Analysis of Probabilistic Systems

Boot camp Lecture 4: The logical characterization of bisimulation

Prakash Panangaden[1]

[1]School of Computer Science
McGill University

Fall 2016, Simons Institute

Outline

The definition

- A set of states S,
- a set of *labels* or *actions*, L or $\mathcal{A}$ and
- a transition relation $\subseteq S \times \mathcal{A} \times S$, usually written

$$\rightarrow_a \subseteq S \times S.$$

  The transitions could be indeterminate (nondeterministic).
- We write $s \xrightarrow{a} s'$ for $(s, s') \in \rightarrow_a$.

A simple example

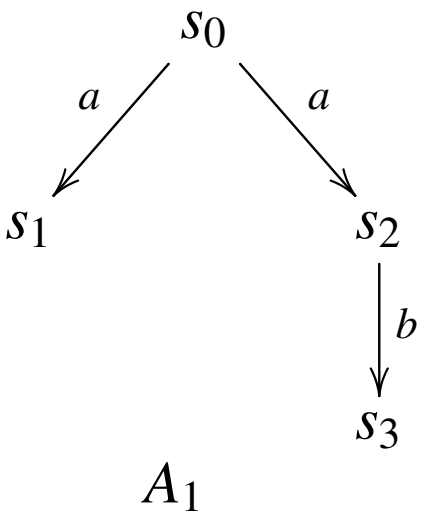

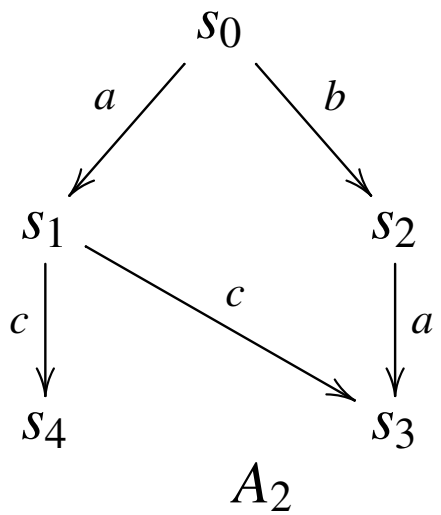

Bisimulation

s and t are states of a labelled transition system. We say s is **bisimilar** to t – written $s \sim t$ – if

$$s \xrightarrow{a} s' \Rightarrow \exists t' \text{ such that } t \xrightarrow{a} t' \text{ and } s' \sim t'$$

and

$$t \xrightarrow{a} t' \Rightarrow \exists s' \text{ such that } s \xrightarrow{a} s' \text{ and } s' \sim t'.$$

Bisimulation relations

- Define *a* (note the indefinite article) bisimulation relation R to be an equivalence relation on S such that
- $$sRt \text{ means } \forall a, s \xrightarrow{a} s' \Rightarrow \exists t', t \xrightarrow{a} t' \text{ with } s'Rt'$$

  and vice versa.
- This is not circular; it is a condition on R.
- We define $s \sim t$ if there is *some* bisimulation relation R with sRt.
- This is the version that is used most often.

What are Labelled Markov Processes?

- Labelled Markov processes are probabilistic versions of labelled transition systems. Labelled transition systems where the final state is governed by a probability distribution - no other indeterminacy.
- All probabilistic data is *internal* - no probabilities associated with environment behaviour.
- We observe the interactions - not the internal states.
- In general, the state space of a labelled Markov process may be a *continuum*.

Markov Kernels

- A *Markov kernel* is a function $h : S \times \Sigma \longrightarrow [0,1]$ with (a) $h(s,\cdot) : \Sigma \longrightarrow [0,1]$ a (sub)probability measure and (b) $h(\cdot, A) : X \longrightarrow [0,1]$ a measurable function.
- Though apparantly asymmetric, these are the probabilistic analogues of binary relations
- and the uncountable generalization of a matrix.

Formal Definition of LMPs

- An LMP is a tuple $(S, \Sigma, \mathsf{L}, \forall \alpha \in \mathsf{L}.\tau_\alpha)$ where
- (S, Σ) is an **analytic space** (what?, why?) with its Borel σ-algebra.
- and $\tau_\alpha : S \times \Sigma \longrightarrow [0, 1]$ the *transition* function is a Markov kernel
- $\forall s : S.\lambda A : \Sigma.\tau_\alpha(s, A)$ is a subprobability measure
  and
  $\forall A : \Sigma.\lambda s : S.\tau_\alpha(s, A)$ is a measurable function.

Probabilistic Bisimulation

Let $\mathcal{S} = (S, i, \Sigma, \tau)$ be a labelled Markov process. An equivalence relation R on S is a **bisimulation** if whenever sRs', with $s, s' \in S$, we have that for all $a \in \mathcal{A}$ and every R-closed *measurable* set $A \in \Sigma$, $\tau_a(s, A) = \tau_a(s', A)$.
Two states are **bisimilar** if they are related by a bisimulation relation.

Logical Characterization

- $$\mathcal{L} ::== \mathsf{T} | \phi_1 \wedge \phi_2 | \langle a \rangle_q \phi$$
- We say $s \models \langle a \rangle_q \phi$ iff

$$\exists A \in \Sigma . (\forall s' \in A . s' \models \phi) \wedge (\tau_a(s, A) > q).$$

- Two systems are bisimilar iff they obey the same formulas of $\mathcal{L}$. [Desharnais, Edalat, P. 1998 LICS, I and C 2002]

That cannot be right?

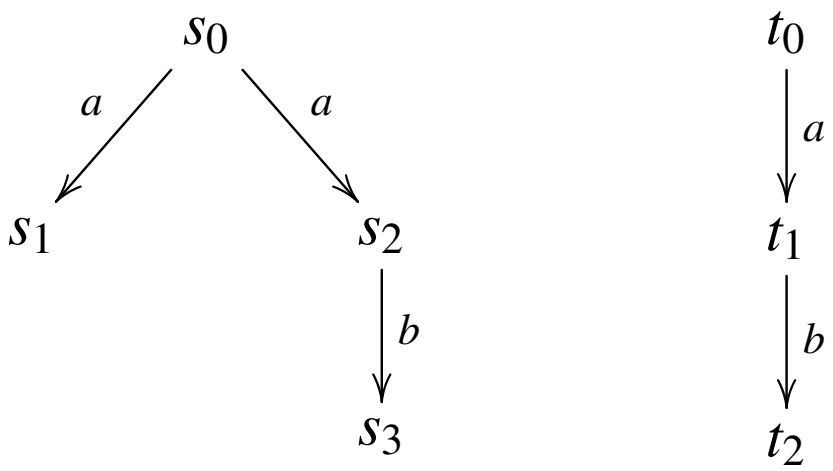

Two processes that cannot be distinguished without negation.
The formula that distinguishes them is $\langle a \rangle (\neg \langle b \rangle \top)$.

But it is!

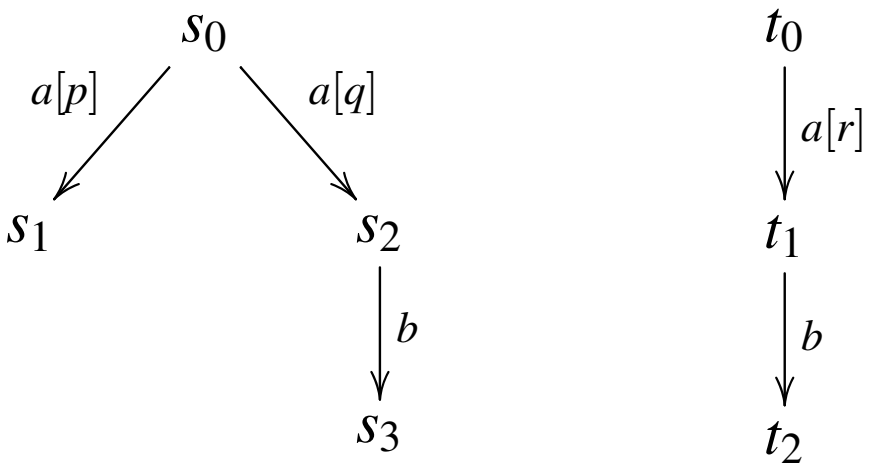

We add probabilities to the transitions.

- If $p + q < r$ or $p + q > r$ we can easily distinguish them.
- If $p + q = r$ and $p > 0$ then $q < r$ so $\langle a \rangle_r \langle b \rangle_1 \top$ distinguishes them.

Bisimulation implies logical equivalence

- Let R be a bisimulation relation on an LMP (S, Σ, τ_a). We prove by induction on ϕ that $\forall \phi \in \mathcal{L}\ \forall s, s' \in S. sRs' \Rightarrow s \models \phi \Leftrightarrow s' \models \phi$.
- Base case trivial.
- $\wedge$ is obvious from Inductive Hypothesis.
- For $\phi = \langle a \rangle_q \psi$ we have that $[\![\psi]\!]$ is R-closed from inductive hypothesis. Thus $\tau_a(s, [\![\psi]\!]) = \tau_a(s', [\![\psi]\!])$ and thus $sRs' \Rightarrow s \models \phi \Leftrightarrow s' \models \phi$.

Polish space

- A topological space is **separable** if it has a countable dense subset.
- A separable metric space has a countable base of open sets.
- A Polish space is the topological space underlying a complete separable metric space.
- Why did topology creep in?
- Measure theory works nicely on Polish spaces: *e.g.* the Borel sets of $X_1 \times X_2$ is the product σ-algebra of the Borel sets of X_1 and X_2 if they are Polish.
- The Gíry monad can be defined on Polish spaces.

Analytic space

- An analytic set A is the image of a Polish space X (or a Borel subset of X) under a continuous (or measurable) function $f : X \to Y$, where Y is Polish. If (S, Σ) is a measurable space where S is an analytic set in some ambient topological space and Σ is the Borel σ-algebra on S.
- Analytic sets do not form a σ-algebra but they are in the completion of the Borel algebra under **any** probability measure. [Universally measurable.]
- Regular conditional probability densities can be defined on analytic spaces.

Amazing Facts about Analytic Spaces

- Given S an analytic space and $\sim$ an equivalence relation such that there is a *countable* family of real-valued measurable functions $f_i : S \longrightarrow \mathbf{R}$ such that

$$\forall s, s' \in S. s \sim s' \iff \forall f_i. f_i(s) = f_i(s')$$

  then the quotient space (Q, Ω) - where $Q = S/\sim$ and Ω is the finest σ-algebra making the canonical surjection $q : S \longrightarrow Q$ measurable - is also analytic.
- If an analytic space (S, Σ) has a sub-σ-algebra Σ_0 of Σ which separates points and is countably generated then Σ_0 is Σ! The Unique Structure Theorem (UST).

Proof Sketch

- Show that the relation "s and s' satisfy exactly the same formulas" is a bisimulation.
- Can easily show that $\tau_a(s, A) = \tau_a(s', A)$ for A of the form $[\![\phi]\!]$.
- Use Dynkin' $\lambda - \pi$ theorem to show that we get a well defined measure on the σ-algebra generated by such sets and the above equality holds.
- Use special properties of analytic spaces to show that this σ-algebra is the same as the original σ-algebra.

The Quotient

- Given (S, Σ, τ_a) an LMP, we define $s \simeq s'$ if s and s' obey exactly the same formulas of $\mathcal{L}_0$.
- The functions $I_{[\![\phi]\!]} : S \longrightarrow \mathbf{R}$ defined by $I_{[\![\phi]\!]}(s) = 1$ if $s \models \phi$ and 0 otherwise are a countable family of measurable functions such that $s \simeq s'$ if and only if all the functions agree on s and s'. Thus the quotient space (Q, Ω) is analytic.
- We define an LMP (Q, Ω, ρ_a) where $\rho_a(t, U) := \tau_a(s, q^{-1}(U))$; $s \in q^{-1}(\{t\})$.

ρ is well defined - I

- Easy to check that $q^{-1}(q([\![\phi]\!])) = [\![\phi]\!]$:

  $s \in q^{-1}(q([\![\phi]\!]))$ implies that $q(s) \in q([\![\phi]\!])$, i.e. $\exists s' \in [\![\phi]\!].s \simeq s'$, so $s \models \phi$ so $s \in [\![\phi]\!]$. Yes, I know this is too small to read.
- Thus $q([\![\phi]\!])$ is measurable.
- Thus the σ-algebra generated -say, Λ - by $q([\![\phi]\!])$ is a sub-σ-algebra of Ω.
- Λ is countably generated and separates points so by UST it *is* Ω. Thus $q([\![\phi]\!])$ generates Ω.

ρ is well defined - II

- The collection $q([\![\phi]\!])$ is a π-system (because $\mathcal{L}_0$ has conjunction) and it generates Ω; thus if we can show that two measures agree on these sets they agree on all of Ω.
- If $q(s) = q(s') = t$ then $\tau_a(s, [\![\phi]\!]) = \tau_a(s', [\![\phi]\!])$ (simple interpolation).
- Thus $\tau_a(s, q^{-1}(q([\![\phi]\!]))) = \tau_a(s', q^{-1}(q([\![\phi]\!])))$ and hence ρ is well defined. We have $\rho_a(q(s), B) = \tau_a(s, q^{-1}(B))$.

Finishing the Argument

- Let X be any $\simeq$-closed subset of S.
- Then $q^{-1}(q(X)) = X$ and $q(X) \in \Omega$.
- If $s \simeq s'$ then $q(s) = q(s')$ and

$$\tau_a(s, X) = \tau_a(s, q^{-1}(q(X))) = \rho_a(q(s), q(X)) =$$

$$\rho_a(q(s'), q(X)) = \tau_a(s', q^{-1}(q(X))) = \tau_a(s', X).$$

Simulation

Let $\mathcal{S} = (S, \Sigma, \tau)$ be a labelled Markov process. A preorder R on S is a **simulation** if whenever sRs', we have that for all $a \in \mathcal{A}$ and every R-closed measurable set $A \in \Sigma$, $\tau_a(s, A) \leq \tau_a(s', A)$. We say s is simulated by s' if sRs' for some simulation relation R.

Logic for simulation?

- The logic used in the characterization has no negation, not even a limited negative construct.
- One can show that if s simulates s' then s satisfies all the formulas of $\mathcal{L}$ that s' satisfies.
- What about the converse?

Counter example!

In the following picture, t satisfies all formulas of $\mathcal{L}$ that s satisfies but t does not simulate s.

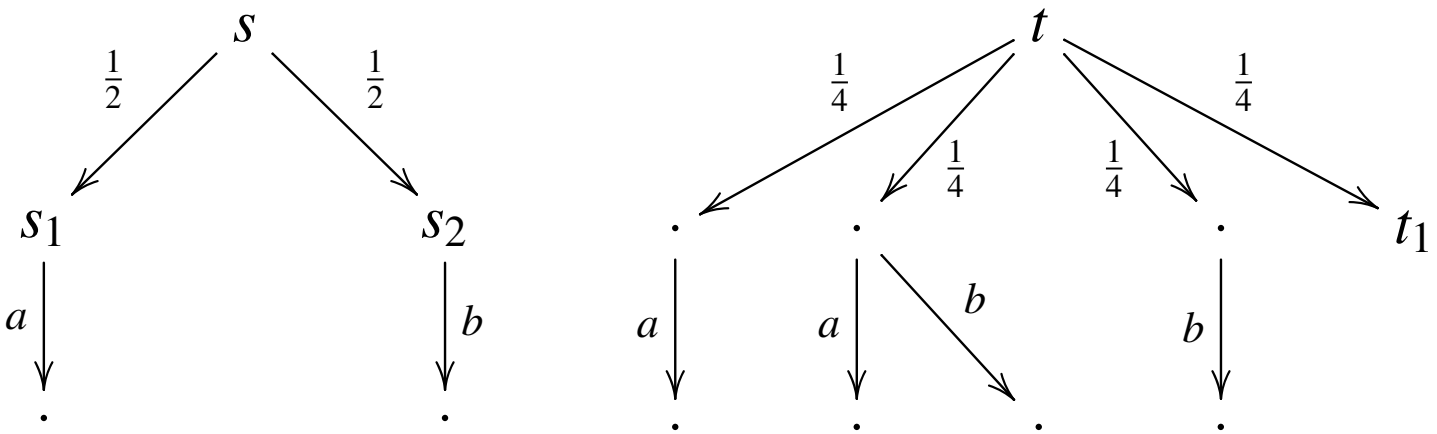

All transitions from s and t are labelled by a.

Counter example (contd.)

- A formula of $\mathcal{L}$ that is satisfied by t but not by s.

$$\langle a \rangle_0 (\langle a \rangle_0 \mathsf{T} \wedge \langle b \rangle_0 \mathsf{T}).$$

- A formula with disjunction that is satisfied by s but not by t:

$$\langle a \rangle_{\frac{3}{4}} (\langle a \rangle_0 \mathsf{T} \vee \langle b \rangle_0 \mathsf{T}).$$

A logical characterization for simulation

- The logic $\mathcal{L}$ does **not** characterize simulation. One needs disjunction.

$$\mathcal{L}_\vee := \mathcal{L} \mid \phi_1 \vee \phi_2.$$

- With this logic we have:
  An **LMP** s_1 simulates s_2 if and only if for every formula ϕ of $\mathcal{L}_\vee$ we have

$$s_1 \models \phi \Rightarrow s_2 \models \phi.$$

- The only proof we know uses domain theory.

Other Logics

$$
\begin{aligned}
\mathcal{L}_{\mathrm{Can}} &:= \mathcal{L}_0 \mid \mathrm{Can}(a) \\
\mathcal{L}_{\Delta} &:= \mathcal{L}_0 \mid \Delta_a \\
\mathcal{L}_{\neg} &:= \mathcal{L}_0 \mid \neg\phi \\
\mathcal{L}_{\vee} &:= \mathcal{L}_0 \mid \phi_1 \vee \phi_2 \\
\mathcal{L}_{\wedge} &:= \mathcal{L}_{\neg} \mid \bigwedge_{i \in \mathbf{N}} \phi_i
\end{aligned}
$$

where

$$
\begin{aligned}
s \models \mathrm{Can}(a) \quad & \text{to mean that } \tau_a(s, S) > 0; \\
s \models \Delta_a \quad & \text{to mean that } \tau_a(s, S) = 0.
\end{aligned}
$$

We need $\mathcal{L}_{\vee}$ to characterise simulation.

Conclusions

- Strong probabilistic bisimulation is characterised by a very simple modal logic with no negative constructs.
- There is a logical characterisation of simulation.
- There is a "metric" on LMPs which is based on this logic.
- Why did the proof require so many subtle properties of analytic spaces? There is a more general definition of bisimulation for which the logical characterisation proof is "easy" but to prove that that definition coincides with this one in analytic spaces requires roughly the same proof as that given here.