Expert Systems with Applications 42 (2015) 3115–3128

Contents lists available at ScienceDirect

Expert Systems with Applications

journal homepage: www.elsevier.com/locate/eswa

Extending knowledge-driven activity models through data-driven learning techniques

Gorka Azkune [a,*], Aitor Almeida [a], Diego López-de-Ipiña [a], Liming Chen [b]

[a] DeustoTech – University of Deusto, Avenida de las Universidades 24, 48007 Bilbao, Spain
[b] School of Computer Science and Informatics – De Montfort University, LE19BH Leicester, United Kingdom

ARTICLE INFO

Article history:
Available online 11 December 2014

Keywords:
Activity recognition
Knowledge-driven
Learning
Activity model

ABSTRACT

Knowledge-driven activity recognition is an emerging and promising research area which has already shown very interesting features and advantages. However, there are also some drawbacks, such as the usage of generic and static activity models. This paper presents an approach to using data-driven techniques to evolve knowledge-driven activity models with a user's behavioral data. The approach includes a novel clustering process where initial incomplete models developed through knowledge engineering are used to detect action clusters which represent activities and aggregate new actions. Based on those action clusters, a learning process is then designed to learn and model varying ways of performing activities in order to acquire complete and specialized activity models. The approach has been tested with real users' inputs, noisy sensors and demanding activity sequences. Initial results have shown that complete and specialized activity models are properly learned with success rates of 100% at the expense of learning some false positive models.

1. Introduction

Human activity recognition has become an important research topic in areas such as pervasive and mobile computing (Choudhury & Consolvo, 2008), ambient assisted living (Philipose & Fishkin, 2004), social robotics (Fong, Nourbakhsh, & Dautenhahn, 2003), surveillance-based security (Fernández-Caballero, 2012) and context-aware computing (Laerhoven & Aidoo, 2001). To perform activity recognition, different kinds of sensors have to be deployed in human-populated environments to monitor inhabitants' behaviors and capture environmental changes generated by human actions. The information provided by those sensors has to be processed through data analysis techniques and/or knowledge representation formalisms to create appropriate activity models and subsequently use them for activity recognition.

The scientific community has developed two main approaches to solve activity recognition, namely the data-driven and knowledge-driven approaches. Data-driven approaches make use of large-scale datasets of sensors to learn activity models using data mining and machine learning techniques. On the other hand, knowledge-driven approaches exploit rich prior knowledge in the domain of interest to build activity models using knowledge engineering and management technologies.

For knowledge-driven activity recognition systems, a widely recognized drawback is that activity models are usually static, i.e. once they have been defined, they cannot be automatically adapted to users' specificities (Chen, Nugent, & Okeyo, 2014). This is a very restrictive limitation, because it is not generally possible to define complete activity models for every user. Domain experts have the necessary knowledge about activities, but this knowledge may not be enough to generate complete models in all the cases. To make knowledge-driven activity recognition systems work in real world applications, activity models have to evolve automatically to adapt to users' varying behaviors. It turns out that model adaptability and evolution are aspects that can be properly addressed by data-driven approaches. Hence, the objective of this paper is to use data-driven techniques to make knowledge-driven activity models evolve automatically based on sensor data generated by specific users.

Let us illustrate our hybrid approach with an example. Fig. 1 shows an initial activity model for the *MakeCoffee* activity, which is composed by the actions *hasCoffee* and *hasContainer*. These are the necessary actions for every person to perform a *MakeCoffee* activity, which highlight the indispensable actions of making a coffee, i.e. the use of coffee and a container to place the coffee in. Nevertheless, some users might add some milk and sugar, others cream etc. The idea of our approach is to create high-level activity models with only these indispensable actions, and then use the

* Corresponding author.
E-mail addresses: gorka.azkune@deusto.es (G. Azkune), aitor.almeida@deusto.es (A. Almeida), dipina@deusto.es (D. López-de-Ipiña), liming.chen@dmu.ac.uk (L. Chen).

http://dx.doi.org/10.1016/j.eswa.2014.11.063
0957-4174/

data generated by a specific user performing the activity to learn those new actions which also configure the personal way of making coffee. In the case of Fig. 1, where the initial activity model will only include coffee and container, the system would learn that *MakeCoffee* is performed in two ways by the user: in the first one, the user adds milk (*hasMilk*) and sugar (*hasFlavor*), while in the second one only sugar is added (*hasFlavor*). Hence two specialized and complete activity models of *MakeCoffee* can be learned. This way, experts, initially, only have to provide incomplete activity models with necessary actions. Afterwards the learning system can analyze a user's behavioral data and learn the specialized and complete models to enrich the knowledge-base, thus improving initial activity models.

Running the proposed learning process periodically with new data generated by a concrete user, a dynamic activity modeling system is achieved. As a user evolves regarding the way (s)he performs certain activities, the learning approach learns new versions of the initial activity models. Hence, activity models can be adapted to users' varying behaviors.

The scientific contributions presented in this paper are:

1. A two-step activity clustering algorithm which uses initial incomplete activity models and context knowledge to recognize action clusters that form an activity and aggregate new actions.
2. A learning algorithm that uses action clusters to learn specialized and complete activity models for every defined activity.

The paper is structured as follows: Section 2 presents the related work. Section 3 introduces the basic theory of ontology-based activity modeling, which is the base of the work presented in this paper. Section 4 describes in detail the proposed approach to learn specialized and complete activity models, followed by the evaluation and results obtained in Section 5. Section 6 discusses the results and the paper concludes in Section 7 with conclusions and future work.

2. Related work

Human activity recognition can be classified into two categories in terms of sensors used for activity monitoring: vision-based and sensor-based activity recognition (Chen, Hoey, Nugent, Cook, & Yu, 2012a). The first is based on the use of visual sensing facilities such as video cameras to monitor an actor's behavior and environmental changes. Vision-based activity recognition has been a research focus for a long period of time, but it is out of the scope of this paper. For a detailed review of vision-based approaches see (Weinland, Ronfard, & Boyer, 2011).

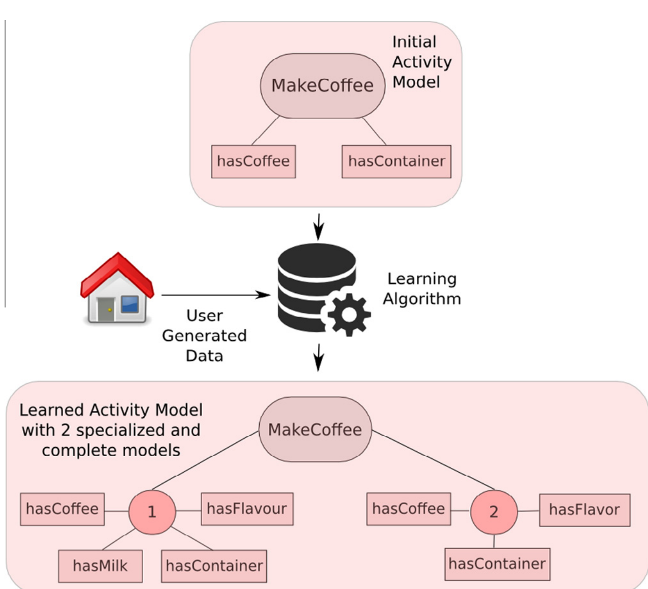

Fig. 1. Illustrative example of the objective of the paper: using the initial incomplete model for MakeCoffee and user generated data, the learning algorithm learns 2 specialized and complete models.

The work presented in this paper lies in the sensor-based activity recognition category, which is based on the use of emerging sensor network technologies for activity monitoring. The principal advantage of sensor-based approaches over vision-based ones are related to privacy and ethics (Yilmaz, Javed, & Shah, 2006) as cameras are generally perceived as recording devices. The generated sensor data from sensor-based monitoring are mainly time series of state changes and/or various parameter values that are usually processed through data fusion, probabilistic or statistical analysis methods and formal knowledge technologies for activity recognition. There are two main approaches for sensor-based activity recognition in the literature: data-driven and knowledge-driven approaches. An exhaustive review can be found in Chen et al. (2012a).

The idea behind data-driven approaches is to use data mining and machine learning techniques to learn activity models. In terms of the modeling approach, two main categories can be found: the generative and the discriminative approaches. The generative approach attempts to build a complete description of the input or data space, usually with a probabilistic model. The simplest possible generative approach is the Naïve Bayes classifier, which has been used with promising results for activity recognition (Bao & Intille, 2004; Brdiczka, Reignier, & Crowley, 2007; Cook & Schmitter-Edgecombe, 2009; Kasteren & Krose, 2007; Maurer & Rowe, 2006; Tapia, Intille, & Larson, 2004). More complex approaches rely on Hidden Markov Models (Galata, Johnson, & Hogg, 1999; Moeslund, Hilton, & Krüger, 2006) or Dynamic Bayesian Networks (Brand, Oliver, & Pentland, 1997; Oliver, Garg, & Horvitz, 2004). In contrast, the discriminative approach only models the mapping from inputs (data) to outputs (activity labels). Discriminative approaches include many supervised learning approaches, such as conditional random fields (Vail, Veloso, & Lafferty, 2007), linear or non-linear discriminative learning (e.g. support vector machines (Brdiczka, 2009)) and online (or incremental) classifiers (Ordóñez, Iglesias, & Toledo, 2013). Finally, there are several approaches that cannot be clearly classified into discriminative or generative categories, but rather use a combination of both (Lester, Choudhury, & Kern, 2005; Omar, Sinn, & Truszkowski, 2010; Pentney, Philipose, & Bilmes, 2008).

One of the problems of data-driven approaches is the need of labeled activity data bases. Rashidi and Cook try to overcome this problem in Rashidi and Cook (2011b, 2013). They use a non-labeled data base, where they extract activity clusters using unsupervised learning techniques. Those clusters are used to train a boosted Hidden Markov Model, which is shown to be able to recognize several activities. Even though the presented approaches overcome the problem of depending on labeled data bases, they still suffer from some other typical problems of data-driven techniques, namely: (i) the *cold-start* problem, because data has to be collected in order to obtain activity models and train the recognizer, and (ii) activity model generalization, since the learned activity models are personal and there are no mechanisms to generalize the extracted knowledge to other users.

On the other hand, transfer learning is being used to make data-driven activity models reusable. Although applying transfer learning techniques to activity recognition is a recent approach, a notable number of research has already been published. An excellent survey on transfer learning for activity recognition has been contributed by Cook, Feuz, and Krishnan (2013). The main idea of