Computing Expected Hitting Times for Imprecise Markov Chains*

Thomas Krak THOMAS.KRAK@UGENT.BE
Jasper De Bock JASPER.DEBOCK@UGENT.BE
ELIS – FLip, Ghent University, Belgium

An important inference problem in the theory of Markov chains is computing the expected hitting time $\mathbb{E}_P[H_A \,|\, X_0]$ of a subset A of the (non-empty, finite) set of all states $\mathscr{X}$ in which a Markov chain P can be. Intuitively, this inference amounts to computing how many steps H_A it will take—in expectation—before the system visits an element of A for the first time, given that it starts in the state X_0 at time 0. It is well known that any homogeneous Markov chain P can be uniquely characterised—up to its initial distribution $P(X_0)$, but this will be irrelevant here—by a single $|\mathscr{X}| \times |\mathscr{X}|$ matrix T that is row-stochastic, meaning that $\sum_{y\in\mathscr{X}} T(x,y) = 1$ and $T(x,y) \geq 0$ for all $x, y \in \mathscr{X}$. When the matrix T is not known precisely, we can instead use an *imprecise Markov chain* to describe the underlying system of interest, and use this to compute tight bounds on the inferential quantities of interest, such as lower and upper expected hitting times.

In particular, an imprecise Markov chain $\mathscr{P}_{\mathscr{T}}$ is a *set* of homogeneous Markov chains, parameterised by a *set* $\mathscr{T}$ of row-stochastic matrices that satisfies some technical closure properties (see [1] for details). That is, $\mathscr{P}_{\mathscr{T}}$ consists of those homogeneous Markov chains whose characterising matrix T is included in $\mathscr{T}$. For this model, we are interested in computing the lower (resp. upper) expected hitting time $\underline{h}(x) := \inf_{P\in\mathscr{P}_{\mathscr{T}}} \mathbb{E}_P[H_A \,|\, X_0 = x]$ (resp. $\overline{h}(x) := \sup_{P\in\mathscr{P}_{\mathscr{T}}} \mathbb{E}_P[H_A \,|\, X_0 = x]$), for all $x \in \mathscr{X}$. Alternatively, we could also consider imprecise Markov chains that include non-Markovian and non-homogeneous processes, or even so-called game-theoretic imprecise Markov chains, all of which can be parametrised by the same set $\mathscr{T}$. However, as we show in [1], this would lead to the exact same lower and upper expected hitting times $\underline{h}(x)$ and $\overline{h}(x)$. In particular, for a given set $\mathscr{T}$, and for all these different types of imprecise Markov chains, we show in [1] that the vector $\underline{h}$ of lower expected hitting times is the minimal non-negative solution to the non-linear system

$$\underline{h} = \mathbb{I}_{A^c} + \mathbb{I}_{A^c} \cdot \underline{T}\,\underline{h}\,, \tag{1}$$

where $\mathbb{I}_{A^c}$ is the indicator of $A^c := \mathscr{X} \setminus A$, and where $\underline{T}$ is the *lower transition operator* corresponding to $\mathscr{T}$, which is defined for all functions $f : \mathscr{X} \to \mathbb{R}$ and all $x \in \mathscr{X}$ as $\left[\underline{T}\,f\right](x) := \min_{T\in\mathscr{T}} \sum_{y\in\mathscr{X}} T(x,y)f(y)$. Note that, because $\mathscr{X}$ is finite, we can also interpret f as a vector in $\mathbb{R}^{|\mathscr{X}|}$ and simply write this using matrix-vector products as $\underline{T}\,f = \min_{T\in\mathscr{T}} Tf$, where the existence of a minimising vector Tf is guaranteed by the aforementioned closure properties of $\mathscr{T}$.

In this short note, we present a method for solving the system (1) under the additional assumption that, for all $x \in A^c$, there is some $n_x \in \mathbb{N}$ such that $\left[\underline{T}^{n_x}\mathbb{I}_A\right](x) > 0$. Intuitively, this ensures that the solution $\underline{h}$ does not diverge to infinity, which substantially simplifies the analysis.

Our method requires a way to compute the vector of expected hitting times h_P, defined by $h_P(x) := \mathbb{E}_P[H_A \,|\, X_0 = x]$ for all $x \in \mathscr{X}$, for any homogeneous Markov chain $P \in \mathscr{P}_{\mathscr{T}}$ with characterising matrix $T \in \mathscr{T}$. Under the conditions stated above, closed-form solutions for this are available in the literature. In particular, $h_P(x) = 0$ for all $x \in A$ and, if we define the $|A^c| \times |A^c|$ matrix F as $F(x,y) := T(x,y)$ for all $x, y \in A^c$, then $h_P(x) = \sum_{y\in A^c} (I - F)^{-1}(x,y)$ for all $x \in A^c$, with I the identity matrix. Our numerical method can now be described as follows:

1. **Initialize:** Pick any $T_{(0)} \in \mathscr{T}$, and let $h_{(0)} := h_{P_{(0)}}$, where $T_{(0)}$ determines $P_{(0)} \in \mathscr{P}_{\mathscr{T}}$.
2. **Iterate:** For $n \in \mathbb{N}$, let $T_{(n)} \in \arg\min_{T\in\mathscr{T}} Th_{(n-1)}$ and $h_{(n)} := h_{P_{(n)}}$, where $T_{(n)}$ determines $P_{(n)} \in \mathscr{P}_{\mathscr{T}}$.
3. **Convergence:** It holds that $\lim_{n\to+\infty} h_{(n)} = \underline{h}$, where $\underline{h}$ is the minimal non-negative solution to (1).

Finally, if in the second step we replace the selection of $T_{(n)}$ to be from an $\arg\max$, then the method instead converges to the vector $\overline{h}$ of *upper* expected hitting times.

References

[1] Thomas Krak, Natan T'Joens, and Jasper De Bock. Hitting Times and Probabilities for Imprecise Markov Chains. https://arxiv.org/abs/1905.08781, 2019.

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