2015 IEEE 14th International Conference on Machine Learning and Applications

Learning Context-based Outcomes for Mobile Robots in Unstructured Indoor Environments

Priyam Parashar
Georgia Tech University
Atlanta, GA 30332
Email: pparasha@andrew.cmu.edu

Robert Fisher
Carnegie Mellon University
Pittsburgh, PA 15213
Email: rwfisher@cs.cmu.edu

Reid Simmons
Carnegie Mellon University
Pittsburgh, PA 15213
Email: reids@cs.cmu.edu

Manuela Veloso
Carnegie Mellon University
Pittsburgh, PA 15213
Email: mmv@cs.cmu.edu

Joydeep Biswas
University of Massachusetts Amherst
Amherst, MA 01003
Email: joydeepb@cs.umass.edu

***Abstract*—We present a method to learn context-dependent outcomes of behaviors in unstructured indoor environments. The idea is that certain features in the environment may be predictive of differences in outcomes, such as how long a mobile robot takes to traverse a corridor. Doing so enables the robot to plan more effectively, and also be able to interact with people more effectively by more accurately predicting when its plans may take longer to execute or may be likely to fail. We use a node-and-edge based map of the environment and treat the traversal time of the robot for each edge as a random variable to be characterized. The first step is to determine whether the distribution of the random variable is multimodal and, if so, we learn to classify the modes using a hierarchy of plan-time features (e.g., time of the day, day of the week) and run-time features (observations of recent traversal times through other corridors). We utilize a cascading regression system that first estimates which mode of the traversal distribution we expect the robot to observe, and then predict the actual traversal time through a corridor. On average, our method produces a mean residual error of less than 2.7 seconds.**

I. INTRODUCTION

Technology has come to a point where cohabitation of living space by humans and machines has become almost commonplace. Examples include self-help kiosks in the supermarket, hospitals and airports, robot receptionists to help people find directions, and robot vacuum cleaners. The next step is when these machines take on a more mobile form for cooperative tasks. Such intelligent mobile robots are being increasingly employed in places such as offices, hospitals, and other institutions. These machines are termed 'service robots', which help take burden off of human beings by taking care of menial tasks involved with the daily workings of a place. For example such robots can handle routine tasks such as automating daily delivery of medicines, taking care of immediate requests for articles, etc.

A major hurdle for such machines is that of path planning or task scheduling in dynamic, unstructured environments. These robots operate in places dominantly inhabited by people, and the presence and activities of people can affect the outcomes of robot actions, such as how long it takes to traverse a corridor, or how the expected time needed to wait for an elevator. It is important to model these outcomes accurately for two main reasons: First, to help the robots create more efficient plans and second to coordinate with people in case the robot predicts it will be late in completing its current task. For example, suppose the shortest path from the office lobby to the HR department involves going through the cafeteria. If a delivery is scheduled around lunch time, it would be advisable for a robot to avoid this crowded junction. Similarly, if a hospital delivery robot finds that one corridor unexpectedly takes very long to traverse, and it knows that means subsequent corridors are also likely to take long, it might decide to notify the nurse that it will likely be delayed.

It seldom happens, however, that these context-dependent action outcomes are programmed, or even known, a priori. Thus, we are investigating learning such models from experience, finding patterns that can be used at plan time to minimize the expected task execution time and at run time to detect potential anomalies – outcomes that differ significantly from prior expectations.

The reasoning behind our intuition lies in the very working order of these places. Human beings follow certain patterns in their daily lives, be it social or workplace related. Especially in organizations such as offices and educational institutions there are certain guidelines to be followed. For example, food courts are always crowded during lunch and breakfast times, in educational establishments the corridors near a lecture hall are always crowded when a lecture ends, etc. These routines result in patterned movement across spaces, which in turn may affect the robot's performance. Our work is to learn, and eventually use, these models to improve the robot's performance.

Our approach involves collecting a stream of data that the robot records as it execute tasks throughout the day. We aggregate data from similar actions (e.g., traversing a particular segment of the corridor) and model the action outcome (e.g., time to traverse) as a random variable. The first step is to analyze the data to determine whether the variable's distribution is unimodal or multimodal and, if the latter, how many modes it has and what are their distributions (assuming a Gaussian mixture model). The data, and the modes, are then fed into a classifier that both predicts the mode and uses regression techniques to predict the actual travel time. Based on just the timestamps (time of day, day of week), we see mode prediction accuracy of around 97% and mean time prediction residuals of 2.6 seconds.

Our model utilizes correlations between contiguous events

978-1-5090-0287-0/15 $31.00 © 2015 IEEE
DOI 10.1109/ICMLA.2015.222

in a sequence of action (such as two nearby corridors whose long and short traversal periods correlate) and use these as additional "medium-range" features. Note that while these features cannot be used at plan time, if we find good correlations between action outcomes then, at run time, unexpected outcomes for one action can be used to improve the predicted action outcomes of correlated actions. This can improve the overall estimate of the execution time of the remaining task, which can be used to trigger contingency plans (such as taking alternate routes or notifying a person that the robot will be delayed).

II. Related Work

Prediction and planning with contextual information is a well-studied topic in the field of robotics. Since the mid 1990s, researchers have been working towards more reactive and proactive models of machines [1]. In recent years, increased instrumentation in consumer technology has driven new research into contextual prediction across a variety of domains, from mobile robots to consumer smartphones and sensor equipped power wheelchairs. Some of this work has leveraged decision-theoretic prediction algorithms [2], [3], while other work has formulated these tasks as a sequential decision making process, leveraging spectral latent variable models [4].

In this paper, we model a qualitative indoor environment in which we observe dynamic movements through the space. Similar work has been done by Haigh et al [5], where they try to leverage observed patterns to find context-dependent map costs. Focusing specifically on the problem in open-ended indoor environments the work done by Bennewitz et al [6] and Kruse et al [7] is relevant as they try to model robot behavior while keeping human dynamics in mind, which is a key factor for success in cohabited environments. More recently, Sehestedt et al [8] looks at the social context in these environment apart from simple human trajectories and tries to deploy a minimally invasive trajectory for service robots.

III. Approach

In this work, our task is to use contextual information to predict how long a robot will require to traverse a path through the environment. We utilize real world data collected over several years by the CoBot project [9], [10] to train and test our methods.

The world map is represented in the form of nodes and connecting edges. These nodes represent landmarks from the real world, such as corridor junctions, the beginning of stairways, etc. and the edges represent the paths connecting them. Our method first tries to isolate those edges of the environment which show multimodal distributions in traversal time (see Fig. 1), like the corridors leading up to lecture halls which are most crowded when lectures end but not otherwise. Some edges demonstrate simple unimodal distributions of traversal time. Estimating travel time on these edges is trivial, and thus we do not consider these edges in our evaluation. We model the multimodal edges to learn feature dependent dynamics of the subspace such as different traversal times associated with different hours of the day. Our feature space consists of the observable data available to a robot during its daily activities including time stamps, sensor data, immediate history, etc. Following a Markovian assumption, we model these multimodal edges, assuming a mixed gaussian distribution, based on the information sampled at the beginning of each edge traversal. We train a classifier from these feature traces, which learns the context-dependent patterns allowing us to make *a priori* predictions for future edge traversals.

A. Identifying Multimodal Edges

With traversal time history for each edge as the variable sample, we use a combination of Mean-shift and Expectation Maximization algorithms to find the generative distribution for it. In order to find multiple modes, we begin by running Gaussian Mean-shift clustering on the data. If clustering results in more than one density center, we use the Expectation Maximization algorithm to fit the traversal time distribution to a gaussian mixture model with number of components equal to the number of density centers identified. We assess the mixture model to make sure that the components are statistically significant contextual-modes by looking at the component weights and distribution over sample points.

Mean-shift [11] is a non-parametric mode-seeking algorithm. It is an iterative algorithm which shifts towards the density center using a kernelized window. In this case we have used a gaussian kernel. The expectation-maximization (EM) algorithm is an iterative method for finding maximum likelihood or maximum a posteriori (MAP) estimates of parameters in statistical models, where the model depends on unobserved latent variables, i.e. gaussian component parameters in our case. It involves two steps, in the E step a function for the expectation of the log-likelihood for the model Z given X using the current estimate for the parameters θ is created.

$$Q(\theta|\theta^t) = E_{Z|X,\theta^t} log[\theta; X, Z]$$

In the M step, parameters for maximizing this log-likelihood function are computed.

$$\theta^{t+1} = \arg\max_{\theta} Q(\theta|\theta^t)$$

B. Learning Contextual Patterns

We employ supervised learning using the mode of each point as labels for feature vectors. We use decision trees to learn the feature patterns corresponding to each mode. Decision trees follow a greedy mechanism, dividing the dataset based on features which result in maximum information gain to reduce the entropy. Information entropy of a random variable X under probability mass function $P(X)$ is defined as,

$$H(X) = E[-ln(P(X))]$$

Information gain for a feature or attribute a, which splits sample set X is given as,

$$IG(X, a) = H(X) - H(X|a)$$

In order to predict the exact traversal time we train regression trees which use predicted modality of an edge as part of its feature space. Regression trees are trained in the same manner as decision trees, but the leaves correspond to real values rather than discreet labels. Regression trees are generally much larger than decision trees, and require much more labeled data.

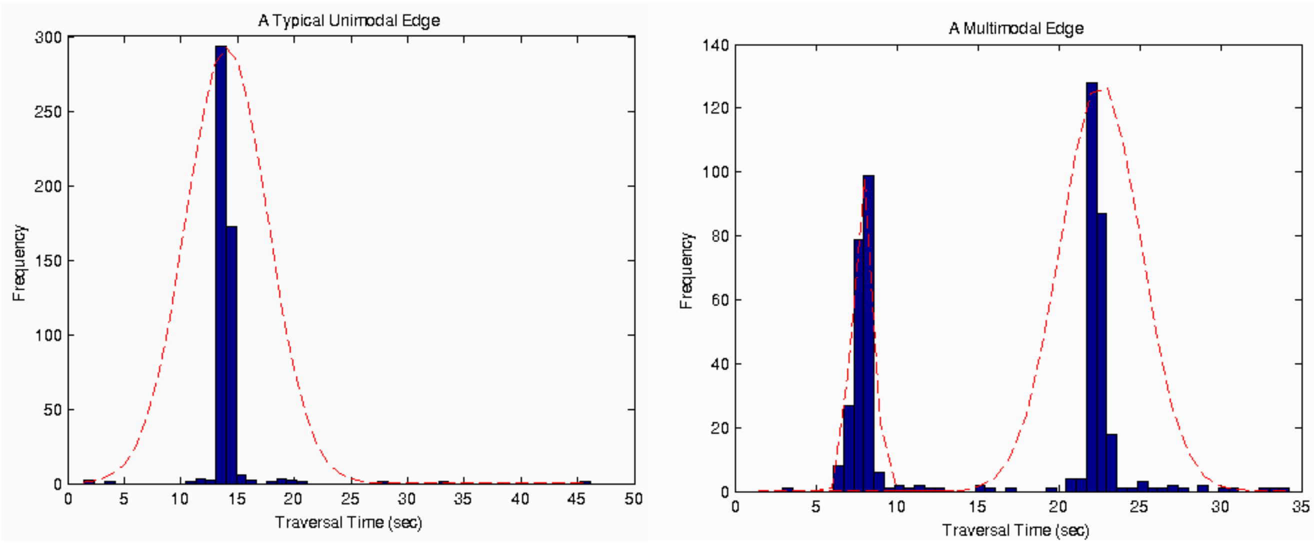

Fig. 1. Graphical Contrast between Unimodal and Multimodal Edges

C. Data

The data that we worked with is associated with CoBot, which is an indoor service robot. The data comprises of all the corridor traversals made by CoBot in different spaces of Carnegie Mellon University since January of 2013, which is approximately one and a half year worth of data. The data is a typical example of patterns associated with an educational institute.

The features that we used as predictors for decision and regression tree have been divided into two categories. First is long range feature, which include temporal data specifically year, month, date, day of the week and time of the day. The second category is that of medium range features, which uses recent traversal history of the agent restricted by an upper bound of being 30 minutes old.

IV. RESULTS

For all experiments presented in this section, we sample from the full dataset to generate training sets of various sizes. We begin with training sets of size 100, and increase this size by increments of 100 until a training set of size 3,400 is created. A disjoint testing set consisting of 457 points is also selected to accompany each training set. This process is repeated with different random for 1,000 iterations, giving us a total of 34,000 evaluations for each of our models.

A. Predicting Traversal Modality

Figure 2 shows the learning curve for predicting traversal time modalities using a decision tree. For reference, a naive classifier that always selected the larger mode in the travel distributions would achieve an accuracy of 63.3%. An empirical 90% confidence interval for the classifier accuracy is also shown. This classifier was trained using timestamp features, hallway ID, and recently observed traversal times in other hallways. The maximum mean classification accuracy achieved is 97.21%.

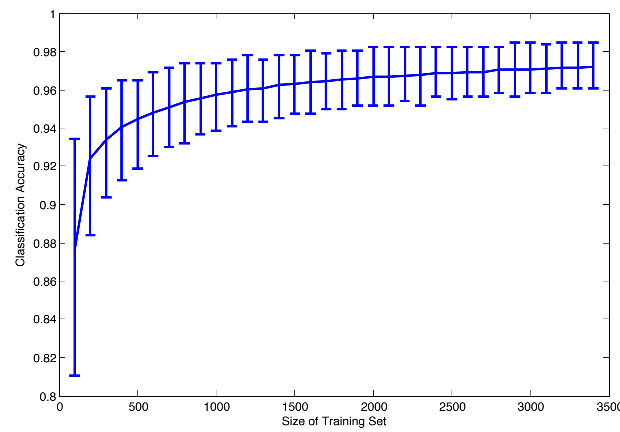

Fig. 2. Empirical results for traversal mode prediction

B. Regression for Traversal Time

Fig. 3 shows the learning curve for two regression models, compared to a baseline model. The baseline model simply computes the mean traversal time for each hallway in the training set, and uses these value to make predictions in the testing set.

The uncorrelated regression tree model shown in Figure 3 was trained using only time stamp information and corridor IDs. The second regression tree was trained with these features, as well the predicted modality produced by the decision tree classifier, and the traversal time of the most recent trip, if the last traversal observation occurred less than 30 minutes prior—this model is detonated as the correlated regression tree. Figure 3 also shows the 90% empirical confidence interval for the correlation regression tree, which indicates the centered 90% quantile of the residual means computed over the 1000 training iterations. With 3400 training points the correlated regression tree achieves an mean residual value of 2.66 seconds.

Corridor	Predicted Traversal Time	Actual Traversal Time
Corridor 5	16.65	17.601
Corridor 10	16.642	46.451
Corridor 4	4.4139	4.6639
Corridor 5	2.269	2.4125
Corridor 10	12.731	12.521
Corridor 7	4.4888	6.4325
Corridor 10	16.782	17.681
Corridor 5	16.241	16.161
Corridor 11	13.511	14.712
Corridor 1	9.8414	10.385
Corridor 10	17.309	18.6

TABLE I. PREDICTED VERSUS ACTUAL TRAVERSAL TIMES

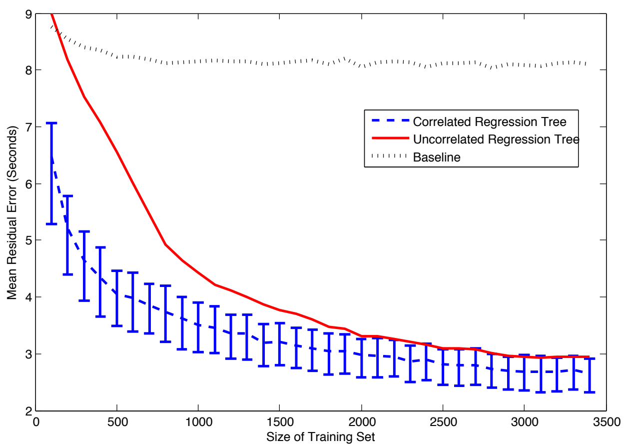

Fig. 3. Empirical results for traversal time regression.

Table I also shows the predicted and actual traversal times for a small sample of the dataset produced using the correlated regression tree.

V. DISCUSSION

A. Identifying Multimodal Corridors

By preprocessing the edge specific edge data first, we were able to eliminate over 90% edges as unimodal and ppinpoint specific areas in the environment where context-dependency was significant. Out of 263 corridors, we identified 12 as being multimodal, thus significantly reducing computation state-space. Apart from this, we think by knowing the status of the corridors around as unimodal or multimodal, new areas of environment being explored can be generalized more accurately thanks to spatial proximity. Also, identifying areas with contextual-dependencies in the environment can help the developers with diagnosis by giving them a more qualitative idea of the exact geographical locations which might be contributing to unwanted or abnormal behavior.

B. Prediction with Regression

We see in figure 3 that both regression models significantly outperform the naive baseline, and the correlated regression tree performs particularly well with very little data. Given enough data (2000 or more points) the uncorrelated regression tree nears the upper 95% percentile of the correlated regression tree. This indicates that with enough data we can accurately predict traversal times fairly well using only plan time features. These estimates can then be refined at travel time using observed traversal times to reduce the average prediction residual by roughly 0.5 seconds.

It is worth noting that the majority of the regression error is caused by a small number of unusual outliers, in which the robot experienced significant unexpected delays during transit, as seen in row 2 of Table I. We have correctly identified some of these outliers using information from recent observations with the correlated regressor. For instance if there is unexpected traffic due to seminar or special event, we will discover this once we see delays in adjacent corridors. However, other unexpected delays such as mechanical failure or disruptive individuals interfering with the robot's operation cannot be predicted using the factors we have described in this work. To discover this sort of unexpected delay, a model would need to incorporate information from onboard sensors or feedback from human observers.

VI. CONCLUSIONS

In this work, we have presented methods for improving plan time and execution time behavior of a mobile robotic system. In particular, we are able to account for contextual and temporal factors that can affect the normal performance of the robot.

REFERENCES

[1] R. C. Arkin, *Behavior-based robotics*. MIT press, 1998.

[2] R. Fisher and R. Simmons, "Smartphone interruptibility using density-weighted uncertainty sampling with reinforcement learning," in *Machine Learning and Applications and Workshops (ICMLA), 2011 10th International Conference on*, vol. 1. IEEE, 2011, pp. 436–441.

[3] S. Rosenthal, A. K. Dey, and M. Veloso, "Using decision-theoretic experience sampling to build personalized mobile phone interruption models," in *Pervasive Computing*. Springer, 2011, pp. 170–187.

[4] R. Fisher, R. Simmons, C.-S. Chung, R. Cooper, G. Grindle, A. Kelleher, H. Liu, and Y. K. Wu, "Spectral machine learning for predicting power wheelchair exercise compliance," in *Foundations of Intelligent Systems*. Springer, 2014, pp. 174–183.

[5] K. Z. Haigh and M. M. Veloso, "Learning situation-dependent costs: Improving planning from probabilistic robot execution," *Robotics and Autonomous Systems*, vol. 29, no. 2, pp. 145–174, 1999.

[6] M. Bennewitz, W. Burgard, and S. Thrun, "Adapting navigation strategies using motions patterns of people," in *Robotics and Automation, 2003. Proceedings. ICRA'03. IEEE International Conference on*, vol. 2. IEEE, 2003, pp. 2000–2005.

[7] E. Kruse, R. Gutsche, and F. Wahl, "Acquisition of statistical motion patterns in dynamic environments and their application to mobile robot motion planning," in *Intelligent Robots and Systems, 1997. IROS '97., Proceedings of the 1997 IEEE/RSJ International Conference on*, vol. 2, Sep 1997, pp. 712–717 vol.2.

[8] S. Sehestedt, S. Kodagoda, and G. Dissanayake, "Robot path planning in a social context," in *Robotics Automation and Mechatronics (RAM), 2010 IEEE Conference on*, June 2010, pp. 206–211.

[9] S. Rosenthal, J. Biswas, and M. Veloso, "An effective personal mobile robot agent through symbiotic human-robot interaction," in *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 2010, pp. 915–922.

[10] J. Biswas and M. Veloso, "Wifi localization and navigation for autonomous indoor mobile robots," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*. IEEE, 2010, pp. 4379–4384.

[11] D. Comaniciu and P. Meer, "Mean shift: A robust approach toward feature space analysis," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 24, no. 5, pp. 603–619, 2002.