Spatiotemporal Abstraction of Stochastic Sequential Processes

Sridhar Mahadevan

Department of Computer Science,
University of Massachusetts,
Amherst, MA 01003, USA
mahadeva@cs.umass.edu

Abstract. Probabilistic finite state machines have become a popular modeling tool for representing *sequential processes*, ranging from images and speech signals to text documents and spatial and genomic maps. In this paper, I describe two hierarchical abstraction mechanisms for simplifying the (estimation) learning and (control) optimization of complex Markov processes: *spatial* decomposition and *temporal* aggregation. I present several approaches to combining spatial and temporal abstraction, drawing upon recent work of my group as well as that of others. I show how spatiotemporal abstraction enables improved solutions to three difficult sequential estimation and decision problems: hidden state modeling and control, learning parallel plans, and coordinating with multiple agents.

1 Introduction

Abstraction has long been viewed as central to artificial intelligence (AI). A popular textbook defines abstraction as the "process of removing detail from a representation" [30]. Many approaches to abstraction have been pursued in the past several decades of research in AI. A common strategy is *constraint relaxation* where the problem is simplified by eliminating some conditions, as illustrated by logic-based planners such as ABSTRIPS [12] and by methods for discovering admissible heuristics [25]. This paper describes some recent work on *probabilistic abstraction* of stochastic sequential processes, which have become a common approach underlying many areas of AI.

Figure 1 characterizes a popular view of AI as the science underlying the design of *agents*: software or hardware artifacts that interact with an external environment through perception and action. What is unique about the agent-centered viewpoint is that it directs attention to the *sequential* interaction between an agent and its environment, and how to model the dynamics of such an interaction. Typically, the interaction is such that decisions (or observations) made earlier can impact later decisions.

Probabilistic finite state machines have become a popular paradigm for modeling sequential processes. In this representation, the interaction between an

S. Koenig and R. Holte (Eds.): SARA 2002, LNAI 2371, pp. 33–50, 2002.

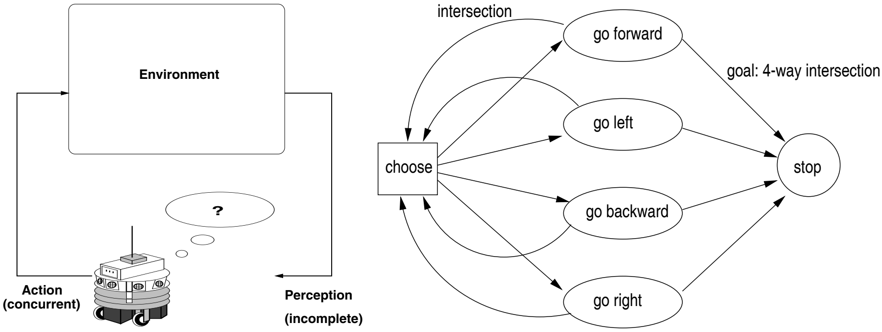

Fig. 1. The perception-action cycle of interaction between an agent and its environment can be modeled as a sequential process. A sequential program for a corridor navigation task represented as a finite state machine on the right. The set of observations generated as the agent executes this machine can be modeled as a Markov process.

agent and its environment is represented as a finite automata, whose *states* partition the past history of the interaction into equivalence classes, and whose *actions* cause (probabilistic) transitions between states. Here, state are a *sufficient statistic* for computing optimal (or best) actions, meaning past history leading to the state can be abstracted. This assumption is usually referred to as the *Markov* property.

Markov processes have become the mathematical foundation for much current work in reinforcement learning [33], decision-theoretic planning [1], information retrieval [7], speech recognition [10], active vision [20], and robot navigation [13]. In this paper, I focus on the abstraction of sequential Markov processes, and present two main strategies for "removing irrelevant detail": state aggregation/decomposition and temporal abstraction. State decomposition methods typically represent states as collections of *factored* variables [1], or simplify the automaton by eliminating "useless" states [3]. Temporal abstraction mechanisms, for example in hierarchical reinforcement learning [34,5,23], encapsulate lower-level observation or action sequences into a single unit at more abstract levels. For a unified algebraic treatment of abstraction of Markov decision processes that covers both spatial and temporal abstraction, the reader is referred to the paper by Ravi and Barto in these proceedings [27].

The main thesis of this paper is that combining spatial and temporal abstraction enables significant advances in solving difficult sequential estimation and decision problems. I focus on three specific problems – multiscale representations of hidden state, learning concurrent plans, and acquiring multiagent coordination strategies – and show how spatiotemporal abstraction is a powerful approach to solving instances of these well-known difficult problems.