Imprecise continuous-time Markov chains

Efficient computational methods with guaranteed error bounds

Alexander Erreygers Jasper De Bock

Ghent University

ISIPTA'17

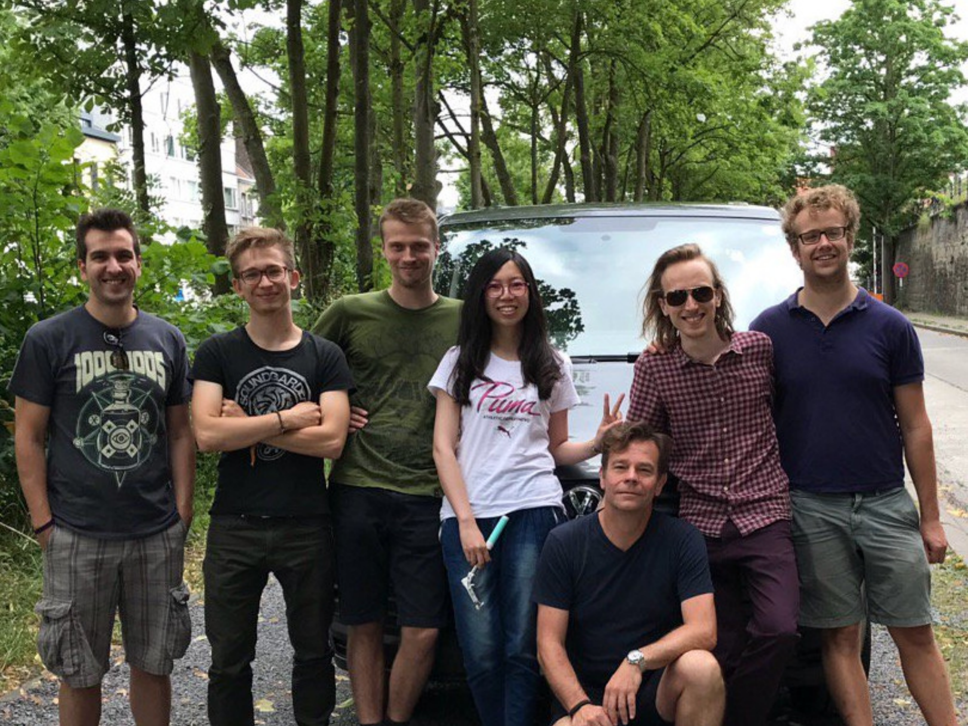
Puma

Imprecise continuous-time Markov chains

Efficient computational methods
with guaranteed error bounds

Precise continuous-time Markov chains

Consider a family of random variables $(X_t)_{t\in\mathbb{R}_{\geq 0}}$ where, for all $t \in \mathbb{R}_{\geq 0}$, X_t takes values in a finite and ordered state space $\mathscr{X}$.
Let $\mathscr{L}(\mathscr{X})$ denote the set of all real-valued functions (vectors) on $\mathscr{X}$.

Precise continuous-time Markov chains

Consider a family of random variables $(X_t)_{t\in\mathbb{R}_{\geq 0}}$ where, for all $t \in \mathbb{R}_{\geq 0}$, X_t takes values in a finite and ordered state space $\mathscr{X}$.
Let $\mathscr{L}(\mathscr{X})$ denote the set of all real-valued functions (vectors) on $\mathscr{X}$.

A stochastic process P is a *precise (continuous-time) Markov chain* (pMC) if

$$\begin{aligned}\mathrm{P}(X_{t+\Delta} = y | X_{t_1} = x_1, \ldots, X_{t_n} = x_n, X_t = x)& \\ = \mathrm{P}(X_{t+\Delta} = y | X_t = x) &=: T_t^{t+\Delta}(x, y).\end{aligned}$$

Precise continuous-time Markov chains

Consider a family of random variables $(X_t)_{t\in\mathbb{R}_{\geq 0}}$ where, for all $t \in \mathbb{R}_{\geq 0}$, X_t takes values in a finite and ordered state space $\mathscr{X}$.
Let $\mathscr{L}(\mathscr{X})$ denote the set of all real-valued functions (vectors) on $\mathscr{X}$.

A stochastic process P is a *precise (continuous-time) Markov chain* (pMC) if

$$\begin{aligned}\mathrm{P}(X_{t+\Delta} = y | X_{t_1} = x_1, \ldots, X_{t_n} = x_n, X_t = x)& \\ = \mathrm{P}(X_{t+\Delta} = y | X_t = x) &=: T_t^{t+\Delta}(x, y).\end{aligned}$$

The *transition matrix* $T_t^{t+\Delta}$ thus defined determines conditional expectations:

$$\begin{aligned}[T_t^{t+\Delta} f](x) &= \mathrm{E}(f(X_{t+\Delta}) | X_t = x) \\ &= \mathrm{E}(f(X_{t+\Delta}) | X_{t_1} = x_1, \ldots, X_{t_n} = x_n, X_t = x).\end{aligned}$$

Precise continuous-time Markov chains

A pMC is called *stationary* if

$$(\forall t, \Delta \in \mathbb{R}_{\geq 0})\; T_t^{t+\Delta} = T_0^{\Delta} =: T_{\Delta}.$$

Precise continuous-time Markov chains

Transition rate matrix

A pMC is called *stationary* if

$$(\forall t, \Delta \in \mathbb{R}_{\geq 0})\; T_t^{t+\Delta} = T_0^{\Delta} =: T_{\Delta}.$$

In this case, there is a unique *transition rate matrix Q*—a matrix with non-negative off-diagonal elements and rows that sum up to zero—such that,

$$(\forall t \in \mathbb{R}_{\geq 0})\; T_t^{t+\delta} = T_\delta \approx (I + \delta Q) \qquad \text{for } \delta \text{ suff. small.}$$

Precise continuous-time Markov chains

Transition rate matrix

A pMC is called *stationary* if

$$(\forall t, \Delta \in \mathbb{R}_{\geq 0})\; T_t^{t+\Delta} = T_0^{\Delta} =: T_{\Delta}.$$

In this case, there is a unique *transition rate matrix* Q—a matrix with non-negative off-diagonal elements and rows that sum up to zero—such that,

$$(\forall t \in \mathbb{R}_{\geq 0})\; T_t^{t+\delta} = T_{\delta} \approx (I + \delta Q) \qquad \text{for } \delta \text{ suff. small.}$$

Similarly, for any non-stationary pMC there is a *time-dependent* transition rate matrix Q_t such that

$$(\forall t \in \mathbb{R}_{\geq 0})\; T_t^{t+\delta} \approx (I + \delta Q_t) \qquad \text{for } \delta \text{ suff. small.}$$

Precise continuous-time Markov chains

Transition rate matrix

A pMC is called *stationary* if

$$(\forall t, \Delta \in \mathbb{R}_{\geq 0})\; T_t^{t+\Delta} = T_0^{\Delta} =: T_{\Delta}.$$

In this case, there is a unique *transition rate matrix* Q—a matrix with non-negative off-diagonal elements and rows that sum up to zero—such that,

$$(\forall t \in \mathbb{R}_{\geq 0})\; T_t^{t+\delta} = T_\delta \approx (I + \delta Q) \qquad \text{for } \delta \text{ suff. small.}$$

Similarly, for any non-stationary pMC there is a *time-dependent* transition rate matrix Q_t such that

$$(\forall t \in \mathbb{R}_{\geq 0})\; T_t^{t+\delta} \approx (I + \delta Q_t) \qquad \text{for } \delta \text{ suff. small.}$$

In practice, a (non-)stationary pMC is characterised by specifying

(i) a state space $\mathscr{X}$,
(ii) an initial distribution $\mathrm{P}(X_0 = x)$, and
(iii) a (time-dependent) transition rate matrix Q (Q_t).

Precise continuous-time Markov chains

Approximation conditional expectations

Combining some properties of stationary pMCs, we find that

$$\begin{aligned}\mathrm{E}(f(X_t)|X_0 = x) = [T_t f](x) &= [T_{\delta_1} \cdots T_{\delta_n} f](x)\\ &\approx [(I + \delta_1 Q) \cdots (I + \delta_n Q) f](x),\end{aligned}$$

where $\delta_{1:n} \coloneqq \{\delta_1, \dots, \delta_n\}$ is a sequence of sufficiently small strictly positive time steps such that $\sum_{i=1}^{n} \delta_i = t$.

Precise continuous-time Markov chains

Approximation conditional expectations

Combining some properties of stationary pMCs, we find that

$$\begin{aligned}\mathrm{E}(f(X_t)|X_0 = x) = [T_t f](x) &= [T_{\delta_1} \cdots T_{\delta_n} f](x) \\ &\approx [(I + \delta_1 Q) \cdots (I + \delta_n Q) f](x),\end{aligned}$$

where $\delta_{1:n} := \{\delta_1, \ldots, \delta_n\}$ is a sequence of sufficiently small strictly positive time steps such that $\sum_{i=1}^{n} \delta_i = t$.

In fact, it is well-known that

$$T_t = e^{tQ} = \lim_{n \to +\infty} \left(I + \frac{t}{n} Q \right)^n,$$

which is the unique time-dependent matrix that satisfies

$$\frac{\mathrm{d}}{\mathrm{d}t} T_t = Q T_t \qquad \text{with initial condition } T_0 = I.$$

Precise continuous-time Markov chains

Approximation conditional expectations

Combining some properties of stationary pMCs, we find that

$$\begin{aligned}\mathrm{E}(f(X_t)|X_0 = x) = [T_t f](x) = [T_{\delta_1} \cdots T_{\delta_n} f](x) \\ \approx [(I + \delta_1 Q) \cdots (I + \delta_n Q) f](x),\end{aligned}$$

where $\delta_{1:n} := \{\delta_1, \dots, \delta_n\}$ is a sequence of sufficiently small strictly positive time steps such that $\sum_{i=1}^{n} \delta_i = t$.

In fact, it is well-known that

$$T_t = e^{tQ} = \lim_{n \to +\infty} \left(I + \frac{t}{n} Q \right)^n,$$

which is the unique time-dependent matrix that satisfies

$$\frac{\mathrm{d}}{\mathrm{d}t} T_t = Q T_t \qquad \text{with initial condition } T_0 = I.$$

In general, obtaining an analytical expression for T_t is infeasible!

Imprecise continuous-time Markov chains

Efficient computational methods
with guaranteed error bounds

Imprecise continuous-time Markov chains

In practice, using a stationary pMC is often not warranted, as

- exactly specifying the transition rate matrix Q is often infeasible, and
- assuming stationarity is not always justified.

Imprecise continuous-time Markov chains

In practice, using a non-stationary pMC is often not warranted, as

- exactly specifying the time-dependent transition rate matrix Q_t is often infeasible, and
- assuming stationarity is not always justified.

Imprecise continuous-time Markov chains

In practice, using a (non-)stationary pMC is often not warranted, as

- exactly specifying the (time-dependent) transition rate matrix Q_t is often infeasible, and
- assuming stationarity is not always justified.

Therefore, we assume that the time-dependent transition rate matrix Q_t of a pMC is only known to be contained in some (non-empty and bounded) *set* of transition rate matrices $\mathscr{Q}$.

Imprecise continuous-time Markov chains

In practice, using a (non-)stationary pMC is often not warranted, as

- exactly specifying the (time-dependent) transition rate matrix Q_t is often infeasible, and
- assuming stationarity is not always justified.

Therefore, we assume that the time-dependent transition rate matrix Q_t of a pMC is only known to be contained in some (non-empty and bounded) *set* of transition rate matrices $\mathscr{Q}$.

Let $\mathbb{P}_{\mathscr{Q}}$ be the *set* of all Markov chains P consistent with $\mathscr{Q}$, in the sense that

$$(\forall t \in \mathbb{R}_{\geq 0})(\exists Q_t \in \mathscr{Q})\, T_t^{t+\delta} \approx I + \delta Q_t \quad \text{for } \delta \text{ suff. small.}$$

Imprecise continuous-time Markov chains

In practice, using a (non-)stationary pMC is often not warranted, as

- exactly specifying the (time-dependent) transition rate matrix Q_t is often infeasible, and
- assuming stationarity is not always justified.

Therefore, we assume that the time-dependent transition rate matrix Q_t of a pMC is only known to be contained in some (non-empty and bounded) *set* of transition rate matrices $\mathscr{Q}$.

Let $\mathbb{P}_{\mathscr{Q}}$ be the *set* of all Markov chains P consistent with $\mathscr{Q}$, in the sense that

$$(\forall t \in \mathbb{R}_{\geq 0})(\exists Q_t \in \mathscr{Q})\; T_t^{t+\delta} \approx I + \delta Q_t \quad \text{for } \delta \text{ suff. small.}$$

Krak et al. (2017) use this set $\mathbb{P}_{\mathscr{Q}}$ to characterise an *imprecise (continuous-time) Markov chain*. They define the *lower transition operator* $\underline{T}_t^{t+\Delta}$ as

$$\begin{aligned}[\underline{T}_t^{t+\Delta} f](x) &\coloneqq \underline{\mathrm{E}}(f(X_{t+\Delta})|X_t = x) \\ &= \underline{\mathrm{E}}(f(X_{t+\Delta})|X_{t_1} = x_1, \ldots, X_{t_n} = x_n, X_t = x),\end{aligned}$$

where $\underline{\mathrm{E}}(\cdot|\cdot)$ is the minimum of the conditional expectations induced by $\mathbb{P}_{\mathscr{Q}}$.

Imprecise continuous-time Markov chains

Krak et al. (2017) prove that, under certain conditions on $\mathscr{Q}$,

$$(\forall t, \Delta \in \mathbb{R}_{\geq 0})\ \underline{T}_t^{t+\Delta} = \underline{T}_0^{\Delta} =: \underline{T}_{\Delta}.$$

For any stationary precise continuous-time Markov chain,

$$T_t^{t+\Delta} = T_0^{\Delta} =: T_{\Delta}.$$

Imprecise continuous-time Markov chains

Krak et al. (2017) prove that, under certain conditions on $\mathscr{Q}$,

$$(\forall t, \Delta \in \mathbb{R}_{\geq 0})\; \underline{T}_t^{t+\Delta} = \underline{T}_0^{\Delta} =: \underline{T}_\Delta .$$

Moreover, they show that $\underline{T}_t$ is the unique operator that satisfies

$$\frac{\mathrm{d}}{\mathrm{d}t}\underline{T}_t = \underline{Q}\,\underline{T}_t \text{ with } \underline{T}_0 = I,$$

where $\underline{Q}$ is the *lower transition rate operator* associated with $\mathscr{Q}$, defined as

$$[\underline{Q}f](x) \coloneqq \min\left\{[Qf](x) \colon Q \in \mathscr{Q}\right\}.$$

For any stationary precise continuous-time Markov chain,

$$T_t^{t+\Delta} = T_0^{\Delta} =: T_\Delta .$$

Moreover, T_t is the unique matrix that satisfies

$$\frac{\mathrm{d}}{\mathrm{d}t}T_t = QT_t \text{ with } T_0 = I,$$

Imprecise continuous-time Markov chains

Krak et al. (2017) prove that, under certain conditions on $\mathscr{Q}$,

$$(\forall t, \Delta \in \mathbb{R}_{\geq 0})\ \underline{T}_t^{t+\Delta} = \underline{T}_0^{\Delta} =: \underline{T}_{\Delta}.$$

Moreover, they show that $\underline{T}_t$ is the unique operator that satisfies

$$\frac{\mathrm{d}}{\mathrm{d}t}\underline{T}_t = \underline{Q}\,\underline{T}_t \text{ with } \underline{T}_0 = I,$$

Krak et al. (2017) show that

$$\underline{T}_t = \lim_{n \to +\infty} \left(I + \frac{t}{n}\underline{Q}\right)^n.$$

For any stationary precise continuous-time Markov chain,

$$T_t^{t+\Delta} = T_0^{\Delta} =: T_{\Delta}.$$

Moreover, T_t is the unique matrix that satisfies

$$\frac{\mathrm{d}}{\mathrm{d}t}T_t = QT_t \text{ with } T_0 = I,$$

Furthermore,

$$T_t = \lim_{n \to +\infty} \left(I + \frac{t}{n}Q\right)^n.$$

Imprecise continuous-time Markov chains

Krak et al. (2017) prove that, under certain conditions on $\mathscr{Q}$,

$$(\forall t, \Delta \in \mathbb{R}_{\geq 0})\ \underline{T}_t^{t+\Delta} = \underline{T}_0^{\Delta} =: \underline{T}_{\Delta}.$$

Moreover, they show that $\underline{T}_t$ is the unique operator that satisfies

$$\frac{\mathrm{d}}{\mathrm{d}t}\underline{T}_t = \underline{Q}\underline{T}_t \text{ with } \underline{T}_0 = I,$$

For any sequence $\delta_{1:n}$ of *sufficiently small* steps that sum up to t,

$$[\underline{T}_t f](x) \approx [(I + \delta_1 \underline{Q}) \cdots (I + \delta_n \underline{Q}) f](x).$$

For any stationary precise continuous-time Markov chain,

$$T_t^{t+\Delta} = T_0^{\Delta} =: T_{\Delta}.$$

Moreover, T_t is the unique matrix that satisfies

$$\frac{\mathrm{d}}{\mathrm{d}t}T_t = QT_t \text{ with } T_0 = I,$$

For any sequence $\delta_{1:n}$ of steps that sum up to t,

$$[T_t f](x) \approx [(I + \delta_1 Q) \cdots (I + \delta_n Q) f](x).$$

Imprecise continuous-time Markov chains

Efficient computational methods
with guaranteed error bounds

Approximation methods

For any sequence $\delta_{1:n}$ of sufficiently small steps that sum up to t,

$$\underline{\mathrm{E}}(f(X_t)|X_0 = x) = [\underline{T}_t f](x) \approx [(I + \delta_1 \underline{Q}) \cdots (I + \delta_n \underline{Q}) f](x).$$

Approximation methods

For any sequence $\delta_{1:n}$ of sufficiently small steps that sum up to t,

$$\underline{\mathrm{E}}(f(X_t)|X_0 = x) = [\underline{T}_t f](x) \approx [(I + \delta_1 \underline{Q}) \cdots (I + \delta_n \underline{Q}) f](x).$$

A convenient framework to compute an approximation g_n of $\underline{T}_t f$ is

$g_0 \leftarrow f$, $\Delta_r \leftarrow \Delta$, $i \leftarrow 0$
while $\Delta_r > 0$ **do**
  $i \leftarrow i + 1$
  select some sufficiently small $0 < \delta_i \leq \Delta_r$
  $g_i \leftarrow g_{i-1} + \delta_i \underline{Q} g_{i-1}$
  $\Delta_r \leftarrow \Delta_r - \delta_i$
return g_n

$\simeq$ numerical integration of a non-linear vector differential equation

Approximation methods

For any sequence $\delta_{1:n}$ of sufficiently small steps that sum up to t,

$$\underline{\mathrm{E}}(f(X_t)|X_0 = x) = [\underline{T}_t f](x) \approx [(I + \delta_1 \underline{Q}) \cdots (I + \delta_n \underline{Q}) f](x).$$

A convenient framework to compute an approximation g_n of $\underline{T}_t f$ is

---

$g_0 \leftarrow f$, $\Delta_r \leftarrow \Delta$, $i \leftarrow 0$
while $\Delta_r > 0$ **do**
  $i \leftarrow i + 1$
  select some sufficiently small $0 < \delta_i \leq \Delta_r$
  $g_i \leftarrow g_{i-1} + \delta_i \underline{Q} g_{i-1}$
  $\Delta_r \leftarrow \Delta_r - \delta_i$
return g_n

$\simeq$ numerical integration of a non-linear vector differential equation

We are interested in approximation methods (i.e., methods that select $\delta_{1:n}$) that *a priori and theoretically* guarantee that the error

$$\|\underline{T}_t f - g_n\| \coloneqq \max\{|\underline{T}_t f(x) - g_n(x)| : x \in \mathscr{X}\}$$

is **lower than the *desired maximal error*** ϵ.

Approximation methods

For any sequence $\delta_{1:n}$ of sufficiently small steps that sum up to t,

$$\underline{\mathrm{E}}(f(X_t)|X_0 = x) = [\underline{T}_t f](x) \approx [(I + \delta_1 \underline{Q}) \cdots (I + \delta_n \underline{Q}) f](x).$$

A convenient framework to compute an approximation g_n of $\underline{T}_t f$ is

$g_0 \leftarrow f$, $\Delta_r \leftarrow \Delta$, $i \leftarrow 0$
while $\Delta_r > 0$ **do**
    $i \leftarrow i + 1$
    select some sufficiently small $0 < \delta_i \leq \Delta_r$
    $g_i \leftarrow g_{i-1} + \delta_i \underline{Q} g_{i-1}$
    $\Delta_r \leftarrow \Delta_r - \delta_i$
return g_n $\pm\epsilon$

$\simeq$ numerical integration of a non-linear vector differential equation

We are interested in approximation methods (i.e., methods that select $\delta_{1:n}$) that *a priori and theoretically* guarantee that the error

$$\|\underline{T}_t f - g_n\| := \max\{|\underline{T}_t f(x) - g_n(x)| : x \in \mathscr{X}\}$$

is **lower than the *desired maximal error*** ϵ.

Uniform approximation method

Krak et al. (2017) propose to use the *same* step size $\delta := t/n$ for every step.

They show that if n is greater than two lower bounds, one is guaranteed that

$$\|\underline{T}_t f - g_n\| \leq \epsilon.$$

Uniform approximation method

Krak et al. (2017) propose to use the *same* step size $\delta := t/n$ for every step.

They show that if n is greater than two lower bounds, one is guaranteed that

$$\|\underline{T}_t f - g_n\| \leq \epsilon.$$

We modify their method slightly on two fronts:

- we show that one of the two lower bounds can be less conservative, and
- we provide a method to compute a *posterior error bound* that is lower than—or at worst equal to—ϵ.

Uniform approximation method

Krak et al. (2017) propose to use the *same* step size $\delta := t/n$ for every step.

They show that if n is greater than two lower bounds, one is guaranteed that

$$\|\underline{T}_t f - g_n\| \leq \epsilon.$$

We modify their method slightly on two fronts:

- we show that one of the two lower bounds can be less conservative, and
- we provide a method to compute a *posterior error bound* that is lower than—or at worst equal to—ϵ.

In practice, the posterior error bound is often significantly smaller than ϵ.

Imprecise continuous-time Markov chains

Efficient computational methods

with guaranteed error bounds

Adaptive approximation method

We propose a method that constructs the sequence $\delta_{1:n}$ adaptively, by increasing the step size δ_i every m iterations.

Adaptive approximation method

We propose a method that constructs the sequence $\delta_{1:n}$ adaptively, by increasing the step size δ_i every m iterations.

The adaptive method is *more efficient* than the uniform method, in the sense that

- it selects step sizes δ_i such that the posterior error bound is as close to the desired maximal error ϵ as possible, and
- it requires (considerably) fewer iteration steps than the uniform approximation method.

Adaptive approximation method

We propose a method that constructs the sequence $\delta_{1:n}$ adaptively, by increasing the step size δ_i every m iterations.

The adaptive method is *more efficient* than the uniform method, in the sense that

- it selects step sizes δ_i such that the posterior error bound is as close to the desired maximal error ϵ as possible, and
- it requires (considerably) fewer iteration steps than the uniform approximation method.

However,

- increasing the step size adds non-negligible computational overhead.

Adaptive approximation method

We propose a method that constructs the sequence $\delta_{1:n}$ adaptively, by increasing the step size δ_i every m iterations.

The adaptive method is *more efficient* than the uniform method, in the sense that

- it selects step sizes δ_i such that the posterior error bound is as close to the desired maximal error ϵ as possible, and
- it requires (considerably) fewer iteration steps than the uniform approximation method.

However,

- increasing the step size adds non-negligible computational overhead.

Want to know how this method actually works?

Come see our poster on friday!

Numerical comparison

We consider the Healthy-Sick imprecise Markov chain introduced in (Krak et al., 2017), and determine $\underline{T}_t f$ up to the desired maximal error $\epsilon = 10^{-4}$ for some t and some f.

		duration of comp.		
method	# iter.	no p.e.b.	with p.e.b.	p. e. b.
uniform	80 000	0.414	1.19	4.29×10^{-5}
adaptive ($m = 1$)	34 360	0.593	0.856	1.00×10^{-4}
adaptive ($m = 10$)	34 369	0.224	0.529	1.00×10^{-4}

Numerical comparison

We consider the Healthy-Sick imprecise Markov chain introduced in (Krak et al., 2017), and determine $\underline{T}_t f$ up to the desired maximal error $\epsilon = 10^{-4}$ for some t and some f.

		duration of comp.		
method	# iter.	no p.e.b.	with p.e.b.	p. e. b.
uniform	80 000	0.414	1.19	4.29×10^{-5}
adaptive ($m = 1$)	34 360	0.593	0.856	1.00×10^{-4}
adaptive ($m = 10$)	34 369	0.224	0.529	1.00×10^{-4}

The adaptive method indeed

- needs fewer iterations than,

the uniform method.

Numerical comparison

We consider the Healthy-Sick imprecise Markov chain introduced in (Krak et al., 2017), and determine $\underline{T}_t f$ up to the desired maximal error $\epsilon = 10^{-4}$ for some t and some f.

method	# iter.	duration of comp. no p.e.b.	duration of comp. with p.e.b.	p. e. b.
uniform	80 000	0.414	1.19	4.29×10^{-5}
adaptive ($m = 1$)	34 360	0.593	0.856	1.00×10^{-4}
adaptive ($m = 10$)	34 369	0.224	0.529	1.00×10^{-4}

The adaptive method indeed

- needs fewer iterations than,
- takes less computational time than, and

the uniform method.

Numerical comparison

We consider the Healthy-Sick imprecise Markov chain introduced in (Krak et al., 2017), and determine $\underline{T}_t f$ up to the desired maximal error $\epsilon = 10^{-4}$ for some t and some f.

method	# iter.	duration of comp. no p.e.b.	duration of comp. with p.e.b.	p. e. b.
uniform	80 000	0.414	1.19	4.29×10^{-5}
adaptive ($m = 1$)	34 360	0.593	0.856	1.00×10^{-4}
adaptive ($m = 10$)	34 369	0.224	0.529	1.00×10^{-4}

The adaptive method indeed

- needs fewer iterations than,
- takes less computational time than, and
- results in a posterior error bound that is closer to ϵ compared to

the uniform method.

Ergodicity

Ergodicity

Definition (De Bock, 2017)

An imprecise Markov chain is ergodic if

$$\lim_{t\to+\infty}[\underline{T}_t f](x) = \underline{\mathrm{E}}_\infty(f) \qquad \text{for all } x \in \mathscr{X} \text{ and all } f \in \mathscr{L}(\mathscr{X}).$$

If this is the case, then $\underline{\mathrm{E}}_\infty$ is called the *limit lower expectation*.

Ergodicity

Definition (De Bock, 2017)

An imprecise Markov chain is ergodic if

$$\lim_{t \to +\infty} [\underline{T}_t f](x) = \underline{\mathrm{E}}_\infty(f) \qquad \text{for all } x \in \mathscr{X} \text{ and all } f \in \mathscr{L}(\mathscr{X}).$$

If this is the case, then $\underline{\mathrm{E}}_\infty$ is called the *limit lower expectation*.

Stationary precise Markov chains are a degenerate case of this definition. Indeed, a stationary precise Markov chain is ergodic if

$$\lim_{t \to +\infty} [T_t f](x) = \mathrm{E}_\infty(f) \qquad \text{for all } x \in \mathscr{X} \text{ and all } f \in \mathscr{L}(\mathscr{X}).$$

In this case, E_∞ is called the *limit (linear) expectation*.

Ergodicity

Definition (De Bock, 2017)

An imprecise Markov chain is ergodic if

$$\lim_{t \to +\infty} [\underline{T}_t f](x) = \underline{\mathrm{E}}_\infty(f) \qquad \text{for all } x \in \mathscr{X} \text{ and all } f \in \mathscr{L}(\mathscr{X}).$$

If this is the case, then $\underline{\mathrm{E}}_\infty$ is called the *limit lower expectation*.

Stationary precise Markov chains are a degenerate case of this definition. Indeed, a stationary precise Markov chain is ergodic if

$$\lim_{t \to +\infty} [T_t f](x) = \mathrm{E}_\infty(f) \qquad \text{for all } x \in \mathscr{X} \text{ and all } f \in \mathscr{L}(\mathscr{X}).$$

In this case, E_∞ is called the *limit (linear) expectation*.

How do we approximate $\underline{\mathrm{E}}_\infty(f)$?

The limit lower expectation $\underline{\mathrm{E}}_{\infty}$

Precise & stationary

Imprecise

ergodic CTMC with TRM Q

$\forall \delta > 0$, $\delta \|Q\| < 2$

ergodic DTMC with TM $(I + \delta Q)$

$$\mathrm{E}_{\infty}(f) = \lim_{n \to +\infty} [(I + \delta Q)^n f](x)$$

The limit lower expectation $\underline{\mathrm{E}}_\infty$

Precise & stationary		**Imprecise**
ergodic CTMC with TRM Q	$\mathscr{Q} = \{Q\}$	ergodic CTMC with LTRO $\underline{Q}$
$\forall \delta > 0,\ \delta \lVert Q \rVert < 2$		new $\forall \delta > 0,\ \delta \lVert \underline{Q} \rVert < 2$?
ergodic DTMC with TM $(I + \delta Q)$	$\mathscr{Q} = \{Q\}$	ergodic DTMC with LTO $(I + \delta \underline{Q})$

$\mathrm{E}_\infty(f) = \lim_{n \to +\infty} [(I + \delta Q)^n f](x)$

$\underline{\mathrm{E}}_\infty(f) = \lim_{n \to +\infty} [(I + \delta \underline{Q})^n f](x)$

The limit lower expectation $\underline{\mathrm{E}}_\infty$

Precise & stationary

ergodic CTMC with TRM Q

$\forall \delta > 0,$ $\delta \|Q\| < 2$

ergodic DTMC with TM $(I + \delta Q)$

$\mathscr{Q} = \{Q\}$

$\mathscr{Q} = \{Q\}$

Imprecise

ergodic CTMC with LTRO $\underline{Q}$

new

$\forall \delta > 0,$ $\delta \|\underline{Q}\| < 2$

X

ergodic DTMC with LTO $(I + \delta \underline{Q})$

$$\mathrm{E}_\infty(f) = \lim_{n \to +\infty} [(I + \delta Q)^n f](x)$$

$$\cancel{\underline{\mathrm{E}}_\infty(f) = \lim_{n \to +\infty} [(I + \delta \underline{Q})^n f](x)}$$

at least not for all δ such that $\delta \|\underline{Q}\| < 2$!

Interested in a more detailed explanation?

Did not understand a word of what I was saying?

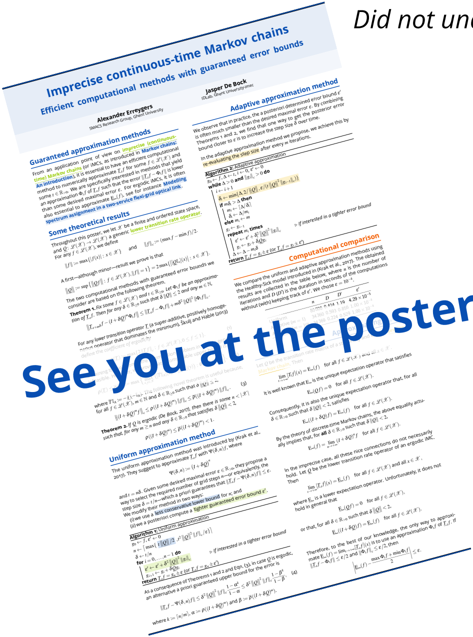

See you at the poster session on friday!

References

Jasper De Bock.
The limit behaviour of imprecise continuous-time Markov chains.
Journal of Nonlinear Science, 27(1):159–196, 2017.

Gert de Cooman, Filip Hermans, and Erik Quaeghebeur.
Imprecise Markov chains and their limit behavior.
Probability in the Engineering and Informational Sciences, 23(4):597–635, 2009.

Thomas Krak, Jasper De Bock, and Arno Siebes.
Imprecise continuous-time Markov chains.
2017. Accepted for publication in IJAR. arXiv:1611.05796 [math.PR].

Damjan Škulj.
Efficient computation of the bounds of continuous time imprecise Markov chains.
Applied Mathematics and Computation, 250:165–180, 2015.

Damjan Škulj and Robert Hable.
Coefficients of ergodicity for Markov chains with uncertain parameters.
Metrika, 76(1):107–133, 2013.