TPLP **20** (1): 147–175, 2020. © The Author(s), 2019. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

doi:10.1017/S1471068419000085 First published online 2 July 2019

*Implementing a Library for Probabilistic Programming Using Non-strict Non-determinism**

SANDRA DYLUS
University of Kiel, Kiel, Germany
(*e-mail:* sad@informatik.uni-kiel.de)

JAN CHRISTIANSEN
Flensburg University of Applied Sciences, Flensburg, Germany
(*e-mail:* jan.christiansen@hs-flensburg.de)

FINN TEEGEN
University of Kiel, Kiel, Germany
(*e-mail:* fte@informatik.uni-kiel.de)

submitted 13 April 2018; revised 23 May 2019; accepted 27 May 2019

Abstract

This paper presents *PFLP*, a library for probabilistic programming in the functional logic programming language Curry. It demonstrates how the concepts of a functional logic programming language support the implementation of a library for probabilistic programming. In fact, the paradigms of functional logic and probabilistic programming are closely connected. That is, language characteristics from one area exist in the other and vice versa. For example, the concepts of non-deterministic choice and call-time choice as known from functional logic programming are related to and coincide with stochastic memoization and probabilistic choice in probabilistic programming, respectively. We will further see that an implementation based on the concepts of functional logic programming can have benefits with respect to performance compared to a standard list-based implementation and can even compete with full-blown probabilistic programming languages, which we illustrate by several benchmarks.

KEYWORDS: probabilistic programming, functional logic programming, non-determinism, laziness, call-time choice

1 Introduction

The probabilistic programming paradigm allows the succinct definition of probabilistic processes and other applications based on probability distributions, for example, Bayesian networks as used in machine learning. A Bayesian network is a visual, graph-based representation for a set of random variables and their dependencies. One of the *hello world*-examples of Bayesian networks is the influence of rain and a sprinkler on

* This is an extended version of a paper presented at the *International Symposium on Practical Aspects of Declarative Languages (PADL 2018)*, invited as a rapid communication in TPLP. The authors acknowledge the assistance of the conference program chairs Nicola Leone and Kevin Hamlen. We are thankful for fruitful discussions with Michael Hanus as well as suggestions of Jan Bracker and Falco Nogatz. Finally, we are thankful for the comments of the anonymous reviewers to improve the readability of this paper.

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

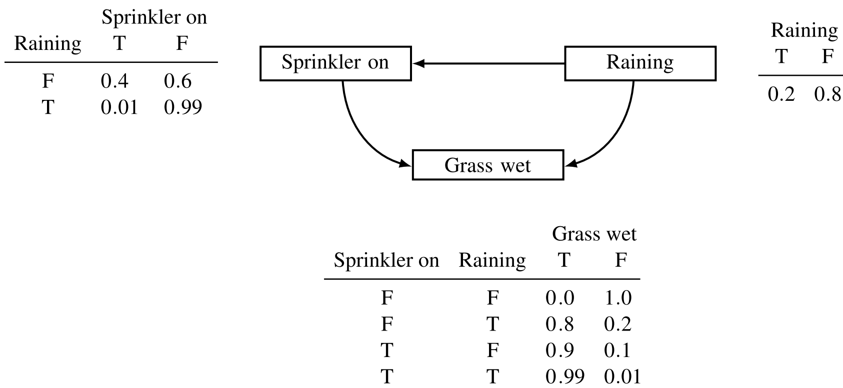

Raining	Sprinkler on T	Sprinkler on F
F	0.4	0.6
T	0.01	0.99

Raining T	Raining F
0.2	0.8

Sprinkler on	Raining	Grass wet T	Grass wet F
F	F	0.0	1.0
F	T	0.8	0.2
T	F	0.9	0.1
T	T	0.99	0.01

Fig. 1. A simple Bayesian network.

wet grass. Figure 1 shows an instance of this example, the concrete probabilities differ between publications. A node in the graph represents a random variable, a directed edge between two nodes represents a conditional dependency. Each node is annotated with a probability function represented as a table. The input values are on the left side of the table and the right side of the table describes the possible output and the corresponding probability. The input values of the function correspond to the incoming edges of that node. For example, the node for sprinkler depends on rain, thus, the sprinkler node has an incoming edge that originates from the rain node. The input parameter rain appears directly in the table that describes the probability function for sprinkler. For the example in Figure 1, the interpretation of the graph reads as follows: it rains with a probability of 20%; depending on the rain, the probability for an activated sprinkler is 40% and 1%, respectively; depending on both these factors, the grass can be observed as wet with a probability of 0%, 80%, 90%, or 99%.

- What is the probability that it is raining?
- What is the probability that the grass is wet, given that it is raining?
- What is the probability that the sprinkler is on, given that the grass is wet?

The general idea of probabilistic programming has been quite successful. There are a variety of probabilistic programming languages supporting all kinds of programming paradigms. For example, the programming languages Church (Goodman *et al*. 2008) and Anglican (Wood *et al*. 2014) are based on the functional programming language Scheme, ProbLog (Kimmig *et al*. 2011) is an extension of the logic programming language Prolog, Probabilistic C (Paige and Wood 2014) is based on the imperative language C, and WebPPL (Goodman and Stuhlmüller 2014), the successor of Church, is embedded in a functional subset of JavaScript. Besides full-blown languages there are also embedded domain-specific languages that implement probabilistic programming as a library. For example, FACTORIE (McCallum *et al*. 2009) is a library for the hybrid programming language Scala, and Erwig and Kollmansberger (2006) present a library for the functional programming language Haskell. We recommend the survey by Gordon *et al.* (2014) about the current state of probabilistic programming for further information.

This paper presents *PFLP*, a library providing a domain-specific language for probabilistic programming in the functional logic programming language Curry (Antoy and

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

Hanus 2010). PFLP makes heavy use of functional logic programming concepts and shows that this paradigm is well-suited for implementing a library for probabilistic programming. In fact, there is a close connection between probabilistic programming and functional logic programming. For example, non-deterministic choice and probabilistic choice are similar concepts. Furthermore, the concept of call-time choice as known from functional logic programming coincides with (stochastic) memoization (De Raedt and Kimmig 2013) in the area of probabilistic programming. We are not the first to observe this close connection between functional logic programming and probabilistic programming. For example, Fischer *et al.* (2009) present a library for modeling functional logic programs in the functional language Haskell. As they state, by extending their approach to weighted non-determinism, we can model a probabilistic programming language.

Besides a lightweight implementation of a library for probabilistic programming in a functional logic programming language, this paper makes the following contributions:

- We investigate the interplay of probabilistic programming with the features of a functional logic programming language. For example, we show how call-time choice and non-determinism interplay with probabilistic choice.
- We discuss how we utilize functional logic features to improve the implementation of probabilistic combinators.
- We present an implementation of probability distributions using non-determinism in combination with non-strict probabilistic combinators that is more efficient than an implementation using lists.
- We illustrate that the combination of non-determinism and non-strictness with respect to distributions has to be handled with care. More precisely, it is important to enforce a certain degree of strictness in order to guarantee correct results.
- In contrast to the conference version of the paper (Dylus *et al.* 2018), we discuss the usage of partial functions in combination with library functions in more detail, reason about laws for two operations of the library, and present performance comparisons between our library, ProbLog and WebPPL.
- Finally, this paper aims at fostering the exchange between the community of probabilistic programming and of functional logic programming. That is, while the connection exists for a long time, there has not been much exchange between the communities. We would like to take this paper as a starting point to bring these paradigms closer together. Thus, this paper introduces the concepts of both, the functional logic and probabilistic programming, paradigms.

Please note that the current state of our library cannot compete against full-blown probabilistic languages or mature libraries for probabilistic programming in terms of features, for example, the library does not provide any sampling mechanisms. Nevertheless, the library is a good showcase for languages with built-in non-determinism, because the functional logic approach can be superior to the functional approach using lists. Furthermore, we want to emphasize that this paper uses non-determinism as an implementation technique to develop a library for probabilistic programming. That is, we are not mainly concerned with the interaction of non-determinism and probabilities as, for example, discussed in the work of Varacca and Winskel (2006) and multiple others. The library

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

we develop in this paper does not combine both effects, but provides combinators for probabilistic programming by leveraging Curry's built-in non-strict non-determinism.

2 Library basics

In this section, we discuss the core of the PFLP library.[1] The implementation is based on a Haskell library for probabilistic programming presented by Erwig and Kollmansberger (2006). We will not present the whole PFLP library, but only core functions. The paper at hand is a literate Curry file. We use the Curry compiler KiCS2,[2] by Braßel *et al.* (2011), for all code examples.

2.1 Modeling distributions

One key ingredient of probabilistic programming is the definition of distributions. A distribution consists of pairs of elementary events and their probability. We model probabilities as *Float* and distributions as a combination of an elementary event and the corresponding probability.[3]

type *Probability* = *Float*
data *Dist a* = *Dist a Probability*

In a functional language like Haskell, the canonical way to define distributions uses lists. Here, we use Curry's built-in non-determinism as an alternative for lists to model distributions with more than one event-probability pair. As an example, we define a fair coin, where *True* represents heads and *False* represents tails, as follows[4]:

$$\begin{aligned}&\mathit{coin} :: \mathit{Dist\ Bool}\\&\mathit{coin} = (\mathit{Dist\ True}\ \tfrac{1}{2})\ ?\ (\mathit{Dist\ False}\ \tfrac{1}{2})\end{aligned}$$

In Curry, the (?)-operator non-deterministically chooses between two given arguments. Non-determinism is not reflected in the type system, that is, a non-deterministic choice has type $a \to a \to a$. Such non-deterministic computations introduced by (?) describe two individual computation branches; one for the left argument and one for the right argument of (?).

We could also define *coin* in Prolog-style by giving two rules for coin.

$$\begin{aligned}&\mathit{coin} :: \mathit{Dist\ Bool}\\&\mathit{coin} = \mathit{Dist\ True}\ \tfrac{1}{2}\\&\mathit{coin} = \mathit{Dist\ False}\ \tfrac{1}{2}\end{aligned}$$

Both implementations can be used interchangeably since the (?)-operator is defined in the Prolog-style using two rules as well.

[1] We provide the code for the library at `https://github.com/finnteegen/pflp`.
[2] We use version `0.6.0` of KiCS2 and the source is found at `https://www-ps.informatik.uni-kiel.de/kics2/`.
[3] The polymorph data type *Dist* is parameterized over a type variable *a*. It has a single constructor also named *Dist* that is of type $a \to \mathit{Probability} \to \mathit{Dist}\ a$. The constructor *Dist* in Curry corresponds to a binary functor in Prolog.
[4] Here, and in the following, we write probabilities as fractions for readability.

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

$(?) :: a \to a \to a$
$x \mathbin{?} y = x$
$x \mathbin{?} y = y$

Printing an expression in the REPL[5] evaluates the non-deterministic computations, thus, yields one result for each branch as shown in the following examples:

$> 1 \mathbin{?} 2$	$> \mathit{coin}$
1	*Dist True* 0.5
2	*Dist False* 0.5

The REPL computes the values using a breadth-first-search strategy to visualize the results. Due to the search strategy, we observe different outputs when changing the order of arguments to (?). However, because Curry's semantics is set-based (Christiansen *et al.* 2011), the order of the results does not matter.

It is cumbersome to define distributions explicitly as in the case of *coin*. Hence, we define helper functions for constructing distributions.[6] Given a list of events and probabilities, *enum* creates a distribution by folding these pairs non-deterministically with a helper function *member*.[7]

$\mathit{member} :: [a] \to a$
$\mathit{member}\ \mathit{xs} = \mathit{foldr}\ (?)\ \mathit{failed}\ \mathit{xs}$
$\mathit{enum} :: [a] \to [\mathit{Probability}] \to \mathit{Dist}\ a$
$\mathit{enum}\ \mathit{xs}\ \mathit{ps} = \mathit{member}\ (\mathit{zipWith}\ \mathit{Dist}\ \mathit{xs}\ \mathit{ps})$

In Curry, the constant *failed* is a silent failure that behaves as neutral element with respect to (?). That is, the expression *True* ? *failed* is equivalent to *True*. Hence, the function *member* takes a list and yields a non-deterministic choice of all elements of the list.

As a shortcut, we define a function that yields a *uniform* distribution given a list of events as well as a function *certainly*, which yields a distribution with a single event of probability one.

$\mathit{uniform} :: [a] \to \mathit{Dist}\ a$
$\mathit{uniform}\ \mathit{xs} = \mathbf{let}\ \mathit{len} = \mathit{length}\ \mathit{xs}\ \mathbf{in}\ \mathit{enum}\ \mathit{xs}\ (\mathit{repeat}\ \frac{1}{\mathit{len}})$
$\mathit{certainly} :: a \to \mathit{Dist}\ a$
$\mathit{certainly}\ x = \mathit{Dist}\ x\ 1.0$

The function *repeat* yields a list that contains the given value infinitely often. Because of Curry's laziness, it is sufficient if one of the arguments of *enum* is a finite list because *zipWith* stops when one of its arguments is empty. We can then refactor the definition of *coin* using *uniform* as follows:

$\mathit{coin} :: \mathit{Dist}\ \mathit{Bool}$
$\mathit{coin} = \mathit{uniform}\ [\mathit{True}, \mathit{False}]$

[5] We visualize the interactions with the REPL using $>$ as prompt.
[6] The definitions of predefined Curry functions like *foldr* are listed in Appendix.
[7] We shorten the implementation of *enum* for presentation purposes; actually, *enum* only allows valid distributions, for example, that the given probabilities sum up to 1.0.

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

In general, the library hides the constructor *Dist*, that is, the user has to define distributions by using the combinators provided by the library.

The library provides additional functions to combine and manipulate distributions. In order to work with dependent distributions, the operator ($\ggg=$) applies a function, which yields a distribution, to each event of a given distribution and multiplies the corresponding probabilities.[8]

$$
\begin{array}{l}
(\ggg=) :: Dist\ a \rightarrow (a \rightarrow Dist\ b) \rightarrow Dist\ b \\
d \ggg= f = \textbf{let}\ Dist\ x\ p = d \\
\qquad\qquad\quad\ \ Dist\ y\ q = f\ x \\
\qquad\qquad \textbf{in}\ Dist\ y\ (p *.\ q)
\end{array}
$$

Intuitively, we have to apply the function f to each event of the distribution d and combine the resulting distributions into a single distribution. In a Haskell implementation, we would use a list comprehension to define this function. In the Curry implementation, we model distributions as non-deterministic computations, thus, the above rule describes the behavior of the function for an arbitrary pair of the first distribution and an arbitrary pair of the second distribution, that is, the result of f.

Using the operator ($\ggg=$) we can, for example, define a distribution that models flipping two coins. The events of this distribution are pairs whose first component is the result of the first coin flip and whose second component is the result of the second coin flip.

$$
\begin{array}{l}
independentCoins :: Dist\ (Bool, Bool) \\
independentCoins = coin \ggg= (\lambda b1 \rightarrow coin \ggg= (\lambda b2 \rightarrow certainly\ (b1, b2)))
\end{array}
$$

In contrast to the example *independentCoins*, we can also use the operator ($\ggg=$) to combine two distributions where we choose the second distribution on the basis of the result of the first. For example, we can define a distribution that models flipping two coins, but in this case, we only flip a second coin if the first coin yields heads.

$$
\begin{array}{l}
dependentCoins :: Dist\ Bool \\
dependentCoins = coin \ggg= (\lambda b \rightarrow \textbf{if}\ b\ \textbf{then}\ coin\ \textbf{else}\ certainly\ False)
\end{array}
$$

The implementation of ($\ggg=$) via **let**-bindings seems a bit tedious, however, it is important that we define ($\ggg=$) as it is. The canonical implementation performs pattern matching on the first argument but uses a **let**-binding for the result of f. That is, it is strict in the first argument but non-strict in the application of f, the second argument. For now it is sufficient to note – and keep in mind – that there is a difference between pattern matching and using **let**-bindings. In order to understand this difference, let us consider the following implementation of *fromJustToList* and an alternative implementation *fromJustToListLet*.[9]

$$
\begin{array}{l}
fromJustToList :: Maybe\ a \rightarrow [a] \\
fromJustToList\ (Just\ x) = x : [\,]
\end{array}
$$

[8] Due to the lack of overloading in Curry, operations on *Float* have a (floating) point suffix, for example, ($*.$), whereas operations on *Int* use the common operation names.

[9] $(:) :: a \rightarrow [a] \rightarrow [a]$ denotes the constructor for a non-empty list – similar to the functor *./2* in Prolog.

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

```
fromJustToListLet :: Maybe a → [a]
fromJustToListLet mx = let Just x = mx in x : []
```

The second implementation, *fromJustToListLet*, is less strict, because it yields a list constructor, (:), without evaluating its argument first. That is, we can observe the difference when passing *failed* and checking if the resulting list is empty or not.

```
> null (fromJustToList failed)
failed
> null (fromJustToListLet failed)
False
```

Due to the pattern matching in the definition of *fromJustToList* the argument *failed* needs to be evaluated, and thus, the function *null* propagates *failed* as result. In contrast, the definition of *fromJustToListLet* postpones the evaluation of its argument to the right-hand side, that is, the argument needs to be evaluated only if the computation demands the value x explicitly. The function *null* does not demand the evaluation of x, because it only checks the surrounding list constructor.

```
null :: [a] → Bool
null []       = True
null (x : xs) = False
```

The same strictness property as for *fromJustToList* holds for a definition via explicit pattern matching using **case** ... **of**. In particular, pattern matching of the left-hand side of a rule desugars to case expressions on the right-hand side.

```
fromJustToListCase :: Maybe a → [a]
fromJustToListCase mx = case mx of
                          Just x → [x]
```

```
> null (fromJustToListCase failed)
failed
```

We discuss the implementation of (≫=) in more detail later. For now, it is sufficient to keep in mind that (≫=) yields a *Dist*-constructor without evaluating any of its arguments. In contrast, a definition using pattern matching or a case expression needs to evaluate its argument first, thus, is more strict.

For independent distributions, we provide the function *joinWith* that combines two distributions with respect to a given function. We implement *joinWith* by means of (≫=).

```
joinWith :: (a → b → c) → Dist a → Dist b → Dist c
joinWith f d1 d2 = d1 ≫= (λx → d2 ≫= (λy → certainly (f x y)))
```

In a monadic setting, this function is sometimes called *liftM2*. Here, we use the same nomenclature as Erwig and Kollmansberger (2006).

As an example, we define a function that flips a coin n times.

```
flipCoin :: Int → Dist [Bool]
flipCoin n | n ≡ 0     = certainly []
           | otherwise = joinWith (:) coin (flipCoin (n − 1))
```

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

When we run the example of flipping two coins in the REPL of KiCS2, we get four events.

$> \mathit{flipCoin}\ 2$
$\mathit{Dist}\ [\mathit{True}, \mathit{True}]\ 0.25$
$\mathit{Dist}\ [\mathit{True}, \mathit{False}]\ 0.25$
$\mathit{Dist}\ [\mathit{False}, \mathit{True}]\ 0.25$
$\mathit{Dist}\ [\mathit{False}, \mathit{False}]\ 0.25$

In the example above, *coin* is non-deterministic, namely, $\mathit{coin} = (\mathit{Dist}\ \mathit{True}\ \frac{1}{2})\ ?\ (\mathit{Dist}\ \mathit{False}\ \frac{1}{2})$. Applying *joinWith* to *coin* and *coin* combines all possible results of two coin tosses.

2.2 Querying distributions

With a handful of building blocks to define distributions available, we now want to query the distribution, that is, calculate the probability of certain events. We provide an operator $(??) :: (a \to \mathit{Bool}) \to \mathit{Dist}\ a \to \mathit{Probability}$ – which we will define shortly – to extract the probability of an event. The event is specified as a predicate, passed as first argument. The operator filters events that satisfy the given predicate and computes the sum of the probabilities of the remaining elementary events. We implement this kind of filter function on distributions in Curry.

```
filterDist :: (a → Bool) → Dist a → Dist a
filterDist pred d = let Dist x p = d
                    in if (pred x) then (Dist x p) else failed
```

The implementation of *filterDist* is a partial identity on the event-probability pairs. Every event that satisfies the predicate is part of the resulting distribution. The function fails for event-probability pairs that do not satisfy the predicate.

Querying a distribution, that is, summing up all probabilities that satisfy a predicate, is a more advanced task in the functional logic approach. Remember that we represent a distribution by chaining all event-probability pairs with (?), thus, constructing non-deterministic computations. These non-deterministic computations introduce individual branches of computations that cannot interact with each other. In order to compute the total probability of a distribution, we have to merge these distinct branches. Such a merge is possible by the encapsulation of non-deterministic computations. Similar to the *findall* construct of the logic language Prolog, in Curry we encapsulate a non-deterministic computation by using a primitive called *allValues*.[10] The function $\mathit{allValues} :: a \to \{a\}$ operates on a polymorphic – and potentially non-deterministic – value and yields a multiset of all non-deterministic values. In order to work with encapsulated values, Curry provides the function $\mathit{foldValues} :: (a \to a \to a) \to a \to \{a\} \to a$ to fold the resulting multiset.

We do not discuss the implementation details behind *allValues* here. It is sufficient to know that, as a library developer, we can employ this function to encapsulate non-

[10] We use an abstract view of the result of an encapsulation to emphasize that the order of encapsulated results does not matter. In practice, we can, for example, use the function $\mathit{allValues} :: a \to [a]$ defined in the library *Findall*.

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

deterministic values and use these values in further computations. However, due to non-transparent behavior in combination with sharing as discussed by Braßel *et al.* (2004), a user of the library should not use *allValues* at all. In a nutshell, inner-most and outer-most evaluation strategies may cause different results when combining sharing and encapsulation.

With this encapsulation mechanism at hand, we can define the query operator (??) as follows:

$$
\begin{aligned}
&prob :: Dist\ a \to Probability\\
&prob\ (Dist\ x\ p) = p\\
&(??) :: (a \to Bool) \to Dist\ a \to Probability\\
&(??)\ pred\ d = foldValues\ (+.)\ 0.0\ (allValues\ (prob\ (filterDist\ pred\ d)))
\end{aligned}
$$

First, we filter the elementary events by some predicate and project to the probabilities only. Afterwards, we encapsulate the remaining probabilities and sum them up. As an example for the use of (??), we may flip four coins and calculate the probability of at least two heads – that is, the list contains at least two *True* values.

$$
\begin{aligned}
&> (\lambda coins \to length\ (filter\ id\ coins) \geqslant 2)\ ??\ (flipCoin\ 4)\\
&0.6875
\end{aligned}
$$

In order to check the result, we calculate the probability by hand. Since there are more events that satisfy the predicate than events that do not, we sum up the probabilities of the events that do not satisfy the predicate and calculate the complementary probability. There is one event where all coins show tails and four events where one of the coins shows heads and all other show tails.

$$
\begin{aligned}
&1-(P(Tails)\cdot P(Tails)\cdot P(Tails)\cdot P(Tails) + 4\cdot P(Heads)\cdot P(Tails)\cdot P(Tails)\cdot P(Tails))\\
&\quad = 1 - (0.5 \cdot 0.5 \cdot 0.5 \cdot 0.5 + 4 \cdot 0.5 \cdot 0.5 \cdot 0.5 \cdot 0.5)\\
&\quad = 1 - (0.0625 + 0.25) = 1 - 0.3125 = 0.6875.
\end{aligned}
$$

3 Library pragmatics

Up to now, we have discussed a simple library for probabilistic programming that uses non-determinism to represent distributions. In this section, we will see that we can highly benefit from Curry-like non-determinism with respect to performance when we compare PFLP's implementation with a list-based implementation. More precisely, when we query a distribution with a predicate that does not evaluate its argument completely, we can possibly prune large parts of the search space. Before we discuss the details of the combination of non-strictness and non-determinism, we discuss aspects of sharing non-deterministic choices. Finally, we discuss details about the implementation of ($\ggg=$) and why PFLP does not allow non-deterministic events within distributions.

3.1 Call-time choice vs. Run-time choice

By default Curry uses call-time choice, that is, variables denote single deterministic choices. When we bind a variable to a non-deterministic computation, one value is chosen and all occurrences of the variable denote the same deterministic choice.

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

Often call-time choice is what you are looking for. For example, this slightly modified definition of *filterDist* makes use of call-time choice.

```
filterDist :: (a → Bool) → Dist a → Dist a
filterDist pred d = let Dist x p = d
                    in if (pred x) then d else failed
```

Due to pattern matching via **let**-binding, the variable d on the right-hand side corresponds to a single deterministic choice for the input distribution, namely, the one that satisfies the predicate and not the non-deterministic computation that was initially passed as second argument to *filterDist*.

Almost as often run-time choice is what you are looking for and call-time choice gets in your way; probabilistic programming is no exception. For example, let us reconsider flipping a coin n times. We parametrize the function *flipCoin* over the given distribution and define the following generalized function:

```
replicateDist :: Int → Dist a → Dist [a]
replicateDist n d | n ≡ 0      = certainly [ ]
                  | otherwise = joinWith (:) d (replicateDist (n − 1) d)
```

Here, we again use guard syntax in order to distinguish several cases depending on the Boolean expression. When we use this function to flip a coin twice, the result is not what we intended.

```
> replicateDist 2 coin
Dist [True, True] 0.25
Dist [False, False] 0.25
```

Because *replicateDist* shares the variable d, we only perform a choice once and replicate deterministic choices. In contrast, top-level nullary functions like *coin* are evaluated every time, thus, exhibit run-time choice, which is the reason why the previously shown *flipCoin* behaves properly.

In order to implement *replicateDist* correctly, we have to enforce run-time choice. We introduce the following type synonym and function to model and work with values with run-time choice behavior.

```
type RT a = () → a
pick :: RT a → a
pick rt = rt ()
```

We can now use the type *RT* to hide the non-determinism on the right-hand side of a function arrow. This way, *pick* explicitly triggers the evaluation of *rt*, performing a new choice for every element of the result list.

```
replicateDist :: Int → RT (Dist a) → Dist [a]
replicateDist n rt | n ≡ 0      = certainly [ ]
                   | otherwise = joinWith (:) (pick rt) (replicateDist (n − 1) rt)
```

In order to use *replicateDist* with *coin*, we have to construct a value of type *RT* (*Dist Bool*). However, we cannot provide a function to construct a value of type *RT*

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

that behaves as intended. Such a function would share a deterministic choice and non-deterministically yield two functions, instead of one function that yields a non-deterministic computation. The only way to construct a value of type RT is to explicitly use a lambda abstraction.

$> replicateDist\ 2\ (\lambda() \rightarrow coin)$
$Dist\ [True, True]\ 0.25$
$Dist\ [True, False]\ 0.25$
$Dist\ [False, True]\ 0.25$
$Dist\ [False, False]\ 0.25$

Instead of relying on call-time choice as default behavior, we could model *Dist* as a function and make run-time choice the default in PFLP. In this case, to get call-time choice we would have to use a special construct provided by the library – as it is the case in many probabilistic programming libraries, for example, *mem* in WebPPL (Goodman and Stuhlmüller 2014).

On the other hand, ProbLog uses a similar concept to call-time choice, namely, stochastic memoization, which reuses already computed results. That is, predicates that are associated with probabilities become part of the memoized result. If a fair coin flip, for example, already resulted in *True*, then the probability of all further coin flips that also result in *True* have probability 1. Due to stochastic memoization the coin is not flipped a second time, but is identified as the same coin as before. Thus, stochastic memoization as used in ProbLog is similar to the extension of tabling in Prolog systems, but adapted to the setting of probabilistic programming that extends predicates with probabilities. Similar to our usage of RT to mimic run-time choice in Curry, we can use a so-called trial identifier, which is basically an additional argument, to circumvent memoization for a predicate like coin in ProbLog. The difference to RT is that the trial identifier needs to be different for each call to the predicate in order to force re-evaluation.

In the end, we have decided to go with the current modeling based on call-time choice, because the alternative would work against the spirit of the Curry programming language. There is a long history of discussions about the pros and cons of call-time choice and run-time choice. It is common knowledge in probabilistic programming (De Raedt and Kimmig 2013) that memoization – that is, call-time choice – has to be avoided in order to model stochastic automata or probabilistic grammars. Similarly, Antoy (2005) observes that you need run-time choice to elegantly model regular expressions in the context of functional logic programming languages. Then again, probabilistic languages need a concept like memoization in order to use a single value drawn from a distribution multiple times.

3.2 Combination of non-strictness and non-determinism

This subsection illustrates the benefits from the combination of non-strictness and non-determinism with respect to performance. More precisely, in a setting that uses Curry-like non-determinism, non-strictness can prevent non-determinism from being "spawned." Let us consider calculating the probability for throwing only sixes when throwing n dice. First, we define a uniform die as follows:

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

Table 1. *Overview of running times for the query allSix n*

No. of dice	5	6	7	8	9	10	100	200	300
Curry ND	<1	<1	<1	<1	<1	<1	48	231	547
Curry List	2	13	72	419	2554	15,394	–	–	–
Curry ND!	52	409	2568	16,382	–	–	–	–	–
Haskell List	1	5	30	210	1415	6538	–	–	–

$\textbf{data}\ Side = One \mid Two \mid Three \mid Four \mid Five \mid Six$
$die :: Dist\ Side$
$die = uniform\ [One, Two, Three, Four, Five, Six]$

We define the following query by means of the combinators introduced so far. The function *all* simply checks that all elements of a list satisfy a given predicate; it is defined by means of the Boolean conjunction $(\wedge)$.

$allSix :: Int \to Probability$
$allSix\ n = (all\ (\equiv Six))\ ??\ (replicateDist\ n\ (\lambda() \to die))$

Table 1 compares running times[11] of this query for different numbers of dice. The row labeled “Curry ND” lists the running times for an implementation that uses the operator $(\ggg=)$. The row “Curry List” shows the numbers for a list-based implementation in Curry, which is a literal translation of the library by Erwig and Kollmansberger. The row labeled “Curry ND!” uses an operator $(\ggg=!)$ instead, which we will discuss shortly. Finally, we compare our implementation to the original list-based implementation, which the row labeled “Haskell List” refers to. The table states the running times in milliseconds of a compiled executable for each benchmark as a mean of three runs. Cells marked with “–” take more than 1 min.

Obviously, the example above is a little contrived. While the query is exponential in both list versions, it is linear in the non-deterministic setting.[12] To illustrate the behavior of the example above, we consider the following application for an arbitrary distribution *dist* of type $Dist\ [Side]$.

$$filterDist\ (all\ (\equiv Six))\ (joinWith\ (:)\ (Dist\ One\ \tfrac{1}{6})\ dist)$$

This application yields an empty distribution without evaluating the distribution *dist*. The interesting point here is that *joinWith* yields a *Dist*-constructor without inspecting its arguments. When we demand the event of the resulting *Dist*, *joinWith* has to evaluate only its first argument to see that the predicate $all\ (\equiv Six)$ yields *False*. The evaluation of the expression fails without inspecting the second argument of *joinWith*. Figure 2 illustrates the evaluation in more detail.

[11] All benchmarks were executed on a Linux machine with an Intel Core i7-6500U (2.50 GHz) and 8 GiB RAM running Fedora 25. We used the Glasgow Haskell Compiler (version `8.0.2`, option `-O2`) and set the search strategy in KiCS2 to depth-first.

[12] Non-determinism causes significant overhead for KiCS2, thus, “Curry ND” does not show linear development, but we measured a linear running time using PAKCS (Hanus 2017).

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

$$
\begin{array}{l}
\mathit{filterDist}\ (\mathit{all}\ (\equiv \mathit{Six}))\ (\mathit{joinWith}\ (:)\ (\mathit{Dist}\ \mathit{One}\ \tfrac{1}{6})\ \mathit{dist})\\
\quad \equiv \quad \{ \text{Definition of } \mathit{joinWith} \}\\
\mathit{filterDist}\ (\mathit{all}\ (\equiv \mathit{Six}))\\
\qquad\qquad (\mathit{Dist}\ \mathit{One}\ \tfrac{1}{6} \ggg\!= (\lambda x \to \mathit{dist} \ggg\!= (\lambda xs \to \mathit{certainly}\ (x : xs))))\\
\quad \equiv \quad \{ \text{Definition of } (\ggg\!=) \text{ (twice)} \}\\
\mathit{filterDist}\ (\mathit{all}\ (\equiv \mathit{Six}))\\
\qquad\qquad (\textbf{let}\ \mathit{Dist}\ x\ p = \mathit{Dist}\ \mathit{One}\ \tfrac{1}{6}\\
\qquad\qquad\qquad \mathit{Dist}\ xs\ q = \mathit{dist}\\
\qquad\qquad\qquad \mathit{Dist}\ ys\ r = \mathit{certainly}\ (x : xs)\\
\qquad\qquad\ \ \textbf{in}\ \mathit{Dist}\ ys\ (p *. (q *. r))\\
\quad \equiv \quad \{ \text{Definition of } \mathit{filterDist} \}\\
\textbf{let}\ \mathit{Dist}\ x\ p = \mathit{Dist}\ \mathit{One}\ \tfrac{1}{6}\\
\quad \mathit{Dist}\ xs\ q = \mathit{dist}\\
\quad \mathit{Dist}\ ys\ r = \mathit{certainly}\ (x : xs)\\
\textbf{in if}\ \mathit{all}\ (\equiv \mathit{Six})\ ys\ \textbf{then}\ \mathit{Dist}\ ys\ (p *. (q *. r))\ \textbf{else}\ \mathit{failed}\\
\quad \equiv \quad \{ \text{Definition of } \mathit{certainly} \}\\
\textbf{let}\ \mathit{Dist}\ x\ p = \mathit{Dist}\ \mathit{One}\ \tfrac{1}{6}\\
\quad \mathit{Dist}\ xs\ q = \mathit{dist}\\
\textbf{in if}\ \mathit{all}\ (\equiv \mathit{Six})\ (x : xs)\ \textbf{then}\ \mathit{Dist}\ (x : xs)\ (p *. (q *. 1.0))\ \textbf{else}\ \mathit{failed}\\
\quad \equiv \quad \{ \text{Definition of } \mathit{all} \}\\
\textbf{let}\ \mathit{Dist}\ x\ p = \mathit{Dist}\ \mathit{One}\ \tfrac{1}{6}\\
\quad \mathit{Dist}\ xs\ q = \mathit{dist}\\
\textbf{in if}\ x \equiv \mathit{Six} \wedge \mathit{all}\ (\equiv \mathit{Six})\ xs\ \textbf{then}\ \mathit{Dist}\ (x : xs)\ (p *. (q *. 1.0))\ \textbf{else}\ \mathit{failed}\\
\quad \equiv \quad \{ \text{Definition of } (\equiv) \text{ and } (\wedge) \}\\
\textbf{let}\ \mathit{Dist}\ x\ p = \mathit{Dist}\ \mathit{One}\ \tfrac{1}{6}\\
\quad \mathit{Dist}\ xs\ q = d\\
\textbf{in if}\ \mathit{False}\ \textbf{then}\ \mathit{Dist}\ (x : xs)\ (p *. (q *. 1.0))\ \textbf{else}\ \mathit{failed}\\
\quad \equiv \quad \{ \text{Definition of } \textbf{if} - \textbf{then} - \textbf{else} \}\\
\mathit{failed}
\end{array}
$$

Fig. 2. Simplified evaluation illustrating non-strict non-determinism.

In case of the example *allSix*, all non-deterministic branches that contain a value different from *Six* fail fast due to the non-strictness. Thus, the number of evaluation steps is linear in the number of rolled dice.

We can only benefit from the combination of non-strictness and non-determinism if we define $(\ggg\!=)$ with care. Let us take a look at a strict variant of $(\ggg\!=)$ and discuss its consequences.

$$
\begin{array}{l}
(\ggg\!=!) :: \mathit{Dist}\ a \to (a \to \mathit{Dist}\ b) \to \mathit{Dist}\ b\\
(\mathit{Dist}\ x\ p) \ggg\!=!\ f = \textbf{case}\ f\ x\ \textbf{of}\\
\qquad\qquad\qquad\qquad\quad \mathit{Dist}\ y\ q \to \mathit{Dist}\ y\ (p *. q)
\end{array}
$$

This implementation is strict in its first argument as well as in the result of the function application. When we use $(\ggg\!=!)$ to implement the *allSix* example, we lose the benefit of Curry-like non-determinism. The row in Table 1 labeled “Curry ND!” shows the running times when using $(\ggg\!=!)$ instead of $(\ggg\!=)$. As $(\ggg\!=!)$ is strict, the function *joinWith* has to evaluate both its arguments to yield a result. Figure 3 shows how the formerly unneeded distribution *dist* now has to be evaluated in order to yield a value. More precisely, using $(\ggg\!=!)$ causes a complete evaluation of *dist*.

Please note that an implementation that is similar to $(\ggg\!=)$ is *not* possible in a list-based implementation. The following definition of *concatMap* is usually used to define the bind operator for lists.

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

$$
\begin{array}{l}
\mathit{filterDist}\;(\mathit{all}\;(\equiv \mathit{Six}))\;(\mathit{joinWith}\;(:)\;(\mathit{Dist}\;\mathit{One}\;\tfrac{1}{6})\;\mathit{dist})\\
\equiv \quad \{\text{ Definition of } \mathit{joinWith} \}\\
\mathit{filterDist}\;(\mathit{all}\;(\equiv \mathit{Six}))\\
\qquad\qquad (\mathit{Dist}\;\mathit{One}\;\tfrac{1}{6} \ggg=!\; (\lambda x \to \mathit{dist} \ggg=!\; (\lambda xs \to \mathit{certainly}\;(x : xs))))\\
\equiv \quad \{\text{ Definition of } (\ggg=!) \}\\
\mathit{filterDist}\;(\mathit{all}\;(\equiv \mathit{Six}))\\
\qquad\qquad (\mathbf{case}\;(\lambda x \to \mathit{dist} \ggg=!\; (\lambda xs \to \mathit{certainly}\;(x : xs)))\;\mathit{One}\;\mathbf{of}\\
\qquad\qquad\qquad \mathit{Dist}\;y\;q \to \mathit{Dist}\;y\;(\tfrac{1}{6} *.\, q))\\
\equiv \quad \{\text{ Evaluation of the scrutinee } \}\\
\mathit{filterDist}\;(\mathit{all}\;(\equiv \mathit{Six}))\\
\qquad\qquad (\mathbf{case}\;\mathit{dist} \ggg=!\; (\lambda xs \to \mathit{certainly}\;(\mathit{One} : xs))\;\mathbf{of}\\
\qquad\qquad\qquad \mathit{Dist}\;y\;q \to \mathit{Dist}\;y\;(\tfrac{1}{6} *.\, q))\\
\equiv \quad \{\text{ Evaluation of } \mathit{dist} \text{ as demanded by the definition of } (\ggg=!) \}\\
\ldots
\end{array}
$$

Fig. 3. Simplified evaluation illustrating strict non-determinism.

$$
\begin{array}{l}
\mathit{concatMap} :: (a \to [\,b\,]) \to [\,a\,] \to [\,b\,]\\
\mathit{concatMap}\;f\;[\,] \qquad\; = [\,]\\
\mathit{concatMap}\;f\;(x : xs) = f\;x \mathbin{+\!\!+} \mathit{concatMap}\;f\;xs
\end{array}
$$

The strict behavior follows from the definition via pattern matching on the list argument. In contrast to ($\ggg=!$) there is, however, no other implementation that is less strict. The pattern matching is inevitable due to the two possible constructors, [ ] and (:), for lists. As a consequence, a list-based implementation has to traverse the entire distribution before we can evaluate the predicate $\mathit{all}\;(\equiv \mathit{Six})$. The consequence is that the running times of “Haskell List” in Table 1 cannot compete with “Curry ND” when the number of dice increases.

Intuitively, we expect similar running times for “Curry ND!” and “Curry List” as the bind operator for lists has to evaluate its second argument as well – similar to ($\ggg=!$). However, the observed running times do not have the expected resemblance. “Curry ND!” heavily relies on non-deterministic computations, which causes significant overhead for KiCS2. We do not investigate these differences here but propose it as a direction for future research.

Obviously, turning an exponential problem into a linear one is like getting only sixes when throwing dice. In most cases we are not that lucky. For example, consider the following query for throwing n dice that are either five or six.

$$
\begin{array}{l}
\mathit{allFiveOrSix} :: \mathit{Int} \to \mathit{Probability}\\
\mathit{allFiveOrSix}\;n = (\mathit{all}\;(\lambda s \to s \equiv \mathit{Five} \lor s \equiv \mathit{Six}))\;??\;(\mathit{replicateDist}\;n\;(\lambda() \to \mathit{die}))
\end{array}
$$

Table 2 lists the running times of this query for different numbers of dice with respect to the four different implementations. As we can see from the running times, this query is exponential in all implementations. Nevertheless, the running time of the non-strict, non-deterministic implementation is much better because we only have to consider two sides – six and five – while we have to consider all sides in the list implementations and the non-deterministic, strict implementation. That is, while the base of the complexity is two in the case of the non-deterministic, non-strict implementation, it is six in all the other cases. As we have observed in the other examples before, we get an overhead in the case of the strict non-determinism compared to the list implementation due to the heavy usage of non-deterministic computations.

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

Table 2. *Overview of running times of the query allFiveOrSix n*

No. of dice	5	6	7	8	9	10
Curry ND	4	7	15	34	76	163
Curry List	2	13	84	489	2869	16,989
Curry ND!	49	382	2483	15,562	–	–
Haskell List	2	5	31	219	1423	6670

3.3 Definition of the bind operator

In this subsection, we discuss our design choices concerning the implementation of the bind operator. We illustrate that we have to be careful about non-strictness, because we do not want to lose non-deterministic results.

First, we revisit the definition of ($\ggg=$) introduced in Section 2.

$$
\begin{aligned}
&(\ggg=) :: \mathit{Dist}\ a \rightarrow (a \rightarrow \mathit{Dist}\ b) \rightarrow \mathit{Dist}\ b \\
&d \ggg= f = \mathbf{let}\ \mathit{Dist}\ x\ p = d \\
&\qquad\qquad\quad\ \ \mathit{Dist}\ y\ q = f\ x \\
&\qquad\qquad\ \mathbf{in}\ \mathit{Dist}\ y\ (p \ast.\ q)
\end{aligned}
$$

We can observe two facts about this definition. First, the definition yields a *Dist*-constructor without matching any argument. Second, if neither the event nor the probability of the final distribution is evaluated, the application of the function f is not evaluated either.

We can observe these properties with some exemplary usages of ($\ggg=$). As a reference, we see that pattern matching the *Dist*-constructor of *coin* triggers the non-determinism and yields two results.

$$
\begin{aligned}
&> (\lambda(\mathit{Dist}\ x\ p) \rightarrow \mathit{True})\ \mathit{coin} \\
&\mathit{True} \\
&\mathit{True}
\end{aligned}
$$

In contrast, distributions resulting from an application of ($\ggg=$) behave differently. This time, pattern matching on the *Dist*-constructor does not trigger any non-determinism.

$$
\begin{aligned}
&> (\lambda(\mathit{Dist}\ x\ p) \rightarrow \mathit{True})\ (\mathit{certainly}\ ()\ \ggg= (\lambda y \rightarrow \mathit{coin})) \\
&\mathit{True} \\
&> (\lambda(\mathit{Dist}\ x\ p) \rightarrow \mathit{True})\ (\mathit{coin} \ggg= \mathit{certainly}) \\
&\mathit{True}
\end{aligned}
$$

We observe that the last two examples yield a single result, because the ($\ggg=$)-operator changes the position of the non-determinism. That is, the non-determinism does not reside at the same level as the *Dist*-constructor, but in the arguments of *Dist*. Therefore, we have to be sure to trigger all non-determinism when we query distributions. Not evaluating non-determinism might lead to false results when we sum up probabilities. Hence, non-strictness is a crucial property for positive pruning effects, but has to be used carefully.

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

Consider the following example usage of $(\ggg=)$, which is an inlined version of *joinWith* applied to the Boolean conjunction $(\wedge)$.

$> (\lambda(\mathit{Dist}\ x\ p) \rightarrow x)\ (\mathit{coin} \ggg= (\lambda b1 \rightarrow \mathit{coin} \ggg= (\lambda b2 \rightarrow \mathit{certainly}\ (b1 \wedge b2))))$
False
True
False

We lose one expected result from the distribution, because $(\wedge)$ is non-strict in its second argument in case the first argument is *False*. When the first *coin* evaluates to *False*, $(\ggg=)$ ignores the second coin and yields *False* straightaway. In this case, the non-determinism of the second *coin* is not triggered and we get only three instead of four results. The non-strictness of $(\wedge)$ has no consequences when using $(\ggg=!)$, because the operator evaluates both arguments, and thus, triggers the non-determinism. In the case of projecting to the event, we do not care about the missing result. However, when we sum up probabilities, we do not want events to get lost.

When we compute the total probability of a distribution, the result should always be 1.0. However, the query above has only three results and every event has a probability of 0.25, resulting in a total probability of 0.75. Here is the good news. While events can get lost when passing non-strict functions to $(\ggg=)$, probabilities never get lost. For example, consider the following application:

$> (\lambda(\mathit{Dist}\ x\ p) \rightarrow p)\ (\mathit{coin} \ggg= (\lambda b1 \rightarrow \mathit{coin} \ggg= (\lambda b2 \rightarrow \mathit{certainly}\ (b1 \wedge b2))))$
0.25
0.25
0.25
0.25

Since multiplication is strict, if we demand the resulting probability, the operator $(\ggg=)$ has to evaluate the *Dist*-constructor and its probability. That is, no values get lost if we evaluate the resulting probability. Fortunately, the query operation (??) calculates the total probability of the filtered distributions, thus, evaluates the probability as the following example shows:

$> \mathit{not}\ ??\ (\mathit{coin} \ggg= (\lambda b1 \rightarrow \mathit{coin} \ggg= (\lambda b2 \rightarrow \mathit{certainly}\ (b1 \wedge b2))))$
0.75

We calculate the probability of the event *False*. While there were only two *False* events when we projected to the event, the total probability of the event *False* is still 0.75, that is, three times 0.25, instead of only 0.5.

All in all, in order to benefit from non-strictness, all operations provided by the library have to use the right amount of strictness, not too much and not too little. For this reason, PFLP does not provide the *Dist*-constructor nor the corresponding projection functions to the user. With this restriction, the library guarantees that no relevant probabilities get lost.

3.4 Non-deterministic events

We assume that all events passed to library functions are deterministic, that is, the library does not support non-deterministic events within distributions. In order to illustrate why this restriction is crucial, we consider an example that breaks this rule.

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

$$
\begin{array}{l}
(\mathit{const\ True})\ ??\ \mathit{die2} \\
\equiv\quad \{\ \text{definition of } (??)\ \} \\
\mathit{foldValues}\ (+.)\ 0.0\ (\mathit{allValues}\ (\mathit{prob}\ (\mathit{filterDist}\ (\mathit{const\ True})\ \mathit{die2}))) \\
\equiv\quad \{\ \text{definition of } \mathit{die2}\ \} \\
\mathit{foldValues}\ (+.)\ 0.0\ (\mathit{allValues}\ (\mathit{prob}\ (\mathit{filterDist}\ (\mathit{const\ True})\ (\mathit{enum}\ [\mathit{side}]\ [\frac{1}{6}])))) \\
\equiv\quad \{\ \text{definition of } \mathit{enum}\ \} \\
\mathit{foldValues}\ (+.)\ 0.0\ (\mathit{allValues}\ (\mathit{prob}\ (\mathit{filterDist}\ (\mathit{const\ True})\ (\mathit{Dist\ side}\ \frac{1}{6})))) \\
\equiv\quad \{\ \text{definition of } \mathit{filterDist}\ \} \\
\mathit{foldValues}\ (+.)\ 0.0\ (\mathit{allValues}\ (\mathit{prob}\ (\mathbf{if}\ \mathit{const\ True\ side}\ \mathbf{then}\ \mathit{Dist\ side}\ \frac{1}{6}\ \mathbf{else}\ \mathit{failed}))) \\
\equiv\quad \{\ \text{definition of } \mathit{const}\ \} \\
\mathit{foldValues}\ (+.)\ 0.0\ (\mathit{allValues}\ (\mathit{prob}\ (\mathbf{if}\ \mathit{True}\ \mathbf{then}\ \mathit{Dist\ side}\ \frac{1}{6}\ \mathbf{else}\ \mathit{failed}))) \\
\equiv\quad \{\ \text{evaluate } \mathbf{if} - \mathbf{then} - \mathbf{else}\ \} \\
\mathit{foldValues}\ (+.)\ 0.0\ (\mathit{allValues}\ (\mathit{prob}\ (\mathit{Dist\ side}\ \frac{1}{6}))) \\
\equiv\quad \{\ \text{definition of } \mathit{prob}\ \} \\
\mathit{foldValues}\ (+.)\ 0.0\ (\mathit{allValues}\ \frac{1}{6}) \\
\equiv\quad \{\ \text{definition of } \mathit{allValues}\ \} \\
\mathit{foldValues}\ (+.)\ 0.0\ \{\ \frac{1}{6}\ \} \\
\equiv\quad \{\ \text{definition of } \mathit{foldValues}\ \} \\
\frac{1}{6}
\end{array}
$$

Fig. 4. Evaluation of a distribution that contains a free variable that is not demanded.

Curry provides free variables, that is, expressions that non-deterministically evaluate to every possible value of its type. When we revisit the definition of a die, we might be tempted to use a free variable instead of explicitly enumerating all values of type *Side*.

We can define a free variable of type *Side* as follows:

```
side :: Side
side = unknown
```

This free variable evaluates as follows:

```
> side
One
Two
Three
Four
Five
Six
```

With this information in mind consider the following alternative definition of a die, which is much more concise than explicitly listing all constructors of *Dist*.

```
die2 :: Dist Side
die2 = enum [side] [1/6]
```

We just use a free variable – the constant *side* – and pass the probability of each event as second parameter. Now, let us consider the following query.

```
> (const True) ?? die2
0.16666667
```

The result of this query is $\frac{1}{6}$ and not 1.0 as expected. Consider Figure 4 for a step-by-step evaluation of this expression in order to understand better what is going on. This example illustrates that probabilities can get lost if we do not use the right amount

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

of strictness. The predicate *const True* does not touch the event at all, thus does not trigger *side* to actually evaluate to all the constructors of *Side*. Then, the definition of (??) directly projects to the probability of *die2* and throws away all non-determinism left in *Dist side* $\frac{1}{6}$. Therefore, we lose probabilities we would like to sum up.

As a consequence for PFLP, non-deterministic events within a distribution are not allowed. If users of the library stick to this rule, it is not possible to misuse the operations and lose non-deterministic results due to non-strictness.

However, one approach to overcome this issue when using *enum* is to use an alternative stricter implementation. That is, we could easily adapt the strictness behavior of *enum* in order to allow a more declarative definition of distributions using free variables without affecting the overall advantage leveraged by non-strict functions.

3.5 Partial functions

Besides not using non-determinism for events, users have to follow another restriction. When using the bind operator ($\ggg=$), the second argument is a function of type $a \to$ *Dist b*, that is, constructs a new distribution. As we have discussed before, distributions need to sum up to a probability of 1.0, and the distributions we create via ($\ggg=$) are no exception. This restriction is violated if we use partial functions as second argument of ($\ggg=$). Recall the definition *coin* that describes a uniform distribution of type *Bool*, and consider the function *partialPattern* that depends on *coin*, but maps *False* to *failed*.

$$
\begin{aligned}
&\mathit{partialPattern} :: \mathit{Dist\ Bool}\\
&\mathit{partialPattern} = \mathit{coin} \ggg= (\lambda b \to \textbf{if}\ b\ \textbf{then}\ \mathit{certainly\ True}\ \textbf{else}\ \mathit{failed})
\end{aligned}
$$

Due to the partial pattern matching in *partialPattern*, the resulting distribution does not sum up to 1.0 anymore, thus, violates the rule for a valid distribution. By performing a query with the predicate *const True*, we can observe this property.

$$
\begin{aligned}
&> (\mathit{const\ True})\ ??\ \mathit{partialPattern}\\
&0.5
\end{aligned}
$$

We only allow to filter distributions when a probability is computed using (??), but not in any other situation. In the current implementation, this restriction on functions when using ($\ggg=$) is neither statically nor dynamically enforced, but a coding convention that users should keep in mind and follow when working with the library.

3.6 Monad laws

When we comply with the restrictions we have discussed above, the operators ($\ggg=$) and *certainly* allow us to formulate probabilistic programs as one would expect. However, there is one obvious question that we did not answer yet. We did not check whether the operator ($\ggg=$) together with *certainly* actually forms a monad as the name of the operator suggests. That is, we have to check whether the following three laws hold for all distributions d and all values x, f, and g of appropriate types.

- $d \ggg= \mathit{certainly} \equiv d$
- $\mathit{certainly}\ x \ggg= f \equiv f\ x$
- $(d \ggg= f) \ggg= g \equiv d \ggg= (\lambda y \to f\ y \ggg= g)$

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

In the previous subsection, we have already observed that the first equality does not hold in general. For example, we have seen that there is a context that is able to distinguish the left-hand side from the right-hand side. For instance, while the expression

$$(\lambda(Dist\ x\ p) \to True)\ coin$$

yields *True* twice, the expression

$$(\lambda(Dist\ x\ p) \to True)\ (coin \ggg= certainly)$$

yields *True* only once. As most Curry semantics are based on sets – and not on multisets, the two sides of the equality would be the same. Notwithstanding, in a multiset semantics the user could still not observe the difference between the two expressions because he does not have access to the *Dist*-constructor. The user cannot pattern match on a *Dist*-constructor, but only use the combinator (??) to inspect a distribution.

In order to discuss the validity of the monad laws more rigorously, we apply equational reasoning to check whether the monad laws might fail.

The first monad law Let $d :: Dist\ \tau$, then we reason as follows about the first monad law:

$d \ggg= certainly$
$\equiv$ { Definition of $(\ggg=)$ }
let $Dist\ x\ p = d$
$\quad Dist\ y\ q = certainly\ x$
in $Dist\ y\ (p *. q)$
$\equiv$ { Definition of *certainly* }
let $Dist\ x\ p = d$
$\quad Dist\ y\ q = Dist\ x\ 1.0$
in $Dist\ y\ (p *. q)$
$\equiv$ { Inlining of $Dist\ y\ q = Dist\ x\ 1.0$ }
let $Dist\ x\ p = d$ **in** $Dist\ x\ (p *. 1.0)$
$\equiv$ { Definition of $(*.)$ }
let $Dist\ x\ p = d$ **in** $Dist\ x\ p$
$\stackrel{?}{\equiv}$
d

Does the last step hold in general? It looks good for the deterministic case with $d = Dist\ evnt\ prb$.

let $Dist\ x\ p = Dist\ evnt\ prb$ **in** $Dist\ x\ p$
$\equiv$
$Dist\ evnt\ prb$

However, the equality **let** $Dist\ x\ p = d$ **in** $Dist\ x\ p \equiv d$ does not hold in general. For instance, let us consider the case $d = failed$.

let $Dist\ x\ p = failed$ **in** $Dist\ x\ p$
$\equiv$
$Dist\ failed\ failed$
$\not\equiv$
$failed$

That is, the left-hand side is more defined than the right-hand side if $d = failed$.

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

Because the user cannot access the *Dist*-constructor she cannot observe this difference. The user can only compare two distributions by using the querying operator (??). Therefore, in the following, we will show that the monad laws hold if we consider a context of the form *pred* ?? *d*, where *pred* is an arbitrary predicate. Recall that we defined the operator (??) as follows:

$$(??) :: (a \to Bool) \to Dist\ a \to Probability$$
$$(??)\ pred\ d = foldValues\ (+.)\ 0.0\ (allValues\ (prob\ (filterDist\ pred\ d)))$$

Fortunately, the monad laws already hold if we consider the context *filterDist pred* for an arbitrary predicate $pred :: a \to Bool$. Therefore, we will show that the following equalities hold for all distributions d, and all values x, *pred*, f, and g of appropriate types.

(1) $filterDist\ pred\ (d \ggg= certainly) \equiv filterDist\ pred\ d$
(2) $filterDist\ pred\ (certainly\ x \ggg= f) \equiv filterDist\ pred\ (f\ x)$
(3) $filterDist\ pred\ ((d \ggg= f) \ggg= g)) \equiv filterDist\ pred\ (d \ggg= (\lambda y \to f\ y \ggg= g))$

In the following, we will first show that equation (1) holds. We reason as follows for all distributions $d :: Dist\ \tau$ and predicates $pred :: \tau \to Bool$.

$filterDist\ pred\ (d \ggg= certainly)$
$\equiv$ { Reasoning above }
$filterDist\ pred\ (\mathbf{let}\ Dist\ x\ p = d\ \mathbf{in}\ Dist\ x\ p)$
$\equiv$ { Definition of *filterDist* }
$\mathbf{let}\ Dist\ y\ q = (\mathbf{let}\ Dist\ x\ p = d\ \mathbf{in}\ Dist\ x\ p)$
$\mathbf{in\ if}\ (pred\ y)\ \mathbf{then}\ (Dist\ y\ q)\ \mathbf{else}\ failed$
$\equiv$ { Inline **let**-declaration }
$\mathbf{let}\ Dist\ y\ q = d$
$\mathbf{in\ if}\ (pred\ y)\ \mathbf{then}\ (Dist\ y\ q)\ \mathbf{else}\ failed$
$\equiv$ { Definition of *filterDist* }
$filterDist\ pred\ d$

The ($\ggg=$)-operator defers the pattern matching to the right-hand side via a **let**-expression. This so-called lazy pattern matching causes the monad laws to not hold without any context. However, because *filterDist* introduces a lazy pattern matching via a **let**-expression as well, observing two distributions via *filterDist* hides the difference between the two sides of the equation.

The second monad law. For the second monad law (2), we reason as follows for all $x :: \tau_1$, and all $f :: \tau_1 \to Dist\ \tau_2$.

$certainly\ x \ggg= f$
$\equiv$ { Definition of ($\ggg=$) }
$\mathbf{let}\ Dist\ y\ p = certainly\ x$
$\quad Dist\ z\ q = f\ y$
$\mathbf{in}\ Dist\ z\ (p *. q)$
$\equiv$ { Definition of *certainly* }
$\mathbf{let}\ Dist\ y\ p = Dist\ x\ 1.0$
$\quad Dist\ z\ q = f\ y$
$\mathbf{in}\ Dist\ z\ (p *. q)$

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

$\equiv$ { Inlining of $Dist\ y\ p = Dist\ x\ 1.0$ }
let $Dist\ z\ q = f\ x$ **in** $Dist\ z\ (1.0 *.\ q)$
$\equiv$ { Definition of $(*.)$ }
let $Dist\ z\ q = f\ x$ **in** $Dist\ z\ q$
$\stackrel{?}{\equiv}$
$f\ x$

Here, we observe the same restrictions as before, for example, if f yields *failed* for any argument x the equality does not hold. Once again, we consider the context *filterDist pred* for all $pred :: \tau_2 \rightarrow Bool$ to reason that the user cannot observe the difference.

filterDist pred (**let** $Dist\ z\ q = f\ x$ **in** $Dist\ z\ q$)
$\equiv$ { Definition of *filterDist* }
let $Dist\ x\ p =$ (**let** $Dist\ z\ q = f\ x$ **in** $Dist\ z\ q$)
in if $(pred\ x)$ **then** $(Dist\ x\ p)$ **else** *failed*
$\equiv$ { Inline **let**-declaration }
let $Dist\ x\ p = f\ x$
in if $(pred\ x)$ **then** $(Dist\ x\ p)$ **else** *failed*
$\equiv$ { Definition of *filterDist* }
filterDist pred $(f\ x)$

Fortunately, the second monad law holds as well in the context of *filterDist*.

The third monad law. We do not discuss the third monad law (3) in detail here, as it holds without restrictions – even without the additional context using *filterDist*. All in all, *certainly* and ($\ggg=$) form a valid monad from the user's point of view.

4 Applications and evaluation

After presenting the basic combinators of the library and motivating the advantages of modeling distributions using non-determinism, we will implement some exemplary applications. We reimplement examples that have been characterized as challenging for probabilistic logic programming by Nampally and Ramakrishnan (2015), who use the examples to discuss the expressiveness of probabilistic logic programming and its cost with respect to performance. The examples focus on properties of random strings and their probabilities. Furthermore, we show benchmarks of these examples and compare them with the probabilistic languages ProbLog (Kimmig *et al.* 2011) and WebPPL (Goodman and Stuhlmüller 2014). These comparisons confirm the advantages of non-strict non-determinism with respect to performance.

4.1 Random strings

In order to compare our library with other approaches for probabilistic programming, we reimplement two examples about random strings that have also been implemented in ProbLog. This implementation can be found online.[13] We generate random strings of a

[13] https://dtai.cs.kuleuven.be/problog/tutorial/various/04_nampally.html.

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

fixed length over the alphabet $\{a, b\}$ and calculate the probability that this string (a) is a palindrome and (b) contains the subsequence *bb*.

First, we define a distribution that picks a character uniformly from the alphabet $\{a, b\}$.

```
pickChar :: Dist Char
pickChar = uniform ['a','b']
```

Based on *pickChar*, we define a distribution that generates a random string of length n, that is, picks a random char n times. We reuse *replicateDist* to define this distribution.

```
randomString :: Int → Dist String
randomString n = replicateDist n (λ() → pickChar)
```

In order to compute the probability that a random string is a palindrome and contains a subsequence *bb*, respectively, we define predicates that test these properties for a given string. A string is a palindrome, if it reads the same forwards and backwards. The following predicate, thus, checks if the reverse of a given string is equal to the original string.

```
palindrome :: String → Bool
palindrome str = str ≡ reverse str
```

The predicate that checks if a string contains two consecutive *b*s can be easily defined via pattern matching and recursion.

```
consecutiveBs :: String → Bool
consecutiveBs str = case str of
                      []                → False
                      ('b' : 'b' : rest) → True
                      (c : rest)        → consecutiveBs rest
```

Now, we are ready to perform some queries. What is the probability that a random string of length 5 is a palindrome?

```
> palindrome ?? (randomString 5)
0.25
```

What is the probability that a random string of length 10 contains two consecutive *b*s?

```
> consecutiveBs ?? (randomString 10)
0.859375
```

In general the above definitions of *palindrome* and *consecutiveBs* are quite naive, and thus, inefficient because all strings of the given length have to be enumerated explicitly. Due to the inefficiency, the ProbLog homepage introduces a more efficient version for both problems. In the following, we will discuss the alternative implementation to compute the probability for a palindrome only. This more efficient version has arguments for the index of the front and back position, picks characters for both ends and then moves the position toward the middle. That is, instead of naively generating the whole string of length n, this version checks each pair of front and back position first and fails straightaway, if

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

they do not match. If the characters do match, the approach continues by moving both indices toward each other. In Curry, an implementation of this idea looks as follows:

$$
\begin{array}{l}
\mathit{palindromeEfficient} :: \mathit{Int} \to \mathit{Dist}\ (\mathit{Bool}, \mathit{String}) \\
\mathit{palindromeEfficient}\ n = \mathit{palindrome}'\ 1\ n \\
\mathit{palindrome}' :: \mathit{Int} \to \mathit{Int} \to \mathit{Dist}\ (\mathit{Bool}, \mathit{String}) \\
\mathit{palindrome}'\ n1\ n2 \mid n1 \equiv n2 \ = \mathit{pickChar} \ggg\!= (\lambda c \to \mathit{certainly}\ (\mathit{True}, [c])) \\
\qquad\qquad\qquad\quad \mid n1 > n2 \ = \mathit{certainly}\ (\mathit{True}, [\,]) \\
\qquad\qquad\qquad\quad \mid \mathit{otherwise} = \mathit{pickChar} \ggg\!= (\lambda c1 \to \\
\qquad\qquad\qquad\qquad\qquad\qquad \mathit{pickChar} \ggg\!= (\lambda c2 \to \\
\qquad\qquad\qquad\qquad\qquad\qquad (\mathit{palindrome}'\ (n1+1)\ (n2-1)) \ggg\!= (\lambda(b, cs) \to \\
\qquad\qquad\qquad\qquad\qquad\qquad \mathit{certainly}\ (c1 \equiv c2 \wedge b, c1 : (cs \mathbin{+\!\!+} [c2])))))
\end{array}
$$

The interesting insight here is that, thanks to the combination of non-determinism and non-strictness, the evaluation of the first query based on *palindrome* behaves similar to the efficient variant in ProbLog. At first, it seems that the query performs poorly, because the predicate *palindrome* needs to evaluate the whole list due to the usage of *reverse*. The good news is, however, that the non-determinism is only spawned if we evaluate the elements of that list, and the elements still evaluate non-strictly, when explicitly triggered by $(\equiv)$. More precisely, because of the combination of *reverse* and $(\equiv)$, the evaluation starts by checking the first and last characters of a string and only continues to check more characters, and spawn more non-determinism, if they match. If these characters do not match, the evaluation fails directly and does not need to check any more characters. In a nutshell, when using PFLP, we get a version competitive with efficient implementations although we used a naive generate and test approach.

4.2 Performance comparisons with other languages

Up to now, the only performance comparisons we discussed were for different implementations of our library in Curry and Haskell. These comparisons showed the advantage of using non-strict non-determinism concepts for the implementation of the library. Next, we want to take a look at the comparison with the full-blown probabilistic programming languages ProbLog and WebPPL. ProbLog is a probabilistic extension of Prolog that is implemented in Python. WebPPL is the successor of Church; in contrast to Church it is not implemented in Scheme but in JavaScript.

In order to try to measure the execution of the programs only, we precompiled the executable for the Curry programs. As Python is an interpreted language, a similar preparation was not available for ProbLog. However, we used ProbLog as a library in order to call the Python[14] interpreter directly. ProbLog is mainly implemented in Python, which allows users to import ProbLog as a Python package.[15] For WebPPL, we used *node.js*[16] to run the JavaScript program as a terminal application. All of the following running times are the mean of 1000 runs as calculated by the Haskell tool bench[17] that we use to run the benchmarks.

[14] We use version 2.7.10 of Python.
[15] https://dtai.cs.kuleuven.be/problog/tutorial/advanced/01_python_interface.html
[16] We use version 8.12.0 of *node.js*.
[17] https://hackage.haskell.org/package/bench

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

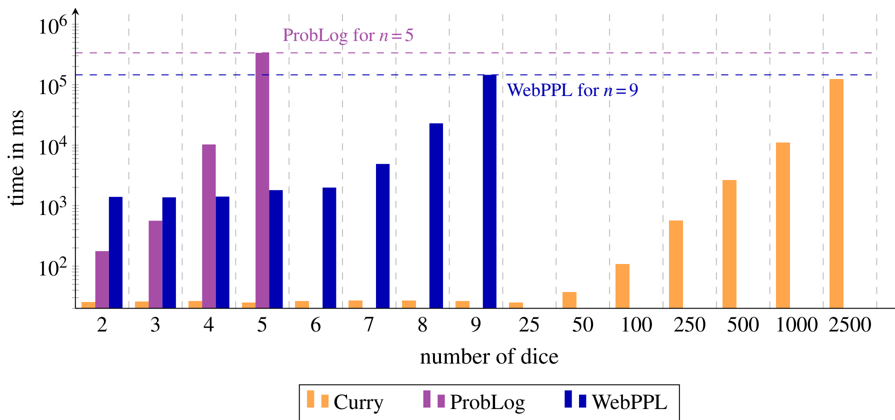

Fig. 5. Getting only sixes when rolling n dice.

We compare the running times based on the two examples we already discussed: the dice rolling example presented in Section 3.2 and the palindrome example from the previous subsection.

Dice Rolling. As discussed before, non-strict non-determinism performs pretty well for the dice rolling example, as a great deal of the search space is pruned early. Figure 5 shows an impressive advantage of our library in comparison with ProbLog and WebPPL. The x-axis represents the number of rolled dice and we present the time in milliseconds in logarithmic scale on the y-axis.

In order to demonstrate that our library outperforms ProbLog and WebPPL by several orders of magnitude for this example, we also run the Curry implementation for bigger values of n that eventually had the same running time as the last tested value for the other languages. The right part of Figure 5 shows the running times for 25–5000 dice. We can see that our library can compute the probability for getting only sixes for 2500 dice in roughly the same time as ProbLog for 5 dice. The running times for WebPPL seem very bad in the beginning, but after a few throws it becomes obvious that there is a constant overhead. In fact, Nogatz *et al.* (2018) observe and discuss this overhead as well. Nevertheless, whereas WebPPL computes the probability for 9 dice, our library can compute the probability for 2500 dice in roughly the same time.

Palindrome. In order to back up the results of the previous example, Figure 6 shows benchmarks for implementations of the naive and the efficient versions in Curry, ProbLog and WebPPL. The x-axis represents the length of the generated palindrome and, once again, we present the time in milliseconds in logarithmic scale on the y-axis.

The figure uses dashed bars for the efficient version of the algorithm and a solid filling for the naive algorithm. The naive algorithm scales pretty bad in ProbLog and WebPPL. The Curry version is still applicable up to 30 as its running time is similar to all three efficient versions. Overall, the efficient versions all perform in a similar time range, but WebPPL shows a slight performance advantage for an increasing length of the string. More precisely, the efficient WebPPL implementation performs a query for strings of length 50 in the same time as the efficient Curry and ProbLog perform a query for

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

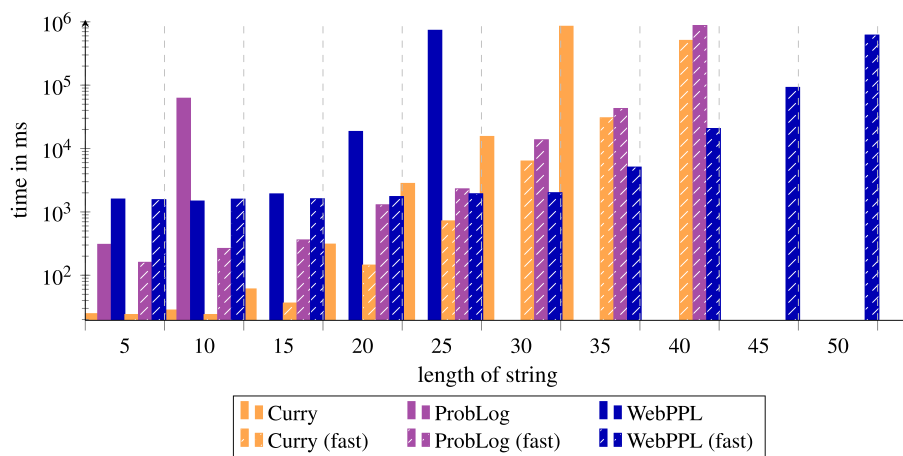

Fig. 6. Palindrome computation for a string of length n.

strings of length 40. That is, the efficient WebPPL implementation outperforms the other implementations by roughly two orders of magnitude.

5 Related and future work

The approach of this paper is based on the work by Erwig and Kollmansberger (2006), who introduce a Haskell library that represents distributions as lists of event-probability pairs. Their library also provides a simple sampling mechanism to perform inference on distributions. Inference algorithms come into play because common examples in probabilistic programming have an exponential growth and it is not feasible to compute the whole distribution. Similarly, Ścibior *et al.* (2015) present a more efficient implementation using a DSL in Haskell. They represent distributions as a free monad and inference algorithms as an interpretation of the monadic structure. Thanks to this interpretation, the approach is competitive to full-blown probabilistic programming languages with respect to performance. PFLP provides functions to sample from distributions as well. However, in this work, we focus on modeling distributions and do not discuss any sampling mechanism. In particular, as future work we plan to investigate whether we can benefit from the improved performance as presented here in the case of sampling. Furthermore, a more detailed investigation of the performance of non-determinism in comparison to a list model is a topic for another paper.

The benefit with respect to the combination of non-strictness and non-determinism is similar to the benefit of property-based testing using Curry-like non-determinism in Haskell (Runciman *et al*. 2008) and Curry (Christiansen and Fischer 2008). In property-based testing, sometimes we want to generate test cases that satisfy a precondition. With Curry-like non-determinism the precondition can prune the search space early, while a list-based implementation has to generate all test cases and filter them afterwards. Both applications, probabilistic programming and property-based testing, are examples, where built-in non-determinism outperforms list-based approaches as introduced by Wadler (1985). In comparison to property-based testing, here, we observe that we can even add a kind of monadic layer on top of the non-determinism that computes additional

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

information and still preserve the demand-driven behavior. However, the additional information has to be evaluated strictly – as it is the case for probabilities, otherwise we might lose non-deterministic results.

There are other more elaborated approaches to implement a library for probabilistic programming. For example, Kiselyov and Shan (2009) extend their library for probabilistic programming in OCaml with a construct for lazy evaluation to achieve similar behavior with respect to efficiency. However, they use lazy evaluation for a concrete application based on importance sampling. Due to the combination of non-strictness and non-determinism, we can efficiently calculate the total probability of the resulting distribution without utilizing sampling.

As future work, we see a high potential for performance improvements for the Curry compiler KiCS2. PFLP serves as a starting point for further studies of functional logic features in practical applications. For example, we would expect the running times of the strict implementation based on non-determinism to be approximately as efficient as a list-based implementation. However, as the numbers in Section 3 show, the list approach is considerably faster.

The library's design does not support the use of non-determinism in events or probabilities of a distribution. In case of deeper non-determinism, we have to be careful to trigger all non-determinism when querying a distribution as shown in Section 3. Hence, the extension of the library with an interface using non-determinism on the user's side is an idea worth studying.

Last but not least, we see an opportunity to apply ideas and solutions of the functional logic paradigm in probabilistic programming. For instance, Christiansen *et al.* (2010) investigate free theorems for functional logic programs. As their work considers non-determinism and sharing, adapting it to probabilistic programming should be easy. As another example, Braßel (2009) presents a debugger for Curry that works well with non-determinism. Hence, it should be possible to reuse these ideas in the setting of probabilistic programming as well.

6 Conclusion

We have implemented a simple library for probabilistic programming in a functional logic programming language, namely Curry. Such a library proves to be a good fit for a functional logic language, because both paradigms share similar features. While other libraries need to reimplement features specific to probabilistic programming, we solely rely on core features of functional logic languages.

The key idea of the library is to use non-determinism to model distributions. We discussed design choices as well as the disadvantages and advantages that result from this approach. In the end, the library provides non-strict probabilistic combinators in order to avoid spawning unnecessary non-deterministic computations. These non-strict combinators have benefits in terms of performance due to early pruning. Using combinators that are too strict leads to a loss of these performance benefits. Fortunately, the user does not have to worry about using the right amount of strictness as long as she only uses the provided combinators.

There are, however, two restrictions the user has to follow when using the library. If the user does not follow these restrictions, a program may behave unexpectedly. Events may

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

not be non-deterministic and the second argument of ($\gg=$)-operator may not be partial. Notwithstanding, we want to emphasize that the restrictions do not affect expressibility. In fact, a programming language like ProbLog shows similar behavior when mixing non-determinism and probabilities as our implementation.

Last but not least, we showed that the library obeys the expected monad laws with respect to observational behavior and reimplemented examples from the probabilistic programming literature to compare the performance of our library with other existing languages.

References

Antoy, S. 2005. Evaluation strategies for functional logic programming. *Journal of Symbolic Computation 40,* 1 875–903.

Antoy, S. and Hanus, M. 2010. Functional logic programming. *Communications of the ACM 53,* 4, 74–85.

Braßel, B. 2009. A technique to build debugging tools for lazy functional logic languages. *Electronic Notes in Theoretical Computer Science 246*, 63–76.

Braßel, B., Hanus, M. and Huch, F. 2004. Encapsulating non-determinism in functional logic computations. *Journal of Functional and Logic Programming 2004,* 6, Special Issue 1.

Braßel, B., Hanus, M., Peemöller, B. and Reck, F. 2011. KiCS2: A New Compiler from Curry to Haskell. In *Proceedings of the 20th International Conference on Functional and Constraint Logic Programming*, H. Kuchen, Ed. WFLP '11. Springer-Verlag, Odense, Denmark, 1–18.

Christiansen, J. and Fischer, S. 2008. EasyCheck – Test Data for Free. In *Proceedings of the International Symposium on Functional and Logic Programming.* Springer, Berlin, Heidelberg, 322–336.

Christiansen, J., Seidel, D. and Voigtländer, J. 2010. Free Theorems For Functional Logic Programs. In *Proceedings of the 4th ACM SIGPLAN Workshop on Programming Languages Meets Program Verification.* ACM, New York, NY, USA, 39–48.

Christiansen, J., Seidel, D. and Voigtländer, J. 2011. An Adequate, Denotational, Functional-Style Semantics for Typed FlatCurry. In *Proceedings of the 19th International Workshop on Functional and (Constraint) Logic Programming.* LNCS. Springer-Verlag, Madrid, Spain.

De Raedt, L. and Kimmig, A. 2013. Probabilistic programming concepts. *arXiv:1312.4328 (preprint).*

Dylus, S., Christiansen, J. and Teegen, F. 2018. Probabilistic Functional Logic Programming. In *Practical Aspects of Declarative Languages.* PADL '18. Springer International Publishing, Los Angeles, CA, USA, 3–19.

Erwig, M. and Kollmansberger, S. 2006. Functional pearls: Probabilistic functional programming in Haskell. *Journal of Functional Programming 16,* 1, 21–34.

Fischer, S., Kiselyov, O. and Shan, C.-c. 2009. Purely Functional Lazy Non-deterministic Programming. In *Proceedings of the 14th ACM SIGPLAN International Conference on Functional Programming.* ACM, New York, NY, USA, 11–22.

Goodman, N. D., Mansinghka, V. K., Roy, D. M., Bonawitz, K. and Tenenbaum, J. B. 2008. Church: A Language for Generative Models. *CoRR*, 220–229.

Goodman, N. D. and Stuhlmüller, A. 2014. The Design and Implementation of Probabilistic Programming Languages. http://dippl.org.

Gordon, A. D., Henzinger, T. A., Nori, A. V. and Rajamani, S. K. 2014. Probabilistic Programming. In *Proceedings of the on Future of Software Engineering.* FOSE '14. ACM, Hyderabad, India, 167–181.

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

HANUS, M. 2017. *PAKCS: The Portland Aachen Kiel Curry System*. Published: Available at http://www.informatik.uni-kiel.de/~pakcs/.

KIMMIG, A., DEMOEN, B., DE RAEDT, L., COSTA, V. S. AND ROCHA, R. 2011. On the implementation of the probabilistic logic programming language ProbLog. *Theory and Practice of Logic Programming 11*, 235–262.

KISELYOV, O. AND SHAN, C.-C. 2009. Embedded probabilistic programming. In *Domain-Specific Languages*. Springer, Berlin, Heidelberg, 360–384.

MCCALLUM, A., SCHULTZ, K. AND SINGH, S. 2009. FACTORIE: Probabilistic Programming via Imperatively Defined Factor Graphs. In *Proceedings of the International Conference on Neural Information Processing Systems*. Vancouver, British Columbia, Canada.

NAMPALLY, A. AND RAMAKRISHNAN, C. R. 2015. Constraint-Based Inference in Probabilistic Logic Programs. In *Workshop on Probabilistic Programming*, 46–56.

NOGATZ, F., FRÜHWIRTH, T. AND SEIPEL, D. 2018. CHR.js: A CHR implementation in JavaScript. In *Rules and Reasoning*. Springer, 131–146.

PAIGE, B. AND WOOD, F. 2014. A Compilation Target for Probabilistic Programming Languages. In *Proceedings of the International Conference on Machine Learning*. Beijing, China.

RUNCIMAN, C., NAYLOR, M. AND LINDBLAD, F. 2008. SmallCheck and Lazy SmallCheck: Automatic Exhaustive Testing for Small Values. In *Proceedings of the First ACM SIGPLAN Symposium on Haskell*. Haskell '08. ACM, Victoria, BC, Canada, 37–48.

ŚCIBIOR, A., GHAHRAMANI, Z. AND GORDON, A. D. 2015. Practical Probabilistic Programming with Monads. In *Proceedings of the 2015 ACM SIGPLAN Symposium on Haskell*. Haskell '15. ACM, Vancouver, BC, Canada, 165–176.

VARACCA, D. AND WINSKEL, G. 2006. Distributing Probability over Non-determinism. *Mathematical Structures in Computer Science 16*, 87–113.

WADLER, P. 1985. How to Replace Failure by a List of Successes. In *Proceedings of the International Conference on Functional Programming Languages and Computer Architecture*. ACM, New York, NY, USA, 113–128.

WOOD, F., MEENT, J. W. AND MANSINGHKA, V. 2014. A new Approach to Probabilistic Programming Inference. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*. Proceedings of Machine Learning Research. PMLR, Reykjavik, Iceland, 1024–1032.

Appendix: Predefined Curry Functions

This section presents all predefined functions that are used in this paper.

$$
\begin{aligned}
&(\wedge) :: \mathit{Bool} \to \mathit{Bool} \to \mathit{Bool}\\
&\mathit{False} \wedge x = \mathit{False}\\
&\mathit{True} \wedge x = x
\end{aligned}
$$

$$
\begin{aligned}
&\mathit{foldr} :: (a \to b \to b) \to b \to [a] \to b\\
&\mathit{foldr}\ f\ z\ []\qquad\ = z\\
&\mathit{foldr}\ f\ z\ (x : \mathit{xs}) = f\ x\ (\mathit{foldr}\ f\ z\ \mathit{xs})
\end{aligned}
$$

$$
\begin{aligned}
&\mathit{zipWith} :: (a \to b \to c) \to [a] \to [b] \to [c]\\
&\mathit{zipWith}\ f\ []\qquad\ \mathit{ys}\qquad = []\\
&\mathit{zipWith}\ f\ (x : \mathit{xs})\ []\qquad = []\\
&\mathit{zipWith}\ f\ (x : \mathit{xs})\ (y : \mathit{ys}) = f\ x\ y : \mathit{zipWith}\ f\ \mathit{xs}\ \mathit{ys}
\end{aligned}
$$

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press

$$
\begin{aligned}
&\mathit{repeat} :: a \rightarrow [\,a\,] \\
&\mathit{repeat}\ x = x : \mathit{repeat}\ x
\end{aligned}
$$

$$
\begin{aligned}
&\mathit{length} :: [\,a\,] \rightarrow \mathit{Int} \\
&\mathit{length}\ [\,] \qquad\ \ = 0 \\
&\mathit{length}\ (x : \mathit{xs}) = 1 + \mathit{length}\ \mathit{xs}
\end{aligned}
$$

$$
\begin{aligned}
&\mathit{filter} :: (a \rightarrow \mathit{Bool}) \rightarrow [\,a\,] \rightarrow [\,a\,] \\
&\mathit{filter}\ p\ [\,] \qquad\ \ = [\,] \\
&\mathit{filter}\ p\ (x : \mathit{xs}) = \mathbf{if}\ p\ x\ \mathbf{then}\ x : \mathit{filter}\ p\ \mathit{xs}\ \mathbf{else}\ \mathit{filter}\ p\ \mathit{xs}
\end{aligned}
$$

$$
\begin{aligned}
&\mathit{all} :: (a \rightarrow \mathit{Bool}) \rightarrow [\,a\,] \rightarrow \mathit{Bool} \\
&\mathit{all}\ p\ \mathit{xs} = \mathit{foldr}\ (\wedge)\ \mathit{True}\ (\mathit{map}\ p\ \mathit{xs})
\end{aligned}
$$

$$
\begin{aligned}
&\mathit{not} :: \mathit{Bool} \rightarrow \mathit{Bool} \\
&\mathit{not}\ \mathit{False} = \mathit{True} \\
&\mathit{not}\ \mathit{True} = \mathit{False}
\end{aligned}
$$

$$
\begin{aligned}
&\mathit{const} :: a \rightarrow b \rightarrow a \\
&\mathit{const}\ x\ y = x
\end{aligned}
$$

https://doi.org/10.1017/S1471068419000085 Published online by Cambridge University Press