Safe Approximation of Probabilities

Eugenio Moggi (University of Genova)
Walid Taha (Halmstad University)

July 26, 2018, Rostock, Germany

Introduction

Reachability analysis is a powerful tool:

- Applies to discrete, continuous, and hybrid systems
- Enables safety verification
- Validated implementations exist (e.g. VNODE, Acumen)

Key-enablers:

- Set-extension is well-defined
- IA provides a computable and correct over-approximation

Problem:

- Engineers describe safety in terms of probabilities and distributions - not just sets
- Can reachability analysis work in this context?

This Paper

Two key questions

- Does set-extension generalize, for example, to distributions?
- Do extensions also apply naturally to non-deterministic and probabilistic systems?

Contributions:

- A mathematical framework for answering such questions
- Leveraging the concepts of monads and monad-transformers

Key insights:

- Discrete probability distributions form a monad
- The non-empty power-set constructor can be turned into a monad-transformer (Caution! Axiom of choice)

Monads

In this work, we will work primarily in the category Set of sets.

Definition (Monad, c.f. Moggi [1])

A monad on the category Set of sets is a triple $(M, \eta, \_^*)$ such that if X:Set then MX:Set, $\eta{:}X \to MX$ is a map (from X to MX), if $f{:}X \to MY$ then $f^*{:}MX \to MY$. Furthermore, η and f^* satisfy the following equational axioms for any $f{:}X \to MY$ and $g{:}Y \to MZ$

1. $\eta_X^* = id_{MX}$
2. $(g^* \circ f)^* = g^* \circ f^*$
3. $f^* \circ \eta_X = f$

Trivial examples: The identity $MX = X$ and the terminal monad $MX = 1$, where 1 is a singleton set.

Monads

Examples:

- Exceptions $X + E$, where $+$ denotes disjoint union
- Powersets (non-determinism) $P(X)$, where $P(X)$ is the set of subsets of X, $\eta(x) = \{x\}$ and $f^*(A) = \bigcup_{x:A} f(x)$; also $P_+(X)$, i.e., $P(X)$ without the empty set, is a monad
- Probability Distributions
  $D(X) = \{p{:}X \to [0,1] | \sum_{x:X} p(x) = 1\}$,
  $\eta(x)(x') = 1$ if $x = x'$ else 0 and
  $f^*(p)(y) = \sum_{x:X} p(x) * f(x)(y)$.
  An equivalent monad is given by the set $D'(X)$ of measures, i.e., $\mu{:}P(X) \to [0,1]$ such that $\mu(X) = 1$ and $\mu(\biguplus_{i:I} A_i) = \sum_{i:I} \mu(A_i)$ for any family $(A_i | i{:}I)$ of disjoint subsets of X. The correspondence between $D(X)$ and $D'(X)$ is $\mu(A) = \sum_{x:A} p(x)$ and $p(x) = \mu(\{x\})$.

Interval for Probability Distributions

The natural order on $[0,1]$ induces a point-wise order on the function space $X \to [0,1]$. This allows to introduce interval notations for subsets of probability distributions in $D(X)$, for instance

$$[\ell, u] = \{p{:}D(X) | \forall x.\ell(x) \leq p(x) \leq u(x)\}$$

where $\ell, u{:}X \to [0,1]$ (not necessarily in $D(X)$).

Another notation is

$$[L, U] = \{p{:}D(X) | \forall (A, \ell_A){:}L.\ell_A \leq p(A) \wedge \forall (B, u_B){:}U.p(B) \leq u_B\}$$

where $L, U{:}(P(X), [0,1])^*$ are finite sequences and p is extended additively to subsets of X, namely $p(A) = \sum_{x:A} p(x)$.

M-Extension

The natural set-extension of $f{:}X \to Y$ is the map $P(f){:}P(X) \to P(Y)$ such that $P(f)(A) = \{f(x) | x{:}A\}$.

It satisfies $P(f)(\{x\}) = \{f(x)\}$, that is it is an extension of f.

Generalizes to any monad, and hinges on the fact that a monad is also a functor:

Definition (Functor)

A functor F on Set maps a set X to a set $F(X)$, and $f{:}X \to Y$ to $F(f){:}F(X) \to F(Y)$ so that

1. $F(id_X) = id_{F(X)}$
2. $F(g \circ f) = F(g) \circ F(f)$

M-Extension

We will also need just one more (standard) concept:

Definition (Natural Transformation)

A natural transformation τ from a functor F to a functor G is a family of maps $\tau_X{:}F(X) \to G(X)$ indexed by X:Set such that for any $f{:}X \to Y$ we have $\tau_Y \circ F(f) = G(f) \circ \tau_X$.

Prop (M-extension)

A monad becomes a functor by defining $M(f) = (\eta_Y \circ f)^*$ *for* $f{:}X \to Y$*, and* $\eta_X{:}X \to M(X)$ *becomes a natural transformation from the identity functor to* M*, i.e.,* $(Mf)(\eta_X(x)) = \eta_Y(f(x))$*.*

When M is the powerset monad P, one recovers as a special case the natural set-extension.

For almost every monad on Set the map η_X is injective, thus one can view X as a subset of $M(X)$ and $M(f)$ as an extension of f.

Application - Approximating Distributions

Given an approximation F of a function $f{:}X \to Y$ and a lower approximation L of a distribution $\mu{:}D'(X)$, we want to compute an approximation $[L', U']$ of $\mu' = D'(f)(\mu){:}D'(Y)$.

To define the algorithm that solves this problem we must first specify the type of approximations involved and the properties that they must satisfy.

Recall that the set-extension of $f{:}X \to Y$ is the map $P(f){:}P(X) \to P(Y)$ such that $P(f)(A) = \{f(x) | x{:}A\}$, and that $\mu' = D'(f)(\mu)$ means that $\mu'(B) = \mu(f^{-1}(B))$.

Application - Approximating Distributions

Inputs: An approximation F of f, namely a map $F{:}P(X) \to P(Y)$ such that $\forall A{:}P(X).P(f)(A) \subseteq F(A)$.

A lower approximation $L = [(A_i, \ell_i)|i{:}n]$ of μ, i.e., $\forall i{:}n.\ell_i \leq \mu(A_i)$, with $(A_i|i{:}n)$ partition of X (thus $\sum_{i:n} \ell_i \leq 1$).

Output: An approximation $[L', U']$ of $\mu' = (D'f)(\mu)$, namely two sequences $L', U'{:}(P(Y) \times [0,1])^*$ such that $\forall (B', l'){:}L'.l' \leq \mu'(B')$ and $\forall (B', u'){:}U'.\mu'(B') \leq u'$.

Application - Approximating Distributions

For convenience, we identify a natural number n with the set $\{i|i < n\}$ of its predecessors.

Algorithm:

1. For $I \subseteq n$, let $A_I = \uplus_{i:I} A_i$, $\ell_I = \sum_{i:I} \ell_i$ and $u_I = 1 - \ell_{I^c}$, where $I^c \subseteq n - I$ is the complement of I. (Note: $\ell_I \leq \mu(A_I) \leq u_I$ holds by the assumption on L, in other words from the lower approximation L we compute its *completion* $[L^\sigma, U^\sigma]$, where $L^\sigma = [(A_I, \ell_I)|I \subseteq n]$ and $U^\sigma = [(A_I, u_I)|I \subseteq n]$, which approximates the same probability distributions, but more explicitly.)
2. Let $B_I = F(A_I)$. (Note: Since $f(A_I) \subseteq F(A_I) = B_I$ we have $A_I \subseteq f^{-1}(B_I)$. Thus, $\ell_I \leq \mu'(B_I)$. Furthermore, $\mu'(B_I^c) \leq u_{I^c}$, as $f^{-1}(B_I^c) = (f^{-1}(B_I))^c \subseteq A_I^c = A_{I^c}$.)
3. $L' = [(B_I, \ell_I)|I \subseteq n]$ and $U' = [(B_I^c, u_{I^c})|I \subseteq n]$.

Monad Transformers

Many of the things we care about in this work are monads. A key question, then, is whether they compose.

If M and M' are monads, then $M' \circ M$ is a functor, $\eta'_{MX} \circ \eta_X{:}X \to M'(MX)$ is a natural transformation, but there is no canonical way to define f^* for $f{:}X \to M'(MY)$.
Unlike monads, monad transformers can be composed.

Definition (Monad Morphism)

A monad morphism is a natural transformation σ from a monad M to a monad M' such that :

$$\eta'_X(x) = \sigma_X(\eta_X(x)) \qquad (\sigma_Y \circ f)^{*'}(\sigma_X(c)) = \sigma_Y(f^*(c))$$

Monad Transformers

We write Mon for the category of monads and monad morphisms.

Definition (Monad Transformer)

A monad transformer consists of a functor T on Mon and a natural transformation $\eta^T_M{:}M \to T(M)$ from the identity functor on Mon to T.

Prop

If M is a monad, then $M(\_ + E)$ and $P_+(M(\_))$ are monads.

Hint Given $F{:}X \to P_+(MY)$, let $\Pi x{:}X.F(x)$ be the set of *choice maps* f st $\forall x{:}X.f(x){:}F(x)$, then $F^*(A) = \{f^*(c) | c{:}A \wedge f{:}\Pi x{:}X.F(x)\}$.

Proving that F^* satisfies the axioms for monads uses crucially the axiom of choice.

(Main) Related Work

Weichselberger [2] (Def 2.2) introduces R-probabilities, namely a pair of maps L and U from a σ-algebra (called σ-field in [2]) $\mathcal{A}$ on a sample space Ω, which bound the probability distributions on Ω, namely $\forall A{:}\mathcal{A}.L(A) \leq p(A) \leq U(A)$.

In this paper we work in a simplified setting: the space Ω is a set X, the σ-algebra $\mathcal{A}$ is the powerset $P(X)$, L and U are finite sequences $L = [(A_i, \ell_i)|i{:}m]$ and $U = [(B_j, u_j)|j{:}n]$ representing maps $L', U'{:}P(X) \to [0, 1]$, namely $L'(A) = \ell_i$ when $A = A_i$ otherwise 0, and $U'(B) = u_j$ when $B = B_j$ otherwise 1.

Conclusions

Specifics

- Set-extension generalizes to (discrete) distributions, and, in fact, to any monad
- Extensions apply at least to non-deterministic systems

More broadly

- Monads facilitate establishing well-definedness of extensions

Future work

- Applying to CDFs on the reals
- Establishing connection to existing implmenetation

Support: KK Foundation, ELLIIT network, and an US NSF (CPS)