


FAKULTÄT FÜR INFORMATIK
DER TECHNISCHEN UNIVERSITÄT MÜNCHEN

Master's Thesis in Informatics

**A Verified Compiler for
Probability Density Functions**

Manuel Eberl




FAKULTÄT FÜR INFORMATIK
DER TECHNISCHEN UNIVERSITÄT MÜNCHEN

Master's Thesis in Informatics

A Verified Compiler for
Probability Density Functions

Verifikation eines Compilers für
Wahrscheinlichkeitsdichten

Author: Manuel Eberl
Supervisor: Prof Tobias Nipkow, PhD
Advisor: Dr Johannes Hölzl
Date: October 6, 2014



I confirm that this Master's thesis is my own work and that I have documented all sources and material used.

Garching, 6 October 2014

Manuel Eberl

Abstract

Bhat *et al.* [BBGR13] developed an inductive compiler that computes density functions for probability spaces described by programs in a probabilistic functional language. In this thesis, we give a verified implementation of such a compiler for a modified version of this language. We use the theorem prover Isabelle to give a formal proof of its soundness w. r. t. the semantics of the source and target language.

This is done in two steps: first, an abstract compiler working with abstract functions modelled directly in the theorem prover's logic is defined and proven sound. This compiler is then refined to a concrete version that returns a target-language expression.

Zusammenfassung

Bhat *et al.* [BBGR13] entwickelten einen induktiven Compiler zur automatischen Berechnung von Dichten bestimmter Wahrscheinlichkeitsräume. Diese werden durch Programme in einer probabilistischen funktionalen Sprache beschrieben. In dieser Arbeit wird solch ein Compiler für eine ähnliche Sprache in dem Theorembeweiser Isabelle implementiert und seine Korrektheit bezüglich der Semantik der zugrundeliegenden Ausgangs- und Zielsprache formal bewiesen.

Dies geschieht in zwei Schritten: Zunächst wird ein abstrakter Compiler definiert und korrekt bewiesen. Dieser arbeitet auf abstrakten Funktionen, die direkt in der Logik des Theorembeweisers modelliert sind. Dieses Ergebnis wird dann durch Refinement auf einen konkreten Compiler übertragen, der Zielsprachenausdrücke zurückliefert.

Contents

1	Introduction	2
1.1	Motivation	2
1.2	Related work	3
1.3	Utilised tools	3
1.4	Outline	4
2	Preliminaries	5
2.1	Notation	5
2.1.1	Typographical notes	5
2.1.2	Deviations from standard mathematical notation	5
2.1.3	Semantics notation	6
2.2	Mathematical basics	6
2.2.1	Basic measure spaces	6
2.2.2	Sub-probability spaces	7
2.2.3	The category Meas	7
2.2.4	The sub-probability functor	7
2.2.5	Giry monad	8
3	Language Syntax and Semantics	10
3.1	Types, values, and operators	10
3.2	Auxiliary definitions	12
3.2.1	Measure embeddings	12
3.2.2	Stock measures	12
3.3	Source language	13
3.4	Deterministic expressions	16
3.5	Target language	16
4	Abstract compiler	19
4.1	Density contexts	19
4.2	Definition	20
4.3	Soundness proof	23
5	Concrete compiler	24
5.1	Approach	24
5.2	Definition	24
5.3	Refinement	26
5.4	Final result	27
5.5	Evaluation	28
6	Conclusion	30
6.1	Breakdown	30
6.2	Difficulties	30
6.3	Future work	31
6.4	Summary	31

Appendix	32
A Notation and Auxiliary Functions	32
A.1 General notation	32
A.2 Semantics notation and general auxiliary functions	33
A.3 Target language auxiliary functions	34
B References	35

Acknowledgements

I would like to thank my advisor, Dr Johannes Hölzl, for the numerous and fruitful discussions I had with him in the course of writing this thesis and in particular for helping me over the initial difficulties. My thanks also go to Prof Tobias Nipkow for assigning to me a topic so perfectly tailored to my interests.

I further thank Kristina Magnussen for extensive proofreading and her invaluable general support and Andrei Popescu for verifying the correctness of the Category Theory section. Furthermore, David Cock and Jeremy Avigad helped me with their quick and concise answers to my questions about their work.

1 Introduction

1.1 Motivation

Random distributions of practical significance can often be expressed as probabilistic functional programs. A simple example would be a board game in which a die is cast and, depending on the result, more dice are cast and their result influences the game in some way.

For instance, consider the Pen and Paper Role Playing Game *Call of Cthulhu*: players have a number of *skills* such as *Archæology*, *Library Usage*, etc. A player's proficiency in a skill is a number between 0 and 100, and when a player uses a skill in an adventure, she must roll a d100 (a die with numbers from 1 to 100). If the result is no higher than her proficiency in that skill, she succeeds; otherwise, she fails. A result between 1 and 5 or below one fifth of her proficiency is a *critical success*, resulting in her receiving a chance to increase her proficiency in the skill: she rolls another d100 and if the result is *above* her proficiency, the skill proficiency is permanently increased by an amount determined through a d10 roll. This can be expressed concisely as a probabilistic program in the following way:

```

let  $x = \text{UniformInt}(1, 100)$ 
in if  $x > \text{skill}$  then
    (FAILURE,  $\text{skill}$ )
else if  $x \leq 5 \vee x \cdot 5 \leq \text{skill}$  then
    let  $\text{increase} = \text{if } \text{UniformInt}(1, 100) > \text{skill} \text{ then } \text{UniformInt}(1, 10) \text{ else } 0$ 
    in (CRITICAL_SUCCESS,  $\text{skill} + \text{increase}$ )
else
    (SUCCESS,  $\text{skill}$ )

```

Finite and discrete examples such as this can be studied with brute force, i. e. simply summing the probabilities for every single case, but more complex random experiments involving an infinite number of outcomes (e. g. Poisson distribution) or continuous distributions require a more refined approach.

When studying a random distribution such as this, it is often desirable to determine its *probability density function* (PDF). This can be used to e. g. determine the expectation or sample the distribution with a sampling method such as *Markov-chain Monte Carlo* (MCMC).

In 2013, Bhat *et al.* presented a compiler that computes the probability distribution function of a program in the probabilistic functional language *Fun* [BBGR13]. They evaluated the compiler on a number of practical problems and concluded that it reduces the amount of time and effort required to model them in an MCMC system significantly compared to hand-written models.

Bhat *et al.* also stated that their eventual goal is the formal verification of such a compiler in a theorem prover [BAVG12]. This has the advantage of providing *guaranteed correctness*, i. e. the result of the compilation is provably a PDF for the source expression, according to the formal semantics. This greatly increases the confidence one has in the compiler.

In this Master’s thesis, we implemented such a compiler for a similar probabilistic functional language in the interactive theorem prover *Isabelle/HOL* and formally proved its correctness. In the process, we uncovered a correctness-compromising issue in a generalisation of one of the compiler rules in the draft of an updated version of the paper by Bhat *et al.* [BBGR].

1.2 Related work

As mentioned before, this work is based on the work of Bhat *et al.* [BAVG12, BBGR13], which presents a density compiler for probability spaces described by expressions in the language *Fun*. This is a small functional language with basic arithmetic, Boolean logic, product and sum types, conditionals, and a number of built-in discrete and continuous distributions. It does, however, not support lists or recursion. The correctness proof is purely pen and paper, but the authors stated that formalisation in a proof assistant such as Coq is the ultimate goal [BAVG12].

Park *et al.* [PPT05] developed a probabilistic extension of Objective CAML called λ_{\circ} . Their work focuses on *sample generation*, i. e. using an infinite stream of random numbers to compute samples of the distribution described by a probabilistic program. While Bhat *et al.* generate density functions of functional programs, Park *et al.* generate *sampling functions*. These are functions that map $[0; 1]^{\infty}$ to the sample space. The sampling function effectively uses a finite number of uniformly-distributed random numbers from the interval $[0; 1]$ and returns a sample of the desired random variable. This approach allows them to handle much more general distributions, even recursively-defined ones and distributions that do not have a density function, but it does not allow precise reasoning about these distributions (such as determining the exact expectation). No attempt at formal verification is made.

There are already existing approaches to formalise semantics of probabilistic programs: Hurd [Hur02] formalises programs as random variables on infinite streams of random bits. Hurd *et al.* [HMM05] and Cock [Coc12, Coc14] both formalise pGCL, an imperative programming language with probabilistic and non-deterministic choice. Audebaud and Paulin-Mohring [APM06] verify probabilistic functional programs in Coq [BC04] using a shallow embedding based on the Giry monad on discrete probability distributions. All these program semantics – even the framework described by Hurd [Hur02], which is based on Measure Theory – allow only discrete distributions.

1.3 Utilised tools

As stated before, we work with the interactive theorem prover *Isabelle/HOL*. *Isabelle* is a generic proof assistant that supports a number of logical frameworks (object logics) such as *Higher-Order Logic* (HOL) or *Zermelo-Fraenkel set theory* (ZF). The most widely used instance is *Isabelle/HOL*, which is what we use.

We heavily rely on Johannes Hölzl’s *Isabelle* formalisation of measure theory [Hö12], which is already part of the *Isabelle/HOL* library. We also use Sudeep Kanav’s formalisation of the Gaussian distribution and a number of libraries from the proof of the Central Limit Theorem by Avigad *et al.* [AHS14], namely the notion of interval and set integrals and the Fundamental Theorem of Calculus for these integrals. All of these have been or will be moved to the *Isabelle/HOL* library by their respective maintainers as well.

1.4 Outline

In Sect. 2.1, we will explain the notation we will use and then give a brief overview of the mathematical basics we require – in particular the Giry Monad – in Sect. 2.2.

Section 3 contains the definition and the semantics of the source and target language. Section 4 defines the abstract compiler and gives a high-level outline of the soundness proof. Section 5 then explains the refinement of the abstract compiler to the concrete compiler and the final correctness result and evaluates the compiler on a simple example.

Section 6 gives an overview of how much work went into the different parts of the project and what difficulties we encountered. It also lists possible improvements and other future work before summarising the results we obtained. Finally, the appendix contains a table of all notation and auxiliary functions used in this work for reference.

2 Preliminaries

2.1 Notation

In the following, we will explain the conventions and the most important notation we will use. For a full overview of all the non-standard notation used, see the appendix.

2.1.1 Typographical notes

We will use the following typographical conventions in mathematical formulæ:

- Constants, functions, datatype constructors, and types will be typeset in regular roman font: `max`, `return`, π , etc.¹
- Free and bound variables (including type variables) are in italics: x , A , dst , etc.
- Isabelle keywords are in bold print: **lemma**, **datatype**, **primrec**, etc.
- σ -algebras are typeset in calligraphic font: \mathcal{A} , \mathcal{B} , \mathcal{M} , etc.
- Categories are typeset in Fraktur: \mathfrak{Set} , \mathfrak{Grp} , \mathfrak{Meas} , etc.
- File names of Isabelle theories are set in monospaced font: `PDF_Compiler.thy`.

2.1.2 Deviations from standard mathematical notation

In order to maintain coherence with the notation used in Isabelle, we will deviate from standard mathematical notation in the following ways:

- As is convention in functional programming (and therefore in Isabelle), function application is often written as $f\ x$ instead of $f(x)$. It is the operation that binds strongest and it associates to the left, i. e. $f\ x + 1$ is $(f\ x) + 1$ and $f\ x\ y$ is $(f\ x)\ y$.
- The integral over integrand f with variable x and the measure μ is written as

$$\int x. f\ x\ \partial\mu \quad \text{instead of} \quad \int f(x)\ d\mu(x).$$

- The special notion of a *non-negative integral*² is used; this integral ‘ignores’ the negative part of the integrand. Effectively:

$$\int x^+ . f\ x\ \partial\mu \quad \hat{=} \quad \int \max(0, f(x))\ d\mu(x)$$

- Lambda abstractions, such as $\lambda x. x^2$, are used to specify functions. The corresponding standard mathematical notation would be $x \mapsto x^2$.
- Some additional, less important differences in notation are explained in the appendix, Sect. A.1.

¹In formulæ embedded in regular text, we will often set types and constants in italics to distinguish them from the surrounding text.

²This is a very central concept in the measure theory library in Isabelle. We will mostly use it with non-negative functions anyway, so the distinction is purely formal.

2.1.3 Semantics notation

We will always use Γ to denote a type environment, i. e. a function from variable names to types, and σ to denote a state, i. e. a function from variable names to values. Note that variable names are always natural numbers as we use de Bruijn indices. The letter Y (Upsilon) will be used to denote *density contexts*, which are used in the compiler.

We also employ the notation $t \bullet \Gamma$ resp. $v \bullet \sigma$ to denote the insertion of a new variable with the type t (resp. value v) into a typing environment (resp. state). This is used when entering the scope of a bound variable; the newly inserted variable then has index 0 and all other variables are incremented by 1. The precise definition is as follows:³

$$(y \bullet f)(x) = \begin{cases} y & \text{if } x = 0 \\ f(x - 1) & \text{otherwise} \end{cases}$$

We use the same notation for inserting a new variable into a set of variables, shifting all other variables, i. e.:

$$x \bullet V = \{x\} \cup \{y + 1 \mid y \in V\}$$

In general, the notation $\Gamma \vdash e : t$ and variations thereof will always mean ‘The expression e has type t in the type environment Γ ’, whereas $Y \vdash e \Rightarrow f$ and variations thereof mean ‘The expression e compiles to f under the context Y ’. For an overview of these predicates, see the appendix, Sect. A.2.

2.2 Mathematical basics

The category theory part of this section is based mainly on a presentation by Ernst-Erich Doberkat [Dob08]. For a more detailed introduction, see his textbook [Dob07] or the original paper by Michèle Giry [Gir82].

2.2.1 Basic measure spaces

There are two basic measure spaces we will use:

Counting space. The counting space on a countable set A has

- the carrier set A
- the measurable sets $\mathcal{P}(A)$, i. e. all subsets of A
- the measure $\lambda X. |X|$, i. e. the measure of a subset of A is simply the number of elements in that set (which can be ∞)

³ Note the analogy to the notation $x \# xs$ for prepending an element to a list. This is because contexts with de Bruijn indices are generally represented by lists, where the n -th element is the entry for the variable n , and inserting a value for a newly bound variable is then simply the prepending operation.

Borel space. The Borel space is a measure space on the real numbers which has

- the carrier set \mathbb{R}
- the Borel σ -algebra as measurable sets, i. e. the smallest σ -algebra containing all open subsets of \mathbb{R}
- the Borel measure as a measure, i. e. the uniquely defined measure that maps real intervals to their lengths

2.2.2 Sub-probability spaces

A sub-probability space is a measurable space (A, \mathcal{A}) with a measure μ such that every set $X \in \mathcal{A}$ has a measure ≤ 1 , or, equivalently, $\mu(A) \leq 1$.

For technical reasons, we also assume $A \neq \emptyset$. This is required later in order to define the *bind* operation in the Giry monad in a convenient way within Isabelle. This non-emptiness condition will always be trivially satisfied by all the measure spaces used in this work.

2.2.3 The category $\mathcal{M}eas$

Note that:

- For any measurable space (A, \mathcal{A}) , the identity function is \mathcal{A} - \mathcal{A} -measurable.
- For any measurable spaces (A, \mathcal{A}) , (B, \mathcal{B}) , (C, \mathcal{C}) , an \mathcal{A} - \mathcal{B} -measurable function f , and a \mathcal{B} - \mathcal{C} -measurable function g , the function $g \circ f$ is \mathcal{A} - \mathcal{C} -measurable.

Therefore, measurable spaces form a category $\mathcal{M}eas$ where:

- the objects of the category are measurable spaces
- the morphisms of the category are measurable functions
- the identity morphism $\mathbf{1}_{(A, \mathcal{A})}$ is the identity function $\text{id}_A : A \rightarrow A, \lambda x. x$
- morphism composition is function composition

2.2.4 The sub-probability functor

Since our programs will return sub-probability measures, we need to define measurability on sub-probability measures. Given a measurable space (A, \mathcal{A}) , we consider the set of all sub-probability measures on (A, \mathcal{A}) :

$$M := \{\mu \mid \mu \text{ is a measure on } (A, \mathcal{A}), \mu(A) \leq 1\}$$

Now we want to find a suitable σ -algebra \mathcal{M} on M in order to turn M into a measurable space. We will write (M, \mathcal{M}) as $\mathcal{S}(A, \mathcal{A})$.

A natural property that this space should satisfy is that measuring a fixed set $X \in \mathcal{A}$ while varying the measure within $\mathbb{S}(A, \mathcal{A})$ should be a Borel-measurable function; formally:

For all $X \in \mathcal{A}$, $f: \mathbb{S}(A, \mathcal{A}) \rightarrow \mathbb{R}$, $\lambda\mu. \mu(X)$ is $\mathbb{S}(A, \mathcal{A})$ -Borel-measurable

We can now simply define $\mathbb{S}(A, \mathcal{A})$ as the smallest measurable space with the carrier set M that fulfils this property, i. e.

$$\mathcal{M} := \{(\lambda\mu. \mu(X))^{-1}(Y) \mid X \in \mathcal{A}, Y \in \mathcal{B}\}$$

where \mathcal{B} is the Borel σ -algebra on \mathbb{R} .

Additionally, for a measurable function f , we define $\mathbb{S}(f) = \lambda\mu. f(\mu)$, where $f(\mu)$ denotes the push-forward measure (or image measure)⁴. Then \mathbb{S} maps objects of $\mathcal{M}\text{eas}$ to objects of $\mathcal{M}\text{eas}$ and morphisms of $\mathcal{M}\text{eas}$ to morphisms of $\mathcal{M}\text{eas}$. We can thus see that \mathbb{S} is an endofunctor in the category $\mathcal{M}\text{eas}$, as $(\text{id}_{(A, \mathcal{A})})(\mu) = \mu$ and $(f \circ g)(\mu) = f(g(\mu))$. We call \mathbb{S} the *sub-probability functor*.

2.2.5 Giry monad

The Giry monad naturally captures the notion of choosing a value according to a (sub-)probability distribution, using it as a parameter for another distribution, and observing the result.

Consequently, *return* (or η) yields a Dirac measure, i. e. a probability measure in which all the ‘probability’ lies in a single element, and *bind* (or $\gg=$) integrates over all the input values to compute one single output measure. Formally, for measurable spaces (A, \mathcal{A}) and (B, \mathcal{B}) , a measure μ on (A, \mathcal{A}) , a value $x \in A$, and an \mathcal{A} - $\mathbb{S}(B, \mathcal{B})$ -measurable function f :

$$\text{return}_{(A, \mathcal{A})} x := \lambda X. \begin{cases} 1 & \text{if } x \in X \\ 0 & \text{otherwise} \end{cases} \quad \mu \gg= f := \lambda X. \int x. f(x)(X) \partial\mu$$

Unfortunately, restrictions due to Isabelle’s type system require us to determine the σ -algebra of the resulting measurable space for *bind* $M f$ since this information cannot be provided by the type of f . This can be done with an additional parameter, but it is more convenient to define *bind* in such a way that it chooses an arbitrary value $x \in M$ and takes the σ -algebra of $f x$ (or the count space on \emptyset if M is empty)⁵.

This choice is somewhat non-standard, but the difference is of no practical significance as we will not use *bind* on empty measure spaces.

To simplify the proofs for *bind*, we instead define the *join* operation (also known as μ in category theory) first and use it to then define *bind*. The *join* operation ‘flattens’ objects, i. e. it maps an element of $\mathbb{S}(\mathbb{S}(A, \mathcal{A}))$ to one of $\mathbb{S}(A, \mathcal{A})$.

⁴ $f(\mu) = \lambda X. \mu(f^{-1}(X))$.

⁵Note that for any \mathcal{A} - $\mathbb{S}(B, \mathcal{B})$ -measurable function f , the σ -algebra thus obtained is independent of which value is chosen.

Such an operation can be naturally defined as:

$$\text{join } \mu = \lambda X. \int \mu'. \mu'(X) \partial \mu$$

Note that in Isabelle, *join* has an additional explicit parameter for the measurable space of the result to avoid the problem we had with *bind*. This makes expressions containing *join* more complicated; this is, however, justified by the easier proofs and will be unproblematic later since we will never use *join* directly, only *bind*.

Now, *bind* can be defined using *join* in the following way, modulo handling of empty measure spaces⁶:

$$\mu \gg= f := \text{join } (\mathbb{S}(f)(\mu)) = \text{join } (f(\mu))$$

The coherence conditions

$$\begin{aligned} \text{join} \circ \mathbb{S}(\text{join}) &= \text{join} \circ \text{join} \\ \text{join} \circ \mathbb{S}(\text{return}_{(A,A)}) &= \text{join} \circ \text{return}_{\mathbb{S}(A,A)} = \mathbf{1}_{\mathbb{S}(A,A)} \\ \text{join} \circ \mathbb{S}(\mathbb{S}(f)) &= \mathbb{S}(f) \circ \text{join} \end{aligned}$$

can be proven easily by unfolding the definitions of the operations and applying the rules for integrals on image measures.

The ‘do’ syntax. For better readability, we employ a Haskell-style ‘do notation’ for operations in the Girly monad. The syntax of this notation is defined recursively, where M stands for a monadic expression and $\langle \text{pattern} \rangle$ stands for arbitrary ‘raw’ text:

$$\mathbf{do} \{M\} \hat{=} M \qquad \mathbf{do} \{x \leftarrow M; \langle \text{pattern} \rangle\} \hat{=} M \gg= (\lambda x. \mathbf{do} \{\langle \text{pattern} \rangle\})$$

Example:

$$\mathbf{do} \{x \leftarrow M; y \leftarrow f x; g x y\} \hat{=} M \gg= (\lambda x. f x \gg= (\lambda y. g x y))$$

⁶ $\mathbb{S}(f)$, i. e. lifting a function on values to a function on measure spaces, is done by the *distr* function in Isabelle. This function was already defined in the library, which simplifies our proofs about *bind*, seeing as some of the proof work has already been done in the proofs about *distr*.

3 Language Syntax and Semantics

The source language used in the formalisation was modelled after the language *Fun* described by Bhat *et al.* [BBGR13]; similarly, the target language is almost identical to the target language used by Bhat *et al.* However, we have made the following changes in our languages:

- Variables are represented by de Bruijn indices to avoid handling freshness, capture-avoiding substitution, and related problems.
- No sum types are supported. Consequently, the **match** ... **with** ... command is replaced with an **IF** ... **THEN** ... **ELSE** ... command. Furthermore, booleans are a primitive type rather than represented as $unit + unit$.
- The type *double* is replaced with the type *real* and it represents a real number with absolute precision as opposed to an IEEE 754 floating point number.
- Beta and Gamma distributions are not included.

In the following sections, we give the precise syntax, typing rules, and semantics of both our source language and our target language.

3.1 Types, values, and operators

The source language and the target language share the same type system and the same operators. Figure 2 shows the types, values, and operators that exist in our languages. Figure 1 defines the semantics of the operators.

All operators are *total*, meaning that for every input value of their parameter type, they return a single value of their result type. This requires some non-standard definitions for inherently non-total operations such as division, the logarithm, and the square root. Non-totality could also be handled by implementing operators in the Giry monad and letting them return either a Dirac distribution with a single result or, when evaluated for a parameter on which they are not defined, the null measure. This, however, would probably complicate many proofs significantly.

To increase readability, we will use the abbreviations shown in the operator semantics table (Fig. 1) and the following additional ones:

- $t_1 \times t_2$ stands for *PRODUCT* $t_1 t_2$.
- *TRUE* and *FALSE* stand for *BoolVal True* and *BoolVal False*, respectively.
- *RealVal*, *IntVal*, etc. will be omitted in expressions when their presence is implicitly clear from the context.
- $a - b$ stands for $a + (-b)$ and a/b for $a \cdot b^{-1}$.

OPERATOR	NOTATION	INPUT TYPE	OUTPUT TYPE	SEMANTICS
Add	$a + b$	INTEG \times INTEG REAL \times REAL	INTEG REAL	$a + b$
Minus	$-a$	INTEG REAL	INTEG REAL	$-a$
Mult	$a \cdot b$	INTEG \times INTEG REAL \times REAL	INTEG REAL	$a \cdot b$
Inverse	a^{-1}	REAL	REAL	$\begin{cases} \frac{1}{a} & \text{for } a \neq 0 \\ 0 & \text{otherwise} \end{cases}$
Sqrt	\sqrt{a}	REAL	REAL	$\begin{cases} \sqrt{a} & \text{for } a \geq 0 \\ 0 & \text{otherwise} \end{cases}$
Exp	$\exp a$	REAL	REAL	e^a
Ln	$\ln a$	REAL	REAL	$\begin{cases} \ln a & \text{for } a > 0 \\ 0 & \text{otherwise} \end{cases}$
Pi	π	UNIT	REAL	π
Pow	a^b	INTEG \times INTEG REAL \times INTEG	INTEG REAL	$\begin{cases} a^b & \text{for } a \neq 0, b \geq 0 \\ 0 & \text{otherwise} \end{cases}$
Fact	$a!$	INTEG	INTEG	$a!$
And	$a \wedge b$	BOOL \times BOOL	BOOL	$a \wedge b$
Or	$a \vee b$	BOOL \times BOOL	BOOL	$a \vee b$
Not	$\neg a$	BOOL	BOOL	$\neg a$
Equals	$a = b$	$t \times t$	BOOL	$a = b$
Less	$a < b$	INTEG \times INTEG REAL \times REAL	BOOL BOOL	$a < b$
Fst	$\text{fst } z$	$t_1 \times t_2$	t_1	a for $z = (a, b)$
Snd	$\text{snd } z$	$t_1 \times t_2$	t_2	b for $z = (a, b)$
Cast REAL	$\langle a \rangle$	BOOL	REAL	$\langle a = \text{TRUE} \rangle$
	implicit	INTEG	REAL	a as a real number
Cast INTEG		BOOL	INTEG	$\langle a = \text{TRUE} \rangle$
		REAL	INTEG	$\lfloor a \rfloor$

Figure 1: The types, abbreviations and semantics for the operators

```

datatype pdf_type =
    UNIT | BOOL | INTEG | REAL | PRODUCT pdf_type pdf_type
datatype val =
    UnitVal | BoolVal bool | IntVal int | RealVal real | <|val, val|>
datatype pdf_operator =
    Fst | Snd | Add | Mult | Minus | Less | Equals | And | Or | Not | Pow |
    Fact | Sqrt | Exp | Ln | Inverse | Pi | Cast pdf_type
    
```

Figure 2: The types and values in the source and target language

3.2 Auxiliary definitions

A number of auxiliary definitions are used in the definition of the semantics; for a full list of auxiliary functions see the appendix. The following two notions require a detailed explanation:

3.2.1 Measure embeddings

A *measure embedding* is the measure space obtained by ‘tagging’ values in a measure space M with some injective function f (in fact, f will always be a datatype constructor). For instance, a set of values of type *REAL* can naturally be measured by stripping away the *RealVal* constructor and using a measure on real numbers (e.g. the Borel measure) on the resulting set of reals. Formally:

$$\text{embed_measure } (A, \mathcal{A}, \mu) f = \left(f(A), \{f(X) \mid X \in \mathcal{A}\}, \lambda X. \mu(f^{-1}(X) \cap A) \right)$$

3.2.2 Stock measures

The *stock measure* for a type t is the ‘natural’ measure on values of that type. This is defined as follows:

- For the countable types *UNIT*, *BOOL*, and *INTEG*: the counting measure over the corresponding type universes
- For *REAL*: the embedding of the Borel measure on \mathbb{R} with *RealVal*
- For $t_1 \times t_2$: the embedding of the product measure

$$\text{stock_measure } t_1 \otimes \text{stock_measure } t_2$$

with $\lambda(v, w). \langle |v, w| \rangle$

Note that in order to save space and increase readability, we will often write $\int x. f \ x \ \partial t$ instead of $\int x. f \ x \ \partial \text{stock_measure } t$ in integrals.

The state measure. Using the stock measure, we can also construct a measure on *states* in the context of a typing environment Γ . A state on the variables V is a function that maps a variable in V to a value. A state σ is *well-formed* w. r. t. to V and Γ if it maps every variable $x \in V$ to a value of type Γx and every variable $\notin V$ to *undefined*.

We now fix Γ and a finite V and consider the set of well-formed states w. r. t. V and Γ . Another representation of these states are tuples in which the i -th component is the value of the i -th variable in V . The natural measure that can be given to such tuples is then the finite product measure of the stock measures of the types of the variables:

$$\text{state_measure } \Gamma V := \bigotimes_{x \in V} \text{stock_measure } (\Gamma x)$$

3.3 Source language

Figures 3 and 4 show the syntax resp. the typing rules of the source language. Figure 6 defines the source language semantics as a primitively recursive function. Similarly to the abbreviations mentioned in Sect. 3.1, we will omit *Val* when its presence is implicitly obvious from the context; e. g. if in some context, e is an expression and c is a constant real number, we will write $e + \text{Val } (\text{RealVal } c)$ as $e + c$.

Figure 5 shows the built-in distributions of the source language, their parameter type and domain, the type of the random variable they describe, and their density functions in terms of their parameter. When given a parameter outside their domain, they return the null measure.

```

datatype expr =
  Var nat | Val val | LET expr IN expr | pdf_operator $ expr | <expr, expr> |
  Random pdf_dist | IF expr THEN expr ELSE expr | Fail pdf_type

```

Figure 3: The source language syntax

$\frac{\text{ET_VAL}}{\Gamma \vdash \text{Val } v : \text{val_type } v}$	$\frac{\text{ET_VAR}}{\Gamma \vdash \text{Var } x : \Gamma x}$	$\frac{\text{ET_FAIL}}{\Gamma \vdash \text{Fail } t : t}$
$\frac{\text{ET_OP}}{\Gamma \vdash e : t \quad \text{op_type } op \ t = \text{Some } t'}{\Gamma \vdash op \$ e : t'}$	$\frac{\text{ET_PAIR}}{\Gamma \vdash e_1 : t_1 \quad \Gamma \vdash e_2 : t_2}{\Gamma \vdash \langle e_1, e_2 \rangle : t_1 \times t_2}$	
$\frac{\text{ET_IF}}{\Gamma \vdash b : \text{BOOL} \quad \Gamma \vdash e_1 : t \quad \Gamma \vdash e_2 : t}{\Gamma \vdash \text{IF } b \text{ THEN } e_1 \text{ ELSE } e_2 : t}$	$\frac{\text{ET_LET}}{\Gamma \vdash e_1 : t_1 \quad t_1 \bullet \Gamma \vdash e_2 : t_2}{\Gamma \vdash \text{LET } e_1 \text{ IN } e_2 : t_2}$	
$\frac{\text{ET_RAND}}{\Gamma \vdash e : \text{dist_param_type } dst}{\Gamma \vdash \text{Random } dst \ e : \text{dist_result_type } dst}$		

Figure 4: The typing rules for source-language expressions

DISTRIBUTION	PARAMETER TYPE	DOMAIN	DISTR. TYPE	DENSITY FUNCTION
Bernoulli	REAL	$p \in [0;1]$	BOOL	$\begin{cases} p & \text{for } x = \text{TRUE} \\ 1 - p & \text{for } x = \text{FALSE} \end{cases}$
UniformInt	INTEG \times INTEG	$p_1 \leq p_2$	INTEG	$\frac{\langle x \in [p_1; p_2] \rangle}{p_2 - p_1 + 1}$
UniformReal	REAL \times REAL	$p_1 < p_2$	REAL	$\frac{\langle x \in [p_1; p_2] \rangle}{p_2 - p_1}$
Gaussian	REAL \times REAL	$p_2 > 0$	REAL	$\frac{1}{\sqrt{2\pi p_2^2}} \exp\left(-\frac{(x - p_1)^2}{2p_2^2}\right)$
Poisson	REAL	$p \geq 0$	INTEG	$\begin{cases} \exp(-p) \cdot p^x / x! & \text{for } x \geq 0 \\ 0 & \text{otherwise} \end{cases}$

Figure 5: The built-in distributions of the source language.

The density functions are given in terms of the parameter p , which is of the type given in the column ‘parameter type’. If p is of a product type, p_1 and p_2 stand for the two components of p . The variable x stands for the parameter of the density function, e.g. the point at which it is evaluated.

```

primrec expr_sem :: state  $\Rightarrow$  expr  $\Rightarrow$  val measure where
  expr_sem  $\sigma$  (Val  $v$ ) = return_val  $v$ 
| expr_sem  $\sigma$  (Var  $x$ ) = return_val ( $\sigma$   $x$ )
| expr_sem  $\sigma$  (LET  $e_1$  IN  $e_2$ ) =
  do {
     $v \leftarrow$  expr_sem  $\sigma$   $e_1$ ;
    expr_sem ( $v \bullet \sigma$ )  $e_2$ 
  }
| expr_sem  $\sigma$  (op $  $e$ ) =
  do {
     $v \leftarrow$  expr_sem  $\sigma$   $e$ ;
    return_val (op_sem op  $v$ )
  }
| expr_sem  $\sigma$   $\langle e_1, e_2 \rangle$  =
  do {
     $v \leftarrow$  expr_sem  $\sigma$   $e_1$ ;
     $w \leftarrow$  expr_sem  $\sigma$   $e_2$ ;
    return_val  $\langle |v, w| \rangle$ 
  }
| expr_sem  $\sigma$  (IF  $b$  THEN  $e_1$  ELSE  $e_2$ ) =
  do {
     $b' \leftarrow$  expr_sem  $\sigma$   $b$ ;
    if  $b' = \text{TRUE}$  then expr_sem  $\sigma$   $e_1$  else expr_sem  $\sigma$   $e_2$ 
  }
| expr_sem  $\sigma$  (Random  $dst$   $e$ ) =
  do {
     $p \leftarrow$  expr_sem  $\sigma$   $e$ ;
    dist_measure  $dst$   $p$ 
  }
| expr_sem  $\sigma$  (Fail  $t$ ) = null_measure (stock_measure  $t$ )

```

Figure 6: The semantics of source-language expressions

3.4 Deterministic expressions

We call an expression e *deterministic* (written as ‘ e det’) if it contains no occurrence of *Random* or *Fail*. Such expressions are of particular interest: if all their free variables have a fixed value, they return precisely one value, so we can define a function expr_sem_rf ⁷ that, when given a state σ and a deterministic expression e , returns this single value.

The definition is obvious and leads to the following equality (assuming that e is deterministic and well-typed and σ is a valid state):

$$\text{expr_sem } \sigma e = \text{return (expr_sem_rf } \sigma e)$$

This property will later enable us to also convert deterministic source-language expressions into ‘equivalent’ target-language expressions.

3.5 Target language

The target language is again modelled closely after the one used by Bhat *et al.* [BBGR13]. The type system and the operators are the same as in the source language, but the key difference is that while expressions in the source language return a measure space on their result type, the expressions in the target language always return a single value.

Since our source language lacks sum types, so does our target language. Additionally, our target language differs from that used by Bhat *et al.* in the following respects:

- Our language has no function types; since functions only occur as integrands and as final results (as the compilation result is a density function), we can simply define integration to introduce the integration variable as a bound variable and let the final result contain a single free variable with de Bruijn index 0, i. e. there is an implicit λ abstraction around the compilation result.
- Evaluation of expressions in our target language can never fail. In the language by Bhat *et al.*, failure is used to handle undefined integrals; we, on the other hand, use the convention of Isabelle’s measure theory library, which returns `o` for integrals of non-integrable functions. This has the advantage of keeping the semantics simple, which makes proofs considerably easier. To illustrate this, note that for integration in a semantics with failure, the evaluation of the integrand only has to succeed *almost everywhere*, i. e. it can fail on a null set. Such semantics would be awkward to define and difficult to use.
- Our target language does not have ‘let’ bindings, since, in contrast to the source language, they would be semantically superfluous here. However, they would still be useful in practice since they yield shorter expressions and can avoid multiple evaluation of the same term; they could be added with little effort.

⁷In Isabelle, the term *randomfree* is used instead of *deterministic*, hence the ‘rf’ suffix. This is in order to emphasise the syntactical nature of the property. Additionally, it is worth noting that a syntactically deterministic expression is not truly deterministic if the variables it contains are randomised over, which is the case sometimes.

Figures 7, 8, and 9 show the syntax, typing rules, and semantics of the target language.

```

datatype cexpr =
  CVar nat | CVal val | pdf_operator $c cexpr | <cexpr, cexpr>c |
  IFc cexpr THEN cexpr ELSE cexpr | ∫c cexpr ∂pdf_type
    
```

Figure 7: The expressions of the target language

$\frac{\text{CET_VAL}}{\Gamma \vdash_c \text{CVal } v : \text{val_type } v}$	$\frac{\text{CET_VAR}}{\Gamma \vdash_c \text{CVar } x : \Gamma x}$
$\frac{\text{CET_OP} \quad \Gamma \vdash_c e : t \quad \text{op_type } op \ t = \text{Some } t'}{\Gamma \vdash_c op \$_c e : t'}$	$\frac{\text{CET_PAIR} \quad \Gamma \vdash_c e_1 : t_1 \quad \Gamma \vdash_c e_2 : t_2}{\Gamma \vdash_c \langle e_1, e_2 \rangle_c : t_1 \times t_2}$
$\frac{\text{CET_IF} \quad \Gamma \vdash_c b : \text{BOOL} \quad \Gamma \vdash_c e_1 : t \quad \Gamma \vdash_c e_2 : t}{\Gamma \vdash_c \text{IF}_c b \text{ THEN } e_1 \text{ ELSE } e_2 : t}$	$\frac{\text{CET_INT} \quad t \bullet \Gamma \vdash_c e : \text{REAL}}{\Gamma \vdash_c \int_c e \partial t : \text{REAL}}$

Figure 8: The typing rules for the target language

```

primrec cexpr_sem :: state ⇒ cexpr ⇒ val where
  cexpr_sem σ (CVal v) = v
| cexpr_sem σ (CVar x) = σ x
| cexpr_sem σ <e1, e2>c = <|cexpr_sem σ e1, cexpr_sem σ e2|>
| cexpr_sem σ (op $c e) = op_sem op (cexpr_sem σ e)
| cexpr_sem σ (IFc b THEN e1 ELSE e2) =
  (if cexpr_sem σ b = TRUE then cexpr_sem σ e1 else cexpr_sem σ e2)
| cexpr_sem σ (∫c e ∂t) =
  RealVal (∫x. extract_real (cexpr_sem (x • σ) e) ∂stock_measure t)
    
```

Figure 9: The semantics of the target language

Auxiliary functions. In our compiler, we will often have to modify target-language expressions and combine them to larger ones. This requires the definition of a number of auxiliary functions e. g. to modify the variables in an expression, substitute an expression for a variable, etc. Using de Bruijn indices saves us the trouble of having to handle variable capture, but it still requires some care when combining expressions: when entering the scope of an integral, a new bound variable is introduced and all other variables are incremented. Therefore, when substituting any expression into the scope of an integral, all its variables must be shifted by 1.

Moreover, as mentioned above, our target language does not contain the notion of a *function*, i. e. no λ abstractions or function applications, and we emulate functions by representing them as expressions with a single free variable with de Bruijn index 0. The drawback of this is that the inability to use abstraction, application, and β reduction requires us to introduce additional auxiliary functions in order to manipulate such expressions, e. g. to compose two functions. For a full list of these functions and other auxiliary functions related to the target language, see Sect. A.3 in the appendix.

Converting deterministic expressions. The auxiliary function *expr_rf_to_cexpr*, which will be used in some rules of the compiler that handle deterministic expressions, is of particular interest. We mentioned earlier that deterministic source-language expressions can be converted to equivalent target-language expressions.⁸ This function does precisely that. Its definition is mostly obvious, apart from the LET case: since our target language does not have a LET construct, the function must expand LET bindings in the source-language expression by substituting in the bound expression, which is done with the auxiliary function *cexpr_subst*.

expr_rf_to_cexpr satisfies the following equality for any deterministic source-language expression *e*:

$$\text{cexpr_sem } \sigma (\text{expr_rf_to_cexpr } e) = \text{expr_sem_rf } \sigma e$$

⁸Bhat *et al.* say that a deterministic expression ‘is also an expression in the target language syntax, and we silently treat it as such’ [BBGR13]

4 Abstract compiler

4.1 Density contexts

First, we define the notion of a *density context*, which holds the accumulated context data the compiler will require to compute the density of an expression. A density context is a tuple $Y = (V, V', \Gamma, \delta)$ that contains the following information:

- The set V of random variables in the current context. These are the variables that are randomised over.
- The set V' of parameter variables in the current context. These are variables that may occur in the expression, but are not randomised over. They are treated as constants.
- The type environment Γ
- A density function δ that returns the common density of the variables V under the parameters V' . Here, δ is a function from $space (state_measure (V \cup V') \Gamma)$ to the extended real numbers.

A density context (V, V', Γ, δ) describes a parametrised measure on the states on V : let $\rho \in space (state_measure V' \Gamma)$ be a parameter state. Then we write

$$dens_ctxt_measure (V, V', \Gamma, \delta) \rho$$

for the measure that we obtain by taking $state_measure V \Gamma$, transforming it by merging a given state σ with the parameter state ρ and finally applying the density δ on the resulting image measure. The Isabelle definition of this is:

definition `dens_ctxt_measure` :: `dens_ctxt` \Rightarrow `state` \Rightarrow `state_measure` **where**

$$dens_ctxt_measure (V, V', \Gamma, \delta) \rho = \\ density (distr (state_measure V \Gamma) (state_measure (V \cup V') \Gamma) (\lambda\sigma. merge V V' (\sigma, \rho))) \delta$$

Informally, *dens_ctxt_measure* describes the measure obtained by integrating over the variables $v_1, \dots, v_m \in V$ while treating the variables $v'_1, \dots, v'_n \in V'$ as parameters. The evaluation of an expression e with variables from $V \cup V'$ in this context is effectively a function

$$\lambda v'_1 \dots v'_n. \int_{v_1}^+ \dots \int_{v_m}^+ \\ expr_sem (v_1, \dots, v_m, v'_1, \dots, v'_n) e \cdot \delta (v_1, \dots, v_m, v'_1, \dots, v'_n) \partial\Gamma v_1 \dots \partial\Gamma v_m .$$

A density context is *well-formed* (implemented in Isabelle as the locale *density_context*) if:

- V and V' are finite and disjoint
- $\delta \sigma \geq 0$ for any $\sigma \in space (state_measure (V \cup V') \Gamma)$
- δ is Borel-measurable w. r. t. $state_measure (V \cup V') \Gamma$
- for any $\rho \in space (state_measure V' \Gamma)$, the measure $dens_ctxt_measure (V, V', \Gamma, \delta) \rho$ is a sub-probability measure

4.2 Definition

As a first step, we have implemented an abstract density compiler as an inductive predicate $Y \vdash_d e \Rightarrow f$, where Y is a density context, e is a source-language expression and f is a function of type $val\ state \Rightarrow val \Rightarrow ereal$. Its first parameter is a state that assigns values to the free variables in e and its second parameter is the value for which the density is to be computed. The compiler therefore computes a density function that is parametrised with the values of the non-random free variables in the source expression.

The compilation rules are virtually identical to those by Bhat *et al.* [BBGR13], except for the following adaptations:

- Bhat *et al.* handle the **IF ... THEN ... ELSE** case with the ‘match’ rule for sum types. As we do not support sum types, we have a dedicated rule for this.
- The use of de Bruijn indices requires shifting of variable sets and states whenever the scope of a new bound variable is entered; unfortunately, this makes some rules somewhat technical.
- We do not provide any compiler support for *deterministic* ‘let’ bindings. They are semantically redundant, as they can always be expanded without changing the semantics of the expression. In fact, they *have* to be unfolded for compilation, so they can be regarded as a feature that adds convenience, but no expressivity.

The following list shows the standard compilation rules adapted from Bhat *et al.* Note that the functions *marg_dens* and *marg_dens2* compute the marginal density of one (resp. two) variables by ‘integrating away’ all the other variables from the common density δ . The function *branch_prob* computes the probability of being in the current branch of execution by integrating over *all* the variables in the common density δ .

HD_VAL	$\frac{\text{countable_type } (\text{val_type } v)}{(V, V', \Gamma, \delta) \vdash_d \text{Val } v \Rightarrow \lambda \rho x. \text{branch_prob } (V, V', \Gamma, \delta) \rho \cdot \langle x = v \rangle}$
HD_VAR	$\frac{x \in V}{(V, V', \Gamma, \delta) \vdash_d \text{Var } x \Rightarrow \text{marg_dens } (V, V', \Gamma, \delta) x}$
HD_PAIR	$\frac{x \in V \quad y \in V \quad x \neq y}{(V, V', \Gamma, \delta) \vdash_d \langle \text{Var } x, \text{Var } y \rangle \Rightarrow \text{marg_dens2 } (V, V', \Gamma, \delta) x y}$
HD_FAIL	$\frac{}{(V, V', \Gamma, \delta) \vdash_d \text{Fail } t \Rightarrow \lambda \rho x. 0}$
HD_LET	$\frac{(\emptyset, V \cup V', \Gamma, \lambda x. 1) \vdash_d e_1 \Rightarrow f}{(0 \bullet V, \{x + 1 \mid x \in V'\}, \text{type_of } \Gamma e_1 \bullet \Gamma, \lambda \rho. (f \cdot \delta)(\rho \circ (\lambda x. x + 1))) \vdash_d e_2 \Rightarrow g}$ $(V, V', \Gamma, \delta) \vdash_d \text{LET } e_1 \text{ IN } e_2 \Rightarrow \lambda \rho. g (\text{undefined} \bullet \rho)$

HD_RANDOM

$$\frac{(V, V', \Gamma, \delta) \vdash_d e \Rightarrow f}{(V, V', \Gamma, \delta) \vdash_d \text{Random } dst \ e \Rightarrow \lambda \rho y. \int^+ x. f \ \rho \ x \cdot \text{dist_dens } dst \ x \ y \ \partial \text{dist_param_type } dst}$$

HD_RANDOM_DET

$$\frac{e \text{ det} \quad \text{free_vars } e \subseteq V'}{(V, V', \Gamma, \delta) \vdash_d \text{Random } dst \ e \Rightarrow \lambda \rho x. \text{branch_prob } (V, V', \Gamma, \delta) \ \rho \cdot \text{dist_dens } dst \ (\text{expr_sem_rf } \rho \ e) \ x}$$

HD_IF

$$\frac{(\emptyset, V \cup V', \Gamma, \lambda \rho. 1) \vdash_d b \Rightarrow f \quad (V, V', \Gamma, \lambda \rho. \delta \ \rho \cdot f \ \text{TRUE}) \vdash_d e_1 \Rightarrow g_1 \quad (V, V', \Gamma, \lambda \rho. \delta \ \rho \cdot f \ \text{FALSE}) \vdash_d e_2 \Rightarrow g_2}{(V, V', \Gamma, \delta) \vdash_d \text{IF } b \ \text{THEN } e_1 \ \text{ELSE } e_2 \Rightarrow \lambda \rho x. g_1 \ \rho \ x + g_2 \ \rho \ x}$$

HD_IF_DET

$$\frac{b \text{ det} \quad (V, V', \Gamma, \lambda \rho. \delta \ \rho \cdot \langle \text{expr_sem_rf } \rho \ b = \text{TRUE} \rangle) \vdash_d e_1 \Rightarrow g_1 \quad (V, V', \Gamma, \lambda \rho. \delta \ \rho \cdot \langle \text{expr_sem_rf } \rho \ b = \text{FALSE} \rangle) \vdash_d e_2 \Rightarrow g_2}{(V, V', \Gamma, \delta) \vdash_d \text{IF } b \ \text{THEN } e_1 \ \text{ELSE } e_2 \Rightarrow \lambda \rho x. g_1 \ \rho \ x + g_2 \ \rho \ x}$$

HD_FST

$$\frac{(V, V', \Gamma, \delta) \vdash_d e \Rightarrow f}{(V, V', \Gamma, \delta) \vdash_d \text{fst } e \Rightarrow \lambda \rho x. \int^+ y. f \ \rho \ \langle |x, y| \rangle \ \partial \text{type_of } \Gamma \ (\text{snd } e)}$$

HD_SND

$$\frac{(V, V', \Gamma, \delta) \vdash_d e \Rightarrow f}{(V, V', \Gamma, \delta) \vdash_d \text{snd } e \Rightarrow \lambda \rho y. \int^+ x. f \ \rho \ \langle |x, y| \rangle \ \partial \text{type_of } \Gamma \ (\text{fst } e)}$$

HD_OP_DISCR

$$\frac{\text{countable_type } (\text{type_of } \Gamma \ (\text{op } \$ e)) \quad (V, V', \Gamma, \delta) \vdash_d e \Rightarrow f}{(V, V', \Gamma, \delta) \vdash_d \text{op } \$ e \Rightarrow \lambda \rho y. \int^+ x. \langle \text{op_sem } \text{op } x = y \rangle \cdot f \ \rho \ x \ \partial \text{type_of } \Gamma \ e}$$

HD_NEG

$$\frac{(V, V', \Gamma, \delta) \vdash_d e \Rightarrow f}{(V, V', \Gamma, \delta) \vdash_d -e \Rightarrow \lambda \rho x. f \ \rho \ (-x)}$$

HD_ADDC

$$\frac{e' \text{ det} \quad \text{free_vars } e' \subseteq V' \quad (V, V', \Gamma, \delta) \vdash_d e \Rightarrow f}{(V, V', \Gamma, \delta) \vdash_d e + e' \Rightarrow \lambda \rho x. f \ \rho \ (x - \text{expr_sem_rf } \rho \ e')}$$

HD_MULTC

$$\frac{c \neq 0 \quad (V, V', \Gamma, \delta) \vdash_d e \Rightarrow f}{(V, V', \Gamma, \delta) \vdash_d e \cdot \text{Val } (\text{RealVal } c) \Rightarrow \lambda \rho x. f \ \rho \ (x/c) / |c|}$$

HD_ADD	$\frac{(V, V', \Gamma, \delta) \vdash_{\text{d}} \langle e_1, e_2 \rangle \Rightarrow f}{(V, V', \Gamma, \delta) \vdash_{\text{d}} e_1 + e_2 \Rightarrow \lambda \rho z. \int^+ x. f \rho \langle x, z - x \rangle \partial \text{type_of } \Gamma e_1}$
HD_INV	$\frac{(V, V', \Gamma, \delta) \vdash_{\text{d}} e \Rightarrow f}{(V, V', \Gamma, \delta) \vdash_{\text{d}} e^{-1} \Rightarrow \lambda \rho x. f \rho (x^{-1}) / x^2}$
HD_EXP	$\frac{(V, V', \Gamma, \delta) \vdash_{\text{d}} e \Rightarrow f}{(V, V', \Gamma, \delta) \vdash_{\text{d}} \text{exp } e \Rightarrow \lambda \rho x. \text{if } x > 0 \text{ then } f \rho (\ln x) / x \text{ else } 0}$

Figure 10: The abstract compilation rules

Congruence rules. There is a slight technical problem, namely that functions in Isabelle must be total. Since the returned density functions have type $\text{val state} \Rightarrow \text{val} \Rightarrow \text{ereal}$, they must also be defined for values outside their domain. E. g. a density function for type REAL must also return some density for an input value of type BOOL. In the above formalisation, this will be some unspecified value; the density computed by the concrete compiler may return some other value outside its domain. The same is true for the common density δ in the density context. Additionally, the concrete density function may deviate from the abstract one on a null set, i. e. it is enough if they are equal almost everywhere.

To account for this, we have three additional rules, which are too technical to be printed here in full. Instead, we will give a brief informal description of what they do:

HD_AE states that if $Y \vdash_{\text{d}} e \Rightarrow f$ and f' is measurable and for any parameter state ρ , the function $f' \rho$ is non-negative and equal to $f \rho$ almost everywhere, then $Y \vdash_{\text{d}} e \Rightarrow f'$ also holds.

HD_DENS_CTXT_CONG states that if $(V, V', \Gamma, \delta) \vdash_{\text{d}} e \Rightarrow f$ and δ is equal to some δ' for all states on $V \cup V'$, then $(V, V', \Gamma, \delta') \vdash_{\text{d}} e \Rightarrow f$ also holds.

HD_CONG is derived from HD_AE; it assumes that f and f' are equal on their entire domain and in turn removes the measurability and non-negativity assumptions.

4.3 Soundness proof

We show the following soundness result for the abstract compiler: ⁹

lemma `expr_has_density_sound` :

assumes $(\emptyset, \emptyset, \Gamma, \lambda\rho. 1) \vdash_d e \Rightarrow f$ **and** $\Gamma \vdash e : t$ **and** $\text{free_vars } e = \emptyset$
shows $\text{has_subprob_density } (\text{expr_sem } \sigma e) (\text{stock_measure } t) (f (\lambda x. \text{undefined}))$

Here, $\text{has_subprob_density } M N f$ is an abbreviation for the following facts:

- f is a non-negative, N -Borel-measurable function
- applying the density f to N yields the sub-probability measure M

We prove this with the following generalised auxiliary lemma:

lemma `expr_has_density_sound_aux` :

assumes $(V, V', \Gamma, \delta) \vdash_d e \Rightarrow f$ **and** $\Gamma \vdash e : t$ **and** $\text{free_vars } e = \emptyset$ **and**
 $\text{density_context } V V' \Gamma \delta$ **and** $\text{free_vars } e \subseteq V \cup V'$
shows $\text{has_parametrized_subprob_density } (\text{state_measure } V' \Gamma)$
 $(\lambda\rho. \text{do } \{\sigma \leftarrow \text{dens_ctxt_measure } (V, V', \Gamma, \delta) \rho; \text{expr_sem } \sigma e\})$
 $(\text{stock_measure } t) f$

The predicate $\text{has_parametrized_subprob_density } R M N f$ simply means that f is Borel-measurable w.r.t. $R \otimes N$ and the predicate $\text{has_subprob_density } M N f$ holds for any parameter state ρ from R .

The proof is by straightforward induction following the inductive definition of the abstract compiler. In many cases, the monad laws for the Giry monad allow restructuring the induction goal in such a way that the induction hypothesis can be applied directly; in the other cases, the definitions of the monadic operations need to be unfolded and the goal is essentially to show that two integrals are equal and that the output produced is well-formed.

The proof given by Bhat *et al.* [BBGR] (which we were unaware of while working on our proof) is analogous to ours, but much more concise due to the fact that side conditions such as measurability, integrability, non-negativity, and so on are not proven explicitly and many important (but uninteresting) steps are skipped or only hinted at.

It should be noted that in the draft of an updated version of their 2013 paper [BBGR], Bhat *et al.* added a scaling rule for real distributions similar to our `HD_MULTC` rule. However, in the process of our formal proof, we found that their rule was too liberal: while our rule only allows multiplication with a fixed constant, their rule allowed multiplication with any deterministic expression, even expressions that may evaluate to 0, but multiplication with 0 always yields the Dirac distribution, which does not have a density function. In this case, the compiler returns a PDF for a distribution that has none, leading to unsoundness. This shows the importance of formal proofs for proving compiler correctness.

⁹Note that since the abstract compiler returns parametrised density functions, we need to parametrise the result with the state $\lambda x. \text{undefined}$, even if the expression contains no free variables.

5 Concrete compiler

5.1 Approach

The concrete compiler is another inductive predicate, modelled directly after the abstract compiler, but returning a target-language expression as the compilation result instead of a HOL function. We will use a standard refinement approach to relate the concrete compiler to the abstract compiler. We thus lift the soundness result on the abstract compiler to an analogous one on the concrete one. This effectively shows that the concrete compiler always returns a well-formed target-language expression that represents a density for the sub-probability space described by the source language.

The concrete compilation predicate is written as

$$(vs, vs', \Gamma, \delta) \vdash_c e \Rightarrow f$$

Here, vs and vs' are lists of variables, Γ is a typing environment, and δ is a target-language expression describing the common density of the random variables vs in the context. It may be parametrised with the variables from vs' .

5.2 Definition

The concrete compilation rules are, of course, a direct copy of the abstract ones, but with all the abstract HOL operations replaced with operations on target-language expressions. Due to the de Bruijn indices and the lack of functions as explicit objects in the target language, some of the rules are somewhat complicated – inserting an expression into the scope of one or more bound variables (such as under an integral) requires shifting the variable indices of the inserted expression correctly.

The following list shows these rules. They heavily use the auxiliary functions listed in Sect. A.3 in the appendix.

EDC_VAL	$\frac{\text{countable_type (val_type } v)}{(vs, vs', \Gamma, \delta) \vdash_c \text{Val } v \Rightarrow \text{ins_var0 (branch_prob_cexpr (vs, vs', \Gamma, \delta))} \cdot \langle \text{CVar } 0 = \text{CVal } v \rangle}$
EDC_VAR	$\frac{x \in \text{set } vs}{(vs, vs', \Gamma, \delta) \vdash_c \text{Var } x \Rightarrow \text{marg_dens_cexpr } \Gamma \text{ } vs \ x \ \delta}$
EDC_PAIR	$\frac{x \in \text{set } vs \quad y \in \text{set } vs}{(vs, vs', \Gamma, \delta) \vdash_c \langle x, y \rangle \Rightarrow \text{marg_dens2_cexpr } \Gamma \text{ } vs \ x \ y \ \delta}$
EDC_FAIL	$\frac{}{(vs, vs', \Gamma, \delta) \vdash_c \text{Fail } t \Rightarrow \text{CVal (RealVal 0)}}$

EDC_LET

$$\frac{(\[], vs @ vs', \Gamma, 1) \vdash_c e \Rightarrow f \quad (0 \# \text{map Suc } vs, \text{map Suc } vs', \text{type_of } \Gamma e \bullet \Gamma, \text{ins_var0 } \delta \cdot f) \vdash_c e' \Rightarrow g}{(vs, vs', \Gamma, \delta) \vdash_c \text{LET } e \text{ IN } e' \Rightarrow \text{del_var0 } g}$$

EDC_RAND

$$\frac{(vs, vs', \Gamma, \delta) \vdash_c e \Rightarrow f}{(vs, vs', \Gamma, \delta) \vdash_c \text{Random } dst \Rightarrow \int_c \text{ins_var1 } f \cdot \text{dist_dens_cexpr } dst \text{ (CVar 0) (CVar 1) } \partial \text{dist_param_type } dst}$$

EDC_RAND_DET

$$\frac{e \text{ det} \quad \text{free_vars } e \subseteq \text{set } vs'}{(vs, vs', \Gamma, \delta) \vdash_c \text{Random } dst \Rightarrow \text{ins_var0 (branch_prob_cexpr (vs, vs', \Gamma, \delta))} \cdot \text{dist_dens_cexpr } dst \text{ (ins_var0 (expr_rf_to_cexpr } e)) \text{ (CVar 0)}}$$

EDC_IF

$$\frac{(\[], vs @ vs', \Gamma, 1) \vdash_c b \Rightarrow f \quad (vs, vs', \Gamma, \delta \cdot \langle \text{cexpr_subst_val } f \text{ TRUE} \rangle) \vdash_c e_1 \Rightarrow f_1 \quad (vs, vs', \Gamma, \delta \cdot \langle \text{cexpr_subst_val } f \text{ FALSE} \rangle) \vdash_c e_2 \Rightarrow f_2}{(vs, vs', \Gamma, \delta) \vdash_c \text{IF } b \text{ THEN } e_1 \text{ ELSE } e_2 \Rightarrow f_1 + f_2}$$

EDC_IF_DET

$$\frac{b \text{ det} \quad (vs, vs', \Gamma, \delta \cdot \langle \text{expr_rf_to_cexpr } b \rangle) \vdash_c e_1 \Rightarrow f_1 \quad (vs, vs', \Gamma, \delta \cdot \langle \neg \text{expr_rf_to_cexpr } b \rangle) \vdash_c e_2 \Rightarrow f_2}{(vs, vs', \Gamma, \delta) \vdash_c \text{IF } b \text{ THEN } e_1 \text{ ELSE } e_2 \Rightarrow f_1 + f_2}$$

EDC_OP_DISCR

$$\frac{\Gamma \vdash e : t \quad \text{op_type } oper \ t = \text{Some } t' \quad \text{countable_type } t' \quad (vs, vs', \Gamma, \delta) \vdash_c e \Rightarrow f}{(vs, vs', \Gamma, \delta) \vdash_c oper \$ e \Rightarrow \int_c \langle oper \$_c \text{CVar 0} = \text{CVar 1} \rangle \cdot \text{ins_var1 } f \ \partial t}$$

EDC_FST

$$\frac{\Gamma \vdash e : t \times t' \quad (vs, vs', \Gamma, \delta) \vdash_c e \Rightarrow f}{(vs, vs', \Gamma, \delta) \vdash_c \text{fst } e \Rightarrow \int_c \text{ins_var1 } f \circ_c \langle \text{CVar 1, CVar 0} \rangle_c \partial t'}$$

EDC_SND

$$\frac{\Gamma \vdash e : t \times t' \quad (vs, vs', \Gamma, \delta) \vdash_c e \Rightarrow f}{(vs, vs', \Gamma, \delta) \vdash_c \text{snd } e \Rightarrow \int_c \text{ins_var1 } f \circ_c \langle \text{CVar 0, CVar 1} \rangle_c \partial t}$$

EDC_NEG

$$\frac{(vs, vs', \Gamma, \delta) \vdash_c e \Rightarrow f}{(vs, vs', \Gamma, \delta) \vdash_c -e \Rightarrow f \circ_c (\text{-CVar 0})}$$

EDC_ADDC

$$\frac{(vs, vs', \Gamma, \delta) \vdash_c e \Rightarrow f \quad e' \text{ det} \quad \text{free_vars } e' \subseteq \text{set } vs'}{(vs, vs', \Gamma, \delta) \vdash_c e + e' \Rightarrow f \circ_c (\text{CVar 0} - \text{ins_var0 (expr_rf_to_cexpr } e'))}$$

EDC_MULTC	$\frac{(vs, vs', \Gamma, \delta) \vdash_c e \Rightarrow f \quad c \neq 0}{(vs, vs', \Gamma, \delta) \vdash_c e \cdot \text{Val} (\text{RealVal } c) \Rightarrow (f \circ_c (\text{CVar } 0 / c)) / c }$
EDC_ADD	$\frac{\Gamma \vdash \langle e, e' \rangle : t \times t \quad (vs, vs', \Gamma, \delta) \vdash_c \langle e, e' \rangle \Rightarrow f}{(vs, vs', \Gamma, \delta) \vdash_c e + e' \Rightarrow \int_c \text{ins_var1 } f \circ_c \langle \text{CVar } 0, \text{CVar } 1 - \text{CVar } 0 \rangle \partial t}$
EDC_INV	$\frac{(vs, vs', \Gamma, \delta) \vdash_c e \Rightarrow f}{(vs, vs', \Gamma, \delta) \vdash_c e^{-1} \Rightarrow (f \circ_c (1/\text{CVar } 0)) / (\text{CVar } 0)^2}$
EDC_EXP	$\frac{(vs, vs', \Gamma, \delta) \vdash_c e \Rightarrow f}{(vs, vs', \Gamma, \delta) \vdash_c \text{exp } e \Rightarrow \text{IF } \text{CVar } 0 > 0 \text{ THEN } (f \circ_c \ln (\text{CVar } 0)) / \text{CVar } 0 \text{ ELSE } 0}$

Figure 11: The concrete compilation rules

5.3 Refinement

The refinement relates the concrete compilation

$$(vs, vs', \Gamma, \delta) \vdash_c e \Rightarrow f$$

to the abstract compilation

$$(\text{set } vs, \text{set } vs', \Gamma, \lambda \sigma. \text{cexpr_sem } \sigma \delta) \vdash_c e \Rightarrow \lambda \rho x. \text{cexpr_sem } (x \bullet \rho) f$$

In words: we take the abstract compilation predicate and

- variable sets are refined to variable lists
- the typing context and the source-language expression remain unchanged
- the common density in the context and the compilation result are refined from HOL functions to target-language expressions (by applying the target language semantics)

The main refinement lemma states that the concrete compiler yields a result that is equivalent to that of the abstract compiler, modulo refinement. Informally, the statement is the following: if e is ground and well-typed under some well-formed concrete density context Y and $Y \vdash_c e \Rightarrow f$, then $Y' \vdash_d e \Rightarrow f'$, where Y' and f' are the abstract versions of Y and f .

The proof for this is conceptually simple – induction over the definition of the concrete compiler; in practice, however, it is quite involved. In every single induction step, the well-formedness of the intermediary expressions needs to be shown, the congruence lemmas for the abstract compiler need to be applied (see Sect. 4.2), and, when integration is involved, non-negativity and integrability have to be shown in order to convert non-negative integrals to Lebesgue integrals and integrals on product spaces to iterated integrals.

Combining this main refinement lemma and the abstract soundness lemma, we can now easily show the concrete soundness lemma:

lemma `expr_has_density_cexpr_sound` :

assumes $([], [], \Gamma, 1) \vdash_c e \Rightarrow f$ **and** $\Gamma \vdash e : t$ **and** $\text{free_vars } e = \emptyset$

shows $\text{has_subprob_density } (\text{expr_sem } \sigma e) (\text{stock_measure } t)$
 $(\lambda x. \text{cexpr_sem } (x \bullet \sigma) f)$

$\Gamma' 0 = t \Longrightarrow \Gamma' \vdash_c f : \text{REAL}$
 $\text{free_vars } f \subseteq \{0\}$

Informally, the lemma states that if e is a well-typed, ground source-language expression, compiling it with an empty context will yield a well-typed, well-formed target-language expression representing a density function on the measure space described by e .

5.4 Final result

We will now summarise the soundness lemma we have just proven in a more concise manner. For convenience, we define the symbol $e : t \Rightarrow_c f$ (read ‘ e with type t compiles to f ’), which includes the well-typedness and groundness requirements on e as well as the compilation result:¹⁰

$$e : t \Rightarrow_c f \iff ((\lambda x. \text{UNIT}) \vdash e : t \wedge \text{free_vars } e = \emptyset \wedge ([], [], \lambda x. \text{UNIT}, 1) \vdash_c e \Rightarrow f)$$

The final soundness theorem for the compiler, stated in Isabelle syntax, is then:¹¹

lemma `expr_compiles_to_sound` :

assumes $e : t \Rightarrow_c f$

fixes $\Gamma \Gamma' \sigma \sigma'$

shows $\text{expr_sem } \sigma e = \text{density } (\text{stock_measure } t) (\lambda x. \text{cexpr_sem } (x \bullet \sigma') f)$
 $\forall x \in \text{type_universe } t. \text{cexpr_sem } (x \bullet \sigma') f \geq 0$
 $\Gamma \vdash e : t$
 $t \bullet \Gamma' \vdash_c f : \text{REAL}$
 $\text{free_vars } f \subseteq \{0\}$

¹⁰In this definition, the choice of the typing environment is, of course, completely arbitrary since the expression contains no free variables.

¹¹To be precise, the lemma statement in Isabelle is slightly different; for better readability, we unfolded one auxiliary definition here and omitted the type cast from *real* to *ereal*.

In words, this result means the following:

Theorem

Let e be a source-language expression. If the compiler determines that e is well-formed and well-typed with type t and returns the target-language expression f , then:

- the measure obtained by taking the stock measure of t and using the evaluation of f as a density is precisely the measure obtained by evaluating e
- f is non-negative on all input values of type t
- e has no free variables and indeed has type t (in any type environment Γ)
- f has no free variable except the parameter (i. e. the variable o) and is a function from t to $REAL$ ¹²

Isabelle's code generator now allows us to execute our inductively-defined verified compiler using the **values** command¹³ or generate code in one of the target languages such as Standard ML or Haskell.

5.5 Evaluation

As an example on which to test the compiler, we choose the same expression that was chosen by Bhat *et al.* [BBGR13]:¹⁴

```
LET Random UniformReal <0,1> IN
LET Random Bernoulli (Var 0) IN
IF Var 0 THEN Var 1 + 1 ELSE Var 1
```

Using symbolic variable names instead of de Bruijn indices, the expression would look like this:

```
LET x = Random UniformReal <0,1> IN
LET y = Random Bernoulli x IN
IF y THEN x + 1 ELSE x
```

We abbreviate this expression with e . We can then display the result of the compilation using the following Isabelle command:

```
values "{(t,f) | t f. e : t ⇒c f}"
```

¹²meaning if its parameter variable has type t , it is of type $REAL$

¹³Our compiler is inherently non-deterministic since it may return zero, one, or many density functions, seeing as an expression may have no matching compilation rules or more than one. Therefore, we must use the **values** command instead of the **value** command and receive a set of compilation results.

¹⁴Val and RealVal were omitted for better readability.

The result is a singleton set which contains the pair (REAL, f) , where f is a very long and complicated expression. Simplifying constant sub-expressions and expressions of the form $\text{fst } \langle e_1, e_2 \rangle$ and again replacing de Bruijn indices with symbolic identifiers, we obtain:

$$\int b. (\text{IF } 0 \leq x - 1 \wedge x - 1 \leq 1 \text{ THEN } 1 \text{ ELSE } 0) \cdot \\ (\text{IF } 0 \leq x - 1 \wedge x - 1 \leq 1 \text{ THEN IF } b \text{ THEN } x - 1 \text{ ELSE } 1 - (x - 1) \text{ ELSE } 0) \cdot \langle b \rangle + \\ \int b. (\text{IF } 0 \leq x \wedge x \leq 1 \text{ THEN } 1 \text{ ELSE } 0) \cdot \\ (\text{IF } 0 \leq x \wedge x \leq 1 \text{ THEN IF } b \text{ THEN } x \text{ ELSE } 1 - x \text{ ELSE } 0) \cdot \langle \neg b \rangle$$

Further simplification yields the following result:

$$\langle 1 \leq x \leq 2 \rangle \cdot (x - 1) + \langle 0 \leq x \leq 1 \rangle \cdot (1 - x)$$

While this result is the same as that which Bhat *et al.* have reached, our compiler generates a much larger expression than the one they printed. The reason for this is that they printed a β -reduced version of the compiler output; in particular, constant sub-expressions were evaluated. While such simplification is, of course, very useful when using the compiler in practice, we have not implemented it since it is not conceptually interesting and outside the scope of this work.




Figure 12: The graph of the density function of the example expression

6 Conclusion

6.1 Breakdown

All in all, the formalisation of the compiler took about three months. It contains a total of roughly 10000 lines of Isabelle code (definitions, lemma statements, proofs, and examples). Figure 13 shows a detailed breakdown of this.

Type system and semantics	2900 lines
Abstract compiler	2600 lines
Concrete compiler	1400 lines
General Measure Theory / auxiliary lemmas	3400 lines

Figure 13: Breakdown of the Isabelle code

As can be seen from this figure, a sizeable portion of the work was the formalisation of results from general measure theory, such as the substitution lemmas for Lebesgue integrals, and auxiliary notions and lemmas, such as measure embeddings. Since the utility of these formalisations is not restricted to this particular project, they will be moved to Isabelle’s Measure Theory library.

6.2 Difficulties

The main problems we encountered during the formalisation were:

Missing background theory. As mentioned in the previous section, a sizeable amount of Measure Theory and auxiliary notions had to be formalised. Most notably, the existing Measure Theory library did not contain integration by substitution. As a side product of their formalisation of the Central Limit Theorem [AHS14], Avigad *et al.* proved the Fundamental Theorem of Calculus and, building thereupon, integration by substitution. However, their integration-by-substitution lemma only supported continuous functions, whereas we required the theorem for general Borel-measurable functions. Using their proof of the Fundamental Theorem of Calculus, we proved such a lemma, which initially comprised almost 1000 lines of proof, but has been shortened significantly thereafter.

Proving side conditions. Many lemmas from the Measure Theory library require measurability, integrability, non-negativity, etc. In hand-written proofs, this is often ‘hand-waved’ or implicitly dismissed as trivial; in a formal proof, proving these can blow up proofs and render them very complicated and technical. The measurability proofs in particular are ubiquitous in our formalisation. The Measure Theory library provides some tools for proving measurability automatically, but while they were quite helpful in many cases, they are still work in progress and require more tuning.

Lambda calculus. Bhat *et al.* use a simply-typed Lambda-calculus-like language with symbolic identifiers as a target language. For a paper proof, this is the obvious choice, since it leads to concise and familiar definitions, but in a formal setting, it always comes with the typical problems of having to deal with variable capture and related issues. For that reason, we chose to use de Bruijn indices instead; however, this makes handling target-language terms less intuitive, since variable indices need to be shifted whenever several target-language terms are combined.

Another issue was the lack of a function type in our target language; allowing first-order functions, as Bhat *et al.* did, would have made many definitions easier and more natural, but would have complicated others significantly due to measurability issues.

Furthermore, we effectively needed to formalise a number of properties of Lambda calculus that can be used implicitly in a paper proof.

6.3 Future work

The following improvements to the Isabelle formalisation could probably be realised with little effort:

- sum types and a `match . . . with . . .` statement
- a compiler rule for deterministic ‘let’ bindings
- ‘let’ bindings for the target language
- a pre-processing stage to allow ‘normal’ variables instead of de Bruijn indices
- a post-processing stage to simplify the density expression as far as possible

The first of these is interesting because it is the only substantial weakness of our formalisation compared to that by Bhat *et al.*; the remaining four merely make using the compiler more convenient or more efficient.

Additionally, in the long term, a *Markov-chain Monte Carlo (MCMC)* sampling method such as the *Metropolis–Hastings algorithm* could also be formalised in Isabelle and integrated with the density compiler. However, this would be a more involved project as it will require the formalisation of additional probability theory in Isabelle.

6.4 Summary

Using Isabelle/HOL, we formalised the semantics of a simple probabilistic functional programming language with predefined probability distributions and a compiler that returns the probability distribution that a program in this language describes. These are modelled very closely after those given by Bhat *et al.* [BBGR13]. We then used the existing formalisations of measure theory in Isabelle/HOL to formally prove the correctness of this compiler w. r. t. the semantics of the source and target languages.

This shows not only that the compiler given by Bhat *et al.* is correct (apart from the minor soundness issue we uncovered), but also that a formal correctness proof for such a compiler can be done with reasonable effort and that Isabelle/HOL in general and its Measure Theory library in particular are suitable for it.

Appendix

A Notation and Auxiliary Functions

A.1 General notation

NOTATION	NAME / DESCRIPTION	DEFINITION
$f x$	function application	$f(x)$
$f ' X$	image set	$f(X)$ or $\{f(x) \mid x \in X\}$
$\lambda x. e$	lambda abstraction	$x \mapsto e$
undefined	arbitrary value	
Suc	successor of a natural number	+ 1
case_nat $x f y$	case distinction on natural number	$\begin{cases} x & \text{if } y = 0 \\ f(y - 1) & \text{otherwise} \end{cases}$
$[]$	Nil	empty list
$x \# xs$	Cons	prepend element to list
$xs @ ys$	list concatenation	
map $f xs$	applies f to all list elements	$[f(x) \mid x \leftarrow xs]$
merge $V V' (\rho, \sigma)$	merging disjoint states	$\begin{cases} \rho x & \text{if } x \in V \\ \sigma y & \text{if } x \in V' \\ \text{undefined} & \text{otherwise} \end{cases}$
$y \bullet f$	add de Bruijn variable to scope	see Sect. 2.1.3
$\langle P \rangle$	indicator function	1 if P is true, 0 otherwise
$\int x. f x \partial \mu$	Lebesgue integral	$\int f(x) d\mu(x)$
$\int^+ x. f x \partial \mu$	Lebesgue integral on non-neg. part	$\int \max(0, f(x)) d\mu(x)$
\mathfrak{Meas}	category of measurable spaces	see Sect. 2.2.3
S	sub-probability functor	see Sect. 2.2.4
return	monadic return (η) in the Giry monad	see Sect. 2.2.5
join	monadic join (μ) in the Giry monad	see Sect. 2.2.5
$\gg=$	monadic bind in the Giry monad	see Sect. 2.2.5
do $\{ \dots \}$	monadic 'do' syntax	see Sect. 2.2.5
density $M f$	measure with density	result of applying density f to M
distr $M N f$	push-forward/image measure	$(B, \mathcal{B}, \lambda X. \mu(f^{-1}(X)))$ for $M = (A, \mathcal{A}, \mu), N = (B, \mathcal{B}, \mu')$

A.2 Semantics notation and general auxiliary functions

NOTATION/FUNCTION	DESCRIPTION
$\Gamma \vdash e : t$	source language typing, see Fig. 4
$\Gamma \vdash_c e : t$	target language typing, see Fig. 8
$Y \vdash_d e \Rightarrow f$	abstract compiler, see Sect. 4.2
$Y \vdash_c e \Rightarrow f$	concrete compiler, see Sect. 5.2
$e : t \Rightarrow_c f$	'compiles to' predicate, see Sect. 5.4
op_sem	operator semantics, see Fig. 1
dist_param_type dst	parameter type of the built-in distribution dst
dist_result_type dst	result type of the built-in distribution dst
dist_measure $dst\ x$	built-in distribution dst with parameter x
dist_dens $dst\ x\ y$	density of the built-in distribution dst w. parameter x at value y
expr_sem	source language semantics, see Fig. 6
expr_sem_rf	semantics for deterministic expressions, see Sect. 3.4
cexpr_sem	target language semantics, see Fig. 9
type_of $\Gamma\ e$	the unique t such that $\Gamma \vdash e : t$
val_type v	the type of value v , e. g. $val_type\ (IntVal\ 42) = INTEG$
type_universe t	the set of values of type t
countable_type t	true iff $type_universe\ t$ is a countable set
free_vars e	the free (i. e. non-bound) variables in the expression e
$e\ det$	true iff e does not contain <i>Random</i> or <i>Fail</i>
extract_real x	returns y for $x = RealVal\ y$ (analogous for int, pair, etc.)
embed_measure $M\ f$	embedding measure M with inj. function f , see Sect. 3.2.1
stock_measure t	stock measure for type t , see Sect. 3.2.2
state_measure $\Gamma\ V$	measure on states, see Sect. 3.2.2
return_val v	$return\ (stock_measure\ (val_type\ v))\ v$
null_measure M	measure with same σ -algebra as M , but measure 0 for all sets
branch_prob	probability of being in the current branch, see Sect. 4.2
marg_dens	marginal density of a random variable, see Sect. 4.2
marg_dens2	marginal density of two random variables, see Sect. 4.2
has_subprob_density $M\ N\ f$	f is a valid density function, applying density f to N yields M , and M is a sub-probability measure. See Sect. 4.2
has_parametrized_subprob_density $R\ M\ N\ f$	As before, but with additional parameter from R . See Sect. 4.2

A.3 Target language auxiliary functions

For the definitions of the following functions, see the corresponding Isabelle theory file `PDF_Target_Semantics.thy`. We will not print them here since they are rather technical.

In this section, x will always be a variable, v will always be a value, and e and e' will always be target language expressions. f and g will always be target language expressions interpreted as functions, i. e. with an implicit λ abstraction around them.

All functions in the table take de Bruijn indices into account, i. e. they will shift variables when they enter the scope of an integral.

FUNCTION	DESCRIPTION
<code>map_vars</code> $h e$	transforms all variables in e with the function h
<code>ins_var0</code> e	$map_vars\ Suc\ e$ (prepares expression e for insertion into scope of new variable)
<code>ins_var1</code> f	$map_vars\ (case_nat\ 0\ (\lambda x. x + 2))\ f$ (prepares function f for insertion into scope of new variable)
<code>del_var0</code> e	$map_vars\ (\lambda x. x - 1)\ e$ (shifts all variables down, deleting the variable 0)
<code>cexpr_subst</code> $x e e'$	substitutes e for x in e'
<code>cexpr_subst_val</code> $e v$	substitutes v for <code>CVar 0</code> in e , effectively applying e interpreted as a function to the argument v .
<code>cexpr_comp</code> $f g$	function composition, i. e. $\lambda x. f (g x)$
$f \circ_c g$	alias for <code>cexpr_comp</code>
<code>expr_rf_to_cexpr</code> e	converts a deterministic source language expression into a target language expression, see Sect. 3.5
<code>integrate_var</code> $\Gamma x e$	integrates e over the free variable x
<code>integrate_vars</code> $\Gamma xs e$	integrates e over the free variables in the list xs
<code>branch_prob_cexpr</code> $(vs, vs', \Gamma, \delta)$	returns an expression that computes <code>branch_prob</code>
<code>marg_dens_cexpr</code> $\Gamma vs x \delta$	returns an expression that computes <code>marg_dens</code>
<code>marg_dens2_cexpr</code> $\Gamma vs x y \delta$	returns an expression that computes <code>marg_dens2</code>
<code>dist_dens_cexpr</code> $dst e e'$	returns an expression that computes the density of the built-in distribution <code>dst</code> , parametrised with e and evaluated at e'

B References

- [AHS14] Jeremy Avigad, Johannes Hölzl, and Luke Serafin. A formally verified proof of the Central Limit Theorem. *CoRR*, abs/1405.7012, 2014.
- [APM06] Philippe Audebaud and Christine Paulin-Mohring. Proofs of randomized algorithms in Coq. In *Mathematics of Program Construction*, volume 4014 of *Lecture Notes in Computer Science*, pages 49–68. Springer Berlin Heidelberg, 2006.
- [BAVG12] Sooraj Bhat, Ashish Agarwal, Richard Vuduc, and Alexander Gray. A type theory for probability density functions. In *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '12, pages 545–556, New York, NY, USA, 2012. ACM.
- [BBGR] Sooraj Bhat, Johannes Borgström, Andrew D. Gordon, and Claudio Russo. Deriving probability density functions from probabilistic functional programs. (full version, submitted for publication).
- [BBGR13] Sooraj Bhat, Johannes Borgström, Andrew D. Gordon, and Claudio Russo. Deriving probability density functions from probabilistic functional programs. In *Tools and Algorithms for the Construction and Analysis of Systems*, volume 7795 of *Lecture Notes in Computer Science*, pages 508–522. Springer Berlin Heidelberg, 2013. Best Paper Award.
- [BC04] Yves Bertot and Pierre Castéran. *Interactive Theorem Proving and Program Development. Coq'Art: The Calculus of Inductive Constructions*. Springer, 2004.
- [Coc12] David Cock. Verifying probabilistic correctness in Isabelle with pGCL. In *Proceedings of the 7th Systems Software Verification*, pages 1–10, November 2012.
- [Coc14] David Cock. pGCL for Isabelle. *Archive of Formal Proofs*, July 2014. <http://afp.sf.net/entries/pGCL.shtml>, Formal proof development.
- [Dob07] Ernst-Erich Doberkat. *Stochastic relations: foundations for Markov transition systems*. Studies in Informatics. Chapman & Hall/CRC, 2007.
- [Dob08] Ernst-Erich Doberkat. Basing Markov transition systems on the Giry monad. http://www.informatics.sussex.ac.uk/events/domains9/Slides/Doberkat_GiryMonad.pdf, 2008.
- [Gir82] Michèle Giry. A categorical approach to probability theory. In *Categorical Aspects of Topology and Analysis*, volume 915 of *Lecture Notes in Mathematics*, pages 68–85. Springer Berlin Heidelberg, 1982.
- [HMM05] Joe Hurd, Annabelle McIver, and Carroll Morgan. Probabilistic guarded commands mechanized in HOL. *Electron. Notes Theor. Comput. Sci.*, 112:95–111, January 2005.
- [Hur02] Joe Hurd. *Formal Verification of Probabilistic Algorithms*. PhD thesis, University of Cambridge, 2002.
- [Hö12] Johannes Hölzl. Construction and stochastic applications of measure spaces in Higher-Order Logic. PhD thesis, Technische Universität München, Institut für Informatik, 2012.
- [PPT05] Sungwoo Park, Frank Pfenning, and Sebastian Thrun. A probabilistic language based upon sampling functions. In *Proceedings of the 32Nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '05, pages 171–182, New York, NY, USA, 2005. ACM.