

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49

Safe Couplings: Coupled Refinement Types

LISA VASILENKO, IMDEA Software Institute, Spain
NIKI VAZOU, IMDEA Software Institute, Spain
GILLES BARTHE, MPI-SP, Germany IMDEA Software Institute, Spain

We enhance refinement types with mechanisms to reason about relational properties of probabilistic computations. Our mechanisms, which are inspired from probabilistic couplings, are applicable to a rich set of probabilistic properties, including expected sensitivity, which ensures that the distance between outputs of two probabilistic computations can be controlled from the distance between their inputs. We implement our mechanisms in the type system of Liquid Haskell and we use them to formally verify Haskell implementations of two classic machine learning algorithms: Temporal Difference (TD) reinforcement learning and stochastic gradient descent (SGD). We formalize a fragment of our system for discrete distributions and we prove soundness with respect to a set-theoretical semantics.

Additional Key Words and Phrases: refinement types, relational types, program verification, Haskell

1 INTRODUCTION

Refinement types provide an appealing mechanism for proving program properties in executable programming languages (including Haskell [Vazou et al. 2014b], Scala [Hamza et al. 2019], and F* [Swamy et al. 2016]). They have been used to good effect for reasoning about functional correctness and termination [Vazou et al. 2014a], resource analysis [Handley et al. 2019], security policies [Lehmann et al. 2021], and other properties of large developments.

However, refinement types do not provide support for reasoning about relational and hyper properties. The main difference between trace properties, which are the usual target of refinement type systems, and relational and hyper properties is that the latter reason about pairs of traces and sets of traces. This generalization allows to account for a wide variety of security, privacy, and robustness properties.

One natural approach to support relational reasoning is to use relational type systems, as proposed for instance in [Barthe et al. 2014; Maillard et al. 2020]. These type systems are similar to classic refinement type systems, but reason about relational assertions. The latter are interpreted over pairs of (typed) values, and therefore capture relational properties in a natural way. Relational refinement type systems retain the feel of refinement types and are particularly effective when reasoning about two executions of the same program, or two programs that follow the same control-flow.

Unfortunately, relational refinement types offer limited support to reason about programs with diverging control-flow. This is due to the fact that relational refinement types are syntax-directed, whereas many examples of relational program verification benefit from or even require non-syntax-directed reasoning. This is in particular the case for reasoning about program optimizations that restructure the control flow of the program, and about probabilistic programs, since their correctness or security proofs often use mathematical arguments that is not reflected in their syntax. Moreover, it remains a challenge to make relational refinement types practical, even in settings where syntax-directed reasoning suffices. This difficulty is perhaps best witnessed by prior work on relational cost [Çiçek et al. 2019]. In this work, Çiçek et al. [2019] develop BiRelCost, a state-of-the-art bi-directional type checker that compares the cost of two programs, or two program executions. In this setting, relational syntax-directed reasoning alternates with non-relational syntax-directed reasoning, in a way that the latter takes over whenever the two program executions

Authors' addresses: Lisa Vasilenko, elizaveta.vasilenko@imdea.org, IMDEA Software Institute, Madrid, Spain; Niki Vazou, niki.vazou@imdea.org, IMDEA Software Institute, Madrid, Spain; Gilles Barthe, gbarthe@mpi-sp.org, gilles.barthe@imdea.org, MPI-SP, Bochum, Germany and IMDEA Software Institute, Madrid, Spain.

no longer have the same control-flow. Unfortunately, controlling such alternations automatically by typing ultimately relies on intricate and partial heuristics. As a consequence, BiRelCost sacrifices predictability and generality, which are some key advantages of refinement types.

A principled approach to overcome the limitations of relational refinement types is to impose a separation between types and relational assertions. This approach, which is realized by Relational Higher-Order Logic (RHOL) [Aguirre et al. 2017], ensures maximal flexibility and expressiveness. However, the approach is not implemented, and as a consequence it remains an open question if Relational Higher-Order Logic can be made practical.

In this paper, we explore a middle ground approach that retains key benefits of refinement types. The crux of our approach is a carefully crafted interface for supporting relational reasoning *within* a unary refinement type system. Our approach is implemented atop Liquid Haskell [Vazou et al. 2014b] and inherits many of its essential features: first, our formal guarantees hold for Haskell programs and these programs can be executed using the existing runtime system and optimized libraries of Haskell. Second, verification is carried using a mature refinement type checker and should be familiar for users of Liquid Haskell. Third, the known techniques of Vazou et al. [2018] for encoding proofs manually remain applicable. Naturally, these benefits come at a cost: concretely, our proofs are less automated than proofs in classic Liquid Haskell. However, proofs remain reasonably short, even for relatively complex examples, demonstrating that our middle ground approach achieves predictable and practical verification, a combination that has not been achieved by any prior relational verification tool.

We realize our approach not only for classic higher-order programs, but also for probabilistic programs, an important class of programs that is used pervasively in cryptography, privacy, machine learning and many other areas. In addition to their numerous applications, probabilistic algorithms are an interesting class of programs to consider in their own right, because they often have intricate specifications and complex proofs. In particular, many properties of interest of probabilistic programs are quantitative, i.e. they reason about probabilities or expectations — or in a relational setting, about differences between probabilities or expectations. Although such forms of quantitative reasoning are *seemingly* out of reach of SMT-based verification, prior work has shown that relational verification of probabilistic programs can be achieved using *probabilistic couplings* [Barthe and Hsu 2020; Lindvall 2002; Thorisson 2000; Villani 2008]. We formalize the main tools from coupling-based reasoning in our framework, and illustrate how these tools can be used to verify two classic examples of probabilistic programs from machine learning. The first example is Temporal Difference (TD) reinforcement learning, for which we show rapid convergence to a stationary distribution independently of its initial input. The second example is Stochastic Gradient Descent (SGD), for which we show algorithmic stability— a classic machine learning property ensuring that a supervised machine learning algorithm generalizes well and does not overfit with respect to its training set. Both properties are captured in our system as instances of *expected sensitivity*, i.e. they upper bound the expected distance between two outputs of the program as a function of the distance between the corresponding inputs. However, both examples use distinct proof tools: the first example is verified using classic techniques from probabilistic couplings. In contrast, this second example uses a more elaborate, quantitative, form of probabilistic couplings which embeds reasoning about the Kantorovich distance between two distributions. Thus, the two examples showcase the different components of our system.

Summary of contributions. Our contributions are the following:

- We define a probabilistic relational refinement type system and encode it into the unary types of Liquid Haskell (§ 3). We choose Liquid Haskell as a mature refinement type checker, but our methodology can be used to encode any relational to any unary refinement type system.

- We use our system to prove two case studies from the literature: TD (§ 4.1) and SGD (§ 4.2).
- We prove soundness of our type system with respect to a denotational semantics (§ 5).

We start with an overview of our approach (§ 2) and conclude with related work (§ 6) and future work (§ 7).

2 OVERVIEW

We start with an overview of our system that uses (unary) refinement types to machine check relational properties of probabilistic, executable programs. First (§ 2.1) we introduce the `PrM` probabilistic monad and our `bins` running example. Next, we encode (§ 2.2) and formally prove (§ 2.3) a relational specification for the returned values of `bins` by axiomatizing probabilistic relational logic as refinement type assumptions. Finally, we follow a similar methodology to encode (§ 2.4) and prove (§ 2.5) a relational property about `bins` distance.

2.1 The Probability Distribution

The common way to implement probability distributions in Haskell is to use a probability monad, see for instance [Ramsey and Pfeffer 2002]. Therefore, our framework is set up as a verification wrapper around any Haskell library that supports a monadic implementation of probabilities. In order to execute our implementations, we wrapped the `probability` library¹; however, our proofs are independent on the choice of the library, and only requires the existence of some type `PrM` that implements the standard interface for a probabilistic monad. This includes `pure` and `bind` and constructors for `bernoulli`, `choice`, and uniform distributions.

```
type PrM a = ... -- defined in Sec 3 using an existing probabilistic Haskell library
```

```
{-@ type Prob = {p:Double | 0 ≤ p ≤ 1} @-}
{-@ pure      :: a → PrM a @-}
{-@ (>>=)    :: PrM a → (a → PrM b) → PrM b @-}

{-@ bernoulli :: Prob → PrM {v:Integer | v == 0 || v == 1} @-}
{-@ choice    :: Prob → PrM a → PrM a → PrM a @-}
{-@ unif      :: {xs:[a] | 0 < len xs} → PrM a @-}
```

We define the Haskell probability monad `PrM` using an interface of an existing library (§ 3). We use the notation `{-@ ... @-}` to define refinement types and refinement type specifications. That is, the `Prob` type is a Haskell `Double`, refined to be between 0 and 1. The specifications for the monadic `pure` and `>>=` are standard. The `bernoulli` function takes an input a probability `p`, *i.e.* a `Double` between 0 and 1, and returns 1 with probability `p`, otherwise 0. The function `choice p d1 d2` returns the distribution `d1` with probability `p`, otherwise `d2`. Finally, `unif` takes an input a non-empty list and returns one of its elements uniformly at random. All these functions are executed using the underlying Haskell implementation, but are left as uninterpreted (later § 3.4 and § 3.5 axiomatized) in our logic.

Probabilistic Programming: The Bins Example. Using the above interface, we can define (and execute) probabilistic programs. For example below we define the `bins` program that models a simple balls and bins process. In this process, `n` balls are thrown into a bin; in each throw, there is a probability that the ball lands outside the bin. The throws are independent and the probability to send any ball in the bin is `p`. The result of the process is the number of balls that lands into the bin. We model the process using the `bernoulli` distribution.

¹ We used the probabilistic functional programming library <https://hackage.haskell.org/package/probability-0.2.7>

```

148   {-@ type Nat = {n:Integer | 0 ≤ n }           @-}
149
150   {-@ bins :: n:Nat → Prob → PrM {r:Nat | r ≤ n} @-}
151   bins 0 _ = pure 0
152   bins n p = do x ← bins (n - 1) p
153               y ← bernoulli p
154               pure (x + y)

```

We can use standard refinement types to verify various *unary* properties of the `bins` function. In particular, Liquid Haskell will use SMT automation to easily verify that `bins` terminates (because the recursive call occurs on smaller `n`). One can also prove that the result is always a natural number (as specified by the refined signature) and that it is not greater than `n`.

Using the theorem proving capabilities of Liquid Haskell [Vazou et al. 2017], we can construct extrinsic proofs that validate probabilistic, unary properties of `bins`. For example, we can define `expect f e` to be the expected value of `e`, for some function `f` that turns the values of `e` to booleans:

```

162   {-@ expect :: (a → Double) → PrM a → Double   @-}
163   {-@ natToD :: n:Nat → {d:Double | d = to_real n} @-}

```

We can extrinsically prove that `expect natToD (bins n p) = n * p`, assuming that the expectation is linear and `expect natToD (bernoulli p) = p`. Note that the Haskell `Double` is represented, by Liquid Haskell, as `real` in SMT, so the function `natToD` converts natural to double numbers while its specification ensures that the value is not changed. Next, we see how to construct extrinsic proofs that establish relational properties.

2.2 Relational Specifications & Lifting

A first relational property of interest is stochastic dominance, a classic property that defines when a real-valued probabilistic process is better than another. Informally, a real-valued probabilistic process is better than another if it always outputs a “higher value” w.r.t. the usual order on real numbers. Interestingly, the intuitive notion of “higher value” is formally defined over two random variables, which makes the definition of stochastic dominance non-trivial. Fortunately, stochastic dominance can be characterized using probabilistic couplings, a classic tool to reason about Markov chains. Informally, couplings are probabilistic equivalent of cartesian products, and can be used to lift relational properties to distributions. Informally, the lifting of a relational property is defined as follows: two distributions satisfy the lifting of a property p if there exists a coupling of the two distributions such that p holds surely, i.e. with probability 1, in this coupling. The formal definition can be found e.g. in [Barthe and Hsu 2020]. For our purposes, it suffices to assume an operator \diamond that transforms a relation over two types into a relation over probabilistic distributions over these two types:

```

185   (◇) :: (a → b → Bool) → PrM a → PrM b → Bool

```

The operator (\diamond) supports compositional reasoning via two axioms: `pureAxiom` states that Dirac distributions of elements related by p are related by the lifted relation $\diamond p$ and `bindAxiom` states that lifted relations are preserved by monadic composition. These axioms, which are expressed below using refinement type signatures, conveniently eschew probabilistic reasoning, and open the possibility of carrying SMT-based verification:

```

192   {-@ assume pureAxiom :: p:(a → b → Bool) → x_l:a → x_r:b → {p x_l x_r}
193               → {◇ p (pure x_l) (pure x_r)} @-}
194
195   {-@ assume bindAxiom :: p:(b → b → Bool) → q:(a → a → Bool)

```

```

197     → el:PrM a → fl:(a → PrM b) → er:PrM a → fr:(a → PrM b)
198     → {◊ q el er}
199     → (xl:a → xr{a | q xl xr} → {◊ p (fl xl) (fr xr)})
200     → {◊ p (el >>= fl) (er >>= fr)} @-}

```

Both `pureAxiom` and `bindAxiom` are defined as Haskell functions that return `()`, but their refinement type signatures encode the desired axioms. The **assume** keyword prevents refinement type checking, since the function definitions do not actually inhabit their type specifications. `pureAxiom` states that for all `p, xl xr`, if `p xl xr`, then `◊ p (pure xl) (pure xr)`. The type `{p xl xr}` is an abbreviation of `{v:() | p xl xr}`. In general, we can write `{q}` instead of `{v:a | q}`, when `v` does not appear in `q`. Similarly, the `bindAxiom` ensures `◊ p (el >>= fl) (er >>= fr)` when it is provided a proof that `◊ q el er` and a (higher-order) proof that for all `xl` and `xr` such that `q xl xr`, `◊ p (fl xl) (fr xr)` holds. In § 3 we discuss the implementation of our library that includes these two assumptions, while later (§ 5) we develop a formalisation that justifies these assumptions, concretely, the `pureAxiom` and `bindAxiom` are respectively encoded in the rules T-RET and T-BIND of Figure 7.

In our `bins` running example, we are interested to show that for two competing throwers sending the same number of balls into bins, the more gifted thrower, i.e. the thrower with a higher probability to send balls into the bins, will have a higher count. More formally, our goal is to show that if `p ≤ q` then `◊ (≤) (bins n p) (bins n q)`, from which one can conclude that `expect naToD (bins n p) ≤ expect naToD (bins n q)` by a simple property of couplings. In our syntax, we formalize our goal as:

```

217 {-@ binsSpec :: p:Prob → {q:Prob | p ≤ q} → n:Nat → {◊ (≤) (bins n p) (bins n q)} @-}

```

In the next paragraph, we use relational refinement types to establish this goal.

2.3 Relational Proofs

Following Handley et al. [2019], relational proofs are (Haskell) inhabitants of the refinement type specification that expresses the relational specification. Such proofs rely on assumptions about relational properties of the probabilistic primitives. For example, `bins` is using `bernoulli`, thus the proof of `binsSpec` relies on the assumption below, which captures `bernoulli`'s relational specification.

```

227 {-@ assume bernoulliAxiom :: p:Prob → {q:Prob | p ≤ q}
228     → {◊ (≤) (bernoulli p) (bernoulli q)} @-}

```

The specification (formalized in rule T-BERN of § 5; Figure 7) states that `bernoulli q` stochastically dominates `bernoulli p` if `p ≤ q`. In our current implementation, this specification is taken as an axiom, although it would be possible to establish this specification from first principles, by making the definition of lifting explicit for finitely supported distributions.

Using the `bernoulliAxiom` we prove `binsSpec` following the structure of `bins` definition:

```

235 binsSpec p q 0
236   = pureAxiom (≤) 0 0 ()
237 binsSpec p q n
238   = bindAxiom (≤) (≤) (bernoulli p) (bins1 n p) (bernoulli q) (bins1 n q)
239     (bernoulliAxiom p q) (\xl xr →
240     bindAxiom (≤) (≤) (bins (n-1) p) (bins2 xl) (bins (n-1) q) (bins2 xr)
241     (binsSpec p q (n-1))) (\yl yr →
242     pureAxiom (≤) (yl + xl) (yr + xr) ())
243 where
244 bins1 n p x = bind (bins (n-1) p) (bins2 x)

```

245

246 bins2 x y = pure (y + x)

247 The proof, as `bins`, is inductively defined on `n`. In the base case, we call `pureAxiom` $(\leq) \ 0 \ 0 \ ()$
 248 to get $\diamond (\leq) \ (\text{pure } 0) \ (\text{pure } 0)$, which concludes the proof, since `bins` $0 \ p == \text{pure } 0$. Such
 249 rewrite steps are automated by Liquid Haskell’s logical evaluation strategy (namely PLE [Vazou
 250 et al. 2017]). Further, our proofs are automated by SMT arithmetic. For example, `pureAxiom`’s last
 251 argument needs to prove that $0 \leq 0$, which is trivially shown by $()$ and SMT automation. The
 252 inductive case starts with a call to `bindAxiom`, again following `bins` inductive definition. There are
 253 two interesting points here. First, the `bins` definition binds `bernoulli p` to a continuation. Since
 254 `bindAxiom`’s 4th and 6th arguments require to explicitly pass these continuations, we use a `where`
 255 clause to name the continuation `bins1`. Second, `bindAxiom` requires two proof terms. The first one
 256 should show that $\diamond (\leq) \ (\text{bernoulli } p) \ (\text{bernoulli } q)$, which is shown by the `bernoulliAxiom`.
 257 The second one, should show $\diamond (\leq) \ (\text{bins1 } n \ p \ x_l) \ (\text{bins1 } n \ q \ x_r)$, for all x_l and x_r that satisfy
 258 $x_l \leq x_r$. We construct such a proof term using a `lambda`². Since the `bins` definition is using another
 259 `bind`, the proof again calls `bindAxiom` with similar arguments. In the last step, `bins` calls `pure`, thus
 260 the proof calls `pureAxiom`, whose proof argument is again $()$, *i.e.* automated by rewriting and SMT.
 261

262 2.4 Quantitative Specifications and Kantorovich lifting

263 So far, we have established that if $p \leq q$ then `bins` `n` `q` stochastically dominates `bins` `n` `p` and
 264 thus `expect natToD (bins n p) ≤ expect natToD (bins n q)`. However, our specification does
 265 not provide quantitative information on `expect natToD (bins n q) - expect natToD (bins n p)`.
 266 In fact, one can use simple properties of expectation to show that if $p \leq q$ then we have `expect`
 267 `natToD (bins n q) - expect natToD (bins n p) ≤ n * (q-p)`, where `natToD` is just the cast
 268 defined in § 2.1. Unfortunately, one cannot prove this fact using the previous approach based on
 269 lifting. The solution is to use a richer notion of lifting, that allows to reason about quantitative
 270 properties, and in particular about the expected distance between two distributions. In order to
 271 accommodate such reasoning, one considers a richer setting where each type is equipped with a
 272 distance `dist`. These distances are defined inductively on the structure of types; for distribution
 273 types, they use the so-called Kantorovich metric [Deng 2015; Villani 2009], which lifts a distance
 274 over some types to a distance over its corresponding distribution type:
 275

```
276 {-@ dist :: Dist a → (a → a → {d:Double | 0 ≤ d}) @-}
277 {-@ kant :: Dist a → Dist (PrM a) @-}
278
279 {-@ kdist :: Dist a → PrM a → PrM a → {d:Double | 0 ≤ d} @-}
280 kdist d = dist (kant d)
```

281 The `kant` function turns a distance into Kantorovich and `kdist` simply composes `kant` with `dist`.
 282 The formal definition of the Kantorovich metric can be found for instance in Deng [2015]; for this
 283 work, it suffices that the Kantorovich metric is also based on couplings and also lends itself to
 284 compositional reasoning. For instance, the following axioms (corresponding to the rules T-RET and
 285 T-BIND § 5; Fig. 7) are valid:

```
286 {-@ assume pureDist :: d:Dist a
287     → x_l:a → x_r:a
288     → { kdist d (pure x_l) (pure x_r) = dist d x_l x_r } @-}
289
290
291 {-@ assume bindDist :: d:Dist b
```

292 ² The actual proof is using a named function with explicit type specification, since Liquid Haskell does not infer preconditions,
 293 but for space, here we use `lambdas`.

```

295         → m:Double → p:(a → a → Bool)
296         → fl:(a → PrM b) → el:PrM a
297         → fr:(a → PrM b) → er::{PrM a | ◊ p el er }
298         → (xl:a → {xr:a | p xl xr}) → { kdist d (fl xl) (fr xr) ≤ m}
299         → { kdist d (el >>= fl) (er >>= fr) ≤ m }

```

300 The first axiom states that the Kantorovich distance of two Dirac distributions is the distance of
301 their generating element. The second axiom upper bounds the Kantorovich distance between two
302 monadic compositions by the maximal Kantorovich distance between f_l x_l and f_r x_r for all x_l and
303 x_r related by an intermediate assertion p , such that e_l and e_r are related by the lifting of p . We
304 emphasize that the `bind` rule for Kantorovich distance is more intricate than the corresponding
305 rules for lifting, and that the “obvious” compositional rule that adds distance between the `es` and
306 the `fs` would not be sound, as discussed in § 5.

307 Then, the distance spec of `bins` is:

```

308 {-@ binsDist :: p:Prob → {q:Prob | p ≤ q} → n:Nat
309         → { kdist distNat (bins n p) (bins n q) ≤ n * (q - p) } @-}
310
311 {-@ distNat :: Dist Nat @-}
312

```

313 That is, for each probabilities p and q , so that $p \leq q$ and each natural number n the Kantorovich
314 distance between `bins n p` and `bins n q` is bounded by $n * (q - p)$. Since `bins` returns natural
315 numbers, the distance is given by `distNat` that defines the distance metric on natural numbers
316 (§ 3.2). Finally, Haskell’s `Doubles` are represented in the SMT logic are SMT reals, *i.e.* there is no
317 reasoning about overflows and precision loss. The `binsDist` specification is well sorted in Z3, since
318 Z3 automatically converts between `int` and reals.

320 2.5 Distance Proofs

321 Unlike the proof of § 2.3, the proof of `binsDist` is not syntax directed. On the contrary, it requires
322 the construction of a “ghost” probabilistic function that splits the distance between `bins n p` and
323 `bins n q`. We call this function `ghost bins (gbins)` and define it as follows:

```

324 gbins :: Nat → Prob → Prob → PrM Nat
325 gbins n p q = do x ← bins (n-1) p
326                y ← bernoulli q
327                pure (x + y)
328

```

329 The ghost `gbins n p q` adds `bins` with probability argument p and `bernoulli` with probability
330 argument q , thus connecting `bins n p` and `bins n q`.

331 Using mostly syntax-directed proofs, we establish the following two lemmata:

```

332 {-@ binsDistL :: p:Prob → {q:Prob | p ≤ q} → n:Nat
333         → { kdist distNat (bins n p) (gbins n p q) ≤ q - p } @-}
334
335 {-@ binsDistR :: p:Prob → {q:Prob | p ≤ q} → n:Nat
336         → { kdist distNat (gbins n p q) (bins n q) ≤ (n-1) * (q - p) } @-}
337

```

338 Both proofs use the distance axioms of § 2.4. The proof of `binsDistR` (inductively) calls `binsDist`,
339 while the proof of `binsDistL` requires the below axiom for `bernoulli`’s distance.

```

340 {-@ assume bernoulliDist :: d:Dist Nat
341         → p:Prob → {q:Prob | p ≤ q}
342         → { kdist d (bernoulli p) (bernoulli q) ≤ p - q } @-}
343

```

| | module name | LoC |
|----|----------------------|-----|
| 1. | Data.Dist | 155 |
| 2. | Monad.PrM | 93 |
| 3. | Monad.PrM.TCB.Axioms | 43 |
| 4. | Monad.PrM.TCB.Dist | 62 |
| 5. | Monad.PrM.Theorems | 100 |
| | Total | 453 |

Table 1. Summary of safe-coupling library. **LoC** is commented lines of Haskell Code.

That is, the expected distance of two bernoulli distributions, is bounded by the distance of the bernoulli's arguments (as formalized in rule T-BERN of § 5; Figure 7).

We prove `binsDist` combining `binsDistL` and `binsDistR` with triangular inequality, which, as explained in § 3.2 is a property (concretely a field) of the `Dist` type. The proof goes by induction:

```

359 binsDist p q 0
360   = pureDist distNat 0 0
361 binsDist p q n
362   = dist d (bins n p) (bins n q) -- Step 1
363     ? trinequality d (bins n p) (gbins n p q) (bins n p) -- Defined in Sec 3.2
364   <= dist d (bins n p) (gbins n p q) + dist d (gbins n p q) (bins n q) -- Step 2
365     ? binsDistL n p q
366     ? binsDistR n p q
367   <= (q - p) + (n-1) * (q - p) -- Step 3
368   <= n * (q - p) -- Step 4
369   *** QED
370   where d = kant distNat
371

```

The base case is merely an application of the `pureDist` axiom. In the inductive case, we use the (in)equational reasoning proof combinators of [Handley et al. 2019]: `l ? j <= r` ensures `l` is not greater than `r` using the justification `? j` which is optional and `*** QED` concludes the proof. The first step is to start from the distance between `bins n p` and `bins n q`. Applying triangular inequality, in the second step, we split the distance using `gbins`. Next, we use the two helper lemmata to bound each of the two distances. Finally, using trivial (SMT-automated) arithmetic, we get the desired bound. We note that the lemma `binsDistR` inductively calls `binsDist` on a smaller `n`, so our proof is inductive, while Liquid Haskell is ensuring (mutually recursive) termination.

The `binsDist` example showcases that our framework can be used to machine-check sophisticated proofs, that require ghost proof objects. To evaluate the expressiveness of our framework, we used it to prove two classic properties of machine learning, probabilistic programs: convergence of TD (§ 4.1) and stability of SGD (§ 4.2).

3 IMPLEMENTATION OF safe-coupling

In this section we present `safe-coupling`, a Haskell library that exports an interface for probabilistic (executable) programming and permits relational probabilistic verification using Liquid Haskell (§ 3.1). Table 1 summarizes the five main modules of `safe-coupling` that *define* distance (§ 3.2) and the probabilistic monad (§ 3.3), *assume* relational (§ 3.4) and distance (§ 3.5) axioms, and *prove* relational theorems (§ 3.6). In section (§ 5), we formally justify the assumptions made by `safe-coupling`.

3.1 Liquid Haskell Preliminaries

Verification with Refinement Types. Refinement types are used to do “light” verification. For example `max` of two probabilities (*i.e.* doubles between 0 and 1) is also a probability:

```
max :: Prob → Prob → Prob
max x y = if x ≤ y then y else x
```

To achieve SMT decidable and automatic verification, a refinement type systems clearly separate the executable code (here, the Haskell definitions) from the logic (here, the predicates on the refinement types). In the `max` example, every caller of `max` knows the type signature (*i.e.* that the result is also `Prob`), but not its implementation (*i.e.* that it returns one of its arguments). This way verification is modular, but, by default, the Haskell function does not exist in the logical fragment.

User Defined Functions in the Logic. An attempt to refer to user defined definitions in the refinement predicates, will lead to an undefined error. For instance, below we define a proof that $\forall x y . x \leq \max x y$ as a function whose arguments encode the quantified `x` and `y` and its result is a unit refined with the desired predicate. (Notation: we simplify $\{v:() \mid p\}$ to $\{p\}$.)

```
{-@ not_found :: x:Prob → y:Prob → {x ≤ max x y} @-} -- ERROR: max is unknown
not_found _ _ = ()
```

Since refinement types clearly separate the executable code from the logic, the above specification leads to an error: `max` is unknown to the logic. There are two ways to lift executable definitions in the logic: 1) axiomatization and 2) reflection.

1) *Axiomatization* of `max` defines a logical *uninterpreted* function that has the same refinement type and returns the same result as the executable `max`, but `max`’s definition is not available in the logic. For example, one can use axiomatization to show $0 \leq \max x y \leq 1$ but not that $x \leq \max x y$:

```
{-@ axiomatize max @-}
{-@ ok    :: x:Prob → y:Prob → {0 ≤ max x y} @-} -- max's specification is known
{-@ error :: x:Prob → y:Prob → {x ≤ max x y} @-} -- max's definition is unknown
```

2) *Reflection* of `max` defines a `max` function in the logic and further makes its definition available:

```
{-@ reflect max @-}
{-@ theorem :: x:Prob → y:Prob → {x ≤ max x y} @-} -- max's definition is known
theorem _ _ = ()
```

Reflection of executable functions permits “deep verification”, *i.e.* reasoning about sophisticated properties like Kantorovich distance of two probabilistic runs. Yet, for decidable refinement type checking, this reasoning requires explicit (user-provided) proofs. Importantly (but not surprisingly) functions can only get reflected, when their definitions consist only of reflected or axiomatized functions. For example, in our setting, functions imported from an unverified Haskell library cannot get reflected.

Proof Terms. Liquid Haskell is using refinement types to encode theorems []. The definitions of these functions can be unit (then the theorem is trivially proved by SMT and existing automation) or can contain inductive calls and combine other proof terms using proof combinators. Usually, such definitions do not have runtime meaning: if executed they will not produce any interesting result. But, since Liquid Haskell is checking for totality and completeness of these definitions they encode mathematical proofs.

We call *theorem* a refinement type specification that has a proof term, like the `theorem` for `max` defined above. As an alternative, Liquid Haskell’s `assume` keyword let’s you assume *axioms*

(refinement types) that one cannot prove. We use such axioms to encode properties of axiomatized functions. For example, when `max` is axiomatized we can define two axioms that describe its behavior.

```

444 {-@ axiomatize max @-}
445 {-@ assume max1 :: x:Prob → y:Prob → {x ≤ max x y} @-}
446 {-@ assume max2 :: x:Prob → y:Prob → {y ≤ max x y} @-}

```

Since `max` is axiomatized, its definition is not available in the logic, so none of the above axioms can be proved.

3.2 Data.Dist: Definition of Distance

We used Liquid Haskell to define the refined data type `Dist a` that encodes a metric, as follows.

```

453 {-@ data Dist a = Dist {
454     dist      :: a → a → {v:Double | 0.0 ≤ v }
455     , identity :: x:a → {dist x x == 0}
456     , symmetry :: x:a → y:a → {dist x y = dist x y}
457     , trinequality :: x:a → y:a → z:a → {dist x z ≤ dist x y + dist y z}
458   } @-}

```

The first field of `Dist` contains the distance function `dist` on any expressions of type `a`: it is a function that given two arguments of type `a` returns a non negative `Double`. The next three fields capture the metric's axioms for identity of indiscernibles, symmetry, and triangle inequality.

In this module, we further defined the distance metric on doubles (`distDouble`) and natural numbers (`distNat`):

```

465 distDouble :: Dist Double
466 distNat    :: Dist Nat

```

These definitions contain both the definition of the function `dist` and the proofs of the metric axioms on the concrete distance. Further, we define a function that computes distance between two same-length lists of a given a `Dist a`:

```

471 {-@ dList :: Dist a → xs:[a] → ys:[a | len xs = len ys] → {v:Double | 0 ≤ v} @-}

```

We proved all the metric axioms of `dList`, yet, since there exists the same-length dependency it is not possible to define a (well-typed) `Dist [a]` function. Still, we use the above definitions to compute distance between natural numbers, doubles and their lists:

```

476 assert (dist distDouble 42.0 40.0 == 2.0)
477 assert (dList distDouble [42] [40.0] == 2.0)

```

3.3 Monad.PrM: Definition of the Probabilistic Monad

The module `Monad.PrM` is essentially a wrapper around an executable Haskell probability monad. We chose the probabilistic functional library `probability`, due to its clear interface. Our development uses `probability` to execute (and test) our probabilistic programs, but our mechanized proofs do not depend on it and could use alternative libraries (e.g. `monad-bayes` [Scibior et al. 2015]).

The underlying `probability` library (here prefixed as `PLib`) exports the type `T prop a`, that essentially maps each `a` to a probability `prop` and defines the monadic and probabilistic primitives.

The data type. Our probability monad type instantiates `prop` to the probability type `Prop`.

```

488 type PrM a = PLib.T Prop a

```

```

491 {-@ assume unifDist :: d:Dist a → xsl: [a] → xsr: { [a] | xsl == xsr }
492     → { kdist d (unif xsl) (unif xsr) == 0 } @-}
493
494 {-@ assume choiceDist :: d:Dist a → p:Prob → el:PrM a → ul:PrM a
495     → q:{Prob | p = q } → er:PrM a → ur:PrM a
496     → { kdist d (choice p el ul) (choice q er ur) ≤
497         p * (kdist d el er) + (1.0 - p) * (kdist d ul ur) } @-}
498
499 {-@ assume pureBindDist :: da:Dist a → db:Dist b → m:Double
500     → fl:(a → b) → el:PrM a
501     → fr:(a → b) → er:PrM a
502     → (xl:a → xr:a → {dist db (fl xl) (fr xr) - dist da xl xr ≤ m})
503     → { kdist db (el >>= (ppure . fl)) (er >>= (ppure . fr)) ≤
504         kdist da el er + m } @-}
505
506
507
508

```

Fig. 1. Distance Axioms of safe-coupling. Encoding rules T-UNIF, T-CHOICE, and T-BIND-RET of fig. 7.

Axiomatized primitives. For each monadic ($\gg=$ and pure) and probabilistic primitive (bernoulli , uniform , and choice) operations we used the probability functions to define the Haskell (executable) function and axiomatized (as described in § 3.1) them in the logic. For example, pure is defined as follows

```

513 {-@ axiomatize pure :: x:a → PrM {v:a | v = x} @-}
514 pure x = PLib.pure x
515

```

We followed this encoding for practical reasons: since PLib is not itself verified with Liquid Haskell (which is a challenging future work) its definitions are not available in the logic. Yet, this encoding leaves us the flexibility to axiomatize the primitives as desired (§ 3.4 and 3.5).

Reflected functions. Using the axiomatized primitives we defined further probabilistic functions, such as mapM .

```

522 {-@ reflect mapM :: (a → PrM b) → [a] → PrM [b] @-}
523

```

Since mapM is reflected, its definition (which is standard) is available in the logic and used to prove (relational) theorems about mapM (§ 3.6).

3.4 TCB.Axioms: Assumption of Relational Axioms

The first trusted computing base (TCB) of safe-coupling is called Axioms and contains the relational specification of each axiomatized primitive. It provides the Haskell axiomatized function (\diamond) and uses it to encode the relational axioms for pure , ($\gg=$), and bernoulli , as presented in § 2.2.

3.5 TCB.Dist: Assumption of Distance Specifications

The second trusted computing base (TCB) of safe-coupling is called Dist and contains the distance specification of each axiomatized primitive. It provides the Haskell function kant that (like \diamond) is axiomatized in the logic and for each distance on a returns the kantovich distance on distributions of a and (as defined in § 2.4) kdist that simply composes kant with dist :

```

537 {-@ axiomatize kant :: Dist a → Dist (PrM a) @-}
538 {-@ kdist :: Dist a → PrM a → PrM a → {d:Double | 0 ≤ d} @-}
539

```

540 This module provides the distance axioms for the axiomatized primitives. In § 2.4 we presented
 541 the axiomatization for `bind` (`bindDist`), `pure` (`pureDist`), and `bernoulli` (`bernoulliDist`). We assume
 542 three more axioms presented in Figure 1. First, `unifDist`, the distance axiom of `uniform`, states that
 543 the Kantorovich distance between two uniform distributions is zero, when the sampling input lists
 544 are equal. Second, `choiceDist`, the distance axiom for `choice`, states that the Kantorovich distance
 545 of two choice expressions `choice p e_l u_l` and `choice q e_r u_r` is p times the Kantorovich distance
 546 of e_l e_r and $1 - p$ times the Kantorovich distance of u_l u_r , when $p = q$. Finally, `pureBindDist` is a
 547 distance axiom for `bind`. For soundness reasons discussed in § 5, the rule is stated only for `bind`
 548 expressions whose second argument is (the monadic lifting of) a pure function. The axiom requires
 549 that the pure function f_l and f_r make the distance between two values grow by at most m ; in order
 550 words, the distance between $f_l x_l$ and $f_r x_r$ cannot exceed the distance between x_l and x_r and some
 551 fixed constant m . Under this assumption, the Kantorovich distance between $(e_l \gg= (\text{ppure} . f_l))$
 552 and $(e_r \gg= (\text{ppure} . f_r))$ is bounded by the Kantorovich distance between e_l and e_r plus m . This
 553 specialized axiom provide a means to upper bound the Kantorovich distance between two `bind`
 554 expressions as a function of the Kantorovich distance of their first arguments, and is instrumental
 555 for our case studies.

557 3.6 Theorems: Proof of Relational Properties

558 This module proves common theorems using our assumed TCB and the defined functions of
 559 `Monad.PrM`. Concretely, it provides a simplified version of the `bindAxiom` when the two `bind` argu-
 560 ments form a bijectional coupling and the a relational specification for the monadic map.

561 All the properties on this module are proved. Next, we provide a formalism that justifies the
 562 assumptions of our two TCB modules.

564 4 CASE STUDIES

565 To evaluate `safe-coupling` we used it to verify two classic machine learning properties conver-
 566 gence of TD (§ 4.1) and stability of SGD (§ 4.2). § 4.3 summarizes our results.

569 4.1 Case Study I: Convergence of TD(0)

570 Our first case study proves convergence for TD(0), a classical algorithm for Reinforcement Learning.

572 *4.1.1 Implementation of TD(0).* In the standard reinforcement learning setting, an agent (*i.e.* the
 573 learning algorithm) repeatedly interacts with the environment, a Markov Decision Process (MDP)
 574 with state space `State` and set of actions `A`. At each step, the MDP reacts to the agent's action by
 575 drawing a new random state and a numeric reward according to a function $t :: \text{State} \rightarrow \text{PrM}$
 576 $(\text{State}, \text{Double})$. The current state i of the process is known to the learner, but the exact function t
 577 is not. Given black-box access to t , the goal of the learner is to find a policy map $\pi :: \text{State} \rightarrow A$
 578 from the state space to the best available action from `A` that maximizes the learner's expected
 579 reward over infinite time.

580 Figure 2 presents the implementation of TD(0), a Temporal Difference (TD) learning algorithm
 581 that estimates the value function $v :: \text{State} \rightarrow \text{Reward}$ of the MDP, *i.e.* the expected reward at
 582 each state if the agent were to repeatedly act according to some assumed policy π . For simplicity
 583 of verification, we defined `State` as $\{s:\text{Nat} \mid s \leq n\}$ and functions on `State` as lists of length n .
 584 The TD learner (*i.e.* `td0`) takes as input the number of iterations n , a transition function t , and an
 585 estimate of v and it iterates through the n states. At each iteration i , the learner runs `sample` that
 586 draws a reward and transition (j, r) from the i th element of t . Then, the estimate $v j$ is updated
 587 by incorporating the observed reward r and the estimated value $v j$ of the new state. Estimated
 588

```

589 import Monad.PrM -- mapM defined here
590
591 td0 :: Int → [PrM (State, Reward)] → [Reward] → PrM [Reward]
592 td0 n t v = iterate n (act t) v
593
594 act :: [PrM (State, Reward)] → [Reward] → PrM [Reward]
595 act t v = mapM (sample t v) [0..length v]
596
597 sample :: [PrM (State, Reward)] → [Reward] → State → PrM Reward
598 sample t v i = do
599   (j, r) ← t !! i
600   pure (update v i j r)
601
602 iterate :: Int → (a → PrM a) → a → PrM a
603 iterate 0 _ x = pure x
604 iterate n f x = f x >>= iterate (n - 1) f
605
606 update :: [Reward] → State → State → Reward → Reward
607 update v i j r = (1 - α) * (v !! i) + α * (r + γ * v !! j)
608
609 type State = Int
610 type Reward = Double
611

```

Fig. 2. Implementation of TD(0).

rewards in the future are reduced by a discount factor $\gamma \in [0, 1)$. Higher γ allows v to converge faster.

4.1.2 Convergence for TD(0). Our goal is to show that td0 converges to a stationary distribution independently of its initial input. This can be achieved by proving that td0 is contractive on v . One potential approach would be to prove that for $k = \alpha * \gamma + (1-\alpha)$,

$$\{-@ \text{td0Goal} :: n:\text{Nat} \rightarrow t:[\text{PrM (State, Reward)}] \rightarrow v_l:[\text{Reward}] \rightarrow v_r:[\text{Reward}] \rightarrow \{k\text{dist dList (td0 n } v_l \text{ t) (td0 n } v_r \text{ t)} \leq k^n * (\text{dist dList } v_l \text{ } v_r)\} @-\}$$

Instead, we prove a stronger property that td0 is contractive for all possible outcomes (via lifting). To do so, first we defined a pure (to be lifted) predicate that bounds the distance (since Liquid Haskell does not permit lambdas in the refinements):

$$\begin{aligned} \text{bounded} &:: \text{Dist } a \rightarrow \text{Double} \rightarrow a \rightarrow a \rightarrow \text{Bool} \\ \text{bounded } d \ m \ v_1 \ v_2 &= \text{dist } d \ v_1 \ v_2 \leq m \end{aligned}$$

Using bounded we define the td0 specification as follows:

$$\{-@ \text{td0Spec} :: n:\text{Nat} \rightarrow t:[\text{PrM (State, Reward)}] \rightarrow v_l:[\text{Reward}] \rightarrow v_r:[\text{Reward}] \rightarrow \{\diamond (\text{bounded dList (} k^n * (\text{dist dList } v_l \text{ } v_r)) (\text{td0 n } v_l \text{ t) (td0 n } v_r \text{ t)}) @-\}$$

The td0Spec specification implies our original td0Goal . This is because a bound property on all outcomes implies an average bound:

$$\{-@ \text{bound} :: d:\text{Dist } a \rightarrow k:\text{Double} \rightarrow e_l:\text{PrM } a \rightarrow e_r:\text{PrM } a \rightarrow \{\diamond (\backslash x_l \ x_r \rightarrow \text{dist } d \ x_l \ x_r \leq k) e_l \ e_r \Rightarrow k\text{dist } d \ e_l \ e_r \leq k \} @-\}$$

```

638 type PDouble = {Double | 0 ≤ m}
639
640 {-@ iterateSpec :: m:PDouble → n:Nat → f:([Reward] → PrM [Reward])
641     → (m:PDouble → xl:[Reward] → xr:[Reward] →
642         { bounded dList m xl xr ⇒ ◇ (bounded dList (m * k)) (f xl) (f xr)}))
643     → rl:[Reward] → rr:[Reward]
644     → { bounded dList m rl rr ⇒ ◇ (bounded dList (m * kn)
645         (iterate n f rl) (iterate n f rr)) @-}
646
647 {-@ actSpec :: m:PDouble → t:[PrM (State, Reward)] → vl:[Reward] → vr:[Reward]
648     → {bounded dList m vl vr ⇒ ◇ (bounded dList (k * m)) (act t vl) (act t vr)} @-}
649
650
651 {-@ sampleSpec :: m:PDouble → t:[PrM (State, Reward)] → vl:[Reward] → vr:[Reward]
652     → i:State
653     → {bounded dList m vl vr ⇒ ◇ (bounded distDouble (k * m))
654         (sample t vl i) (sample t vr i)} @-}
655
656 {-@ updateSpec :: vl:[Reward] → vr:[Reward] → i:State → j:State → r:Reward
657     → {distD (update vl i j r) (update vr i j r) ≤
658         k * max (distD (vl !! i) (vr !! i)) (distD (vl !! j) (vr !! j))} @-}
659
660

```

Fig. 3. Relational Lemmas for the $td0Spec$ Proof; where $distD\ x\ y = dist\ distDouble\ x\ y$.

The reverse implication does not hold: two distributions can have a Kantorovich distance that is upper bounded by k , and do not satisfy the lifting of $bounded\ d\ k$. As a counterexample, assume $e_l = [(3, 0.5), (77, 0.5)]$ and $e_r = [(7, 0.5), (75, 0.5)]$, then for $k = 3$ and d the distance of natural numbers, the right side is true (*i.e.* $kdist\ distNat\ e_l\ e_r \leq 3$), but the left side does not hold. Thus, and to our surprise, we could prove a stronger property than originally anticipated, and in a simpler system with plain, non-quantitative, liftings.

4.1.3 Proof of Convergence for $TD(0)$. We proved $td0Spec$ in 128 lines of (Liquid) Haskell code and, as summarized in table 2, we used five lemmas; one for each used function. We named each lemma by postfixing the name of the function with *Spec*. The specification $mapMSpec$ comes from the *safe-coupling* library (§ 3.6), while the rest lemmas are presented in fig. 3.

The proof of each lemma, like $binsSpec$ of § 2.3, is following the structure of the function definition. Concretely, $td0Spec$ is proved by $iterateSpec$, using $actSpec$ as the proof requirement; $iterateSpec$ is proved by induction, using the pure and bind axioms; $actSpec$ is proved by $mapMSpec$, using $sampleSpec$ as the proof argument; $sampleSpec$ is using the axioms and $updateSpec$, which is proved using the linearity and triangular inequality of the distance on doubles.

In short, the great challenge was to come up with the correct invariant for $td0Spec$, after which the proof follows the structure of the $td0$'s implementation.

4.2 Case Study II: Stability of SGD

Supervised machine learning algorithms are algorithms that aim to select the best fitting model from a class of parametric models by iteratively refining some initial parameter, based on some training set. A good measure of the quality of these algorithms is their, so called, generalization

```

687 import Monad.PrM
688 import Data.Derivative (grad)
689
690 {-@ sgd :: zs:{[DataPoint] | 2 ≤ length zs} → w0:Weight → αs:[StepSize]
691     → f:(DataPoint → Weight → Double) → PrM Weight @-}
692 sgd _ w0 [] _ = pure w0
693 sgd zs w0 (α:αs) f
694   = choice (1 / natToD (length zs))
695     (pure (head zs) >>= sgdRec zs w0 αs f)
696     (unif (tail zs) >>= sgdRec zs w0 αs f)
697   where sgdRec zs w0 αs f z = do
698     w ← sgd zs w0 αs f
699     pure (update z α f w)
700
701 update :: DataPoint → StepSize → (DataPoint → Weight → Double)
702     → Weight → Weight
703 update z α f w = w - α * grad (f z) w
704
705 {-@ type StepSize = {v:Double | 0 ≤ v} @-}
706 type StepSize     = Double
707 type DataPoint    = (Double, Double)
708 type Weight       = Double
709
710

```

Fig. 4. Implementation of SGD.

error, which measures how they perform on previously unseen data. One sufficient condition for an algorithm to have a controlled generalization error is to be algorithmically stable [Bousquet and Elisseeff 2002]. Informally, a supervised machine learning algorithm is algorithmically stable if the output of the algorithm is not overly dependent on any single element in the training set. More formally, a supervised machine learning algorithm is ϵ -stable if the Kantorovich distance between the parameters obtained by running the algorithm on the same initial parameter and two adjacent training sets (i.e. training sets that differ in a single element) is upper bounded by ϵ . In this case study, we show that Stochastic Gradient Descent, the *de facto* backpropagation algorithm for deep learning, is algorithmically stable. The proof follows the steps of [Hardt et al. 2016].

4.2.1 Implementation of SGD. Figure 4 presents our `sgd` implementation, a variant of Stochastic Gradient Descent. The algorithm takes as input a training set `zs`, modeled as a list of data points, an initial weight `w0`, and a list of learning step sizes `αs` and loss function `f`. In the general setting, the loss function `f` in the definition of SGD would be a vector function $Z \times \mathbb{R}^d \rightarrow \mathbb{R}$, where Z is the data set and \mathbb{R}^d is the weight space. For the sake of simplicity, in our implementation `Weight` is defined as a single value of type `Double` which corresponds to $d = 1$.

The function `sgd` recursively computes a sequence of so-called weights (or classifiers) starting from the initial parameter `w0`, by updating at each step the current weight `w` into `update z α f w`, where `z` is sampled uniformly from data set `zs` and the learning step `α` represents the influence of each iteration in the final result.

In our implementation we unfold the definition of uniform sampling based on the operator `choice`. With probability $1 / \text{length } zs$ we sample the head element of the data set. Otherwise,

we sample z from the tail part. In both cases, we proceed by updating the result of a recursive call with respect to z .

To implement update, we assume a partial function $\text{grad} :: (\text{Weight} \rightarrow \text{Double}) \rightarrow \text{Weight} \rightarrow \text{Double}$ which computes the gradient of f . Our proof does not rely on grad 's definition, therefore it can be imported from any automatic differentiation library.

4.2.2 Stability of SGD. The stability statement bounds Kantorovich distance of two runs of sgd on data sets z_{s_l} and z_{s_r} which differ in exactly one element. Without loss of generality, we assume that the differing element is the first. We express the stability statement as a refinement type specification, as follows.

```

{-@ sgdDist :: d:Dist Double → αs:[StepSize] → f:(DataPoint → Weight → Double)
    → zsl:[DataPoint] → zsr::{[DataPoint] | tail zsl = tail zsr}
    → wsl:Weight → wsr:Weight
    → { kdist d (sgd zsl wsl αs f) (sgd zsr wsr αs f)
      ≤ dist d wsl wsr + estab (length zsl) αs } @-}

estab :: Nat → [StepSize] → Double
estab l αs = 2.0 * lip / natToD l * sum αs

```

That is, the Kantorovich distance between two runs of sgd is bounded by the distance of the different weights plus an ϵ , defined by the helper function estab . Note that estab does not depend on the different inputs of sgd , but depends on lip , which represents the Lipschitz constant L and is axiomatized in our proof.

4.2.3 Assumptions. Proofs of algorithmic stability are traditionally based on strong assumptions on the loss function $f : Z \times \mathbb{R}^d \rightarrow \mathbb{R}$, namely:

- (1) *L-Lipschitz*: $\|f(z, w_1) - f(z, w_2)\| \leq L\|w_1 - w_2\|$;
- (2) *Convex*: $f(z, w_1) \geq f(z, w_2) + \langle \nabla f(z, w_2), w_1 - w_2 \rangle$; and
- (3) *β -smooth*: $\|\nabla f(z, w_1) - \nabla f(z, w_2)\| \leq \beta\|w_1 - w_2\|$ and for all step sizes $\alpha, \alpha\beta < 1$.

where $\|\cdot\|$ and $\langle \cdot \rangle$ are the norm and the scalar product on \mathbb{R}^d respectively.

Instead of directly encoding these assumptions, which require advanced mathematical machinery that is not readily available in Liquid Haskell, we assume two properties of the update function. These properties follow from the assumptions on the loss function and are presented in fig. 5.

The contractiveness axiom for update `contractive` states that the distance of update on the same data point is equal to the distance of the weights. The boundedness of the difference of gradients bounded states that the distance of update on different data points is equal to the distance of the weights plus $2 * \text{lip} * \alpha$.

4.2.4 Proof of Stability of SGD. Using the assumptions of update, we proved `sgdDist` in 157 lines of Liquid Haskell code. The proof proceeds by induction on n . We use the structure of the implementation to distinguish between the case where the two algorithms sample the same element, for which we apply the contractiveness property of update, and the case where the two algorithms sample the element in which we datasets differ, for which we apply boundedness of the difference of gradients. The `choiceDist` axioms then guarantees that Kantorovich distance increases by $2 * \text{lip} * \alpha / n$ at each iteration, from which we conclude by induction. Other than `choiceDist`, our proof is using the `pureDist` and `pureBindDist` axioms, respectively in the base and inductive case.


```

785 {-@ assume contractive :: d:Dist Double → α:StepSize
786                       → f:(DataPoint → Weight → Double)
787                       → z:DataPoint
788                       → wl:Weight → wr:Weight
789                       → {dist d (update z α f wl) (update z α f wr)
790                          = dist d w w'} @-}
791
792 {-@ assume bounded    :: d:Dist Double → α:StepSize
793                       → f:(DataPoint → Weight → Double)
794                       → zl:DataPoint → zr:DataPoint
795                       → wl:Weight → wr:Weight
796                       → {dist d (update zl α f wl) (update zr α f wr)
797                          = dist d wl wr + 2 * lip * α} @-}
798
799
800
801

```

Fig. 5. Assumptions of the sgdDist Proof.

| Case study | LoC | Proof LoC | Lemmas | Axioms | Time (sec) |
|-------------------|-----|-----------|--------|--------|------------|
| bins Spec (§ 2.3) | 21 | 28 | 3 | 0 | 8 |
| bins Dist (§ 2.5) | | 143 | 7 | 0 | 113 |
| td0 (§ 4.1) | 31 | 128 | 5 | 0 | 24 |
| sgd (§ 4.2) | 26 | 157 | 3 | 2 | 46 |
| Total | 78 | 456 | 18 | 2 | 191 |

Table 2. Summary of case studies. **LoC** is lines of executable Haskell code to define the implementations. **Proof LoC** is lines of commented Liquid Haskell code to define refinement type specifications and their proofs. **Lemmas** is the number of total lemmas proved. **Axioms** is the assumed specifications. **Time** is verification time in seconds on a 2,8 GHz Dual-Core Intel Core i5, RAM 8 GB 1600 MHz DDR3.

4.3 Quantitative Summary

Table 2 summarizes the effort required to verify the three examples that we presented through the paper: bins, td0, and sgd. In total, we used 456 lines of proof code, *i.e.* commented (Liquid) Haskell lines that express refinement types specifications and their proofs, to verify 78 lines of executable Haskell code, giving an executable-to proof-ratio of almost 6, which is high but expected, given that our proofs are extrinsic. Most of our proofs directly follow the structure of the definitions, thus are easy, once the proper specification is set. The exception is the distance proof for bins (§ 2.5) which required the construction of a “ghost” proof distribution, providing evidence that such sophisticated proofs are feasible in our framework. The same proof is an outlier for our verification time: it requires almost 2 minutes, while the rest of the proofs need less than 1 minute. We note that for td0 we distributed the proof over multiple Haskell modules to allow fast and interactive proof development, since Liquid Haskell verifies per-module and provides local proof-error messages. All our proofs are submitted as Anonymized Supplementary Material.

5 PROOF SYSTEM

To justify the axioms of our implementation, in this section, we define λ^{RP} , a core probabilistic λ -calculus, with a set of relational proof rules. Although each proof rule of the relational program

| | | | |
|-----|---------------------|---|-------------------------|
| 834 | Types | $t, s ::= \text{nat} \mid \text{bool} \mid \text{list } t$ | <i>Discrete Types</i> |
| 835 | | $\mid \text{real} \mid t \rightarrow t \mid \text{prM } t$ | <i>Continuous Types</i> |
| 836 | Constants | $c ::= \text{nil} \mid \text{cons} \mid n \in \mathbb{N} \mid a \in \mathbb{R}^+ \mid +\infty \mid \text{true} \mid \text{false}$ | |
| 837 | | $\mid + \mid - \mid \cdot \mid / \mid = \mid < \mid \wedge \mid \vee \mid \Rightarrow \mid \neg$ | |
| 838 | Expressions | $e, u ::= \text{unif } e \mid \text{bern } x \mid \text{choice } x e e$ | <i>Probabilistic</i> |
| 839 | | $\mid \text{ret } e \mid \text{bind } e e$ | <i>Monadic</i> |
| 840 | | $\mid \lambda x. e \mid e e \mid c \mid x$ | <i>Pure</i> |
| 841 | | $\mid \text{case } e \text{ of } \{\text{nil} \rightarrow e; \text{cons } x x \rightarrow e\}$ | |
| 842 | | $\mid \text{let } x = e \text{ in } e$ | |
| 843 | Assertions | $p ::= e \mid p[\wedge, \vee, \Rightarrow]p \mid \neg p$ | <i>Boolean</i> |
| 844 | | $\mid \forall x : t. p \mid \exists x : t. p$ | <i>Quantifiers</i> |
| 845 | | $\mid \diamond_e p e e$ | <i>Lifting</i> |
| 846 | Environments | | |
| 847 | Typing | $\Gamma ::= \emptyset \mid x:t, \Gamma$ | |
| 848 | Predicate | $\Phi ::= \emptyset \mid p, \Phi$ | |

Fig. 6. Syntax of λ^{RP} . (Variables include r_l, r_r , and d .)

logic is encoded independently from others as an axiom in Liquid Haskell, we follow the same style of presentation as Aguirre et al. [2017] and treat our set of proof rules as a proof system. This treatment is primarily motivated by our desire to prove a crisp statement for soundness. We also note that for the purpose of establishing soundness, we only consider discrete distributions.

This section is organized as follows: we define the syntax (§ 5.1) and type system (§ 5.2) of λ^{RP} ; then, the axioms (§ 5.3) and the proof rules (§ 5.4) of our logic; and finally, the denotational semantics of λ^{RP} (§ 5.5) which we use to show soundness (§ 5.6).

5.1 Syntax

We consider a typed probabilistic λ -calculus with algebraic datatypes and distributions (fig. 6). Types are built from base types using the usual function space constructor and type constructors for lists (`list`) and probability distributions (`prM`; which encodes the Haskell type `PrM` of § 3.3).

λ^{RP} features a rich set of *constants* that include natural (n) and real (a) numbers, the special constant $+\infty$, the `true` and `false` booleans, the `nil` and `cons` list constructors. Furthermore, λ^{RP} features constants for arithmetic operations, boolean operations, equality, and inequality.

Variables in λ^{RP} include three special variables d , r_l , and r_r that respectively model distance and the left and right result of computations.

Expressions are built from constants and variables using the standard constructions from λ -calculus and monadic constructions. The former include constructions for lambda abstraction ($\lambda x. e$), application ($e e$), case analysis (`case e of {nil → e; cons x x → e}`), and structurally recursive definitions (`let x = t in e`). The latter include the probabilistic primitives `unif e` that probabilistically returns an element of its input list e , with uniform distribution, `bern x` that returns 1 with probability x , otherwise 0, and `choice x e1 e2` that returns the distribution e_1 with probability x and e_2 with probability $1 - x$. These primitives model our implementation interface (§ 3.3).

Assertions include (arbitrary, but boolean typed) expressions of λ^{RP} and boolean operators. To allow reduction to RHOL, assertions also include quantifiers even though our system does not explicitly use them. Finally, λ^{RP} has the lift assertion $\diamond_k p e_l e_r$ that encodes the combination of (\diamond) and `kdist` of our implementation (§ 3). Here k is a real typed expression, p is a relational assertion that depends two arguments of type t_l and t_l respectively, and e_l and e_r are probability distributions

883 over t_l and t_r respectively. The assertion ensures that p holds for the distribution coupling, and
 884 that the Kantorovich distance between the two distributions is bounded by k .

885 We adopt standard conventions, e.g. $g \cdot f$ stands for $\lambda x. g (f x)$. By abuse of notation, we write
 886 $\text{ret } \cdot e$ as shorthand for $\lambda x. \text{ret } (f x)$ and $\diamond_k p$ as syntactic sugar for $\diamond_k (\lambda r_l r_r. p) r_l r_r$, i.e. when the
 887 lifting happens on the special variables r_l and r_r . By convention, we also write $\diamond p$ as a shorthand
 888 for $\diamond_{+\infty} p$. As usual, we also let $e[e_x/x]$ denote the capture-free substitution of e_x for x in e .

889 5.2 Type system

891 We equip our language with a simple type system which serves three purposes: first, it ensures that
 892 expressions respect the type signatures of operators; second, it ensures that recursive definitions
 893 are structurally terminating. Our logic is agnostic to the mechanism used to enforce structural
 894 termination, so we leave this mechanism abstract. Finally, our type system restricts the use of
 895 distribution types to discrete types, so that types and expressions of our language can be given a
 896 set-theoretic interpretation—we discuss the case of continuous distributions in the § 7.

897 5.3 Axioms

899 The special variable d encodes distance that we assume satisfies the axioms of an (extended) metric:

900 *Definition 5.1 (Metric Axioms).* For every type t and every $x, y, z : t$:

- 901 (1) *Identity:* $d(x, y) = 0 \Leftrightarrow x = y$.
 902 (2) *Symmetry:* $d(x, y) = d(y, x)$.
 903 (3) *Triangular Inequality:* $d(x, z) \leq d(x, y) + d(y, z)$.

904 In addition, we assume that the distance d is defined for all types of λ^{RP} : For nat and real is
 905 defined as $d(x, y) = |x - y|$ and further satisfies linearity. For booleans it is defined as $d(x, y) = 1$ if
 906 $x \neq y$, i.e. it coincides with the discrete distance on booleans. For lists of equal length, we assume
 907 the distance is the maximum of distances between elements at the same positions. When the
 908 length is different, the distance is infinite. For functions, it is defined as the maximum distance over
 909 all function domains and for distributions, as the Kantorovich distance (of Villani [2009]). These
 910 assumptions are required for soundness.

911 5.4 Proof System

912 Our proof system uses two judgments to decide logical implication and relational typing. The first
 913 judgment $\Gamma; \Phi \vdash p$ states that the assertion p is valid under the assumptions of Φ , where all the free
 914 variables of Φ and p appear in Γ . This first judgment is similar to [Aguirre et al. 2017] and the proof
 915 rules are thus omitted.

916 The second judgment $\Gamma; \Phi \vdash e_l : t_l \sim e_r : t_r \mid p$ states that the expressions e_l and e_r , respectively of
 917 types t_l and t_r under Γ , satisfy the predicate p , under the assumptions of Φ . To encode this property
 918 the predicate p might refer to two special variables r_l and r_r (i.e. not bound in the environment Γ)
 919 that respectively refer to the expressions e_l and e_r . For example $\emptyset; \emptyset \vdash 1 : \text{nat} \sim 2 : \text{nat} \mid r_l < r_r$
 920 holds, since $1 < 2$. In general, the relational typing means that $\Gamma; \Phi \vdash p[e_l/r_l][e_r/r_r]$.

921 *Contexts.* Contexts are pairs consisting of a typing environment (Γ) that maps variables to their
 922 types and a logical environment (Φ) consisting of assertions.

923 *Proof Rules.* Figure 7 defines selected proof rules $\Gamma; \Phi \vdash e_l : \text{prM } t_l \sim e_r : \text{prM } t_r \mid p$. The proof
 924 rules are given for monadic expressions, in which case the relational assertions are of the form $\diamond_k p$.
 925 Informally, these assertions state that 1) the two related expressions satisfy the lifted predicate p
 926 and 2) the Kantorovich distance between the two expressions is upper bounded by k .

Relational Probabilistic Typing

$$\boxed{\Gamma; \Phi \vdash e_l : \text{prM } t_l \sim e_r : \text{prM } t_r \mid \diamond_k p}$$

$$\begin{array}{c}
\text{T-BOT} \\
\frac{}{\Gamma; \Phi \vdash e_l : \text{prM } t_l \sim e_r : \text{prM } t_r \mid \diamond_{+\infty} \text{true}} \\
\text{T-WEAKEN} \\
\frac{r_l : t_l, r_r : t_r, \Gamma; \Phi \vdash p \Rightarrow p_w \quad r_l : t_l, r_r : t_r, \Gamma; \Phi \vdash k \leq k_w \quad \Gamma; \Phi \vdash e_l : \text{prM } t_l \sim e_r : \text{prM } t_r \mid \diamond_k p}{\Gamma; \Phi \vdash e_l : \text{prM } t_l \sim e_r : \text{prM } t_r \mid \diamond_{k_w} p_w} \\
\text{T-UNIF} \\
\frac{\Gamma; \Phi \vdash e_l : \text{list } t \sim e_r : \text{list } t \mid r_l = r_r}{\Gamma; \Phi \vdash \text{unif } e_l : \text{prM } t \sim \text{unif } e_r : \text{prM } t \mid \diamond_0 r_l = r_r} \\
\text{T-BERN} \\
\frac{\Gamma; \Phi \vdash x_l : \text{real} \sim x_r : \text{real} \mid 0 \leq r_l \leq r_r \leq 1}{\Gamma; \Phi \vdash \text{bern } x_l : \text{prM } \text{nat} \sim \text{bern } x_r : \text{prM } \text{nat} \mid \diamond_{|x_r - x_l|} 0 \leq r_l \leq r_r \leq 1} \\
\text{T-CHOICE} \\
\frac{\Gamma; \Phi \vdash x_l : \text{real} \sim x_r : \text{real} \mid r_l = r_r \quad \Gamma; \Phi \vdash e_l : \text{prM } t_l \sim e_r : \text{prM } t_r \mid \diamond_{k_e} p \quad \Gamma; \Phi \vdash u_l : \text{prM } t_l \sim u_r : \text{prM } t_r \mid \diamond_{k_u} p}{\Gamma; \Phi \vdash \text{choice } x_l e_l u_l : \text{prM } t_l \sim \text{choice } x_r e_r u_r : \text{prM } t_r \mid \diamond_{x_l \cdot k_e + (1 - x_l) \cdot k_u} p} \\
\text{T-RET} \\
\frac{\Gamma; \Phi \vdash e_l : t \sim e_r : t \mid p \wedge d(r_l, r_r) \leq k}{\Gamma; \Phi \vdash \text{ret } e_l : \text{prM } t \sim \text{ret } e_r : \text{prM } t \mid \diamond_k p} \\
\text{T-BIND} \\
\frac{\Gamma; \Phi \vdash e_l : \text{prM } s_l \sim e_r : \text{prM } s_r \mid \diamond_q \quad x_l : s_l, x_r : s_r, \Gamma; q[x_l/r_l][x_r/r_r], \Phi \vdash u_l x_l : \text{prM } t_l \sim u_r x_r : \text{prM } t_r \mid \diamond_k p}{\Gamma; \Phi \vdash \text{bind } e_l u_l : \text{prM } t_l \sim \text{bind } e_r u_r : \text{prM } t_r \mid \diamond_k p} \\
\text{T-BIND-RET} \\
\frac{x_l : s_l, x_r : s_r, \Gamma; q[x_l/r_l][x_r/r_r], \Phi \vdash f_l x_l : t_l \sim f_r x_r : t_r \mid p \wedge d(r_l, r_r) \leq m \cdot d(x_l, x_r) + k_u \quad \Gamma; \Phi \vdash e_l : \text{prM } s_l \sim e_r : \text{prM } s_r \mid \diamond_{k_e} q}{\Gamma; \Phi \vdash \text{bind } e_l (\text{ret} \cdot f_l) : \text{prM } t_l \sim \text{bind } e_r (\text{ret} \cdot f_r) : \text{prM } t_r \mid \diamond_{m \cdot k_e + k_u} p}
\end{array}$$

Fig. 7. Typing of λ^{RP} , where k ranges over distance (real typed expressions). We use the syntactic sugar $g \cdot f \doteq \lambda x. g (f x)$ and $\diamond_k p \doteq \diamond_k (\lambda r_l r_r. p) r_l r_r$.

Rules T-BOT and T-WEAKEN are non-syntax-directed rules that can be applied to any probabilistic expression. Rule T-BOT states that any two probabilistic expressions $e_l \sim e_r$ satisfy the (lifted) true predicate and have expected distance bounded by $+\infty$. Rule T-WEAKEN weakens the lifted predicate $\diamond_k p$ to $\diamond_{k_w} p_w$, provided $p \Rightarrow p_w$ and $k \leq k_w$. These two requirements are established by pure logic.

The rules T-UNIF, T-BERN, and T-CHOICE are used to relate (the same) probabilistic primitives. Rule T-UNIF states that the uniform distributions of two *semantically* equal lists produce equal values and have zero expected distance. To semantically equate two (pure) lists e_l and e_r we use (pure) relational typing. The lifted predicate $\diamond_0 r_l = r_r$ ensures that the two distributions are the same and accordingly their expected distance is zero. We note that more general rules exist, but are omitted here since they are not used in our examples.

The rule T-BERN states that the bernoulli distribution $\text{bern } x_l$ dominates $\text{bern } x_r$, when $x_l \leq x_r$. The condition is expressed by the pure predicate $r_l \leq r_r$ in the premise, while the conclusion of the rule lifts the same predicate to encode dominance. Finally, it ensures that the distance of the two distributions is bounded by $|x_r - x_l|$.

The rule T-CHOICE relates two choice expressions $\text{choice } x_i e_i u_i$ (with i being l or r). To do so, it requires that x_l and x_r are equal, and that the pairs of distributions e_l and e_r and respectively u_l and u_r satisfy the same lifted predicate $\diamond p$ and respectively have distances k_e and k_u . It then ensures that choice will also satisfy the predicate $\diamond p$, while the distance is $x_l \cdot k_e + (1 - x_l) \cdot k_u$.

The rules T-RET and T-BIND are used to relate (the same) monadic primitives. The rule T-RET relates $\text{ret } e_l$ with $\text{ret } e_r$. Using pure relational typing, it requires that e_l and e_r satisfy the predicates p (in which the relational variables can freely appear) and their distance is bounded by k . Note that bounding distance in the pure setting is encoded as a logical statement; while a weakening rule can bring the premise in the required syntactic form, potentially with infinite distance. The rule concludes that the return expressions satisfy the lifted p (since the relational variables are not monadic) and their distance is bounded by k . The rule T-BIND relates two expressions $\text{bind } e_i u_i$ (with i being l or r). In the first premise, it assumes that e_l and e_r satisfy some lifted predicate q . In the second premise, the predicate p is assumed in the predicate environment to check the application $u_i x_i$ where x_i is the value of the probabilistic argument e_i , *i.e.* satisfies the predicate q . The predicate and distance of the bind expressions are the same as these of the second premise.

The final rule T-BIND-RET relates two bind expressions where the second argument is a pure function, composed with ret (we use the syntactic sugar for composition $g \cdot f \doteq \lambda x. g (f x)$). Such expressions could be typed by the rule T-BIND, but this special case permits more precise reasoning and is used critically by our case studies. The first premise of the rule is the same as of the T-BIND, but the second premise is now using pure relational typing to bound the distance of f 's result as a function of the distance of its input, *i.e.* $d(r_l, r_r) \leq m \cdot d(x_l, x_r) + k_u$, so the distance of the conclusion ($m \cdot k_e + k_u$) depends on the distance of the e arguments (k_e).

We conclude this section with a brief, high-level, explanation for the need of having two rules for bind. In contrast with other existing quantitative relational logics, such as those used for differential privacy [Barthe et al. 2012], there is no “obvious” composition rule for bind when considering the Kantorovich metric. Specifically, given two distributions e_l and e_r whose Kantorovich distance is upper bounded by k , there is no obvious condition on continuations f_l and f_r , such that applying these continuations results in two distributions whose Kantorovich distance is upper bounded by k and some additional argument. At a high level, the problem arises from the fact that lifted assertions guarantee that the Kantorovich (*i.e.* average) distance between the two distributions is upper bounded by k , whereas assumptions in context would require that there exists a coupling such that the distance between e_l and e_r is bounded by k . Since they are complex and not needed for our purposes, we leave the study of more general rules for future work.

5.5 Denotational Semantics

Here we define a denotational semantics of λ^{RP} by defining the denotations of types and typing environments (§ 5.5.1); expressions (§ 5.5.2); and assertions and predicate environments (§ 5.5.3). Our denotational semantics only considers discrete distributions, so that expressions have a straightforward set-theoretic interpretation.

5.5.1 Denotations of Types and Typing Environments. Definition 5.3 inductively defines the denotations of types. The interesting case is prM that gives a probability distribution, as per definition 5.2.

Definition 5.2 (Discrete Probability Distribution). A probability distribution over a discrete set C is a function $\mu : C \rightarrow [0, 1]$ such that $\sum_{x \in \text{supp } \mu} \mu(x) = 1$, where the support of μ is defined as $\text{supp } \mu \doteq \{x \mid x \in C \wedge \mu(x) \neq 0\}$. We denote the set of discrete distributions over C as $D(C)$.

Definition 5.3 (Denotations of Types). For each type t , $\llbracket t \rrbracket$ is inductively defined as follows:

$$\begin{aligned} \llbracket \text{bool} \rrbracket &\doteq \mathbb{B} & \llbracket \text{nat} \rrbracket &\doteq \mathbb{N} & \llbracket \text{list } t \rrbracket &\doteq \text{list}_{\llbracket t \rrbracket} \\ \llbracket \text{prM } t \rrbracket &\doteq D(\llbracket t \rrbracket) & \llbracket \text{real} \rrbracket &\doteq \mathbb{R}^+ & \llbracket t_x \rightarrow t \rrbracket &\doteq \llbracket t_x \rrbracket \rightarrow \llbracket t \rrbracket \end{aligned}$$

Using the denotations of types, definition 5.4 defines the denotation of a typing environment Γ as a set of *models* ρ that map each binder ($x : t$) in Γ to an element in the denotation of t .

Definition 5.4 (Denotation of Type Environments). $\llbracket \Gamma \rrbracket \doteq \{\rho \mid \forall (x : t) \in \Gamma. \rho(x) \in \llbracket t \rrbracket\}$

5.5.2 *Denotations of Expressions.* The denotation of an expression e is defined for a fixed model ρ as $\llbracket e \rrbracket_\rho$ in definition 5.5. The denotations of the pure fragment are standard, where the model is used to define the definition of a variable as $\rho(x)$ and we skip the verbose definitions for case and let. The probability and return cases use standard distributions, while monadic bind maps to distribution composition, defined in definition 5.6.

Definition 5.5 (Denotations of Expressions). For each expression e and model ρ , $\llbracket e \rrbracket_\rho$ is defined as:

$$\begin{array}{ll} \llbracket \text{unif } e \rrbracket_\rho & \doteq U_{\llbracket e \rrbracket_\rho} & \llbracket x \rrbracket_\rho & \doteq \rho(x) \\ \llbracket \text{bern } x \rrbracket_\rho & \doteq B_{\llbracket x \rrbracket_\rho} & \llbracket c \rrbracket_\rho & \doteq c \\ \llbracket \text{choice } x \ e \ u \rrbracket_\rho & \doteq \llbracket x \rrbracket_\rho \cdot \llbracket e \rrbracket_\rho + (1 - \llbracket x \rrbracket_\rho) \cdot \llbracket u \rrbracket_\rho & \llbracket \lambda x. e \rrbracket_\rho & \doteq \lambda x. \llbracket e \rrbracket_\rho \\ \llbracket \text{bind } e \ u \rrbracket_\rho & \doteq \text{scomp } \llbracket e \rrbracket_\rho \llbracket u \rrbracket_\rho & \llbracket e \ u \rrbracket_\rho & \doteq \llbracket e \rrbracket_\rho \llbracket u \rrbracket_\rho \\ \llbracket \text{ret } e \rrbracket_\rho & \doteq \delta_{\llbracket e \rrbracket_\rho} \end{array}$$

where δ_x represents the Dirac distribution at x , U_{xs} represents the uniform distribution over a non-repeating list xs , and B_p represents the p -biased Bernoulli distribution on $\{0, 1\}$.

Definition 5.6 (Sequential Composition Distribution). Let $\mu \in D(C)$ and $f : C \rightarrow D(C_2)$. Sequential composition distribution $\text{scomp } \mu f$ is defined as:

$$\text{scomp } \mu f(y) \doteq \sum_{x \in C} \mu(x) \cdot f(x)(y)$$

5.5.3 *Denotations of Assertions and Predicate Environments.* Finally, we inductively define denotations of assertions in Definition 5.10. Most of the cases are standard except from the case for lifting that relies on Kantorovich couplings (also known as expectation couplings), which we define in definition 5.9 and, in turn, relies on basic definitions of expectation and marginals.

Definition 5.7 (Expectation). For all $\mu \in D(C)$ and $f : C \rightarrow \mathbb{R}$, the expected value $E_{x \leftarrow \mu}[f(x)]$ (or $E_\mu[f]$) of f is a partial function defined as:

$$E_{x \leftarrow \mu}[f(x)] = \sum_{x \in C} \mu(x) \cdot f(x)$$

Definition 5.8 (Marginals). For all $\mu \in D(C_1 \times C_2)$, the first and second marginals of μ are respectively the distributions $\pi_1(\mu) \in D(C_1)$ and $\pi_2(\mu) \in D(C_2)$ defined as:

$$\pi_1(\mu)(x) \doteq \sum_{y \in C_2} \mu(x, y) \quad \pi_2(\mu)(y) \doteq \sum_{x \in C_1} \mu(x, y)$$

Intuitively, the marginals are the “projections” of the couplings, *i.e.* “dependent products” over distributions. Using these, we define Kantorovich coupling, *i.e.* the conditions that render $\diamond_k R$ valid.

Definition 5.9 (Kantorovich coupling). For all $\mu_1 \in D(C_1)$, $\mu_2 \in D(C_2)$, $R \subseteq C_1 \times C_2$, $d : C_1 \times C_2 \rightarrow \mathbb{R}^+$, and $k \in \mathbb{R}^+$, $\models \diamond_k R(\mu_1, \mu_2)$ iff there exists $\mu \in D(A_1 \times A_2)$ such that:

$$(1) \pi_1(\mu) = \mu_1 \text{ and } \pi_2(\mu) = \mu_2 \quad (2) \text{supp}(\mu) \subseteq R \quad (3) E_{(x,y) \leftarrow \mu}[d(x, y)] \leq k$$

As usual, $\models p$, means logical validity of p . The first clause corresponds to the standard definition of coupling. The second clause corresponds to the definition of R -coupling, and is connected to lifting by the following equivalence: $\diamond R$ iff there exists an R -coupling. The last clause is specific to Kantorovich coupling, and states that the expected distance with respect to the definition μ is upper bounded by k . Since the Kantorovich distance corresponds to the minimum expected distance over all possible couplings, it follows that k is an upper bound for the Kantorovich distance.

We use the Kantorovich coupling to define the denotation of λ^{RP} 's lifted predicate, while the rest of the denotations are standard.

1079 *Definition 5.10 (Denotation of Assertions).* For each assertion p and model ρ , $\llbracket p \rrbracket_\rho$ is defined as:

$$\begin{array}{ll}
 1080 \quad \llbracket e \rrbracket_\rho & \doteq \llbracket e \rrbracket_\rho & \llbracket p_1[\wedge, \vee, \Rightarrow] p_2 \rrbracket_\rho & \doteq \llbracket p_1 \rrbracket_\rho[\wedge, \vee, \Rightarrow] \llbracket p_2 \rrbracket_\rho \\
 1081 \quad \llbracket \neg p \rrbracket_\rho & \doteq \neg \llbracket p \rrbracket_\rho & \llbracket \forall x : t. p \rrbracket_\rho & \doteq \forall x. \llbracket p \rrbracket_\rho \\
 1082 \quad \llbracket \diamond_k p \ e \ u \rrbracket_\rho & \doteq \diamond_{\llbracket k \rrbracket_\rho} \llbracket p \rrbracket_\rho(\llbracket e \rrbracket_\rho, \llbracket u \rrbracket_\rho) & \llbracket \exists x : t. p \rrbracket_\rho & \doteq \exists x. \llbracket p \rrbracket_\rho
 \end{array}$$

1084 The denotation of a predicate environment is the conjunction of the denotation of all its predicates.

1085 *Definition 5.11 (Denotation of Predicate Environment).* $\llbracket \Phi \rrbracket_\rho \doteq \bigwedge_{p \in \Phi} \llbracket p \rrbracket_\rho$

1087 5.6 Soundness

1088 Our proof system is sound with respect to its denotational semantics. That is, for every model
 1089 of the typing environment that renders the predicate environment valid, the denotation of the
 1090 predicate, with the special variables r_l and r_r substituted by the typed expressions, is valid.

1092 **THEOREM 5.12 (SOUNDNESS).** *If $\Gamma; \Phi \vdash e_l : t_l \sim e_r : t_r \mid p$, then for every $\rho \in (\Gamma)$ such that $\models \llbracket \Phi \rrbracket_\rho$,
 1093 we have $\models \llbracket p[\llbracket e_l \rrbracket_\rho / r_l][\llbracket e_r \rrbracket_\rho / r_r] \rrbracket_\rho$.*

1094 In the appendix we prove soundness of the monadic rules, while soundness for the pure fragment
 1095 follows from Aguirre et al. [2017].

1096 As a corollary of soundness and definition 5.9, we get soundness of the probabilistic fragment.
 1097 Informally, our judgement relates e_l and e_r when there exists a coupling μ of e_l and e_r such that
 1098 predicate p holds for all samples with non-zero probability and the expected distance between
 1099 samples is less than k .

1101 **COROLLARY 5.13 (SOUNDNESS OF PROB. FRAGMENT).** *If $\Gamma; \Phi \vdash e_l : \text{prM } t_l \sim e_r : \text{prM } t_r \mid \diamond_k p$, then
 1102 for every $\rho \in (\Gamma)$ such that $\models \llbracket \Phi \rrbracket_\rho$, there exists an $\mu \in D(\llbracket t_l \rrbracket \times \llbracket t_r \rrbracket)$ such that:*

- 1103 (1) $\pi_1(\mu) = \llbracket e_l \rrbracket_\rho$ and $\pi_2(\mu) = \llbracket e_r \rrbracket_\rho$;
- 1104 (2) $\text{supp}(\mu) \subseteq \{(x_l, x_r) \mid x_l \in \llbracket t_l \rrbracket, x_r \in \llbracket t_r \rrbracket, \models \llbracket p \rrbracket_\rho \ x_l \ x_r\}$; and
- 1105 (3) $E_{(x,y) \leftarrow \mu}[d(x,y)] \leq \llbracket k \rrbracket_\rho$.

1107 5.7 Continuous distributions

1108 Our denotational semantics and soundness claim are established for the fragment of the language
 1109 with discrete distributions. We stress that, while it is possible to give a denotational semantics for
 1110 continuous distributions using Quasi-Borel Spaces (QBS) [Heunen et al. 2017; Vákár et al. 2019].
 1111 However, prior work in [Aguirre et al. 2021a] suggests that it may be challenging to give a sound
 1112 denotational semantics for relational preexpectations for a higher-order language.

1114 6 RELATED WORK

1115 *First Order, Imperative, Probabilistic Languages.* There is a large body of work that builds and
 1116 applies program logics to reason about probabilistic programs. Many of these works are based on
 1117 first-order imperative languages. Broadly speaking, there exists two main lines of work: the first line
 1118 of work focuses on non-relational properties, and can be traced back to early work by Kozen [Kozen
 1119 1985] and Morgan and McIver [Morgan et al. 1996]. Many of these works focus on establishing
 1120 sound foundations, and have not been implemented in practice. There are however, some noticeable
 1121 exceptions [Hölzl 2016; Hurd 2003], including an active line of work in mechanizing or automating
 1122 proofs of expected cost of probabilistic programs [Avanzini et al. 2020; Ngo et al. 2018; Tassarotti
 1123 and Harper 2018].

1124 The second line of work focuses on relational properties; this line of work has been initiated
 1125 in [Barthe et al. 2009], and its mechanization in the Coq proof assistant or as a self-standing proof
 1126 assistant [Barthe et al. 2011] have been used to verify formally concrete security of cryptographic

1128 constructions. Similar approaches have been developed by the FCF [Petcher and Morrisett 2015] and
 1129 CryptHOL[Basin et al. 2020]. These logics have been further extended to reason about differential
 1130 privacy [Barthe et al. 2012] and expected sensitivity [Barthe et al. 2018]. Our work is most closely
 1131 related to the latter, which introduces the notion of expectation coupling and formally verifies
 1132 algorithmic stability of SGD. More recently, Wang et al [Wang et al. 2020] develop an alternative
 1133 formalism which supports a richer class of termination behaviors. Our work is also closely related
 1134 to [Aguirre et al. 2021b] in which the authors develop a relational weakest precondition calculus to
 1135 reason about expected sensitivity. Their calculus avoids some of the peculiarities of [Barthe et al.
 1136 2018], and is used to verify our two main examples, but is not implemented. We also note that
 1137 their proof of convergence is significantly different and uses Kantorovich couplings rather than the
 1138 simpler, vanilla couplings used in our proof.

1139
 1140 *Relational Reasoning for Higher Order Languages.* Nanevski et al. [2011] were among the first
 1141 to explore relational reasoning for higher-order languages. Their work defines Relational Hoare
 1142 Type Theory (RHTT), a powerful program logic for proving relational properties of stateful higher-
 1143 order programs. RHTT is implemented in the Coq proof assistant, and used to verify intricate
 1144 information flow properties. The main difference with our work is that RHTT operates on a
 1145 shallow embedding of programs and does not support probabilistic reasoning. BiRelCost [Çiçek
 1146 et al. 2019] is a bidirectional type checker that automatically performs relational and unary cost
 1147 analysis requiring minimal user annotations. However, the checker is incomplete and relies on
 1148 example-driven heuristics. It is unlikely that the used heuristics would suffice to prove our case
 1149 studies. Handley et al. [2019], like us, address this incompleteness by encoding the relational rules
 1150 in Liquid Haskell and requiring explicit, user-provided, extrinsic proofs; sacrificing automation in
 1151 the name of expressiveness and predictability. Our work applies the technique of Handley et al.
 1152 [2019] to reason about probabilistic programs and quantitative specifications.

1153
 1154 *Verification of Higher Order Probabilistic Programs.* There have been two lines of work that
 1155 verify relational properties on executable probabilistic programs in F^* [Swamy et al. 2016]. First,
 1156 rF^* [Barthe et al. 2014] is an extension of F^* that supports relational reasoning via relational
 1157 refinement types. As with F^* , type checking generates SMT queries that are discharged by Z3.
 1158 Although rF^* supports probabilistic programs, reasoning is constrained by syntax-directed typing,
 1159 for example the `binsDist` of § 2.5 cannot be type-checked. We overcome this limitation by support-
 1160 ing a richer set of axioms and extrinsic proofs. Second, Grimm et al. [2018] present an alternative
 1161 approach to reason about relational properties in F^* that, similar to our work, uses extrinsic proofs
 1162 to reason about programs. Their approach supports probabilistic reasoning, and is used to prove
 1163 probabilistic non-interference of Shannon’s classic cryptographic one-time pad. However, their
 1164 support for probabilistic reasoning is limited and quantitative specifications are not supported.

1165
 1166 *Proof Systems of Higher Order Probabilistic Programs.* Aguirre et al. [2021a] present several
 1167 proof systems for higher-order stateful probabilistic programs. One key novelty of their proof
 1168 systems is the support of interactive adversarial computations. Such rules can be used to prove
 1169 that programs verify relational properties for all possible (well-typed) adversarial computations.
 1170 Unfortunately, these rules are only proved sound for boolean relational properties and it is an open
 1171 problem whether these rules remain sound for quantitative properties modeled using Kantorovich
 1172 lifting. Their approach builds on Relational Higher-Order Logic (RHOL) [Aguirre et al. 2017],
 1173 which achieves full expressiveness and bypasses limitations of syntax-directed reasoning via an
 1174 embedding into HOL. However, there is no implementation of RHOL and of its extensions.

1175
 1176

1177 *Differential privacy*. Differential privacy is a mathematical notion of privacy that can be un-
1178 derstood as a form of probabilistic sensitivity w.r.t. some pseudo-distance induced by a specific
1179 f -divergence. As such, differential privacy is directly related to our work.

1180 Fuzz [Reed and Pierce 2010] was the first type-based, approach to verify distance of higher order
1181 programs, in the domain of differential privacy. Since introduced, Fuzz has been extended in various
1182 ways. DFuzz [Gaboardi et al. 2013] introduces recursion; Adaptive Fuzz [Winograd-Cort et al.
1183 2017] supports dynamic data analysis; Fuzzi [Zhang et al. 2019] extend Fuzz with APRHL [Barthe
1184 et al. 2012] to prove trusted primitives (like Laplace) Duet [Near et al. 2019] extends Fuzz with
1185 support for advanced variants of differential privacy via a dual type system; HOARE2 [Barthe et al.
1186 2015] provides a relational refinement system that embeds Fuzz; and Bunched Fuzz [june wunder
1187 et al. 2022] extends the system with bunches to, like us, reason about distances of probability
1188 distributions. Like our system, most of these approaches use typing rules to trace distance and have
1189 some support of probabilistic programs. However, they lack the flexibility to reason about expected
1190 sensitivity for most advanced examples such as those considered here.

1191 In a work most closely related to ours, DPella [Lobo-Vesga et al. 2021] and Solo [Abuah et al.
1192 2021] use Haskell’s dependent types to encode differential privacy. Both these systems are similar to
1193 ours since they do verify executable Haskell code. The critical difference is that they use Haskell’s
1194 dependent types while we use the refinement type extension of Liquid Haskell. One benefit of our
1195 approach is that we can delegate arithmetic reasoning to SMT. At a more general level, the two
1196 approaches are incomparable: we are able to reason finely about expected sensitivity, whereas they
1197 can reason about differential privacy using standard composition theorems, and about accuracy
1198 using a fine-grained approach that exploits a slick combination of information flow typing and
1199 concentration inequalities.

1200

1201

7 CONCLUSION & FUTURE WORK

1202 We have enhanced Liquid Haskell with support for reasoning about relational properties of prob-
1203 abilistic computations. We have also demonstrated that the resulting framework is sufficiently
1204 expressive for proving formally convergence and algorithmic stability properties of classic machine
1205 learning algorithms from the literature. An important direction for future work is to verify the
1206 implementation of our library and complement it with mechanisms to reason about non-relational,
1207 expectation based properties. This would allow us to reason about convergence to the optimal
1208 parameters, which is another property of concern for machine learning algorithms. Armed with
1209 the mechanisms developed here and mechanisms for non-relational expectation-based reasoning,
1210 it would be possible to develop a formally verified library of machine-learning algorithms.

1211 Another important direction for future work is to improve automation of relational proofs.
1212 The main challenge is to conceive a set of effective mechanisms to combine syntax-directed,
1213 synchronous, reasoning about the two programs, with syntax-directed, asynchronous reasoning
1214 about a single program (typically when the two programs do not follow the same control-flow),
1215 and non-syntax-directed reasoning, like our `binsDist` proof. One potential approach is to combine
1216 relational refinement type checking with the idea of “game-hopping” used in security proofs:
1217 informally, to help relational verification by introducing a sequence of intermediate programs, such
1218 that relational verification of two consecutive programs can be fully automated, and the overall
1219 verification goal follows directly from combining the results of these intermediate verifications.

1220

1221

ACKNOWLEDGMENTS

1222

REFERENCES

1223

1224

1225

Chike Abuah, David Darais, and Joseph P. Near. 2021. Solo: Enforcing Differential Privacy Without Fancy Types. *CoRR* abs/2105.01632 (2021). arXiv:2105.01632 <https://arxiv.org/abs/2105.01632>

- 1226 Alejandro Aguirre, Gilles Barthe, Marco Gaboardi, Deepak Garg, Shin-ya Katsumata, and Tetsuya Sato. 2021a. Higher-order
 1227 probabilistic adversarial computations: categorical semantics and program logics. *Proc. ACM Program. Lang.* 5, ICFP
 1228 (2021), 1–30. <https://doi.org/10.1145/3473598>
- 1229 Alejandro Aguirre, Gilles Barthe, Marco Gaboardi, Deepak Garg, and Pierre-Yves Strub. 2017. A relational logic for
 1230 higher-order programs. *Proc. ACM Program. Lang.* 1, ICFP (2017), 21:1–21:29. <https://doi.org/10.1145/3110265>
- 1231 Alejandro Aguirre, Gilles Barthe, Justin Hsu, Benjamin Lucien Kaminski, Joost-Pieter Katoen, and Christoph Matheja.
 1232 2021b. A pre-expectation calculus for probabilistic sensitivity. *Proc. ACM Program. Lang.* 5, POPL (2021), 1–28. <https://doi.org/10.1145/3434333>
- 1233 Martin Avanzini, Georg Moser, and Michael Schaper. 2020. A modular cost analysis for probabilistic programs. *Proc. ACM
 1234 Program. Lang.* 4, OOPSLA (2020), 172:1–172:30. <https://doi.org/10.1145/3428240>
- 1235 Gilles Barthe, Thomas Espitau, Benjamin Grégoire, Justin Hsu, and Pierre-Yves Strub. 2018. Proving expected sensitivity of
 1236 probabilistic programs. *Proc. ACM Program. Lang.* 2, POPL (2018), 57:1–57:29. <https://doi.org/10.1145/3158145>
- 1237 Gilles Barthe, Cédric Fournet, Benjamin Grégoire, Pierre-Yves Strub, Nikhil Swamy, and Santiago Zanella Béguelin. 2014.
 1238 Probabilistic relational verification for cryptographic implementations. In *The 41st Annual ACM SIGPLAN-SIGACT Sym-
 1239 posium on Principles of Programming Languages, POPL '14, San Diego, CA, USA, January 20-21, 2014*, Suresh Jagannathan
 1240 and Peter Sewell (Eds.). ACM, 193–206. <https://doi.org/10.1145/2535838.2535847>
- 1241 Gilles Barthe, Marco Gaboardi, Emilio Jesús Gallego Arias, Justin Hsu, Aaron Roth, and Pierre-Yves Strub. 2015. Higher-Order
 1242 Approximate Relational Refinement Types for Mechanism Design and Differential Privacy. In *Proceedings of the 42nd
 1243 Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January
 1244 15-17, 2015*, Sriram K. Rajamani and David Walker (Eds.). ACM, 55–68. <https://doi.org/10.1145/2676726.2677000>
- 1245 Gilles Barthe, Benjamin Grégoire, and Santiago Zanella Béguelin. 2009. Formal certification of code-based cryptographic
 1246 proofs. In *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL
 1247 2009, Savannah, GA, USA, January 21-23, 2009*, Zhong Shao and Benjamin C. Pierce (Eds.). ACM, 90–101. <https://doi.org/10.1145/1480881.1480894>
- 1248 Gilles Barthe, Benjamin Grégoire, Sylvain Heraud, and Santiago Zanella Béguelin. 2011. Computer-Aided Security Proofs
 1249 for the Working Cryptographer. In *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa
 1250 Barbara, CA, USA, August 14-18, 2011. Proceedings (Lecture Notes in Computer Science)*, Phillip Rogaway (Ed.), Vol. 6841.
 1251 Springer, 71–90. https://doi.org/10.1007/978-3-642-22792-9_5
- 1252 Gilles Barthe and Justin Hsu. 2020. Probabilistic couplings from program logics. *Foundations of Probabilistic Programming
 1253 (2020)*, 145.
- 1254 Gilles Barthe, Boris Köpf, Federico Olmedo, and Santiago Zanella Béguelin. 2012. Probabilistic relational reasoning for
 1255 differential privacy. In *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages,
 1256 POPL 2012, Philadelphia, Pennsylvania, USA, January 22-28, 2012*, John Field and Michael Hicks (Eds.). ACM, 97–110.
 1257 <https://doi.org/10.1145/2103656.2103670>
- 1258 David A. Basin, Andreas Lochbihler, and S. Reza Sefidgar. 2020. CryptHOL: Game-Based Proofs in Higher-Order Logic. *J.
 1259 Cryptol.* 33, 2 (2020), 494–566. <https://doi.org/10.1007/s00145-019-09341-z>
- 1260 Olivier Bousquet and André Elisseff. 2002. Stability and Generalization. 2 (2002), 499–526. [http://www.jmlr.org/papers/v2/
 1261 bousquet02a.html](http://www.jmlr.org/papers/v2/bousquet02a.html)
- 1262 Ezgi Çiçek, Weihao Qu, Gilles Barthe, Marco Gaboardi, and Deepak Garg. 2019. Bidirectional type checking for relational
 1263 properties. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation,
 1264 PLDI 2019, Phoenix, AZ, USA, June 22-26, 2019*, Kathryn S. McKinley and Kathleen Fisher (Eds.). ACM, 533–547. <https://doi.org/10.1145/3314221.3314603>
- 1265 Yuxin Deng. 2015. *Semantics of Probabilistic Processes: An Operational Approach*. <https://doi.org/10.1007/978-3-662-45198-4>
- 1266 Marco Gaboardi, Andreas Haeberlen, Justin Hsu, Arjun Narayan, and Benjamin C. Pierce. 2013. Linear Dependent Types for
 1267 Differential Privacy. In *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming
 1268 Languages (POPL '13)*. <https://doi.org/10.1145/2429069.2429113>
- 1269 Niklas Grimm, Kenji Maillard, Cédric Fournet, Catalin Hritcu, Matteo Maffei, Jonathan Protzenko, Tahina Ramananandro,
 1270 Aseem Rastogi, Nikhil Swamy, and Santiago Zanella Béguelin. 2018. A monadic framework for relational verification:
 1271 applied to information security, program equivalence, and optimizations. In *Proceedings of the 7th ACM SIGPLAN
 1272 International Conference on Certified Programs and Proofs, CPP 2018, Los Angeles, CA, USA, January 8-9, 2018*, June
 1273 Andronick and Amy P. Felty (Eds.). ACM, 130–145. <https://doi.org/10.1145/3167090>
- 1274 Jad Hamza, Nicolas Voirol, and Viktor Kunčák. 2019. System FR: Formalized Foundations for the Stainless Verifier. *Proc.
 1275 ACM Program. Lang.* 3, OOPSLA, Article 166 (oct 2019), 30 pages. <https://doi.org/10.1145/3360592>
- 1276 Martin A. T. Handley, Niki Vazou, and Graham Hutton. 2019. Liquidate Your Assets: Reasoning about Resource Usage in
 1277 Liquid Haskell. *Principles of Programming Languages* (2019). <https://doi.org/10.1145/3371092>
- 1278 Moritz Hardt, Ben Recht, and Yoram Singer. 2016. Train faster, generalize better: Stability of stochastic gradient descent,
 1279 Vol. 48. *JMLR.org*, 1225–1234. <http://jmlr.org/proceedings/papers/v48/hardt16.html>

- 1275 Chris Heunen, Ohad Kammar, Sam Staton, and Hongseok Yang. 2017. A convenient category for higher-order probability
1276 theory. In *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23,*
1277 *2017*. IEEE Computer Society, 1–12. <https://doi.org/10.1109/LICS.2017.8005137>
- 1278 Johannes Hölzl. 2016. Formalising Semantics for Expected Running Time of Probabilistic Programs. In *Interactive Theorem*
1279 *Proving - 7th International Conference, ITP 2016, Nancy, France, August 22-25, 2016, Proceedings (Lecture Notes in Computer*
1280 *Science)*, Jasmin Christian Blanchette and Stephan Merz (Eds.), Vol. 9807. Springer, 475–482. [https://doi.org/10.1007/978-](https://doi.org/10.1007/978-3-319-43144-4_30)
1281 [3-319-43144-4_30](https://doi.org/10.1007/978-3-319-43144-4_30)
- 1282 Joe Hurd. 2003. Verification of the Miller-Rabin probabilistic primality test. *J. Log. Algebraic Methods Program.* 56, 1-2 (2003),
1283 3–21. [https://doi.org/10.1016/S1567-8326\(02\)00065-6](https://doi.org/10.1016/S1567-8326(02)00065-6)
- 1284 june wunder, Arthur Azevedo de Amorim, Patrick Baillot, and Marco Gaboardi. 2022. Bunched Fuzz: Sensitivity for Vector
1285 Metrics. *arXiv:cs.PL/2202.01901*
- 1286 Dexter Kozen. 1985. A Probabilistic PDL. 30, 2 (1985), 162–178.
- 1287 Nico Lehmann, Rose Kunkel, Jordan Brown, Jean Yang, Niki Vazou, Nadia Polikarpova, Deian Stefan, and Ranjit Jhala. 2021.
1288 STORM: Refinement Types for Secure Web Applications. In *15th USENIX Symposium on Operating Systems Design and*
1289 *Implementation (OSDI 21)*. USENIX Association, 441–459. [https://www.usenix.org/conference/osdi21/presentation/](https://www.usenix.org/conference/osdi21/presentation/lehmann)
1290 [lehmann](https://www.usenix.org/conference/osdi21/presentation/lehmann)
- 1291 Torgny Lindvall. 2002. *Lectures on the coupling method*. Courier Corporation.
- 1292 Elisabet Lobo-Vesga, Alejandro Russo, and Marco Gaboardi. 2021. A Programming Language for Data Privacy with Accuracy
1293 Estimations. *ACM Trans. Program. Lang. Syst.* 43, 2, Article 6 (jun 2021), 42 pages. <https://doi.org/10.1145/3452096>
- 1294 Kenji Maillard, Catalin Hritcu, Exequiel Rivas, and Antoine Van Muylder. 2020. The next 700 relational program logics.
1295 *Proc. ACM Program. Lang.* 4, POPL (2020), 4:1–4:33. <https://doi.org/10.1145/3371072>
- 1296 Carroll Morgan, Annabelle McIver, and Karen Seidel. 1996. Probabilistic Predicate Transformers. 18, 3 (1996), 325–353.
- 1297 Aleksandar Nanevski, Anindya Banerjee, and Deepak Garg. 2011. Verification of Information Flow and Access Control
1298 Policies with Dependent Types. In *32nd IEEE Symposium on Security and Privacy, S&P 2011, 22-25 May 2011, Berkeley,*
1299 *California, USA*. 165–179.
- 1300 Joseph P. Near, David Darais, Chike Abua, Tim Stevens, Pranav Gaddamadugu, Lun Wang, Neel Somani, Mu Zhang,
1301 Nikhil Sharma, Alex Shan, and Dawn Song. 2019. Duet: An Expressive Higher-Order Language and Linear Type System
1302 for Statically Enforcing Differential Privacy. *Proc. ACM Program. Lang.* 3, OOPSLA, Article 172 (oct 2019), 30 pages.
1303 <https://doi.org/10.1145/3360598>
- 1304 Van Chan Ngo, Quentin Carbonneaux, and Jan Hoffmann. 2018. Bounded expectations: resource analysis for probabilistic
1305 programs. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation,*
1306 *PLDI 2018, Philadelphia, PA, USA, June 18-22, 2018*, Jeffrey S. Foster and Dan Grossman (Eds.). ACM, 496–512. <https://doi.org/10.1145/3192366.3192394>
- 1307 Adam Petcher and Greg Morrisett. 2015. The Foundational Cryptography Framework. In *Principles of Security and Trust -*
1308 *4th International Conference, POST 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software,*
1309 *ETAPS 2015, London, UK, April 11-18, 2015, Proceedings (Lecture Notes in Computer Science)*, Riccardo Focardi and Andrew C.
1310 Myers (Eds.), Vol. 9036. Springer, 53–72. https://doi.org/10.1007/978-3-662-46666-7_4
- 1311 Norman Ramsey and Avi Pfeffer. 2002. Stochastic lambda calculus and monads of probability distributions. In *Conference*
1312 *Record of POPL 2002: The 29th SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Portland, OR, USA,*
1313 *January 16-18, 2002*, John Launchbury and John C. Mitchell (Eds.). ACM, 154–165. <https://doi.org/10.1145/503272.503288>
- 1314 Jason Reed and Benjamin C. Pierce. 2010. Distance Makes the Types Grow Stronger: A Calculus for Differential Privacy. In
1315 *Proceedings of the 15th ACM SIGPLAN International Conference on Functional Programming (ICFP '10)*. <https://doi.org/10.1145/1863543.1863568>
- 1316 Adam Scibior, Zoubin Ghahramani, and Andrew D. Gordon. 2015. Practical Probabilistic Programming with Monads.
1317 *SIGPLAN Not.* 50, 12 (aug 2015), 165–176. <https://doi.org/10.1145/2887747.2804317>
- 1318 Nikhil Swamy, Catalin Hritcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargava,
1319 Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean-Karim Zinzindohoué, and Santiago Zanella-Béguelin.
1320 2016. Dependent Types and Multi-Monadic Effects in F*. In *43rd ACM SIGPLAN-SIGACT Symposium on Principles of*
1321 *Programming Languages (POPL)*. ACM. <https://www.fstar-lang.org/papers/mumon/>
- 1322 Joseph Tassarotti and Robert Harper. 2018. Verified Tail Bounds for Randomized Programs. In *Interactive Theorem Proving -*
1323 *9th International Conference, ITP 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 9-12,*
1324 *2018, Proceedings (Lecture Notes in Computer Science)*, Jeremy Avigad and Assia Mahboubi (Eds.), Vol. 10895. Springer,
1325 560–578. https://doi.org/10.1007/978-3-319-94821-8_33
- 1326 Hermann Thorisson. 2000. *Coupling, Stationarity, and Regeneration*.
- 1327 Matthijs Vákár, Ohad Kammar, and Sam Staton. 2019. A domain theory for statistical probabilistic programming. *Proc. ACM*
1328 *Program. Lang.* 3, POPL (2019), 36:1–36:29. <https://doi.org/10.1145/3290349>

- 1324 Niki Vazou, Joachim Breitner, Rose Kunkel, David Van Horn, and Graham Hutton. 2018. Theorem Proving for All: Equational
1325 Reasoning in Liquid Haskell (Functional Pearl). *SIGPLAN Not.* 53, 7 (sep 2018), 132–144. [https://doi.org/10.1145/3299711.](https://doi.org/10.1145/3299711.3242756)
1326 3242756
- 1327 Niki Vazou, Eric L. Seidel, and Ranjit Jhala. 2014a. LiquidHaskell: Experience with Refinement Types in the Real World. 49,
12 (sep 2014), 39–51. <https://doi.org/10.1145/2775050.2633366>
- 1328 Niki Vazou, Eric L. Seidel, Ranjit Jhala, Dimitrios Vytiniotis, and Simon Peyton-Jones. 2014b. Refinement Types for Haskell.
1329 In *Proceedings of the 19th ACM SIGPLAN International Conference on Functional Programming (ICFP '14)*. Association for
1330 Computing Machinery, New York, NY, USA, 269–282. <https://doi.org/10.1145/2628136.2628161>
- 1331 Niki Vazou, Anish Tondwalkar, Vikraman Choudhury, Ryan G. Scott, Ryan R. Newton, Philip Wadler, and Ranjit Jhala. 2017.
1332 Refinement Reflection: Complete Verification with SMT. *Proc. ACM Program. Lang.* POPL (2017). [https://doi.org/10.1145/](https://doi.org/10.1145/3158141)
1333 3158141
- 1334 Cédric Villani. 2008. *Optimal transport: old and new*.
- 1335 Cédric Villani. 2009. *Optimal Transport, old and new*. Springer. <https://link.springer.com/book/10.1007/978-3-540-71050-9>
- 1336 Peixin Wang, Hongfei Fu, Krishnendu Chatterjee, Yuxin Deng, and Ming Xu. 2020. Proving expected sensitivity of
1337 probabilistic programs with randomized variable-dependent termination time. *Proc. ACM Program. Lang.* 4, POPL (2020),
1338 25:1–25:30. <https://doi.org/10.1145/3371093>
- 1339 Daniel Winograd-Cort, Andreas Haeberlen, Aaron Roth, and Benjamin C. Pierce. 2017. A Framework for Adaptive Differential
1340 Privacy. *Proc. ACM Program. Lang.* ICFP, Article 10 (aug 2017), 29 pages. <https://doi.org/10.1145/3110254>
- 1341 Hengchu Zhang, Edo Roth, Andreas Haeberlen, Benjamin C. Pierce, and Aaron Roth. 2019. Fuzzi: A Three-Level Logic for
1342 Differential Privacy. *Proc. ACM Program. Lang.* 3, ICFP, Article 93 (jul 2019), 28 pages. <https://doi.org/10.1145/3341697>
- 1343
- 1344
- 1345
- 1346
- 1347
- 1348
- 1349
- 1350
- 1351
- 1352
- 1353
- 1354
- 1355
- 1356
- 1357
- 1358
- 1359
- 1360
- 1361
- 1362
- 1363
- 1364
- 1365
- 1366
- 1367
- 1368
- 1369
- 1370
- 1371
- 1372