

AN INFORMATION-THEORETIC APPROACH FOR AUTOMATICALLY DETERMINING THE NUMBER OF STATE GROUPS WHEN AGGREGATING MARKOV CHAINS

Isaac J. Sledge¹, *Member, IEEE* José C. Príncipe^{2,3}, *Life Fellow, IEEE*

¹ Advanced Signal Processing and Automated Target Recognition Branch, Naval Surface Warfare Center—Panama City

² Department of Electrical and Computer Engineering, University of Florida

³ Department of Biomedical Engineering, University of Florida

ABSTRACT

A fundamental problem when aggregating Markov chains is the specification of the number of state groups. Too few state groups may fail to sufficiently capture the pertinent dynamics of the original, high-order Markov chain. Too many state groups may lead to a non-parsimonious, reduced-order Markov chain whose complexity rivals that of the original. In this paper, we show that an augmented value-of-information-based approach to aggregating Markov chains facilitates the determination of the number of state groups. The optimal state-group count coincides with the case where the complexity of the reduced-order chain is balanced against the mutual dependence between the original- and reduced-order chain dynamics.

Index Terms—Aggregation, model reduction, Markov chains, information theory

1. INTRODUCTION

Markov models have been widely adopted in a variety of disciplines. Part of their appeal is that the application and simulation of such models is rather efficient, provided that the corresponding state-space has a small to moderate size. Dealing with large state spaces is often troublesome, in comparison, as it may not be possible to adequately and efficiently simulate the underlying models [1–5].

A means of rendering the simulation of large-scale models tractable is crucial for many applications. One way to do this is by reducing the overall size of the Markov-chain state space via aggregation [6]. Aggregation typically entails either defining and utilizing a function to partition nodes in the probability transition graph associated with the large-scale chain. Groups of nodes, which are related by their inter-state transition probabilities and have strong interactions, are combined and treated as a single aggregated node in a new graph. This results in a lower-order chain with a reduced state space. A stochastic matrix for the lower-order chain is then specified, which describes the transitions from one super-state to another. This stochastic matrix should roughly mimic the dynamics of the original chain despite the potential loss in information [7–9].

There are a variety of methods for aggregating Markov chains. Some of the earliest work exploited the strong-weak interaction structure of nearly completely decomposable Markov chains to obtain reduced-order approximations [10–13]. Both uncontrolled [14, 15] and controlled Markov chains [16, 17] have been extensively studied in the literature.

In this paper, we consider an approach for aggregating nearly-completely-decomposable Markov chains, which is composed of two information-theoretic processes [18]. The first process entails quantifying the dissimilarity of nodes in the original and reduced-order probability transition graphs, despite the difference in state space sizes. The second process involves iteratively partitioning similar nodes without explicit knowledge of the number of groups. For this second process, we consider the use of an information-theoretic criterion known as the value of information [19, 20] to efficiently partition the probability transition graph. The value of information is a constrained, modified-free-energy-difference criterion that describes the maximum benefit associated with a

given quantity of information in order to minimize the average distortion [21–23]. It is an optimal, non-linear conversion between information, usually in the Shannon sense [24], and either costs or utilities, in the von-Neumann-Morgenstern sense [25].

Optimizing the value of information in a grouped-coordinate-descent manner yields a single free parameter that represents the effect of the information bound. Increasing this parameter from some base value yields a hierarchy of partitions that monotonically decrease the modified-free-energy difference. Each hierarchy element corresponds to a partition with an increasing information bound amount and hence a potentially increasing number of state groups. Finer-scale group structure in the transition matrix is captured as the parameter value rises. After some free-parameter value threshold, however, there are diminishing returns on the aggregation quality. Finding this threshold, in a completely data-driven fashion, is thus crucial: over-partitioning the original chain leads to a non-parsimonious representation that can be computationally expensive to evaluate.

To automatically discern the optimal number of state groups for arbitrary Markov chains, we modify the value-of-information cost function and hence its updates. For this modified criterion, the first update step trades off the divergence of the original and aggregated chain transition probabilities against the mutual dependence between the original and aggregated chain dynamics. The remaining updates then trade off between how well the reduced-order model dynamics are compressed versus how much information about the original-chain dynamics is retained. That is, it follows an information-bottleneck model [26] with the added effect that the resulting Markov chain resembles the original. Through update schedule, we are guaranteed that the negative modified-free-energy-difference curve is, in general, convex. We can therefore find a maximal value on the curve. This value corresponds to the situation where the aggregation process begins to overfit to noise in the Markov-chain transition dynamics versus the underlying nearly-completely-decomposable structure if more state groups are added.

2. METHODOLOGY

Our approach for aggregating Markov chains can be described as follows. Given a stochastic matrix $\Pi \in \mathbb{R}_+^{n \times n}$ of transition probabilities between n states, we seek to partition this matrix to produce a reduced-size stochastic matrix $\Phi \in \mathbb{R}_+^{m \times m}$ with m states. Since there are many possible Φ 's that can be formed, we would like one with the least divergence to Π for some measure. Due to the different sizes of Π and Φ , though, directly assessing divergence is not possible. To facilitate this comparison, we consider the construction of a joint-model stochastic matrix $\Theta \in \mathbb{R}_+^{m \times n}$ that encodes the dynamics of Φ .

The definition of Θ relies on finding a partition matrix $\Psi \in \mathbb{R}_+^{m \times n}$ for determining which states in Π should be combined to create an aggregated state in Φ . We show that an optimal Ψ and Θ , and hence Φ , can be found by solving a modified value-of-information criterion. We then characterize the errors associated with estimating this modified value-of-information for Markov chains with a finite number of states. Removing the errors penalizes marginal improvements in the divergence associated with including more state groups m . This non-linearly transforms the value of information such that it contains a global maximum for a single state group count. We then specify an expression for this value, which coincides where the aggregation complexity is balanced against the preserved information.

This work was funded by ONR grants N00014-15-1-2013 and N00014-14-1-0542. The first author was additionally funded by a UF Research Fellowship, a UF Robert C. Pittman Fellowship, and an ASEE NREIP Fellowship.

$$\min_{\Psi^{(k)} \in \mathbb{R}_+^{n \times m}, \Theta^{(k)} \in \mathbb{R}_+^{m \times n}} \left(\sum_{i=1}^n \sum_{j=1}^m \gamma_i \psi_{i,j}^{(k)} \sum_{p=1}^n \pi_{i,p} \log(\pi_{i,p} / \vartheta_{j,p}^{(k)}) \mid \begin{array}{l} 0 \leq \vartheta_{i,q}^{(k)}, \psi_{i,q}^{(k)} \leq 1, \sum_{q=1}^m \vartheta_{i,q}^{(k)} = 1, \\ \sum_{q=1}^m \psi_{i,q}^{(k)} = 1 \forall i, j, k \end{array} \right), \forall k \quad (1)$$

$$\min_{\Psi^{(k)} \in \mathbb{R}_+^{n \times m}} \left(\begin{array}{l} \mathbb{I}_{(k=0)} \sum_{i=1}^n \sum_{j=1}^m \gamma_i \psi_{i,j}^{(k)} \sum_{p=1}^n \pi_{i,p} \log(\pi_{i,p} / \vartheta_{j,p}^{(k)}) + \\ \mathbb{I}_{(k>0)} \sum_q \sum_{j=1}^m \omega_q \kappa_{q,j}^{(k)} \log(\kappa_{q,j}^{(k)} / \alpha_j^{(k)}) \end{array} \mid \begin{array}{l} \sum_{j=1}^m \alpha_j^{(k)} \sum_{i=1}^n \psi_{i,j}^{(k)} \log(\psi_{i,j}^{(k)} / \gamma_i) \leq r \\ \kappa_{q,j}^{(k)} = \sum_{i=1}^n \gamma_i \eta_{q,i} \psi_{i,j}^{(k)} / \alpha_j^{(k)} \\ 0 \leq \vartheta_{i,q}^{(k)}, \psi_{i,q}^{(k)} \leq 1, \sum_{q=1}^m \vartheta_{i,q}^{(k)} = 1, \sum_{q=1}^m \psi_{i,q}^{(k)} = 1 \end{array} \right), \forall k \quad (2)$$

2.1 Preliminaries

For our approach, we consider a first-order, homogeneous Markov chain defined on a finite state space. We assume that is nearly-completely decomposable.

Definition 2.1. The transition model of a first-order, homogeneous, nearly-completely-decomposable Markov chain is a weighted, directed graph R_π given by the three-tuple (V_π, E_π, Π) with:

(i) A set of n vertices $V_\pi = v_1^1 \cup \dots \cup v_n^n$ representing the states of the Markov chain.

(ii) A set of $n \times n$ edge connections $E_\pi \subset V_\pi \times V_\pi$ between reachable states in the Markov chain.

(iii) A stochastic transition matrix $\Pi \in \mathbb{R}_+^{n \times n}$. Here, $[\Pi]_{i,j} = \pi_{i,j}$ represents the non-negative transition probability between states i and j . We impose the constraint that the probability of experiencing a state transition is independent of time. Moreover, for a block-diagonal matrix Π^* with zeros along the diagonal, we have that $\Pi = \Pi^* + \varepsilon C$. $\Pi^* \in \mathbb{R}_+^{n \times n}$ is a completely-decomposable stochastic matrix with m indecomposable sub-matrix blocks Π_i^* of order n_i . The matrix $C \in \mathbb{R}_+^{n \times n}$ satisfies $\sum_{k=1}^{n_i} c_{p_i,k_i} = -\sum_{j \neq i} \sum_{q=1}^{n_j} c_{p_i,q_j} \forall p_i$, for blocks Π_i^* and Π_j^* .

Throughout, we assume that all Markov chains are irreducible and aperiodic. As a consequence, there is a unique invariant probability distribution γ associated with the chain such that $\gamma^\top \Pi = \gamma^\top$.

We are interested in comparing pairs of nearly-completely-decomposable Markov chains. A means to do this is by considering given rows of the stochastic transition matrix Π with those of a reduced-order chain's stochastic matrix Φ . We will perform this comparison via the negative Kullback-Leibler divergence. It coincides with the Donsker-Varadhan rate function appearing in the large-deviations theory of Markov chains [27, 28] and measures the dissimilarity between chains defined on the same discrete state space.

Since we are considering the problem of chain aggregation, the discrete state spaces will be different. One chain R_π will have n states while another R_φ will have m states. The dimensionalities of given rows in the corresponding transition matrices will hence not be equivalent, which precludes a direct comparison using the discrete Kullback-Leibler divergence. To resolve this issue, we consider construction of a joint model R_ϑ . This joint model defines a joint state space composed of the states from R_π and R_φ . It, however, re-defines the edge set along with the weighting matrix. This weighting matrix Θ has the same number of columns as Π and the same dynamics as Φ , which facilitates comparisons using conventional divergences.

The joint model relies on the specification of a partition matrix Ψ for mapping states from the reduced-order model R_φ to the original model R_π . Here, we consider probabilistic partition functions so as to capture the inherent uncertainty in the state combination.

Definition 2.2. Let $R_\pi = (V_\pi, E_\pi, \Pi)$ and $R_\varphi = (V_\varphi, E_\varphi, \Phi)$ be transition models of two Markov chains over n and m states, respectively. A probabilistic partition function ψ is a surjective mapping

between two state index sets, $\mathbb{Z}_{1:n}$ and $\mathbb{Z}_{1:m}$, such that $\psi^{-1}(\mathbb{Z}_{1:m})$ is a partition of $\mathbb{Z}_{1:n}$, which has a given probabilistic chance of occurring. That is, $\psi^{-1}(j) \subset \mathbb{Z}_{1:n} \times \mathbb{R}_+^n$ is not empty and where $\psi^{-1}(1) \cup \dots \cup \psi^{-1}(m) = \mathbb{Z}_{1:n} \times \mathbb{R}_+^n$, with the real-valued responses being non-negative and summing to one.

The probabilistic partition of a state index set induces a probabilistic partition matrix $[\Psi]_{i,j} = \psi_{i,j}$, where $\psi_{i,j} = \zeta$ if $i \in \psi^{-1}(j)$ occurs with probability ζ . The set of all probabilistic partition matrices is $\{\Psi \in \mathbb{R}_+^{n \times m} \mid [\Psi]_{i,j} = \psi_{i,j} \in [0, 1], \sum_{j=1}^m \psi_{i,j} = 1\}$.

Definition 2.3. Let $R_\pi = (V_\pi, E_\pi, \Pi)$ and $R_\varphi = (V_\varphi, E_\varphi, \Phi)$ be transition models of two Markov chains over n and m states, respectively, where $m < n$. $R_\vartheta = (V_\vartheta, E_\vartheta, \Theta)$ is a joint model, with $m+n$ states, that is defined by

(i) A vertex set $V_\vartheta = V_\pi \cup V_\varphi$, which is the union of all state vertices in R_π and R_φ .

(ii) An edge set $E_\vartheta \subset V_\vartheta \times V_\vartheta$, which are one-to-many mappings from the states in the original transition model R_π to the reduced-order transition model R_φ .

(iii) A weighting matrix $\Theta \in \mathbb{R}_+^{(m+n) \times (m+n)}$. The partition function ψ provides a relationship between the stochastic matrices Φ and Θ of R_φ and R_ϑ , respectively. This is given by $\Theta = \Theta \Psi$, or, rather, $\varphi_{i,j} = \sum_{k=1}^n \vartheta_{k,j} \psi_{k,i} \forall i, j$, where Ψ is a probabilistic partition matrix; here, $\psi_{i,j} = p(v_\varphi^j | v_\pi^i)$ and $\varphi_{i,j} = p(v_\varphi^j | v_\vartheta^i)$.

Our corresponding technical report should be consulted for further details about these choices and illustrations of the various concepts [29]. This report also contains proofs for many of the ensuing claims.

2.2 Aggregating Markov Chains

For any given transition model R_π , we would like to find, by way of the joint model R_ϑ , another transition model R_φ with fewer states that resembles the dynamics encoded by R_π . At the very least, a model R_ϑ should be sought with a weighting matrix Θ that has the least expected divergence with respect to the transition matrix Π of R_π for some partition.

Definition 2.4. Let $R_\pi = (V_\pi, E_\pi, \Pi)$, $R_\varphi = (V_\varphi, E_\varphi, \Phi)$, and $R_\vartheta = (V_\vartheta, E_\vartheta, \Theta)$ be transition models of two Markov chains over n and m states and the joint model over $n+m$ states, respectively. The least expected divergence between Π and Θ , and hence Π and Φ , is given in (1). Here, $\gamma_i = p(v_\pi^i)$ and $\pi_{i,k} = p(v_\pi^i | v_\pi^k)$.

The least expected distortion possesses too few constraints to make discerning the number of state groups m viable, though.

To help discern m , we impose that the partitions should minimize the information loss associated with the state quantization process. That is, the mutual dependence between states in the high-order and low-order chains should be maximized with respect to a supplied bound. Simultaneously, the least expected divergence, for this bound, should be achieved while also seeking a maximal compression of the dynamics. Realizing each of these competing objectives leads to a combined value-of-information and information-bottleneck aggregation process.

$$\alpha_j^{(0)} \leftarrow \sum_{i=1}^n \gamma_i \psi_{i,j}^{(0)}, \quad \psi_{i,j}^{(1)} \leftarrow \left(\alpha_j^{(0)} e^{-\beta \sum_{r=1}^n \pi_{i,r} \log(\pi_{i,r} / \vartheta_{j,r}^{(1)})} \right) / \left(\sum_{p=1}^m \alpha_p^{(0)} e^{-\beta \sum_{r=1}^n \pi_{i,r} \log(\pi_{i,r} / \vartheta_{p,r}^{(1)})} \right) \quad (3)$$

$$\alpha_j^{(k)} \leftarrow \sum_{i=1}^n \gamma_i \psi_{i,j}^{(k)}, \quad \kappa_{q,j}^{(k)} \leftarrow \sum_{i=1}^n \gamma_i \eta_{q,i} \psi_{i,j}^{(k)} / \alpha_j^{(k)}, \quad \psi_{i,j}^{(k+1)} \leftarrow \left(\alpha_j^{(k)} e^{-\beta \sum_q \eta_{q,i} \log(\eta_{q,i} / \kappa_{q,j}^{(k)})} \right) / \left(\sum_{p=1}^m \alpha_p^{(k)} e^{-\beta \sum_q \eta_{q,i} \log(\eta_{q,i} / \kappa_{q,p}^{(k)})} \right) \quad (4)$$

$$\min_{\Psi^{(k)} \in \mathbb{R}_+^{n \times m}} \left(\begin{array}{c} \mathbb{I}_{(k=0)} \sum_{i=1}^n \sum_{j=1}^m \gamma_i \psi_{i,j}^{(k)} \sum_{p=1}^n \pi_{i,p} \log(\pi_{i,p} / \vartheta_{j,p}^{(k)}) + \mathbb{I}_{(k>0)} \sum_q \sum_{j=1}^m \omega_q \kappa_{q,j}^{(k)} \log(\kappa_{q,j}^{(k)} / \alpha_j^{(k)}) \\ + \mathbb{I}_{(k=0)} \sum_{g=2}^{\infty} \frac{(-1)^g}{\log(2)(g^2-g)} \sum_{j=1}^m \mathbb{E}[(\sum_{i=1}^n \bar{\gamma}_i \psi_{i,j}^{(k)})^g] / (\alpha_j^{(k)})^{g-1} \\ - \mathbb{I}_{(k>0)} \sum_{g=2}^{\infty} \frac{(-1)^g}{\log(2)(g^2-g)} \sum_q \sum_{j=1}^m \frac{\omega_q^{g-1} \mathbb{E}[(\sum_{i=1}^n \bar{\eta}_{q,i} \rho_{q,i}^{(k)})^g] - (\rho_{q,j}^{(k)})^{g-1} \mathbb{E}[(\sum_{i=1}^n \gamma_y \bar{\eta}_{q,i})^g]}{(\rho_{q,j}^{(k)})^{g-1} \omega_q^{g-1}} \end{array} \right), \forall k \quad (5)$$

Definition 2.5. Let $R_\pi = (V_\pi, E_\pi, \Pi)$, $R_\varphi = (V_\varphi, E_\varphi, \Phi)$, and $R_\vartheta = (V_\vartheta, E_\vartheta, \Theta)$ be transition models of two Markov chains over n and m states and the joint model over $n+m$ states, respectively. An optimal reduced-order transition model R_φ with respect to the original model R_π can be found as follows:

(i) Optimal partitioning: Find a probabilistic partition matrix Ψ that solves (2) for some $r \in \mathbb{R}_+$; r has an upper bound of $-\sum_{i=1}^n \gamma_i \log(\gamma_i)$. As well, find the corresponding weighting matrix for R_ϑ : $\Theta = U^\top \Pi$, where $[U]_{i,j} = \gamma_i \psi_{i,j} / \sum_{p=1}^n \gamma_p \psi_{p,j}$.

(ii) Transition matrix construction: Obtain the transition matrix for R_φ via: $\varphi_{i,j} = \sum_{k=1}^n \vartheta_{k,j} \psi_{k,i}$ using the optimal weights Θ and the probabilistic partition matrix Ψ from step (i).

Here, $\alpha_j = p(v_\varphi^j)$, $\omega_q = p(q)$, $\eta_{q,i} = p(q|v_\pi^i)$, and $\kappa_{q,j} = p(q|v_\varphi^j)$ for some intermediate random variable q .

A reason for considering this combined model is that simply trading off between the preserved information and the model complexity, as in the information bottleneck, does not encode notions of the underlying weighted-graph geometry. This is because Shannon information is geometrically invariant: it yields an infinite set of degenerate solutions that do not minimize a specific divergence measure.

Global solutions to (2) can be found via the initial condition in (3) and the expectation-maximization-like updates in (4).

Proposition 2.1. Let $R_\pi = (V_\pi, E_\pi, \Pi)$ and $R_\varphi = (V_\varphi, E_\varphi, \Phi)$ be transition models of two Markov chains over n and m states. The optimal partition matrix Ψ that globally solves (2) can be found via the alternating update in (3) for iteration $k=0$ and the alternating updates in (4) for iterations $k=1, 2, \dots$. Here, $\beta \in \mathbb{R}_+$ is a Lagrange multiplier for handling the constraint bound $r \in \mathbb{R}_+$.

The following proposition shows that the number of reduced-model state groups increases once β reaches certain critical values.

Proposition 2.2. Let $R_\pi = (V_\pi, E_\pi, \Pi)$ and $R_\varphi = (V_\varphi, E_\varphi, \Phi)$ be transition models of two Markov chains over n and m states. For some β_0 , suppose Θ_{β_0} , the matrix Θ for that value of β_0 , satisfies the following inequality $d^2/d\epsilon^2 F(\Psi, K, \alpha; \Pi, \Theta_{\beta_0} + \epsilon Q, H, \gamma)|_{\epsilon=0} > 0$, for the modified value-of-information Lagrangian from (2). Here, $Q \in \mathbb{R}_+^{m \times n}$ is such that $\sum_{k=1}^m \top q_{k,1:n} q_{k,1:n} = 1$ and $\sum_{j=1}^n q_{i,j} = 0 \forall i$. A critical value β_c satisfies

$$\beta_c = \min_{\beta > \beta_0} (d^2/d\epsilon^2 F(\Psi, K, \alpha; \Pi, \Theta_\beta + \epsilon Q, H, \gamma)|_{\epsilon=0} \leq 0).$$

The number of rows in Θ and columns in Ψ needs to be increased, by one, once $\beta > \beta_c$, since a phase change occurs.

Sweeping over critical values of β yields a hierarchy of probabilistic partitions Ψ and hence reduced-order-model transition matrices Φ with different numbers of state groups.

This proposition does not, however, reveal a way to find the optimal number of state groups. To do that, we will determine when the modified value-of-information begins to overfit to the transition-dynamics noise. This leads to additional terms that we can subtract from the criterion to

essentially regularize the aggregation and penalize for using too many or too few state groups than can be resolved in the transition dynamics.

2.2.1 Aggregated State Group Count

We assume, for the modified value-of-information, that overfitting to noise is a byproduct of dealing with finite-state-space Markov chains and hence introducing estimation errors into the joint probabilities. Characterizing the error effects and removing them yields a version of the value-of-information curve that is either convex or monotonically non-decreasing then monotonically non-increasing.

Definition 2.6. Let $R_\pi = (V_\pi, E_\pi, \Pi)$, $R_\varphi = (V_\varphi, E_\varphi, \Phi)$, and $R_\vartheta = (V_\vartheta, E_\vartheta, \Theta)$ be transition models of two Markov chains over n and m states and the joint model over $n+m$ states, respectively. An optimal reduced-order transition model R_φ , with respect to the original model R_π , can be found as follows:

(i) Optimal partitioning: Find a probabilistic partition matrix Ψ that solves (5) for some $r \in \mathbb{R}_+$. Equation (5) has the same constraint set as (2). Find the corresponding weighting matrix for R_ϑ : $\Theta = U^\top \Pi$, where U is defined above.

(ii) Transition matrix construction: Obtain the transition matrix for R_φ via: $\varphi_{i,j} = \sum_{k=1}^n \vartheta_{k,j} \psi_{k,i}$ using the optimal weights Θ and the probabilistic partition matrix Ψ from step (i).

The systematic underestimation/overestimation error in (2) is g th-order minimized when solving (5).

Proposition 2.3. Let $R_\pi = (V_\pi, E_\pi, \Pi)$ and $R_\varphi = (V_\varphi, E_\varphi, \Phi)$ be transition models of two Markov chains over n and m states. The optimal partition matrix Ψ for R_π that globally solves (5) and facilitates the construction of R_φ can be found via the updates in (6) for iteration $k=0$ and the update in (7) for iterations $k=1, 2, \dots$. For (6) and (7), the denominators are such that the partition-matrix entries are normalized to become probabilities.

Here, $\alpha_j = p(v_\varphi^j)$, $\gamma_i = p(v_\pi^i)$, $\psi_{i,j} = p(v_\varphi^j | v_\pi^i)$, $\eta_{q,i} = p(q|v_\pi^i)$, $\kappa_{q,j} = p(q|v_\varphi^j)$, $\tau_{i,j} = p(v_\pi^i | v_\varphi^j)$, and $\rho_{q,j} = p(q, v_\varphi^j)$. The terms $\bar{\kappa}_{q,j}$ and $\bar{\eta}_{q,i}$ represent the estimation errors associated with $\kappa_{q,j}$ and $\eta_{q,i}$, while $\bar{\gamma}_i$ is the error associated with approximating γ_i . We assume that the average error is zero: $\mathbb{E}[\bar{\eta}_{q,i}] = 0$ and $\mathbb{E}[\bar{\gamma}_i] = 0$.

Due to the monotonicity properties of (5), we can construct upper and lower bounds for it and characterize their rate of change with respect to β . This permits algebraically discerning where maxima of the equivalent dual problem (5) occur.

Proposition 2.4. Let $R_\pi = (V_\pi, E_\pi, \Pi)$, $R_\varphi = (V_\varphi, E_\varphi, \Phi)$, and $R_\vartheta = (V_\vartheta, E_\vartheta, \Theta)$ be transition models of two Markov chains over n and m states and the joint model over $n+m$ states, respectively. The dual problem to (5) achieves a maximum for the Lagrange multiplier value $\beta^* = 2 \sum_{i=1}^n \sum_{j=1}^m \gamma_i \psi_{i,j} \log(\psi_{i,j} / \alpha_j) / 2n$.

As before, the partitioning process defined by (6) and (7) undergoes a series of phase changes whereby the number of state groups increase once β exceeds some critical value. The value of β^* specified by the preceding proposition coincides with a partitioning where a certain

$$\psi_{i,j}^{(1)} \leftarrow \left(\sum_{y=1}^n \gamma_y \psi_{y,j}^{(0)} \exp \left(-\beta \left(\sum_{p=1}^n \pi_{i,p} \log \left(\frac{\pi_{i,p}}{\vartheta_{j,p}^{(0)}} \right) \right) + \sum_{g=2}^{\infty} (-1)^g \sum_{p=1}^m \frac{\mathbb{E}[(\sum_{s=1}^n \bar{\gamma}_s \psi_{s,p}^{(0)})^g]}{g(\alpha_p^{(0)})^g} - \frac{\mathbb{E}[\bar{\gamma}_i (\sum_{s=1}^n \bar{\gamma}_s \psi_{s,p}^{(0)})^{g-1}]}{(g-1)\gamma_i (\alpha_p^{(0)})^{g-1}} \right) \right) / \left(\dots \right) \quad (6)$$

$$\alpha_j^{(k)} \leftarrow \sum_{i=1}^n \gamma_i \psi_{i,j}^{(k)}, \quad \kappa_{q,j}^{(k)} \leftarrow \sum_{i=1}^n \gamma_i \eta_{q,i} \psi_{i,j}^{(k)} / \alpha_j^{(k)}, \quad \bar{\kappa}_{q,j}^{(k)} \leftarrow \sum_{i=1}^n \bar{\eta}_{q,i} \tau_{i,j}^{(k)},$$

$$\psi_{i,j}^{(k+1)} \leftarrow \left(\alpha_j^{(k)} \exp \left(-\beta \left(\sum_q \eta_{q,i} \log \left(\frac{\eta_{q,i}}{\kappa_{q,j}^{(k)}} \right) \right) - \beta \sum_{g=2}^{\infty} (-1)^g \left(\sum_q \frac{\mathbb{E}[(\bar{\kappa}_{q,j}^{(k)})^g]}{\eta_{q,i}^{-1} g (\kappa_{q,j}^{(k)})^g} - \frac{\mathbb{E}[\bar{\eta}_{q,i} (\bar{\kappa}_{q,j}^{(k)})^{g-1}]}{(g-1)(\kappa_{q,j}^{(k)})^{g-1}} \right) \right) \right) / \left(\dots \right) \quad (7)$$

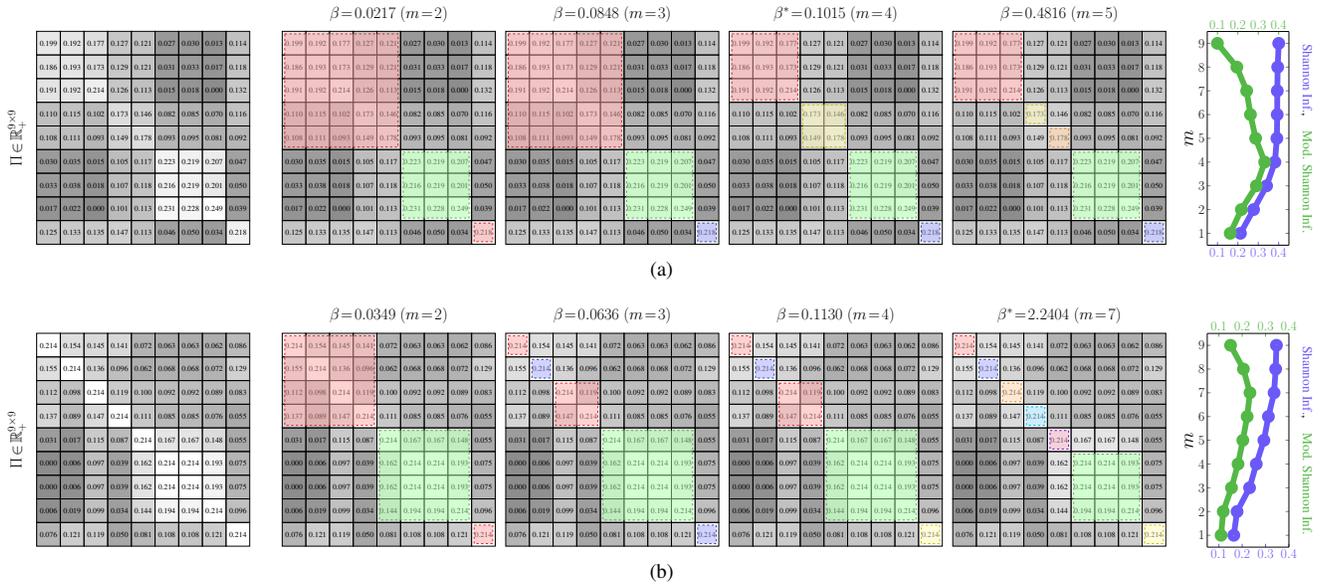


Figure 1: Modified value-of-information-based partitioning results for 9-state nearly-completely-decomposable Markov chains with: (a) four discernible state groups and (b) seven discernible state groups. In both (a) and (b), we show the original stochastic matrix Π with hardened versions of the partitions Ψ overlaid for four critical values of β ; as noted in the previous section, m can be inferred from each value of β . The unique colors in the partition plots correspond to which state group in Φ a state in Π is most likely to be associated. The right-most plots highlight the Shannon information in blue and the error-subtracted Shannon information in green. After a certain number of clusters, the Shannon information plateaus, indicating that there is negligible benefit for including more state groups in Φ . The modified Shannon information begins to decrease when this occurs; the results align with the discernible number of groups for these chains.

number of state groups are defined. It can be taken as the point in (5) where the complexity of the dynamics compression is balanced against the information contained about the original Markov chain in the reduced-order Markov chain. For (2), β^* often marks the beginning of the value-of-information's asymptotic region where there are diminishing returns for including more clusters.

3. SIMULATIONS

In this section, we assess the empirical performance of the modified value-of-information criterion. We gauge how well the predicted ‘optimal’ free-parameter value aligns with the discernible number of state groups for nearly-completely-decomposable Markov chains.

3.1 Simulation Protocols

We adopted the following simulation protocols for our aggregation framework. We initialized the aggregation process with a partition matrix of all ones, $\Psi = [1]_{9 \times 1}$, signifying that each state belongs to a single group. This is the global optimal solution of the aggregation problem and coincides with a parameter value β of zero for the modified value-of-information. We then found the subsequent critical values of β and increased the column count of Ψ . We determined which state group would be further split and modified both the new column and an existing column of Ψ to randomly allocate the appropriate states. This initialization process bootstraps the quantization for the new cluster and typically achieves convergence in only a few iterations. It also permits the value of information to reliably track the global minimizer as β increases.

For certain problems, a priori specifying a fixed amount of partition updates may not permit finding a steady-state solution. We therefore run the alternating updates until no entries of the partition matrix change across two iterations.

3.2 Simulation Results

We established the performance of the aggregation partitioning process through two examples. The first, shown in figure 1(a), corresponds to a Markov chain with nine states and four state groups with strong intra-group interactions and weak inter-group interactions. This is a relatively simple aggregation problem. The second example, presented in figure 1(b), is of a nine-state Markov chain with a single dominant state group and six outlying states with near-equal transition probabilities. This is a more challenging problem than the first, as the outlying states

cannot be reliably combined without adversely impacting the mutual dependence.

In figures 1(a) and 1(b), we provided partitions for four critical values of the free parameter β . The ‘optimal’ value of β , as predicted by our modified value of information, leads to four and seven state groups for the first and second examples, respectively. Here, we considered a second-order approximation of the underestimation/overestimation error, i.e., $g=2$. Higher-order approximations did not modify the shape of the resulting Shannon-information curves much compared to those in figures 1(a) and 1(b); the resulting ‘optimal’ number of state groups remained the same in these cases.

For both examples, the associated partitions for the ‘optimal’ values of β align well with an inspection of the dynamics of the stochastic matrix: the partitions separate states that are more likely to transition to each other from those that are not. Those partitions for ‘non-optimal’ β 's either over- or under-quantize the chain states. That is, for critical β 's before the ‘optimal’ value, there is a moderate increase in the Shannon information, while the remaining β 's only yield modest increases. The ‘optimal’ value of β for both examples, in contrast, lies at the ‘knee’ of this curve, which is where the divergence minimization is balanced against the competing objective of state-mutual-dependence maximization with respect to some bound. It is also where the complexity of the aggregation result is in harmony with the information it contains about the original Markov chain.

To lend credence to these results, we considered forty reformulated graph-based partition validity measures [30, 31] across a hundred Monte Carlo simulations. For the Markov chain in figure 1(a), thirty-three indices specified that there were four state groups while seven indices indicated that there were three state groups. The results for the chain in figure 1(b) had more uncertainty, which was due to the high number of outlying states and the ability to combine them in multiple ways to reduce the models’ divergence. Ten indices suggested that there were two state groups, eight indices chose three state groups, while the remaining twenty-two favored either six or seven state groups almost equally.

REFERENCES

- [1] E. F. Arruda and M. D. Fragoso, “Standard dynamic programming applied to time aggregated Markov decision processes,” in *Proceedings of the IEEE Conference on Decision and Control*

- (CDC), Shanghai, China, December 15-18 2009, pp. 2576–2580. Available: <http://dx.doi.org/10.1109/CDC.2009.5400692>
- [2] R. W. Aldhaheri and H. K. Khalil, “Aggregation of the policy iteration method for nearly completely decomposable Markov chains,” *IEEE Transactions on Automatic Control*, vol. 36, no. 2, pp. 178–187, 1991. Available: <http://dx.doi.org/10.1109/9.67293>
 - [3] Z. Ren and B. H. Krogh, “Markov decision processes with fractional costs,” *IEEE Transactions on Automatic Control*, vol. 50, no. 5, pp. 646–650, 2005. Available: <http://dx.doi.org/10.1109/TAC.2005.846520>
 - [4] T. Sun, Q. Zhao, and P. B. Luh, “Incremental value iteration for time-aggregated Markov decision processes,” *IEEE Transactions on Automatic Control*, vol. 52, no. 11, pp. 2177–2182, 2007. Available: <http://dx.doi.org/10.1109/TAC.2007.908359>
 - [5] Q. S. Jia, “On state aggregation to approximate complex value functions in large-scale Markov decision processes,” *IEEE Transactions on Automatic Control*, vol. 56, no. 2, pp. 333–334, 2011. Available: <http://dx.doi.org/10.1109/TAC.2010.2052697>
 - [6] M. Aoki, “Some approximation methods for estimation and control of large scale systems,” *IEEE Transactions on Automatic Control*, vol. 23, no. 2, pp. 173–182, 1978. Available: <http://dx.doi.org/10.1109/TAC.1978.1101705>
 - [7] K. Deng, Y. Sun, P. G. Mehta, and S. P. Meyn, “An information-theoretic framework to aggregate a Markov chain,” in *Proceedings of the American Control Conference (ACC)*, St. Louis, MO, USA, June 10-12 2009, pp. 731–736. Available: <http://dx.doi.org/10.1109/ACC.2009.5160607>
 - [8] K. Deng, P. G. Mehta, and S. P. Meyn, “Optimal Kullback-Leibler aggregation via spectral theory of Markov chains,” *IEEE Transactions on Automatic Control*, vol. 56, no. 12, pp. 2793–2808, 2011. Available: <http://dx.doi.org/10.1109/TAC.2011.2141350>
 - [9] B. C. Geiger, T. Petrov, G. Kubin, and H. Koeppl, “Optimal Kullback-Leibler aggregation via information bottleneck,” *IEEE Transactions on Automatic Control*, vol. 60, no. 4, pp. 1010–1022, 2015. Available: <http://dx.doi.org/10.1109/TAC.2014.2364971>
 - [10] R. W. Aldhaheri and H. K. Khalil, “Aggregation and optimal control of nearly completely decomposable Markov chains,” in *Proceedings of the IEEE Conference on Decision and Control (CDC)*, Tampa, FL, USA, December 13-15 1989, pp. 1277–1282. Available: <http://dx.doi.org/10.1109/CDC.1989.70343>
 - [11] G. Kotsalis and M. Dahleh, “Model reduction of irreducible Markov chains,” in *Proceedings of the IEEE Conference on Decision and Control (CDC)*, Maui, HI, USA, December 9-12 2003, pp. 5727–5728. Available: <http://dx.doi.org/10.1109/CDC.2003.1271917>
 - [12] H. A. Simon and A. Ando, “Aggregation of variables in dynamic systems,” *Econometrica*, vol. 29, no. 2, pp. 111–138, 1961. Available: <http://dx.doi.org/10.2307/1909285>
 - [13] P. J. Courtois, “Error analysis in nearly-completely decomposable stochastic systems,” *Econometrica*, vol. 43, no. 4, pp. 691–709, 1975. Available: <http://dx.doi.org/10.2307/1913078>
 - [14] A. A. Pervozvanskii and I. N. Smirnov, “Stationary-state evaluation for a complex system with slowly varying couplings,” *Kibernetika*, vol. 3, no. 1, pp. 45–51, 1974. Available: <http://dx.doi.org/10.1007/BF01071538>
 - [15] V. G. Gaitsgori and A. A. Pervozvanskii, “Aggregation of states in a Markov chain with weak interaction,” *Kibernetika*, vol. 4, no. 1, pp. 91–98, 1975. Available: <http://dx.doi.org/10.1007/BF01069471>
 - [16] D. Teneketzis, S. H. Javid, and B. Sridhar, “Control of weakly-coupled Markov chains,” in *Proceedings of the IEEE Conference on Decision and Control (CDC)*, Albuquerque, NM, USA, December 10-12 1980, pp. 137–142. Available: <http://dx.doi.org/10.1109/CDC.1980.272033>
 - [17] F. Delebecque and J.-P. Quadrat, “Optimal control of Markov chains admitting strong and weak interactions,” *Automatica*, vol. 17, no. 2, pp. 281–296, 1981. Available: [http://dx.doi.org/10.1016/0005-1098\(81\)90047-9](http://dx.doi.org/10.1016/0005-1098(81)90047-9)
 - [18] J. C. Príncipe, *Information Theoretic Learning*. New York City, NY, USA: Springer-Verlag, 2010.
 - [19] R. L. Stratonovich, “On value of information,” *Izvestiya of USSR Academy of Sciences, Technical Cybernetics*, vol. 5, no. 1, pp. 3–12, 1965.
 - [20] R. L. Stratonovich and B. A. Grishanin, “Value of information when an estimated random variable is hidden,” *Izvestiya of USSR Academy of Sciences, Technical Cybernetics*, vol. 6, no. 1, pp. 3–15, 1966.
 - [21] I. J. Sledge and J. C. Príncipe, “Analysis of agent expertise in Ms. Pac-Man using value-of-information-based policies,” *IEEE Transactions on Computational Intelligence and Artificial Intelligence in Games*, 2018, (accepted, in press). Available: <http://dx.doi.org/10.1109/TG.2018.2808201>
 - [22] I. J. Sledge, M. S. Emigh, and J. C. Príncipe, “Guided policy exploration for Markov decision processes using an uncertainty-based value-of-information criterion,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 29, no. 6, pp. 2080–2098, 2018. Available: <http://dx.doi.org/10.1109/TNNLS.2018.2812709>
 - [23] I. J. Sledge and J. C. Príncipe, “An analysis of the value of information when exploring stochastic, discrete multi-armed bandits,” *Entropy*, vol. 20, no. 3, pp. 155(1–34), 2018. Available: <http://dx.doi.org/10.3390/e20030155>
 - [24] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY, USA: John Wiley and Sons, 2006.
 - [25] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*. Princeton, NJ, USA: Princeton University Press, 1953.
 - [26] N. Tishby, F. C. Pereira, and W. Bialek, “The information bottleneck method,” in *Proceedings of the Allerton Conference on Communication, Control and Computing*, Monticello, IL, USA, September 22-24 1999, pp. 368–377.
 - [27] M. D. Donsker and S. R. S. Varadhan, “Asymptotic evaluation of certain Markov process expectations for large time I.” *Communications of Pure and Applied Mathematics*, vol. 28, no. 1, pp. 1–47, 1975. Available: <http://dx.doi.org/10.1002/cpa.3160280102>
 - [28] ———, “Asymptotic evaluation of certain Markov process expectations for large time II.” *Communications of Pure and Applied Mathematics*, vol. 28, no. 2, pp. 279–301, 1975. Available: <http://dx.doi.org/10.1002/cpa.3160280206>
 - [29] I. J. Sledge and J. C. Príncipe, “Reduction of Markov chains using a value-of-information-based approach,” University of Florida, Gainesville, FL, USA, Computational NeuroEngineering Laboratory Report CNEL-2018-10-1, 2018. Available: <https://www.dropbox.com/s/rvgt1cwbz9o6zb/Sledge-UFTechRep-2018-1col.pdf?dl=0>
 - [30] I. J. Sledge, J. C. Bezdek, T. C. Havens, and J. M. Keller, “Relational generalizations of cluster validity indices,” *IEEE Transactions on Fuzzy Systems*, vol. 18, no. 4, pp. 771–786, 2010, (under review). Available: <http://dx.doi.org/10.1109/TFUZZ.2010.2048114>
 - [31] I. J. Sledge, T. C. Havens, J. C. Bezdek, and J. M. Keller, “Relational duals of cluster-validity functions for the c -means family,” *IEEE Transactions on Fuzzy Systems*, vol. 18, no. 6, pp. 1160–1170, 2010, (under review). Available: <http://dx.doi.org/10.1109/TFUZZ.2010.2079331>