Discrete Markov Chains - PCCA+ and TPT

May 22, 2015

Formulation of the PCCA-Problem

There are infinitely many transformations $\mathbf{A}$ of the eigenvectors resulting in a soft membership matrix χ satisfying the positivity and partition of unity constraints. Consequently, we have to determine the transformation $\mathbf{A}$ that satisfies some optimality condition.
The starting point is the question whether it is possible to define some simplified dynamics on the coarse grained state space. Can we replace the original transition matrix $\mathbf{T}$ by a smaller transition matrix $\mathbf{T}_c \in \mathbb{R}^{n_c \times n_c}$ that propagates probability densities on the space of macrostates in a correct way? First of all, we note that densities on the original state space $\mathbf{p}_f \in \mathbb{R}^n$ can be restricted to a density $\mathbf{p}_c \in \mathbb{R}^{n_c}$ on the coarse state space by the **restriction** operator:

$$\begin{aligned} \mathbf{p}_c(i) &= \sum_{k=1}^{n} \mathbf{p}_f(k)\chi_i(k) \\ \mathbf{p}_c^T &= \mathbf{p}_f^T \chi \\ \mathbf{p}_c^T &= \mathbf{p}_f^T \mathbf{R}^T. \end{aligned}$$

In particular, we can define a coarse-grained stationary density by

$$\pi_c^T = \pi^T \mathbf{R}^T.$$

Ideally, we would like to define $\mathbf{T}_c$ such that restriction and propagation commute:

$$\mathbf{p}_f^T \mathbf{R}^T \mathbf{T}_c = \mathbf{p}_f^T \mathbf{T} \mathbf{R}^T. \tag{1}$$

We will see that generally, this is impossible, as the coarse graining process leads to a loss of information that is irreversible. Let us understand this in more detail: We start by computing the joint probability of being in macrostate i and in macrostate j one step later:

$$\begin{aligned} \mathbf{C}(i,j) &= \sum_{k,l} \chi_i(k) \left[\Pi \mathbf{T}\right](k,l) \chi_j(l) \\ &= \left[\chi^T \Pi \mathbf{T} \chi\right](i,j). \end{aligned}$$

This matrix can be normalized to become a row-stochastic matrix via

$$\begin{aligned}\mathbf{T}_c &= \overline{D}^{-1}\chi^T\Pi\mathbf{T}\chi \\ &= \mathbf{I}^T\mathbf{T}\mathbf{R}^T. \\ \overline{D} &= \mathrm{diag}(\pi_c) \\ \mathbf{I}^T &= \overline{D}^{-1}\chi^T\Pi.\end{aligned}$$

With these definitions, the matrix $\mathbf{T}_c$ consists of three steps: An interpolation onto the fine states, propagation by the original transition matrix, and restriction. We can observe that the interpolation of the fine density is always contained in the space of dominant left eigenvectors, as the columns of $\chi$are contained in the space of right eigenvectors. Consequently, restriction and propagation will only commute if the starting distribution is already contained in the space of dominant left eigenvectors, but not in general.
In order to arrive at a better solution, let us make the following two observations:

$$\begin{aligned}\mathbf{T}\mathbf{R}^T &= \mathbf{T}\mathbf{X}\mathbf{A} \\ &= \mathbf{X}\mathbf{\Lambda}\mathbf{A} \\ &= \mathbf{R}^T\mathbf{A}^{-1}\mathbf{\Lambda}\mathbf{A}. \\ \mathbf{R}^T\mathbf{T}_c &= \mathbf{R}^T\mathbf{I}^T\mathbf{T}\mathbf{R}^T \\ &= \mathbf{R}^T\mathbf{I}^T\mathbf{R}^T\mathbf{A}^{-1}\mathbf{\Lambda}\mathbf{A} \\ &= \mathbf{R}^T\left(\mathbf{R}\mathbf{I}\right)^T\mathbf{A}^{-1}\mathbf{\Lambda}\mathbf{A}.\end{aligned}$$

These two equalities show that there is an extra term $(\mathbf{R}\mathbf{I})^T$ involved if we first restrict and then propagate. Therefore, it was suggested to re-define $\mathbf{T}_c$ by

$$\mathbf{T}_c = \left(\mathbf{R}\mathbf{I}\right)^{-T}\mathbf{I}^T\mathbf{T}\mathbf{R}^T. \tag{2}$$

By construction, the commutation relation Eq. (1) is now satisfied. Moreover, the new $\mathbf{T}_c$ can be demonstrated to maintain the major properties of the original chain:

Lemma 1. *The coarse-grained propagator $\mathbf{T}_c$ defined by Eq. (2) possesses the same eigenvalues as the dominant eigenvalues of $\mathbf{T}$. Moreover, it has the stationary distribution π_c.*

Proof. The statement about eigenvalues follows from

$$\begin{aligned}\mathbf{T}_c &= \left(\mathbf{R}\mathbf{I}\right)^{-T}\left(\mathbf{R}\mathbf{I}\right)^T\mathbf{A}^{-1}\mathbf{\Lambda}\mathbf{A} \\ &= \mathbf{A}^{-1}\mathbf{\Lambda}\mathbf{A}.\end{aligned}$$

Similarly,

$$\begin{aligned}\pi_c^T\mathbf{T}_c &= \pi^T\mathbf{R}^T\mathbf{A}^{-1}\mathbf{\Lambda}\mathbf{A} \\ &= \pi^T\mathbf{X}\mathbf{\Lambda}\mathbf{A} \\ &= \pi^T\mathbf{T}\mathbf{X}\mathbf{A} \\ &= \pi^T\mathbf{R}^T \\ &= \pi_c^T.\end{aligned}$$

□

In summary, the matrix $\mathbf{T}_c$ inherits most of the original structure, but it is not a real transition matrix anymore. The inverse $(\mathbf{RI})^{-T}$ can introduce negative elements. Nevertheless, $\mathbf{T}_c$ allows us to formulate a useful objective function. The objective function is usually chosen in such a way that the factor $\mathbf{RI}$ is as close to the identity matrix as possible, because that means that $\mathbf{T}_c$ is close to a true transition matrix. Using the fact that $\mathrm{tr}(\mathbf{RI}) \leq n_c$, which can be shown via the definition of $\mathbf{RI}$, one frequently chooses

$$F(\mathbf{A}, \mathbf{X}, \pi) \quad = \quad n_c - \mathrm{tr}(\mathbf{RI})$$

in order to enforce this. Indeed, it can be shown that this objective function as well as the constraints can be expressed entirely in terms of $\mathbf{A}$. For more details on the algorithm, as well as different choice of the objective function, see the papers [1, 4].

Transition Path Theory (TPT)

The final step in our analysis is based on the following questions: Say we have determined two interesting subsets $A, B \subset S$ of the state space, s.t. $A \cap B = \emptyset$. These sets could have been obtained as metastable sets by PCCA+ or they might correspond to biologically interesting arrangements of a molecule. Can we determine the average time it takes to transition from A to B, and can we even determine a typical or probable pathway for this transition? For a detailed presentation of the theory, see the references [2, 3].

Hitting Probabilities and Committors

The starting point for these questions is a quantity we have already encountered, the hitting time of a set A, and the corresponding hitting probability:

$$\begin{aligned} H_A &= \min\{k \geq 0 : X_k \in A\}, \\ h_A(i) &= \mathbb{P}_i(H_A(i) < \infty). \end{aligned}$$

The vector h_A of hitting probabilities solves the following problem:

Lemma 2. *The vector h_A is the minimal solution of the problem*

$$h_A(i) \quad = \quad \begin{cases} 1, & i \in A \\ \sum_{j=1}^{n} \mathbf{T}_{ij} h_A(j), & \textit{otherwise.} \end{cases}$$

Proof. The statement for $i \in A$ is clear. The second statement follows by

$$
\begin{aligned}
h_A(i) &= \mathbb{P}_i(H_A < \infty) \\
&= \sum_{j=1}^{n} \mathbb{P}_i(H_A < \infty,\, X_1 = j) \\
&= \sum_{j=1}^{n} \mathbb{P}_i(H_A < \infty |\, X_1 = j)\mathbb{P}_i(X_1 = j) \\
&= \sum_{j=1}^{n} \mathbf{T}_{ij} h_A(j).
\end{aligned}
$$

We do not show the minimality of h_A here, the proof is rather technical and can be found in the old lecture notes. □

The hitting probabilities allow us to characterize the central quantity for TPT, the **committor probability**:

Definition 3. The forward committor of sets A, B is defined as the probability to hit set B next rather than A:

$$
q_i^+ = \mathbb{P}_i(H_B < H_A).
$$

Lemma 4. *The vector of forward committors satisfies the system of equations:*

$$
q_i^+ = \begin{cases} 1, & i \in B \\ 0, & i \in A \\ \sum_{j=1}^{n} \mathbf{T}_{ij} q_j^+, & \textit{otherwise.} \end{cases}
$$

Proof. All we have to do is define the so-called A-absorbing process:

$$
\tilde{\mathbf{T}}_{ij} = \begin{cases} \delta_{ij}, & i \in A \\ \mathbf{T}_{ij} & \text{otherwise.} \end{cases}
$$

The committor probability equals the hitting probability of set B under the A-absorbing process. The statement follows directly from the previous lemma. □

Definition 5. The backward committor q_i^- of sets A, B is defined as the probability to come from A rather than from B.

Lemma 6. *For reversible systems, the backward committor can be computed by*

$$
q_i^- = \begin{cases} 0, & i \in B \\ 1, & i \in A \\ \sum_{j=1}^{n} \mathbf{T}_{ij} q_j^-, & \textit{otherwise.} \end{cases}
$$

Proof. The backward committor can be computed in the same way as the forward committor using the time-reversed propagator $\hat{\mathbf{T}}_{ij} = \mathbb{P}(X_0 = j | X_1 = i)$. But for detailed balance, we have

$$\begin{aligned} \hat{\mathbf{T}}_{ij} &= \frac{\mathbb{P}(X_0 = j,\, X_1 = i)}{\pi_i} \\ &= \frac{\pi_j \mathbf{T}_{ji}}{\pi_i} \\ &= \mathbf{T}_{ij}. \end{aligned}$$

The assertion then follows from the previous lemma. □

Fluxes and Transition Rates

The forward and backward committor are sufficient to calculate transition rates between the sets A and B. First, we can define the **probability current** between states i, j:

$$f_{ij}^{AB} = \begin{cases} \pi_i q_i^- \mathbf{T}_{ij} q_j^+, & i \neq j \\ 0, & \text{otherwise.} \end{cases}$$

The probability current yields the average number of **reactive trajectories** (those going from A to B, without entering A before reaching B) that transition from i to j per time unit. Moreover, we can compute the **effective probability current** from i to j by

$$f_{ij}^+ = \max\left(f_{ij}^{AB} - f_{ji}^{AB}, 0\right).$$

Now, we can determine the average total number of trajectories going from A to B per time unit via

$$\begin{aligned} F^{AB} &= \sum_{i \in A} \sum_{j \in S} f_{ij}^{AB} \\ &= \sum_{i \in A} \sum_{j \in S} f_{ij}^{+}. \end{aligned}$$

Finally, the **transition rate** is the average fraction of reactive trajectories by the total number of trajectories that are going forward from state A:

$$\kappa_{AB} = \frac{F^{AB}}{\sum_{j \in S} \pi_j q_j^-}.$$

The inverse of the transition rate,

$$\tau_{AB} = \kappa_{AB}^{-1}$$

is called the **mean first-passage-time**.

Dominant Pathways

Finally, we briefly sketch how a characteristic reaction pathway can be identified from the effective probability currents. To this end, we make use of graph theory again, and define a directed graph G where the nodes are given by the states and the edges are weighted by the effective fluxes. To make this work, we have to assume that all effective currents are distinct. Let $w = (i_0, i_1, \ldots, i_K)$ be a simple reaction pathway, i.e. all i_j are different and $i_0 \in A$, $i_K \in B$, $i_1, \ldots, i_{K-1} \in (A \cup B)^c$. The **capacity** or **min-current** of the pathway is the minimum probability current

$$c(w) \quad = \quad \min_{(i,j) \in w} f_{ij}^+,$$

and the edge (i, j) where that minimum occurs is called the **bottleneck** of the pathway. The "best" pathway is then defined as the one which maximizes the min-current, let us call its bottleneck the global bottleneck. Usually, there will be more than one pathway that uses the global bottleneck. Let us call all of these pathways **dominant reaction pathways**. In order to remove the ambiguity and find a single best pathway, we can proceed as follows:

1. Determine the global bottleneck (b_1, b_2).
2. Determine all nodes and edges of dominant reaction pathways and reduce G to these nodes and edges.
3. The bottleneck divides the graph G into two disjoint parts of nodes L ("all nodes to the left of b_1") and R ("all nodes to the right of b_2").
4. Recursively find the best pathways in L with B replaced by $\{b_1\}$ and in R with A replaced by $\{b_2\}$.

References

[1] S. Kube and M. Weber. A coarse graining method for the identification of transition rates between molecular conformations. *J. Chem. Phys.*, 126(2):024103+, 2007.

[2] P. Metzner, C. Schütte, and E. Vanden-Eijnden. Transition Path Theory for Markov Jump Processes. *Multiscale Model. Simul.*, 7:1192–1219, 2009.

[3] F. Noé, C. Schütte, E. Vanden-Eijnden, L. Reich, and T.R. Weikl. Constructing the full ensemble of folding pathways from short off-equilibrium simulations. *Proc. Natl. Acad. Sci. USA*, 106:19011–19016, 2009.

[4] S. Röblitz and M. Weber. Fuzzy spectral clustering by pcca+: application to markov state models and data classification. *Adv. Data Anal. Classif.*, 7:147–179, 2013.