International Journal of Approximate Reasoning 80 (2017) 137–173

Contents lists available at ScienceDirect

International Journal of Approximate Reasoning

www.elsevier.com/locate/ijar

Computing lower and upper expected first-passage and return times in imprecise birth–death chains

Stavros Lopatatzidis *, Jasper De Bock, Gert de Cooman

Ghent University, Data Science Lab, Belgium

ARTICLE INFO

Article history:
Received 13 April 2016
Received in revised form 26 August 2016
Accepted 26 August 2016
Available online 5 September 2016

Keywords:
Birth–death chain
First-passage time
Return time
Imprecise probability
Lower expectation
Submartingale

ABSTRACT

We provide simple methods for computing exact bounds on expected first-passage and return times in finite-state birth–death chains, when the transition probabilities are imprecise, in the sense that they are only known to belong to convex closed sets of probability mass functions. In order to do that, we model these so-called imprecise birth–death chains as a special type of time-homogeneous imprecise Markov chain, and use the theory of sub- and supermartingales to define global lower and upper expectation operators for them. By exploiting the properties of these operators, we construct a simple system of non-linear equations that can be used to efficiently compute exact lower and upper bounds for any expected first-passage or return time. We also discuss two special cases: a precise birth–death chain, and an imprecise birth–death chain for which the transition probabilities belong to linear-vacuous mixtures. In both cases, our methods simplify even more. We end the paper with some numerical examples.

1. Introduction

Birth–death chains [27, Section 9.4] are a special type of time-homogeneous Markov chains, where transitions from a given state are only possible to that state or to adjacent ones. They are used in various scientific fields, including evolutionary biology [1, Chapter 3] and queueing theory [16]. We consider the generalised case of an imprecise birth–death chain, which, basically, is a birth–death chain whose transition probabilities are not specified exactly, but are only known to belong to some given closed convex set of probability mass functions. This may be the case because the transition probabilities are based on partial expert knowledge or limited data, or for the purposes of conducting a sensitivity analysis. Similar models have already been considered in Reference [6], which presented results on limiting conditional distributions for imprecise birth–death chains with one absorbing state. Imprecise birth–death chains are themselves a special case of so-called (time-homogeneous) imprecise Markov chains, which were studied in—amongst others—References [11,18,25].

This paper focuses on—upward and downward—first-passage and return times.[1] For precise birth–death chains, these have for example been studied in Reference [23]. For the more general case of imprecise birth–death chains, we are not aware of any previous discussion in the literature. Our most important contributions are simple methods for computing exact lower and upper bounds for any expected first-passage and return time in finite-state imprecise birth–death chains. We also consider two special cases: precise birth–death chains and imprecise birth–death chains whose local models are linear-vacuous mixtures. In those cases, our methods lead to closed-form expressions.

* Corresponding author.
E-mail addresses: stavros.lopatatzidis@UGent.be (S. Lopatatzidis), jasper.debock@UGent.be (J. De Bock), gert.decooman@UGent.be (G. de Cooman).

[1] These are often called recurrence times as well.

http://dx.doi.org/10.1016/j.ijar.2016.08.008

We start in Section 2 with a brief introduction to the framework of imprecise probabilities, restricting our attention to the simple case of finite state spaces. We introduce credal sets as closed and convex sets of probability mass functions, provide a number of examples, and establish a connection with lower and upper probabilities. This section also introduces the concept of a lower and upper expectation operator, and explains that these operators are mathematically equivalent to credal sets.

Next, we explain what a birth–death chain is in Section 3, and then introduce an imprecise version in Section 4. As we will see, an imprecise birth–death chain is just a special type of a time-homogeneous imprecise Markov chain [11], which, basically, can be regarded as birth–death chain whose local models are credal sets. However, as we will explain, this should not be taken to mean that an imprecise birth–death chain is a collection of birth–death chains. Instead, an imprecise birth–death chain can be regarded as a set of probability trees, only some of which are birth–death chains.

Section 5 defines the global lower and upper expectations that correspond to these imprecise birth–death chains, using the notions of sub- and supermartingales. For real-valued functions that only depend on a finite number of variables, these lower and upper expectations are just the minimum and maximum expectations of this function, with respect to the probability trees in the imprecise birth–death chain. For more general—possibly extended real-valued—functions on the infinite sequence of all variables, the expressions become more intricate. We also recall some convenient properties of the definitions that we adopt, including a global Markov property and a generalised version of the law of iterated expectations.

With all this machinery in place, Section 6 then finally introduces our main topic of interest: return and—upward and downward—first-passage times, and in particular, their lower and upper expectations. As we will see, these lower and upper expected first-passage and return times satisfy a relatively simple system of non-linear equations. The next three sections of the paper are concerned with solving this system, and by doing so, we obtain a simple method for computing the lower and upper expected first-passage and return times that we are interested in. In Sections 7 and 8, we develop recursive methods for computing lower and upper expected upward and downward first-passage times, respectively, and Section 9 explains how these results can be used to compute lower and upper expected return times.

The next two sections are concerned with special cases. Section 10 considers the special case of precise birth–death chains and establishes closed-form expressions for their expected first-passage and return times. It also proves that even though an imprecise birth–death chain is more than just a collection of precise birth–death chains, nevertheless, the lower and upper value of an expected first-passage or return time can always be obtained by a precise birth–death chain. Section 11 discusses the special case where the local models are linear-vacuous mixtures, that is, when they are ϵ-contaminated. Here too, we are able to obtain closed-form expressions.

Section 12 presents some numerical results. We apply our methods for lower and upper first-passage times to a general example, and illustrate our methods for lower and upper return times on an example with local models that are linear-vacuous mixtures.

Finally, Section 13 briefly concludes the paper and mentions some possible avenues for future research. The proofs of our main results are gathered in Appendix A.

2. A brief introduction to imprecise probabilities

We start by presenting some basic concepts from the theory of imprecise probabilities. For more information, we refer the reader to Walley's book [26], and to more recent textbooks [2,24].

Consider a variable X that takes values in some non-empty finite set $\mathscr{X}$. A common approach to describe a subject's uncertainty about the actual value of X is then to consider a probability mass function p on $\mathscr{X}$, that is, an element of the set

$$\Sigma_{\mathscr{X}} := \left\{ p \in \mathbb{R}^{\mathscr{X}} : \sum_{x \in \mathscr{X}} p(x) = 1 \text{ and } (\forall x \in \mathscr{X})\, p(x) \geq 0 \right\}.$$

For any real-valued function f on $\mathscr{X}$—also called a *gamble*—the corresponding expectation of f is then given by

$$E_p(f) := \sum_{x \in \mathscr{X}} p(x) f(x).$$

If we now let $\mathscr{G}(\mathscr{X})$ be the set of all gambles, then the expectation operator $E_p : \mathscr{G}(\mathscr{X}) \to \mathbb{R}$ can be regarded as an alternative, equivalent representation for p. Indeed, E_p can clearly be inferred from p and, conversely, if we know E_p, then for all $x \in \mathscr{X}$, $p(x)$ is equal to the expectation $E_p(\mathbb{I}_x)$ of the indicator $\mathbb{I}_x \in \mathscr{G}(\mathscr{X})$ of x, as defined by

$$\mathbb{I}_x(y) := \begin{cases} 1 & \text{if } y = x \\ 0 & \text{if } y \neq x \end{cases} \text{ for all } y \in \mathscr{X}.$$

The success of this approach depends crucially on the assumption that our uncertainty about X can be described by a probability mass function p, and furthermore requires that p is specified precisely. However, in practice, eliciting such a probability function can be difficult, especially if it is based on—possibly disagreeing—expert opinions, or when it has to be learned from small amounts of data. Whenever this is the case, the theory of imprecise probabilities [2,24,26] does not