Explanation-Based Learning for Mobile Robot Perception

Tom M. Mitchell
Joseph O'Sullivan
Sebastian Thrun[1]

School of Computer Science
Carnegie Mellon University

Abstract

Explanation-based neural network learning (EBNN) has recently been proposed as a method for reducing the amount of training data required for reliable generalization, by relying instead on approximate, previously learned knowledge. We present first experiments applying EBNN to the problem of learning object recognition for a mobile robot. In these experiments, a mobile robot traveling down a hallway corridor learns to recognize distant doors based on color camera images and sonar sensations. The previously learned knowledge corresponds to a neural network that recognizes nearby doors, and a second network that predicts the state of the world after travelling forward in the corridor. Experimental results show that EBNN is able to use this approximate prior knowledge to significantly reduce the number of training examples required to learn to recognize distant doors. We also present results of experiments in which networks learned by EBNN (e.g., "there is a door 2 meters ahead") are then used as background knowledge for learning subsequent functions (e.g., "there is a door 3 meters ahead").

1. Introduction

One crucial problem for robot learning is scaling up. Whereas various research has shown how robots can learn relatively simple strategies from very little initial knowledge (e.g., [4, 6, 5, 9], theoretical and experimental results indicate that simple learning approaches will require unrealistic amounts of training data to learn significantly more complex functions. One approach to scaling up is to rely on human training (e.g., [2, 3, 4, 11]). We are interested here in methods by which the robot can use previously learned knowledge to reduce the need for new data in subsequent learning.

This paper presents early experiments with a robot architecture that utilizes previously learned knowledge about the effects of its actions, to reduce the difficulty of learning to recognize objects that it approaches as it travels through its environment. This architecture is based on explanation-based neural network learning (EBNN) [8, 15], an explanation-based algorithm that represents both prior knowledge and the current target function using neural networks. Initial experiments with EBNN demonstrated its ability to learn robot control knowledge in simulated [8] and real-world robot domains [14]. Here we examine its ability to reduce training data required for learning robot perception skills. These results show that EBNN is able to significantly reduce the amount of training data required for learning, compared to purely inductive methods such as Backpropagation, when an appropriate domain theory is available.

---

[1]Sebastian Thrun is currently at Bonn University.

2. Explanation-Based Neural Network Learning

Explanation-based neural network learning (EBNN) is a method for using approximate prior knowledge to improve the learner's ability to generalize correctly from limited training data. Prior knowledge consists of a collection of previously learned neural networks (the *domain theory networks*), with the function to be learned represented by an additional neural network (the *target network*). The EBNN algorithm uses the domain theory networks to analyze each training example, in order to extract information about the relevance of the different features of the example. This relevance information is then used, together with standard neural network induction, to constrain the weights of the target network. In this way, the target network is constrained both *inductively* by the training data, and *analytically* by the knowledge implicit in the domain theory.

How can the domain theory networks be used to analyze training examples so that information is extracted in a form useful for refining the target network? The key lies in the TangentProp algorithm [13], an extension to Backpropagation [12] that is able to adjust network weights to minimize the error in both the *values* and the *derivatives* of the function computed by the target network. Consider some target function, F, and a training example, X, consisting is a vector of components x_i. The TangentProp algorithm iteratively adjusts the weights of the target network, NET, to minimize an error measure containing two terms. The first error term is the difference between the target function value, F(X), and the value predicted by the network, NET(X). The second error term is the difference between the partial derivatives of the target function, $dF(X)/dx_i$, and the partial derivatives of the function represented by the target network, $dNET(X)/dx_i$. EBNN obtains estimates of target values for $dF(X)/dx_i$ by using its domain theory to analyze training examples.

EBNN analyzes a training example, X, by first using the domain theory to predict the value of F(X). This requires that it be possible to estimate the target function by chaining together the domain theory networks (as illustrated in Figure 4-1). The weights and activations of the domain theory networks in this computation are then examined, in order to analytically extract the partial derivative of F(X) with respect to each component feature x_i. Notice this derivative contains information about the relevance of each feature, x_i: if the value of x_i is irrelevant to the value of F it will have a derivative of 0, whereas features with large derivatives will be highly relevant. In this way, the information about feature relevance summarized in the domain theory networks is extracted from the analysis of the example, and used to provide the target derivatives to the TangentProp algorithm, which updates the weights of the target network. In order to minimize the impact of incorrect domain theories, the degree to which TangentProp seeks to fit these target derivatives is reduced by the degree to which the domain theory networks err in predicting the observed trainiing value, F(X).

The analytical component of learning in EBNN is similar to earlier explanation-based learning methods based on symbolic representations [1, 7]: given a target function, F, the domain theory is used as an alternative method for computing F, and the dependencies are extracted from this computation. The EBNN algorithm is described in greater detail in [8, 15], and summarized here in Figure 2-1.

For each training example:

1. *Explain* how the training example satisfied the target function. Chain together the domain theory networks to predict the value of the target function for the training example.
2. *Analyze* this explanation to determine feature relevance. Examine the weights and activations of the domain theory networks in the above explanation, to extract the partial derivatives of the explained target function value, with respect to each training example feature.
3. *Refine* the target function network. Update the target function weights to fit both the observed target function value (inductive component), and the target derivatives extracted from the explanation (analytical component).

Figure 2-1: Summary of the EBNN Algorithm.

3. The Problem

Here we consider the problem of learning to recognize objects that a robot is approaching as it travels through a hallway corridor. In this task, it is easy to recognize the object in predetermined locations and orientations. Forcing the recognizer to be invariant to robot movement adds significant complexity to this task. The experiments reported here involve learning functions of the form:

$$\text{Door-ahead}_d: S \rightarrow \{\text{True}, \text{False}\}$$

where $s \in S$ is the current robot sensor input (color vision and sonar readings), and Door-ahead_d is True if and only if, after traveling forward d meters in the corridor, the robot will be adjacent to a door. Notice this learning task is a special case of the more general problem: learning to recognize states in which a particular action will reliably lead to subsequent states with some desired property.

Figure 3-1 shows the mobile robot, Xavier [10], used in these experiments. The sensors include the color camera visible on top of the robot, and a ring of sonar sensors visible approximately one third of the way down the robot torso. Although the camera is mounted on a pan/tilt unit controllable by the robot, in these experiments the camera was maintained in a single position pointed 15 degrees downward and 30 degrees to the right, as the robot traveled straight down the center of the corridor.

In order to limit the complexity of the sensor data for this learning task, Xavier summarizes its vision and sonar sensations in a vector of just 25 sensor readings, as shown in Figure 3-2. A row of 10 rectangular regions is extracted from the image, as highlighted in the figure. For each such rectangular region, two values are computed: the average red intensity and the average blue intensity. Thus, the image is reduced to a vector of just 20 intensity values, as displayed at the bottom right of the figure. Here, the two leftmost vector values correspond to the red and blue intensities of the leftmost rectangle, and so on. The area of the white box in the vector diagram

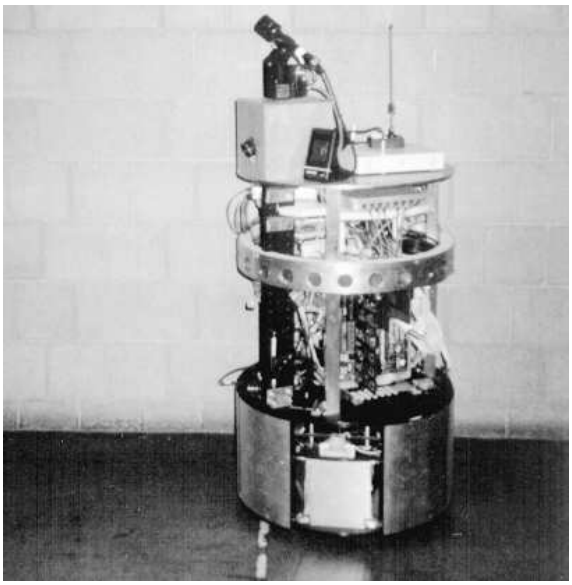

Figure 3-1: The CMU Learning Lab Robot, Xavier.

indicates the magnitude of the corresponding vector value. The sonar data is also subsampled, to include only the 5 sonar readings shaded in dark gray in the sonar display. These are the sonars that face somewhat forward and toward the right, and therefore can be expected to contain relevant information for recognizing upcoming doorways on the right. The length of these sonar values (clipped to a maximum of 300 cm) is displayed in the row of 5 white boxes directly under the sonar display in the figure. Together, the 5 sonar readings, and 20 "fat pixels" from the color image, form the summary of robot sensations used in these experiments.

4. Approach

Consider the task of learning the target function Door-ahead$_1$: S $\rightarrow$ {True, False}, whose value is True when there is a door 1 meter ahead. One way to learn this function is to inductively learn a neural network whose inputs are the vector of 25 sensor readings, and whose single output indicates the boolean value of the function. Given a significant number of training examples, such a network can be learned to a reasonable accuracy using the Backpropagation algorithm [12]. We are interested here in comparing such purely inductive methods with EBNN.

To apply EBNN to the task of learning Door-ahead$_1$, we must have available prior knowledge capable of explaining observed training examples of the target function. We therefore allow Xavier to first learn two neural networks. The first represents the function

Forward-1M: S $\rightarrow$ S

which, given the current state (represented by the vector of 25 sensor readings) predicts the

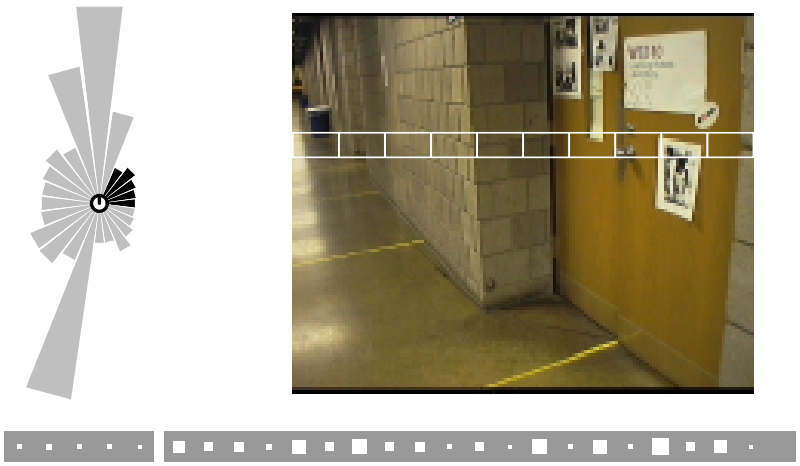

Figure 3-2:

Typical sensor input and its representation. *This image was collected by Xavier while it proceeded forward down the center of a corridor. The circle of gray bars on the left indicate the 24 sonar distance echoes detected by Xavier's sonars at the same time the camera image was collected. The long sonar reading extending toward the top of the figure corresponds to the sonar facing forward down the corridor, and the shorter readings to the side reflect the position of the wall. This sonar and image data is subsampled and summarized in the vector of 25 values displayed at the bottom of the figure, as explained in the text.*

sensor readings (next state) that can be expected if the robot drives 1 meter forward in the

corridor. Notice that this function is independent of the particular perception learning task, and therefore the cost of learning it can be amortized over a variety of target functions.

The second domain theory network represents the function

Door-Here?: S $\rightarrow$ {True, False}

which is True only for states in which the robot is directly adjacent to a door. Although the function Door-Here? is specific to recognizing doors, it is significantly easier to learn than the target function which requires recognizing doors at a distance.

Both of these functions are learned (approximately) from previous experience, in this case using the purely inductive Backpropagation algorithm. These two learned networks are then used by EBNN to explain individual training examples of the target function Door-ahead$_1$, as illustrated in Figure 4-1.

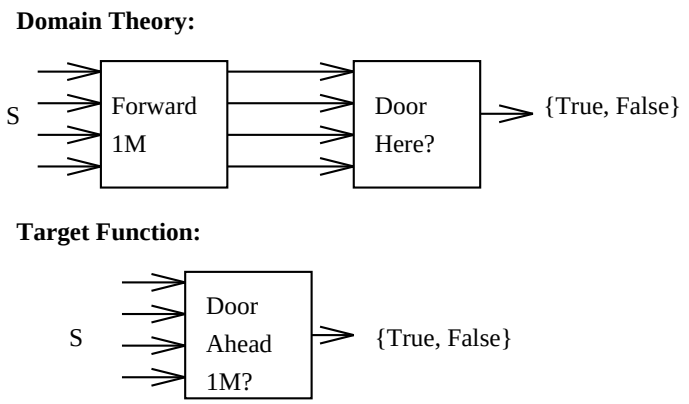

Figure 4-1:

A Typical EBNN Explanation. *Each training example of the target function, Door-ahead$_1$, is explained in terms of the previously learned functions Forward-1M and Door-Here?.*

To illustrate the application of EBNN, consider analyzing a single training example. Figure 4-2 shows the explanation for a single training example of the target function -- a particular state, S_{902}, for which there is no door one meter ahead. The training example state is represented by the vector of 25 sensor readings on the left of the diagram. The first domain theory network, applied to this state, predicts the state one meter forward. The second domain theory network is then applied to this predicted state, to estimate whether a door is present in this predicted next state. Taken together, these predictions form EBNN's explanation of the training example.

EBNN uses this explanation to determine the relevance of each feature of the training example, in order to guide its learning of the target function. More specifically, it uses the explanation to compute the partial derivative of the target function with respect to each training example feature. Irrelevant inputs (such as camera pixels that provide no information about doors on the

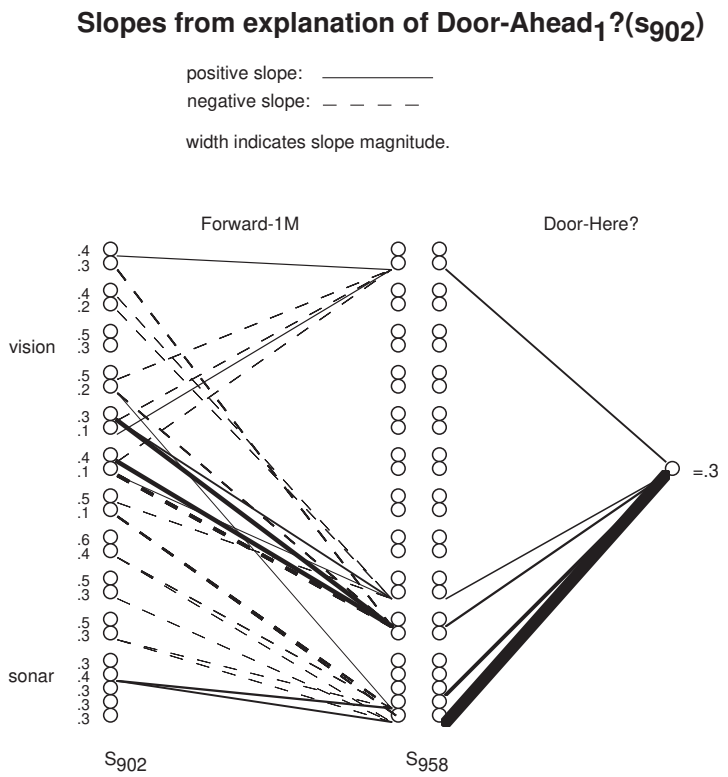

Figure 4-2:

Derivatives Extracted from an Explanation. *Sensed state S_{902} is a training example for which target function Door-Ahead$_1$ is False. EBNN explains this value for the target function by using its previously learned neural networks Forward-1M, and Door-Here?. The derivatives of these individual networks, indicated by solid and dashed lines in the diagram, represent information about the presumed relevance of network input features to the network output. Derivatives of individual networks are combined to compute the derivative of the entire explanation. For example, the derivative of the output of Door-Here?, with respect to the second sonar value of S_{902} is .48.*

right) will have partial derivatives of zero (provided the domain theory is correct!). In contrast, inputs with large derivatives are known to be highly relevant. The partial derivatives of the explanation output with respect to each input (i.e., each training example feature) are computed by applying the chain rule to the derivatives of each individual network in the explanation. The lines in Figure 4-2 indicate the sign and magnitude of the most significant derivatives for each individual network in the domain theory. Consider, for example, the rightmost network in the explanation: Door-Here?. The thickest solid line indicates that the output of this network increases rapidly with the value of its bottommost input (the fifth sonar). Because doors in this corridor are typically recessed, a long sonar reading in this direction is a strong indicator that the robot is directly by a door. Similarly, the knowledge captured in the first network indicates that

this fifth sonar reading depends positively on the second sonar value of the initial state. This is reasonable, because as the robot drives forward, the object that provided the second sonar echo moves into the field of view of the fifth sonar. Coupled together, these two dependencies indicate that the larger the second sonar value in the current state, the more likely that there is a door one meter ahead.

Using the chain rule, EBNN combines the matrix of partial derivatives from each step in the explanation to compute the derivative of the target function with respect to each training example feature. It then updates the weights of the target network, to fit both these target slopes, and the observed training value. The former provides an analytical component of learning, transferring knowledge from the domain theory networks into the target network. The later provides an inductive component, tuning the weights of the target network to fit the observed value independent of prior knowledge.

5. Experimental Results

This section considers two questions. First, can EBNN generalize more correctly than purely inductive methods, by relying on its previously learned knowledge? Second, can knowledge acquired by EBNN be successfully built upon, by using it as part of the domain theory for learning subsequent target functions?

The experimental setup was as follows: Xavier was first allowed to travel an entire corridor, with one region held aside to test generalization capability (see Figure 5-1). From this experience, it acquired 403 sensor "snapshots", one approximately every 12 cm. The correct value of Door-Here? was manually provided for each snapshot. The Forward-1M and Door-Here? domain theory networks were then trained by applying the inductive Backpropagation algorithm to this data.

5.1. Generalization Accuracy

To explore the generalization accuracy of EBNN and Backpropagation, we trained the robot to learn the target function Door-ahead$_1$, while varying training set size. Training sets were composed of N contiguous examples, the final examples being directly adjacent to one of the doors. Values of N were 5, 10, 15, 20, 25, 30, 35, 40, 45, and 52. Given the spacing between successive examples and between doors, only one door was in view over the entire set. Positive examples of Door-ahead$_1$ were examples 10 through 21, inclusive.

Backpropagation and EBNN were trained on identical data sets using a learning rate of 0.1, and cross validation to determine the number of training epochs. Generalization accuracy was then measured by applying the learned networks to the evaluation set (see Figure 5-1).

Figure 5-2 summarizes the result of this experiment, plotting generalization accuracy of both learning methods against the number of training examples available. As can be seen, EBNN generalizes more accurately than the inductive Backpropagation algorithm over the entire range of training set size. It is interesting to compare the performance of the two approaches on very small sets of training data. The first five training examples were all negative examples in this case. As a result, Backpropagation learned a network that always predicted there is no door ahead. This lead to an accuracy of 79% on the evaluation data set. Surprisingly, EBNN was

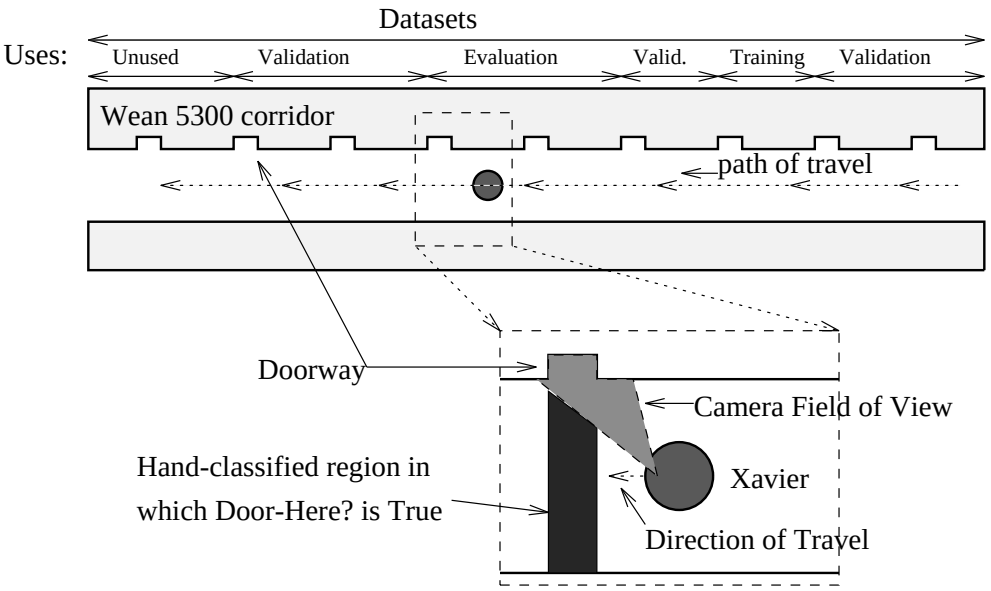

Figure 5-1:

Corridor Enviroment *For these experiments, Xavier traversed an entire corridor, gathering samples approximately every 12cms. Each sample included sonar and vision information. These samples were then hand-classified and divided into 3 distinct sets dependant on location along the path; a set from which training sets were extracted, a evaluation set to test performance, with the remaining samples being used for validation*

able to learn a more accurate network from this data, because its use of the dependencies extracted from its explanations captured the information that changes to certain features values would make it more likely that a door would be present, even though none were present in the training examples. Notice the dip in performance for both methods at 20 examples. At this point, the data was particularly misleading because the heavy number of positive examples was misrepresentative of their distribution in the evaluation set.

5.2. Building on Acquired Knowledge

To test the ability of EBNN to learn progressively more difficult function and to bootstrap itself by using its earlier learned networks as domain theory for subsequent learning, we extended the previous experiment by additionally learning Door-ahead$_2$, Door-ahead$_3$, and Door-ahead$_4$.

The experimental setup was as above, using the same set of training examples. The explanation used by EBNN to learn Door-ahead$_d$ was formed by chaining together the networks for Forward-1M, and Door-ahead$_{d-1}$. For example, EBNN explained examples of Door-ahead$_2$ by first predicting the state one meter ahead (using the Forward-1M network), then estimating

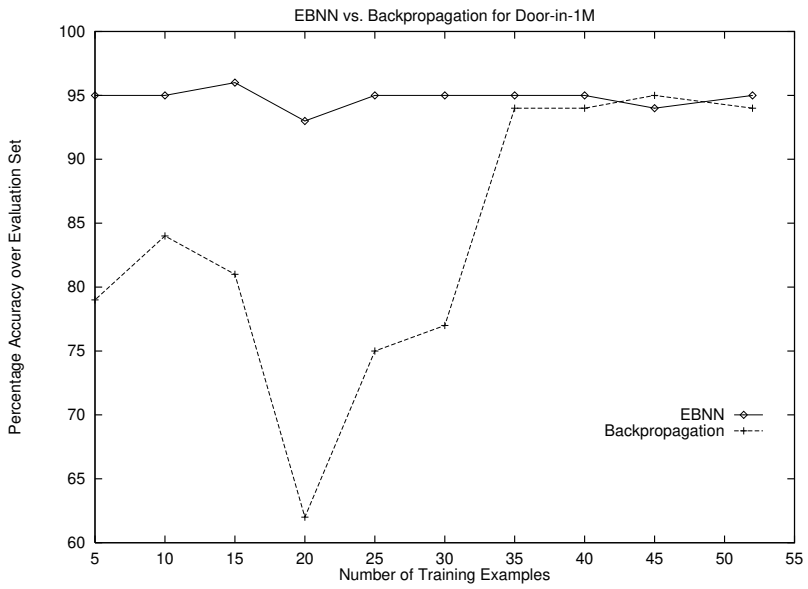

Figure 5-2:

Generalization accuracy for Door-Ahead$_1$. *The dotted line indicate the performance of purely inductive Backpropagation. The solid line indicate the performance of EBNN.*

whether a door was present one meter ahead of that state (using the Door-ahead$_1$ network learned in the first experiment).

Figure 5-3 shows the generalization accuracy after training on 35 examples, for Door-ahead$_d$, where d=1,2,3 and 4 meters[2]. As in the first experiment, results are shown for both EBNN and Backpropagation. As the figure shows, the generalization accuracy decreases for both methods as *d* increases. EBNN generalizes more accurately than Backpropagation for each of these learning tasks.

6. Conclusions

This paper presents first results applying explanation-based neural network learning (EBNN) to the problem of mobile robot object recognition. These initial experiments demonstrate the ability of EBNN to use approximate, previously learned knowledge about the effects of its actions (traveling forward down a corridor) and about the appearance of a very close object (being directly in front of a door), to guide the learning of neural networks to recognize the object from further distances. The ability of EBNN to utilize information from its approximate

[2]In this draft of the paper, the generalization accuracy for this plot is evaluated over all scenes from the corridor. In the final copy, this will be replaced by a plot evaluating performance only over a test set.

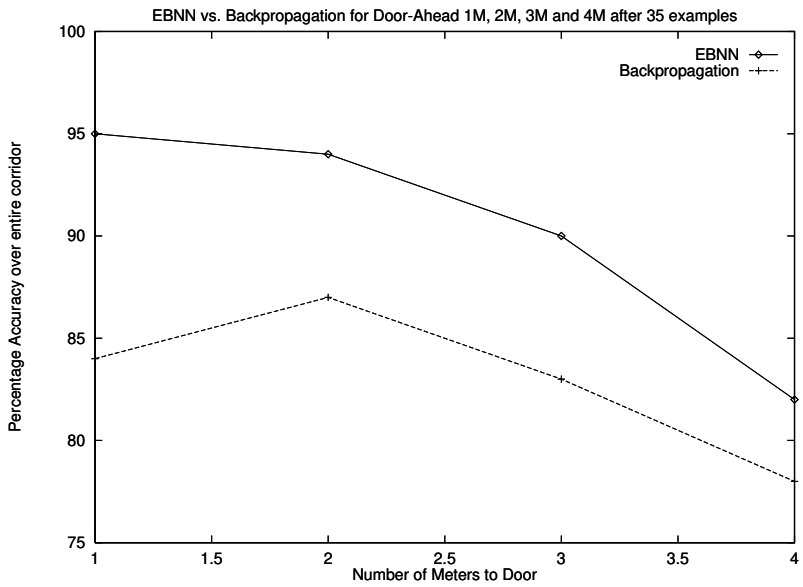

Figure 5-3:

Generalization accuracy Door-Ahead$_d$. *Each point corresponds to a different learning task: learning Door-Ahead$_d$ for some value of d in meters. For each experiment, 35 training examples were provided.*

domain theory is demonstrated by the fact that it outperforms the standard inductive neural network learning method, Backpropagation, across a variety of data set sizes, and a family of related target functions. While further experimentation is warranted, these results suggest the potential role of EBNN in helping to scale up robot learning by moving away from purely inductive, tabula rasa, learning techniques, toward techniques that bootstrap themselves by leveraging previously learned knowledge to simplify subsequent learning.

7. Acknowledgments

This research is sponsored in part by the National Science Foundation under award IRI-9313367, and by the Wright Laboratory, Aeronautical Systems Center, Air Force Materiel Command, USAF, and the Advanced Research Projects Agency (ARPA) under grant number F33615-93-1-1330.

References

1. DeJong, G., and Mooney, R. "Explanation-Based Learning: An Alternative View". *Machine Learning 1*, 2 (1986), 145-176.

2. Katsushi Ikeuchi and Takashi Suehiro. Towards an assembly plan from observation. Proceedings of IEEE Int. Conf. on Robotics and Automation, May, 1992.

3. J. Laird, E. Yager, C. Tuck, and M. Hucka. Learning in tele-autonomous systems using Soar. Proceedings of the 1989 NASA Conference of Space Telerobotics, 1989.

4. Long-Ji Lin. *Reinforcement Learning for Robots using Neural Networks*. Ph.D. Th., Carnegie Mellon University, School of Computer Science, Pittsburgh, PA 15232, 1993. Technical Report CMU-CS-93-103.

5. Pattie Maes and Rodney Brooks. Learning to Coordinate Behaviors. Proceedings of AAAI-90, Boston, August, 1990.

6. Sridhar Mahadevan and Jonathan Connell. Scaling Reinforcement Learning to Robotics by Exploiting the Subsumption Architecture. Proceedings of Machine Learning Workshop '91, July, 1991.

7. Mitchell, T.M., Keller, R.K., and Kedar-Cabelli, S. "Explanation-Based Generalization: A Unifying View". *Machine Learning 1*, 1 (1986).

8. Tom M. Mitchell and Sebastian B. Thrun. Explanation-Based Neural Network Learning for Robot Control. In *Advances in Neural Information Processing Systems 5*, J. E. Moody and S. J. Hanson and R. P. Lipmann, Ed., Morgan Kaufmann, 1993.

9. Andrew W. Moore. *Efficent Memory-based Learning for Robot Control*. Ph.D. Th., Trinity Hall, University of Cambridge, England, 1990, 1990.

10. Joseph O'Sullivan. Xavier Manual. Carnegie Mellon University, Learning Robot Lab Internal Document - contact josullvn@cs.cmu.edu.

11. Dean Pomerleau. *Neural Network Perception for Mobile Robot Guidance*. Ph.D. Th., Carnegie Mellon University, School of Computer Science, Pittsburgh, PA 15232, 1992. Technical Report CMU-CS-92-115.

12. David E. Rumelhart, Geoffery E.Hinton and Ronald J. Williams. Learning internal representations by error propagation. In David E. Rymelhart and J. L. McClelland, Ed., *Parallel Distributed Processing. Vol I+II*, MIT Press, 1986.

13. Patrice Simard, Bernard Victorri, Yann LeCun, and John Denker. Tangent prop -- a formalism for specifying selected invariances in an adaptive network. In *Advances in Neural Information Processing Systems 4*, J. E. Moody and S. J. Hanson and R. P. Lipmann, Ed., Morgan Kaufmann, 1992, pp. 895--903.

14. Sebastian B. Thrun. An Approach to Learning Robot Navigation. Proceedings of the IEEE/RSJ/GI International Conference on Intelligent Robots and Systems, September, 1994. (to appear).

15. Sebastian B. Thrun, and Tom M. Mitchell. Integrating inductive neural network learning and explanation-based learning. Proceedings of IJCAI-93, IJCAI, Chamberry, France, July, 1993.