

THE DUALITY OF STATE AND OBSERVATIONS

Monica Dinculescu

Supervised by Prakash Panangaden and Doina Precup

School of Computer Science, McGill University

Introduction

Partially observable Markov decision processes (POMDPs) are an important model used in the AI literature for planning under uncertainty. They differ from MDPs in that one has only partial information, through so-called observations, about the state.

A few years ago, Chris Hundt, Prakash Panangaden, Doina Precup and Joelle Pineau discovered a duality construction for automata that interchanges the role of state and observations. They extended this to probabilistic systems, including POMDPs, and showed that the so called double dual gives a minimal representation of the behaviour of the original system.

In this work we discuss the problem of comparing the behavioural equivalence of partially observable systems with observations. We examine different types of equivalence relations induced by this duality, and discuss their inter-relations.

Hierarchy of Equivalence Relations

- A **test** t is a sequence of actions followed by a single observation. For example, $a_1 a_2 a_3 \omega$ is a test.
- An **experiment** e is a non-empty sequence of tests, $e = t_1 t_2 \dots t_n$.
- The probability of observing a test in a given state, denoted by $\langle s|t \rangle$ is the probability of observing the observations in the test/experiment at appropriate points in the action sequence.
- A similar definition exists for the probability of observing an experiment.

Linear Equivalence Relations:

- Equivalence relations between states that are required to agree on all linear behaviours (linear paths between a start and an end state)
- Two states s_1, s_2 are **t-equivalent** if they observe all tests with the same probability:

$$\forall \text{ tests } t, \langle s_1|t \rangle = \langle s_2|t \rangle$$

- Similarly, two states are **e-equivalent** if they observe all experiments with the same probability.

$$\forall \text{ experiments } e, \langle s_1|e \rangle = \langle s_2|e \rangle$$

- Two states s_1, s_2 are **trace-equivalent** if, performing the same sequence of actions we observe an identical sequence of observations with the same probability:

$$P(\omega_1 \dots \omega_n | s_1, a_1 \dots a_n) = P(\omega_1 \dots \omega_n | s_2, a_1 \dots a_n)$$

Branching Equivalence Relations:

- Equivalence relations between states that are required to agree on all possible behaviour trees between the two states.
- Two states s_1, s_2 are related by a **bisimulation** relation R if

$$s_1 R s_2 \Rightarrow \forall a \in \mathcal{A}, \forall \omega \in \mathcal{O}, \forall C \in \mathcal{S}/R, P(C, \omega | s_1, a) = P(C, \omega | s_2, a)$$

- Bisimulation requires not only the two starting states to exhibit the same behaviour, but each of their successors as well.

Results

- bisimulation \Rightarrow e-equivalence \Rightarrow trace equivalence \Rightarrow t-equivalence
- bisimulation \Leftarrow e-equivalence \Leftarrow trace equivalence \Leftarrow t-equivalence
- Branching relations are stronger than linear ones.
- Traces of the system are not enough to fully characterize its behaviour.

Duality for Automata

A **Deterministic Kripke Automaton** (DKA) is a quintuple

$$\mathcal{K} = (S, \mathcal{A}, \mathcal{O}, \delta : S \times \mathcal{A} \rightarrow S, \gamma : S \rightarrow 2^{\mathcal{O}}).$$

where S is a set of states, \mathcal{A} is a set of actions, \mathcal{O} is a set of observations, δ is a transition function and γ is an observation function associated with the states.

A **Partially Observable Probabilistic Automaton** (POPA) is a quintuple

$$\mathcal{H} = (S, \mathcal{A}, \mathcal{O}, \tau : S \times \mathcal{A} \times S \rightarrow [0, 1], \gamma : S \times \mathcal{O} \rightarrow [0, 1])$$

where $\tau(s, a, \cdot)$ defines a probability distribution on possible final states and $\gamma(s, \omega)$ is the probability of observing ω in state s .

POPAs are a special case of **Partially Observable Markov Decision Processes** (POMDPs), usually found in AI literature.


The **dual machine** \mathcal{K}' is constructed by interchanging the states and observations of the original automaton. In the dual machine, the observations one can make of a state $[t]$ are exactly the equivalence classes of states of the primal \mathcal{K} that satisfy the test t , i.e. $[t]$.

The **double dual** \mathcal{K}'' is constructed by repeating the same operation. The states of \mathcal{K}'' are observations of the dual, \mathcal{K}' and the observations of \mathcal{K}'' are the observations of \mathcal{K} .

The double dual construction produces a machine which is not only isomorphic to the original one, but is also its minimal version.


Duality for Deterministic Automata

Consider the finite automaton below, where $\mathcal{S} = \{A, \dots, F\}$, $\mathcal{A} = \{0, 1\}$ and $\mathcal{O} = \{G, R, Y\}$.



We can identify 3 equivalence classes of states, namely $t_1 = \{A, D\}$, $t_2 = \{B, E\}$ and $t_3 = \{C, F\}$.

Dual:




- The states of the dual are equivalence classes of states in the original automaton.
- The observations of the dual are the states of the original automaton.

Double Dual:

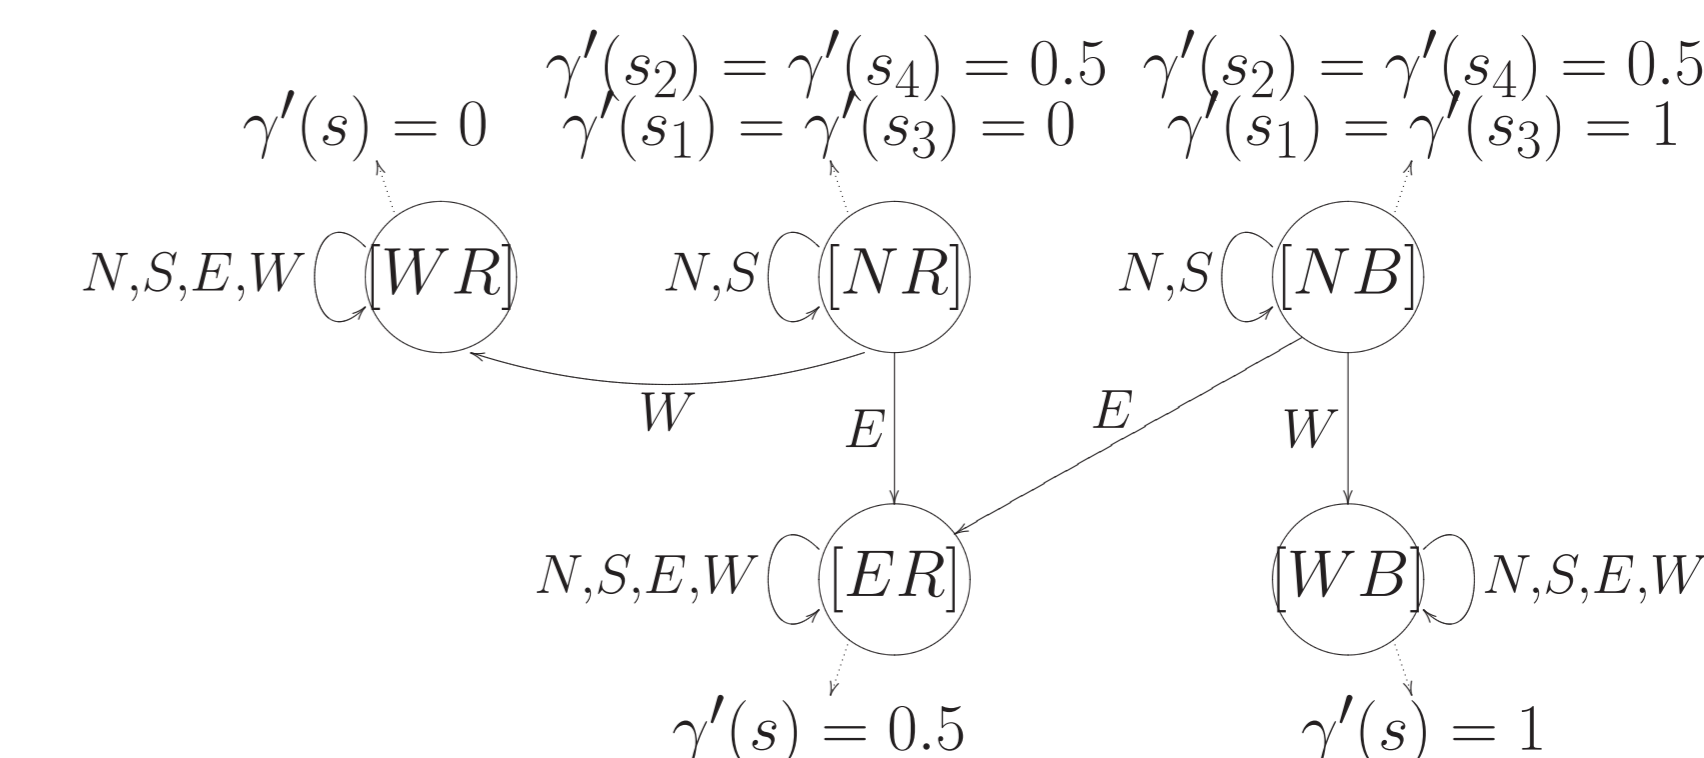
- In this special case, the original automaton was already minimized.
- The double dual has an isomorphic state-transition structure to the primal, but of course, the observations will be different.

Duality for Partially Observable Probabilistic Automata

Consider the POPA below, where $\mathcal{S} = \{s_1, \dots, s_4\}$, $\mathcal{A} = \{N, S, W, E\}$ and with real valued observations on the interval $[0, 1]$.

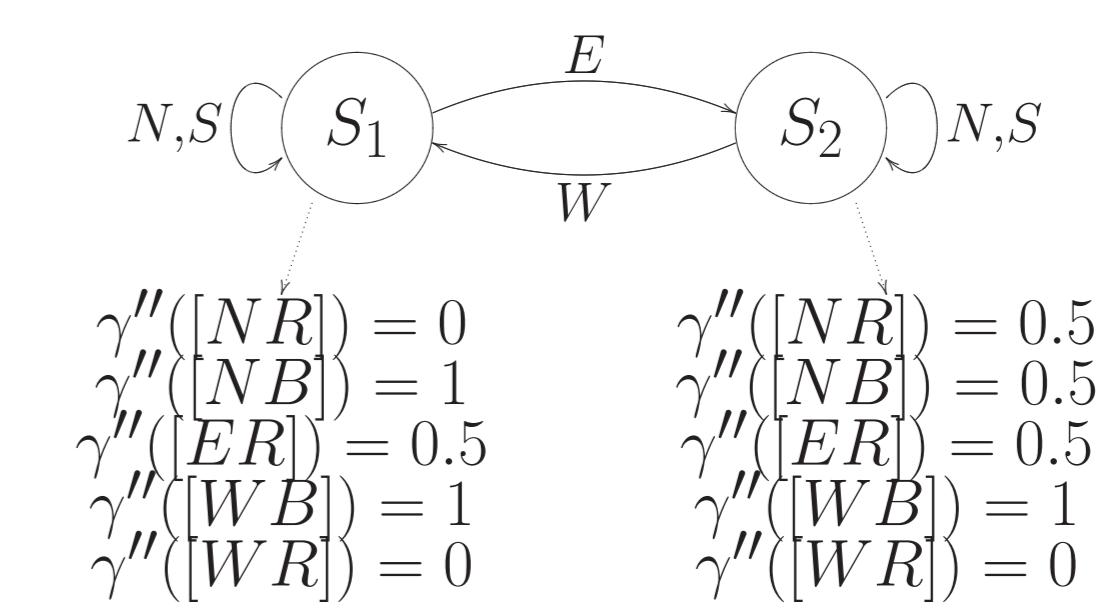


Dual:



- The states of the dual are equivalence classes of states that satisfy the same tests.
- The observations are the states of the original machine.
- The duality construction has made the transition function completely deterministic. Thus, the dual system is deterministic.

Double dual:



- The states are \mathcal{H}' -equivalence classes of formulas.
- The double dual of a probabilistic automaton is also deterministic.
- It is a minimal representation of the original system and completely predicts all its possible experimental outcomes.

Conclusions and Future Work

Summary

- There is a duality between state and observation, or, more precisely, between state and experiments.
- The double dual construction has a deterministic transition structure, and no hidden state.
- Bisimulation is the strongest equivalence relation on states.

Future Work

- Algorithms for planning and learning from the double dual
- Developing an approximation theory as well as the possibility of constructing more compact representations based on identifying nearby experiments.
- Extending the theory to continuous state spaces, and continuous observations.