J''' \\ \vdash P'N: \left\{ \begin{matrix} \sigma \\ \frac{1}{r} \cdot J \end{matrix} \right\} \end{matrix} (APP)$$

because if  $J' \times^{\mathbb{I}} J'' \times^{\mathbb{I}} J''' \sqsubseteq J$  then  $\frac{1}{r} \cdot J' \times^{\mathbb{I}} J'' \times^{\mathbb{I}} J''' \sqsubseteq \frac{1}{r} \cdot J$ .

The proof for the other evaluation contexts, i.e., where P = if(P', M, N), P = VP', P = score(P'), or  $P = \underline{f}(r_1, \ldots, r_{i-1}, P', N_{i+1}, \ldots, N_{|f|})$  for some  $P' \rightarrow_r P''$ , are all analogous to the above.

**Lemma D.3** (Zero-Weighted Subject Reduction). Let *P* be any program such that  $\vdash P : \begin{cases} \sigma \\ I \end{cases}$  and  $P \rightarrow_0 P'$ . Then

1. 
$$\vdash P' : \begin{pmatrix} \sigma \\ J' \end{pmatrix}$$
 for some J', and  
2.  $0 \in J$ 

*Proof.* The proof that  $\vdash P' : {\sigma \atop J'}$  for some J' is analogous to the proof of Lemma D.2 with fewer restrictions on the weight. The claim  $0 \in J$  follows by observing that  $P \rightarrow_0 P'$  is only possible if the redex in P is score( $\underline{0}$ ). In case  $P = \text{score}(\underline{0})$ , the claim follows directly from (Score). If  $P = E[\text{score}(\underline{0})]$ , it is a simple induction on the structure of the evaluation context E.

**Theorem 5.1.** Let  $\vdash P : \mathbf{R}$  be a simply-typed program. If  $\vdash P : \left\{ \begin{smallmatrix} [a,b] \\ [c,d] \end{smallmatrix} \right\}$  and  $(P, \mathbf{s}, 1) \to^* (\underline{r}, \langle \rangle, w)$  for some  $\mathbf{s} \in \mathbb{T}$  and  $r, w \in \mathbb{R}$ , then  $r \in [a,b]$  and  $w \in [c,d]$ .

Proof. Let

 $(P, \mathbf{s}, 1) = (P_0, \mathbf{s}_0, w_0) \longrightarrow \cdots \longrightarrow (P_n, \mathbf{s}_n, w_n) = (\underline{r}, \langle \rangle, w)$ 

be the reduction sequence of (P, s, 1). By definition of  $\rightarrow_w$  it is easy to see that we get

$$P = P_0 \longrightarrow_{\tilde{w}_1} P_1 \longrightarrow_{\tilde{w}_2} \cdots \longrightarrow_{\tilde{w}_n} P_n = \underline{P}_1$$

for unique  $\tilde{w}_1, \ldots, \tilde{w}_n$ . Note that  $w_i = \prod_{j=1}^i \tilde{w}_j$ .

We first assume that  $w \neq 0$  (so  $\tilde{w}_i \neq 0$  for all *i*). We claim that for each  $0 \leq i \leq n$ , we can type

$$\vdash P_i : \left\{ \begin{bmatrix} [a,b] \\ \frac{1}{w_i} \cdot [c,d] \end{bmatrix} \right\}$$
(7)

Equation (7) follows by simple induction: the base case i = 0holds by the assumption  $\vdash P_0 : {[a, b] \\ [c, d]}$  and because  $w_0 = 1$ . For the inductive case, we assume that  $\vdash P_i : {[a, b] \\ \frac{1}{w_i} \cdot [c, d]}$ . We apply Lemma D.2 to  $P_i \rightarrow_{\tilde{w}_i} P_{i+1}$  and finish the induction step using  $w_{i+1} = w_i \tilde{w}_{i+1}$ :

$$\vdash P_{i+1} : \left\{ \begin{bmatrix} [a,b] \\ \frac{1}{\tilde{w}_{i+1}} \cdot \frac{1}{w_i} \cdot [c,d] \right\} = \left\{ \begin{bmatrix} [a,b] \\ \frac{1}{w_{i+1}} \cdot [c,d] \right\}$$

Equation (7) thus implies  $\vdash \underline{r} : \left\{ \begin{bmatrix} [a, b] \\ \frac{1}{w} \cdot [c, d] \end{bmatrix} \right\}$ . W.l.o.g., we can assume that this type judgment has the form

$$\frac{\overline{\left( \begin{array}{c} \hline r \\ r \end{array} \right)}}{\left( \begin{array}{c} \hline r \\ r \end{array} \right)} \left( \begin{array}{c} \hline r \\ r \end{array} \right)} \\ + \frac{r}{\underline{r} : \left\{ \begin{array}{c} [a,b] \\ \frac{1}{w} \cdot [c,d] \right\}} \end{array} \left( \begin{array}{c} \text{SUB} \end{array} \right)}$$

This implies  $r \in [a, b]$  and  $w \in [c, d]$  (because  $1 \in \frac{1}{w} \cdot [c, d]$ ), as required.

Now consider the case w = 0, so at least one  $\tilde{w}_i = 0$ . The claim that  $r \in [a, b]$  follows easily as above (now using Item 1 of Lemma D.3 to handle the weight 0 reduction steps). Let  $i^*$  be the smallest index such that  $\tilde{w}_{i^*} = 0$ . Using the same argument as above on the (possibly empty) prefix up to index  $i^* - 1$  (where all  $\tilde{w}_i$  are non-zero) we find

$$\vdash P_{i^*-1} : \left\{ \begin{bmatrix} [a,b] \\ \frac{1}{w_{i^*-1}} \cdot [c,d] \right\}$$

Note that this is well defined because  $\tilde{w}_j > 0$  for  $j < i^*$ . Item 2 of Lemma D.3 shows  $0 \in \frac{1}{w_{i^*-1}} \cdot [c, d]$ , which already implies  $w = 0 \in [c, d]$ , as required.

## D.2 Weak Completeness

**Proposition 5.2.** Let  $\vdash P : \alpha$  be a simply-typed program. There exists a weighted interval type  $\mathcal{A}$  such that  $\vdash P : \mathcal{A}$ .

*Proof.* For every *simple* type  $\alpha$ , we define a weighted type  $\mathcal{A}_{\alpha}$  and weightless type  $\sigma_{\alpha}$  by mutual recursion as follows.

$$\sigma_{\mathbf{R}} := [-\infty, \infty]$$
$$\sigma_{\alpha \to \beta} := \sigma_{\alpha} \to \mathcal{A}_{\beta}$$
$$\mathcal{A}_{\alpha} := \left\{ \begin{matrix} \sigma_{\alpha} \\ [0, \infty] \end{matrix} \right\}$$

That is, we insert  $[-\infty, \infty]$  for values and  $[0, \infty]$  for weights in all locations. We claim that if  $\vdash P : \alpha$  in the simple type system, then  $\vdash P : \mathcal{A}_{\alpha}$  in the weight-aware interval type system. For the proof, we strengthen the induction hypothesis by claiming: if  $\Gamma \vdash P : \alpha$  in the simple type system then  $\Gamma^{\uparrow \sigma} \vdash P : \mathcal{A}_{\alpha}$  where

$$\Gamma^{\uparrow \sigma} := \{ x : \sigma_{\alpha} \mid x : \alpha \in \Gamma \}.$$

The claim can be proved by simple induction on the derivation of  $\Gamma \vdash P : \alpha$  using the respective typing rule for the interval type system possibly followed by (SUB) for typing rules that yield proper subtypes of  $\mathcal{A}_{\alpha}$ .

Note that this typing derivation does not contain any useful information to improve the precision of GuBPI.  $\hfill \Box$ 

# D.3 Constraint-based Type Inference

In this section we formalize the constraint-based type inference algorithm and sketch our constraints-solving method based on worklist and widening. The overarching idea is to substitute intervals with variables  $v_i$ , called interval variables, and to encode typability as a constraint system. As we work in the restricted interval domain (as opposed to e.g. full first-order refinements), the resulting constraints can be solved very efficiently, which is crucial to the practicality of our tool.

*Symbolic types.* Symbolic types are defined by the following grammar:

$$\kappa := v_i \mid \kappa \to \mathscr{A} \qquad \mathscr{A} := \begin{cases} \kappa \\ v_j \end{cases}$$

where  $v_i$ ,  $v_j$  are interval variables. Symbolic types are identical to interval types but use interval variables instead of intervals as first-order types and in the weight bound.

*Constraints.* Constraints on interval variables come in three forms:

$$c := v_n \equiv [a, b] \mid \kappa_1 \sqsubseteq \kappa_2 \mid v_n \equiv \underline{f}(v_{n_1}, \dots, v_{n_{|f|}})$$

where *f* is a primitive function and [a, b] an interval. That is, a constraint can either equate an interval variable to a particular interval, require a symbolic type  $\kappa_1$  to be a subtype of a type  $\kappa_2$ , or equate an interval variable to the result of a function applied to interval variables. Note that due to the compositional nature of our subtype relation  $\sqsubseteq$  (which extends to symbolic types) we can restrict ourself to constraints of the form  $v_1 \sqsubseteq v_2$  because each constraint of the form  $\kappa_1 \sqsubseteq \kappa_2$  or  $\mathscr{A}_1 \sqsubseteq \mathscr{A}_2$  (with identical base types) can be reduced to an equivalent set of constraints on interval variables by the definition of the subtype relation.

*Symbolic type system.* In the presentation of our symbolic type inference system, we aim to stay as close as possible to the implementation. Thus we describe it as an impure type system, meaning that our typing rules have side effects. In our case, typing rules can generate fresh interval variables.

For a simple type  $\alpha$ , we write  $fresh(\alpha)$  for the symbolic weightless type obtained by replacing every base type **R** with a fresh interval variable  $v_n$  (and adding weights given

by fresh interval variables where needed). We write fresh() for  $fresh(\mathbf{R})$ . For a symbolic type  $\kappa$  we write  $base(\kappa)$  for the underlying simple type (defined in the obvious way).

Our constraint generation system is given in Fig. 10. Judgments have the form  $\Gamma \vdash M$  :  $\mathscr{A}, C$  where  $\Gamma$  maps variables to weightless symbolic types,  $\mathscr{A}$  is a weighted symbolic type and C a list of constraints on the interval variables. The rules follow the structure of the system in Fig. 4 but replace all operations on intervals with interval variables and constraints. The term structure directly determines the symbolic typing derivation; there are no choices to be made, contrary to Fig. 4, which requires "cleverness", for example to find a suitable interval for an argument in the fixpoint rule. Note that in our system, we assume that the *simple* types of arguments of abstractions and fixpoints are given. These types can be determined by a simple prior run of any standard type inference algorithm.

**From symbolic to concrete types.** An assignment *A* is a mapping from interval variables to concrete intervals. Given a symbolic type  $\kappa$ , we define the interval type  $\kappa^A$  by replacing every interval variable in  $\kappa$  with the concrete interval assigned to it in *A*. For a weighted symbolic type  $\mathscr{A}$ , we define  $\mathscr{A}^A$  in the same way. Given a set of constraints *C*, we say that *A* satisfies *C*, written  $A \models C$  if all constraints in *C* are satisfied (defined in the obvious way).

## **Theorem D.4.** If $\vdash M : \mathcal{A}, C \text{ and } A \models C \text{ then } \vdash M : \mathcal{A}^A$ .

*Proof.* This can be shown by induction on the structure of the term, which also determines the symbolic typing derivation. From this, we obtain a valid interval typing derivation by replacing interval variables in the derivation with the concrete intervals assigned to them in A, and by applying the (SUB) rule in places where subtyping constraints are introduced.

This theorem states that solutions to our constraints directly give us valid judgments in our interval type system, which allows us to invoke Theorem 5.1.

**Solving Constraints.** To solve the resulting constraints, we employ known techniques from abstract interpretation [14]. Again, note that due to the simplicity of our constraints, our approach avoids expensive calls to a theorem prover. When solving a set of constraints *C*, we are interested in the smallest solution, i.e. an assignment *A* with  $A \models C$  where the intervals in  $\mathcal{A}$  are smallest possible w.r.t.  $\sqsubseteq$ .

**Naïve algorithm.** A naïve attempt to find a satisfying assignment for a set of constraints would be to iterate over the constraints and to extend the current assignment (initially chosen to map all elements to the bottom element  $\perp$  in the interval domain, i.e. an empty interval) whenever needed. For example, if  $v_i \subseteq v_j$  is not satisfied by assignment *A*, we can update *A* by mapping  $v_j$  to  $A(v_j) \sqcup A(v_i)$ . As is well known from abstract interpretation, this naïve

$$\frac{x:\kappa \in \Gamma \quad v = fresh()}{\Gamma + x: \binom{\kappa}{\nu}, \{v \equiv 1\}} \qquad \frac{\kappa = fresh(\alpha) \quad v = fresh() \quad \Gamma; x:\kappa + M: \mathscr{A}, C}{\Gamma + \lambda x^{\alpha} \cdot M: \binom{\kappa \to \mathscr{A}}{\nu}, C \cup \{v \equiv 1\}} \qquad \frac{\Gamma + M: \binom{v_1}{v_2}, C \cup \{v \equiv v_1 \times v_2\}}{\Gamma + score(M): \binom{v_1}{\nu}, C \cup \{v \equiv v_1 \times v_2\}} \\ \frac{v_1 = fresh() \quad v_2 = fresh()}{\Gamma + \varepsilon: \binom{v_1}{v_2}, \{v_1 \equiv [r, r], v_2 \equiv 1\}} \qquad \frac{\kappa = fresh(\alpha) \quad \kappa_1 = fresh(\beta) \quad v, v_1 = fresh() \quad \Gamma; \varphi: \kappa \to \binom{\kappa_1}{v_1}; x: \kappa + M: \binom{\kappa_2}{v_2}, C \\ \Gamma + \mu_x^{\varphi;\alpha \to \beta} \cdot M: \binom{\kappa \to \binom{\kappa_2}{v_2}}{\nu_2}, C \cup \{v \equiv 1, \kappa_2 \equiv \kappa_1, v_2 \equiv v_1\} \\ \frac{v, v' = fresh()}{\Gamma + sample: \binom{v}{v'}, \{v \equiv [0, 1], v' \equiv 1\}} \qquad \frac{\Gamma + M: \binom{\kappa_1}{v_3}, C_1 \quad \Gamma + N: \binom{\kappa_3}{v_3}, C_2 \quad v = fresh() \\ \frac{\Gamma + M: \binom{v_1}{v_2}, C_M \quad \Gamma + N: \binom{\kappa_1}{v_3}, C_N \quad \Gamma + P: \binom{\kappa_2}{v_4}, C_P \quad \kappa = fresh(base(\kappa_1)) \quad v, v' = fresh() \\ \Gamma + if(M, N, P): \binom{\kappa}{v}, C_M \cup C_N \cup C_P \cup \{\kappa_1 \equiv \kappa, \kappa_2 \equiv \kappa, v \equiv v_2 \times v', v_3 \equiv v', v_4 \equiv v'\} \\ \frac{\Gamma + M: \binom{v_1}{v'_1}, C_1 \quad \cdots \quad \Gamma + M_{[f]}: \binom{v_{[f]}}{v'_{[f]}}, C_{[f]} \quad v, v' = fresh() \\ \frac{\Gamma + M: \binom{v_1}{v'_1}, C_1 \quad \cdots \quad \Gamma + M_{[f]}: \binom{v_{[f]}}{v'_1}, C_{[f]} \mid v, v' = fresh() \\ \frac{\Gamma + f(M_1, \dots, M_{[f]}): \binom{v}{v'}, \bigcup (f_1) \in (v \equiv f(v_1, \dots, v_{[f]}), v' \equiv \Pi)} \end{cases}$$

Figure 10. Symbolic weight-aware type system.

approach may not terminate because the interval domain is not chain-complete. For instance, consider the constraints  $C = \{v_1 = [0,0], v_2 = [1,1], v_1 \sqsubseteq v_3, v_3 \equiv v_3 + v_2\}$ . The minimal solution is  $\{v_1 \mapsto [0,0], v_2 \mapsto [1,1], v_3 \mapsto [0,\infty]\}$ , but the algorithm never terminates and instead assigns the ascending chain  $[0,0], [0,1], [0,2], \dots$  to  $v_3$ .

*Widening.* To remedy the above problem, we use *widening*, a standard approach to ensure termination of abstract interpretation on domains with infinite chains [14]. Let  $\nabla$  be a *widening operator* for intervals. This means that  $I_1 \sqcup I_2 \sqsubseteq I_1 \nabla I_2$  for all intervals  $I_1, I_2$  and for every chain  $I_0 \sqsubseteq I_1 \sqsubseteq I_2 \sqsubseteq \cdots$ , the chain  $(I_i^{\nabla})_{i \in \mathbb{N}}$  defined by  $I_0^{\nabla} \coloneqq I_0$  and  $I_i^{\nabla} \coloneqq I_{i-1} \nabla I_i$  for  $i \ge 1$  stabilises eventually. A trivial widening operator is given by only allowing intervals to extend to infinity, defined as follows:

As soon as the upper bound increases or lower bound decreases, the bound is directly set to  $= \infty$  or  $-\infty$  respectively.

By using the widening operator in each update step of our naïve algorithm, we break infinite increasing chains and the resulting algorithm is guaranteed to converge to a satisfying assignment (if one exists).

GuBPI solves constraints by using a standard worklist algorithm [15, 23], combined with the previous widening operator.

# E Supplementary Material for Section 6

## E.1 Extensions to Linear Splitting

In this section we give additional information about how our linear optimisation of the interval-trace semantics can be extended to non-uniform samples and non-linear scoring values.

**Beyond uniform samples.** To allow for non-uniform samples, we can combine the standard interval trace semantics with the linear optimisation. That is, in addition to bounding linear score functions, we also split and bound each non-uniform sample (as in the standard interval trace semantics). Suppose that  $\alpha_i$  is sampled from some continuous distribution *D*. We then split the real line into chunks (the size and number of which is selected by a heuristic depending on *D*). For each such chunk [a, b], we compute the volume and multiply by the lower and upper bounds of the pdf of *D* on [a, b]. In this way, we can even approximate integrals where

not all variables are sampled from a uniform distribution (without needing to resort to inverse cumulative distribution functions). Our experiments show, that as long as the guards are still linear, this approach is advantageous compared to the naive interval-based semantics.

**Beyond linear functions.** While guards on conditionals are often linear, this is rarely the case for score expressions (as one usually observes values from some non-uniform distribution with a non-linear pdf). Consider the pedestrian example again. While all guards are linear, the score expression has the form  $pdf_{Normal(1.1,0.1)}(\mathcal{V})$ . We can handle such *non-linear* score values by applying linear optimisation to the linear subexpressions and interval arithmetic for the nonlinear parts. Formally, we assume that each  $W_i \in \Xi$  (each symbolic value we score with) has the form  $W_i = f_i(\mathcal{Z}_1^i, \ldots, \mathcal{Z}_i^{m_i})$  where  $\mathcal{Z}_i^j$  for  $1 \leq j \leq m_i$  denote *linear* functions of the sample variables and  $f : \mathbb{R}^{m_i} \to \mathbb{R}$  is a possibly non-linear function. Every score expression can be written in this way. For instance, in the pedestrian example, we have  $f = pdf_{Normal(1.1,0.1)}$ .

Let  $\Xi' = \{Z_i^j \mid i \in \{1, ..., n\}, j \in \{1, ..., m_i\}\}$  be the set of all such linear functions. We bound each linear function in  $\Xi'$ using linear optimisation as before. We obtain a box b (which now has dimension  $|\Xi'|$  instead of  $|\Xi|$ ) and define the weight weight(b) by applying the interval liftings  $f_i^{\mathbb{I}}$  of the nonlinear functions  $f_i$  to the bounds for each argument. Formally, weight(b) =  $\prod_{i=1}^{k} f_i^{\mathbb{I}}(b_i^1, \dots, b_i^{m_i})$  where  $b_i^j$  is the interval bound on  $\mathcal{Z}_{i}^{j}$ . Note that this strictly generalizes the approach outlined before since we can choose  $f_i$  as the identity if  $W_i$ is already linear. The definition of approx(b) with the new weight definition still satisfies Proposition 6.4. This way, we can even approximate integrals over non-linear functions by means of simple volume computations. As our experiments (e.g. on the pedestrian example) show, this approximation is precise enough to obtain useful bounds on the posterior. It is important to note that while we can deal with non-linear score values, we cannot handle non-linear guards and instead use the standard semantics for such cases.

# F Supplementary Material for Section 7

Our experiments were performed on a server running Ubuntu 18.04 with an 8core Intel(R) Xeon(R) CPU E3-1271 v3 @ 3.60GHz CPU with 32Gbp of RAM. The current version of GuBPI is not parallelised and makes no use of the additional cores. The running times on a Macbook Pro with Apple M1 were comparable, and sometimes even faster.

## F.1 Pyro's HMC samples for the pedestrian example

The HMC samples plotted in Figs. 1 and 7 were generated with the probabilistic programming system Pyro [5]. Since the original pedestrian program has infinite expected running time, we introduced a stopping condition in the random walk: if the distance traveled exceeded 10, the loop was exited. (This has a negligible effect on the posterior distribution because the weight of such a trace is at most  $pdf_{Normal(1,1,0,1)}(10) < 10^{-1700}$ .)

We used Pyro's HMC sampler to compute 10 Markov chains with 1000 samples each for this program. We set the hyperparamaters to 0.1 for the step size and 50 for the number of steps. We also tried the NUTS sampler, which aims to automatically estimate good values for the hyperparameters, but it performed worse than the manually chosen values. The running time for the chains varied significantly: some took around one minute, others almost an hour. This depended on whether the Markov chain got "stuck" in a long trace. (The length of the traces varied between 2 and about 200.)

We discarded chains with very low acceptance rates (under 1%), aggregated the remaining chains, which had acceptance rates of over 50%, and used their histogram in Figs. 1 and 7.

# F.2 Details on Probability Estimation

In Table 1 (results of our tool for the probability estimation benchmark), we omitted the query for space reasons. Complete information including the query can be found in Table 4.

# F.3 Simulation-based Calibration

We implemented SBC for both likelihood-weighted importance sampling and Pyro's HMC. As hyperparameters for SBC, we picked L = 63 samples per simulation (following the suggestion in [60] to take one less than a power of two) and N simulations with N = 10L (also following the paper's suggestion). Note that the number of samples is much less than the 10000 samples used for Figs. 1 and 7 (10 chains with 1000 samples each). But setting L = 1000 would be at least 100 times slower because N has to increase proportionally to L. Also note that for Pyro, we ran HMC with L warmup steps before generating L samples. Both importance samples and HMC samples exhibited significant autocorrelation. As suggested in [60], we applied thinning to reduce its effect, choosing a thinning factor of around  $\frac{L}{L_{eff}}$  where  $L_{eff}$  is the effective sample size.

**Pedestrian example.** For importance sampling, the rank histogram looks fairly uniform (Fig. 11a), which means that SBC does not detect an issue with the inference and thus increases the confidence in the validity of the importance samples. For Pyro's HMC, simulation-based calibration is very slow: the rank histogram (Fig. 11b) took 32 hours (!) to produce. (Note that only 332 simulations could be used, the rest were discarded because the acceptance rate was too low.) The spikes at the boundary indicate that the samples have high autocorrelation and in fact, the effective sample size  $L_{eff}$  was at most  $\frac{L}{10}$ , often much lower (depending on the chain). The suggestion in [60] is thus to apply thinning, with a factor of  $\frac{L}{L_{eff}}$ , which is at least 10 in our case. This would increase

Table 4. Evaluation on selected benchmarks from [56]. We give the times (in seconds) and bounds computed by [56] and
GuBPI. The table agrees with Table 1 but spells out the full problem name and the exact query.

			То	ool from [56]		GuBPI
Program	Q	Query	t	Result	t	Result
tug-of-war	Q1	$total\_a\_b < total\_t\_s$	1.29	[0.6126, 0.6227]	0.72	[0.6134, 0.6135]
tug-of-war	Q2	<pre>total_a_s &lt; total_b_t</pre>	1.09	[0.5973, 0.6266]	0.79	[0.6134, 0.6135]
beauquier-etal-3	Q1	<i>count</i> < 1	1.15	[0.5000, 0.5261]	22.5	[0.4999, 0.5001]
example-book-simple	Q1	$count \geq 2$	8.48	[0.6633, 0.7234]	6.52	[0.7417, 0.7418]
example-book-simple	Q2*	$count \ge 4$	10.3	[0.3365, 0.3848]	8.01	[0.4137, 0.4138]
example-cart	Q1	$count \ge 1$	2.41	[0.8980, 1.1573]	67.3	[0.9999, 1.0001]
example-cart	Q2	$count \ge 2$	2.40	[0.8897, 1.1573]	68.5	[0.9999, 1.0001]
example-cart	Q3	$count \ge 4$	0.15	[0.0000, 0.1150]	67.4	[0.0000, 0.0001]
example-ckd-epi-simple	Q1*	$f_1 \le 4.4 \land f \ge 4.6$	0.15	[0.5515, 0.5632]	0.86	[0.0003, 0.0004]
example-ckd-epi-simple	Q2*	$f_1 \ge 4.6 \land f \le 4.4$	0.08	[0.3019, 0.3149]	0.84	[0.0003, 0.0004]
example-fig6	Q1	<i>c</i> ≤ 1	1.31	[0.1619, 0.7956]	21.2	[0.1899, 0.1903]
example-fig6	Q2	$c \leq 2$	1.80	[0.2916, 1.0571]	21.4	[0.3705, 0.3720]
example-fig6	Q3	<i>c</i> ≤ 5	1.51	[0.4314, 2.0155]	24.7	[0.7438, 0.7668]
example-fig6	Q4	<i>c</i> ≤ 8	3.96	[0.4400, 3.0956]	27.4	[0.8682, 0.9666]
example-fig7	Q1	$x \le 1000$	0.04	[0.9921, 1.0000]	0.18	[0.9980, 0.9981]
example4	Q1	x + y > 10	0.02	[0.1910, 0.1966]	0.31	[0.1918, 0.1919]
example5	Q1	x + y > z + 10	0.06	[0.4478, 0.4708]	0.27	[0.4540, 0.4541]
herman-3	Q1	<i>count</i> < 1	0.47	[0.3750, 0.4091]	124	[0.3749, 0.3751]









(a) Importance sampling for the (b) HMC for the pedestrian exam- (c) NUTS for the binary Gaussian (d) NUTS for the two-dimensional tor 100.

pedestrian example, thinning fac- ple, no thinning. (Only 332 simu- mixture model, thinning factor 10. binary Gaussian mixture model, lations were used because the rest had too low acceptance rates.)

thinning factor 10.

Figure 11. Simulation-based calibration: rank histogram plots (630 simulations with 63 samples each).

the running time of SBC by the same factor, to at least 300 hours, but probably 600 or more. We did not consider it a good use of resources to carry out this experiment.

Binary Gaussian Mixture Model. We also considered the binary GMM (Fig. 5c) and a two-dimensional version of the same model. The spikes at the boundary could again be a sign of high autocorrelation, but in this case, we already applied thinning with a factor of 10 (again based on the effective sample size). Instead, as discussed in [60], this symmetric

U-shape indicates that the computed data-averaged posterior is underdispersed relative to the prior distribution. This interpretation is consistent with our knowledge about the model: HMC only finds one mode in the posterior distribution and misses the others. Hence SBC successfully detects the issue, and in the case of the higher-dimensional model, it does so in less time than GuBPI (cf. Table 3). For the other models, including the pedestrian example, GuBPI is faster.

# Contents

Abst	ract	1
1	Introduction	1
1.1	Guaranteed Bounds	2
1.2	Contributions	2
1.3	Scope and Limitations	3
2	Background	3
2.1	Basic Probability Theory and Notation	3
2.2	Statistical PCF (SPCF)	3
2.3	Trace Semantics	4
3	Interval Trace Semantics	4
3.1	Interval Arithmetic	4
3.2	Interval Traces and Interval SPCF	4
3.3	Bounds from Interval Traces	5
4	Soundness and Completeness	6
4.1	Soundness	6
4.2	Completeness	6
5	Weight-aware Interval Type System	7
5.1	Interval Types	7
5.2	Type System	8
5.3	Constraint-based Type Inference	8
6	Symbolic Execution and GuBPI	8
6.1	Symbolic Execution	8
6.2	GuBPI	9
6.3	Standard Interval Trace Semantics	9
6.4	Linear Interval Trace Semantics	10
7	Practical Evaluation	10
7.1	Probability Estimation	10
7.2	Exact Inference	11
7.3	Recursive Models	12

7.4	Comparison with Statistical Validation	13
7.5	Limitations and Future Improvements	13
8	Related Work	13
9	Conclusion	14
Refe	rences	15
А	Supplementary Material for Section 3	17
A.1	Intervals as a lattice	17
A.2	Lifting Functions to Intervals	17
A.3	Properties of Interval Reduction	17
A.4	Additional Possible Reduction Rules	17
В	Symbolic Execution	18
С	Supplementary Material for Section 4	19
C.1	Infinite Trace Semantics	19
C.2	Exhaustivity and Soundness	20
C.3	Assumptions for Completeness	20
C.4	Completeness Proof	22
D	Supplementary Material for Section 5	25
D.1	Soundness	25
D.2	Weak Completeness	27
D.3	Constraint-based Type Inference	28
E	Supplementary Material for Section 6	29
E.1	Extensions to Linear Splitting	29
F	Supplementary Material for Section 7	30
F.1	Pyro's HMC samples for the pedestrian	
	example	30
F.2	Details on Probability Estimation	30
F.3	Simulation-based Calibration	30
Cont	tents	32