PROBLEM SET 1

Emmanuel Munguia Tapia
emunguia@mit.edu

Part A: Topics of fascination

Dynamic Bayesian Networks and Continuous Time Bayesian Networks

Dynamic Bayesian networks are nowadays the method of choice for supervised sequential learning. One particular case of DBN is hidden markov models that are widely used in different domain such as speech recogntion. One extension to DBNs are continuous time Bayesian networks. Continuous time Bayesian networks (CTBN) describe structured stochastic processes with finitely many states that evolve over continuous time. A CTBN is a directed (possibly cyclic) dependency graph over a set of variables, each of which represents a finite state continuous time Markov process whose transition model is a function of its parents. CTBNs are relatively new in the field of machine learning, they provide a way to do inference on continuous state spaces without having to discretize the state space as usually done in dynamic Bayesian networks. I would like to learn more about them and probably implement a reasoning system using DBNs or CTBNs.

Filtering and state space estimation with particle filters

Particle filters are nowadays very popular and probably the method of choice for doing inference on Bayesian networks and dynamic Bayesian networks. The reason is because they are an approximated way of doing inference based on monte-carlo sampling techniques and their accuracy depend on the number of particles used (samples of the system). Also, particle filters are widely used in position estimation since they can represent and solve non-linear problems that the Kalman filter and EKF cannot solve or takes excessive time to solve. I would like to gain practical experience by using this technique applied to inference on large Bayesian networks and state spaces.

Reinforcement learning

There is a lot of research done in applying machine learning algorithms and pattern recognition techniques to infer human activity by smart homes and robots. However there is very little research on how to fix the models of human activity in real-time if they are not performing well. More importantly, most of these ML techniques require extensive training data to perform with a reasonable accuracy and most of the time, in the home setting, it is not feasible to assume that the user will provide 50-100 examples of activities to recognize. I believe that novel reinforcement learning techniques could provide a robust way to fix the models in real-time by just having to provide answers such as ‘Yes’ or ‘No’ whenever the system correctly or incorrectly predicts the human activities.

Part B. Your cognitive robot

My dream scenario is a stationary robot: your own home. Nowadays, homes do not have even the minimal amount of intelligence to know what its inhabitants are up to. If a home was able to understand people’s activities, it could monitor the health of its inhabitants, automate tasks, interrupt the user at convenient times as well as present the information at the right time or teach somebody something at the right moment.

Scenario: Imagine the scenario of an elderly person leaving alone. One of the main difficulties for relatives is to identify when this person needs help performing daily tasks because they are, for example, no longer able to clean the house or go shopping. Another important aspect is to know when an elderly person stops performing activities of daily living which are essential for self survival such as preparing meals, going to the toilet, taking a bath and taking medication.

In order for a robot (home) to be able to monitor and identify when an elderly people is no longer able to perform a specific task or activity, the home needs to infer the person’s activities. One way to do it, is to place small wireless sensors in different everyday objects and have machine learning and pattern recognition algorithms that infer the person’s activities based on the sensor activations and sensor values.

Another application, is to infer when a person is behaving differently from the usual behavior pattern. In this case, the robot needs to first establish a baseline and a model of how the person’s activities look like over extended periods of time. After this baseline has been established, the robot could infer that the person is behaving differently by comparing the baseline against the new activity patterns detected. Activity changes are important because medical professionals believe that people’s activity patterns change right before people get sick (elderly dementia, schizophrenia and even a simple flu).

Diagram of software architecture

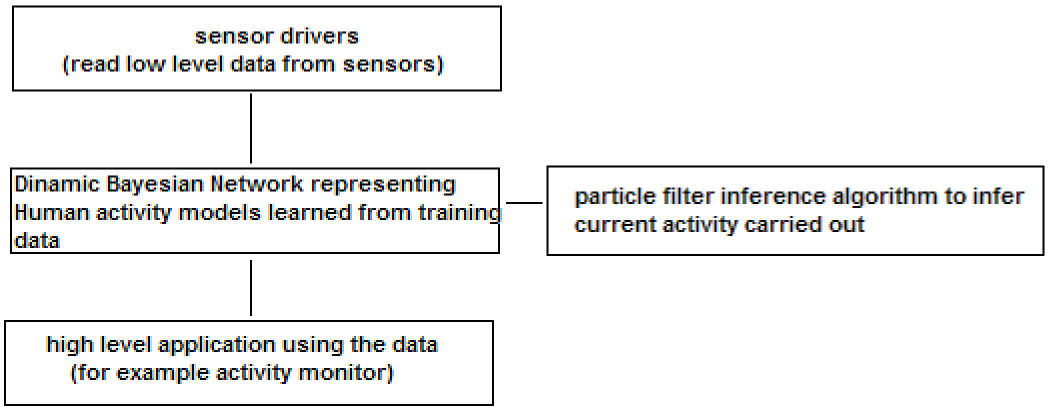

Part C: Making your cognitive robot

I think that the most important form of reasoning that my system needs to act capably is dynamic Bayesian networks in combination with particle filtering techniques to infer what is the most likely activity being carried out by the person. Ideally, I would have one DBN model or network per activity to be recognized (bathing, toileting, preparing lunch, etc). Then, based on the sensors firing being detected, I would use particle filters to infer what the most likely activity is.

Part D: Researching a critical Reasoning method

The three papers selected are the following:

Philipose, M., K. Fishkin, et al. (2003). "Guide: Towards Understanding Daily Life via Auto-Identification and Statistical Analysis." UbiHealth Workshop at Ubicomp 2003.

In this paper, a system for recognizing activities in the home setting is proposed by using a combination of dynamic Bayesian networks and RFID tags attached to everyday objects in the environment. I think that this is a very interesting approach to the problem of figuring out what a person is doing in real-time. The authors assume that a person using such a system in the home would wear a wearable RFID glove in their hand that will read the RFID tags every time the person interact with different everyday objects. The main idea is the following: You wake up, touch the toothpaste, then the toothbrush and also open the cold water faucet. If you have sensors in these different objects and an inference algorithm, a computer or robot could infer that you are brushing your teeth.

Another novelty of the paper is the fact that they automatically extract human activity models from the web by mining websites such as ehow.com that have “recipes” on how to carry out different human activities. Thus, this approach do not require extensive examples of activities in order to train the machine learning algorithms. The authors compile the models of human activities from the web in dynamic Bayesian networks and use particle filters to infer what the most likely model is given the trace of RFID sensor activations. I basically have two different criticisms: one is that we cannot assume that a person at home will always use a wearable glove that is the basis of the system to work. Furthermore, I think that the approach of compiling human activities mined from the web into DBNs is promising, however, I believe that it would be important to provide a way to fix or retrain the system in case the system is not recognizing the activities accurate enough. I think this is the main flaw of this particular approach. This paper is relevant to the cognitive robotic system I am planning to build because I will use a similar approach. Instead of making the person wear a glove, I will place the sensors in the environment. Small sensors everywhere that could provide me information about object’s usability.

Wilson, D. "Simultaneous Tracking & Activity Recognition (STAR) Using Many Anonymous, Binary Sensors." Pervasive 2005.

In this paper, a system for recognizing activities in the home setting and for tracking people’s location in a room base level is proposed. The system uses sensors commonly found in current security systems such as contact reed switches located in doors and windows, as well as motion sensors. For the inference, the author proposes an inference algorithm based on layers of dynamic Bayesian networks in order to infer the inhabitant’s activities. The beauty of using layers of DBNs is that the lower layers of the DBNs can be retrained to compensate for changes in sensor readings due to noise or different home settings while keeping the upper layers with no changes. Given the limited amount of sensors that the author suggests, the granularity of the system in terms of the different activities that it is able to recognize is not big enough. The author proposes to recognize a subset of 5 different everyday activities as well as to track one persons location in the environment. This paper is relevant because it uses sensors in the environment to track people’s location as well as people’s activities. I like the idea of using different layers of dynamic Bayesian networks to perform inference at different levels of granularity about human activities. I will probably implement something similar for my final project and probably add another layer for correcting the system whenever a misclassification is performed. Finally the system proposes to train the system by providing explicit training examples which I think might be impractical for some activities. I believe that the combination of prior information about human activities (probably mined from the web) as well as from explicit training examples could prove to be particularly powerful.

Patterson, D., D. Fox, et al. (2003). "Expressive, Tractable and Scalable Techniques for Modeling Activities of Daily Living." UbiHealth 2003: The 2nd International Workshop on Ubiquitous Computing for Pervasive Healthcare Applications.

This paper is relevant since it touches on different aspects of tractability and scalability in methods for inferring and modeling human activity. The methods proposed are similar to dynamic Bayesian networks and the main inference algorithm proposed is again a particle filter. Prior information about human activities is provided in the form of activity models. The author proposes some changes to the conventional way probabilities are used in DBN systems and proposes the concept of mutually exclusive probabilities. This paper is relevant because it also uses similar reasoning techniques to infer people's activities but touches on issues of scalability and tractability.

Part E: A Simple Project for your Cognitive Robot

I think that for my final project I will implement a stationary robot embedded in the home environment that will be able to reason about its inhabitants activities and take intelligent decisions according. One possible approach could be to infer the right time to interrupt the user to provide some information or reminder. Another interesting approach could be to try to do some action execution monitoring for elderly people. Some people suffering from dementia are have problems performing activities such as preparing coffee or tea because they forget if they have already put sugar in the coffee. A system that would provide assistance is such domain should be able to take readings from sensors placed in different objects, infer the current step in an activity that the user is in, and provide a plan of action towards what the next step in the activity sequence should be.

This problem combines interesting aspects: A reasoning algorithm to infer the current system state, and an action/planning algorithm to find what the next step or state of the system should be.