International Journal of Approximate Reasoning 61 (2015) 16–32

Contents lists available at ScienceDirect

International Journal of Approximate Reasoning

www.elsevier.com/locate/ijar

Marginal filtering in large state spaces

Martin Nyolt *, Frank Krüger, Kristina Yordanova, Albert Hein, Thomas Kirste

University of Rostock, Albert-Einstein-Straße 22, 18051 Rostock, Germany

ARTICLE INFO

Article history:
Received 10 July 2014
Received in revised form 2 April 2015
Accepted 8 April 2015
Available online 15 April 2015

Keywords:
Activity recognition
Plan recognition
Particle filter
Bayesian inference
Symbolic models
Marginal filter

ABSTRACT

Recognising everyday activities including information about the context requires to handle large state spaces. The usage of wearable sensors like six degree of freedom accelerometers increases complexity even more. Common approaches are unable to maintain an accurate belief state within such complex domains. We show how marginal filtering can overcome limitations of standard particle filtering and efficiently infer the context of actions. Symbolic models of human behaviour are used to recognise activities in two different settings with different state space sizes. Based on these scenarios we compare the marginal filter to the standard particle filter. An evaluation shows that the marginal filter performs comparably in small state spaces but outperforms the particle filter in large state spaces.

1. Introduction

Recognising activities from noisy and ambiguous sensor data is an important but also challenging task. Knowledge about the activity a person is executing is a necessary prerequisite for deploying assistant technologies, for instance in the context of Ambient Assisted Living [19,30]. In addition, the context of an activity like the location of used objects is of interest in order to provide assistance. Identifying the actual action sequence from a set of plans given a set of observations is known as plan recognition [36]. While early approaches use plan libraries that list all possible plans, the current trend is to generate them from symbolic descriptions [17,18,23,29,33].

According to Denney [8], two distinct styles of symbolic descriptions exist. Algebraic descriptions model the system as combinations of actions. The system's state is defined by equivalence relations between action sequences and not by a collection of state variables. Model-based descriptions, in contrast, model actions in terms of changes to the system state (i.e. context or environmental state). As additional advantage this also allows to employ traditional planning techniques, for instance to assist the user in achieving a certain goal and propose (or automatically execute) further actions. Model-based approaches support complex behaviour including fine-grained activities in rich scenarios that were not possible before. However, by developing complex models with detailed activities and context information, the state space increases by orders of magnitude, imposing challenges to the inference. In the rest of this paper, we use the term *symbolic model* to refer to model-based descriptions.

Inference in symbolic models can also be thought of as program execution. In the sense of Schwering et al. [34], finding the action sequence executed means to determine the simulation that matches best the observation. A Dynamic Bayesian Network (DBN) is usually employed to cope with uncertainties, but other representations such as Markov Logic Networks [33] or possibility theory [32] have also been adopted. One challenge while handling large state spaces is maintaining

* Corresponding author.
E-mail address: martin.nyolt@uni-rostock.de (M. Nyolt).

http://dx.doi.org/10.1016/j.ijar.2015.04.003

an accurate belief state. As exact inference is not tractable in sufficiently complex, non-linear, non-Gaussian models [6], only a limited number of possible states are usually represented. Particle filters and other Monte Carlo methods are common techniques for approximative inference in models based on DBNs. In addition, contradictory sensor observations might reject correct hypotheses, which results in the inability of further state tracking due to resampling. We empirically show that particle filters (even using the optimal proposal function [12]) have difficulties to maintain the belief state of complex symbolic models and discuss the drawbacks in detail. We present the marginal filter to overcome these problems by efficiently representing more states, thus maintaining a more accurate belief state.

A typical scenario for fine-grained activity recognition is preparing and consuming a meal in a kitchen environment [31,37]. We instructed seven persons to cook carrots while observing their body motions with six degree of freedom (6DOF) inertial measurement units (IMUs). The experiment included the tasks setting up the table, eating the meal, cleaning up, and washing the dishes. To recognise fine-grained activities in this scenario, we set up a detailed symbolic action model comprised of 99 actions and 60 state features describing properties of eleven involved objects and five possible locations.

An evaluation based on two scenarios is presented. We first show that the proposed method yields similar results as the particle filter based approach when small belief states have to be handled. A meeting scenario [23] with about 70 000 states serves as this baseline model. In the second part we use the kitchen activities model, which has about 146 million states, to compare both marginal filtering and particle filtering. We evaluate the filters with a simulated dataset and real IMU sensor data. This evaluation shows that the particle filter is unable to handle the kitchen activities model sufficiently, while the marginal filter reaches 80% accuracy. We present empirical results on recognising human activities within a kitchen using wearable sensors only.

The remainder of this paper is structured as follows. We first describe the objective of our study and illustrate the inference methods used in Section 2. The results are presented in Section 3 and discussed in Section 4, including an analysis why the particle filter performs poorly in our setting. Section 5 gives an overview of recent work on plan recognition using symbolic modelling of human behaviour (a general overview of latest activity recognition results has been compiled by Chen et al. [5]). Section 6 concludes the work.

2. Methods and experimental setup

For activity recognition, we use symbolic action models of human behaviour that will be translated into probabilistic models. The current activity (and the context) is then inferred by using Bayesian filtering. Recently, Krüger et al. [23] showed that Bayesian filtering can be used to recognise activities from noisy sensor data based on symbolic action models. They showed that particle filters can successfully incorporate durative actions and goal-directed behaviour. Both features are important in our settings. Participants work towards a defined goal (for example having eaten carrots), and successfully recognising real-life activities requires to explicitly model action durations [20].

Here, we present the filtering methods and the experimental settings we have used in our study. First, we present our modelling approach. Then, we describe two variants of the particle filter, which is a standard approach for approximate state estimation in complex systems. Subsequently, the marginal filter is introduced in detail. Finally, we present the experimental settings we used to compare the performance of both inference algorithms.

2.1. Modelling approach

We use a generative modelling approach. As usual, our models are split into a system model and the observation model. The system model is a symbolic action model based on common sense and specified only by prior knowledge [38]. The observation model can either be specified manually by prior knowledge or can be learned by a few samples. Our underlying rationale is manifold:

- Using a generative approach allows not only *recognition*, but also *prediction*. This is especially important for (dynamic) environments where sensors may be temporarily unavailable.
- Building on symbolic descriptions allows to apply planning, and thus execute assistive actions based on the current situation. When the same model is used for recognition and planning, there is no need to implement two models.
- Separating the system model from the observation model makes it easy to adapt to changing environments. Often only the observation model needs to be learned from few samples.
- The common-sense system model is flexible, extendible, and re-usable.
- Our experience shows that supervised learning approaches and discriminative models, which classify the activity directly from the observation data (such as QDA, HMMs, SVMs, and artificial neural networks) achieve inferior results on complex activity structure [22].

The symbolic action models are based on actions described by preconditions and effects. Preconditions and effects are formulas over state predicates, and state predicates model the current context. The syntax is based on PDDL [15], an excerpt of an action can be found in Listing 1. This symbolic description is then compiled into a probabilistic model as introduced by Krüger et al. [23]. This model can be represented as the DBN in Fig. 1.