International Journal of Approximate Reasoning 54 (2013) 323–342

Contents lists available at SciVerse ScienceDirect

International Journal of Approximate Reasoning

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

Bayesian networks for mathematical models: Techniques for automatic construction and efficient inference

Catherine G. Enright [a], Michael G. Madden [a,*], Niall Madden [b]

[a] College of Engineering and Informatics, National University of Ireland, Galway, Ireland
[b] School of Mathematics, Statistics and Applied Mathematics, National University of Ireland, Galway, Ireland

ARTICLE INFO

Article history:
Received 9 February 2012
Received in revised form 19 October 2012
Accepted 22 October 2012
Available online 12 November 2012

Keywords:
Ordinary differential equations
Dynamic Bayesian networks
Adaptive time stepping
Particle filtering

ABSTRACT

Expert knowledge in the form of mathematical models can be considered sufficient statistics of all prior experimentation in the domain, embodying generic or abstract knowledge of it. When used in a probabilistic framework, such models provide a sound foundation for data mining, inference, and decision making under uncertainty.

We describe a methodology for encapsulating knowledge in the form of ordinary differential equations (ODEs) in dynamic Bayesian networks (DBNs). The resulting DBN framework can handle both data and model uncertainty in a principled manner, can be used for temporal data mining with noisy and missing data, and can be used to re-estimate model parameters automatically using data streams.

We propose an alternative to the fixed time step inference used in standard DBNs. In our algorithm, the DBN automatically adapts the time step lengths to suit the dynamics in each step. The resulting system allows us to efficiently infer probable values of hidden variables using multiple time series of evidence, some of which may be sparse, noisy or incomplete.

We evaluate our approach with a DBN based on a variant of the van der Pol oscillator, and demonstrate an example where it gives more accurate results than the standard approach, but using only one tenth the number of time steps.

We also apply our approach to a real-world example in critical care medicine. By incorporating knowledge in the form of an existing ODE model, we have built a DBN framework for efficiently predicting individualised patient responses using the available bedside and lab data.

1. Introduction

In many domains, for example, engineering and medicine, expert knowledge in the form of mathematical models is readily available. We assert that knowledge in such a form can be considered sufficient statistics of all prior experimentation in the domain, embodying generic or abstract knowledge of it. When used in a probabilistic framework, such models provide a sound foundation for data mining, inference, and decision making under uncertainty.

Ordinary Differential Equation (ODE) models are generally available in mathematical, engineering and biological textbooks and research publications, and typically describe general population-level behaviours. In order to describe individuals, model parameters must be re-calibrated using observations of the individual. However, in most real-life situations, these observations contain uncertainty. They can be subject to measurement error or simple transcription errors. Data may be missing or observations may be sparse relative to the dynamics of the underlying system thus making it difficult to individualise the parameters.

A key contribution of this paper is a methodology for encapsulating existing knowledge in the form of ordinary differential equations in dynamic Bayesian networks (DBNs). This is important because the proposed DBN framework can handle both

* Corresponding author.
E-mail addresses: cathenright@gmail.com (C.G. Enright), michael.madden@nuigalway.ie (M.G. Madden), niall.madden@nuigalway.ie (N. Madden).

0888-613X/$ - see front matter © 2012 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.ijar.2012.10.004

data and model uncertainty in a principled manner, can be used for temporal data mining with noisy and missing data, and can be used to re-estimate model parameters automatically using data streams. From the DBN perspective, we can significantly reduce the knowledge engineering effort involved in building DBNs by basing them on readily available ODE models.

A second contribution of this paper is a new adaptive-time particle filtering algorithm for performing inference on these DBNs. The importance of this contribution lies in the improved efficiency of the inference task, especially where the dynamics of the system are not uniform over time or the underlying ODEs are "stiff".

In Section 2 we discuss related research and in Section 3 we provide a brief introduction to DBNs. In Section 4 we propose a direct automatic mapping from the ODEs to a DBN by incorporating a first order Euler approximation. In Section 5 we show how higher order solvers can be used, and discuss if and when this may be more appropriate. We examine the options available for "stiff" systems. To generate stable solutions to such problems, one must employ either an *implicit* scheme and so solve a set of nonlinear algebraic equations at each time step, or an *explicit* scheme with very small time steps. The former greatly increases the cost of implementation, since many iterations of the nonlinear solver would be required at each step. On the other hand, using very small time steps when evidence is sparse is also wasteful. A natural compromise is to use an explicit method with automatic stepsize control. This forms the basis of the new adaptive-time particle filtering algorithm introduced in Section 7.

A standard assumption when performing inference in DBNs is that time steps are fixed. Generally, the time step chosen is small enough to capture the dynamics of the most rapidly changing variable. This can result in DBNs having a natural time step that is very short, leading to inefficient inference; this is particularly an issue for DBNs derived from ODEs and for systems where the dynamics are not uniform over time.

Therefore we propose a new adaptive-time particle filtering algorithm as an alternative to the standard fixed time step particle filtering. In adaptive-time particle filtering, the DBN automatically adapts the time step lengths to suit the dynamics in each step. The resulting system allows us to efficiently infer probable values of hidden variables using multiple time series of evidence, some of which may be sparse, noisy or incomplete.

In Section 8 we evaluate our approach with a DBN based on a variant of the van der Pol oscillator with parameters chosen to create a stiff problem. We show how the DBN framework can be used to track the dynamic system and re-estimate model parameters using observations from the true solution. We also demonstrate an example where adapting the time step gives more accurate results than the standard fixed time step approach, but using only one tenth the number of time steps (see Section 8.4).

Finally in Section 9 we apply our approach to a real-world example in critical care medicine. Insulin is prescribed to critically ill patients to regulate glycaemia, i.e., plasma glucose levels. Determining the correct dosage is difficult. We incorporate an existing system of ODEs into a DBN in order to build a system to assist clinicians prescribe the correct insulin dosage. Using only the sparse noisy data streams available at the bedside, the DBN framework efficiently predicts glucose levels, and the predictions are more accurate than an existing methodology.

In this paper we apply our techniques to the van der Pol model, because of its applicability to domains as diverse as electrical circuits, seismology [1] and biology [2, Chapter 6]. We also apply our techniques to the medical domain by modelling glycaemia. However, the methods proposed here are relevant to many engineering and science domains where noisy temporal data must be analysed and mathematical models are readily available.

2. Related research

In this section, we discuss research relating to the two main contributions of this paper: a methodology for incorporating ODE models in DBNs, and a new particle filtering algorithm that allows for adaptive time steps.

Much work has previously been carried out to represent dynamic systems in both Bayesian Networks and DBNs. Bellazzi et al. [3] provide a good comparison of some of these methods. While some focus on predicting individual model parameters which are then used off-line [4], others discretise the state-space [5] and so do not explicitly incorporate the model equations. Voortman et al. [6] propose building causal graphs from time-series data and exploiting the ODEs to impose constraints on the model structure.

With an approach that is somewhat similar to ours, Evers and Lucas [7] recently proposed constructing DBNs for Linear Dynamic Systems. They too recognise that using existing models can significantly reduce the knowledge engineering effort required when building DBNs. Such problems are amenable to efficient simulation, since the exact solution is available in a closed form. In contrast, we are interested in non-linear systems which typically cannot be solved exactly, and so must be treated using numerical solvers.

Andersen and Højbjerre [8] proposed an approach whereby they first derive a system of stochastic differential equations (SDEs) from the ODEs and then encode these equations in a DBN. This bears some similarity to the ideas we consider. In the SDE formulation, noise appears explicitly as a term in the equation. The solution is understood to be itself a stochastic process. To simulate such solutions numerically, one may use methods such as the Euler–Maruyama approach (see, e.g., [9]) to generate approximate solutions for a given set of random walks representing the Wiener processes. Our approach is a more direct one: the model is given in the form of deterministic ODEs and mapped directly to the DBN which then incorporates the effects of noise, and generates solutions using a numerical technique such as the (standard) Euler approach. This does