Computing Expected Hitting Times for Imprecise Markov Chains


Thomas Krak and Jasper De Bock

Foundations Lab for Imprecise Probabilities, Ghent University, Belgium

Abstract

In recent work (Krak *et al.* 2019) we proved that the minimal non-negative solution $\underline{h}$ to the non-linear system

$$\underline{h} = \mathbb{I}_{A^c} + \mathbb{I}_{A^c} \cdot \underline{T}\,\underline{h},$$

corresponds to the vector of **lower expected hitting times** of a subset A of the set $\mathscr{X}$ of possible states of any type of imprecise Markov chain with lower transition operator $\underline{T}$ corresponding to a set $\mathscr{T}$ of transition matrices that satisfies some technical closure properties.

We here present two methods to numerically find the solution of this system; one based on the method of **iterated lower expectations**, and one **novel algorithm**. We present an initial comparison of their efficiency, and show that the novel algorithm appears to substantially outperform the other method.

Expected hitting times, detailed

Expected hitting times for a stochastic process P are the **conditional expectations** $\mathbb{E}_P[H_A \,|\, X_0]$ of a function H_A describing the number of steps until A is visited. To formalise this, note that P is a probability measure on a measurable space $(\Omega, \mathscr{F})$, where the sample space Ω is the set of all paths $\omega : \mathbb{N}_0 \to \mathscr{X}$ that the process can take. The function H_A is then defined, for all $\omega \in \Omega$, as

$$H_A(\omega) := \inf\{t \in \mathbb{N}_0 : \omega(t) \in A\}.$$

A convenient representation of H_A is obtained through the **finite horizon approximations** $H_A^{(n)}$, which are defined, for all $n \in \mathbb{N}_0$ and all $\omega \in \Omega$, as

$$H_A^{(n)}(\omega) := \begin{cases} H_A(\omega) & \text{if } H_A(\omega) \leq n, \text{ and} \\ n+1 & \text{otherwise.} \end{cases} \tag{1}$$

Then $\lim_{n\to+\infty} H_A^{(n)} = H_A$, and it can be shown that $\lim_{n\to+\infty} \mathbb{E}_P[H_A^{(n)} \,|\, X_0 = x] = \mathbb{E}_P[H_A \,|\, X_0 = x]$ for all $x \in \mathscr{X}$.

Iterated lower expectation

A useful property of the finite horizon approximations $H_A^{(n)}$ is that their lower (or upper) expectations can be computed through **iterated lower (or upper) expectation**. Specifically, it turns out that there is a very **elegant recursive form** that allows us to do this.

To obtain this, define $\underline{h}_A^{(0)} := \mathbb{I}_{A^c}$ and, for all $n \in \mathbb{N}_0$, let

$$\underline{h}_A^{(n+1)} := \mathbb{I}_{A^c} + \mathbb{I}_{A^c} \cdot \underline{T}\,\underline{h}_A^{(n)}. \tag{2}$$

It can then be shown that, for the game-theoretic model whose local models are described by $\underline{T}$, it holds that

$$\underline{\mathbb{E}}^{\mathrm{V}}_{\mathscr{T}}[H_A^{(n)} \,|\, X_0 = x] = \underline{h}_A^{(n)}(x) \quad \text{for all } x \in \mathscr{X} \text{ and } n \in \mathbb{N}_0,$$

and, moreover, when we take limits in n it holds that

$$\begin{aligned} \underline{\mathbb{E}}^{\mathrm{V}}_{\mathscr{T}}[H_A \,|\, X_0 = x] &= \lim_{n\to+\infty} \underline{\mathbb{E}}^{\mathrm{V}}_{\mathscr{T}}[H_A^{(n)} \,|\, X_0 = x] \\ &= \lim_{n\to+\infty} \underline{h}_A^{(n)} =: \underline{h}_A^*. \end{aligned}$$

This result immediately provides us with **a numerical scheme** to compute lower expected hitting times for imprecise Markov chains: **simply iterate the scheme** (2) until some appropriate stopping condition is satisfied.

Recall that the (lower and upper) expected hitting times are the same for any type of imprecise Markov chain, so despite being derived for game-theoretic models, this method also **works for any other type** of model. Moreover, this method works to compute the upper expected hitting times, simply replacing the lower transition operator $\underline{T}$ by its conjugate upper transition operator $\overline{T}$.

A reachability condition

Our novel method requires that the imprecise Markov chain has a lower transition operator $\underline{T}$ such that, for all $x \in A^c$, there is some $n_x \in \mathbb{N}$ for which $\left[\underline{T}^{n_x}\mathbb{I}_A\right](x) > 0$.

This says that $\underline{P}(X_{n_x} \in A \,|\, X_0 = x) > 0$, or, in words, every process $P \in \mathscr{P}_{\mathscr{T}}$ must be able—have positive probability—to move from any state $x \in A^c$ to some element of A in a finite number of steps n_x.

Novel algorithm

The novel method that we present here provides an alternative to computing lower and upper expected hitting times through the method of **iterated lower** (or upper) **expectations**. However, it requires a **reachability condition** of the states in A, which makes the method slightly more restrictive. Nevertheless, a preliminary comparison of the efficiency of these methods indicates that **this novel method is preferable** in those cases where it can be used.

We recall that the closure conditions on $\mathscr{T}$ guarantee that, for any $f : \mathscr{X} \to \mathbb{R}$, we can write $\underline{T}f = \min_{T\in\mathscr{T}} Tf$. In other words, there is a minimising vector Tf, implying that the associated $\arg\min$ is non-empty. With this observation, the algorithm can be described as follows:

Initialise: Pick any $T_{(0)} \in \mathscr{T}$, and let $h_{(0)}(x) := \mathbb{E}_{P_{(0)}}[H_A \,|\, X_0 = x]$ for all $x \in \mathscr{X}$, where $T_{(0)}$ determines the homogeneous Markov chain $P_{(0)} \in \mathscr{P}^{\mathrm{H}}_{\mathscr{T}}$.

Iterate: For all $n \in \mathbb{N}$, let

$$T_{(n)} \in \operatorname*{arg\,min}_{T\in\mathscr{T}} T h_{(n-1)}, \tag{3}$$

and define $h_{(n)}(x) := \mathbb{E}_{P_{(n)}}[H_A \,|\, X_0 = x]$ for all $x \in \mathscr{X}$, where $T_{(n)}$ determines $P_{(n)} \in \mathscr{P}^{\mathrm{H}}_{\mathscr{T}}$.

Convergence: It holds that

$$\lim_{n\to+\infty} h_{(n)}(x) = \underline{\mathbb{E}}^{\mathrm{H}}_{\mathscr{T}}[H_A \,|\, X_0 = x] = \underline{h}_A^*(x) \quad \text{for all } x \in \mathscr{X}.$$

Note that this result is stated for the imprecise Markov chain represented as a set of homogeneous Markov chains. However, because the lower and upper expected hitting times are the same for any type of imprecise Markov chain, this method also **works for any other type** of model. Moreover, this method works to compute the upper expected hitting times, simply replacing the selection in (3) to instead be from an $\arg\max$.

Some first properties

Monotonically improving inner approximation: At each $n \in \mathbb{N}_0$ the vector $h_{(n)}$ represents the hitting times of a precise homogeneous Markov chain $P_{(n)} \in \mathscr{P}^{\mathrm{H}}_{\mathscr{T}}$. Hence, each $h_{(n)}$ is an upper bound on the actual lower expected hitting time, i.e., **an inner approximation** of the quantity of interest. Moreover, the sequence $h_{(n)}$ is **monotonically improving** (i.e. decreasing towards the lower expected hitting times). The same holds, *mutatis mutandis*, when computing upper expected hitting times.

Simple stopping criterion: If for some $n \in \mathbb{N}_0$ it holds that $h_{(n)} = h_{(n+1)}$, then it can be shown that $h_{(n)} = \underline{h}_A^*$, i.e. that the algorithm has converged after n iterations.

Convergence in finite number of iterations: If the set $\mathscr{T}$ has only a finite number K of extreme points, a minor specialisation of the algorithm is **guaranteed to converge after a finite number of iterations**. In particular, we then require that in (3), we take $T_{(n)}$ to be an extreme point of $\mathscr{T}$; note that this is always possible. If we do this, then it can be shown that $h_{(n)} = h_{(n+1)} = \underline{h}_A^*$ for some $n \leq K$. Moreover, empirically we observe that this tends to happen for $n << K$.

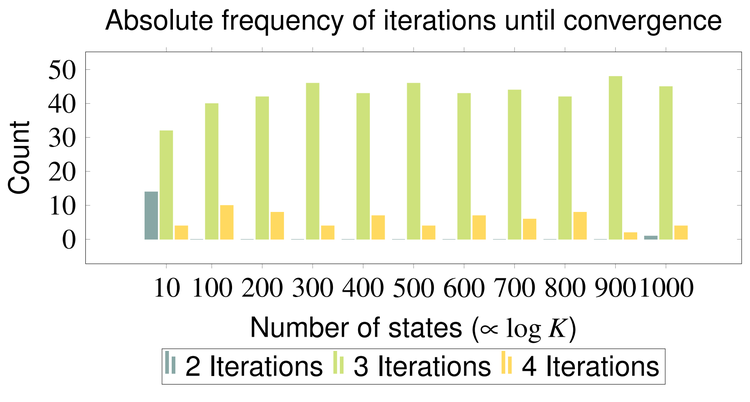

Computing precise hitting times

The computational method that we present here requires an **auxiliary method** to compute the expected hitting times for any (precise) homogeneous Markov chain P with transition matrix $T \in \mathscr{T}$. There are various methods in the literature for doing this.

Here, computing the expected hitting times reduces to **solving a linear system** due to the required **reachability condition**. In particular, it follows immediately from the definition that $\mathbb{E}_P[H_A \,|\, X_0 = x] = 0$ for all $x \in A$. For the states in A^c, we define the $|A^c| \times |A^c|$ matrix F, as

$$F(x,y) := T(x,y) \quad \text{for all } x, y \in A^c.$$

It can be shown that the matrix $(I - F)$ is invertible, where I denotes the identity matrix. Moreover, for any $x \in A^c$, it holds that

$$\mathbb{E}_P[H_A \,|\, X_0 = x] = \sum_{y\in A^c} (I-F)^{-1}(x,y).$$

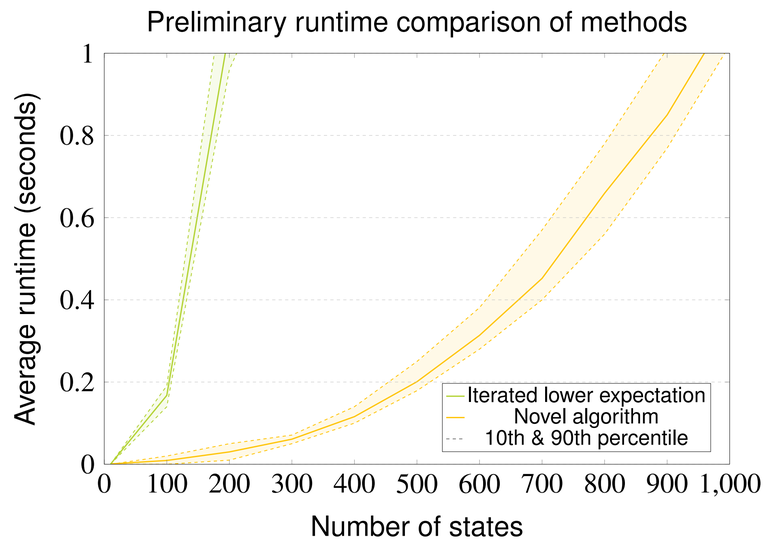

Relation to MDPs

It is well known that there is a connection between the theory of imprecise Markov chains and the theory of **Markov Decision Processes** (MDPs). An important distinction between these theories is the semantics: with imprecise Markov chains, we optimise over a set of precise inferences with the aim to obtain robust bounds on quantities of interest. In contrast, MDPs are typically formulated to obtain "policies" $T \in \mathscr{T}$ that optimise some operational reward function. In other words, the possible variation of $T \in \mathscr{T}$ is either a quantification of uncertainty about some "true" T, for imprecise Markov chains; or, in an MDP context, a controllable choice from a set of feasible parameters of the system.

In the present context, the connection between these two theories appears to run deeper, and it allows us to place the algorithms presented here in an MDP context. Specifically, the characterising equation of the upper expected hitting times of an imprecise Markov chain is strongly related to the **equation of optimality** for an infinite horizon undiscounted Markov decision process under the reward function $\mathbb{I}_{A^c}$, *viz.* $\overline{h} = \max_{T\in\mathscr{T}} \mathbb{I}_{A^c} + T\overline{h}$.

When possible, this equation for MDPs is typically solved using either the **value iteration** algorithm or the **policy iteration** algorithm. However, for infinite horizon undiscounted MDPs the applicability of these algorithms seems to be somewhat of an edge case in the literature. Nevertheless, remarkably, these algorithms appear **formally equivalent** to the method of **iterated lower expectation**, and our **novel algorithm**, respectively.

We express our sincere gratitude to an anonymous reviewer for pointing out this connection. Further investigation of this connection is a matter of ongoing research.

This work was partially supported by H2020-MSCA-ITN-2016 UTOPIAE, grant agreement 722734.