

X-trees: An unsupervised learning method for use in developmental agents

Brandon Rohrer

Intelligent Systems, Robotics, and Cybernetics Group

Sandia National Laboratories, Albuquerque, NM 87185-1010

Email: brrohre@sandia.gov

I. INTRODUCTION

In developmental agents, creating abstract representations is an important aspect of learning. Unless an agent starts out with a pre-programmed set of concepts, it must generate them from its experience. One method for doing this is to group observations into clusters which are then assumed to be similar. Clustering state observations is a type of unsupervised learning. (A review of unsupervised learning methods can be found in [1, ch. 10].) Unsupervised learning allows agents to create abstractions from their experiences. Expressing raw experiences in abstract terms greatly decreases the size of the agent's input space. This in turn lets the agent use more powerful learning, perception, and planning algorithms that can only operate effectively in modest state spaces.

In this work, state observations are grouped into a small number of clusters. Each dimension in the state space represents a separate state variable. What separates this work from previous unsupervised learning methods is that it does not assume that state dimensions are ordered fields. That is, it does not assume that 1 is less than 2 or that the distance from 1 to 2 is less than the distance from 1 to 1000. As a result, this method can be applied to a wide variety of data, including continuous, discretized, categorical, binary, and hybrid data types.

The algorithm for clustering non-ordered fields is a decision tree. Initially the entire state space represents a single cluster. The state space is then repeatedly subdivided into regions, with some of the regions representing clusters. Each subdivision is represented by splitting a leaf node into two new leaves. A subset of the tree's leaves are the clusters. The dimension along which to divide each leaf is chosen so as to maintain clusters as much as possible. A tree that chooses when and where to divide solely on the basis of the data distribution is said to have the X-property. [2] In order to distinguish between this and other tree-based tools (including Decision Q-Trees [3], S-trees [4], and T-Trees, an extension of Continuous U-Trees [5]) it will be referred to as an X-tree.

In this paper, I present a detailed description of X-trees, together with two implementations demonstrating its operation. One X-tree is used with digital images to find visual primitives similar to those found in human cortical area V1. A second X-tree is used to create abstractions of visual inputs to a mobile robot.

II. METHOD

X-Trees are conceptually similar to other decision tree based algorithms, but differ in the specifics of their implementation.

Although the most well known implementations of decision trees are CART [6] and C4.5 [7], a concise overview of a wide variety of decision trees is given in [8]. A decision tree is defined by how it answers the three questions 1) When to split a node, 2) Where to split that node, and 3) When to stop splitting a node. (Some trees also answer a fourth question—When and how to prune branches from the tree—but X-Trees do not use pruning.) After describing the structure and populating process of a leaf, this section answers those questions for X-trees.

A. Leaf structure

Each leaf represents an explicitly bounded region of the state space. When defining the initial leaf that encompasses the entire state space, upper and lower bounds on all input dimensions must be known or assumed. As a leaf is divided, the bounds of its children are created using the decision boundary. In this way, the leaves form a non-overlapping set of regions that completely cover the state space. (See Figure 1) Each leaf is a hyperbox, defined by minimum and maximum values in each dimension of the state space.

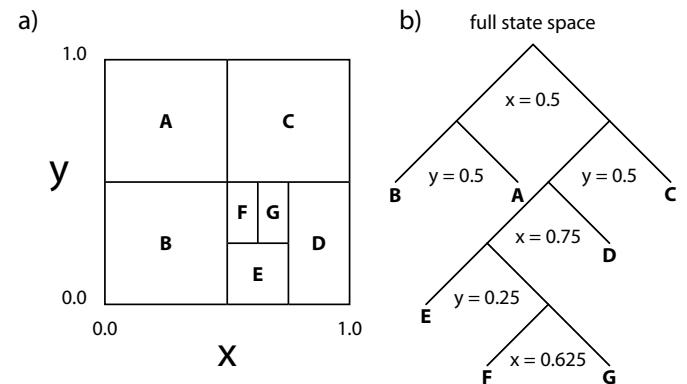


Fig. 1. a) A partitioned state space and b) the decision tree representation of that partitioning. The top node in the tree represents the entire state space. It also indicates the division of the space into its left and right halves along the line $x = 0.5$. Leaf nodes in the tree represent individual regions as labeled. Any internal node on the tree represents the union of all of the leaf node regions that fall under it.

Each leaf represents a single region. Each node on the tree represents a super-region, the combination of the regions defined by its children, which themselves may be either leaves or branch nodes. This family tree of regions allows for a rough measure of distance between them. The number of generations that must be ascended and descended to traverse from one node to another provides a quantitative measure of their similarity. Note that this traversal distance is not a state space distance metric and does not impose an assumption of ordered fields on the state space. Nodes may be adjacent in state space, yet have an arbitrarily large traversal distance.

Most tree-based learning tools are supervised learning algorithms, rather than unsupervised. Observations of input states are typically accompanied by categories, in classification problems, or values, in regression problems. X-trees differ from these. The leaves of an X-tree do not provide an estimate of a reward, value function, or class membership. Instead, X-trees identify locally dense regions of state activity. In addition, X-trees subdivide clusters once they exceed a predetermined size. The motivating assumption behind X-trees' partitioning heuristic is that useful concepts capture a roughly fixed number of observations. If a concept represents very few observations, it provides only limited abstraction, and if it represents a very large number of observations, it may be too general to be useful.

B. When to split a node

An X-tree node becomes eligible to split when the number of observed states it contains exceeds a user-defined threshold, M . Each observation that falls within a given region of state space is a *member* of that region. When more a node has more than M members, it is considered ripe for bisection. However, each member has a finite lifetime, L , also defined by the user. As a result, a node is ready to split when more than M of the last L state observations fall within its region of state space.

C. Where to split a node

The objective of an X-tree is to find clusters of observations and bound them tightly. Each cluster represents a concept, or a closely related group of states. Tighter bounds give a more concise concept definition. Therefore, when a node is bisected, the dimension along which to split and the value at which to split it are chosen such that the cluster is preserved as much as possible. For all admissible splits, the split which maximizes the disparity between the number of members in each child is selected. If more than one candidate split achieves the maximum value, then the split value is randomly selected from among them.

More formally, for an admissible split (d, x) along dimension d at value x , the quality of the split, q is given by

$$q_{d,x} = |m_a - m_b| \quad (1)$$

where m_a is the number of members from the parent node that would fall under one of its children for that split and m_b is the number of members that would fall under the other child

node. Quality is symmetric with respect to m_a and m_b . The winning split fulfills the condition

$$\max_{d \in D} (\max_{x \in X_d} q_{d,x}) \quad (2)$$

where D is the set of all state space dimensions and X_d is the set of all admissible split values for dimension d . The admissible split values in X_d divide the region along dimension d roughly in half. X_d is generated by taking a number of evenly-spaced values along d . More specifically, if d_{\min} and d_{\max} are the limits of d in the region to be divided, then N admissible split values can be generated using natural numbers $i | i \leq N$:

$$d_{\text{low}} = d_{\min} + \frac{1}{4}(d_{\max} - d_{\min}) \quad (3)$$

$$d_{\text{high}} = d_{\min} + \frac{3}{4}(d_{\max} - d_{\min}) \quad (4)$$

$$x_i = d_{\text{low}} + (d_{\text{high}} - d_{\text{low}}) \frac{i}{N+1} \quad (5)$$

The resulting values of X_d fall between, but do not include, one quarter and three quarters of the range of d . N is a user-selected constant.

D. When to stop splitting a node

Due to the fact that members have a limited lifetime, an explicit stopping criterion is typically unnecessary. The splitting condition of accumulating more than M members within any interval L becomes harder to achieve as regions get smaller. With any continuous distribution of observations, the member accumulation rate will approach zero as region size approaches zero. However, when processing discontinuous distributions (as with categorical or discretized data), observations may be grouped at a single value, prompting endless repeated splits. This can occur whenever a region's range in a given dimension exceeds the discrete resolution in that dimension. A method for handling this degenerate case is to specify a threshold for the minimum region size in each dimension. This can either be a constant ϵ that is the same for each dimension or a vector E with separate values for each dimension. Further divisions are not allowed along any dimension for which $d_{\max} - d_{\min} < \epsilon$.

E. Clusters

X-trees seek to isolate groups of members into tightly-grouped clusters. The condition for becoming a cluster is being a leaf node while having the number of members in a node exceed $M/2$. Thus, when a node is divided, one of the children will become a cluster and the other will not (although it may eventually become one), and the parent node ceases to be a cluster. Once a node is designated a cluster, it remains one, even if the number of members it contains falls below $M/2$. Cluster nodes represent concepts, the categories into which most of the data fall.

Once clusters have been formed, they serve as a means of interpreting raw data. A single observation is a point in state space; if it falls within the hyperbox of a cluster, the

data point can be represented more abstractly as that concept. But typically, clusters do not provide complete coverage of the state space. In order to find a conceptual interpretation of points that do not fall within a cluster, a distance measure can be used to identify the nearest clusters. The distance measure X-trees use reflect the fraction of dimensions in which an observation matches a cluster. The distance δ between a state s and a cluster c in an n -dimensional space is given by

$$\delta = 1 - \frac{\sum_{i=1}^n \text{in}(s, c, i)}{n} \quad (6)$$

where

$$\begin{aligned} \text{in}(s, c, i) &= 1, \text{ if } s_i \geq c_{i \min} \text{ and} \\ &\quad s_i < c_{i \max} \\ &= 0, \text{ otherwise} \end{aligned} \quad (7)$$

where s_i is the value of s , $c_{i \min}$ is the minimum value of c , and $c_{i \max}$ is the maximum value of c , all in the i^{th} dimension. If s falls within c for all dimensions, and $\delta = 0$. If s falls within c in only half of the dimensions, $\delta = 1/2$. The distance measure takes no account of how far outside a cluster a state may fall. Only matching and non-matching are reflected in the distance. It should be noted that this is a departure from many other unsupervised learning approaches, which use the Euclidean distance measure. X-trees' distance measure avoids making the assumption that the state space is well scaled and allows it to be applied to a wider range of problems.

F. Simulation 1: Clustering example

The first simulation reported here serves to illustrate the operation of X-trees. A pair of arbitrarily selected two-dimensional normal distributions were used to generate random points in a two-dimensional state space. The defining parameters for the X-tree were chosen as follows:

Anomaly lifetime,	$L = 100$
Maximum number of members,	$M = 40$
Number of split candidates per dim.,	$N = 101$
Minimum category size,	$\epsilon = 0$

G. Simulation 2: Vision

In the second simulation, X-trees were applied to grayscale images to find clusters of common pixel groups. Three-by-three pixel groups were randomly selected from a set of images and processed by an X-tree. The nine-dimensional state space was partitioned and clusters were identified. The defining parameters for the X-tree were chosen as follows:

Anomaly lifetime,	$L = 1000$
Maximum number of members,	$M = 70$
Number of split candidates per dim.,	$N = 1$
Minimum category size,	$\epsilon = 1.0$

Choosing $N = 1$ ensured that all splits would be perfect bisections, and choosing $\epsilon = 1.0$ ensured that no dimension would be split more than once. The resulting state space was then had $2^9(512)$ distinct states.

III. RESULTS

The results for simulation 1 are shown in Figure 4. The two normal distributions produced two groups of points, and the recursive partitioning of the space produced cluster nodes corresponding to the densest portion of each group. The tree in Figure 3 shows two major branches, one corresponding to each of the clusters. Simulation 1 was run until a stable tree structure was achieved, that is, no new nodes were created for a substantial period of time. Simulation 1 provides an illustration of X-tree's operation on a simple problem and is intended to provide an intuitive basis for understanding the results of simulation 2.

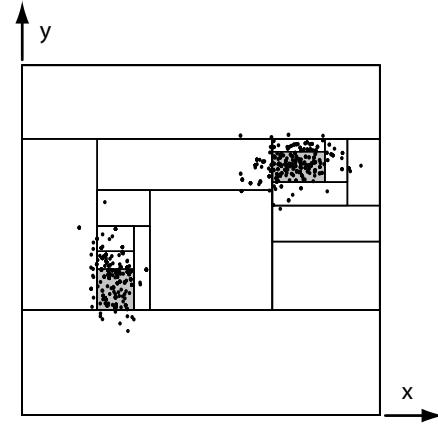


Fig. 2. Two dimensional state space as partitioned by an X-tree. The two normal distributions produced clearly separated groups of points. Each bisection is shown by a line on the bisection boundary. Each region represents a leaf node. Cluster nodes are shaded gray and correspond to the densest part of each data distribution.

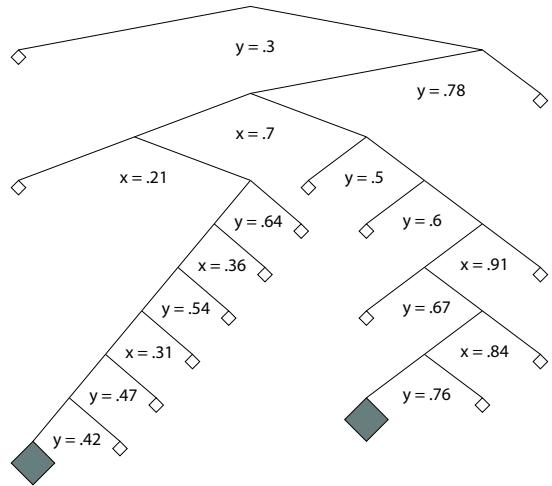


Fig. 3. Tree representation of the partitions shown in Figure 4. The root node of the tree represents the entire state space. Each branching represents a split at the line given by the corresponding equation. The right branch represents the region to the right or above the split. Unfilled boxes on leaf nodes show regions that are not clusters. The two gray filled boxes show the clusters.

The result of simulation 2 was a set of clusters in the nine-dimensional state space, shown in Figure 5. These can be

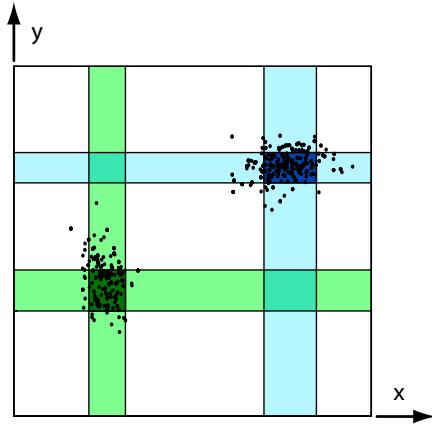


Fig. 4. Cluster assignment vectors across the state space. Every point in the state space can be said to have a similarity to each cluster of $1 - \delta$ between 0 and 1 as described in Equation 6. For the white regions, those similarities are $[0, 0]$. For the other regions they are $[1, 0]$ (dark green), $[.5, 0]$ (light green), $[0, 1]$ (dark blue), $[0, .5]$ (light blue), and $[.5, .5]$ (turquoise).

notionally compared to the receptive fields identified in the human visual cortex, which have exhibited strong characteristics of directionality. (See Figure 6.)

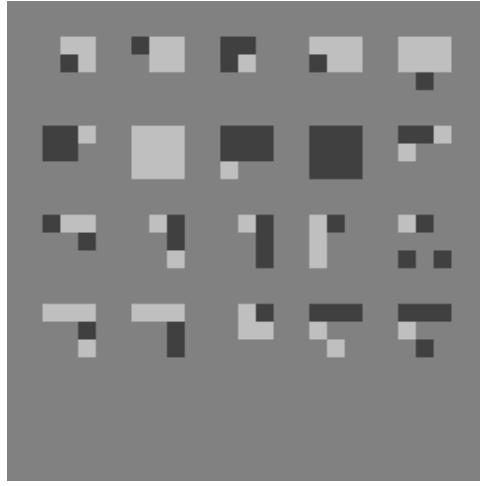


Fig. 5. Graphical representation of the 20 clusters found in the space of 9 pixel groups.

An image can be reconstructed using the cluster set by breaking it into 9-pixel tiles, finding the similarity of the tile with each of the clusters, linearly combining the clusters weighted by their similarity, and substituting them for the original tile. The results of one such reconstruction are shown in Figures 7 and 8. The reconstructed image is a visual representation of the information that the X-tree retains in its similarity vector.

IV. DISCUSSION

Concept acquisition is an important aspect of development and learning, clearly evident in the cognitive processes of humans and animals. In this work, concepts are defined as clusters of lower-level observations. While this definition

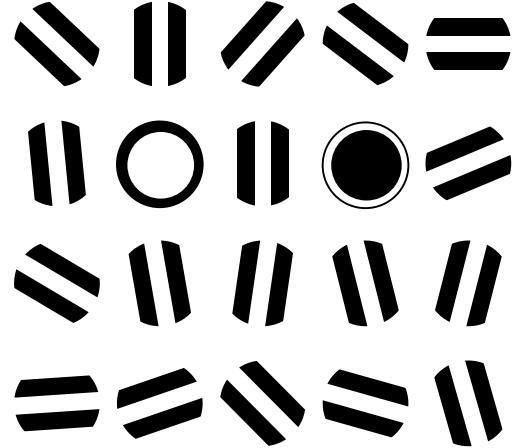


Fig. 6. Strongly oriented receptive fields that would be associated with each of the clusters in Figure 5.

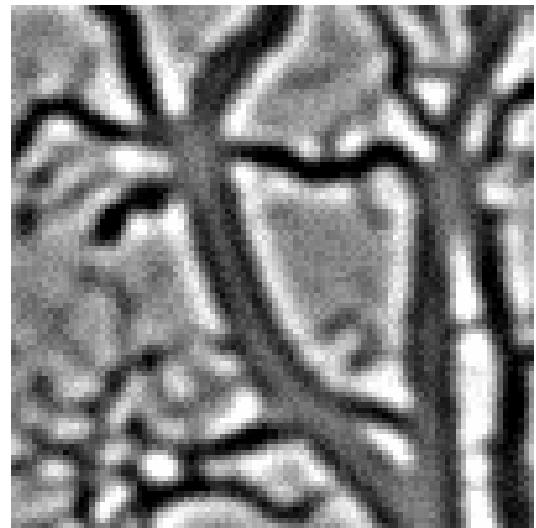


Fig. 7. An original grayscale image.

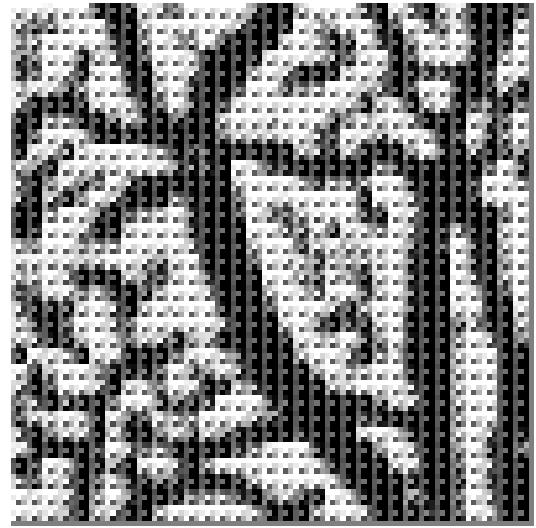


Fig. 8. A reconstructed image, based on the clusters.

does not satisfy all theories of what a concept is, it does provide a useful basis for beginning an investigation of the topic. Formulated this way, the problem of concept acquisition becomes a clustering or unsupervised learning problem. Acquired concepts can be high level and very abstract, or low level and quite specific. In either case, clusters of state observations are identified and used to represent the underlying states.

A. Biological motivation

The splitting criteria and cluster definitions have been chosen carefully to result in an algorithm with adaptive behavior. If an X-tree is exposed to the same types of observations for a significant period of time, it will converge on a stable representation, with a fixed number of clusters (concepts). However, if that exposure should shift later to focus on smaller portion of the state space, the density of observations in that region will increase, resulting in a finer conceptual representation locally.

B. Relation to other work

A decision tree is a natural representation for the iterative division of categories. Each leaf represents a terminal category. Each node represents both the super-category formed by combining its two children, as well as the decision boundary used to differentiate between them. (See Figure 1) Several tree-based adaptive partitioning algorithms have been previously proposed. The Parti-Game algorithm [9] uses a greedy controller to crawl through a partitioned state space. When the controller fails, the last visited subspace is divided. Parti-Game is specifically geared to geometric path planning; it assumes that all paths through the state space are continuous. The G Algorithm [10], U-Tree [11], Continuous U-Tree [12], AMPS [13], and decision trees of Pyeatt and Howe [14] are all approaches used in conjunction with dynamic programming methods or the popular temporal difference method, Q-learning [15], to estimate the value function across the state space. Whenever a subspace's value estimate is shown to be inadequate, the subspace is divided. The G Algorithm handles binary data, U-Trees handle discrete data, and Continuous U-Trees, AMPS, and Pyeatt and Howe's approach handle continuous data.

Taking a broader view, X-trees are a member of the set of unsupervised learning method, that is, methods that group data that is unlabeled and unclassified. Also referred to as clustering algorithms, there are many unsupervised learning methods developed with different sets of assumptions, but X-trees provide a novel collection of characteristics. X-trees are an on-line, hierarchical, divisive clustering algorithm. They are stable in the sense that cluster boundaries do not move. X-trees are most notable for what they don't assume. Unlike many clustering algorithms, X-trees operate without assuming 1) how many categories exist in the underlying data, 2) that prior probabilities of any category are known, 3) that prior probabilities are stationary, and 4) that the forms of the class conditional probabilities are known. Once the dimensionality

and range of the state space is defined, X-trees always start from the same initial conditions, so there is no need to carefully pick initial values for parameters. One of X-trees' greatest strengths is that it does not assume that the input space is well scaled or even linear. In fact, because X-trees use such a weak distance metric, they do not even assume that each of their input dimensions is an ordered field. As a result, X-trees can handle inputs consisting of enumerated types, such as $\{\text{vanilla} = 1, \text{chocolate} = 2, \text{strawberry} = 3, \text{pistachio} = 4\}$. And their computational demands are low.

X-trees' strength comes at a cost, of course. They bound clusters with hyperboxes, so clusters that are not well fit by a hyperbox may not be well represented. If M is large, X-trees can require a relatively large amount of data before useful clusters are generated. They are unstable in the sense that the total number of clusters created can grow without bound if not artificially limited. Like other unsupervised learning methods in which the number of clusters is not specified, the validity of the clusters found depends entirely on the appropriateness of the measure of cluster goodness, the node-splitting criterion in the case of X-trees. And, like other unsupervised learning algorithms of its class, there are no theoretical performance guarantees.

C. Relation to other biologically motivated methods

There are several other approaches to the problem of concept acquisition that are motivated by theories of human psychology and neuroscience. Hierarchical Temporal Memory (HTM) [16] uses on a hierarchy of Bayesian classification elements to cluster data into concepts ("causes") and to infer likely causes for a given set of inputs. Adaptive Resonance Theory (ART) [17] uses neural networks to bound clusters of states within hyperboxes. Like X-trees, both HTMs and ART are online methods, meaning that they can incorporate new state observations and update their results incrementally. However, HTMs and ART differ from X-trees in that they use a Euclidean distance metric to determine the similarity between states. For Euclidean distance to be valid, the state space must be well scaled, that is, a unit step in each dimension must represent the same magnitude of change. This is a somewhat restrictive condition that X-trees do not need to meet. Because X-trees do not use Euclidean distances (or any p -norm), they do not require well-scaled state spaces.

ACKNOWLEDGMENTS

This work was supported by the Laboratory Directed Research and Development program at Sandia National Laboratories. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000.

REFERENCES

- [1] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed. Wiley Interscience, 2001.
- [2] L. Devroye, L. Györfi, and G. Lugosi, *A Probabilistic Theory of Pattern Recognition*, ser. Applications of Mathematics: Stochastic Modeling and Applied Probability. New York: Springer-Verlag, 1996, no. 31.

- [3] N. Gilbert, M. den Besten, A. Bontovics, B. G. W. Craenen, F. Divina, A. E. Eiben, R. Griffioen, G. Hévízi, A. Lörincz, B. Paechter, S. Schuster, M. C. Schut, C. Tzolov, P. Vogt, and L. Yang, “Emerging artificial societies through learning,” *Journal of Artificial Societies and Social Simulation*, vol. 9, no. 2, 2006.
- [4] B. Rohrer, J. Morrow, F. Rothganger, and P. Xavier, “Concepts from data,” in *Proceedings of the Brain-Inspired Cognitive Architectures Symposium, AAAI Fall Symposium Series*, 2009.
- [5] W. T. B. Uther and M. M. Veloso, *Adaptive Agents in Multi-Agent Systems: Adaptation and Multi-Agent Learning*, ser. Lecture Notes in Computer Science. Springer, 2003, vol. 2636, ch. T Tree: Tree-based state generalization with temporally abstract actions, pp. 266–296.
- [6] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone, *Classification and Regression Trees*. Belmont, California: Wadsworth International Group, 1984.
- [7] J. R. Quinlan, *C4.5: Programs for Machine Learning*. Morgan Kaufmann, 1993.
- [8] L. Rokach and O. Maimon, *Data Mining with Decision Trees: Theory and Applications*, ser. Machine perception and artificial intelligence, H. Bunke and P. S. P. Wang, Eds. Singapore: World Scientific Publishing Co., 2008, vol. 69.
- [9] A. W. Moore and C. G. Atkeson, “The parti-game algorithm for variable resolution reinforcement learning in multidimensional state-spaces,” *Machine Learning*, vol. 21, pp. 199–233, 1995.
- [10] D. Chapman and L. P. Kaelbling, “Input generalization in delayed reinforcement learning,” in *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, 1991, pp. 726–731.
- [11] A. K. McCallum, “Reinforcement learning with selective perception and hidden state,” Ph.D. dissertation, University of Rochester, Computer Science Department, 1995.
- [12] W. T. B. Uther and M. M. Veloso, “Tree based discretization for continuous state space reinforcement learning,” in *Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI-98)*, Madison, WI, 1998.
- [13] M. J. Kochenderfer, “Adaptive modelling and planning for learning intelligent behaviour,” Ph.D. dissertation, University of Edinburgh, School of Informatics, 2006.
- [14] L. D. Pyeatt and A. E. Howe, “Decision tree function approximation in reinforcement learning,” in *Proceedings of the Third International Symposium on Adaptive Systems: Evolutionary Computation and Probabilistic Graphical Models*, 2001, pp. 70–77.
- [15] C. J. C. H. Watkins and P. Dayan, “Technical note: Q-learning,” *Machine Learning*, vol. 8, no. 3-4, pp. 279–292, May 1992.
- [16] D. George and J. Hawkins, “A hierarchical bayesian model of invariant pattern recognition in the visual cortex,” in *Proceedings of the IEEE Intl Joint Conference on Neural Networks*, vol. 3, 31 July-4 Aug 2005, pp. 1812–1817.
- [17] G. Carpenter and S. Grossberg, “A massive parallel architecture for a self-organizing neural pattern recognition machine,” *Computer Vision, Graphics, and Image Processing*, vol. 37, no. 1, pp. 54–115, 1987.