Behavior Recognition in Assisted Cognition

Lin Liao, Donald Patterson, Dieter Fox, Henry Kautz
Department of Computer Science & Engineering
University of Washington
Seattle, WA 98195

Introduction

If a computer could recognize users' behaviors automatically, it could interact with users in a more efficient and friendly manner. The problem of behavior recognition, however, has been an active research topic for a long time (Kautz 1987; Charniak & Goldman 1993; Pynadath 1999; Bui 2003) and still remains very challenging. While most early systems were focused on simulations or toy examples, more recent research has begun to build behavior recognition systems that work with real-world data (Pollack *et al.* 2002; Patterson *et al.* 2004a; Liao, Fox, & Kautz 2004). There are three main technical advancements that have made these systems possible. First, various sensing technologies such as the Global Positioning System (GPS), Radio Frequency Identification (RFID) tags, digital cameras, ultrasound sensors, infrared sensors, light sensors, and motion sensors are maturing and becoming widely used. Thus, it becomes much easier for computers to sense the physical world. Second, great advances have been achieved in probabilistic reasoning, especially for large and complex systems. Third, large amount of background knowledge can be acquired through the World Wide Web and other related technologies.

Behavior recognition is one of the central pieces in the Assisted Cognition project at University of Washington (Kautz *et al.* 2002). By combining Artificial Intelligence and ubiquitous computing technologies, the Assisted Cognition project aims to augment human capabilities, with a particular emphasis on increasing the independence of people suffering from cognitive limitations, *e.g.,* Alzheimer's disease patients. In this paper, we explain our recent progress on probabilistic behavior recognition and discuss future research directions. In particular, we focus on the outdoor activities; see (Patterson *et al.* 2004a) for more about indoor behavior recognition.

Probabilistic Behavior Recognition

Human behaviors involve many uncertain factors and are hard to model deterministically. Probabilistic reasoning thus becomes a promising approach for behavior recognition. Even though we may regard probabilistic behavior recognition just as a probabilistic inference system, it has some distinct features.

- First, many factors that affect behaviors—the preferences and capabilities of the people as well as the constraints from physical world—are very difficult to characterize.
- Second, usually a big gap exists between the low level sensor measurements (*what we can observe*) and high level behaviors (*what we want to know*).
- Third, human behaviors often have inherent complex structures and relations. For example, Kautz pointed out two basic structures for behaviors: *decomposition* and *abstraction*(Kautz 1987).
- Fourth, people often make mistakes. Recognizing users' errors is especially important when we want to assist cognitively-impaired people.

Behavior Decomposition

Behaviors can usually be decomposed into a number of components. For example, a simple trip "going to campus" may include several segments such as "walking to the bus stop," "taking the bus to campus," and "walking to the office." Each segment may be further decomposed. We can represent such relations using a *decomposition hierarchy*. The hierarchical structure is very important for behavior recognition, because it bridges the gap between low level measurements and high level behaviors.

Bui introduced the abstract hidden Markov model, which used hierarchical representations to infer a person's indoor behaviors(Bui 2003). Efficient inference algorithms were designed using Rao-Blackwellised particle filters. Recently, we developed a hierarchical structure to model people's transportation routines (Patterson *et al.* 2003; Liao, Fox, & Kautz 2004). Using this hierarchy, we can infer users' future goals and transportation modes (*i.e.,* car, bus or foot) from raw GPS measurements. Furthermore, we showed how to learn the parameters of the model in a completely unsupervised manner using Expectation-Maximization (EM).

Behavior Abstraction

For behavior abstraction, we mean two things. First, behaviors can be categorized into a set of classes. For example, "visiting friend John" and "visiting friend Mike" could both be abstracted as "visiting friends." In particular, people's preference or patterns are often characterized in abstract ways, *e.g.,* "Lin usually visits friends on Friday evening." Abstraction allows us to learn general patterns with limited training data and predict novel events. For instance, if our system identifies Lin's behavior pattern on Friday evening as we just metioned, it might be able to recognize his behavior even when he is visiting a friend whom he has not visited before. Second, we could generalize a common behavior pattern from individuals' behaviors. A general model such as this could prove to be essential in order to help a user who has not had a chance to provide sufficient training data to the computer.

Such abstractions could be represented using Probabilistic Relational Models (Koller & Pfeffer 1998; Sanghai, Domingos, & Weld 2003). These models use *relational schemas* to explicitly characterize the relations between classes and share parameters among different objects. Learning in this framework is a challenging task (Getoor *et al.* 2001). In some situations, the relations may depend on complex features that span a period of time in which case discriminative learning might prove to be more appropriate than generative models (Taskar, Abbeel, & Koller 2002).

External Knowledge

Knowledge plays an important role when people try to recognize others' behaviors. When used properly, it could help computers as well. In (Patterson *et al.* 2003), we showed how to use the knowledge about street maps, bus stops and bus routes to help with inferring transportation modes. Today, the World Wide Web has become a huge knowledge base that could be very helpful. For example, (Perkowitz *et al.* 2004) showed how to mine models of human activities from the Web. In (Patterson *et al.* 2004b), we used dynamic information from the web, such as real time bus schedules, to guide users back on track after the system detected user errors.

Error Detection

Behavior recognition is hard; user errors make it even harder. In (Liao, Fox, & Kautz 2004), we used a simple model selection technique to detect user errors or novelties. More specifically, we use two different trackers simultaneously. The first tracker assumes the user is behaving according to his personal historical trends and uses the learned hierarchical model for tracking. The second tracker assumes abnormal activities and thus uses a prior model that accounts for general physical constraints but is not adjusted to the user's past routines. The trackers are run in parallel and the probability of each model given the observations is calculated. When the user is following his ordinary routine, the first tracker should have a higher probability; when the user does something unexpected, the second model should become more likely.

The above approach is simple, but unable to distinguish errors from deliberate novel behavior, or discriminate kinds of errors. In order to do so we must employ explicit error models. Because the set of possible errors is enormous (much larger than the set of normative behaviors), a key research challenge is to induce rules for *generating* error models from general principles of purposeful behavior and a very sparse set of examples.

References

Bui, H. H. 2003. A general model for online probabilistic plan recognition. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Charniak, E., and Goldman, R. P. 1993. A bayesian model of plan recognition. In *Proc. of the National Conference on Artificial Intelligence (AAAI)*.

Getoor, L.; Friedman, N.; Koller, D.; and Taskar, B. 2001. Learning probabilistic models of relational structure. In *Proc. of the International Conference on Machine Learning (ICML)*.

Kautz, H.; Arnstein, L.; Borriello, G.; Etzioni, O.; and Fox, D. 2002. An overview of the assisted cognition project. In *AAAI Workshop on Automation as Caregiver: The Role of Intelligent Technology in Elder Care*.

Kautz, H. 1987. *A Formal Theory of Plan Recognition*. Ph.D. Dissertation, University of Rochester.

Koller, D., and Pfeffer, A. 1998. Probabilistic frame-based systems. In *Proc. of the National Conference on Artificial Intelligence (AAAI)*.

Liao, L.; Fox, D.; and Kautz, H. 2004. Learning and inferring transportation routines. In *Proc. of the National Conference on Artificial Intelligence (AAAI)*.

Patterson, D. J.; Liao, L.; Fox, D.; and Kautz, H. 2003. Inferring high-level behavior from low-level sensors. In *International Conference on Ubiquitous Computing (UbiComp)*.

Patterson, D. J.; Fox, D.; Kautz, H.; Fishkin, K.; Perkowitz, M.; and Philipose, M. 2004a. Contextual computer support for human activities. In *AAAI Spring Symposium: Interaction between Humans and Autonomous Systems over Extended Operations*.

Patterson, D. J.; Liao, L.; Gajos, K.; Collier, M.; Livic, N.; Olson, K.; Wang, S.; Fox, D.; and Kautz, H. 2004b. Opportunity knocks: a system to provide cognitive assistance with transportation services. In *International Conference on Ubiquitous Computing (UbiComp)*. submitted.

Perkowitz, M.; Philipose, M.; Patterson, D. J.; and Fishkin, K. 2004. Mining models of human activities from the web. In *Proc. of The International World Wide Web Conference (WWW)*.

Pollack, M. E.; .McCarthy, C. E.; Ramakrishnan, S.; Tsamardinos, I.; Brown, L.; Carrion, S.; Colbry, D.; Orosz, C.; and Peintner, B. 2002. Autominder: A planning, monitoring, and reminding assistive agent. In *7th International Conference on Intelligent Autonomous Systems*.

Pynadath, D. V. 1999. *Probabilistic Grammars for Plan Recognition*. Ph.D. Dissertation, University of Michigan.

Sanghai, S.; Domingos, P.; and Weld, D. S. 2003. Dynamic probabilistic relational models. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Taskar, B.; Abbeel, P.; and Koller, D. 2002. Discriminative probabilistic models for relational data. In *Proc. of the Conference on Uncertainty in Artificial Intelligence (UAI)*.