



Inference of Probabilistic Programs with Moment-Matching Gaussian Mixtures

FRANCESCA RANDONE, IMT School for Advanced Studies Lucca, Italy

LUCA BORTOLUSSI, University of Trieste, Italy

EMILIO INCERTO, IMT School for Advanced Studies Lucca, Italy

MIRCO TRIBASTONE, IMT School for Advanced Studies Lucca, Italy

Computing the posterior distribution of a probabilistic program is a hard task for which no one-fit-for-all solution exists. We propose Gaussian Semantics, which approximates the exact probabilistic semantics of a bounded program by means of Gaussian mixtures. It is parametrized by a map that associates each program location with the moment order to be matched in the approximation. We provide two main contributions. The first is a universal approximation theorem stating that, under mild conditions, Gaussian Semantics can approximate the exact semantics arbitrarily closely. The second is an approximation that matches up to second-order moments analytically in face of the generally difficult problem of matching moments of Gaussian mixtures with arbitrary moment order. We test our second-order Gaussian approximation (SOGA) on a number of case studies from the literature. We show that it can provide accurate estimates in models not supported by other approximation methods or when exact symbolic techniques fail because of complex expressions or non-simplified integrals. On two notable classes of problems, namely collaborative filtering and programs involving mixtures of continuous and discrete distributions, we show that SOGA significantly outperforms alternative techniques in terms of accuracy and computational time.

CCS Concepts: • **Theory of computation** → **Denotational semantics**; • **Mathematics of computing** → **Probabilistic reasoning algorithms**.

Additional Key Words and Phrases: probabilistic programming, inference, Gaussian mixtures

ACM Reference Format:

Francesca Randone, Luca Bortolussi, Emilio Incerto, and Mirco Tribastone. 2024. Inference of Probabilistic Programs with Moment-Matching Gaussian Mixtures. *Proc. ACM Program. Lang.* 8, POPL, Article 63 (January 2024), 31 pages. <https://doi.org/10.1145/3632905>

1 INTRODUCTION

Probabilistic programming languages are programming languages augmented with primitives expressing probabilistic behaviours [Gordon et al. 2014]. Examples are random assignments (“program variable x is distributed according to the probability distribution D ”), probabilistic choices (“do P_1 with probability p else P_2) or conditioning (“variable x is distributed according to D , under the constraint that it can only take positive values”). This has enabled a variety of applications such as the analysis of randomized algorithms, machine learning and biology [Gordon et al. 2014].

Given a probabilistic program, there are different equivalent ways in which its semantics can be defined [Kozen 1983]. Following Kozen’s Semantics 2 [Kozen 1979], in this paper we see a program as a transformer: given an initial joint distribution over the program variables, each instruction

Authors’ addresses: [Francesca Randone](#), IMT School for Advanced Studies Lucca, Italy, francesca.randone@imtlucca.it; [Luca Bortolussi](#), University of Trieste, Italy, lbortolussi@units.it; [Emilio Incerto](#), IMT School for Advanced Studies Lucca, Italy, emilio.incerto@imtlucca.it; [Mirco Tribastone](#), IMT School for Advanced Studies Lucca, Italy, mirco.tribastone@imtlucca.it.



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

© 2024 Copyright held by the owner/author(s).

ACM 2475-1421/2024/1-ART63

<https://doi.org/10.1145/3632905>

in the program transforms that joint distribution into a possibly different one, for example, due to the presence of probabilistic assignments or conditional statements. In this framework, we are interested in the *inference problem*: given a program P and an initial distribution D over program variables, what is the distribution over program variables after executing P ? Borrowing from Bayesian inference, we will sometimes refer to the initial distribution D as the *prior distribution* over program variables and to the distribution obtained after executing P as the *posterior distribution*. Then, the inference problem boils down to computing the posterior.

Over the years, many approaches have tackled this problem: numerical methods based on Monte Carlo Markov chain (MCMC) sampling [Chaganty et al. 2013; Goodman et al. 2008; Hastings 1970; Mansinghka et al. 2014; Nori et al. 2014; Pfeffer 2001], variational inference (VI) [Bingham et al. 2019; Jordan et al. 1999; Kucukelbir et al. 2015], symbolic execution [Gehr et al. 2016; Narayanan et al. 2016; Saad et al. 2021], volume computation [Fileri et al. 2013; Holtzen et al. 2020; Huang et al. 2021], and approaches based on moment-based invariants [Barthe et al. 2016; Bartocci et al. 2020; Chakarov and Sankaranarayanan 2014; Katoen et al. 2010; Moosbrugger et al. 2022].

1.1 Motivating Example

```

1:  $x = 2, y = -1$ 
2: for  $i$  in  $\text{range}(n)$  do
3:    $x = x + \text{gauss}(0, 1)$ 
4:    $y = y + \text{gauss}(0, 1)$ 
5: end for
6:  $\text{dist} = x^2 + y^2$ 
7: if  $\text{dist} > 10$  then
8:    $\text{out} = 1$ 
9: else
10:   $\text{out} = 0$ 
11: end if
12: if  $\text{out} = 1$  then
13:   $\text{obs\_dist} = 10$ 
14: else
15:   $\text{obs\_dist} = \text{gauss}(\text{dist}, 1)$ 
16: end if
17: observe( $\text{out} == 1$ )
18: return  $\text{obs\_dist}$ 

```

As a motivating example, let us consider the *Tracking_n* model reported in left inset and adapted from Wu et al. [2018]. It describes a Gaussian process evolving on a bi-dimensional space for n steps and starting from coordinates (2, -1) (lines 1-5). A radar is positioned in (0,0), and can sense the process if it is at a squared distance (*dist*, line 6) of less than 10 units from the radar. Therefore the process can be either out of scope (*out* = 1, line 8) or in scope (*out* = 0, line 10). When the process is out of scope, the radar returns an observed distance of 10 (line 13) and a noisy measurement of the true distance else (line 15). Therefore, the distribution over *obs_dist* is a mixture of δ_{10} , i.e., a Dirac delta centered in 10, and a Gaussian with mean *dist*. However, if we observe that the process is out of scope (line 17), the posterior over *obs_dist* is just δ_{10} because any continuous distribution puts zero mass on a single point. Therefore, the exact posterior over *obs_dist* is a distribution placing probability 1 on 10. While this program may seem quite simple, performing inference may be challenging.

Using PSI [Gehr et al. 2016], an exact symbolic solution returns a formula for the posterior mean of *obs_dist* in less than a second, which, however, contains several non-simplified integrals. This is because, in line 6, computing *dist* requires computing the probability density function (pdf) of the product of two continuous distributions, and this requires symbolic integration. Attempting to integrate it numerically using Mathematica [Wolfram Research, Inc.] did not terminate after 30 minutes on common machine. One can resort to approximate approaches; however, many methods, such as AQUA's quantization [Huang et al. 2021] and STAN's MCMC sampling [Carpenter et al. 2017] and Pyro's VI [Bingham et al. 2019] do not support discrete posteriors, therefore this particular program cannot be encoded in their syntax. BLOG is a probabilistic programming language relying on probabilistic relational model representation and likelihood weighting sampling [Milch et al. 2004], that has been extended by Wu et al. [2018] for mixtures of continuous and discrete distribution such as the one in our example. It computes the exact posterior in 0.516 s for $n = 1$ and about 5 s for $n = 100$.




Fig. 1. Left part: general approximation scheme used in Gaussian Semantics. In program location P_i the exact semantics of the program transforms a GM G into a non-GM distribution D . In the same program location, Gaussian Semantics transform G into another GM G_D , that approximates D using moment-matching. Right part: concrete example. The Gaussian distribution G is transformed by the exact semantics into a truncated Gaussian D and by the second-order Gaussian Semantics into the red Gaussian distribution G_D .

A similar behavior is exhibited by applying Pyro’s variable elimination [Obermeyer et al. 2019], which computes the exact posterior in 0.192 s for $n = 1$ and about 9 s for $n = 100$ (see Section 7.3).

1.2 Proposed Approach

The difficulty in performing inference on the previous program stems from various factors: PSI’s exact engine returns non-simplified integrals, requiring computationally expensive numerical integration. STAN’s MCMC, Pyro’s VI and AQUA’s quantization cannot be applied in this case, but in general, can incur long computational times and out-of-memory errors (see Section 7). BLOG’s and Pyro’s ad hoc sampling perform best, but increasing the number of steps hinders scalability.

To complement all these techniques, we present a new *approximate analytical* method that does not require integration or sampling and that relies on a compact representation of the joint distribution using moment-matching Gaussian mixtures (GMs). Our choice of representation is based on some desirable properties of GMs, and in particular the following three: i) they can encode both continuous and discrete distributions (using degenerate GMs); ii) their moments can be computed exactly and efficiently; iii) they are universal approximators, so we can always increase the number of components in our representation to get a better approximation. These considerations lead to the definition of a family of approximating semantics called *Gaussian Semantics*.

More in detail, we define Gaussian Semantics so that it is closed with respect to the class of (degenerate) GMs, meaning that, at every program location, the Gaussian Semantics of a program transforms a GM into a GM. In particular, we proceed as in the general approximation scheme proposed by Boyen and Koller [1998]: given a GM G , the exact semantics of a program location would transform it into a different distribution D , which is not necessarily a GM. However, we approximate D with a new GM G_D and define the Gaussian Semantics as semantics that transform G into G_D at that program location. This process is represented in Figure 1. Performing this at every program location approximates the whole program semantics. In particular, we choose to approximate D with G_D using moment-matching, meaning that G_D is a GM having the same moments of D up to a certain order r . This is convenient for two reasons: first, it avoids computing the full pdf of D , as only its first r moments are needed to find G_D ; second, since D is obtained as a transformation of a GM, it can be expressed as a linear combination of transformed Gaussians, and its moments can be computed analytically using the results summarized in Table 1.

To sum up, in a Gaussian Semantics each program location is associated with an integer r , and the semantics acts on a GM G performing two steps: first, it computes the first r order moments of the transformed distribution D , using the results in Table 1; then, it finds a new GM G_D having same moments as D up to order r . More than one moment-matching GM G_D can exist, therefore,

Table 1. Summary of the theoretical results used to compute the moments of transformed Gaussian Mixtures. $\mathcal{N}(\mu, \Sigma)$ denotes a Gaussian distribution with mean μ and covariance matrix Σ , c is any real constant and a, b are vectors in \mathbb{R}^d defining the hyper-rectangle $[a, b] = \{x \in \mathbb{R}^d : a_i \leq x_i \leq b_i\}$.

Operation	Theoretical Result	Computes moments for:
Sum of Gaussians	Closed w.r.t linear transformations [Billingsley 2008]	$c_1\mathcal{N}(\mu_1, \Sigma_1) + c_2\mathcal{N}(\mu_2, \Sigma_2)$
Conditioning Gaussians to $x_i = c$	Closed w.r.t. conditioning [Bishop and Nasrabadi 2006]	$\mathcal{N}(\mu, \Sigma \mid x_i = c)$
Conditioning Gaussians to $x \in [a, b]$	Iterative formulas [Kan and Robotti 2017]	$\mathcal{N}(\mu, \Sigma \mid x \in [a, b])$
Product of Gaussian	Isserlis' Theorem [Wick 1950]	$\mathcal{N}(\mu_1, \Sigma_1)\mathcal{N}(\mu_2, \Sigma_2)$

we give a heuristic to determine a unique G_D for any r . In particular, we base our heuristics, called *max entropy matching*, on the maximum entropy principle [Kullback and Leibler 1951].

Our first technical contribution is theoretical: we provide a universal approximation result stating that, under mild conditions, when the order of moments matched at each program location grows, the family of Gaussian Semantics converges to the exact probabilistic semantics. While our result exploits the well-known universal approximation power of GMs [Lo 1972], it is a non-trivial consequence of it. The density of GMs guarantees the existence of a GM arbitrarily close to a target distribution; however, for a probabilistic program the target distribution is generally not known. Here we give a *constructive method* to build the approximating GM.

Besides the definition of Gaussian Semantics, we look at how they can be practically computed. Unfortunately, it turns out that while the formulas in Table 1 allow us to compute moments up to any order, finding a moment-matching GM is a hard task. In fact, finding a moment-matching GM for moment orders higher than two requires the solution of a constrained system of polynomial equations, for which no analytical solution is known [Lasserre 2009]. Despite this, when only the first two orders of moments are matched, our matching boils down to using a single Gaussian distribution with a given mean and covariance, and no system of equations needs to be solved.

We call this particular instance *Second Order Gaussian Approximation* (SOGA) and present an algorithm that implements it. In our motivating example, at line 6, to approximate the distribution of *dist* after the assignment $dist = x^2 + y^2$, SOGA proceeds as follows. It first computes the means and covariance matrices of x^2 and y^2 using Isserlis' theorem [Wick 1950] (observe that x and y are Gaussian, but x^2 and y^2 are not). Then, it approximates the distributions of x^2 and y^2 with two Gaussians having the computed means and covariance matrices. Finally, it exploits the closedness of Gaussians with respect to sum to approximate the distribution of *dist* with the sum of the Gaussians approximating x^2 and y^2 . Therefore, while in the exact semantics, after line 6, *dist* does not have a GM distribution, in SOGA it does. This significantly simplifies the subsequent computations. Indeed, when entering the if statement at line 7, *dist* is conditioned to $dist > 10$. Performing conditioning in the exact semantics requires computing the integral of the pdf of *dist* over the set of vectors satisfying $dist > 10$. Instead, in SOGA *dist* is Gaussianly distributed, therefore we can compute the moments of the conditioned distribution using the formulas from Kan and Robotti [2017], and then approximate the conditioned distribution with a Gaussian having given mean and covariance matrix. Overall, for Algorithm 1 SOGA computes the output, which in this case is exact, in 0.042 s for $n = 1$ and in 0.192 s for $n = 100$, performing significantly better than BLOG and Pyro.

In general, the posterior computed by SOGA is a GM whose number of components grows exponentially in the number of conditional statements. To help cope with this, we introduce a pruning strategy that keeps the number of components in the GMs below a user-specified threshold by merging components with minimal cost. Using a prototype implementation, we compare SOGA on a corpus of benchmarks against state-of-the-art tools representative of different inference methods: MCMC sampling (STAN), symbolic execution (PSI), quantization (AQUA), VI (Pyro). Even when it is not the best-performing method, it still provides the flexibility to model both continuous and discrete posteriors, unlike STAN, Pyro and AQUA, which only support the former. Additionally, it enables reaching numerical solutions in reasonable runtimes when PSI returns non-simplified integrals that demand computationally prohibitive times for numerical integration. When applied to the analyzed benchmarks, pruning significantly reduced the computational time without incurring noticeable approximation errors.

Importantly, we highlight that SOGA is particularly useful for performing inference on two classes of programs: those involving mixtures of continuous and discrete distributions and collaborative filtering models. Most state-of-the-art approaches do not support the first class, even though it is known that this kind of distribution arises in various application domains [Gao et al. 2017; Kharchenko et al. 2014; Pierson and Yau 2015]. Thanks to its GM representation, SOGA can easily encode these distributions. When tested on benchmarks introduced specifically for this problem, SOGA is able to perform inference faster than dedicated methods such as Wu et al. [2018], while identifying the exact posterior. Collaborative filtering models are an established framework to model recommendation systems and have been extensively investigated in the machine learning community [Koren et al. 2021]. SOGA can deal with a large number of variables without incurring large computational times or out-of-memory errors, as happens with alternative methods.

Contributions. In summary, the contribution of this paper is threefold:

- i) From the theoretical point of view, we define a family of approximating semantics called Gaussian Semantics and prove that they approximate the exact semantics of a bounded probabilistic program arbitrarily well.
- ii) From the practical perspective, we present an implementation of a particular instance of Gaussian Semantics, called SOGA, and evaluate it against other state-of-the-art implementations of alternative techniques (PSI [Gehr et al. 2016], STAN [Carpenter et al. 2017], AQUA [Huang et al. 2021], Pyro [Bingham et al. 2019]) on a set of benchmarks taken from the literature. While not always the best performer, SOGA can handle models with discrete posteriors while STAN, AQUA and Pyro only support continuous ones. On the other hand, SOGA can provide accurate and computationally tractable approximations when symbolic analysis by PSI may fail due to complex formulas or non-simplified integrals that cause high computational costs for their numerical integration.
- iii) We focus on two classes of models taken from the machine learning literature – collaborative filtering and inference involving mixtures of continuous and discrete distributions – where SOGA clearly outperforms the other methods, complementing the current state-of-the-art.

Paper Structure. Notation and background notions are presented in Section 2. Control-flow syntax and exact semantics are introduced in Section 3. Gaussian Semantics is introduced in Section 4, while the universal approximation theorem is presented in Section 5. We present SOGA in Section 6 and evaluate it in Section 7. We cover further related works in Section 8, while conclusions and future works are drawn in Section 9.




Fig. 2. Plot of: a) a non-degenerate 2-dimensional Gaussian; b) a degenerate Gaussian whose covariance matrix has rank 1; c) a degenerate Gaussian with null covariance matrix (Dirac delta). Considering mixtures of possibly degenerate Gaussians, allow us to capture (mixtures of) both continuous and discrete distributions.

2 BACKGROUND

We now introduce some of the notation and the concepts that will be used in the rest of the paper (we refer the reader to the Supplementary Material for additional background material).

Notation. Given a Boolean value B , $\neg B$ denotes its negation. For a vector $x \in \mathbb{R}^d$, $x \setminus x_i$ denotes a vector in \mathbb{R}^{d-1} obtained from x by suppressing the i -th component; $x[x_i = E(x)]$ denotes a vector $x \in \mathbb{R}^d$ in which the i -th component is replaced with the expression $E(x)$; $\|x\|$ denotes the 2-norm; $\text{diag}(d_1, \dots, d_s)$ denotes the $\mathbb{R}^{s \times s}$ diagonal matrix having d_1, \dots, d_s as diagonal elements.

Probability Distributions. We always deal with distributions over \mathbb{R}^d and use $x \sim D$ to denote that the stochastic vector x is distributed according to distribution D . We always assume that a distribution D can be specified by its probability density function (pdf) $f_D : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$. For $x \sim D$ and a set $A \subseteq \mathbb{R}^d$, the probability of A under D , denoted by $P_D(A)$, can be expressed as the Lebesgue integral $P_D(A) = \int_A f_D(x) dx$. Sometimes, we will find it more convenient to refer to the probability measure induced by D on the measurable space $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$, where $\mathcal{B}(\mathbb{R}^d)$ is the Borel σ -algebra on \mathbb{R}^d . By probability measure, we mean a function $m : \mathcal{B}(\mathbb{R}^d) \rightarrow [0, 1]$ that satisfies the following two properties: i) $m(\emptyset) = 0$ and $m(\mathbb{R}^d) = 1$; ii) $m(\cup_{i \in \mathbb{N}} A_n) = \sum_{i \in \mathbb{N}} m(A_n)$ for any countable collection of disjoint sets $A_n \subset \mathbb{R}^d$. For a distribution D , the associated measure $m_D(A)$ is given by $m_D(A) = P_D(A) = \int_A f_D(x) dx$ for every $A \in \mathcal{B}(\mathbb{R}^d)$. Moreover, due to the presence of conditional branches and observe statements in a probabilistic program, we consider distributions conditioned to subsets of \mathbb{R}^d . Letting \mathbb{I}_A be the characteristic function of a set $A \subset \mathbb{R}^d$ such that $P_D(A) > 0$, $D|A$ will denote the *distribution of D truncated (or conditioned) to A* , whose pdf is given by $f_{D|A} = \frac{1}{P_D(A)} f_D \mathbb{I}_A$. Observe that $f_{D|A}$ is obtained by setting f_D to 0 outside A , and then, by dividing it by $P_D(A)$, so that the induced measure is still a probability measure. Given a d -dimensional random vector $x \sim D$ and a subvector $x' = (x_{i_1}, \dots, x_{i_s})$ with $i_1, \dots, i_s \in \{1, \dots, d\}$, we denote by $\text{Marg}_{x'}(D)$ the marginal distribution of D , obtained integrating out the components not in x' , i.e. $x' \sim \text{Marg}_{x'}(D) = \int_{\mathbb{R}^{d-s}} f_D(x) d(x \setminus x')$.

Gaussian Distributions and Mixtures. Gaussian distributions with mean μ and covariance matrix Σ are denoted by $\mathcal{N}(\mu, \Sigma)$. We assume that Σ can be singular, which corresponds to Gaussian distributions having support in a subspace of \mathbb{R}^d , as shown in Figure 2. In particular, when Σ is the null matrix we consider the associated random variable to be a Dirac delta centered in μ , referred

to as δ_μ . Gaussian distributions enjoy many useful properties; some that we will use are listed in Table 1 while we refer to [Bishop and Nasrabadi 2006, Chapter 2.3] for a detailed treatment.

We refer to *mixtures* as the scalar products of two vectors (p_1, \dots, p_C) and (D_1, \dots, D_C) such that $\sum_{i=1}^C p_i = 1$, $0 < p_i \leq 1$, and D_i is the distribution of the i -th *component*, with $i = 1, \dots, C$. The numbers p_i are called *mixing coefficients*. We will denote a mixture as $M = p_1 D_1 + \dots + p_C D_C$, thus indicating that M has pdf $f_M(x) = p_1 f_{D_1}(x) + \dots + p_C f_{D_C}(x)$. When $C = 1$, we recover the case of a single distribution. A special case is given by Gaussian Mixtures (GMs) in which $D_i = \mathcal{N}(\mu_i, \Sigma_i)$ for $i = 1, \dots, C$, with mean vectors and covariance matrices μ_i, Σ_i . We assume $(\mu_i, \Sigma_i) \neq (\mu_j, \Sigma_j)$ for $i \neq j$. The set of GMs is dense in the set of probability distributions with respect to the weak topology [Lo 1972], meaning that for any probability distribution one can always find a GM that approximates it arbitrarily closely with respect to a particular metric, the Levy-Prokhorov distance. Since we consider Dirac deltas as particular Gaussian distributions, discrete distributions over a finite set of values are included in the set of GMs.

Distributions Determined by Their Moments. Let $x \sim D$. For $r = (r_1, \dots, r_d)$, with $r_i \in \mathbb{N}$, $\forall i$ define

$$\mathbb{E}[x^r] = \int_{\mathbb{R}^d} x_1^{r_1} \dots x_d^{r_d} f_D(x) dx.$$

Letting r vary over all vectors in \mathbb{N}^d such that $r_1 + \dots + r_d = r$ we obtain the set of all r -th order central moments of D . Observe that, for any D , $\mathbb{E}[x^0] = 1$. Since the construction of our semantics relies on the Method of Moments, we need to assure that this converges to the correct distribution. This is true only if no other distribution has all moments equal to those of the target one [Billingsley 2008]. We say that in this case, the target distribution is *determined by its moments*, formalized next.

Definition 2.1. A distribution D is determined by its moments, if for any other distribution D' such that for all $r_1, \dots, r_d \geq 0$

$$\int_{\mathbb{R}^d} x_1^{r_1} \dots x_d^{r_d} f_D(x) dx = \int_{\mathbb{R}^d} x_1^{r_1} \dots x_d^{r_d} f_{D'}(x) dx$$

it holds that $m_D = m_{D'}$.

3 SYNTAX AND EXACT PROBABILISTIC SEMANTICS

3.1 Syntax

Following Kozen [1979], we will consider probabilistic programs as transformers over distributions D defined over a vector of variables taking values in \mathbb{R}^d . Similarly to Chaudhuri and Solar-Lezama [2011], we represent programs in a control flow-graph (cfg) syntax [Cousot and Cousot 1977]. We use as explanatory example the simple program in Algorithm 1.

A program is a directed graph $P = (V, E)$ where V is set of nodes and E is the set of edges. Specifically, we consider directed acyclic graphs (DAGs) of bounded depth. Each node belongs to one of five types in $\Gamma = \{\text{entry}, \text{state}, \text{test}, \text{observe}, \text{exit}\}$. We denote the fact that a node $v \in V$ is of a given type $\gamma \in \Gamma$ with $v : \gamma$. The nodes satisfy the following properties.

- A node v : *entry* has no incoming edge and one outgoing edge.
- A node v : *state* has any number of incoming edges and one outgoing edge. A function *cond* is defined on the set of state nodes, such that $\text{cond} : \{v \in V \mid v : \text{state}\} \rightarrow \{\text{true}, \text{false}, \text{none}\}$ and $\text{cond}(v) = \text{none}$ if and only if the parent of v is not a test node.
- A node v : *test* has one incoming edge and two outgoing edges toward state nodes v_1, v_2 such that $\text{cond}(v_1) = \text{true}$ and $\text{cond}(v_2) = \text{false}$.
- A node v : *observe* has one incoming edge and one outgoing edge.
- A node v : *exit* has any number of incoming edges and no output edge.

Moreover, for each program, there is exactly one $v \in V$ such that v : *entry* and one $v \in V$ such that v : *exit*, and they correspond to the root and the only leaf of the DAG representing the program, respectively. The control-flow syntax for Algorithm 1 is represented in Figure 3.

Algorithm 1 Example

```

1: entry                                { $v_0$  : entry}
2:  $x_1 = \text{gauss}(0, 1)$                     { $v_1$  : state}
3: if  $x_1 > 0$                             { $v_2$  : test}
4:    $x_2 = 2x_1 + 1 + \text{gauss}(0, 0.1)$  { $v_3$  : state}
5: else
6:    $x_2 = -2x_1 + 1 + \text{gauss}(0, 0.1)$  { $v_4$  : state}
7: end if
8: skip                                    { $v_5$  : state}
9: observe( $x_2 < 3$ )                       { $v_6$  : observe}
10: exit                                   { $v_7$  : exit}

```

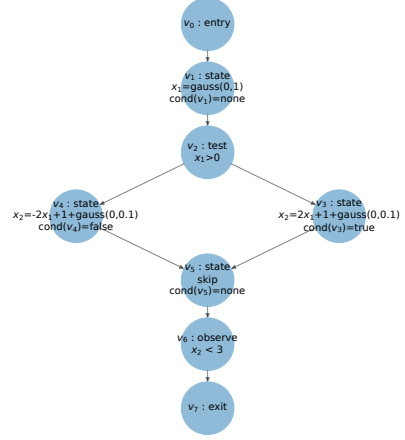


Fig. 3. Cfg representation of Algorithm 1

Variables are defined as

$$z := x \mid g \quad g := \text{gm}([\pi_1, \dots, \pi_s], [\mu_1, \dots, \mu_s], [\sigma_1, \dots, \sigma_s]) \mid \text{gauss}(\mu, \sigma)$$

where x is an *output variable*, i.e., a variable on which to compute the posterior distribution, and g denotes a fresh *read-only variable* distributed according to univariate GMs with mixing coefficients π_i , means μ_i and variances σ_i^2 , $i = 1, \dots, s$. For the sake of brevity, in our examples, we will also use read-only variables denoted by $\text{gauss}(\mu, \sigma)$ which is syntactic sugar for $\text{gm}([1], [\mu], [\sigma])$. We use read-only variables to perform random assignments, as it is done in lines 2, 4, and 6 of Algorithm 1 and to encode Boolean conditions depending on arbitrary distributions.

The vector of output variables is denoted by $x = (x_1, \dots, x_d)$. The vector z augmented with read-only variables is denoted by $z = (x_1, \dots, x_d, g_1, \dots, g_{n-d})$. We denote the distribution of the augmented vector with D_z . We assume read-only variables are dropped after the assignment is performed or the condition is evaluated, marginalizing them out. For example, in line 4 of Algorithm 1, the assignment $x_2 = 2x_1 + 1 + \text{gauss}(0, 0.1)$ is performed by augmenting the vector $x = (x_1, x_2)$ to $z = (x_1, x_2, g)$, with g being an independent standard Gaussian, and assigning x_2 with $2x_1 + 1 + g$. After the new posterior on z is computed, g is marginalized out, returning to the vector $x = (x_1, x_2)$.

State nodes are labelled by either skip or assignment instructions of the type $x_i = E(z)$, where $E(z)$ is an expression of the following form:

$$E(z) := c_1 \cdot z_1 + \dots + c_n \cdot z_n + c \mid z_{i_1} \cdot z_{i_2} \quad (1)$$

where c, c_1, \dots, c_n are scalar.

Test and observe nodes are labelled by linear Boolean conditions (LBCs) of the following form:

$$B(z) := \text{true} \mid \text{false} \mid c_1 \cdot z_1 + \dots + c_n \cdot z_n \bowtie c \mid z_i \square c \quad (2)$$

where c, c_1, \dots, c_n are scalar constants, $\bowtie \in \{<, \leq, \geq, >\}$ and $\square \in \{=, !=\}$. We associate an LBC with the set, defined on the space of augmented variables,

$$\llbracket B(z) \rrbracket = \{z \in \mathbb{R}^n \text{ s.t. } B(z) \text{ holds}\}. \quad (3)$$

where $\llbracket \text{true} \rrbracket = \mathbb{R}^n$ and $\llbracket \text{false} \rrbracket = \emptyset$. Observe that an expression or LBC can have at most d output variables but any finite number of read-only variables.

3.2 Supported Programs

Our syntax rules out general distributions depending on non-constant parameters, unbounded loops, and non-polynomial functions. We briefly comment on the limitations of this approach, how they can be mitigated, and when they are shared by other techniques.

Probabilistic Assignments. Probabilistic assignments are performed by assigning univariate read-only variables to output variables. This is not a limitation since dependence between variables can be encoded using multiple assignments. For what concerns the restriction on GM distributions, instead, we exploit the already discussed density of GMs in the space of distribution [Lo 1972], and assign a GM arbitrarily close to the target distribution. In this sense, we will assume that we are approximating non-GM distributions with a GM whenever we refer to non-GM distributions.

Finally, probabilistic assignments will involve only distributions depending on constant parameters. This restriction is more difficult to overcome and is shared with other tools based on moment-based techniques, such as Bartocci et al. [2020] and Moosbrugger et al. [2022]. This is because it is not always possible to derive how the moments change if one or more parameters of a distribution are probabilistic. As in Moosbrugger et al. [2022], this limitation can be mitigated by performing suitable reparametrizations (see Supplementary Material).

Iterations. We restrict our attention to loops bounded by deterministic constants (as in our illustrating example in Algorithm 1), similarly to Gehr et al. [2016], Huang et al. [2021], Holtzen et al. [2020], Albarghouthi et al. [2017] and Nori et al. [2014]. If guarantees on almost sure termination can be given, the true distribution of the loop could be approximated by a bounded unrolled loop with a sufficiently large number of iterations [Kozen 1979].

Polynomial Programs. Differently from Carpenter et al. [2017]; Gehr et al. [2016]; Huang et al. [2021], we consider programs involving only the arithmetic operations $+$, $-$, $*$, and $^$. This assumption is common to other approaches relying on moment-based techniques such as Bartocci et al. [2020] and Moosbrugger et al. [2022], due to the fact that non-polynomial functions (such as the logarithm) may generate distributions that are not determined by their parameters. We remark that from expressions such as (1) and (2), general polynomial expressions and Boolean conditions can be obtained, respectively, by chaining state nodes and nesting conditional statements. A probabilistic choice, i.e., x is assigned e_1 with probability p or e_2 with probability $1 - p$, is encoded using the LBC $y < q$ where y is a standard Gaussian and q is the Gaussian p -quantile.

3.3 Exact Probabilistic Semantics

The “exact” semantics follows Kozen’s *Semantics 2* [Kozen 1979]. Since we are using the control-flow syntax of Cousot and Cousot [1977], we are close to the *collecting semantics* in Chaudhuri and Solar-Lezama [2011]: we combine the semantics of the nodes to define the semantics of the paths, then define the semantics of the program as a sum over the semantics of the paths. This semantics is particularly convenient for our method because it gives the posterior distribution as a mixture, similarly to Zhou et al. [2020].

Given a program $P = (V, E)$, we define a path π as an ordered sequence of nodes $\pi = v_0 \cdots v_n$ with v_0 : *entry*, v_n : *exit* and $(v_{i-1}, v_i) \in E, \forall i = 1, \dots, n$. The successor of node v_i in path π is denoted as $s_\pi(v_i) = v_{i+1}$. The set of all paths of P is denoted by Π^P . We define the semantics of a path π , denoted by $\llbracket \pi \rrbracket$, as a pair (p, D) , where $p \geq 0$ and D is a probability distribution on \mathbb{R}^d . The semantics $\llbracket \pi \rrbracket$ composes the semantics of the nodes along path π , i.e., $\llbracket \pi \rrbracket = \llbracket v_n \rrbracket_\pi \circ \dots \circ \llbracket v_0 \rrbracket_\pi$. The semantics of each node is defined as follows.

- The *entry* node outputs the pair $(1, \delta_0)$, δ_0 being a Dirac delta centered on the zero vector:

$$\text{if } v: \textit{entry} \text{ then } \llbracket v \rrbracket_\pi = (1, \delta_0).$$

- A *state* node v takes as input a pair (p, D) and returns a pair (p, D') depending on its label. If it is labelled by *skip*, it returns (p, D) . If it is labelled by $x_i = E(z)$, it returns (p, D') with D' the distribution of the vector $x[x_i = E(z)]$:

$$\text{if } v: \textit{state} \text{ then } \llbracket v \rrbracket_\pi(p, D) = \begin{cases} (p, D) & \text{if } v \text{ is labelled by skip} \\ (p, D') & \text{if } v \text{ is labelled by } x_i = E(z). \end{cases}$$

- A *test* node v labelled by $B(z)$ takes as input a pair (p, D) returns (p', D') depending on the value of $\text{cond}(s_\pi(v))$. In particular, first, the augmented vector z and its distribution D_z are considered. If $\text{cond}(s_\pi(v)) = \textit{true}$, then the node computes the probability of the Boolean condition evaluating to true, i.e. $P_{D_z}(\llbracket B(z) \rrbracket)$. Then, it conditions the current distribution to such event, i.e. $D_z \mid \llbracket B(z) \rrbracket$. The result is the output pair $(p \cdot P_{D_z}(\llbracket B(z) \rrbracket), \text{Marg}_x(D_z \mid \llbracket B(z) \rrbracket))$. Similarly, if $\text{cond}(s_\pi(v)) = \textit{false}$ the output is $(p \cdot P_{D_z}(\llbracket \neg B(z) \rrbracket), \text{Marg}_x(D_z \mid \llbracket \neg B(z) \rrbracket))$. To overcome conditioning with respect to zero-probability events we assume that whenever $P_{D_z}(\llbracket B(z) \rrbracket) = 0$ (resp. $P_{D_z}(\llbracket \neg B(z) \rrbracket) = 0$) the output pair is $(0, D)$:

$$\text{if } v: \textit{test} \text{ then } \llbracket v \rrbracket_\pi(p, D) = \begin{cases} (p \cdot P_{D_z}(\llbracket B(z) \rrbracket), \text{Marg}_x(D_z \mid \llbracket B(z) \rrbracket)) & \text{cond}(s_\pi(v)) = \textit{true} \wedge P_{D_z}(\llbracket B(z) \rrbracket) \neq 0 \\ (p \cdot P_{D_z}(\llbracket \neg B(z) \rrbracket), \text{Marg}_x(D_z \mid \llbracket \neg B(z) \rrbracket)) & \text{cond}(s_\pi(v)) = \textit{false} \wedge P_{D_z}(\llbracket \neg B(z) \rrbracket) \neq 0 \\ (0, D) & \textit{else}. \end{cases}$$

- For an *observe* node v labelled by $B(z)$ we condition the current distribution to $\llbracket B(z) \rrbracket$. Observe that if $B(z)$ only contains read-only variables, conditioning does not affect the distribution of the output variables x . If $B(z) = x_i == c$ the node returns a distribution D' having pdf $\frac{1}{I} f_D(x, x_i = c) \delta_c(x_i)$ with $I = \int_{\mathbb{R}^{d-1}} f_D(x, x_i = c) d(x \setminus x_i)$. In all other cases conditioning is treated as usual:

$$\text{if } v: \textit{observe} \text{ then } \llbracket v \rrbracket_\pi(p, D) = \begin{cases} (p \cdot I, D') & B(z) = x_i == c \\ (p \cdot P_{D_z}(\llbracket B(z) \rrbracket), D \mid \llbracket B(z) \rrbracket) & B(z) \neq x_i == c \wedge P_{D_z}(\llbracket B(z) \rrbracket) > 0 \\ (0, D) & \textit{else} \end{cases}$$

- The *exit* node v takes as input (p, D) and outputs the same pair (p, D) :

$$\text{if } v: \textit{exit} \text{ then } \llbracket v \rrbracket_\pi(p, D) = (p, D).$$

The semantics of the program P is then defined as:

$$\llbracket P \rrbracket = \sum_{\substack{(p, D) = \llbracket \pi \rrbracket \\ \pi \in \Pi^P}} \frac{p}{\sum_{\substack{(p', D') = \llbracket \pi \rrbracket \\ \pi \in \Pi^P}} p'} D. \quad (4)$$

Example 3.1. For the program in Algorithm 1 we have only two paths, $\pi_T = v_0v_1v_2v_3v_5v_6v_7$ and $\pi_F = v_0v_1v_2v_4v_5v_6v_7$ corresponding to evaluations of the conditional statement as true or false, respectively. To compute the semantics of π_T we start from v_0 , that outputs $(1, \delta_0)$. This pair is taken as input by v_1 , which is a state node assigning $\text{gauss}(1, 0)$ to x_1 , so its output is $(1, \mathcal{N}(0, \Sigma_1))$ with $\Sigma_1 = \text{diag}(1, 0)$ (corresponding to the distribution in Figure 2b). This pair is taken as input by the test node v_2 . Since we are considering π_T , for which $s_{\pi_T}(v_2) = v_3$, and $\text{cond}(v_3) = \text{true}$, the semantics of v_2 in this path conditions $\mathcal{N}(0, \Sigma_1)$ to $x_1 > 0$. Therefore in this path the output of v_2 is $(0.5, \mathcal{N}(0, \Sigma_1) \llbracket x_1 > 0 \rrbracket)$, where $0.5 = P_{\mathcal{N}(0, \Sigma_1)}(\llbracket x_1 > 0 \rrbracket)$. This pair is taken as input by v_3 , which updates the distribution of x_2 and therefore outputs a new pair $(0.5, D_3)$. We can proceed until we compute the output of v_7 , which gives the final pair $\llbracket \pi_T \rrbracket = (p_T, D_T)$. In the same way, we compute $\llbracket \pi_F \rrbracket = (p_F, D_F)$, and finally, the semantics of the whole program as the mixture $(p_T D_T + p_F D_F) / (p_T + p_F)$.

4 GAUSSIAN SEMANTICS

Gaussian Semantics is a family of semantics closed with respect to GMs. Each node takes as input and returns a GM over the program variables. This is done by composing the exact semantics of a node with an operator T_r^{GM} acting on the output distribution of $\llbracket v \rrbracket_\pi$. In particular, T_r^{GM} transforms any distribution D into a GM G , having the same moments of D up to order r . Therefore, we call T_r^{GM} the *moment-matching operator*. Formally, we use a map $R : V \rightarrow \mathbb{N}_0$ to associate each node $v \in V$ with the highest order of moments that will be matched at v . The semantics of a node is then:

$$\llbracket v \rrbracket_\pi^R = \begin{cases} \llbracket v \rrbracket_\pi & \text{if } v : \text{entry, exit} \\ (\mathbb{I}, T_{R(v)}^{GM}) \circ \llbracket v \rrbracket_\pi & \text{if } v : \text{state, test, observe} \end{cases} \quad (5)$$

where \mathbb{I} is the identity acting on the first element of the pair (p, D) , and $T_{R(v)}^{GM}$ is the moment-matching operator. The Gaussian Semantics of paths and programs are defined similarly to the exact semantics as:

$$\llbracket \pi \rrbracket^R = \llbracket v_n \rrbracket_\pi^R \circ \dots \circ \llbracket v_0 \rrbracket_\pi^R \quad \text{and} \quad \llbracket P \rrbracket^R(D_0) = \sum_{\pi \in \Pi^P} \frac{p}{\sum_{\substack{(p', D') = \llbracket \pi \rrbracket^R \\ \pi \in \Pi^P}} p'} D.$$

Example 4.1. We have seen in Algorithm 1 that the exact semantics is not closed with respect to GMs. For example, node v_2 takes as input a Gaussian but returns a truncated Gaussian, which is not a GM. If, instead, we consider Gaussian Semantics with $R(v_2) = 3$ the output of v_2 will be a GM matching the first three order moments of $\mathcal{N}(0, \Sigma_1) \llbracket x > 0 \rrbracket$. Two steps are required to compute $\llbracket v_2 \rrbracket_\pi^R$. In the first step, we compute the first $R(v_2) = 3$ order moments of the output distribution of $\llbracket v_2 \rrbracket_\pi$. Therefore, we compute the first order moments $\mathbb{E}[x_1] = 0.7979$ and $\mathbb{E}[x_2] = 0$, the second order moments $\mathbb{E}[x_1^2] = 1$, $\mathbb{E}[x_1 x_2] = \mathbb{E}[x_2^2] = 0$ and the third order moments $\mathbb{E}[x_1^3] = 1.5958$, $\mathbb{E}[x_1^2 x_2] = \mathbb{E}[x_1 x_2^2] = \mathbb{E}[x_2^3] = 0$. Observe that, thanks to the results in Table 1, this is significantly easier than computing the whole pdf of the output distribution. The second step involves finding a GM having the computed moments. This is generally more complex and is performed by the operator T_r^{GM} . In the rest of the section, we will assume that the computation of the moments is done using the aforementioned formulas, and we focus on the definition of the operator T_r^{GM} and the derivation of its properties.

Moment-Matching Operator. In general, the operator $T_r^{GM}(D)$ acts on distributions $D = \sum_{i=1}^C p_i D_i$ that are mixtures (possibly of a single component). The moments of D are computed as a linear combination of the moments of its components: if $x \sim D = \sum_{i=1}^C p_i D_i$ and $x_i \sim D_i$, then $\mathbb{E}[x^n] =$

$\sum_{i=1}^C p_i \mathbb{E}[x_i^n]$. Therefore, when computing the moments of D , we first compute the moments of every component D_i . Then, it makes sense to define $T_r^{GM}(D)$ so that when it acts on a mixture D ($C > 1$), it recursively acts on each component of the mixture, moment-matching each of them. When, instead, T_r^{GM} acts on a non-mixture distribution D ($C = 1$), it returns a GM having moments up to order r equal to those of D . This second action is encoded by a second operator **match** _{r} .

$$T_r^{GM}(D) = \begin{cases} p_1 T_r^{GM}(D_1) + \dots + p_C T_r^{GM}(D_C) & \text{if } D = p_1 D_1 + \dots + p_C D_C \text{ and } C > 1 \\ \mathbf{match}_r(D) & \text{otherwise.} \end{cases} \quad (6)$$

We require that **match** _{r} (D) satisfies the following two conditions:

- R1) for any distribution D , **match** _{r} (D) is a GM;
- R2) **match** _{r} (D) has central moments up to order r equal to those of D .

The existence of the operator **match** _{r} is guaranteed by the following result, derived from [Schmüdgen \[2017, Theorem 17.2\]](#), stating that that, for any finite sequence of moments, there exists a moment-matching discrete distribution putting positive mass on a number of points smaller than or equal to the number of matched moments. Since discrete distributions are GMs, the Proposition holds. For detailed proof, see the Supplementary Material.

PROPOSITION 4.2. *For any $r \in \mathbb{N}_0$, there exists an operator **match** _{r} satisfying R1 and R2.*

Example 4.3. In our example, we want to match a total of 10 moments, a zeroth-order moment (which is always 1), two first-order, three second-order, and four third-order moments. Theorem 17.2 in [Schmüdgen \[2017\]](#) ensures that there exists at least one discrete distribution (therefore a GM) with $C \leq 10$ components that has exactly the given moments.

Proposition 4.2 ensures the existence of at least one GM matching the moments of D up to order r . In general, letting $GM_r(D)$ denote the set of all finite GMs matching the moments of D up to order r , we may have that $|GM_r(D)| > 1$. For **match** _{r} to be well-defined, we need to uniquely identify a moment-matching GM in $GM_r(D)$. This can be done using different heuristics: we propose one based on the principle of maximum entropy, which we call the max entropy matching (MEM).

Max Entropy Matching. MEM can be summed up as follows: if $|GM_r(D)| > 1$ we choose the GM G having the least number of components (in order to minimize the number of parameters to be fit) and minimizing a certain cost function. Any remaining tie is resolved by comparing the vectors of parameters P that identify the GMs with respect to lexicographic ordering (we give an ordering on the parameters of GMs in the Supplementary Material). We select our cost function as the sum of the opposite of the differential entropy plus a penalty term, where the differential entropy for a distribution D with pdf f_D is defined as [\[Cover 1999\]](#)

$$H(D) = - \int_{\mathbb{R}^d} f_D(x) \log(f_D(x)) dx. \quad (7)$$

Intuitively, the *principle of maximum entropy* asserts that the distribution maximizing entropy is the one that minimizes the number of assumptions on the distribution [\[Cover 1999\]](#). Therefore, maximizing $H(D)$ we are choosing the most general moment-matching distribution. We add to $H(D)$ a penalty term to avoid uncontrolled growth of the parameter values.

Then, the procedure to compute **match** _{r} (D) is the following.

- 1) Find $c^* = \min \{c : \exists G = \sum_{s=1}^c p_s \mathcal{N}(\mu_s, \Sigma_s) \in GM_r(D)\}$.
- 2) Find the set \mathcal{P} such that $P = (p_1, \dots, p_{c^*}, \mu_1, \dots, \mu_{c^*}, \Sigma_1, \dots, \Sigma_{c^*}) \in \mathcal{P}$ if and only if the GM with parameters P matches the moments of D up to the r -th order.

3) Find the set \mathcal{P}^* given by:

$$\mathcal{P}^* = \arg \min_{\mathcal{P}} \left\{ -H \left(\sum_{i=1}^{c^*} p_i \mathcal{N}(\mu_i, \Sigma_i) \right) + \sum_{i=1}^{c^*} (\|\mu_i\|_2^2 + \|\Sigma_i\|_2^2) \right\}. \quad (8)$$

4) If $|\mathcal{P}^*| > 1$ choose $P^* \in \mathcal{P}^*$ maximum with respect to lexicographic ordering.

The following proposition guarantees that MEM leaves us with a well-defined operator **match_r**. It is again proved using Theorem 17.2 from [Schmüdgen \[2017\]](#), and by noticing that \mathcal{P} is a compact set, therefore Eq. (8) is well-defined. Again, we defer a detailed proof to the Supplementary Material while we explain how MEM works using an example.

PROPOSITION 4.4. *For any D and r the max entropy matching uniquely identifies **match_r**(D).*

We remark that the choice of MEM is arbitrary, as other cost functions could be introduced. However it has various benefits. (i) To guarantee that Proposition 4.4 holds, one needs a bounded cost function. (ii) Using entropy leads to a parallelism with VI: SOGA itself can be seen as a form of VI since it involves the minimization of the reverse differential entropy [[Kullback and Leibler 1951](#)]. However, correspondence with VI is lost when higher-order moments are considered, because the minimizer of the reverse differential entropy is not analytically expressible for GMs. (iii) In the spirit of minimizing the number of assumptions made on the approximating distribution, the approach looks more pleasing mathematically.

Example 4.5. While [Schmüdgen \[2017\]](#) ensures that we can find a moment-matching GM with 10 components, it is easy to check that $c^* = 2$ is the minimum number of components required to match three order moments. In fact, $c^* > 1$, since for a single Gaussian, given the mean and the covariance matrix, all the other moments are fixed (so, we can match the first two order moments but not the third). For $c = 2$ instead we can consider the GM G with pdf $p\mathcal{N}(m, S) + (1-p)\mathcal{N}(m', S')$ such that m, m', S, S' satisfy the following system:

$$\begin{aligned} pm_i + (1-p)m'_i &= \mathbb{E}[x_i] && \text{for } i = 1, 2 \\ p(m_i^2 + S_{i,i}) + (1-p)(m_i'^2 + S_{i,i}') &= \mathbb{E}[x_i^2] && \text{for } i = 1, 2 \\ p(m_1m_2 + S_{1,2}) + (1-p)(m_1'm_2' + S_{1,2}') &= \mathbb{E}[x_1x_2] \\ p(m_i^3 + 3m_iS_{i,i}) + (1-p)(m_i'^3 + S_{i,i}') &= \mathbb{E}[x_i^3] && \text{for } i = 1, 2 \\ p(m_1S_{2,2} + 2m_2S_{1,2} + m_1m_2^2) + (1-p)(m_1'S_{2,2}' + 2m_2'S_{1,2}' + m_1'm_2'^2) &= \mathbb{E}[x_1x_2^2] \\ p(2m_1S_{1,2} + m_2S_{1,1} + m_1^2m_2) + (1-p)(2m_1'S_{1,2}' + m_2'S_{1,1}' + m_1'^2m_2') &= \mathbb{E}[x_1^2x_2] \\ 0 < p < 1, \quad S_{1,1}S_{2,2} - S_{1,2}^2 \geq 0, \quad S_{1,1}'S_{2,2}' - S_{1,2}'^2 \geq 0 \end{aligned}$$

In the system we equate the moments of G (l.h.s) with those of $x \sim D$ (r.h.s, computed in Example 4.1). Moreover, we look for solutions such that $0 < p < 1$ and S, S' are positive semidefinite (last line). Since the system is polynomial, using SMT solvers over reals we can check it is satisfiable; therefore, $c^* = 2$. Now we should determine the set \mathcal{P} of all solutions and find those that minimize the cost function. Since finding all solutions is generally impossible, we directly proceed to optimize our cost function numerically, constraining the variables to satisfy the previous system. We find the approximate solution $p = 0.572$, $m = (0.471, 0)$, $m' = (1.236, 0)$ and $S = \text{diag}(0.066, 0)$, $S' = \text{diag}(0.426, 0)$. The approximating GM is shown in the green line of Figure 4, while the blue line shows the true non-Gaussian distribution.

The example shows that the difficult step in computing a Gaussian Semantics of arbitrary order is 2). Finding the parameters of a moment-matching GM requires the solution of a system of

polynomial equations, like the one in the example. This problem is notoriously hard to solve, as no analytical solution exists [Lasserre 2009]. Performing numerical optimization can solve the problem approximately, but is in general numerically unstable and requires relatively long computational times (in our example, using `scipy` [Virtanen et al. 2020], it took around 7 s to match a single Gaussian!). While we leave open the problem of solving 2) efficiently in the general case, the following lemma gives two important properties of \mathbf{match}_r , which will be used to derive our second-order approximation. The proof is quite trivial and reported in the Supplementary Material.

LEMMA 4.6. *The following two properties hold:*

- i) when $r = 2$, $\mathbf{match}_2(D)$ is a single Gaussian distribution with mean and covariance matrix equal to those of D ;
- ii) if D is Gaussian, for any $r \geq 2$ $\mathbf{match}_r(D) = D$.

We conclude the section with a consequence of Lemma 4.6. It follows from ii) that T_r^{GM} has the desirable property of leaving GMs unaltered, i.e. if M is a GM $T_r^{GM}(M) = M$ for all $r \geq 2$. As a consequence, Gaussian Semantics coincides with the exact semantics for programs only involving GMs, and in particular, for programs involving only discrete distributions (mixtures of deltas).

PROPOSITION 4.7. *Let $P = (V, E)$ be such that every read-only variable in the program is a finite discrete distribution. Then, for any R , $\llbracket P \rrbracket^R = \llbracket P \rrbracket$.*

PROOF. Since truncations, linear combination and products of discrete distributions are discrete distributions, only discrete distributions are generated in the execution of $\llbracket P \rrbracket$. By Lemma 4.6 applying the moment-matching operator T_r^{GM} to them leaves them unaltered, so conclusion follows. \square

Example 4.8. Consider a second Gaussian semantics that maps v_2 to $R(v_2) = 2$. In this case, we want to match only the first two order moments, namely $\mathbb{E}[x_1], \mathbb{E}[x_2], \mathbb{E}[x_1^2], \mathbb{E}[x_2^2], \mathbb{E}[x_1x_2]$. As noticed before, in this case $c^* = 1$, since we can take the Gaussian $\mathcal{N}(\mu, \Sigma)$ with $\mu = (0.7979, 0)$ and $\Sigma = \text{diag}(1, 0)$ and it will have required moments. Observe that we do not need to solve any system or optimization problem.

In general, for a fixed number of moments matched, we expect Gaussian Semantics to approximate reasonably well the matched moments but not necessarily the whole distribution (see Section 7.2 for further discussion). Indeed, let us compare the exact posterior distribution with the ones obtained distribution when $R = 2, 3, 4$, respectively, using Kullback-Leibler (KL) divergence [Kullback and Leibler 1951]. The KL divergence between two distributions P and Q is a standard way to evaluate the error committed in approximating P with Q . In our case, we take as P the truncated Gaussian and as Q the GMs obtained matching different order moments. The respective values are 2.29 ($R=2$), 1.71 ($R=3$) and 1.36 ($R=4$), so indeed higher-order Gaussian semantics improve the approximation. Figure 4 compares the true marginal pdf of x_1 at node v_2 (blue solid line) with the second- (orange dashed), third- (green dash-dotted) and fourth (red dotted) approximations. The advantages of fitting only a small number of moments are mainly computational. Indeed, it can be seen from the legend that as the number of moments matched grows, the increased KL accuracy comes with an increased computational cost.

5 UNIVERSAL APPROXIMATION THEOREM

Our main convergence result states that, for well-behaved programs, it is possible to find a map R so that the output distribution yielded by the semantics $\llbracket \cdot \rrbracket^R$ is arbitrarily close to that of $\llbracket \cdot \rrbracket$ in the Levy-Prokhorov metric [Ethier and Kurtz 2009]. By “well-behaved” we mean that the distributions in the exact semantics are determined by their moments and that they can measure continuously




Fig. 4. Marginal pdf of x_1 at node v_2 in Algorithm 1 given by exact semantics (*true*) and the Gaussian Semantics with $R = 2, 3$ and 4 . In the legend we report the KL divergence with the true distribution and the time need to compute the approximating GM.

sets in the form (3). To formalize the latter requirement we introduce m -continuity sets. Since both this definition and that of Levy-Prokhorov metric are borrowed from measure theory we refer the reader to the Supplementary Material for a more detailed background on these concepts.

Definition 5.1. Given a measure m , a set $A \subseteq \mathbb{R}^d$ is called an m -continuity set if $m(\partial A) = 0$ where, ∂A is the boundary of set A , defined as the closure of the set A minus its interior.

Then, we can state our main theorem.

THEOREM 5.2. *Assume that $P = (V, E)$ is a program such that for each $v \in V$ and each path $\pi \in \Pi^P$ the output distribution D of $\llbracket v \rrbracket_\pi$ satisfies the following:*

- H1) D is determined by its moments;
- H2) if D is the input distribution for a test or observe node v' , then the set defined by the LBC labelling v' is an m_{D_z} -continuity set.

Then there exists a sequence of maps $(R^k : V \rightarrow \mathbb{N}_0)_{k \in \mathbb{N}}$ such that:

$$\llbracket P \rrbracket^{R^k} \xrightarrow{k \rightarrow \infty} \llbracket P \rrbracket. \quad (9)$$

where the convergence is intended in the weak topology, or equivalently, in the Levy-Prokhorov metric.

5.1 Satisfaction of the Hypotheses

Before giving an outline of the proof, we briefly comment on the hypotheses.

First, observe that H1 and H2 are sufficient but not necessary. In particular, if the hypotheses of Proposition 4.7 are satisfied convergence holds trivially, even when H1 or H2 are violated.

Hypothesis H1 is common to other works considering moment-based approximation, such as in Bartocci et al. [2020] and Moosbrugger et al. [2022] and is needed to guarantee that the method of moments converges to the true distribution [Billingsley 2013]. For a given program, it is possible to perform static analysis to check whether the arising distributions are determined by their moments, exploiting known results on moment determinacy (see, for example, the moment-generating function characterization in [Billingsley 2013]). Notably, to apply these results is not necessary to compute the exact pdf of the arising distributions, but it is sufficient to keep track of their

type. In fact, for some classes of distributions, moment-determinacy is established: this is true for finite discrete distributions, Gaussians, uniforms, Poissons, exponentials, truncations and mixtures thereof [Billingsley 2013]. The case studies analyzed in this paper feature such distributions. On the contrary, log-normal distributions are not determined by their moments. However, as long as moments are computable, Gaussian Semantics can still be applied: in this case no formal guarantee of convergence towards the true distribution is given, but the method still provides an analytical approximation for the moments.

Hypothesis H2 guarantees that when distributions are conditioned to sets in the form (3), weak convergence is preserved. This requirement can be falsified if D has degenerate components that place positive probability mass on the boundary of the set defined by the LBC. This could happen, for instance, if a component is a Dirac measure centered on any point of $\partial\llbracket B(x) \rrbracket$. For example, consider line 12 of *Tracking_n* in Section 1.1, where *out* can be 1 or 0 with probability > 0 . This falsifies H2. However, such cases can be statically detected and the program can be transformed into one that uses the equivalent condition as $\textit{out} > 0.5$, so that H2 holds. More in general, continuity corrections such as those performed in Laurel and Misailovic [2020] can be adopted.

Example 5.3. Algorithm 1 satisfies H1 since the joint at each location is either a mixture of Gaussians or truncated Gaussians, for which moment-determinacy is known. Moreover, the two LBCs checked in the program are $x_1 > 0$ (line 3) and $x_2 < 3$ (line 9). Before checking $x_1 > 0$, x_1 has non-degenerate Gaussian distribution, and therefore the set $x_1 = 0$ (border of $\llbracket x_1 > 0 \rrbracket$) has measure 0. Similarly, before line 9, x_2 is Gaussian-distributed with $\sigma > 0$, therefore $x_2 = 3$ (border of $\llbracket x_2 < 3 \rrbracket$) has again probability 0. We conclude that also H2 is verified.

For *Tracking_n* in Section 1.1, each marginal is obtained by Gaussians, performing sums, squares, or conditioning. Here, to check moment-determinacy, we use a result in Billingsley [2008], which states that if the moment-generating function (mgf) of a distribution is defined in an interval of 0, then the distribution is determined by its moments. Using symbolic integration, we can compute the mgf of the product of two Gaussians and verify that it is defined in an interval of 0. Therefore, the distribution is determined by its moments and H1 holds. For H2, we have already shown how to correct the condition in line 12 so that H2 is verified. For the if statement in line 7, observe that before entering it, the marginal w.r.t. to *dist* is the distribution of $x^2 + y^2$, which is continuous, and therefore the point 10 has measure 0 with respect to it. Therefore also H2 holds.

5.2 Outline of the Proof

The proof of Theorem 5.2 is rather technical and involves a number of results from measure theory. Thus, here we provide a sketch, and refer the reader to Supplementary Material for a detailed proof.

First of all, recall that the semantics of a program is defined as a (finite) mixture of the semantics of the paths in the program. Therefore, (9) holds if it holds for every path $\pi \in \Pi^P$. Moreover, it can be shown that we can consider only the paths such that $\llbracket \pi \rrbracket = (p, D)$ with $p \neq 0$, since paths for which $p = 0$ contribute to the semantics of the program nor in the exact, neither in the Gaussian Semantics. So, the proof amounts to showing that for every path $\pi \in \Pi^P$ such that $\llbracket \pi \rrbracket = (p, D)$ with $p \neq 0$ we can choose a sequence of maps R^k such that $\llbracket \pi \rrbracket^{R^k} \rightarrow \llbracket \pi \rrbracket$.

We can build the sequence of maps R^k by specifying $R^k(v)$ for each $v \in \pi$. In particular, for v_0 we can choose any value of R^k . For v_1 , if it is not the exit node, we can assume that $\llbracket v_1 \rrbracket$ transforms the pair $(1, \delta_0)$ into (p_1, D_1) with $p_1 \neq 0$. Then, by definition of Gaussian Semantics, $\llbracket v_1 \rrbracket^{R^k}$ transforms $(1, \delta_0)$ into the pair $(p_1, T_{R^k(v_1)}^{GM}(D_1))$. To ensure convergence in this case, we use Theorem 5.4 of Billingsley [2008], which we state below.

THEOREM 5.4 (BILLINGSLEY [2008]). *Suppose $X \sim D$, $X_n \sim D_n$, and D is a distribution determined by its moments, while D_n has have moments of all orders. If for all $r > 0$*

$$\lim_n \mathbb{E}[X_n^r] = \mathbb{E}[X]$$

then $D_n \rightarrow D$ in the Levy-Prokhorov metric.

Then, if $R^k(v_1)$ is an increasing sequence in k (for example $R^k(v_1) = k$) we can use Theorem 30.2 of Billingsley [2013] and H1 to say that $\llbracket v_1 \rrbracket^{R^k}(1, \delta_0) \rightarrow \llbracket v_1 \rrbracket(1, \delta_0)$.

For v_2 and for any node v_i after that, we assume that, in the exact semantics, the node takes as input a pair (p, D) such that $p \neq 0$, while in the Gaussian Semantics associated with R^k , it takes as input (p_k, D_k) such that $p_k \rightarrow p$ in \mathbb{R} and $D_k \rightarrow D$ in the Levy-Prokhorov metric. Then, we need to prove that we can choose $R^k(v_i)$ so that $\llbracket v_i \rrbracket^{R^k}(p_k, D_k) \rightarrow \llbracket v_i \rrbracket(p, D)$.

We do this in three steps. First, we prove that the exact semantics preserves the convergence, i.e. $\llbracket v_i \rrbracket(p_k, D_k) = (p'_k, D'_k) \rightarrow (p', D') = \llbracket v_i \rrbracket(p, D)$ for each node type. Then, we use again Theorem 5.4 and H1 to say that $T_r^{GM}(D'_k) \rightarrow D'$ as $r \rightarrow \infty$. Finally, since $\llbracket v \rrbracket^{R^k}(p_k, D_k)$ is obtained from (p'_k, D'_k) applying the operator $(\mathbb{I}, T_{R^k(v_i)}^{GM})$, we can use a diagonal argument to ensure that we can fix $R^k(v_i)$ for each k so that $\llbracket v_i \rrbracket^{R^k} = (p'_k, T_{R^k(v_i)}^{GM}(D'_k)) \rightarrow (p', D')$.

The most industrious step is the first one, i.e. proving that $(p'_k, D'_k) = \llbracket v_i \rrbracket(p_k, D_k) \rightarrow (p', D') = \llbracket v_i \rrbracket(p, D)$ for each type of node. To do this, we use the Mapping Theorem [Billingsley 2008].

THEOREM 5.5 (MAPPING THEOREM). *Suppose $h : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is measurable and that the set H of its discontinuities is such that $m_D(H) = 0$. If $D_n \rightarrow D$ in the Levy-Prokhorov metric, then $h(D_n) \rightarrow h(D)$.*

In particular, when $v_i : \text{state}$, an assignment is of the type $x_j = E(x)$ with $E(x)$ in the form (1) (here we do not consider read-only variables for simplicity, but they can easily taken into account). Since $E(x)$ is continuous, convergence follows from Theorem 5.5 by taking

$$[h(x)]_k = \begin{cases} x_k & \text{if } k \neq j \\ E(x) & \text{if } k = j. \end{cases}$$

For $v_i : \text{test}$ and $v_i : \text{observe}$ we use again the Mapping Theorem, in two different ways. First, we use it with $h = \mathbb{I}_{\llbracket B(x) \rrbracket}$, where $B(x)$ is the LBC labelling the node, to show that $p'_k \rightarrow p'$. Then, we fix a point x^* such that $m_{D_k}(x^*) = 0$ for all k and we choose h to be:

$$[h(x')]_k = \begin{cases} x' & \text{if } x' \in \llbracket B(x) \rrbracket \\ x^* & \text{else} \end{cases}$$

which proves that $D'_k = D_k | \llbracket B(x) \rrbracket \rightarrow D' = D | \llbracket B(x) \rrbracket$. In both cases, H2 is fundamental in guaranteeing that the Mapping Theorem still holds.

6 SECOND ORDER GAUSSIAN APPROXIMATION

As discussed in Examples 4.5 and 4.8, while implementing an arbitrary order Gaussian Semantics may be difficult, it is straightforward to compute the Gaussian Semantics associated with $R = 2$ (i.e., such that at each node of the control-flow graph matches the first two order moments (mean and covariance matrix). We propose *Second Order Gaussian Approximation (SOGA)*, an algorithm that implements this particular case.

A prototype implementation can be found at <https://zenodo.org/records/10026970>.

Algorithm 2 *SOGA*(*node*)

```

1: input_list = []
2: for par in node.parents do
3:   input_list.append((par.p, par.dist))
4: end for
5: input_p, input_dist = merge_dist(input_list)
6: node.p, node.dist = node_semantics(node, input_p, input_dist)
7: for child in node.children do
8:   SOGA(child)
9: end for

```

6.1 Overview

In our implementation, SOGA accepts programs in a Python-like syntax, then compiled into a formal control-flow graph. SOGA recursively visits the nodes of the control-flow graph in a breadth-first fashion to compute the semantics of all paths. Furthermore, each node has two attributes, p and $dist$: p is a non-negative scalar proportional to the probability of reaching that node, while $dist$ stores the output distribution (in the form of a GM) computed by the semantics of that node.

The procedure applied for each node is summarized in Algorithm 2, where we assume that each node stores its parents and children in suitable attributes. When entering a new node, SOGA retrieves the pairs $(p, dist)$ computed by the parents of the current node and merges them in a single pair $(input_p, input_dist)$ using the function *merge_dist* (line 1-5). Then, it computes how the node semantics transforms the latter pair and stores the result in the attributes $node.p, node.dist$ (line 6). Finally, it calls itself recursively on the children node (lines 7-9). When the *exit* node is reached the algorithm ends, leaving the posterior distribution stored in its attribute $node.dist$.

The core of the algorithm is the function *node_semantics*, that, for each node type, transforms the pair $input_p, input_dist$ into a new pair $p, dist$. When $v: entry, exit$ *node_semantics* leaves the pair unaltered; when $v: state$ or $v: test, observe$ the functions *apply_rule*($input_dist, expr$) and *approx_trunc*($input_dist, trunc$) are invoked, respectively. We detail the functions below.

Function apply_rule. It implements the semantics of a *state* node. In particular, it takes as input the current mixture $input_dist$ and an expression $expr$ of type (1). It returns a new distribution $dist$ obtained applying $expr$ and T_2^{GM} to $input_dist$. To compute the moments of the transformed distribution $dist$, and therefore its second-order approximation, it uses the results in Table 1: when $expr$ is a linear transformation, it applies the formulas for the sum of multivariate Gaussians [Billingsley 2013]. When $expr$ involves products, it applies Isserlis' theorem [Wick 1950].

Function approx_trunc. It implements the semantics of a *test* or an *observe* node. It takes as input the current mixture $input_dist$ and a set $trunc$ defined by an LBC of type (2). It returns the probability mass p , given by the probability that $input_dist$ satisfies $trunc$, and a new mixture distribution $dist$, representing the GM approximating $input_dist$ conditioned to $trunc$. Again, it applies the results in Table 1 to compute $dist$: in particular, when the LBC expresses inequality constraints the formulas in Kan and Robotti [2017] are used; when instead the LBC has the form $x_i == c$ it uses the formulas from Bishop and Nasrabadi [2006].

Function merge_dist. Merging is performed whenever a node is accessed prior to applying its semantics: *merge_dist* collects all the output pairs (p, D) computed at the parent nodes, and merges them together in a single GM. Given the set of s parents' pairs $(p_1, D_1), \dots, (p_s, D_s)$, the function

returns probability mass $input_p = p_1 + \dots + p_s$ and a new GM $input_dist = \frac{1}{input_p} (p_1 D_1 + \dots + p_s D_s)$. For an *exit* node, the output of this function is the output distribution of the program.

6.2 Distributivity of Transfer Functions

SOGA explores the control-flow graph in a breadth-first fashion, performing merges when required. On the other end, the exact and the Gaussian semantics are defined as a sum over all execution paths, leading to an apparent discrepancy. To ensure that SOGA indeed computes the Gaussian Semantics associated with the map $R(v) = 2$ we show in Proposition 6.1 that the transfer function of the exact semantics is distributive with respect to the merge operation.

To do this, for a set of pairs (p_i, D_i) we define the merge operator

$$merge((p_1, D_1), \dots, (p_s, D_s)) = \left(\sum_{i=1}^s p_i, \sum_{i=1}^s \frac{p_i}{\sum_{j=1}^s p_j} D_i \right) = (P, D).$$

We show that computing the semantics of a node after performing a merge gives the same output distribution as computing the semantics of each pair and then merging the results. This distributivity transfers straightforwardly to Gaussian Semantics, since the latter is computed by composing the exact semantics with the operator T_r^{GM} , which is distributive with respect to merging by Definition 6. This, in turn, justifies computing the semantics exploring the control-flow graph in a breadth-first fashion as SOGA does.

PROPOSITION 6.1. *Let (p_i, D_i) be pairs with $p_i \geq 0$ and D_i a distribution for $i = 1, \dots, s$. Let v be a node of type state, test, observe or exit. Then*

$$merge(\llbracket v \rrbracket(p_1, D_1), \dots, \llbracket v \rrbracket(p_s, D_s)) = \llbracket v \rrbracket(merge((p_1, D_1), \dots, (p_s, D_s))) \quad (10)$$

PROOF. Let $(\tilde{p}_i, \tilde{D}_i) = \llbracket v \rrbracket(p_i, D_i)$ for $i = 1, \dots, s$. Then the L.H.S. of Equation 10 becomes

$$merge(\llbracket v \rrbracket(p_1, D_1), \dots, \llbracket v \rrbracket(p_s, D_s)) = \left(\sum_{i=1}^s \tilde{p}_i, \sum_{i=1}^s \frac{\tilde{p}_i}{\sum_{j=1}^s \tilde{p}_j} \tilde{D}_i \right) = (\tilde{P}, \tilde{D}).$$

For the R.H.S. let:

$$\llbracket v \rrbracket(merge((p_1, D_1), \dots, (p_s, D_s))) = \llbracket v \rrbracket(P, D) = (\hat{P}, \hat{D}).$$

Let us show for each type of node that $\tilde{P} = \hat{P}$ and $\tilde{D} = \hat{D}$. We observe that for v : *state*, with v labeled by *skip*, and for v : *exit* conclusion follows trivially. We examine the remaining cases separately.

- Let v : *state* and suppose v is labelled by $x_k := E(z)$ then $\tilde{p}_i = p_i$ and \tilde{D}_i is the distribution of $x[x_k = E(z)]$ where $x \sim D_i$. Then $\tilde{P} = P$ and $\tilde{D} = \sum_{i=1}^s \frac{p_i}{P} \tilde{D}_i$. On the other hand $\hat{P} = P = \tilde{P}$ and \hat{D} is the distribution of $y[y_k = E(z)]$ where $y \sim D = \sum_{i=1}^s \frac{p_i}{P} D_i$. Therefore $\hat{D} = \sum_{i=1}^s \frac{p_i}{P} \tilde{D}_i = \tilde{D}$.
- Let v : *test*. To ease the notation, let us assume $cond(s_\pi(v)) = true$ and $B(z) = B(x)$, but the argument works analogously in the other cases. In this case $\tilde{p}_i = p_i \cdot P_{D_i}(\llbracket B(x) \rrbracket)$ and $\tilde{D}_i = D_i \mid \llbracket B(x) \rrbracket$. Therefore $\tilde{P} = \sum_{i=1}^s p_i \cdot P_{D_i}(\llbracket B(x) \rrbracket)$ and $\tilde{D} = \sum_{i=1}^s \frac{p_i \cdot P_{D_i}(\llbracket B(x) \rrbracket)}{\tilde{P}} (D_i \mid \llbracket B(x) \rrbracket)$. On the other hand $\hat{P} = P \cdot P_D(\llbracket B(x) \rrbracket) = P \cdot \sum_{i=1}^s \frac{p_i}{P} P_{D_i}(\llbracket B(x) \rrbracket) = \sum_{i=1}^s p_i P_{D_i}(\llbracket B(x) \rrbracket) = \tilde{P}$. Moreover, $\hat{D} = D \mid \llbracket B(x) \rrbracket$. Therefore, \hat{D} has density

$$\frac{f_D(x) \mathbb{I}_{\llbracket B(x) \rrbracket}}{P_D(\llbracket B(x) \rrbracket)} = \sum_{i=1}^s \frac{p_i P_{D_i}(\llbracket B(x) \rrbracket)}{P_D(\llbracket B(x) \rrbracket)} \frac{f_{D_i}(x) \mathbb{I}_{\llbracket B(x) \rrbracket}}{P_{D_i}(\llbracket B(x) \rrbracket)}$$

Table 2. Function implementing node semantics in SOGA and SOGAprune. The input arguments *input_p*, *input_dist* are retrieved by the parent nodes' attributes *p*, *dist*. The input arguments *expr*, *trunc* and *K* are stored in node attributes when the cfg is compiled from the program script.

Type	Function	Input	Computing
<i>state</i>	<i>apply_rule</i>	<i>input_dist</i> , <i>expr</i>	First two order moments of the components of the distribution obtained applying <i>expr</i> to <i>input_dist</i>
<i>text</i> , <i>observe</i>	<i>approx_trunc</i>	<i>input_dist</i> , <i>trunc</i>	Probability mass (or normalization constant) and first two order moments of the components of the distribution obtained truncating <i>input_dist</i> to <i>trunc</i>
<i>prune</i>	<i>prune_dist</i>	<i>input_dist</i> , <i>K</i>	Distribution <i>input_dist</i> iteratively pruned until the number of components is $\leq K$

which is the same density as the one of $\sum_{i=1}^s \frac{p_i P_{D_i}(\llbracket B(x) \rrbracket)}{\tilde{p}} (D_i \mid \llbracket B(x) \rrbracket) = \tilde{D}$. Observe that we have assumed that for at least one i $\tilde{p}_i \neq 0$. However, if that is not the case $P_D(\llbracket B(x) \rrbracket) = 0$ and the conclusion still holds.

- Let v : *observe*. If $B(x)$ has a probability greater than zero conclusion follows as in the previous case. If $B(x)$ has the form $x_k == c$ we can use the same argument, but we need to replace $P_{D_i}(\llbracket B(x) \rrbracket)$ with the normalization constant $I_i = \int_{\mathbb{R}^{d-1}} f_{D_i}(x, x_i = c) d(x \setminus x_i)$.

□

6.3 SOGAprune

To improve the scalability of SOGA we propose a second version of the algorithm, called SOGAprune, in which the user can introduce at script level the instruction **prune**(K), K being an integer number. When the script is compiled in a cfg, the **prune** instruction is compiled in a new node of type *prune*. When accessed, the function `node_semantics` invokes the function `prune_dist(input_dist, K)`.

Function prune_dist. It prunes the current distribution *input_dist* to keep the number of its components below a user-specified bound K . The pruning is performed similarly to [Chaudhuri and Solar-Lezama \[2010\]](#). In particular, for each pair of components i, j in the input distribution *input_dist*, having mixing coefficients π_i, π_j , means μ_i, μ_j and covariance matrices Σ_i, Σ_j , we compute the mean $\mu' = \frac{\pi_i \mu_i + \pi_j \mu_j}{\pi_i + \pi_j}$ and the cost $cost(i, j) = \pi_i \|\mu' - \mu_i\| + \pi_j \|\mu' - \mu_j\|$.

After computing the cost for all pairs (i, j) such that $i \neq j$ and $i, j < K$, the pair (i, j) with minimal cost is substituted with a single component having mean μ' and covariance matrix Σ' with

$$\Sigma' = \frac{\pi_i}{\pi_i + \pi_j} \left(\Sigma_i + \mu_i^T \mu_i \right) + \frac{\pi_j}{\pi_i + \pi_j} \left(\Sigma_j + \mu_j^T \mu_j \right) - \mu'^T \mu'$$

Observe that μ' and Σ' are exactly the mean and the covariance matrix of the mixture $\frac{\pi_i}{\pi_i + \pi_j} \mathcal{N}(\mu_i, \Sigma_i) + \frac{\pi_j}{\pi_i + \pi_j} \mathcal{N}(\mu_j, \Sigma_j)$. This produces the best possible approximation of the two components [[Chaudhuri and Solar-Lezama 2010](#)]. The procedure is iterated until the number of components is less than K (observe that after the first two components have been merged into a new one, we need to recompute the cost only for the pairs in which the new component appears).

A summary of how the semantics of each node is implemented is reported in Table 2, while detailed algorithms for SOGA implementation can be found in the Supplementary Material.

6.4 Computational Cost

We first compute the computational cost without pruning, then we discuss how pruning affects it.

Cost without pruning. Let $|V|$ denote the total number of nodes, $|T|$ the number of *test* nodes, $|TO|$ the number of *test* and *observe* nodes and $|S|$ the number of *state* nodes. W.l.o.g. we assume for simplicity that all read-only variables are pushed to an initial distribution D_0 over \mathbb{R}^n ; thus the output of the entry node is $(1, D_0)$ and all assignments only use output variables. By doing this we compute an upper bound on the true computational cost since the dimensions corresponding to read-only variables are dropped after marginalization. Letting C_0 denote the number of components of D_0 , the output distribution will have at most $C \leq C_{max} = 2^{|T|} C_0$ components.

We consider the cost to access a node and perform elementary operations, such as assignments and products, constant. Expressions *expr* and *trunc* are assumed to be stored in suitable data structures accessible in constant time, so storage and reading of them are also considered elementary operations. Overall, elementary operations contribute to the total computational cost with a term $O(|V|)$, which is however dominated by the computational cost of executing *approx_trunc*, *apply_rule* and *merge_dist*. We examine their cost separately.

The function *approx_trunc* is invoked once when an *observe* node is accessed and twice when a *test* node is accessed, for the true and the false branch respectively. When $B(z)$ is in the form $c_1 \cdot z_1 + \dots + c_n \cdot z_n \triangleright c$ a singular value decomposition is performed to change coordinates, so that in the new set of coordinates the truncation set is a hyper-rectangle (cost $O(n^3)$, [Gu and Eisenstat 1995]). Then, a new mixing coefficient has to be computed for each component to convert the truncated GM into a mixture of truncated Gaussians (cost $O(n)$). Finally, for each truncated Gaussian, the first two order moments are computed using the formulas in Kan and Robotti [2017] (cost $O(n^4)$, for a detailed account see Supplementary Material). When $B(z)$ is in the form $x_i == c$, to apply the formulas in Bishop and Nasrabadi [2006], matrix multiplication must be performed, amounting to cost $O(n^3)$ [Skiena 2008]. Overall, we have a cost of $O(|TO|C_{max}n^4)$.

The function *apply_rule* is invoked every time a *state* node is accessed. Since affine transformations require matrix multiplication (cost $O(n^3)$), the total cost is $O(|S|C_{max}n^3)$.

Finally, the function *merge_dist* is invoked whenever a node is accessed and performs a scalar product. It contributes for a cost $O(|V|C_{max})$.

The total cost of SOGA is therefore

$$O(|TO|C_{max}n^4) + O(|S|C_{max}n^3) + O(|V|C_{max}) \leq O(|V|2^{|T|}C_0n^4), \quad (11)$$

that is, linear in the number of nodes $|V|$ and in the initial number of components C_0 , polynomial in the dimensionality of the augmented input space n and exponential in the number of *test* nodes $|T|$, i.e., linear in the number of paths.

Effect of pruning. Let us now consider the effect of introducing **prune**(K) instructions. Let $|P|$ be the number of *pruning* nodes and $|T|_{bet}$ be the maximum number of subsequent *test* nodes without *pruning* instructions between them. Then $|T|_{bet} \leq |T|$ and $|T|_{bet} = |T|$ if no pruning instructions have been introduced in the program. Then, the maximum number of components a mixture can have before pruning occurs is $C_{max} = K2^{|T|_{bet}}$ (assuming w.l.o.g. $|C_0| < K$).

The function *prune_dist* is invoked at most $|P|$ times. When invoked, it first computes the cost for all possible pairs of components, which is at most $C_{max}(C_{max} - 1)$. The computation of the cost function for each pair has cost $O(n)$, while the computation of the covariance matrix (cost $O(n^2)$) is performed for a single pair. At its first iteration, the computational cost of *prune_dist* is, therefore, $O(C_{max}^2 n)$. After this, new costs are computed for at most $C_{max} - K$ times, but each time only for $C < C_{max}$ pairs of components. The whole cost of the function is therefore $O(C_{max}^2 n)$.

Substituting in (11) one gets that the computational cost with pruning is bounded by:

$$O(|V|K2^{|T|_{bet}}n^4) + O(K^22^{2|T|_{bet}}n) \leq O(|V|K^22^{2|T|_{bet}}n^4). \quad (12)$$

Comparing (11) with (12) one can conclude that pruning is only effective in reducing the computational cost when the overhead introduced by pruning ($K^22^{2|T|_{bet}}$) is less demanding than dealing with the full space of paths ($C_02^{|T|}$). To keep the overhead contained one could use small values of K while keeping also $|T|_{bet}$ small (e.g. by introducing many *pruning* instructions). However, this introduces an additional level of approximation which can hinder the accuracy of SOGA.

7 NUMERICAL EVALUATION

We split the numerical evaluation into four parts. In Section 7.1 we compare SOGA with four baseline tools representative of different inference methods for estimating the posterior mean: STAN for MCMC [Carpenter et al. 2017], PSI for exact symbolic analysis [Gehr et al. 2016], AQUA for quantization of posterior distributions [Huang et al. 2021] and Pyro for VI [Bingham et al. 2019]. In Section 7.2 we compare SOGA against Pyro in performing Maximum A Posteriori (MAP) estimation [Gelman et al. 2013], to test how well our method is able to capture the posterior distribution, in addition to its moments. Finally, in Sections 7.3 and 7.4, we evaluate SOGA's performance on two applications that have been extensively studied in the literature, owing to their significant practical impact. The first application is inference on models involving mixtures of continuous and discrete distributions, as in Gao et al. [2017]; Kharchenko et al. [2014]; Pierson and Yau [2015]; the second application is Bayesian inference on collaborative filtering. [Zhao et al. 2013].

7.1 Posterior Mean Estimation

We start by comparing SOGA with STAN for MCMC [Carpenter et al. 2017], PSI for exact symbolic analysis [Gehr et al. 2016], and AQUA for quantization of posterior distributions [Huang et al. 2021] and Pyro for VI [Bingham et al. 2019]. We consider the case studies from these tool's reference papers [Carpenter et al. 2017; Gehr et al. 2016; Huang et al. 2021], excluding those which could not be encoded in our syntax. This choice is intended to stress SOGA in the analysis of programs that were not designed to enhance its properties. Overall out of 31 total models, 13 were left out: 9 because of non-parametrizable distributions depending on variable parameters, and 4 because of the presence of non-polynomial functions (taken from: STAN - 1, PSI - 3, AQUA - 8, Pyro - 1). The remaining 18 models can be found in Carpenter et al. [2017] (*Bernoulli*), Gehr et al. [2016] (*BayesPointMachine*, *Burglar*, *ClickGraph*, *ClinicalTrial*, *CoinBias*, *DigitRecognition*, *Grass*, *MurderMystery*, *NoisyOr*, *SurveyUnbias*, *TrueSkills*, *TwoCoins*) and Huang et al. [2021] (*Altermu*, *Altermu2*, *NormalMixtures*, *RadarQuery*, *TimeSeries*).

The considered programs are listed in Table 3. Pruning was applied after every test and observe nodes (repeating it only once if they occur subsequently) for programs whose computation time was greater than 1 s and at least ten times larger than the worst performing tool. We set $K = 0.015C_{max}$ except for *NormalMixtures*; there, since C_{max} exceeded the tens of thousands, we set $K = 30$. With this strategy, the pruning algorithm was invoked only in 4 out of the 18 considered programs.

The experiments were performed on a laptop equipped with a 2.8 GHz Intel i7 quad-core processor and 16 GB RAM, CmdStan v2.30.1 and Wolfram Mathematica 13.1 (Wolfram Research, Inc.), setting a time-out threshold at 600 s.

7.1.1 Results. Table 3 collects the results where *time* refers to the average runtimes (in seconds) out of 10 executions and *value* refers to the computed expected value of a target variable in the model. For each model we specify the kind of distributions involved: B=Bernoulli, Be=Beta, D=Discrete, G(*) =Gaussian (with non-constant mean), U=Uniform. For STAN, we indicate the

Table 3. Results using STAN, PSI, AQUA and SOGA. ‘—’: discrete posterior not supported; ‘mem’: out of memory error; ‘err’: tool returns error state. For SOGA, C: final number of components; d: dimensionality of the output vector.

Model	Dist.	STAN	AQUA	Pyro (VI)	PSI	SOGA		
		time	time	time	time	time		
		value	value	value	value	value		
						C		
						d		
Bernoulli	B,U	0.17	0.84	4.51	0.38	1.28*	27	
	G*	51.0	0.056	60.49	err	2.20	1	
BayesPointMachine	B	—	—	—	0.12	0.003	4	
	B,U	102	0.540	3.13	1.10	0.614	35	
ClickGraph	B,U	—	—	—	0.97	0.755	23	
	B,U	—	—	—	0.34	0.417	64	
ClinicalTrial	B,Be	0.07	0.420	0.91	err	4.46	10	
	D	—	—	—	0.08	0.708	28	
DigitRecognition	B	—	—	—	0.12	0.016	2	
	B	—	—	—	0.16	0.814	256	
Grass	B	—	—	—	18.5	0.800	128	
	B	—	—	—	0.05	104.7	1	
MurderMystery	B	—	—	—	0.10	0.333	3	
	B	—	—	—	to	0.16	0.000	1
NoisyOr	B	—	—	—	284	0.155	4	
	B	—	—	—	to	50.4*	0.298	30
SurveyUnbias	B,G,U	0.10	0.800	2.89	18.5	0.800	4	
	G*	0.04	104.0	1.30	to	0.05	104.7	1
Trueskills	B	—	—	101.4	0.10	0.333	3	
	B	—	—	—	to	0.16	0.000	1
TwoCoins	G*	19.0	0.009	33.1	to	0.16	0.000	1
	G*,U	15.0	0.170	5.50	284	0.155	4	
Altermu	G*,U	0.38	0.286	104.89	to	0.36	0.156	4
	G*,U	144	5.000	0.295	7.75	6.333	2016	
NormalMixtures	B,G*,U	0.37	-1.600	err	to	3.79	-1.590	19
	G*,U	0.37	-1.600	26.15	to	3.79	-1.590	19
RadarQuery	B,G*,U	0.90	6.333	err	7.75	6.333	2016	
	G*,U	1.67	-1.575	26.15	to	3.79	-1.590	19
Timeseries	B,G*,U	0.90	6.333	err	7.75	6.333	2016	
	G*,U	1.67	-1.575	26.15	to	3.79	-1.590	19

time needed to obtain a 5% confidence interval whose amplitude is contained in 1% of the mean (up to a maximum of 10^5 samples). For PSI we report the sum of the time needed to generate the symbolic formula and that needed to integrate it when in the presence of non-simplified integrals (observe that in the original paper, only the time for symbolic computation was considered). For VI, due to high sensitivity with respect to the hyperparameters [Hoffman et al. 2013], we proceed using three different learning rates (0.01, 0.005, 0.001), and we report the most accurate estimation (detailed results can be found in the Supplementary Material). The number of iterations of the stochastic gradient descent is increased from a minimum of 100 to a maximum of 10k, stopping the optimization if the difference between the estimated mean posterior and the mean posterior estimated 100 steps before is less than 1% of the current estimation. For SOGA, runtimes labeled with * indicate that the pruning algorithm was invoked. Finally, we highlight the fastest method with a grey background. For accuracy evaluation, we consider PSI's results as ground truth when available (i.e., when PSI terminates and the integration is successfully computed within the timeout threshold). We made this choice since PSI is an exact method and the only guaranteed to be exact among the evaluated tools.

Only in one example, *BayesPointMachine*, SOGA performs poorly in terms of accuracy, estimating a value of 0.011 for a parameter estimated by STAN as $0.054 \pm 2e-4$. We remark, however, that this program turned out to be particularly difficult to solve for AQUA (which issued an out-of-memory error) and PSI (which was not able to complete the symbolic computation of the posterior). On the other examples, SOGA yields very good accuracy, with a relative error below 7% across all comparable models. We now discuss a detailed comparison of runtimes against each tool method.

STAN. STAN does not support discrete posteriors; hence it could not analyze eight models. For the models that can be analyzed by both, SOGA outperforms STAN in terms of runtimes on *Altermu*, *Altermu2*, and *RadarQuery*. By contrast, STAN outperforms SOGA in *Bernoulli* and *NormalMixtures*. We attribute this to the presence of non-Gaussian priors and a large number of observations, resulting in a high number of components and truncations to be computed. Both have similar performance on the remaining models.

AQUA. SOGA is more flexible than AQUA in that it supports discrete posteriors. On *ClickGraph*, and *TrueSkills* AQUA issued an out-of-memory error while SOGA could approximate the posterior mean. We ascribe this issue to the fact that AQUA uses tensors, whose dimension rapidly increases with the number of distributions. In particular, in AQUA each distribution must be stored in the tensor, while SOGA can use fresh read-only variables which are dropped once the marginal over the output variables is evaluated. Notably, SOGA outperforms AQUA also on *Altermu*, *Altermu2* and *TimeSeries* proposed in the AQUA paper [Huang et al. 2021]. Instead, AQUA is more efficient than SOGA in *RadarQuery*, *Bernoulli* and *NormalMixtures*, for the same reasons explained for STAN.

Pyro. Being a gradient-based method, Pyro's VI offers limited support for discrete variables,¹ so that, similarly to STAN and AQUA, we were not able to encode models with discrete posterior. In addition, we found that the encoding of *RadarQuery* incurred runtime errors. For the remaining models, VI is comparable to SOGA, when not less accurate, and taking longer runtimes. Noticeable exceptions are *BayesPointMachine*, where, as already noticed, SOGA is not able to achieve a good accuracy and *ClickGraph*, where SOGA incurs long runtimes even with pruning. On the other hand, VI exhibits a significant sensitivity with respect to the choice of the hyperparameters, which can result in non-convergence and sloppy approximations for poor choices of the parameters (results for all the tested learning rates can be found in the Supplementary Material).

¹<https://pyro.ai/examples/enumeration.html>

Table 4. Comparison between Pyro and SOGA for MAP estimation. Models with ‘(P)’ were pruned when SOGA was applied. True values are derived optimizing the exact posterior, or from samples (denoted with ‘**’).

<i>Model</i>	<i>Pyro</i>		<i>SOGA</i>		<i>True value</i>
	<i>value</i>	<i>time</i>	<i>value</i>	<i>time</i>	
<i>Bernoulli</i>	0.200	0.20	0.220	12.0	0.200
<i>Bernoulli (P)</i>	0.200	0.20	0.290	1.28	0.200
<i>BayesPointMachine</i>	0.000	7.83	0.011	2.20	0.032 ± 0.002*
<i>ClickGraph</i>	0.501	2.98	0.861	208	1.000
<i>CoinBias</i>	0.400	0.63	0.493	0.61	0.400
<i>SurveyUnbias</i>	0.964	3.46	0.755	1.56	1.000
<i>TrueSkills</i>	101.6	0.99	104.7	0.05	104.8 ± 0.681*
<i>Altermu</i>	not converged		0.000	0.16	0.114 ± 0.092*
<i>NormalMixtures (P)</i>	0.236	49.5	0.276	50.4	0.275 ± 0.005*
<i>TimeSeries</i>	-1.564	55.4	-1.494	3.69	-1.694 ± 0.021

PSI. *PSI* outperforms *SOGA* on *Bernoulli*, *ClickGraph*, and *ClinicalTrial*. However, on six models (*SurveyUnbias*, *TrueSkills*, *Altermu*, *Altermu2*, *RadarQuery*, and *TimeSeries*) *PSI* timed out or resulted in long runtimes. This behavior can be explained by the presence of distributions dependent on variable parameters (*SurveyUnbias*, *Trueskills*, *Radar*) or by the high number of observations (*Altermu*, *Altermu2*, *TimeSeries*). In *Altermu*, *PSI* could not compute a symbolic formula within the time-out threshold, while in *BayesPointMachine*, *TrueSkills*, and *TimeSeries* the formula contained non-simplified integrals, whose integration in *Mathematica* took longer than the time-out threshold. Notably, for models involving only *Bernoulli* distributions (*Burglar*, *Grass*, *MurderMystery*, *NoisyOr*, *TwoCoins*), for which both tools are exact, their performance is comparable.

7.1.2 Performance of Pruning.

In the right inset we report the runtimes, values, and number of components *C* for *SOGA* without pruning applied to the four models that required the application of pruning (*Bernoulli*, *ClickGraph*, *ClinicalTrial*, *NormalMixtures*). All models share the occurrence of *GM* distributions with more than 1000 components. For *Bernoulli*, pruning allowed comparable runtimes with respect to the best-performing tool, while base *SOGA* was about 9 times slower (11.97 s). In addition, base *SOGA* computed an output indistinguishable from the pruned version up to the third decimal digit. For the other three cases, base *SOGA* was unable to compute a numerical result within the time-out threshold. For these, the number of components *C* is the one reached before timing out. Applying pruning allowed *SOGA* to complete the computation within the time-out threshold while achieving excellent accuracy with respect to the ground truth.

<i>Model</i>	<i>Time</i>	<i>Value</i>	<i>C</i>
<i>Bernoulli</i>	11.97	0.252	1774
<i>ClickGraph</i>	to		2304
<i>ClinicalTrial</i>	to		1508
<i>NormalMixtures</i>	to		97714

For *Bernoulli*, pruning allowed comparable runtimes with respect to the best-performing tool, while base *SOGA* was about 9 times slower (11.97 s). In addition, base *SOGA* computed an output indistinguishable from the pruned version up to the third decimal digit. For the other three cases, base *SOGA* was unable to compute a numerical result within the time-out threshold. For these, the number of components *C* is the one reached before timing out. Applying pruning allowed *SOGA* to complete the computation within the time-out threshold while achieving excellent accuracy with respect to the ground truth.

7.2 Maximum a Posteriori Estimation

Since *SOGA* approximates the posterior with a Gaussian mixture, it can also compute the Maximum a Posteriori (MAP) estimate by simply returning the mean of the *GM* component with the largest mixing coefficient. Here we compare its performance against *Pyro*, in which MAP estimation can be performed using a different parametrizing distribution than the one used for the mean posterior

inference.² To get a baseline for the MAP value, we first generate the symbolic posterior using PSI and then optimize it numerically. For models in which PSI is not able to compute the exact posterior, we estimate the ground truth by taking 10k samples from the posterior and binning them into 50 intervals; then, MAP is the midpoint of the interval with the most samples. We tested the same models with continuous posterior reported in Table 3, except *Altermu2*, since, by visual inspection, we found that it has a flat posterior.

Results are reported in Table 4. Due to Pyro’s sensitivity to hyperparameters observed in the previous section, we tested three different values of learning rate. Table 4 only reports the closest estimation to the baseline; full results are available in the Supplementary Material, confirming the sensitivity issues. These experiments show that SOGA performs relatively worse than in the estimation of the posterior mean. This is expected because SOGA is designed to match means and variances, but it does not necessarily approximate the whole distribution. However, compared to Pyro, it is still able to obtain the closest estimation for *BayesPointMachine*, *ClickGraph*, *TrueSkills*, *Altermu* and *NormalMixtures*, while it is outperformed by Pyro in *Bernoulli*, *CoinBias*, *SurveyUnbias* and *TimeSeries*. Finally, we note that analyzing *Bernoulli* with SOGAprune degrades the MAP estimation, unlike in the posterior mean.

7.3 Mixtures of Continuous and Discrete Distributions

Mixtures of continuous distributions and discrete probability masses appear in different domains such as in Gao et al. [2017]; Kharchenko et al. [2014]; Pierson and Yau [2015]. Languages such as STAN and AQUA do not support them. Ad hoc methods have been proposed in Tolpin et al. [2016] and Nitti et al. [2016]. More recently Wu et al. [2018] extended the sampling techniques used in BLOG for more accurate inference. We test SOGA on the three benchmarks proposed by Wu et al. [2018] and compare its runtimes against PSI, BLOG, and variable elimination (VE) as implemented in Pyro [Obermeyer et al. 2019]. *IndianGPA* and *Scale*

<i>Model</i>	<i>Runtimes (s)</i>			
	<i>SOGA</i>	<i>PSI</i>	<i>BLOG</i>	<i>VE</i>
<i>IndianGPA</i>	0.099	0.180	0.516	0.192
<i>Scale</i>	0.013	0.120	0.810	0.150
<i>Tracking_1</i>	0.042	to	0.803	0.143
<i>Tracking_5</i>	0.046	to	1.044	0.394
<i>Tracking_10</i>	0.046	to	1.330	0.670
<i>Tracking_50</i>	0.110	to	2.886	4.095
<i>Tracking_100</i>	0.192	to	5.054	8.885
<i>Tracking_150</i>	0.271	to	6.602	13.723

are reported exactly as in the original paper, while the *Tracking_n* example from Section 1.1 is adapted since it was originally cast as a control problem. All examples have a Dirac delta posterior, which is computed exactly by all. However, SOGA is the fastest and the one which scales better as the number of steps n increases.

7.4 Bayesian Inference for Collaborative Filtering

Collaborative filtering models are well-known in machine learning for applications to recommendation systems [Koren et al. 2021]. We target the problem of Bayesian inference on the latent factor model proposed in Hofmann and Puzicha [1999], which arises after a singular value decomposition and serves as the basis for solving an optimization problem [Zhao et al. 2013]. The model assumes noisy observations sampled from $\mathcal{N}(cf_k, 1)$ where cf_k has the form $cf_k = a_1b_1 + \dots + a_kb_k + c$, where a_i , b_i , and c are unknown latent variables. As noticed in Nishihara et al. [2013], performing Bayesian inference on these models is particularly difficult due to non-identifiability [Tsiatis 1975] and symmetry [Neal 1999] of the parameters. For example, switching the distributions of a_i and

²https://pyro.ai/examples/mle_map.html

Table 5. Runtimes (in seconds) and computed values for the collaborative filtering model $\mathcal{N}(cf_k, 1)$.

k	<i>Ground truth</i>	SOGA		STAN		AQUA		VI	
		<i>time</i>	<i>value</i>	<i>time</i>	<i>value</i>	<i>time</i>	<i>value</i>	<i>time</i>	<i>value</i>
1	2	0.16	1.86	0.94	1.90	1.64	1.83	18.90	1.79
2	25	0.18	24.28	4.87	24.00	mem		25.10	23.93
3	-5	0.19	-5.79	7.63	-5.80	mem		26.40	-5.82
5	-30	0.22	-31.98	7.22	-32.00	mem		23.19	-31.47
10	151	0.30	149.75	5.42	150.00	mem		20.04	146.39
20	70	0.60	73.76	14.10	74.00	mem		23.18	69.92

b_i will result in the same distribution for cf_k , which is the only one observed. In some cases, one may still want to model each parameter separately to allow for more flexibility. In this particular case, though not solving the problem of symmetry and non-identifiability, SOGA can estimate the distribution of cf_k faster than its competitors. Results are shown in Table 5 for various values of k . PSI results are not reported because the tool was able to produce a symbolic formula only up to $k = 3$; however, even in these cases, numerical integration of the non-simplified integrals required more than 600 s. Although STAN’s estimates are accurate and close to SOGA’s ones, its runtimes are longer due to the increased cost of sampling, which is exponential in the number of variables. As above, we attribute AQUA’s out-of-memory error to its tensor based representation. For VI, we report results for the learning rate 0.005, which we found to be the one performing best in average, among the tested ones. A full set of experiment results can be found in the Supplementary Material. VI exhibits an accuracy comparable to SOGA’s, but significantly longer runtimes. We observe, however, that thanks to vectorization, VI’s runtimes do not significantly increase with k . Overall, the excellent runtime performance of SOGA is due the particular structure of the models, which exhibit Gaussian posteriors on variables combined in a scalar product without introducing truncations that could slow down the computations.

8 FURTHER RELATED WORK

Inference. In addition to the techniques discussed earlier in this paper, volume computation can be quite efficient for discrete models [Fileri et al. 2013; Holtzen et al. 2020]; however, it cannot be applied to continuous distributions. All the mentioned methods use a pdf representation of the distributions. More recently, representations using generating functions have been investigated, but only for discrete distributions [Chen et al. 2022]. Finally, some approaches use moment-based invariants [Barthe et al. 2016; Bartocci et al. 2020; Chakarov and Sankaranarayanan 2014; Katoen et al. 2010; Moosbrugger et al. 2022]. While they share the idea of computing moments up to a certain order, they differ both with respect to the supported programs and the computed information, making a direct comparison difficult.

Universal Approximators. Our approach can be ascribed to the practice, common in many branches of mathematics, of studying *universal approximators*, whereby one shows that a given function belonging to a certain class is shown to be approximated, arbitrarily closely, by another family of (parameterized) functions. Notable examples are polynomials [Pérez and Quintana 2006], and neural networks [Hornik et al. 1989; Zhou 2020].

Gaussian Approximators. The approximation-by-Gaussian approach is also common to Laplace approximation [Tierney and Kadane 1986]. Laplace approximation is a mode matching strategy

and is more expensive computationally than VI, as it is based on an optimization process to find the mode. Generally, however, it is inferior to VI [Bishop and Nasrabadi 2006]. Another kind of Gaussian approximation is Gaussian Smoothing [Chaudhuri and Solar-Lezama 2010, 2011], although it does target neither probabilistic programs nor the inference problem.

9 CONCLUSIONS

Gaussian Semantics is a family of approximations parameterized by the moment order to match against a Gaussian mixture at each location of a probabilistic program. The universal approximation theorem states that such a family converges to the true semantics. Although, in principle, any program location could be treated with different moment-order matching, in practice this is a difficult problem that requires the solution of a system of nonlinear equations. While the system is guaranteed to have a solution, finding it using SMT solvers over reals or numerical methods yields poor results, due to long computational times and numerical instability. Therefore we leave open the general problem of implementing Gaussian Semantics for any order of moments. However, we provide an analytical method that matches second-order moments of the exact probabilistic semantics (SOGA). The numerical results for the case studies demonstrate high quality of the approximation and that SOGA complements state-of-the-art methods for probabilistic inference and in particular for inference on models with mixtures of discrete and continuous distributions and for Bayesian inference on collaborative filtering models. Due to the efficiency shown by SOGA, we believe that in these cases our method can effectively be used as an alternative to sampling.

As regards future work, while SOGA performed satisfactorily on all tested benchmarks, it could not be applied to some of the models from the same repositories, due to the limitations of our syntax. Extending the latter to include general distributions depending on non-constant parameters, unbounded loops and non-polynomial functions would widen its scope of applicability. A possible way to overcome the former restriction could be learning offline the approximating distributions as a function of the variable parameters, but how to do this efficiently is currently not clear, even though of great interest. For what concerns unbounded loops, we observed that for almost surely terminating programs, the loops can be unrolled for a finite number of iterations so that the error committed in the approximation is arbitrarily small. This suggests that increasing the number of unrolled iterations together with the number of moments matched should preserve our convergence theorem, even in the case of almost surely terminating unbounded programs. Similarly, one could exploit convergence results for polynomial approximations to extend the convergence result to sequences of polynomial programs that approximate programs featuring non-polynomial functions. We leave the possibility to explore these extensions of our convergence result in future work.

Finally, one might devise algorithms for higher-order moments. While an extension to *exact* higher-order moment matching seems hard, a relaxed moment problem could be defined as an optimization problem [Hansen 2010].

DATA AVAILABILITY STATEMENT

A replication package is openly available at <https://zenodo.org/records/10026970>.

ACKNOWLEDGMENT

This work was partially supported by the projects SERICS (PE00000014) and by Investment 1.5 Ecosystems of Innovation, Project Tuscany Health Ecosystem (THE, B83C22003920001) and Interconnected North-East Innovation Ecosystem (iNEST, ECS_00000043) under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU. We would like to thank Joost-Pieter Katoen for his feedback on a preliminary version of this paper and the anonymous reviewers for their valuable comments.

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), 1–30.
- Gilles Barthe, Thomas Espitau, Luis María Ferrer Fioriti, and Justin Hsu. 2016. Synthesizing probabilistic invariants via Doob's decomposition. In *International Conference on Computer Aided Verification*. Springer, 43–61.
- Ezio Bartocci, Laura Kovács, and Miroslav Stankovič. 2020. Mora-automatic generation of moment-based invariants. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 492–498.
- Patrick Billingsley. 2008. *Probability and measure*. John Wiley & Sons.
- Patrick Billingsley. 2013. *Convergence of probability measures*. John Wiley & Sons.
- Eli Bingham, Jonathan P Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul Szerlip, Paul Horsfall, and Noah D Goodman. 2019. Pyro: Deep universal probabilistic programming. *The Journal of Machine Learning Research* 20, 1 (2019), 973–978.
- Christopher M Bishop and Nasser M Nasrabadi. 2006. *Pattern Recognition and Machine Learning*. Vol. 4. Springer.
- Xavier Boyen and Daphne Koller. 1998. Tractable inference for complex stochastic processes. In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*. 33–42.
- Bob Carpenter, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. Stan: A probabilistic programming language. *Journal of Statistical Software* 76, 1 (2017).
- Arun Chaganty, Aditya Nori, and Sriram Rajamani. 2013. Efficiently sampling probabilistic programs via program analysis. In *Artificial Intelligence and Statistics*. PMLR, 153–160.
- Aleksandar Chakarov and Sriram Sankaranarayanan. 2014. Expectation invariants for probabilistic program loops as fixed points. In *International Static Analysis Symposium*. Springer, 85–100.
- Swarat Chaudhuri and Armando Solar-Lezama. 2010. Smooth interpretation. *ACM Sigplan Notices* 45, 6 (2010), 279–291.
- Swarat Chaudhuri and Armando Solar-Lezama. 2011. Smoothing a program soundly and robustly. In *International Conference on Computer Aided Verification*. Springer, 277–292.
- Mingshuai Chen, Joost-Pieter Katoen, Lutz Klinckenberg, and Tobias Winkler. 2022. Does a program yield the right distribution? Verifying probabilistic programs via generating functions. In *International Conference on Computer Aided Verification*. Springer, 79–101.
- Patrick Cousot and Radhia Cousot. 1977. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proceedings of the 4th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*. 238–252.
- Thomas M Cover. 1999. *Elements of information theory*. John Wiley & Sons.
- Stewart N Ethier and Thomas G Kurtz. 2009. *Markov processes: characterization and convergence*. John Wiley & Sons.
- Antonio Filieri, Corina S Păsăreanu, and Willem Visser. 2013. Reliability analysis in symbolic pathfinder. In *2013 35th International Conference on Software Engineering (ICSE)*. IEEE, 622–631.
- Weihao Gao, Sreeram Kannan, Sewoong Oh, and Pramod Viswanath. 2017. Estimating mutual information for discrete-continuous mixtures. *Advances in Neural Information Processing Systems* 30 (2017).
- Timon Gehr, Sasa Misailovic, and Martin Vechev. 2016. PSI: Exact symbolic inference for probabilistic programs. In *International Conference on Computer Aided Verification*. Springer, 62–83.
- Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. 2013. *Bayesian data analysis*. CRC press.
- Noah D Goodman, Vikash K Mansinghka, Daniel Roy, Keith Bonawitz, and Joshua B Tenenbaum. 2008. Church: a language for generative models. In *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence*. 220–229.
- Andrew D Gordon, Thomas A Henzinger, Aditya V Nori, and Sriram K Rajamani. 2014. Probabilistic programming. In *Future of Software Engineering Proceedings*. 167–181.
- Ming Gu and Stanley C Eisenstat. 1995. A divide-and-conquer algorithm for the symmetric tridiagonal eigenproblem. *SIAM J. Matrix Anal. Appl.* 16, 1 (1995), 172–191.
- Lars Peter Hansen. 2010. Generalized method of moments estimation. In *Macroeconometrics and Time series Analysis*. Springer, 105–118.
- W Keith Hastings. 1970. Monte Carlo sampling methods using Markov chains and their applications. (1970).
- Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. 2013. Stochastic variational inference. *Journal of Machine Learning Research* (2013).
- Thomas Hofmann and Jan Puzicha. 1999. Latent class models for collaborative filtering. In *IJCAI*, Vol. 99.
- Steven Holtzen, Guy Van den Broeck, and Todd Millstein. 2020. Scaling exact inference for discrete probabilistic programs. *Proceedings of the ACM on Programming Languages* 4, OOPSLA (2020), 1–31.
- Kurt Hornik, Maxwell Stinchcombe, and Halbert White. 1989. Multilayer feedforward networks are universal approximators. *Neural Networks* 2, 5 (1989), 359–366.

- 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*. Springer, 229–246.
- Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. 1999. An introduction to variational methods for graphical models. *Machine Learning* 37, 2 (1999), 183–233.
- Raymond Kan and Cesare Robotti. 2017. On moments of folded and truncated multivariate normal distributions. *Journal of Computational and Graphical Statistics* 26, 4 (2017), 930–934.
- Joost-Pieter Katoen, Annabelle K McIver, Larissa A Meinicke, and Carroll C Morgan. 2010. Linear-invariant generation for probabilistic programs. In *International Static Analysis Symposium*. Springer, 390–406.
- Peter V Kharchenko, Lev Silberstein, and David T Scadden. 2014. Bayesian approach to single-cell differential expression analysis. *Nature Methods* 11, 7 (2014), 740–742.
- Yehuda Koren, Steffen Rendle, and Robert Bell. 2021. Advances in collaborative filtering. *Recommender systems handbook* (2021), 91–142.
- Dexter Kozen. 1979. Semantics of probabilistic programs. In *20th Annual Symposium on Foundations of Computer Science (FOCS 1979)*. IEEE, 101–114.
- Dexter Kozen. 1983. A probabilistic PDL. In *Proceedings of the fifteenth annual ACM Symposium on Theory of computing*. 291–297.
- Alp Kucukelbir, Rajesh Ranganath, Andrew Gelman, and David Blei. 2015. Automatic variational inference in Stan. *Advances in Neural Information Processing Systems* 28 (2015).
- Solomon Kullback and Richard A Leibler. 1951. On information and sufficiency. *The Annals of Mathematical Statistics* 22, 1 (1951), 79–86.
- Jean Bernard Lasserre. 2009. *Moments, positive polynomials and their applications*. Vol. 1. World Scientific.
- Jacob Laurel and Sasa Misailovic. 2020. Continualization of probabilistic programs with correction. In *European Symposium on Programming*. Springer, Cham, 366–393.
- J Lo. 1972. Finite-dimensional sensor orbits and optimal nonlinear filtering. *IEEE Transactions on information theory* 18, 5 (1972), 583–588.
- Vikash Mansinghka, Daniel Selsam, and Yura Perov. 2014. Venture: a higher-order probabilistic programming platform with programmable inference. *arXiv preprint arXiv:1404.0099* (2014).
- Brian Milch, Bhaskara Marthi, and Stuart Russell. 2004. BLOG: Relational modeling with unknown objects. In *ICML 2004 workshop on statistical relational learning and its connections to other fields*. 67–73.
- Marcel Moosbrugger, Miroslav Stanković, Ezio Bartocci, and Laura Kovács. 2022. This is the moment for probabilistic loops. *Proceedings of the ACM on Programming Languages* 6, OOPSLA2 (2022), 1497–1525.
- 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*. Springer, 62–79.
- Radford M Neal. 1999. Erroneous results in “Marginal likelihood from the Gibbs output”. *minimeo, University of Toronto* (1999).
- Robert Nishihara, Thomas Minka, and Daniel Tarlow. 2013. Detecting parameter symmetries in probabilistic models. *arXiv preprint arXiv:1312.5386* (2013).
- Davide Nitti, Tinne De Laet, and Luc De Raedt. 2016. Probabilistic logic programming for hybrid relational domains. *Machine Learning* 103, 3 (2016), 407–449.
- Aditya Nori, Chung-Kil Hur, Sriram Rajamani, and Selva Samuel. 2014. R2: An efficient MCMC sampler for probabilistic programs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 28.
- Fritz Obermeyer, Eli Bingham, Martin Jankowiak, Neeraj Pradhan, Justin Chiu, Alexander Rush, and Noah Goodman. 2019. Tensor variable elimination for plated factor graphs. In *International Conference on Machine Learning*. PMLR, 4871–4880.
- Dilcia Pérez and Yamilet Quintana. 2006. A survey on the Weierstrass approximation theorem. *arXiv preprint math/0611038* (2006).
- Avi Pfeffer. 2001. IBAL: A probabilistic rational programming language. In *IJCAI*. Citeseer, 733–740.
- Emma Pierson and Christopher Yau. 2015. ZIFA: Dimensionality reduction for zero-inflated single-cell gene expression analysis. *Genome Biology* 16, 1 (2015), 1–10.
- Feras A Saad, Martin C Rinard, and Vikash K Mansinghka. 2021. SPPL: probabilistic programming with fast exact symbolic inference. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. 804–819.
- Konrad Schmüdgen. 2017. *The moment problem*. Vol. 9. Springer.
- SS Skiena. 2008. *The Algorithm Design Manual*. Springer Publishing Company.
- Luke Tierney and Joseph B Kadane. 1986. Accurate approximations for posterior moments and marginal densities. *J. Amer. Statist. Assoc.* 81, 393 (1986), 82–86.

- David Tolpin, Jan-Willem van de Meent, Hongseok Yang, and Frank Wood. 2016. Design and implementation of probabilistic programming language anglican. In *Proceedings of the 28th Symposium on the Implementation and Application of Functional programming Languages*. 1–12.
- Anastasios Tsiatis. 1975. A nonidentifiability aspect of the problem of competing risks. *Proceedings of the National Academy of Sciences* 72, 1 (1975), 20–22.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- Gian-Carlo Wick. 1950. The evaluation of the collision matrix. *Physical Review* 80, 2 (1950), 268.
- Wolfram Research, Inc. [n. d.]. *Mathematica*. <https://www.wolfram.com/mathematica>
- Yi Wu, Siddharth Srivastava, Nicholas Hay, Simon Du, and Stuart Russell. 2018. Discrete-continuous mixtures in probabilistic programming: Generalized semantics and inference algorithms. In *International Conference on Machine Learning*. PMLR, 5343–5352.
- Xiaoxue Zhao, Weinan Zhang, and Jun Wang. 2013. Interactive collaborative filtering. In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management*. 1411–1420.
- Ding-Xuan Zhou. 2020. Universality of deep convolutional neural networks. *Applied and Computational Harmonic Analysis* 48, 2 (2020), 787–794.
- Yuan Zhou, Hongseok Yang, Yee Whye Teh, and Tom Rainforth. 2020. Divide, conquer, and combine: a new inference strategy for probabilistic programs with stochastic support. In *International Conference on Machine Learning*. PMLR, 11534–11545.

Received 2023-07-11; accepted 2023-11-07