Adaptive Tile-Coding for Reinforcement Learning

Shimon Whiteson and Peter Stone

Abstract: Difficult problems in reinforcement learning typically require function approximators to effectively estimate value functions. Many different kinds of function approximators are currently in use, including neural networks, radial basis functions, and instance-based methods. Among the most successful are tile-codings (or CMACs), which consist of piecewise-constant approximations formed by discretizing the state space into disjoint tiles and aggregating values from multiple, slightly offset tilings. However, to make tile-coding work well in practice, a human expert must manually design the tile-coding representation, i.e. the size and shape of each tile. We present a new method, called adaptive tile coding, which automates this process. Borrowing the idea of "complexification" from methods that learn representations for neural network function approximators [1], this approach begins with simple representations with few tiles and adds new complexity during learning by splitting existing tiles into smaller ones. Local estimates of Bellman error are used to determine which tiles should be split. This approach, in addition to automatically discovering effective representations, provides a natural way to reduce the function approximator's level of generalization over time, a factor known to be critical in tile-coding [2]. Empirical results in multiple domains confirm that adaptive tile coding can perform better and faster than fixed representations.

References

[1] Shimon Whiteson and Peter Stone. Evolutionary Function Approximation for Reinforcement Learning. Journal of Machine Learning Research, 7(May):877–917, 2006.

[2] Alexander Sherstov and Peter Stone. Function approximation via tile coding: Automating parameter choice. In Proc. Symposium on Abstraction, Reformulation, and Approximation (SARA-05), Edinburgh, Scotland, UK, July 26-29, 2005.