

CIGAR: Concurrent and Interleaving Goal and Activity Recognition

Derek Hao Hu and Qiang Yang

Department of Computer Science and Engineering
Hong Kong University of Science and Technology
Clear Water Bay, Kowloon, Hong Kong
{derekhh, qyang}@cse.ust.hk

Abstract

In artificial intelligence and pervasive computing research, inferring users' high-level goals from activity sequences is an important task. A major challenge in goal recognition is that users often pursue several high-level goals in a concurrent and interleaving manner, where the pursuit of goals may spread over different parts of an activity sequence and may be pursued in parallel. Existing approaches to recognizing multiple goals often formulate this problem either as a single-goal recognition problem or in a deterministic way, ignoring uncertainty. In this paper, we propose *CIGAR* (Concurrent and Interleaving Goal and Activity Recognition) - a novel and simple two-level probabilistic framework for multiple-goal recognition where we can recognize both concurrent and interleaving goals. We use skip-chain conditional random fields (SCCRF) for modeling interleaving goals and we model concurrent goals by adjusting inferred probabilities through a correlation graph, which is a major advantage in that we are able to reason about goal interactions explicitly through the correlation graph. The two-level framework also avoids the high training complexity when modeling concurrency and interleaving together in a unified CRF model. Experimental results show that our method can effectively improve recognition accuracies on several real-world datasets collected from various wireless and sensor networks.

Introduction

In recent years, goal recognition, or activity recognition¹, has been drawing increasing interest in AI and pervasive computing communities. This research is particularly active given the fast advances in wireless and sensor networks, due to the prospect of directly recognizing users' goals and activities from sensor readings. Typical applications of goal recognition range from services for helping the elderly people to identifying significant activities and places from GPS traces. (Pollack *et al.* 2003; Liao, Fox, & Kautz 2007).

Beyond goal recognition, other similar fields may include activity recognition, behavior recognition or intent recognition. As pointed out in (Liao 2006), although these terms may emphasize different aspects of human activities, their essential goals are the same. Therefore, in this paper, we use

the term goal recognition and do not distinguish the minor differences among the different terms mentioned above.

Historically, goal recognition was done through logic and consistency based approaches (Kautz 1987). In the past few years, probabilistic approaches have been developed that are capable of handling uncertainty. Many of these approaches (Patterson *et al.* 2005; Vail, Veloso, & Lafferty 2007) assume that users achieve one goal at a time, and that goals are achieved through a consecutive sequence of actions (See Figure 1, top). However, in many real-world situations, users may accomplish multiple goals within a single sequence of actions where goals are achieved concurrently and the actions that achieve them are interleaving. We call this problem the *multiple-goal recognition* problem. Previous approaches will have problems in this situation.

Two real-world examples help explain the necessities and difficulties of modeling concurrent and interleaving goals in the multiple-goal recognition problem, respectively. Consider a professor who is leaving his office to achieve the goal of "printing some research papers", he then goes to the seminar room for achieving the goal of "presentation". If the printing room is on his way to the seminar room, then the professor can be considered as pursuing two goals through an observed activity sequence, i.e. the goals of printing and presentation, *concurrently*. In another example, an individual gets up early in the morning and boils water on the kettle. The kettle boils while he is having his breakfast. To attend to the boiling water, he has to pause the process of having breakfast to finish the "water-boiling" goal, by turning off the stove and pouring the hot water. Then he can resume his goal of "having-breakfast". In this example, the user is pursuing two goals in an *interleaving* way, where one goal is paused and then resumed after executing some activities for pursuing a different goal.

Generally speaking, in real-world scenarios, there are five basic goal composition types in activity sequences, which are illustrated in Figure 1.

MG-Recognizer in (Chai & Yang 2005) tries to tackle the multiple-goal recognition problem, by creating finite state machines to model transitions between states of various goals in a deterministic way. Thus, the approach has trouble handling uncertainty, which is a major drawback. Another drawback is that the *MG-Recognizer* system did not explicitly consider the correlations between different goals.

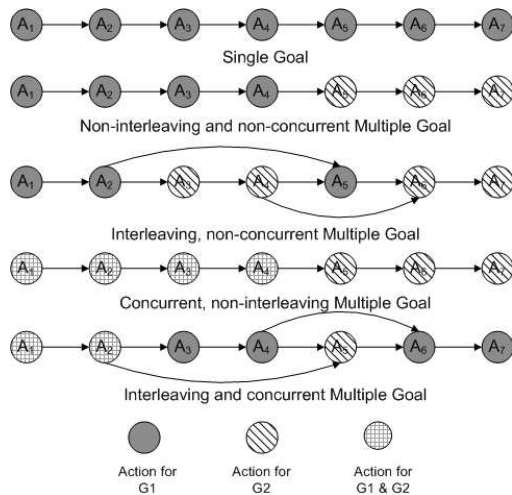


Figure 1: Goal composition types in activity sequences

In real-world situations, when we know that a user is pursuing one goal that has strong correlation with some other goals, there is high probability that he is pursuing these correlated goals at the same time. Hence, exploiting correlations between goals can help improve the accuracy of recognizing multiple goals. However, the *MG-Recognizer* system, as well as many previous approaches, did not handle this case either.

The main contribution of this paper is that we propose a novel two-level probabilistic and goal-correlation framework that deals with both concurrent and interleaving goals from observed activity sequences. Both single-goal recognition and multiple-goal recognition are supported by our solution. In order to reason about goals that can pause and continue through activities in the course of observations, we exploit skip-chain conditional random fields (SCCRF) (Sutton & McCallum 2004) at the lower level to estimate the probabilities of whether each goal is being pursued given a newly observed activity. To further consider the correlation between goals, a graph that represents the correlation between different goals is learned at the upper level. This goal graph allows us to infer goals in a “collective classification” manner. The probability inferred from the lower level is adjusted by minimizing a loss function via quadratic programming (QP) to derive a more accurate probability of all goals, taking the correlation graph into consideration. We show experimental results using several real-world data sets to demonstrate that our recognition algorithm is effective and accurate than several state-of-the-art methods.

Related Work

Activity recognition aims to recognize a user’s high-level behavior from the level of actions or sensor readings. Goal recognition, as a special case of activity recognition, refers to the task of identifying the goals that an agent is trying to accomplish. In this area, there are two major approaches in solving the problem of goal recognition: logic-based approaches and probabilistic-based approaches.

Logic-based approaches keep track of all logically consistent explanations of the observed activities. (Kautz 1987) provided a formal theory of plan recognition. However, logic-based approaches have limitations in distinguishing among consistent plans and have problems to handle uncertainty and noise in sensor data.

In probabilistic approaches, state-space models are especially attractive with the underlying assumption that there exist hidden states (e.g. activities and goals) of the world, and that the hidden states are evolving. State-space models enable the inference of hidden states given the observations up to the current time. They are suitable for modeling high-level hidden concepts given the low-level observations. Here we just name a few examples of state-space models in goal recognition: aggregate dynamic Bayesian networks (Patterson *et al.* 2005) and conditional random fields with its many variants (Vail, Veloso, & Lafferty 2007).

An alternative approach was proposed in (Chai & Yang 2005) to solve the multiple-goal recognition problem. In their approach, a finite state machine model is used for transitions between different states of goals. This framework was shown to be able to handle goals that are paused for a while in the middle of their achievements and then resumed, i.e. interleaving goals. However, as stated in the last section, the deterministic goal model cannot handle uncertainty.

In our algorithm, we will exploit SCCRf, which is especially attractive for our problem in meeting the requirement of goal recognition where goals can interleave. First proposed by (Lafferty, McCallum, & Pereira 2001), there has been an explosion of interest in CRFs these years, with successful applications in areas such as natural language processing, bioinformatics, information extraction, web page classification and computer vision. CRFs directly represent the conditional distribution over hidden states given the observations. Different from Hidden Markov Models, CRFs make no assumption on the dependence structure between observations, thereby making it very suitable for modeling complex relationships, specifying the relations between different labels and labeling all the data collectively.

In the past, researchers have been using CRFs and its Relational Markov Network (RMN) extension to solve goal recognition problems. (Liao, Fox, & Kautz 2005) defined a general framework for activity recognition using RMNs. (Liao, Fox, & Kautz 2007) uses a hierarchically structured conditional random field to extract a person’s activities and significant places from GPS data. (Wu, Lian, & Hsu 2007) proposed an algorithm using factorial conditional random field (FCRF) for recognizing multiple concurrent activities. This model can handle concurrency but cannot model interleaving activities and cannot scale up easily. So far, to the best of our knowledge, no algorithm has been proposed that deals with *both concurrent and interleaving activities* under uncertainty in a unified framework.

Our Proposed Method

We formally define our multiple-goal recognition problem. We assume that, as training data, a set S of observed activity sequences is given, without loss of generality, each sequence consists of T observed actions in the form of

$\{A_1, A_2, \dots, A_T\}$. We also assume that there are m goals which are used to label the activity sequences in all. Our objective is to train a model that can decide which subset of the m goals are being pursued given newly observed actions.

Modeling interleaving goals via SCCRF

Our model for interleaving goals is illustrated by the following example. Consider a professor who goes to the general office to get the projector for the “seminar” goal, he then goes to the printing room to pick up the printing material out for the “printing” goal. Finally, the professor may go down a corridor towards the seminar room. Through this example, we can observe that the “get projector” and “go towards seminar room” activities may have *long-distance dependencies* because the “seminar” goal is paused when the professor goes to the printing room. Generally, a goal G may be paused after an action a in a time slice t_i , and then resumed at a later time slice t_j with action b , where actions b and a are separated by several other actions for other goals.

We choose SCCRF proposed in (Sutton & McCallum 2004) to model the interleaving goal issue for the following reasons. Firstly, SCCRF has deep roots in Natural Language Processing (NLP). In NLP, the problem of Named Entity Recognition (NER) has similarities with the multiple-goal recognition problem, which needs to model the correlation between non-consecutive identical words in the text. Secondly, being a probabilistic graphical model, SCCRF has its advantage in modeling uncertainty in a natural and convenient way. Thirdly, the key issue in SCCRF is how to add skip edges. We use the posterior probabilities from the training data to add the skip edges. Based on the above reasons, we believe that SCCRF would be a model appropriate for handling the interleaving property of multiple goals.

At the lower level of our two-level framework, we consider recognizing each goal separately. Each SCCRF will infer whether an individual goal is active or not at each newly observed activity (see Figure 2). To model long-distance

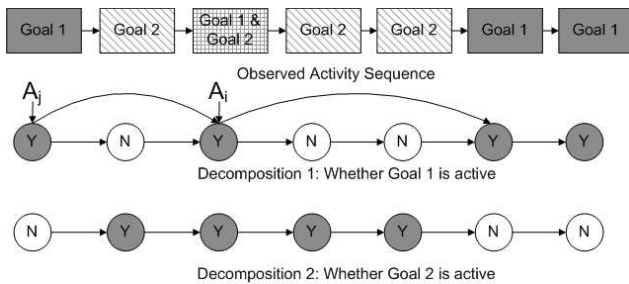


Figure 2: Decomposition into goal sequences

dependencies, for each goal G_k , we first infer the action-transition probability $P(A_i|A_j, G_k)$ (G_k is shown as Goal 1 in Figure 2), which stands for the probability of the following situation: given that the goal being pursued is G_k , the last action in the process of pursuing goal G_k is A_j and the next activity being A_i . This probability can be learned by the standard statistical method of Maximum Likelihood Estimation (MLE) or maximum a posteriori (MAP) estimation,

where a prior distribution is known. To simplify this preprocessing step, we assume the prior distribution is uniform and then employ the MAP approach in a standard way.

The main characteristic of a SCCRF model over the commonly used linear-chain CRF models is that the SCCRF model added a second type of potential, which was represented using long-distance edges, to the linear-chain model. For each of the m goals under consideration, we build a corresponding SCCRF model, with the i^{th} SCCRF being used to infer whether goal G_i is active given the set of observed activity sequences as training data.

Formally, for the k^{th} SCCRF model which is used to infer the probability of G_k being active, let y_t be a random variable whose value represents whether goal G_k is active or not given activity A_t , which occurs at time slice t . Let x_t be the observed activity at time slice t . For the factor graph $\mathcal{G} = \langle V, E \rangle$, it is essentially a linear-chain CRF with additional long-distance edges between activities A_i and A_j such that $P(A_i|A_j, G_k) > \theta$ (Refer to Figure 3 for an illustration). θ is a parameter that can be tuned to adjust the confidence of such long-distance dependencies. We will experimentally verify that small modifications of θ will not affect accuracy greatly.

For an observation sequence \mathbf{x} , let \mathcal{I} be the set of all pairs of activities for which there are skip edges connected with each other. Then the probability of a label sequence \mathbf{y} given an observation activity sequence \mathbf{x} is:

$$p(\mathbf{y}|\mathbf{x}) = \frac{1}{Z(\mathbf{x})} \prod_{t=1}^n \Psi_t(y_t, y_{t-1}, \mathbf{x}) \prod_{(u,v) \in \mathcal{I}} \Psi_{uv}(y_u, y_v, \mathbf{x}). \quad (1)$$

In Equation 1, Ψ_t are the factors for linear-chain edges and Ψ_{uv} are the factors over the skip edges. (Also refer to Figure 3 for illustration) $Z(\mathbf{x})$ is the normalization factor. We define the potential functions Ψ_t and Ψ_{uv} in Equation 2 and Equation 3 as:

$$\Psi_t(y_t, y_{t-1}, \mathbf{x}) = \exp \left(\sum_k \lambda_{1k} f_{1k}(y_t, y_{t-1}, \mathbf{x}, t) \right) \quad (2)$$

$$\Psi_{uv}(y_u, y_v, \mathbf{x}) = \exp \left(\sum_k \lambda_{2k} f_{2k}(y_u, y_v, \mathbf{x}, u, v) \right) \quad (3)$$

λ_{1k} are the parameters of the linear-chain template and λ_{2k} are the parameters of the skip-chain template. Each of them factorize according to a set of features f_{1k} or f_{2k} .




Figure 3: Illustration of the SCCRF model

Exact inference in CRFs maybe intractable as the time complexity is exponential in the size of the largest clique in the junction tree of the graph, and that there may be long and overlapping loops in the model. Loopy Belief Propagation (LBP) is used widely for performing approximate inference in CRFs and experiments show that LBP has been effective. Therefore, we set a maximum number of iterations, after which we can calculate the marginal probability of nodes.

Learning the weights λ_{1k} and λ_{2k} for the SCCRF model can be achieved by maximizing the log-likelihood of the training data, which requires calculating the partial derivative and optimization techniques. We omit the details of inference and parameter estimation. Interested readers can consult (Lafferty, McCallum, & Pereira 2001; Sutton, Rohanimanesh, & McCallum 2007) for technical details.

Modeling concurrent goals via correlation graph

In order to model correlations between concurrent goals, we need to know how similar and correlated two goals are. Similarly, correlation can also tell when one goal is being pursued (e.g. academic-related goals), other goals may be unlikely to be pursued at the same time (e.g. sports-related goals). Therefore, we wish to use the training data to build a correlation graph of goals, where two goals are related by an edge with a large positive weight in $[0, 1]$ if they have strong positive correlations. We omit the considerations of negative correlations here, which we leave for future work.

Note that a full-fledged Bayesian network can be built to model more complex correlations between goals in the form of conditional dependencies which are dependent on multiple random variables, such as $P(G_i|G_j, G_k)$. Furthermore, it is also possible to model concurrency and interleaving together in a CRF framework. However, in real-world situations, such kinds of complex dependencies between goals usually may not occur frequently, resulting the training data acquired to be too sparse to model such a probability and that the learned probability may be highly biased. Another reason is that usually the correlation between goals will not be known as prior knowledge, and such unknown structure adds expensive cost to training. In particular, combining the model of interleaving goals and concurrent goals via a CRF framework will make the training time intolerable, for which we will explain at the end of this section.

Therefore, we only model the probability $P(G_i|G_j)$ using our goal graph explicitly. We show in the experimental section that a factorial conditional random field (FCRF) (Wu, Lian, & Hsu 2007), which represents fully-connected goals through a Bayesian network structure and where goals are modeled in the CRF model, often does not perform as well as our correlation graph-based inference.

From the training data, we can infer the posterior probability $P(G_i|G_j)$ and use it as the initial similarity matrix. The reason why we do not take the currently observed activity into consideration and calculate posterior probability $P(G_i|G_j, A_t)$ is that in real-world situations, the activity sequence usually is not explicitly given as prior knowledge and should be inferred from sensor readings. Therefore, the activity inferred may have noise or bias, which may hinder the inference of probability of goal correlations. Another rea-

son is that we want to model the correlations between goals under a more general environment and assumption.

After calculating the posterior probability of each pair of goals, we take this value and define an $m \times m$ initial similarity matrix S as $S[i, j] = P(G_i|G_j)$. Since the training data may be sparse, the posterior probability we get from the training data may not be so reliable. (Blondel *et al.* 2004) proposed a method for computing the similarity matrix between vertices of different graphs. We adapted their method for modeling the similarity between vertices of the same graph. We build a directed graph $\mathcal{G} = \langle V, E \rangle$, where the vertices V indicate different goals and $e = \langle G_a, G_b \rangle$ indicates that a goal G_a and a goal G_b have some kind of connection, so that when G_a appears, G_b is also likely to appear. The similarity matrix is updated through iterations of $S_{k+1} = AS_kA^T + A^T S_k A$, where A denotes the adjacency matrix of the similarity graph, where $A[i, j] = 0$ if $P(G_i|G_j) = 0$, otherwise $A[i, j] = 1$. Here S_0 is the initial similarity matrix as defined above. (Blondel *et al.* 2004) proved the convergence property of this update function. When the iteration procedure converges, some of the edge weights will become zero.

Given m goals, we infer m initial posterior probabilities $P'_i, i = 1, 2, \dots, m$ from the SCCRF model, which means the probability of goal i being active given a particular observed activity. Then we create the similarity matrix using the probability and the update function mentioned above. After creating the similarity matrix S , we can model concurrent goals by minimizing the differences between strong correlated goals (i.e. $P(G_i|G_j)$ is rather large), to ensure that they will appear together, and minimizing the differences between adjusted posterior probability of a goal P_i and its initial posterior probability P'_i from its individual SCCRF, since this probability carries the observed evidence.

As a result, our top level inference consists of minimizing the following objective function with similarity matrix S and initially inferred probabilities $\mathbf{P}' = \{P'_1, P'_2, \dots, P'_m\}$ and our desired output $\mathbf{P} = \{P_1, P_2, \dots, P_m\}$.

$$\min \sum_{i,j \in \{1, \dots, m\}} (P_i - P_j)^2 S_{ij} + \mu (P_i - P'_i)^2 \quad (4)$$

The new probabilities P_i are then used as our predictions. Considering the similarity matrix \mathbf{S} , as $S[i, j]$ increases towards 1, the difference between P_i and P_j needs to decrease in order to minimize the objective function. The parameter μ can be tuned to reflect the importance of the initial posterior probability learnt from the SCCRF model.

Next, we show that the optimization problem mentioned above can be formulated as a quadratic programming (QP) problem and solved using standard techniques in QP.

We define vector $\mathbf{P} = [P_1, P_2, \dots, P_m]$ and vector $\mathbf{P}' = [P'_1, P'_2, \dots, P'_m]$. Then the problem can be expressed as an optimization problem:

$$\begin{aligned} \min_{\mathbf{P}} \quad & P^T (L_S + \mu I) P - 2\mu (\mathbf{P}')^T P \\ \text{s.t.} \quad & 0 \leq P_i \leq 1, \quad i \in \{1, 2, \dots, m\} \end{aligned} \quad (5)$$

Algorithm 1 Multiple Goal Recognition: *CIGAR*

Input: T is the length of an observed activity sequence $\mathbf{A} = \{A_1, A_2, \dots, A_T\}$ and m goals G_1, G_2, \dots, G_m .

Output: P_j^i is the probability of goal G_j to be active given observed activity A_i .

- 1: **for** $i = 1$ to m **do**
 - 2: Learn the posterior probability $P(A_j|A_k, G_i)$ for every pair of actions A_j and A_k .
 - 3: Add a skip edge between y_j and y_k in the i^{th} SCCRF if $P(A_j|A_k, G_i) > \theta$.
 - 4: Train the corresponding i^{th} SCCRF model.
 - 5: **end for**
 - 6: **for** $i = 1$ to T **do**
 - 7: **for** $j = 1$ to m **do**
 - 8: Infer probability $P_j^{i'}$, which represents whether goal G_j is active at time slice i .
 - 9: **end for**
 - 10: Adjust the initial inferred probability and get the adjusted inferred probability P_j^i with QP.
 - 11: **end for**
-

L_S is the Laplacian of S , and is defined as $L_S = D - S$. D is a diagonal matrix where $D[i, i] = \sum_{j=1}^m S_{i,j}$. Equation 4 can be shown to lead to Equation 5. Also, it is evident to show that $(L_S + \mu I)$ is positive definite. Thus the above QP formulation is convex and always has only one global optimum. Furthermore, many state-of-the-art methods can solve QP problems efficiently.

Putting the above together, our main *CIGAR* algorithm is shown in Algorithm 1. We analyze the time complexity of our algorithm and compare with the referenced FCRF method (Wu, Lian, & Hsu 2007). Assume that training a CRF with T nodes requires time $O(V)$. Therefore, training m CRFs with T nodes only require a complexity of $O(mV)$. However, worst case analysis shows that training a FCRF with mT nodes require a complexity of $O(V^{mT})$, where m is the number goals and T is the number of activities. Therefore, another advantage of our algorithm over the FCRF method is that our algorithm is more scalable than the FCRF method. The reason also applies to why we did not model concurrent goals explicitly in the CRF model. In the future, we plan to use other methods for training CRF, like the Virtual Evidence Boosting (Liao *et al.* 2007), hoping that we can achieve better accuracy as well as improved training time with the new method.

Experimental Results²

In previous sections, we have described our *CIGAR* approach for recognizing multiple goals in an activity sequence to allow concurrent and interleaving goals. In this section, we will present experimental results of our model to demonstrate that it is both accurate and effective. We compare our algorithm *CIGAR* to the following competing methods. (1) *SCCRF*: interleaving but not concurrent goal recognizer,

²All the datasets used in this section can be downloaded via <http://www.cse.ust.hk/~derekhh/>.

which applies SCCRF model without correlation graph; (2) *MG-Recognizer*: multiple goal-recognition algorithm presented in (Chai & Yang 2005), with several finite state machines which have different states indicating whether a goal is evolving, suspending or terminating; (3) *FCRF*: which builds a factorial conditional random field (FCRF) over the observed activity sequence, as presented in (Wu, Lian, & Hsu 2007). We show that our *CIGAR* algorithm can outperform these baseline algorithms. We use three datasets in a cross validation setting to get the recognition accuracy against the baseline methods. Recognition accuracy is defined as the percentage of correctly recognized goals over all goals across all time slices for all the activity sequences.

The first domain is from (Chai & Yang 2005) where the observations are obtained directly from sensor data and the activities correspond to that of a professor walking in a university office area. In this data set, nine goals of a professor’s activities are recorded, 850 single-goal traces, 750 two-goal traces and 300 three-goal traces are collected so that the dataset can evaluate both multiple-goal recognition and single-goal recognition. We used three-fold cross validation for training and testing. Table 1 shows the comparison in recognition accuracy for both single and multiple-goal recognition tasks. We also tested the performance of our algorithm with different parameter settings. As we can see, *CIGAR* achieves the best performance among all baseline methods, also, small modifications of the parameters θ and μ won’t change the recognition accuracy much. Note that FCRF performs much worse in the multiple goal dataset. This is because FCRF did not model interleaving goals.

Algorithm	Single	Multi
<i>MG-Recognizer</i>	94.6%(3.3)	91.4%(4.7)
<i>FCRF</i>	93.6%(5.7)	74.4%(3.8)
<i>SCCRF</i> ($\theta = 0.7$)	94.0%(2.5)	93.5%(2.6)
<i>SCCRF</i> ($\theta = 0.8$)	94.9%(2.8)	93.1%(3.9)
<i>SCCRF</i> ($\theta = 1$)	94.8%(2.9)	91.6%(2.9)
<i>CIGAR</i> ($\theta = 0.7, \mu = 0.4$)	94.0%(2.7)	95.3%(3.4)
<i>CIGAR</i> ($\theta = 0.7, \mu = 0.5$)	94.8%(2.7)	94.5%(3.7)
<i>CIGAR</i> ($\theta = 0.7, \mu = 0.6$)	94.2%(2.7)	94.4%(3.2)

Table 1: Comparison in office dataset

We also used the dataset collected in (Patterson *et al.* 2005) to further test the accuracy of our algorithm. In this dataset, routine morning activities which used common objects interleavingly are detected through sensors and recorded as sensor data. In this domain, there are a lot of interleaving activities, but there are no concurrent activities. Ten-fold cross-validation is used for testing on this dataset. Table 2 shows the comparison in recognition accuracy for this dataset. As we can see, *SCCRF* and *CIGAR* performs the best amongst all other methods. Note that there are no concurrent goals in this domain, QP actually does no adjustment of the inferred probabilities from the *SCCRF*.

The last dataset we are using is the MIT PlaceLab dataset from (Intille *et al.* 2006) and also used for the activity recognition experiment in (Wu, Lian, & Hsu 2007). We used the *PLIAI* dataset, which was recorded on Friday March 4, 2005

Algorithm	Accuracy (Variance)
<i>MG-Recognizer</i>	85%(4.6)
<i>FCRF</i>	83%(3.3)
<i>SCCRF</i> ($\theta = 0.7$)	92%(5.4)
<i>SCCRF</i> ($\theta = 0.8$)	91%(6.2)
<i>SCCRF</i> ($\theta = 1$)	91%(5.9)
<i>CIGAR</i> ($\theta = 0.7, \mu = 0.4$)	92%(5.2)
<i>CIGAR</i> ($\theta = 0.7, \mu = 0.5$)	92%(5.4)
<i>CIGAR</i> ($\theta = 0.7, \mu = 0.6$)	92%(5.0)

Table 2: Comparison in (Patterson *et al.* 2005) dataset

from 9AM to 1PM with a volunteer in the MIT PlaceLab. Note that in this dataset, we are using the location information to predict what activity the user is currently pursuing. Since the original dataset may not contain many concurrent activities, we follow the method in (Wu, Lian, & Hsu 2007) to cluster the 89 activities into six categories where each category corresponds to a new goal. In this way, both interleaving and concurrent activities can be modeled. Table 3 shows the comparison in recognition accuracy for the MIT PlaceLab dataset. In this dataset, *CIGAR* performs much better than the baseline methods.

Algorithm	Accuracy(Variance)
<i>MG-Recognizer</i>	68% (4.1)
<i>FCRF</i>	73% (3.8)
<i>SCCRF</i> ($\theta = 0.7$)	80%(3.1)
<i>SCCRF</i> ($\theta = 0.8$)	80%(3.3)
<i>SCCRF</i> ($\theta = 1$)	79%(4.5)
<i>CIGAR</i> ($\theta = 0.7, \mu = 0.4$)	84%(4.3)
<i>CIGAR</i> ($\theta = 0.7, \mu = 0.5$)	86%(3.0)
<i>CIGAR</i> ($\theta = 0.7, \mu = 0.6$)	85%(3.3)

Table 3: Comparison in MIT PlaceLab dataset

Hence, from the above experiments, we show that *CIGAR* can perform significantly better than baseline methods, and that *CIGAR* can better model concurrent and interleaving goals in real-world situations.

Conclusions

In this paper, we proposed a two-level framework for inferring the user’s high-level goals from activity sequences, meeting the real-world requirement that goals are often interleaving and concurrent. We improve previous algorithms with probabilistic transitions and considered the advantage of exploiting correlations between different goals. Experimental results show that our algorithm achieves better accuracy than baseline methods.

Our work can be extended in several directions. One is that our models can be adapted into an online inference algorithm such that the real-world requirement is better modeled. Also, the effect of negative or more complex correlations between goals may be considered. Another is that we could try to use some other CRF training methods for better accuracy and training complexity.

Acknowledgements

The authors are supported by a grant from Hong Kong CERG Project 621606. We thank Yufeng Li for providing valuable suggestions when writing this paper. We also thank Xiaoyong Chai for providing relevant code and datasets for our experiment.

References

- Blondel, V. D.; Gajardo, A.; Heymans, M.; Senellart, P.; and Dooren, P. V. 2004. A measure of similarity between graph vertices: Applications to synonym extraction and web searching. *SIAM Rev.* 46(4):647–666.
- Chai, X., and Yang, Q. 2005. Multiple-goal recognition from low-level signals. In *AAAI*, 3–8.
- Intille, S. S.; Larson, K.; Tapia, E. M.; Beaudin, J.; Kaushik, P.; Nawyn, J.; and Rockinson, R. 2006. Using a live-in laboratory for ubiquitous computing research. In *Pervasive*, 349–365.
- Kautz, H. 1987. *A Formal Theory of Plan Recognition*. Ph.D. Dissertation, University of Rochester.
- Lafferty, J. D.; McCallum, A.; and Pereira, F. C. N. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *ICML*, 282–289.
- Liao, L.; Choudhury, T.; Fox, D.; and Kautz, H. A. 2007. Training conditional random fields using virtual evidence boosting. In *IJCAI*, 2530–2535.
- Liao, L.; Fox, D.; and Kautz, H. A. 2005. Location-based activity recognition using relational markov networks. In *IJCAI*, 773–778.
- Liao, L.; Fox, D.; and Kautz, H. A. 2007. Extracting places and activities from gps traces using hierarchical conditional random fields. *IJRR* 26(1):119–134.
- Liao, L. 2006. *Location-Based Activity Recognition*. Ph.D. Dissertation, University of Washington.
- Patterson, D. J.; Fox, D.; Kautz, H. A.; and Philipose, M. 2005. Fine-grained activity recognition by aggregating abstract object usage. In *ISWC*, 44–51.
- Pollack, M. E.; Brown, L. E.; Colbry, D.; McCarthy, C. E.; Orosz, C.; Peintner, B.; Ramakrishnan, S.; and Tsamardinos, I. 2003. Autominder: an intelligent cognitive orthotic system for people with memory impairment. *Robotics and Autonomous Systems* 44(3-4):273–282.
- Sutton, C., and McCallum, A. 2004. Collective segmentation and labeling of distant entities in information extraction. Technical Report TR 04-49, University of Massachusetts Amherst.
- Sutton, C. A.; Rohanimanesh, K.; and McCallum, A. 2007. Dynamic conditional random fields: Factorized probabilistic models for labeling and segmenting sequence data. *JMLR* 8:693–723.
- Vail, D. L.; Veloso, M. M.; and Lafferty, J. D. 2007. Conditional random fields for activity recognition. In *AAMAS*, 1331–1338.
- Wu, T.; Lian, C.; and Hsu, J. Y. 2007. Joint recognition of multiple concurrent activities using factorial conditional random fields. In *AAAI Workshop PAIR 2007*.