2003

TR-2003008: Learning and Applying Temporal Patterns through Experience

Esther Lock

Abstract

Learning and Applying Temporal Patterns through Experience

By

Esther Lock

Adviser: Professor Susan Epstein

The thesis of this work is that the way patterns are formed over time can be learned during experience in a domain, and subsequently be used to improve decision making in that domain. Learning from temporal patterns can aid in the transition towards expertise, because it can increase the speed of decision making and improve the quality of the decisions made. While many researchers have explored ways to learn the significant patterns in a domain, the idea of focusing on the order with which these patterns are formed is novel.

This thesis investigates three learning methods that acquire and apply temporal patterns for game playing (i.e., sequences of game playing actions). *TP-Rote* (Temporal Patterns by Rote) is a rote-learning, caching scheme that hones in on frequently-used segments of a search space and memorizes them to reuse them later. Experiments with TP-Rote indicate a significant speedup in play, simulating the gradual shift to "play without thought" seen in human game players at various points in a game. The second method, *TP-Context* (Temporal Patterns through Context), generalizes states, retaining only the *context* (the motivating description) of a sequence. Such a generalization may cover many states and thereby expand the applicability of a temporal pattern. This is especially important in large state spaces, where identical states may not often recur. TP-Context exploits the knowledge inherent in the sequences actually experienced in game-

playing contests, to discover context. A learned context is associated with a sequence that TP-Context suggests as a course of action whenever the context is found on a given state. Experimental results show that TP-Context can learn to play two games, lose tic-tac-toe and five mens morris, based upon this sequence knowledge. The third method, *TP-Sitact* (Temporal Patterns through Context using a situation-action representation), is an extension of TP-Context. Although based on sequence knowledge, TP-Sitact suggests individual actions rather than sequences. TP-Sitact significantly improves TP-Context's game playing prowess. With these three methods, this thesis has successfully demonstrated that temporal patterns can be successfully learned by machine learning programs and used to decrease execution speed, help make decisions and aid in context discovery.