Distributed Collaborative Activity Recognition

Qiang Yang, qyang@cse.ust.hk, http://www.cse.ust.hk/~qyang
Department of Computer Science and Engineering
Hong Kong University of Science and Technology, Hong Kong

1. Participant Background and Experience

Yang's group has been actively working in the field of activity recognition based on mobile devices. Their work includes transfer learning for activity recognition and activity recognition for indoor daily-living activity recognition via mobile devices (see http://www.cse.ust.hk/~qyang/byarea.htm#Activity_Recognition) at location, action and goal levels. Yang was an invited speaker at IJCAI Conference in 2009 to give a talk on activity recognition where he gave an overview of the dynamic research field (http://videolectures.net/ijcai09_yang_fllshli/). One of the systems developed by Yang's group can transform the received sensor signals to predicted higher-level user actions and goals such as 'taking a bus', 'turning on the stove', etc. Then, once the action sequences are partially known, the system can also infer even higher-level user goals to be achieved, even when these goals are interleaving and concurrent. Using detected activities. Yang's group gave an overview of their work in a video and via a software system known as VTrack. Recently, they have applied transfer learning to activity recognition in a system known as cross-domain activity recognition. Yang and his collaborators have designed a system that combines collaborative filtering in recommender systems, GPS and various other sources of data, into an integrated intelligent activity recognition system to recommend activities and venues for visitors in a large city. His work also covers mobile phone based abnormal activity recognition and activity profiling for the elderly, in order to provide e-health support. He has also been working with his students on the topic of collaborative filtering and large-scale cloud-based computation. In this area, Yang and his students have been active in research on recommendation systems and Web mining. They also developed a collaborative filtering approach that addresses the item-ranking problem directly by modeling user preferences derived from the ratings, and a novel transfer-learning methods for knowledge transfer between different collaborative filtering tasks (see http://www.cse.ust.hk/~qyang)

2. Our Vision: Distributed Collaborative Activity Recognition

Consider the following scenario: another old parent lives alone at his apartment in a big city. His cell phone has software systems installed to allow for *long-term, population-based activity-recognition (AR)*. As he wanders in an unusual pattern, his movement pattern is recorded and transmitted to the distributed system, where the pattern is extracted and compared to patterns mined from many users to identify a similar group. The system then exploits distributed collaborative filtering to query different systems to compare his current trajectory and movement patterns to decide whether he is lost, shows early sign of AD, and/or needs assistance. An alert message can be automatically sent to the hospital with his identity and location information. An emergency medical crew arrives in good time to take the person to hospital for diagnosis, treatment and home.

This home-based healthcare scenario requires several technological advances to support (see Figure 1). First, since individual users' data are likely to be very limited, any mining results from an individual's activity is unlikely to generate any significant patterns. To resolve this issue, we need to develop a collaborative system for activity recognition, where the collaboration comes from the user dimension. Similar to collaborative filtering in recommender systems, a collaborative activity recognition system will borrow the useful data or knowledge from users with similar behaviors to address the lack of data problem for each individual user. We think a transfer learning or a multitask framework will be effective in filtering in the useful knowledge and boost the performance.

Second, we envision that the user behavioral data are stored on many different servers (Figure 1). For example, the data from each elderly care center may be stored on one server. To exchange the whole data set between all different servers would be unreasonable and undesirable due to concerns such as privacy. Thus, to enable collaborative activity recognition, some kind of distributed data mining algorithms must be designed so that only useful knowledge is transferred between the different servers, instead of the raw data themselves. As shown in Figure 1, a master node can handle the coordination of transfer of knowledge between slave servers.

Finally, accurate activity recognition techniques need be redesigned to take into account a massive number of potentially dynamic data characterized in several dimensions: spatial, temporal, users and devices. Traditional algorithms for activity recognition need to be re-examined to ascertain their ability to handle activity recognition in massive data streams.

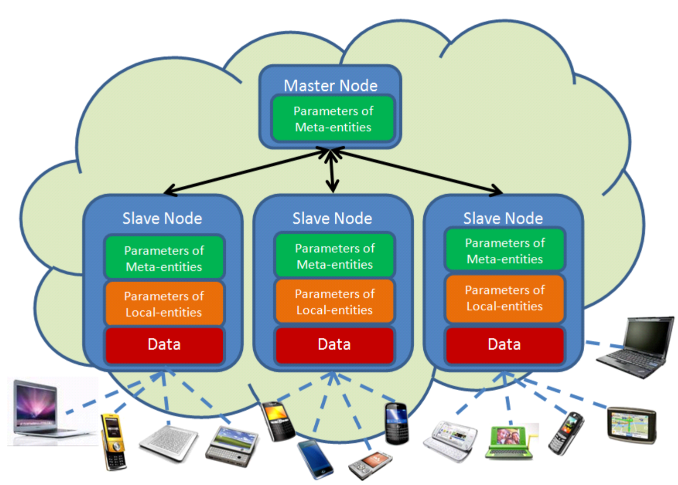

Figure 1. An illustration of a distributed collaborative activity recognition system