

# Function Approximation Using Tile and Kanerva Coding For Multi-Agent Systems

Cheng Wu  
Northeastern University  
Boston, MA, U.S.A.  
cwu@ece.neu.edu

Waleed Meleis  
Northeastern University  
Boston, MA, U.S.A.  
meleis@ece.neu.edu

## ABSTRACT

Function approximation can improve the ability of a reinforcement learner. Tile coding and Kanerva coding are two classical methods for implementing function approximation, but these methods may give poor performance when applied to large-scale, high-dimensional instances. In the paper, we evaluate a collection of hard instances of the predator-prey pursuit problem, a classic multi-agent reinforcement learning problem, to compare these two methods and their optimization techniques. We first show that Kanerva coding gives better results than Tile coding when the dimension of the instances increases. We then describe a feature optimization mechanism and show that it can increase the fraction of solved instances by both Tile coding and Kanerva coding. Finally, we demonstrate that a fuzzy approach to function approximation can further increase the fraction of instances. We show that our fuzzy approach to Kanerva coding outperforms fuzzy Tile coding when feature optimization is applied. on large-scale, high-dimensional multi-agent problems.

## Categories and Subject Descriptors

I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search

## General Terms

Algorithms, Experimentation

## Keywords

Reinforcement Learning, Function Approximation, Fuzzy Logic

## 1. INTRODUCTION

Q-learning [10] has emerged as one of the most successful reinforcement learning strategies. The algorithm works by combining state space exploration and exploitation to learn the value of each state-action pair. A key limitation of Q-learning is the size of the table needed to store the state-action values. The requirement that an estimated value be stored for every state-action pair limits the size of the learning problems that can be solved. The Q-learning table is

**Cite as:** Function Approximation Using Tile and Kanerva Coding For Multi-Agent Systems, C. Wu and W. Meleis, *Adaptive Learning Agents (ALA) workshop in 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, May, 10–15, 2009, Budapest, Hungary.

Copyright © 2009, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

typically large because of high dimensionality of the state-action space, or because the state or action space is continuous. Function approximation [2], which stores an approximation of the entire table, is one way to solve this problem.

In recent years, many function approximation techniques exist. They can be divided into two categories: feature-based function approximation and basis-function-based function approximation. Feature-based function approximation includes coarse coding [4], Tile coding [1] (known as CMAC [10]). Coarse coding, for example, uses a collection of overlapping regions within the state-action space, each of which corresponds to a binary natural feature. Basis-function-based function approximation includes Radial Basis Function Networks (RBFNs) [7] and Kanerva coding [6]. For example, RBFNs uses a series of radial basis functions that vary smoothly and are differentiable, not natural features.

Tile coding is a special case of coarse coding. In Tile coding,  $k$  tilings are selected, each of which partitions the state-action space into tiles. The receptive field of each binary feature corresponds to a tile, and a  $\theta$  value is maintained for each tile. A state-action pair  $p$  is *adjacent* to a tile if the receptive field of the tile includes  $p$ . The Q-value of a state-action pair is equal to the summation of the  $\theta$  values of all adjacent tiles. A limitation of this technique is that they cannot handle the state-action spaces with high dimensions. In general, the number of tilings needed to achieve good results grows exponentially with the number of dimensions in the state-action space [9].

Kanerva coding is a special case of RBFNs. It uses the architecture of sparse distributed memories [5] that can also reduce the amount of memory needed to store the state-action value table. This approach is particularly well-suited to problem domains with high dimensionality. A collection of  $k$  *prototype state-action pairs*, (prototypes) is selected, each of which again corresponds to a binary feature. A state-action pair  $s$  and a prototype  $p_i$  are said to be *adjacent* if their bit-wise representations differ by no more than a threshold number of bits. Normally we set the threshold as 1 bit. We define the *membership grade*  $\mu_i(s)$  of  $s$  with respect to  $p_i$  to be equal 1 if  $s$  is adjacent to  $p_i$ , and equal to 0 otherwise. A value  $\theta(i)$  is maintained for the  $i$ th feature, and an approximation of the value of a state-action pair is then the sum of the  $\theta$  values of the adjacent prototypes, that is  $\sum_i \theta(i) \mu_i(s)$ . In this way, Kanerva coding can greatly reduce the size of the value table that needs to be stored.

A tile in Tile coding represents information about one or some, but not all, dimensions of the state-action space. This allows that computational complexity increases expo-

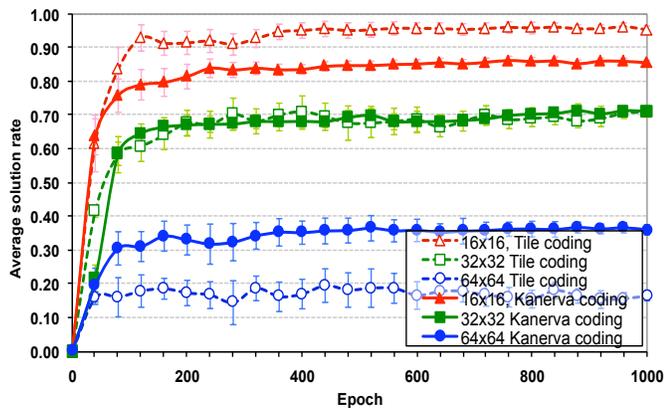


Figure 1: Average solution rate of traditional Tile coding and traditional Kanerva coding.

nentially with the number of dimensions. A prototype in Kanerva coding contains information about all dimensions. This allows that the computational complexity of approximate function depends completely on the number of prototypes, not the number of dimensions of state-action space.

Several researchers have investigated the use of Tile coding and Kanerva coding within reinforcement learning [8, 11]. However there have been few published attempts to compare them for multi-agent problems. Furthermore, optimizing features in Tile coding and Kanerva coding [8] has been shown to reduce the number of features needed to achieve a good solution rate. However no work has been done to directly compare the effect of this adaptive function approximation across Tile coding and Kanerva coding.

## 2. EXPERIMENTAL DESIGN

Multi-agent problems can be difficult to solve by traditional machine learning techniques because the state space is often very large. The predator-prey pursuit problem [3] is a classic example. A general version of the predator-prey pursuit problem takes place on a rectangular grid with one or more predator agents and one or more prey agents. Each grid cell is either open or closed, and an agent can only occupy open cells. Each agent has an initial position. The problem is played in a sequence of time periods. In each time period, each agent can move to a neighboring open cell one horizontal or vertical or diagonal step from its current location, or it can remain in its current cell. All moves occur simultaneously, and no more than one predator agent occupy the same cell at the same time. Each predator agent can observe the location of each prey agent. If a predator agent is in the same cell as a prey agent at the end of a period, then that target has been caught. The predator agents’ goal is to catch all prey agents in the shortest time.

In our experiments, pursuit takes place on a rectangular grid with size of  $n \times n$  and  $n$  is 16, 32 and 64. There are  $n$  closed blocks that are distributed randomly. The agent receives a reward of 1 when it reaches the goal cell, and receives a reward of 0 in every other cell. The agent attempts to reach a fixed goal cell. In each epoch, we apply the Q-learning with Tile coding or Kanerva coding to 40 random training instances followed by 40 random test instances. The exploration rate  $\epsilon$  is set to 0.3, which we found experimen-

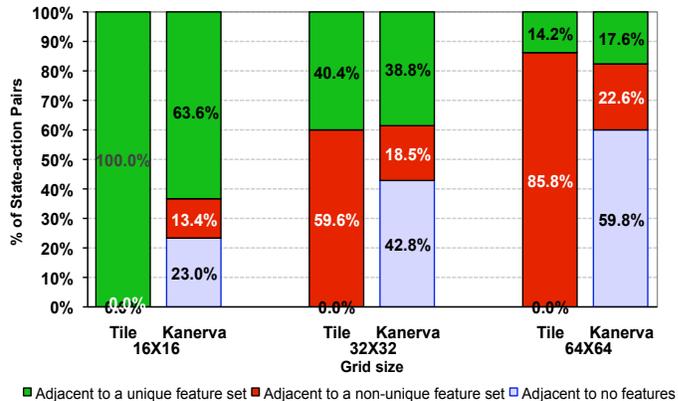


Figure 2: Distribution of the number of collisions per prototype using Tile and Kanerva coding in a sample run.

tally to give the best results. The initial learning rate  $\alpha$  is set to 0.8, and it is decreased by a factor of 0.995 after each epoch. For every 40 epochs, we record the average fraction of test instances solved during those epochs. Each experiment is performed 5 times and we report the means and standard deviations of the recorded values.

## 3. COMPARISON OF TRADITIONAL TILE CODING AND KANERVA CODING

We begin by comparing Q-learning with traditional Tile coding and traditional Kanerva-based function approximation as the dimension of the state-action space increases. We implement Tile coding by representing the state-action pair as a binary vector. Each tiling corresponds to a 3-tuple of bit positions. For  $n = 16$ , the number of possible tilings is 364 which produces a total of 2912 tiles. We fix the number of tilings to 364 across all grid sizes. In this way we can evaluate the effect of a fixed tiling on the performance of Q-learning as the dimension of the state-action space increases. In a similar way, we implement Kanerva coding by randomly selecting 2912 prototype states from the state-action space. The number of prototype states is fixed across all grid sizes.

Figures 1 shows the results obtained by traditional Tile coding and traditional Kanerva as the size of the grid varies from 16 to 64. For a grid size of 16, Tile coding solves 96% of the test instances while Kanerva coding solves only 86% of the test instances, after 1000 epochs. However, for a grid size of 64, Tile coding solves only 17% of the test instances while Kanerva Coding solves only 37% of the test instances.

The results show that when the number of dimensions is small, traditional Tile coding outperforms traditional Kanerva coding. However, as the number of dimensions increases, Tile coding’s performance degrades faster than Kanerva coding when the number of prototypes is fixed. We conclude that Kanerva coding performs better than Tile coding when the dimension of the state-action space is large.

Function approximation works best when each state-action pair is adjacent to unique subset of features (that is, tiles in Tile coding, and prototypes in Kanerva coding). If features are not well distributed across the state-action space, many state-action pairs will either not be adjacent to any features,

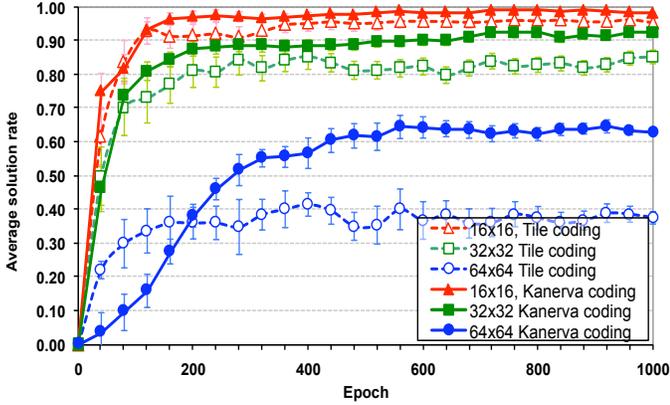


Figure 3: Average solution rate of Tile coding and Kanerva coding with feature optimization.

or be adjacent to identical sets of features. Such feature collisions reduce the quality of the results because the solver will not be able to distinguish distinct state-action pairs, and the Q-values of such state-action pairs will be equal.

Figure 2 shows the fraction of state-action pairs that are adjacent to no features, more than one feature, and exactly one feature when traditional Tile coding and Kanerva coding are applied to sample predator-prey instances of varying sizes. For Tile coding, the fraction of state-action pairs that are adjacent to exactly one tile varies from 100% to 14% as the size of the grid increases, so up to 86% of the state-action pairs cause collisions. For Kanerva coding, the fraction of state-action pairs that are adjacent to exactly one prototype varies from 64% to 18% as the size of the grid increases, so up to 82% of the state-action pairs cause collisions.

These results show that when the dimension of the state-action space is small, Tile coding causes fewer collisions than Kanerva coding, which explains the good performance of Tile coding. Kanerva coding gives slightly fewer collisions when the dimension of the state-action space is large.

In both cases, the number of collisions increases sharply as the dimension of the state-action space increases, which explains the reduction in performance. It is therefore necessary to consider optimization mechanisms that adjust the distribution of features to reduce the number of collisions as the dimension of the state-action space increases.

#### 4. EFFECT OF FEATURE OPTIMIZATION ON TILE AND KANERVA CODING

We say that a feature is visited during Q-learning if it is adjacent to the current state-action pair, and we optimize features based on these visit frequencies. Initial features are selected randomly. After a fixed number of iterations, features are updated using the following mechanisms.

Features that are rarely visited do not contribute to the solution of instances and indirectly cause collisions. We periodically delete each feature with probability equal to an exponential function of the visit frequency. I.e. the probability  $p_{del}$  of deleting a feature whose visit frequency is  $v$  is  $p_{del} = \lambda e^{-\lambda v}$ , where  $\lambda = 1$ . Features that are visited frequently are also likely to cause many collisions. We reduce the number of collisions by splitting frequently visited fea-

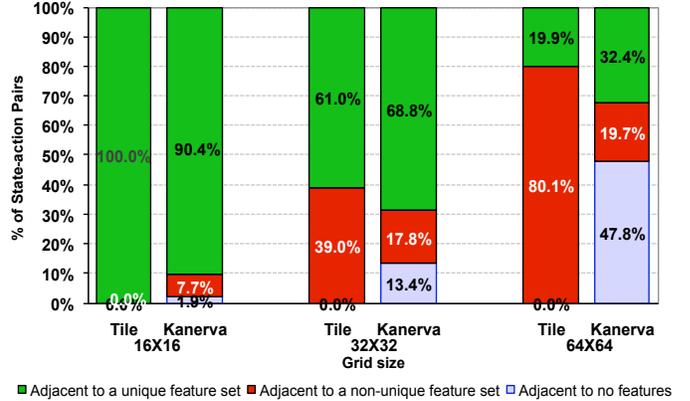


Figure 4: Distribution of the number of collisions per prototype using Tile and Kanerva coding with feature optimization in a sample run.

tures. A feature  $s_1$  that has been visited the most times is selected, and a new feature  $s_2$  is created based on  $s_1$ . For Tile coding,  $s_2$  is constructed by randomly selecting and retaining two of the bit positions in  $s_1$ , and selecting a new third bit position. For Kanerva coding,  $s_2$  is constructed by inverting a fixed number of bits in  $s_1$ . In both cases, the feature  $s_1$  remains unchanged.

The effect of feature optimization is shown in Figure 3 as the size of the grid varies from 16 to 64. The graph in Figure 3 shows the average solution rate obtained by Tile coding and Kanerva coding with feature optimization. Tile coding with feature optimization solves an average of 85% of the test instances with a grid size of 32, and 37% with a grid size of 64. Kanerva coding with feature optimization solves an average of 92% of the test instances with a grid size of 32, and 63% with a grid size of 64.

These results show that feature optimization can significantly improve the average solution rate when using Tile coding or Kanerva coding. However, feature optimization causes a larger absolute improvement in the solution rate with Kanerva coding (26%) than with Tile coding (20%). We also observe that Tile coding gives a larger deviation in the solution rate than Kanerva coding. This shows that Tile coding with feature optimization is less stable than Kanerva coding with optimization.

Figure 4 shows the fraction of state-action pairs that are adjacent to no features, more than one feature, and exactly one feature when Tile coding and Kanerva coding with feature optimization are applied to sample predator-prey instances. For Tile coding, the fraction of state-action pairs that are adjacent to exactly one tile varies from 100% to 20% as the size of the grid increases, so up to 80% of the state-action pairs cause collisions. For Kanerva coding, the fraction of state-action pairs that are adjacent to exactly one tile varies from 90% to 32% as the size of the grid increases, so up to 68% of the state-action pairs cause collisions.

These results show that when the dimension of the state-action space is large, feature optimization causes a larger absolute reduction in collisions using Kanerva coding (15%) than using Tile coding (6%). However, in both cases the number of collisions again increases sharply as the dimension of the state-action space increases.

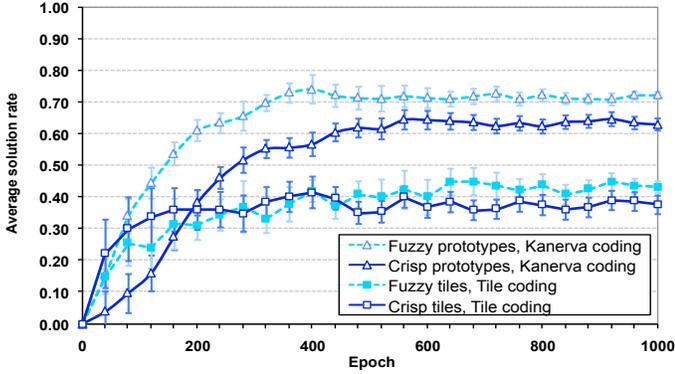


Figure 5: Average solution rate of fuzzy Tile and Kanerva coding with feature optimization.

## 5. COMPARISON OF FUZZY TILE AND KANERVA CODING

Feature receptive fields with crisp boundaries can cause frequent collisions. A more flexible approach to computing state-action values is to use receptive fields with fuzzy boundaries. For Tile coding, we allow a state-action pair to update  $\theta$  values of all tiles within a tiling, instead of a single tile. For Kanerva coding, a state-action pair updates  $\theta$  values of all prototypes, instead of just adjacent prototypes.

Instead of being binary values, membership grades vary continuously between 0 and 1 across features. Such fuzzy membership grades are larger for adjacent features and smaller for more distant features. Since feature collisions occur only when two state-action pairs have identical membership vectors, collisions are less likely.

In the fuzzy approach to Tile coding, given a state-action pair  $s$ , the  $i$ th tile of the  $j$ th tiling  $t_{i,j}$ , and a constant variance  $\sigma$ , the membership grade of  $s$  with respect to the  $t_{i,j}$  is  $\mu_{i,j} = \exp\left(-\frac{\|s - t_{i,j}\|^2}{2\sigma^2}\right)$ , where  $\|s - t_{i,j}\|$  represents the bit difference between  $s$  and  $t_{i,j}$ . Note that the membership grade of a tile with respect to an identical state-action pair is 1. In the fuzzy approach to Kanerva coding, given a state-action pair  $s$ , the  $i$ th prototype  $p_i$ , and a constant variance  $\sigma$ , the membership grade of  $s$  with respect to  $p_i$  is  $\mu_i = \exp\left(-\frac{\|s - p_i\|^2}{2\sigma^2}\right)$ , where  $\|s - p_i\|$  represents the bit difference between  $s$  and  $p_i$ . Note that the membership grade of a prototype with respect to an identical state-action pair is 1, and the membership grade of a state-action pair and a completely different prototype approaches 0.

As with traditional Kanerva coding, a value  $\theta(i)$  is maintained for the  $i$ th feature and an approximation of the value of a state-action pair is computed in the same way as before. The effect of an update to a prototype's  $\theta$ -value is now a continuous function of the bit difference between the state-action pair and the prototype, and an update can have a large effect on immediately adjacent prototypes, and a smaller effect on more distant prototypes.

The effect of fuzzy function approximation with a grid size of  $64 \times 64$  is shown in Figure 5. The graph shows the average solution rate obtained by fuzzy and crisp function approximation using Tile coding and Kanerva coding with feature optimization. Fuzzy function approximation increases the

average solution rate from 37% to 43% using Tile coding, and from 63% to 72% using Kanerva coding. The fuzzy approach improves the performance of both approaches, and the absolute improvement in the average solution rate is larger using Kanerva coding than using Tile coding.

## 6. CONCLUSION

In this paper we have evaluated the performance of several function approximation techniques by applying them to large-scale, high-dimensional instances of predator-prey pursuit instances. We first showed that traditional Tile coding and Kanerva coding give poor performance and demonstrate that Kanerva coding gives better results than Tile coding when the dimension of the instances is large. We then showed that this poor performance is a result of the inefficient distribution of features which causes frequent collisions. We described a feature optimization mechanism and showed that it can increase the fraction of instances that are solved by both Tile coding and Kanerva coding. Finally, we showed that a fuzzy approach to function approximation can further reduce the number of collisions. We showed that our fuzzy approach to Kanerva coding outperforms fuzzy Tile coding when feature optimization is applied. We conclude that for large-scale, high-dimensional multi-agent optimization problems, discrete and fuzzy Kanerva coding represent powerful function approximation techniques that can outperform discrete and fuzzy Tile coding.

## 7. REFERENCES

- [1] J. Albus. *Brains, Behaviour, and Robotics*. McGraw-Hill, 1981.
- [2] L. Baird. Residual algorithms: Reinforcement learning with function approximation. In *Proc. of the 12th Intl. Conf. on Machine Learning*. Morgan Kaufmann, 1995.
- [3] M. Benda, V. Jagannathan, and R. Rodhiawalla. On optimal cooperation of knowledge sources. *Technical Report, Boeing Computer Services*, 1985.
- [4] G. Hinton. Distributed representations. *Technical Report, Department of Computer Science, Carnegie-Mellon University, Pittsburgh*, 1984.
- [5] P. Kanerva. *Sparse Distributed Memory*. MIT Press, 1988.
- [6] K. Kostiadis and H. Hu. Kabage-rl: kanerva-based generalisation and reinforcement learning for possession football. In *Proc. of IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2001.
- [7] T. Poggio and F. Girosi. Networks for approximation and learning. In *Proc. IEEE, vol78, no. 9, pp. 1481-1497*, 1990.
- [8] B. Ratitch and D. Precup. Sparse distributed memories for on-line value-based reinforcement learning. In *Proc. of the European Conf. on Machine Learning*, 2004.
- [9] R. Sutton and A. Barto. *Reinforcement Learning: An Introduction*. Bradford Books, 1998.
- [10] C. Watkins and P. Dayan. Q-learning. *Machine Learning*, 8:279–292, 1989.
- [11] C. Wu and W. Meleis. Adaptive kanerva-based function approximation for multi-agent systems. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2008.