

Lifelong Learning for Mobile Robotics Applications

A Half-day Workshop for IROS 2012

Rudolph Triebel¹ and Luciano Spinello²

Abstract—Online and lifelong learning approaches are key functions for a robot to adapt knowledge and behavior over time. Crucial topics in mobile robotics such as mapping, planning, navigation, target tracking can be performed more reliably when formulated as online and lifelong learning problems. The goal of this workshop is to bring together top researchers from mobile robotics and other related disciplines to share knowledge and encourage discussion about the usage of these techniques. A mixture of invited talks and presentations of submitted contributions is scheduled to aim for fruitful discussions and to highlight many different aspects of the topic.

I. MOTIVATION AND OBJECTIVES

In recent years robots left controlled laboratories settings and start to address real world scenarios and situations. In order to tackle 'the wild', robots need to adapt their knowledge and their behaviour over time. Lifelong and online learning techniques are currently being established as one of the major research goals in mobile robotics. The rising interest in lifelong learning stems from the fact that all major mobile robot tasks, including mapping [5, 4, 8, 7], navigation [1], manipulation [2], target tracking [10], planning [3] or learning in general [9] can be performed more accurately and more reliably when formulated and addressed as online learning problems. Given that modern mobile robot systems are being employed in highly dynamic environments such as cities [6], the capability of improving and refining the internal representation and performance of algorithms for mobile robots during their operation becomes a crucial requirement.

The aim of this workshop is to specifically show how the formulation as lifelong learning problems can be done, to give solutions and show results to these problems, and to create a platform for interchanging ideas between the different areas of mobile robotics on the topic of online and lifelong learning. Some of the leading researchers from robotics, computer vision and machine learning will give invited talks, but there will also be enough time for presentations of submitted and peer-reviewed contributions.

II. LIST OF TOPICS

The topics that will be covered in the workshop include, but are not limited to:

¹R. Triebel is with the Department of Engineering Science, University of Oxford, OX1 3PJ Oxford, UK rudi@robots.ox.ac.uk

²L. Spinello is with the Department of Computer Science, Albert-Ludwigs-University Freiburg, 79110 Freiburg, Germany spinello@informatik.uni-freiburg.de

- Unsupervised/weakly supervised online learning techniques
- Bayesian state estimation
- Unsupervised/semi supervised object discovery
- Evaluation of online/offline methods
- Multiple target tracking with adaptive target models
- Online and adaptive outlier detection
- Lifelong mapping and localization
- Adaptive/online transfer learning and domain adaptation
- Online learning by imitation (grasping, manipulation)
- Planning in continuously changing environments
- Change detection in video/data streams

III. INVITED SPEAKERS

In the workshop, we will have a number of experts in the field as invited speakers. These are:

- *Oliver Brock*
Technische Universität Berlin, Germany
- *Wolfram Burgard*
University of Freiburg, Germany
- *Giorgio Grisetti*
La Sapienza University Rome, Italy
- *Kurt Konolige*
Industrial Perception, USA
- *Maxim Likhachev*
Carnegie Mellon University, USA
- *Francesco Orabona*
University of Chicago, USA

IV. INTENDED AUDIENCE

The interested audience include robotics researchers and scientists that are interested in the topics of online/lifelong methods applied to navigation, planning, computer vision, target tracking and machine learning.

V. SCHEDULE

Table I shows the workshop schedule. There are two sessions with invited speakers, one session with oral presentations on focused work on the workshop topic, and one poster session. Just before the start of the poster session, there will be spot light presentations for each poster.

VI. TALK DETAILS

In the following, we give the details of the oral presentations that will be given in the workshop. At the time when this document was printed, not all talk details were available, we recommend to consult the workshop webpage for updates: http://ais.informatik.uni-freiburg.de/lifelonglearning_ws12iros

TABLE I
TENTATIVE SCHEDULE

14:00 - 14:10	Introduction and Motivation	Rudolph Triebel and Luciano Spinello
Session 1: Theory and Application to Mapping and Planning		
14:10 - 14:35	Efficient and Principled Online Algorithms for Lifelong Learning	Francesco Orabona
14:35 - 15:00	Using RANSAC to Detect Dynamic Changes in Graph-Based SLAM	Giorgio Grisetti
15:00 - 15:25	TBD	Kurt Konolige
15:25 - 15:50	Coffee Break	
Session 2: Applications to Manipulation and Navigation		
15:50 - 16:15	Lifelong Search-based Planning: From Incremental Planning to Planning with Experience	Maxim Likhachev
16:15 - 16:40	About Lifelong Learning in Autonomous Manipulation	Oliver Brock
16:40 - 17:05	TBD	Wolfram Burgard
Session 3: Selected Contributions		
17:10 - 17:25	A Generic Robot Database and its Application in Fault Analysis and Performance Evaluation	Tim Niemüller
17:25 - 17:40	Learning Temporal Co-Observability Relationships for Lifelong Robotic Mapping	Nicholas Carlevaris-Bianco
Session 4: Spotlight and Poster Presentations		
17:40 - 17:45	Spotlight Presentations	G.-D. Tipaldi, R. Triebel
17:45 - 18:15	Poster Session	

A. Oliver Brock – About Lifelong Learning in Autonomous Manipulation

Manipulation has remained a tremendous challenge in robotics research: the real world offers overwhelming and unpredictable variability, uncertainty lurks in sensor signals and actions, and as soon as the robot makes contact with the world, the line between success and failure is determined by tiny differences. Any attempt of achieving manipulation will therefore have to rely on perception and feedback. Life-long learning could be employed quite naturally to continuously improve performance in this continuous domain. But in this presentation I would like to discuss the opportunities for life-long learning at the symbolic, relational level instead. How can life-long relational learning advance autonomous manipulation? What are the challenges? What is the relationship between the continuous and the relational domains? I will attempt to provide some answers to these questions in the context of an interactive perceptual learning system for the

extraction and actuation of kinematic degrees of freedom in the environment.

B. Giorgio Grisetti – Using RANSAC to Detect Dynamic Changes in Graph-Based SLAM

Having an up-to-date model of the environment is an essential precondition for a robot executing complex tasks in large and articulated environments. State of the art graph-based SLAM approaches can be effectively used to estimate a static model from a sequence of observations, however this data acquisition process, is a tedious and error prone that needs to be repeated each time something changes in the environment. Alternatively, the robot could constantly refine the estimate of the environment during its mission by means of a SLAM algorithm. However, this strategy suffers of a computational load that increases with the time of operation, and is subject to data association errors caused by changed portions of the environment. Being able to detect these changes, to discard outdated information is

one of the building blocks of a long-life SLAM system. Consensus algorithms, like RANSAC are successfully used address alignment problems in presence of a high degree of outliers. In this talk, we will discuss the advantages and the shortcomings of this family of techniques in detecting changed regions of the environment.

C. Maxim Likhachev – Lifelong Search-based Planning: From Incremental Planning to Planning with Experience

Planners running on physical robots often have to solve similar planning problems over and over again. For example, a path planner for robots operating in partially-known or dynamic environments has to repeatedly re-plans paths in response to changes in the map of the environment. Similarly, a motion planner for mobile manipulators performing everyday tasks such as picking up objects and opening doors often has to compute motions that are very similar to motions it has computed previously. It is therefore beneficial to study planning techniques that can exploit the redundancy in planning efforts whenever possible. In this talk, I will give overview of two very different approaches to doing this. One is based on incremental graph search techniques such as Lifelong Planning A*, D* Lite, Adaptive A* and Anytime D*. These methods are well suited for efficient re-planning. The second approach is based on a recently developed idea of bootstrapping graph search with Experience Graphs. The latter approach is well-suited for high-dimensional planning that uses experience to learn how to plan motions efficiently for repetitive tasks. I will illustrate pros and cons of these approaches based on their applications to autonomous navigation, autonomous flight and autonomous mobile manipulation.

D. Francesco Orabona – Efficient and Principled Online Algorithms for Lifelong Learning

For how long do you have to train a robot? The current mainstream approaches assume a supervised training phase, where the human labels the data, well separated from the actual unsupervised deployment of the robot in the world. Training usually stops when a pre-defined level of performance on a validation set of data is achieved, or when the user decides it. This approach is known to be doomed to fail: environment changes continuously over time and it is impossible to predict how. Continuous adaptation and learning is the only possibility to address real world scenarios. In other words, the transition from fully supervised to unsupervised should be smooth, robot driven, and reversible. I will present a formal theory, almost without any assumption, that allows to design and analyze online semi-supervised learning algorithms. When needed these algorithms can switch back from the unsupervised phase to the supervised one, triggered by the inability to explain the new incoming data. They also have formal guarantees on the performance and on the maximum amount of queried supervised data. Moreover, being based on the regularized least squares algorithm, they are easy and efficient to implement. Empirical results will be shown as well to validate the approach.

E. Nicholas Carlevaris-Bianco – Learning Temporal Co-Observability Relationships for Lifelong Robotic Mapping

This work reports on a method that learns the temporal co-observability relationships for exemplar views of a dynamic environment collected during long-term robotic mapping. These relationships are efficiently captured using a Chow-Liu tree approximation and allow one to predict which exemplars will be observed by the robot given the robots recent observations. For example, these learned relationships can encode scene dependent changes in lighting due to time of day and weather, without explicitly modeling them. Preliminary experimental results are shown using images from 17 fixed locations collected hourly over the course of 116 days.

Note: This presentation is accompanied with an article submitted to the workshop, which is attached at the end of this document.

F. Tim Niemüller – A Generic Robot Database and its Application in Fault Analysis and Performance Evaluation

During operation of robots large amounts of data are produced and processed for instance in perception, actuation, or decision making. Nowadays this data is typically volatile and disposed right after use. But this data can be valuable and useful later. Therefore we propose a database system that taps into common robot middleware to record any and all data produced at run-time. We present two examples using this data in fault analysis and performance evaluation and describe real-world experiments run on the domestic service robot HERB.

VII. POSTER PRESENTATIONS

In the workshop there will be the following poster presentations:

- Gian Diego Tipaldi: “Lifelong Localization and Dynamic Map Estimation in Changing Environments”
- Rudolph Triebel: “Parsing Outdoor Scenes from Streamed 3D Laser Data Using Online Clustering and Incremental Belief Updates”
- Nicholas Carlevaris-Bianco: “Learning Temporal Co-Observability Relationships for Lifelong Robotic Mapping”

REFERENCES

- [1] W. Churchill and P. Newman, “Practice makes perfect? managing and leveraging visual experiences for lifelong navigation,” in *Int. Conf. on Robotics and Automation (ICRA)*, 2012, to appear.
- [2] Y. P. Dov Katz and O. Brock, “Learning to manipulate articulated objects in unstructured environments using a grounded relational representation,” in *Proc. of the Robotics: Science and Systems (RSS)*, 2008.
- [3] S. Koenig, M. Likhachev, and D. Furcy, “Lifelong planning A*,” *Artificial Intelligence Journal (AIJ)*, vol. 155, no. 1-2, pp. 93–146, 2004.

- [4] K. Konolige and J. Bowman, "Towards lifelong visual maps," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2009.
- [5] H. Kretschmar, G. Grisetti, and C. Stachniss, "Life-long map learning for graph-based simultaneous localization and mapping," in *KI - Kuenstliche Intelligenz*, 2010.
- [6] J. Maye, R. Triebel, L. Spinello, and R. Siegwart, "Bayesian on-line learning of driving behaviors," in *Int. Conf. on Robotics and Automation (ICRA)*, 2011.
- [7] D. Meyer-Delius, J. Hess, G. Grisetti, and W. Burgard, "Temporary maps for robust localization in semi-static environments," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2012.
- [8] P. Newman, G. Sibley, M. Cummins, M. Smith, A. Harrison, C. Mei, and I. Reid, "Toward life long mapping," in *Proc. of the Int. Symposium of Robotics Research (ISRR)*, 2009.
- [9] F. Orabona and K. Crammer, "New adaptive algorithms for online classification," in *Advances in Neural Information Processing Systems (NIPS)*, 2010.
- [10] J. Santner, C. Leistner, A. Saffari, T. Pock, and H. Bischof, "Prost: Parallel robust online simple tracking," in *Proc. of the IEEE Conf. on Comp. Vis. and Pat. Rec. (CVPR)*, 2010.

Learning Temporal Co-Observability Relationships for Lifelong Robotic Mapping

Nicholas Carlevaris-Bianco and Ryan M. Eustice

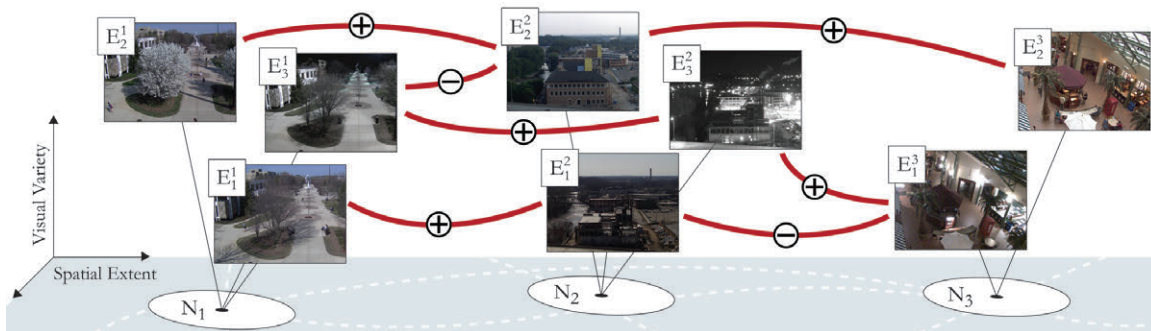


Fig. 1: Illustration of an “exemplar” based map representation. Locations are modeled as “metric” neighborhoods, N_i for $i = [1 \dots n]$. The visual appearance of each neighborhood is represented by a set of exemplar views, E_j^i for $j = [1 \dots m_i]$, which seek to represent the possible variation in appearance of a location due to temporal changes in the environment. **Our proposed method learns the temporal observability relationships between exemplars allowing one to predict which exemplars will be observed based on recent measurements.** The temporal observability relationships are illustrated as red links with the sign indicating the sign of the correlation.

Abstract—This paper reports on a method that learns the temporal co-observability relationships for exemplar views of a dynamic environment collected during long-term robotic mapping. These relationships are efficiently captured using a Chow-Liu tree approximation and allow one to predict which exemplars will be observed by the robot given the robot’s recent observations. For example, these learned relationships can encode scene dependent changes in lighting due to time of day and weather, without explicitly modeling them. Preliminary experimental results are shown using images from 17 fixed locations collected hourly over the course of 116 days.

I. INTRODUCTION

Recent work has sought to address the challenge of robotic simultaneous localization and mapping (SLAM) in dynamic environments by using a map representation that encodes both spatial and temporal information. Promising examples of these methods elect to represent locations in the map as a collection of views [1], environmental states [2], or view sequences [3], corresponding to how the environment appeared at various observation instances. These “exemplar” based methods are capable of representing many different types of temporal variation. It is our opinion that these methods provide a

promising way forward toward developing SLAM systems for long-term autonomy in dynamic environments.

Current methods do not attempt to understand the temporal relationship between map exemplars. Data association is performed by comparing the current sensor measurement against all possible exemplars to determine which best explains the current measurement. Essentially, the robot considers all possible temporal variations of a location every time it attempts to measure the environment. This results in a map that grows both with spatial extent and with the amount of dynamic variation in the environment.

In this paper we present a method that learns the temporal observability relationships between exemplars in the map. These relationships can be used to predict which exemplars will be observed given what the robot has recently observed. Conceptually this can be thought of as producing a maximum likelihood (ML) temporal map.

The learned relationships encode many different types of temporal changes, without explicitly modeling them. For example, the proposed method learns groupings of exemplars caused by changes in lighting between morning, noon, and night, even though the time of day is not a modeled input to the system. The structure of the correlation learned by the proposed method is determined by the properties of the sensor system. In the case of vision the correlation structure is primarily defined by changes in lighting. This would not be true in the case of light detection and ranging (LIDAR), which is invariant to changes in illumination and would be more affected by changes in geometric structure and occlusion.

Our motivation for this work is a centralized multi-agent system in which a common map is shared between many

This work was supported in part by the National Science Foundation under award IIS-0746455, and in part by the Naval Sea Systems Command (NAVSEA) through the Naval Engineering Education Center under award N65540-10-C-0003.

N. Carlevaris-Bianco is with the Department of Electrical Engineering & Computer Science, University of Michigan, Ann Arbor, MI 48109, USA carlevar@umich.edu

R. Eustice is with the Department of Naval Architecture & Marine Engineering, University of Michigan, Ann Arbor, MI 48109, USA eustice@umich.edu

agents. We advocate in Section V that the proposed algorithm is particularly well suited for this scenario. The remainder of the paper is outlined as follows. In Section II we review related work. The proposed method is described in Section III. Finally, our experimental results are presented and discussed in Sections IV and V, respectively. Section VI concludes.

II. RELATED WORK

The majority of SLAM solutions, to date, make the assumption that the environment is static. For many robotic applications this assumption is not extremely detrimental. This is especially true for systems that exploit robust data association methods, notably image and laser range registration, as these methods are capable of successful measurements in the presence of short-term dynamic effects, such as partial scene occlusion and moving objects.

Several methods have been proposed that try to filter out dynamic elements of the environment while maintaining the assumption that the underlying environment is static, for example [4]–[7], among others. It has also been shown that it is possible to explicitly identify and track specific classes of dynamic objects in the environment such as people [8].

Recently, several works have proposed methods that explicitly model dynamic changes in the map. In [9], [10], Biber and Duckett represent the environment by a collection of sample-based maps, each of which incorporates new samples and forgets old samples at a different rate. This rate determines the timescale of each map. During localization the robot tests all timescales to determine which best agrees with the current observation and then performs measurement updates against that map. In some respects this representation is similar to the exemplar-based methods that this paper expands upon, as each map could be thought of as a type of exemplar. They are not exactly equivalent though, due to the fact that in this case each timescale-sampled map is a collection of data measured at multiple instances in time, and not a collection of a single-time-instance example views of the environment. The idea of sample-based maps is continued by Dayoub and Duckett in [11], which uses a short-term versus long-term memory model to update the collection of visual features that represent the appearance of a location. In [12] a method is presented which uses a “dynamic occupancy grid” based on a hidden Markov model to capture dynamic changes in the environment.

The works most relevant to our proposed method can be considered as “exemplar-based” methods [1]–[3]. In [2] Stachniss and Burgard propose a method to learn exemplar configurations of an indoor environment from 2D laser scan data using fuzzy k-means clustering. They then use these exemplar configurations in a particle-filter based localization framework. Konolige and Bowman present a vision-based method in [1]. This method works within the context of vision based pose-graph SLAM [13], where the pose-graph is divided into metric neighborhoods of views bounded by physical location and view attitude. Within these neighborhoods each view is an example of how the environment looked at the time it was collected. They then present a least-recently-used view

deletion algorithm, which limits the number of exemplars per neighborhood to a fixed number and encourages a long-term equilibrium with the minimum set of exemplars that explains the visual variation of that neighborhood. In [3], Churchill and Newman use sequences of views, termed “experiences”, as the basic unit of the temporal map instead of individual views. New experiences are added to the map when the existing experiences are unable to explain the current observations.

The map representation used in this paper, including the idea of exemplars and their use in metric neighborhoods, was first introduced in [1]. Our contribution beyond existing exemplar-based methods, [1]–[3], is an algorithm to learn the temporal observability relationships between exemplars throughout the map. This allows one to *predict which exemplar within a yet unseen neighborhood is most likely to be observed* given the most recent exemplar observations from other neighborhoods, which avoids having to compare the robot’s current view against all exemplars at each neighborhood.

III. LEARNING TEMPORAL OBSERVABILITY RELATIONSHIPS BETWEEN EXEMPLARS

The proposed algorithm uses a map representation that models each location as a “spatial” or “metric” neighborhood, N_i for $i = [1 \dots n]$. Each neighborhood represents a physical space in the environment. The nature of this space is defined by the choice of sensing modality so that measurements within the neighborhood should overlap sufficiently for registration. For example, in the case of omnidirectional vision or 2D LIDAR a neighborhood may be defined as a max translation between any two views. For monocular vision, one would also need to include the vantage point of the camera. The appearance, geometric shape, or state of each neighborhood is represented by a set of “exemplar” views, E_j^i for $j = [1 \dots m_i]$, which seek to represent the possible variation in appearance or structure of a location due to temporal changes in the environment, as illustrated in Fig. 1.

In order to learn the temporal observability relationships between map exemplars, and to predict the observability of each exemplar in the map, we wish to estimate the joint distribution of the probability that each exemplar will be observed during a given time window of length λ . Let Z_k be a binary random variable representing if the k^{th} exemplar was observed in a given time window. Therefore, we wish to estimate,

$$P(\mathbf{Z}) = P(Z_1, Z_2, \dots, Z_M), \quad (1)$$

where $M = \sum_{i=1}^n m_i$ is the total number of exemplars in the map.

As the robot builds the map over time, it records each of its attempts to localize against the exemplars in the map. This produces a vector of observations \mathbf{z} . We subdivide this set of all previous observations into two groups

$$\mathbf{z} = \mathbf{z}_p \cup \mathbf{z}_c \quad (2)$$

$$\mathbf{z}_p = [z_i : t_i \neq t] \quad (3)$$

$$\mathbf{z}_c = [z_i : t_i = t] \quad (4)$$




Fig. 3: Illustration of the Chow-Liu tree approximation. The magnitude of mutual information between variables is indicated by line thickness. The original distribution $P(x_1, x_2, x_3, x_4)$ is approximated as $P(x_1)P(x_3|x_1)P(x_2|x_3)P(x_4|x_3)$.

where \mathbf{z}_p includes all previous measurements observed in a time window, t_i , other than the current, t . Conversely, \mathbf{z}_c represents all current observations collected in the current time window. Our aim then is to learn the joint distribution, (1), from \mathbf{z}_p , and then, as the robot moves through the environment, infer the probability of observing each exemplar in the current time window given the current observations, \mathbf{z}_c .

The time window parameter, λ , can be thought of as a lower bound on the timescale of the map. Two exemplars, E_a and E_b , are considered co-observed if they were observed within the same time window t_i of length λ . Therefore, our method learns the temporal correlation only for events with a timescale greater than λ . In our experiments λ is 1 hour, low enough to capture changes in lighting throughout the day, but long enough to exclude short-term dynamics such as people and moving vehicles.

Unfortunately, as the size of the map increases so will the number of exemplar views, and estimating the full distribution with all possible conditional dependencies quickly becomes intractable. Simply treating the probability of observing each exemplar as independent would solve the tractability problem. However, this would prevent any understanding of the temporal correlation in observability between exemplars. Without this correlation the most recent observations, \mathbf{z}_c , will not provide information about future observations in the same time window.

As a compromise between tractability and the ability to capture the temporal correlation between exemplar views, we have elected to approximate the joint distribution with a Chow-Liu tree [14]. The Chow-Liu tree approximates the full distribution as the product of pairwise conditional distributions. The pairwise conditional distributions are selected such that the Kullback-Leibler divergence between the original distribution and the Chow-Liu tree approximation is minimized. Therefore, we can approximate (1) as

$$P(\mathbf{Z}) \approx P(Z_r) \prod_{i=2}^M P(Z_i|Z_{p(i)}) \quad (5)$$

where Z_r is the root variable of the Chow-Liu tree and $Z_{p(i)}$ is the parent of Z_i .

The Chow-Liu tree approximation is constructed by finding the mutual information between each pair of variables. This produces a fully connected graph where the weight of each edge is determined by the mutual information between the

connected variables. From this graph the Chow-Liu tree is computed by finding the maximum spanning tree (Fig. 3). The mutual information between two variables is defined as

$$I(Z_i, Z_j) = \sum_{z_i \in 0,1} \sum_{z_j \in 0,1} P(z_i, z_j) \log \frac{P(z_i, z_j)}{P(z_i)P(z_j)}. \quad (6)$$

Therefore, the Chow-Liu approximation can be learned directly from the previous observation data, \mathbf{z}_p , by estimating each marginal distribution and each pairwise joint distribution.

Finding the maximum *a posteriori* (MAP) estimate of these quantities amounts to frequency counting based on the number of times in which each exemplar was observed, and the number of times in which each pair was successfully co-observed in the same time window. For the marginal distributions

$$\hat{P}(Z_i = z) = \frac{N_{[Z_i=z]} + (\alpha - 1)}{N_{[Z_i]} + (\alpha + \beta - 2)}, \quad (7)$$

where $N_{[Z_i=z]}$ is the number of observations where $Z_i = z$, $N_{[Z_i]}$ is the number of times Z_i was observed either positively or negatively, and α and β are parameters of a beta prior that encourages the estimate away from the extremes, 0.0 and 1.0, and toward an empirically determined target mode. In our experiments the beta prior's mode was set to 0.2, with the constraint that $\alpha + \beta = 4$. The joint distributions were estimated as

$$\hat{P}(Z_i = z_i, Z_j = z_j) = \frac{N_{[Z_i=z_i \& Z_j=z_j]} + (\alpha - 1)}{N_{[Z_i \& Z_j]} + (\alpha + \beta - 2)}, \quad (8)$$

where $N_{[Z_i=z_i \& Z_j=z_j]}$ is the number of times $Z_i = z_i$ and $Z_j = z_j$ in the same time window, and $N_{[Z_i \& Z_j]}$ is the number of times both exemplars were co-observed in the same time window, regardless of value. In this case the beta prior parameters are automatically chosen so that the mode is equal to $\hat{P}(Z_i = z_i)\hat{P}(Z_j = z_j)$, with the constraint that $\alpha + \beta = 4$. This essentially assumes independence until there is sufficient data to indicate otherwise. Similarly, the conditional distributions that form the Chow-Liu tree were estimated as

$$\hat{P}(Z_i = z_i|Z_j = z_j) = \frac{N_{[Z_i=z_i \& Z_j=z_j]} + (\alpha - 1)}{N_{[Z_i \& Z_j=z_j]} + (\alpha + \beta - 2)}. \quad (9)$$

Again the beta prior parameters are automatically chosen to assume independence, with the mode of the prior equal to $\hat{P}(Z_i = z_i)$ with the constraint that $\alpha + \beta = 4$.

It is important to note that the robot may not encounter all neighborhoods during each time window. Therefore, for the exemplars within an unvisited neighborhood we have no evidence of their observability. It is important not to count this absence of observation as a negative observation. Finally, the robot may never have attempted to co-observe some pairs in the same time window. In this case we again assume their independence with $\hat{P}(z_i, z_j) = \hat{P}(z_i)\hat{P}(z_j)$.

When building the Chow-Liu tree we enforce an additional constraint that no edges should be included between exemplars within the same neighborhood. This forces the tree to only include dependencies between neighborhoods. Our focus is to make predictions between neighborhoods, not within them.




Fig. 4: Depiction of webcam locations at GVSU.

Having learned an approximation of the underlying observability distribution using observations from all previous time windows, \mathbf{z}_p , we now wish to predict the probability of observing each exemplar given the previous observations made during the current time window, \mathbf{z}_c . We can use the sum-product algorithm [15], [16] to efficiently perform inference because the joint distribution (1) has been approximated using a tree structure (5).

IV. EXPERIMENTAL RESULTS

A. Experimental Setup

In order to evaluate the proposed algorithm we performed experiments using data collected at 17 fixed locations over the course of 116 days (Fig. 4). Data was collected approximately hourly from March 9th, 2012 through July 3rd, 2012 by automatically downloading images from web-accessible security cameras at Grand Valley State University’s campus [17]. The data includes both indoor and outdoor scenes under dynamic changes including lighting, weather, people, cars, moving furniture, seasonal changes and construction. Additionally, for each set of images the time of day and weather conditions [18] were automatically recorded for later comparison. This time and weather meta data was not in any way used by the algorithm, and was only used after-the-fact as ground-truth to examine what types of relationships were being learned by the algorithm.

Considering each of the 17 locations as a metric neighborhood, approximately 1000 simulated robot trajectories were generated by visiting each neighborhood in a random order. The time between each synthetic trajectory was also randomized uniformly between 1 and 5 hours. In this experiment we assume that the neighborhood for each view is known. For the experiment the map was built as follows. The first trajectory was used to initialize a set of metric neighborhoods, each with a single visual exemplar view. Each following trajectory was then processed according to Algorithm 1.

Prior to processing each trajectory, we calculate the Chow-Liu approximation [14] of the joint distribution of exemplar

Algorithm 1 Experiment Map Update Algorithm

- 1: Given current map, observation history \mathbf{z}_p , and new trajectory T consisting of a set of views V_i for $1 = [1 \dots v]$
 - 2: // Localization Loop
 - 3: Calculate Chow-Liu approximation of $P(\mathbf{Z})$ using \mathbf{z}_p
 - 4: **for all** Views, V_i , in T **do**
 - 5: Evaluate $P(\mathbf{Z}|\mathbf{z}_c)$ using Sum-Product
 - 6: Predict matching results using $P(\mathbf{Z}|\mathbf{z}_c)$
 - 7: Visual matching between V_i and exemplars in N_i
 - 8: Update \mathbf{z}_c with matching results
 - 9: **end for**
 - 10: // Map Update Loop
 - 11: **for all** Views, V_i , in T **do**
 - 12: **if** no match between V_i and exemplars in N_i **then**
 - 13: **if** $m_i == m_{\max}$ **then**
 - 14: Forget least-recently-used exemplar, $m_i = m_i - 1$
 - 15: **end if**
 - 16: Add V_i as new exemplar in N_i , $m_i = m_i + 1$
 - 17: **end if**
 - 18: **end for**
-

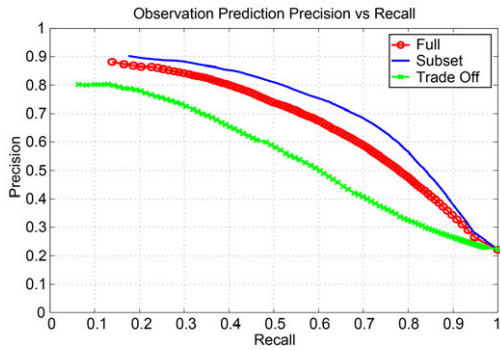
observations, $P(\mathbf{Z})$. We then consider each view in the trajectory. At each neighborhood, prior to performing matching, the probability of observing each of the neighborhood’s exemplars is estimated given the observations made so far in this trajectory, $P(\mathbf{Z}|\mathbf{z}_c)$. We then perform visual matching between the current view and each of the neighborhood’s exemplars in order to compare the results with the prediction and to add more observations to \mathbf{z}_c .

After processing each trajectory, the neighborhoods in the map are updated using a least-recently-used update rule, similar in concept but simplified with respect to that proposed in [1]. In each neighborhood, if the new view does not match against any of the existing exemplars, then the new view is added as a new exemplar. If the number of exemplars in a neighborhood exceeds a threshold (in our experiments $m_{\max} = 9$), the exemplar that has been least recently matched against is removed.

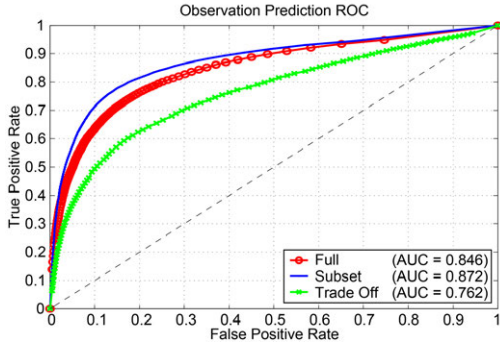
For this experiment a match occurs if the number of inlier features is greater than 20, and the mean squared reprojection error is less than 10 pixels. Note that it is possible for the current view to match more than one exemplar in a neighborhood. The exact definition of a match depends strongly on the sensing modality and SLAM front-end implementation used. Matching should be defined by the same criteria used in determining which constraints are added to the SLAM solution and in a full implementation may include additional criteria such as consistency with SLAM-derived motion priors.

B. Prediction Accuracy

Precision-recall and receiver operating characteristic (ROC) curves for the predicted observability in this experiment are shown in Fig. 5. These curves are generated by varying the probability threshold used to predict if an exemplar will be observable. Ground-truth is known as matching was performed



(a) Precision-Recall Curve



(b) ROC Curve

Fig. 5: Precision-recall curve and ROC curve for observability prediction. Results for all predictions made throughout the experiments are shown in red and labeled “Full.” Results that exclude predictions made on newly added exemplars (less than 2 days old), and predictions made with a small number of current observations (less than 4 neighborhoods observed), are shown in blue and labeled “Subset.” Results that consider all predictions, but only use training data from the top two most likely and two randomly chosen “test” exemplars, are shown in green and labeled “Trade Off.” AUC indicates the area under the curve.

after the prediction to verify the results. It is interesting to note that there are two situations that occur during the experiment where one would not expect the algorithm to perform well. The first occurs when a new exemplar is added to the map. The estimated probabilities for new exemplars with few observations are based largely on beta priors that assume independence from other exemplars. Therefore, we do not expect the algorithm to accurately predict the observability of exemplars until sufficient observations have been made to estimate the inter-exemplar correlation. The second occurs at the start of each trajectory where, before any current observations have been made, the predictions are based on the learned probabilities of observation for each exemplar.

For comparison in Fig. 5, curves are shown for all predictions made throughout the experiment, shown in red and labeled “Full,” and for a subset that excludes predictions made under the aforementioned conditions where less observation data is available, shown in blue and labeled “Subset.” As one would expect, excluding the predictions made with few observations improves performance. More interestingly, this

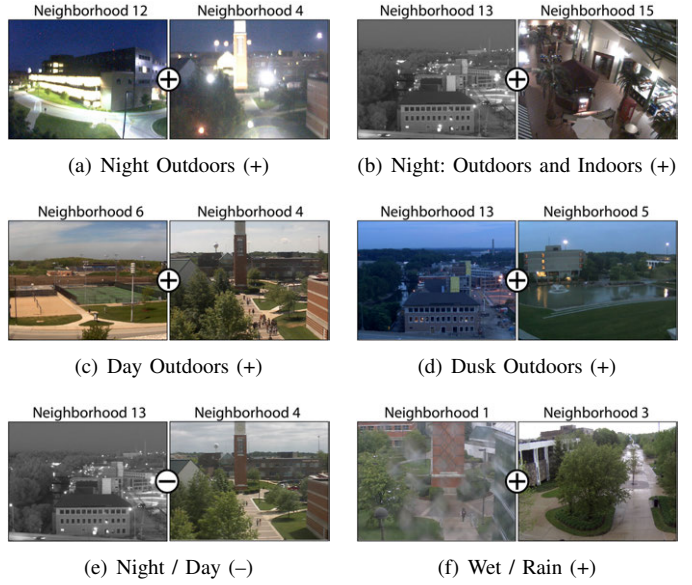


Fig. 6: Sample exemplar pairs with high positive or negative correlation that were included in the Chow-Liu tree approximation. Sign between pairs indicates sign of correlation. These pairs are representative of the types of relationships automatically found in the Chow-Liu approximation.

illustrates the interaction between the proposed observation prediction algorithm and the exemplar updating scheme. Exemplar updating schemes that quickly settle to a stable set of exemplars per neighborhood are ideal, while update rules that continuously add new and forget old exemplars will be detrimental.

It is important to note that there is an inherent trade-off between reducing the data association effort and producing more data to learn the Chow-Liu approximation. In the “Full” and “Subset” experiments, visual matching was performed between the current view and each exemplar in the neighborhood regardless of the predicted probability of observation. All of these observations were then provided to subsequent predictions. Therefore, there was no gain in terms of the effort required for data association. To investigate this trade-off, an experiment was performed in which the match data from only the top two most likely and two randomly chosen “test” exemplars were used to train the Chow-Liu tree and to perform inference. These results are shown in green and labeled “Trade Off” in Fig. 5. As only a maximum of four exemplars are tested per neighborhood, this represents a data association effort reduction of greater than 50% at steady-state (when each neighborhood has reached its maximum of nine exemplars).

C. Learned Temporal Relationships

In order to visualize the type of co-observability relationships being learned by the Chow-Liu approximation several representative examples of highly correlated exemplar pairs (from different neighborhoods) that were learned by the Chow-Liu tree approximation are depicted in Fig. 6. It appears



Fig. 7: Sample imagery for ML exemplars for 7 not-yet-observed neighborhoods (1,2,3,7,8,9,13) as other neighborhoods in the map are observed (5,12,14,16,17). The upper right shows the first 5 observations (in order from left to right) made during this trajectory. In the grid below, the first row shows the true realization of each of the neighborhoods, which have yet to be observed. The remaining rows show the maximum likelihood exemplars after the first 0, 1, 3, and 5 neighborhoods are observed (0 being the prior). Red boxes highlight exemplars that have changed given the new observations for that row.

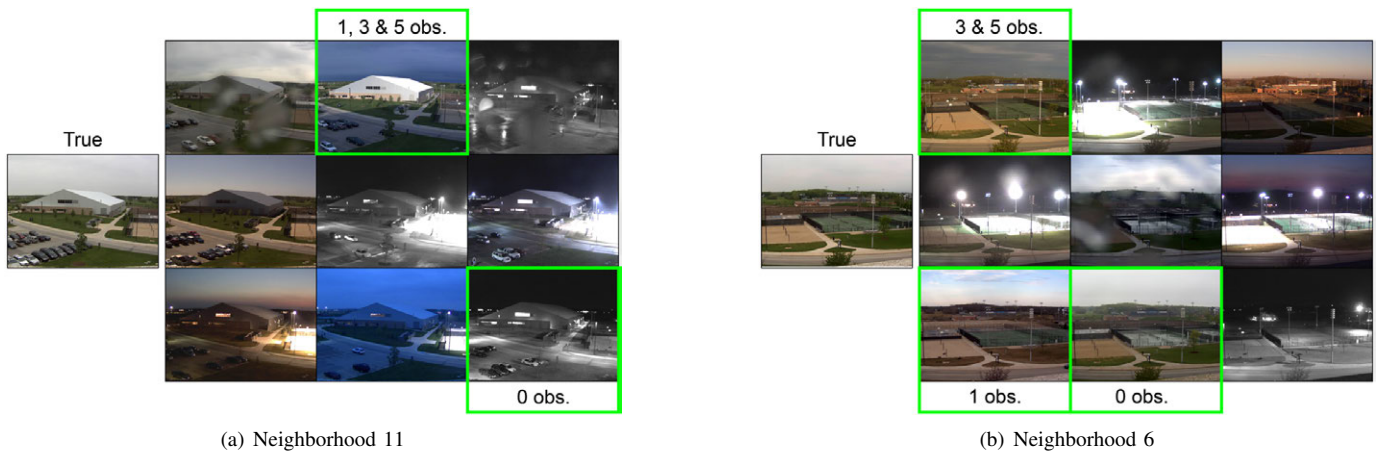


Fig. 8: Sample imagery showing the nine exemplars from two not-yet-observed neighborhoods. To the left of the exemplars the true realization of each of the neighborhoods is shown. The ML exemplars after 0, 1, 2, and 5 other neighborhoods are observed (0 being the prior) are highlighted with green boxes.

that the structure of Chow-Liu tree is driven primarily by lighting variation, with weather and sky conditions playing a much smaller role. Many of the strongest positive correlations were found between night exemplars, most likely due to the consistent artificial lighting in these scenes and their complete visual dissimilarity with daytime exemplars. Almost all cases of negative correlation were between day and night exemplars.

The imagery in Fig. 7 and in Fig. 8 shows examples of how the ML exemplars for several sample neighborhoods change as new observations are made. With no observations the ML exemplars are based on the learned prior probability of observation and in this example contain a mixture of night and day exemplars. After observing the first neighborhood, the night exemplars are replaced with daytime exemplars. As more neighborhoods are observed, the ML exemplars change more subtly toward less sunny exemplars, which better match the true realization.

For illustration purposes, Fig. 7 shows the ML exemplars for multiple neighborhoods at the same time. However, during the actual experiment the ML exemplar for each neighborhood is only calculated immediately before attempting to match exemplars in that neighborhood, and the likelihood is maximized only locally over the exemplars in that neighborhood. To truly estimate a maximum likelihood map over multiple neighborhoods, it would be more appropriate to use the max-product algorithm [16] instead of the sum-product to find the set of exemplars that *jointly* maximize the likelihood.

The temporal changes learned in the Chow-Liu approximation are further illustrated in Fig. 9. By considering all observations that meet a specific time-of-day condition, we can calculate the ratios of the time-of-day conditions of the successfully matched exemplars for these observations. As one would expect, we see in Fig. 9(a) that visual matching is dependent on lighting conditions, which change throughout the day. Observations are more likely to match against exemplars collected under similar time-of-day conditions. At no time does the proposed algorithm use any external information about the time of day, yet it effects visual matching; we would like the algorithm to propose appropriate exemplars, following the relationships in Fig. 9(a). In Fig. 9(b) we see that this is the case. The ratios of conditions of the predicted ML exemplars closely match those in Fig. 9(a) indicating that the algorithm has learned the underlying co-observability relationships caused by changes in lighting. For observations made during the morning, midday, afternoon and night, our algorithm learns to propose exemplars from the appropriate time of day.

Similarly, in Fig. 9(c) we see that the current sky condition also effects visual matching, though to a lesser extent than the time-of-day. Fig. 9(d) indicates our algorithm learns to predict exemplars according to the underlying co-observability relationships caused by changes between clear and overcast skies.

Again it is important to note, that the algorithm does not use any external knowledge of lighting or weather conditions, and that correlation between the observability of exemplars could

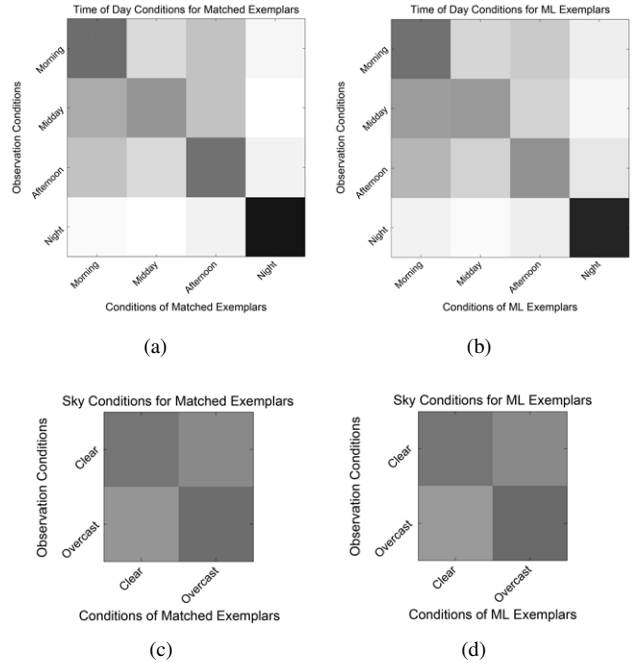


Fig. 9: Comparison between the conditions of the matched exemplars and the ML exemplars for a given observation condition. In (a) and (b) each row corresponds to the set of observations that occurred at a specific time of day. The shade of each cell represents the ratio of exemplars from that time of day, where white is 0.0 and black is 1.0. Similarly, in (c) and (d) exemplars collected under clear and overcast sky conditions are compared.

be developed by many different types of dynamic changes that affect visual matching.

V. DISCUSSION AND FUTURE WORK

There are several important aspects to consider with respect to the applicability and utility of the proposed algorithm:

- There is an inherent trade-off between the data association effort and producing more data to learn the Chow-Liu approximation. The more attempts made to match against exemplars, the more accurate the predictions become. However, the fewer the matching attempts, the larger the reduction in data association effort.
- Once built, performing inference using the Chow-Liu approximation is very efficient due to its tree structure. Unfortunately, building the Chow-Liu approximation itself is computationally expensive (i.e., $\mathcal{O}(M^2 \log M)$ where M is the number of exemplars).
- One might argue that using data association methods based on fast bag-of-words (BoW) place recognition techniques [19]–[21], as demonstrated in [1], reduces the cost of producing data association hypotheses sufficiently such that it is not computationally necessary to predict which exemplars will be observed. Instead, one can simply consider all exemplars at every data association attempt.

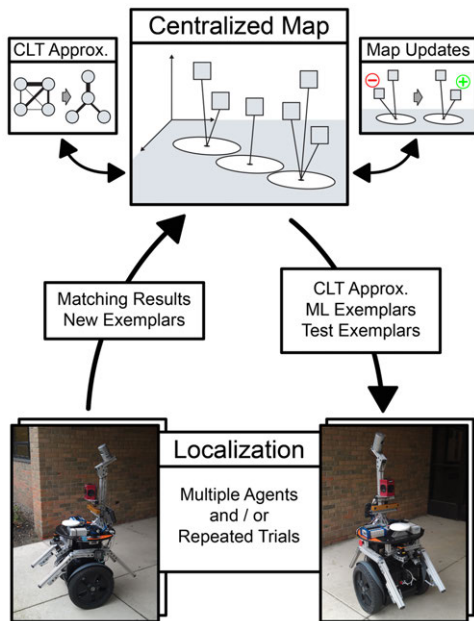


Fig. 10: Illustration of a centralized multi-agent system. In this scenario a common map is shared between many agents. The centralized map provides the current Chow-Liu approximation, ML exemplars, and possibly additional test exemplars to the localizing agents. The localizing agents return match results and possible new exemplars. Updating the Chow-Liu approximation and performing map updates are carried out as offline batch processes.

Based on these concerns we understand that the present algorithm may not seem practical for online, single-agent, map building. It is our belief, though, that if we consider a centralized multi-agent scenario, where a common map is shared between many agents (Fig. 10), then the proposed algorithm could provide substantial benefits:

- By sharing a map between multiple agents the burden of testing exemplar observability can be spread out among all agents while improving the reduction in data association effort.
- Building the Chow-Liu approximation can be done as a batch offline process and the result transferred to all agents where it can be used for efficient inference.
- Being able to predict the ML exemplars allows only a small subset of exemplars to be transferred between the centralized map and the individual agents.

Much work remains to fully demonstrate the effectiveness of the proposed algorithm. Most importantly, we intend to apply it to a large-scale real-world robotic data set. Additionally, exploring methods to balance the trade-off between data association effort and training data production, such as actively selecting test exemplars, and exploring the interaction between exemplar update rules and the ability to predict exemplar observability, could also improve the algorithm further.

VI. CONCLUSIONS

We presented a method that learns the temporal observability relationships between views of a dynamic environment.

We have demonstrated that these relationships can be used to predict which views of the environment will be observed by the robot given what the robot has recently observed. We have demonstrated that the learned relationships can encode many different types of temporal changes, without explicitly modeling them, including changes in lighting due to time of day and weather.

REFERENCES

- [1] K. Konolige and J. Bowman, "Towards lifelong visual maps," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, 2009, pp. 1156–1163.
- [2] C. Stachniss and W. Burgard, "Mobile robot mapping and localization in non-static environments," in *Proc. AAAI Nat. Conf. Artif. Intell.*, 2005, pp. 1324–1329.
- [3] W. Churchill and P. Newman, "Practice makes perfect? managing and leveraging visual experiences for lifelong navigation," in *Proc. IEEE Int. Conf. Robot. and Automation*, 2012, pp. 4525–4532.
- [4] D. Fox, W. Burgard, and S. Thrun, "Markov localization for mobile robots in dynamic environments," *J. Artif. Intell. Res.*, vol. 11, pp. 391–427, Nov. 1999.
- [5] D. Hähnel, R. Triebel, W. Burgard, and S. Thrun, "Map building with mobile robots in dynamic environments," in *Proc. IEEE Int. Conf. Robot. and Automation*, 2003, pp. 1557–1563.
- [6] C.-C. Wang, C. Thorpe, and A. Suppe, "LADAR-based detection and tracking of moving objects from a ground vehicle at high speeds," in *IEEE Intelligent Vehicles Symp.*, June 2003, pp. 416–421.
- [7] W. Burgard, A. B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun, "Experiences with an interactive museum tour-guide robot," *Artificial Intelligence*, vol. 114, pp. 3–55, 2003.
- [8] M. Montemerlo, S. Thrun, and W. Whittake, "Conditional particle filters for simultaneous mobile robot localization and people-tracking," in *Proc. IEEE Int. Conf. Robot. and Automation*, 2002, pp. 695–701.
- [9] P. Biber and T. Duckett, "Dynamic maps for long-term operation of mobile service robots," in *Proc. Robot.: Sci. & Syst. Conf.*, 2005, pp. 17–24.
- [10] P. Biber and T. Duckett, "Experimental analysis of sample-based maps for long-term SLAM," *Int. J. Robot. Res.*, vol. 28, no. 1, pp. 20–33, 2009.
- [11] F. Dayoub and T. Duckett, "An adaptive appearance-based map for long-term topological localization of mobile robots," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, 2008, pp. 3364–3369.
- [12] G. D. Tipaldi, D. Meyer-Delius, M. Beinhofer, and W. Burgard, "Life-long localization and dynamic map estimation in changing environments," in *RSS Workshop on Robots in Clutter*, 2012.
- [13] K. Konolige, J. Bowman, J. Chen, P. Mihelich, M. Calonder, V. Lepetit, and P. Fua, "View-based maps," *Int. J. Robot. Res.*, vol. 29, no. 8, pp. 941–957, 2010.
- [14] C. Chow and C. N. Liu, "Approximating discrete probability distributions with dependence trees," *IEEE Trans. on Info. Theory*, vol. 14, pp. 462–467, 1968.
- [15] J. Pearl, "Reverend Bayes on inference engines: A distributed hierarchical approach," in *Proc. AAAI Nat. Conf. Artif. Intell.*, 1982, pp. 133–136.
- [16] F. Kschischang, B. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. on Info. Theory*, vol. 47, no. 2, pp. 498–519, Feb. 2001.
- [17] Grand Valley State University Webcams. Grand Valley State University. Allendale, MI, USA. Accessed March 9, 2012 through July 3, 2012. [Online]. Available: <http://www.gvsu.edu/webcams.htm>
- [18] Weather Underground API. Weather Underground. Ann Arbor, MI, USA. Accessed March 9, 2012 through July 3, 2012. [Online]. Available: <http://www.wunderground.com/weather/api/>
- [19] J. Sivic and A. Zisserman, "Video Google: A text retrieval approach to object matching in videos," in *Proc. IEEE Int. Conf. Comput. Vis.*, vol. 2, Oct. 2003, pp. 1470–1477.
- [20] D. Nistér and H. Stewénius, "Scalable recognition with a vocabulary tree," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, vol. 2, 2006, pp. 2161–2168.
- [21] M. Cummins and P. Newman, "FAB-MAP: Probabilistic localization and mapping in the space of appearance," *Int. J. Robot. Res.*, vol. 27, no. 6, pp. 647–665, Jun. 2008.