

Clustering-based Mode Reduction for Markov Jump Systems

Zhe Du

Necmiye Ozay

Laura Balzano

*Department of Electrical Engineering and Computer Science
University of Michigan Ann Arbor, MI 48109, USA **

ZHEDU@UMICH.EDU

NECMIYE@UMICH.EDU

GIRASOLE@UMICH.EDU

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

While Markov jump systems (MJSs) are more appropriate than LTI systems in terms of modeling abruptly changing dynamics, MJSs (and other switched systems) may suffer from the model complexity brought by the potentially sheer number of switching modes. Much of the existing work on reducing switched systems focuses on the state space where techniques such as discretization and dimension reduction are performed, yet reducing mode complexity receives few attention. In this work, inspired by clustering techniques from unsupervised learning, we propose a reduction method for MJS such that a mode-reduced MJS can be constructed with guaranteed approximation performance. Furthermore, we show how this reduced MJS can be used in designing controllers for the original MJS to reduce the computation cost while maintaining guaranteed suboptimality.

Keywords: Markov Jump Systems, System Reduction, Clustering

1. Introduction

As the control and machine learning communities build tools to model ever more complex dynamical systems, it will become increasingly important to identify redundant aspects of a model and remove them using various unsupervised learning techniques. State dimensionality reduction has long been common in control systems, using PCA and similar techniques. In this paper we consider the setting where switched systems have redundant modes, and we apply clustering – the other most fundamental unsupervised learning technique – to remove redundancies.

Switched systems generalize time-invariant systems and can be used to model abrupt changes in the environment (e.g. weather and road surfaces), controlled plants (e.g. functioning statuses of different components), disturbances, or even control goals (e.g. cost functions in the optimal control). Switched system models are used in a variety of applications including controlling a Mars rover exploring an unknown heterogeneous terrain, solar power generation, investments in financial markets, and communications with packet losses [Blackmore et al. \(2005\)](#); [Cajueiro \(2002\)](#); [Loparo and Abdel-Malek \(1990\)](#); [Svensson et al. \(2008\)](#); [Ugrinovskii and Pota \(2005\)](#); [Sinopoli et al. \(2005\)](#); [Truong et al. \(2021\)](#). Using a pool of modes and allowing them to switch brings versatility but also new challenges: the number of modes can grow easily during modeling. For example, for controlled plants composed with multiple components, if we model each combination of health statuses, e.g. working and faulty, of all components as a mode, then the number of modes grow exponentially with the number of components. Even analysis such as stability verification

* The full version with extended proofs are provided in [Du et al. \(2022a\)](#).

can become computationally intractable when the number of components is large, thus finding a systematic and theoretically guaranteed way to reduce the number of modes in switched systems is imperative.

Existing work on (switched) system reduction mainly focuses on reducing the state dimension (Zhang et al., 2003) or constructing finite abstractions for the continuous state space Zamani and Abate (2014). Reducing the mode complexity, however, is still an uncharted field. A closely related line of research is the reduction of hidden Markov models (HMM) (Abate et al., 2014; Lun et al., 2018) and Markov chains (Zhang and Wang, 2019; Bittracher and Schütte, 2021), but HMM and Markov chain models are much simplified versions of switched dynamical models.

In this work, we study how one can perform mode reduction for Markov jump systems (MJS), a class of switched systems with the dynamics of mode switching governed by a Markov chain, and each individual mode characterized by an linear time-invariant (LTI) model. Our main contributions are the following:

- As an extension of Du et al. (2019), by treating the dynamics of each mode as features, we propose a clustering-based method that constructs a mode-reduced MJS.
- The reduced MJS is shown to provably well approximate the original MJS in terms of the trajectory difference.
- We show that LQR controllers designed with the reduced MJS can achieve guaranteed performance on the original MJS while significantly reducing the computational cost.

Our work adds a new dimension, i.e., reduction of modes, to the research of switched system reduction. This framework can be generalized to other problems such as stability analysis, robust and optimal control, invariance analysis, partially observed systems, etc. Other than constructing and analyzing the reduced MJS, the technical tools we develop in this work regarding perturbations can be applied to cases when there are model mismatches, e.g. system estimation errors incurred when dynamics are learned in identification or data-driven adaptive control as in Sattar et al. (2021).

2. Related Work

The work on reduction for stochastic switched systems can be roughly divided into three categories: bisimulation, symbolic abstraction, and order reduction.

Bisimulation: To evaluate the equivalency between two stochastic switched systems, notions of (approximate) probabilistic bisimulation are proposed in (Larsen and Skou, 1991; Desharnais et al., 2002, 2004). Approximation metrics from different perspectives (Abate, 2013) are developed to compare two systems, e.g. one(multi)-step transition kernels (Abate et al., 2011) and trajectories (Girard and Pappas, 2007; Tkachev and Abate, 2014; Julius and Pappas, 2009). Unlike existing work which typically defines the exact and approximate bisimulation on the state space, our work takes the mode space into consideration. Existing work on bisimulation is mainly conceptual and more for analysis purposes: they focus on developing bisimulation notions and approximation metrics, and “the majority of the examined approaches assume to be given two similar processes to compare, only a few put forward procedures for model approximation or abstraction with quantified quality” (Abate, 2013). This shortcoming is the work on abstractions address.

Symbolic Abstraction: Given a system with continuous state space, abstraction (Alur et al., 2000) considers discretizing the state space and then constructing a finite state symbolic model, which can be used as a surrogate for model verification (Clarke Jr et al., 2018; Kurshan, 2014) or controller synthesis (Maler et al., 1995). The work on abstraction for stochastic hybrid systems starts with the

autonomous cases. Under uniform discretization, [Abate et al. \(2010, 2011\)](#) provide approximation guarantees in terms of the discretization width. An adaptive partition scheme is proposed in [Soudjani and Abate \(2011\)](#), which mitigates the curse of dimensionality suffered by uniform sampling. Since the systems under consideration are autonomous, these work mainly serves verification purposes, but fall short toward controller synthesis goals. [Zamani and Abate \(2014\)](#) addresses this by allowing inputs in the systems. The idea of partitioning the continuous state space is similar to our work except that our partition is performed on the mode space, a.k.a. the discrete state space in hybrid systems, which provides a new yet closely related dimension to existing abstraction work.

Order Reduction: Another important line of research on system reduction is order reduction ([Gugercin and Antoulas, 2004](#)), where one seeks to reduce the dimension of the state space under certain criterion. With the help of linear matrix inequalities (LMIs), various methods have been applied for MJS, including \mathcal{H}_∞ reduction ([Zhang et al., 2003](#)), balanced truncation ([Kotsalis and Rantzer, 2010](#)), and \mathcal{H}_2 reduction ([Sun and Lam, 2016](#)), etc. Order reduction is also applied to models with adversarial conditions such as time-varying delay ([Zhang et al., 2015a](#)) and partial transition probability observation ([Zhang et al., 2015b](#); [Shen et al., 2019](#)).

3. Preliminaries and Problem Setup

For a matrix \mathbf{E} , $\mathbf{E}(i, :)$ denotes the i th row of \mathbf{E} , and $\mathbf{E}(i, j:k)$ denotes the i th row preserving only the j th to k th columns. For any index set A , $\mathbf{E}(i, A)$ denotes the i th row preserving columns given by A . Let $\sigma_i(\mathbf{E})$ ($\lambda_i(\mathbf{E})$) denote its i th largest singular (eigen) value. For any $s \in \mathbb{N}$, we let $[s] := \{1, 2, \dots, s\}$. We say $\Omega_{1:r} := \{\Omega_1, \dots, \Omega_r\}$ is a r -cluster partition of $[s]$ if $\bigcup_{i=1}^r \Omega_i = [s]$, $\Omega_i \cap \Omega_j = \emptyset$ for any $i \neq j$, and $\Omega_i \neq \emptyset$. We let $\Omega_{(i)}$ denote the cluster with i th largest cardinality. For a sequence of variables X_0, X_1, \dots, X_N , let $X_{0:N} := \{X_i\}_{i=0}^N$.

3.1. Preliminaries

In this work, we consider Markov jump systems (MJSs) with dynamics given by

$$\Sigma := \{\mathbf{x}_{t+1} = \mathbf{A}_{\omega_t} \mathbf{x}_t + \mathbf{B}_{\omega_t} \mathbf{u}_t, \quad \omega_t \sim \text{Markov Chain}(\mathbf{T})\} \quad (1)$$

where $\mathbf{x}_t \in \mathbb{R}^n$ and $\mathbf{u}_t \in \mathbb{R}^p$ denote the state and input at time t . The switching nature of the dynamics is characterized by s modes $\{\mathbf{A}_i, \mathbf{B}_i\}_{i=1}^s$ where $\mathbf{A}_i \in \mathbb{R}^{n \times n}$ and $\mathbf{B}_i \in \mathbb{R}^{n \times p}$ are state and input matrices for mode i . The active mode at time t is indexed by $\omega_t \in [s]$, and the mode switching sequence $\omega_{0:t}$ follows a Markov chain with Markov matrix $\mathbf{T} \in \mathbb{R}^{s \times s}$, i.e., $\mathbb{P}(\omega_{t+1} = j \mid \omega_t = i) = \mathbf{T}(i, j)$. We assume the Markov chain \mathbf{T} is ergodic. By properties of ergodicity, \mathbf{T} has a unique stationary distribution $\boldsymbol{\pi} \in \mathbb{R}^s$, and we let π_{\max} and π_{\min} denote the largest and smallest element in $\boldsymbol{\pi}$. In the remaining of the paper, we use $\Sigma := \text{MJS}(\mathbf{A}_{1:s}, \mathbf{B}_{1:s}, \mathbf{T})$ to denote the groundtruth MJS in (1) that we want to study, and similarly use notation $\text{MJS}(\cdot, \cdot, \cdot)$ to parameterize any MJS with expressions given in (1). We introduce the following two special types of Markov chain, which are closely tied to the main focus of this work.

Definition 1 (Lumpability and Aggregatability ([Buchholz, 1994](#))) *Markov matrix \mathbf{T} is lumpable w.r.t. partition $\Omega_{1:r}$ on $[s]$ if for any $k, l \in [r]$, and $i, i' \in \Omega_k$, $\sum_{j \in \Omega_l} \mathbf{T}(i, j) = \sum_{j \in \Omega_l} \mathbf{T}(i', j)$. As a special case, it is further aggregatable if $\mathbf{T}(i, :) = \mathbf{T}(i', :)$.*

Lumpability of a Markov chain coincides with the definition of probabilistic bisimulation in [Desharnais et al. \(2002\)](#), which describes an equivalence relation on $[s]$, i.e., two members are equivalent if they belong to the same cluster. For a Markov chain \mathbf{T} that is lumpable with respect to partition $\Omega_{1:r}$, we use $\zeta_t \in [r]$ to index the active cluster at time t , i.e., $\zeta_t = i$ if and only if $\omega_t \in \Omega_i$, and use $\zeta_{0:t}$ to denote the active cluster switching sequence.

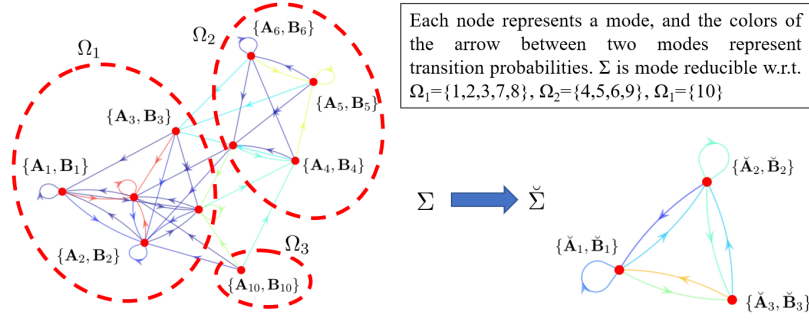


Figure 1: Illustration of reduction under mode-reducibility condition.

3.2. Mode-reducibility Conditions

We first define an equivalence relation between two MJSs with different number of modes via a surjection from each mode of the larger MJS to the smaller one, which is characterized by a partition. This extends the bijection idea in [Julius and Pappas \(2009\)](#); [Zhang et al. \(2003\)](#), which can only be used to compare two switched systems with the same number of modes.

Definition 2 (Equivalence between MJSs) Consider two MJSs Σ_1 and Σ_2 with the same state and input dimensions n, p , but different number of modes s_1 and s_2 respectively. WLOG, assume $s_1 > s_2$. Let $\{\mathbf{x}_t^{(1)}, \mathbf{u}_t^{(1)}, \omega_t^{(1)}\}$ and $\{\mathbf{x}_t^{(2)}, \mathbf{u}_t^{(2)}, \omega_t^{(2)}\}$ denote their respective state, input, and mode index. Σ_1 and Σ_2 are equivalent if there exists a partition $\Omega_{1:s_2}$ on $[s_1]$ such that Σ_1 and Σ_2 have the same transition kernels, i.e., for any t , any $k, k' \in [s_2]$, any $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^n$, and any $\mathbf{u} \in \mathbb{R}^p$

$$\begin{aligned} & \mathbb{P}\left(\omega_{t+1}^{(1)} \in \Omega_{k'}, \mathbf{x}_{t+1}^{(1)} = \mathbf{x}' \mid \omega_t^{(1)} \in \Omega_k, \mathbf{x}_t^{(1)} = \mathbf{x}, \mathbf{u}_t^{(1)} = \mathbf{u}\right) \\ &= \mathbb{P}\left(\omega_{t+1}^{(2)} = k', \mathbf{x}_{t+1}^{(2)} = \mathbf{x}' \mid \omega_t^{(2)} = k, \mathbf{x}_t^{(2)} = \mathbf{x}, \mathbf{u}_t^{(2)} = \mathbf{u}\right). \quad (2) \end{aligned}$$

If one views the discrete mode and the continuous state $\{\omega_t^{(1)}, \mathbf{x}_t^{(1)}\}$ as a *hybrid state* ([Abate et al., 2011](#)), then Definition 2 generalizes Definition 1 from Markov chains to MJSs. A sufficient condition in terms of the dynamics parameters, which guarantees that an MJS can be reduced to a smaller MJS with equivalency between them, is given by the following.

Definition 3 (Mode-reducibility Condition) Σ is mode-reducible with respect to a partition $\Omega_{1:r}$, if its Markov chain \mathbf{T} is lumpable with respect to $\Omega_{1:r}$, and modes within the same cluster have the same dynamics, i.e., for any $k \in [r]$, any $i, i' \in \Omega_k$, we have $\mathbf{A}_i = \mathbf{A}_{i'}$, $\mathbf{B}_i = \mathbf{B}_{i'}$.

If this condition holds for Σ , we construct a mode-reduced MJS given by $\check{\Sigma} := \text{MJS}(\check{\mathbf{A}}_{1:r}, \check{\mathbf{B}}_{1:r}, \check{\mathbf{T}})$ such that for any $k, l \in [r]$, any $i \in \Omega_k$, $\check{\mathbf{A}}_k = \mathbf{A}_i$ and $\check{\mathbf{B}}_k = \mathbf{B}_i$, $\check{\mathbf{T}} \in \mathbb{R}^{r \times r}$ with $\check{\mathbf{T}}(k, l) = \sum_{j \in \Omega_l} \mathbf{T}(i, j)$, which is illustrated in Fig. 1. Let $\{\check{\mathbf{x}}_t, \check{\mathbf{u}}_t, \check{\omega}_t\}$ denote the state, input, and mode index for the reduced $\check{\Sigma}$. Then, the following result shows that $\check{\Sigma}$ and Σ are equivalent according to Definition 2. This is more of a teaser result for the ideal mode-reducible case, and formal problems regarding practical cases when Σ is not mode-reducible will be formulated later.

Proposition 4 Suppose Σ is mode-reducible and $\check{\Sigma}$ is constructed as above. Consider the case when the two MJSs have (i) initial mode distributions satisfy $\mathbb{P}(\omega_0 \in \Omega_k) = \mathbb{P}(\check{\omega}_0 = k)$ for all $k \in [r]$, (ii) the same initial states ($\mathbf{x}_0 = \check{\mathbf{x}}_0$), and (iii) the same input sequences ($\mathbf{u}_{0:t-1} = \check{\mathbf{u}}_{0:t-1}$). Then, these two MJSs have the same mode and state transition kernels, i.e., $\mathbb{P}(\omega_t \in \Omega_k, \mathbf{x}_t = \mathbf{x}) = \mathbb{P}(\check{\omega}_t = k, \check{\mathbf{x}}_t = \mathbf{x})$ for all t , all $k \in [r]$ and $\mathbf{x} \in \mathbb{R}^n$. Particularly, there exists a special type of reduced $\check{\Sigma}$ such that the modes are synchronized: for all t , $\check{\omega}_t = \zeta_t$. In this case, $\check{\mathbf{x}}_t = \mathbf{x}_t$ for all t .

Proposition 4 first shows the equivalency between Σ and $\check{\Sigma}$ for the transition kernels, then if certain synchrony, exists between $\zeta_{0:t}$ and $\check{\omega}_{0:t}$, the equivalency also exists for the trajectories. The condition $\check{\omega}_t = \zeta_t$ in Proposition 4 essentially establishes a coupling between the Markov chains $\omega_{0:t}$ and $\check{\omega}_{0:t}$ such that $\mathbb{P}(\omega_t \in \Omega_k, \check{\omega}_t = k) = \mathbb{P}(\omega_t \in \Omega_k) = \mathbb{P}(\check{\omega}_t = k)$. Establishing coupling between the stochastic systems usually allows for stronger equivalency and approximation metrics result. Similar coupling scheme is implicitly used in Julius and Pappas (2009); Zhang et al. (2003); an optimal coupling by minimizing Wasserstein distance is discussed in Tkachev and Abate (2014); and a weaker coupling using the idea of HMM is discussed in Shen et al. (2019).

In Definition 3 and Proposition 4, one can view $\{\omega_t, \mathbf{x}_t\} \in [s] \times \mathbb{R}^n$ as a hybrid state, then \mathbf{T} being lumpable guarantees the existence of an equivalence relation in the discrete domain $[s]$, while state/input matrices being the same guarantees this in the continuous domain \mathbb{R}^n .

3.3. Problem Formulation

The mode-reducibility condition in Definition 3 and Proposition 4 provide principles under which one could construct a reduced MJS that is equivalent to the original one. Though seemingly strong, the mode-reducibility condition may hold *approximately* in many practical scenarios that allow for mode reduction. For example, in hybrid power systems, solar panels that are geographically close tend to have similar levels of solar radiation thus similar power dynamics. Hence, we would like to study the following: suppose the mode-reducibility condition holds *approximately*, how can we construct a reduced MJS? From here, we formulate the following two concrete problems.

Problem P1 (Lumpable Case) *Assume the dynamics of $\Sigma = \text{MJS}(\mathbf{A}_{1:s}, \mathbf{B}_{1:s}, \mathbf{T})$ are known. Suppose there exists a partition $\Omega_{1:r}$ on $[s]$ such that: for any Ω_k, Ω_l and any $i, i' \in \Omega_k$,*

$$\|\mathbf{A}_i - \mathbf{A}_{i'}\|_F \leq \epsilon_{\mathbf{A}}, \quad \|\mathbf{B}_i - \mathbf{B}_{i'}\|_F \leq \epsilon_{\mathbf{B}}, \quad (3)$$

$$\frac{1}{2} \sum_{l \in [r]} \left| \sum_{j \in \Omega_l} \mathbf{T}(i, j) - \sum_{j \in \Omega_l} \mathbf{T}(i', j) \right| \leq \epsilon_{\mathbf{T}}. \quad (4)$$

Then, we seek to estimate the partition $\Omega_{1:r}$, construct a reduced MJS, and provide guarantees on the behavior difference incurred by the reduction.

Throughout this work, we will refer to $\epsilon_{\mathbf{A}}$, $\epsilon_{\mathbf{B}}$, and $\epsilon_{\mathbf{T}}$ as *perturbations* (that lead to violation of mode-reducibility condition in Definition 3 and prevents the existence of $\check{\Sigma}$ as in Proposition 4). For the special aggregatable case, we separately formulate a similar problem in Problem P2. Even though solutions to P1 automatically solve P2, as we shall see in the next section, algorithm theoretical guarantees for P2 require milder and more interpretable assumptions than P1.

Problem P2 (Aggregatable Case) *In Problem P1, replace (4) with $\|\mathbf{T}(i, :)^T - \mathbf{T}(i', :)^T\|_1 \leq \epsilon_{\mathbf{T}}$.*

In P2, $\mathbf{A}_i, \mathbf{B}_i$, and $\mathbf{T}(i, :)$ can be thought of as features and provide the position for mode i in the feature space, and pairwise distances between all the modes reflect the partition $\Omega_{1:r}$. The recovery of the partition $\Omega_{1:r}$ resembles the classical clustering problems studied in the machine learning community. However, this resemblance does not directly apply to the weak lumpability condition (4) in P1, as the similarity in the transition probabilities is encoded by the partition $\Omega_{1:r}$. This makes the already non-convex clustering problem even more challenging.

In the aforementioned hybrid power system example where the geographical closeness serves as expert knowledge, one may be tempted to naively attempt construction of a small scale system from the beginning, but this makes it impossible to know how well it approximates the true dynamics.

Our approach would be to instead approximate the dynamics as well as possible with many modes and then use clustering to reduce the overall system. Hence, other than constructing a reduced MJS, we also seek to provide guarantees on the behavior difference and performance degradation incurred by the reduction. These guarantees play a critical role in safety-critical environments.

4. Clustering-based Mode Reduction for MJS

Algorithm 1: System Reduction for MJS

Input: $\mathbf{A}_{1:s}, \mathbf{B}_{1:s}, \mathbf{T}, r, \boldsymbol{\pi}, \alpha_{\mathbf{A}}, \alpha_{\mathbf{B}}, \alpha_{\mathbf{T}}$

- 1 Construct feature matrix Φ :
- 2 **case Problem P2 do**
- 3 $\Phi(i, :) = [\alpha_{\mathbf{A}} \mathbf{vec}(\mathbf{A}_i)^\top, \alpha_{\mathbf{B}} \mathbf{vec}(\mathbf{B}_i)^\top, \alpha_{\mathbf{T}} \mathbf{T}(i, :)], \forall i \in [s]$
- 4 **case Problem P1 do**
- 5 $\mathbf{H} = \mathbf{diag}(\boldsymbol{\pi})^{\frac{1}{2}} \mathbf{T} \mathbf{diag}(\boldsymbol{\pi})^{-\frac{1}{2}}$
- 6 $\mathbf{W}_r \leftarrow$ top r left singular vectors of \mathbf{H}
- 7 $\mathbf{S}_r = \mathbf{diag}(\boldsymbol{\pi})^{-\frac{1}{2}} \mathbf{W}_r$
- 8 $\Phi(i, :) = [\alpha_{\mathbf{A}} \mathbf{vec}(\mathbf{A}_i)^\top, \alpha_{\mathbf{B}} \mathbf{vec}(\mathbf{B}_i)^\top, \alpha_{\mathbf{T}} \mathbf{S}_r(i, :)], \forall i \in [s]$
- 9 $\mathbf{U}_r \leftarrow$ top r left singular vectors of Φ
- 10 Solve k-means problem: $\hat{\Omega}_{1:r}, \hat{\mathbf{c}}_{1:r} = \arg \min_{\hat{\Omega}_{1:r}, \hat{\mathbf{c}}_{1:r}} \sum_{k \in [r]} \sum_{i \in \hat{\Omega}_k} \|\mathbf{U}_r(i, :) - \hat{\mathbf{c}}_k\|^2$
- 11 Construct $\hat{\Sigma}, \forall k, l \in [r]$
- $\hat{\mathbf{A}}_k = \frac{1}{|\hat{\Omega}_k|} \sum_{i \in \hat{\Omega}_k} \mathbf{A}_i, \quad \hat{\mathbf{B}}_k = \frac{1}{|\hat{\Omega}_k|} \sum_{i \in \hat{\Omega}_k} \mathbf{B}_i, \quad \hat{\mathbf{T}}(k, l) = \frac{1}{|\hat{\Omega}_k|} \sum_{i \in \hat{\Omega}_k, j \in \hat{\Omega}_l} \mathbf{T}(i, j)$

Output: $\hat{\Sigma} : \text{MJS}(\hat{\mathbf{A}}_{1:r}, \hat{\mathbf{B}}_{1:r}, \hat{\mathbf{T}})$

The clustering-based MJS reduction method is presented in Algorithm 1. We treat the estimation of partition $\Omega_{1:r}$ essentially as a mode clustering problem with dynamics matrices $\mathbf{A}_i, \mathbf{B}_i$ and transition distribution $\mathbf{T}(i, :)$ serving as features for mode i . We first construct feature matrix Φ from Line 2 to Line 8, with $\Phi(i, :)$ denoting the features of mode i . For the aggregatable case in Problem P2, we simply stack the vectorized $\mathbf{A}_i, \mathbf{B}_i$ and $\mathbf{T}(i, :)$, and use $\alpha_{\mathbf{A}}, \alpha_{\mathbf{B}}, \alpha_{\mathbf{T}}$ to denote their weights respectively. One way to choose these weights is to use a normalization, e.g. $\alpha_{\mathbf{A}} = \frac{1}{\max_i \|\mathbf{A}_i\|}$, so that these three features would have the same scales. By the definition of aggregatibility, similarities of $\mathbf{T}(i, :)$ among different modes are direct indicators for the groundtruth partition $\Omega_{1:r}$. For the lumpable case, however, this is not an option since two modes belonging to the same cluster can still have different transition probabilities $\mathbf{T}(i, :)$ even if $\epsilon_{\mathbf{T}} = 0$. According to (4), the groundtruth partition $\Omega_{1:r}$ is only embodied in the mode-to-cluster transition probabilities $\sum_{j \in \Omega_l} \mathbf{T}(i, j)$ constructed using the groundtruth partition itself. This leaves us in a ‘‘chicken-and-egg’’ dilemma. To deal with this, from Line 5 to 8, we compute the first r left singular vectors \mathbf{W}_r of matrix $\mathbf{diag}(\boldsymbol{\pi})^{\frac{1}{2}} \mathbf{T} \mathbf{diag}(\boldsymbol{\pi})^{-\frac{1}{2}}$, and then weight it by $\mathbf{diag}(\boldsymbol{\pi})^{-\frac{1}{2}}$ to obtain matrix $\mathbf{S}_r \in \mathbb{R}^{s \times r}$, which is used to construct features in Φ for the lumpable case. We will later justify using \mathbf{S}_r as features by showing row similarities in \mathbf{S}_r reflect the partition under certain assumptions.

With the feature matrix Φ , to recover the partition, we resort to k-means: in Line 10, k-means is applied to the first r left singular vector \mathbf{U}_r of Φ . Extracting the low-rank structure \mathbf{U}_r de-noises the impact of the perturbations. Based on the k-means solution $\hat{\Omega}_{1:r}$, we construct the reduced $\hat{\Sigma}$ by averaging modes within the same estimated cluster. When one picks $\alpha_{\mathbf{A}} = \alpha_{\mathbf{B}} = 0$, i.e. only the Markov matrix \mathbf{T} is used to cluster the modes, our clustering scheme under the aggregatable case P2 is equivalent to the one in Zhang and Wang (2019). The lumpable case P1, on the other hand, is based on the preliminary analysis in Meilă and Shi (2001).

4.1. Theoretical Guarantees for Clustering

We assume a $(1 + \epsilon)$ solution (Lei et al., 2015) to the k-means problem in Algorithm 1 can be obtained, i.e., $\sum_{k \in [r], i \in \hat{\Omega}_k} \|\mathbf{U}_r(i, :) - \hat{\mathbf{c}}_k\|^2 \leq (1 + \epsilon) \min_{\Omega'_{1:r}, \mathbf{c}'_{1:r}} \sum_{k \in [r], i \in \Omega'_k} \|\mathbf{U}_r(i, :) - \mathbf{c}'_k\|^2$. We later show how ϵ affects the overall clustering performance. To evaluate the performance of partition estimation, we define misclustering rate (MR) as $\text{MR}(\hat{\Omega}_{1:r}) = \min_{h \in \mathcal{H}} \sum_{k \in [r]} \frac{|\{i: i \in \Omega_k, i \notin \hat{\Omega}_{h(k)}\}|}{|\Omega_k|}$, where \mathcal{H} is the set of all bijections from $[r]$ to $[r]$ so that the comparison finds the best cluster label matching. We define the following averaged feature matrix $\bar{\Phi}$ based on the underlying partition $\Omega_{1:r}$: for all $i \in [s]$ (suppose $i \in \Omega_k$ for some $k \in [r]$), $\bar{\Phi}(i, :) = \frac{1}{|\Omega_k|} \sum_{i' \in \Omega_k} \Phi(i', :)$. By construction, there are up to r unique rows in $\bar{\Phi}$, hence $\text{rank}(\bar{\Phi}) \leq r$. We first present the clustering guarantee for Problem P2, i.e., the aggregatable case.

Theorem 5 Consider Problem P2 and Algorithm 1. Suppose $\hat{\Omega}_{1:r}$ is a $(1 + \epsilon)$ k-means solution. Let $\epsilon_{\text{Agg}} := \alpha_{\mathbf{A}} \epsilon_{\mathbf{A}} + \alpha_{\mathbf{B}} \epsilon_{\mathbf{B}} + \alpha_{\mathbf{T}} \epsilon_{\mathbf{T}}$. Then, if $\text{rank}(\bar{\Phi}) = r$ and $\epsilon_{\text{Agg}} \leq \frac{\sigma_r(\bar{\Phi}) \sqrt{|\Omega_{(r)}| + |\Omega_{(1)}|}}{8\sqrt{s(2+\epsilon)|\Omega_{(1)}|}}$, we have

$$\text{MR}(\hat{\Omega}_{1:r}) \leq \frac{64(2 + \epsilon)s}{\sigma_r^2(\bar{\Phi})} \epsilon_{\text{Agg}}^2. \quad (5)$$

Additionally, if $\epsilon_{\text{Agg}} \leq \frac{\sigma_r(\bar{\Phi})}{8\sqrt{s(2+\epsilon)|\Omega_{(1)}|}}$, then $\text{MR}(\hat{\Omega}_{1:r}) = 0$.

We see the misclustering rate can be bounded by the inner-cluster spread ϵ_{Agg} when it is small enough. As for the inter-cluster distance, i.e., dissimilarity of $\mathbf{A}_i, \mathbf{B}_i$ and $\mathbf{T}(i, :)$ of modes in different clusters, it affects the guarantee through the term $\sigma_r(\bar{\Phi})$. This is because when two modes belonging different clusters have similar features, their rows in the averaged feature matrix $\bar{\Phi}$ will also be similar, which could lead to small $\sigma_r(\bar{\Phi})$ and larger error bound.

The clustering guarantee for the lumpable case in Problem P1 is more involved than the aggregatable case. We first provide a few more background notions and definitions that can help the exposition. We say a Markov matrix \mathbf{T} is *reversible* if there exists a distribution $\boldsymbol{\pi} \in \mathbb{R}^s$ such that $\boldsymbol{\pi}(i)\mathbf{T}(i, j) = \boldsymbol{\pi}(j)\mathbf{T}(j, i)$ for all $i, j \in [s]$. This condition for reversibility translates to $\text{diag}(\boldsymbol{\pi})\mathbf{T} = \mathbf{T}^\top \text{diag}(\boldsymbol{\pi})$ when \mathbf{T} is ergodic with stationary distribution $\boldsymbol{\pi}$. For a reversible Markov matrix that is also lumpable, we have the following property.

Lemma 6 (Appendix A in Meilă and Shi (2001)) For a reversible Markov matrix \mathbf{T} that is also lumpable with respect to partition $\Omega_{1:r}$, it is diagonalizable with real eigenvalues. Let $\mathbf{S} \in \mathbb{R}^{s \times s}$ denote an arbitrary eigenvector matrix of \mathbf{T} . Then, there exists an index set $\mathcal{A} \subseteq [s]$ with $|\mathcal{A}| = r$ such that for all $k \in [r]$, for all $i, i' \in \Omega_k$, we have $\mathbf{S}(i, \mathcal{A}) = \mathbf{S}(i', \mathcal{A})$.

We say \mathbf{T} in Lemma 6 has *informative spectrum* if $\mathcal{A} = [r]$ and $|\lambda_r(\mathbf{T})| > |\lambda_{r+1}(\mathbf{T})|$, which implies that the uniquely defined r leading eigenvalues have eigenvectors that carry partition information in Lemma 6. For lumpable Markov matrices, we define the $\epsilon_{\mathbf{T}}$ -neighborhood of \mathbf{T} :

$$\mathcal{L}(\mathbf{T}, \Omega_{1:r}, \epsilon_{\mathbf{T}}) := \left\{ \mathbf{T}_0 \in \mathbb{R}^{s \times s} : \mathbf{T}_0 \text{ is Markovian, } \forall k, l \in [r], \forall i \in \Omega_k, \sum_{j \in \Omega_l} \mathbf{T}_0(i, j) = \frac{1}{|\Omega_k|} \sum_{i' \in \Omega_k, j \in \Omega_l} \mathbf{T}(i', j), \|\mathbf{T}_0 - \mathbf{T}\|_\infty \leq \epsilon_{\mathbf{T}} \right\}. \quad (6)$$

Then we provide the clustering guarantee for the lumpable case.

Theorem 7 Consider Problem P1 and Algorithm 1. Let $\gamma_1 := \sum_{i=2}^s \frac{1}{1 - \lambda_i(\mathbf{T})}$, $\gamma_2 = \min\{\sigma_r(\mathbf{H}) - \sigma_{r+1}(\mathbf{H}), 1\}$, $\gamma_3 = \frac{16\gamma_1 \sqrt{\pi_{\max}} \|\mathbf{T}\|_{\text{F}}}{\gamma_2 \pi_{\min}^2}$, and $\epsilon_{\text{Lmp}} := \alpha_{\mathbf{A}} \epsilon_{\mathbf{A}} + \alpha_{\mathbf{B}} \epsilon_{\mathbf{B}} + \gamma_3 \alpha_{\mathbf{T}} \epsilon_{\mathbf{T}}$. Assume there exists an

ergodic and reversible $\mathbf{T}_0 \in \mathcal{L}(\mathbf{T}, \Omega_{1:r}, \epsilon_{\mathbf{T}})$ with informative spectrum. Suppose $\hat{\Omega}_{1:r}$ is a $(1 + \epsilon)$ k -means solution. Then, if $\text{rank}(\bar{\Phi})=r$, $\epsilon_{\mathbf{T}} \leq \frac{\pi_{\min}}{\gamma_1}$, and $\epsilon_{Lmp} \leq \frac{\sigma_r(\bar{\Phi})\sqrt{|\Omega_{(r)}|+|\Omega_{(1)}|}}{8\sqrt{s(2+\epsilon)|\Omega_{(1)}|}}$, we have

$$\text{MR}(\hat{\Omega}_{1:r}) \leq \frac{64(2 + \epsilon)s}{\sigma_r^2(\bar{\Phi})} \epsilon_{Lmp}^2. \quad (7)$$

Additionally, if $\epsilon_{Lmp} \leq \frac{\sigma_r(\bar{\Phi})}{8\sqrt{s(2+\epsilon)|\Omega_{(1)}|}}$, then $\text{MR}(\hat{\Omega}_{1:r}) = 0$.

The difference of using \mathbf{T} and \mathbf{S}_r to construct features for lumpable and aggregatable cases is also reflected in the comparison between Theorem 5 and 7. They are almost identical with an additional γ_3 term. This term describes how much the lumpability perturbation $\epsilon_{\mathbf{T}}$ on \mathbf{T} affects the row equalities of its spectrum-related matrix \mathbf{S}_r in Lemma 6. The assumption on the existence of \mathbf{T}_0 with informative spectrum guarantees (i) the partition $\Omega_{1:r}$ information is carried by the leading eigenvectors of \mathbf{T}_0 as introduced in Lemma 6, and (ii) this information can still be preserved in \mathbf{S}_r as long as \mathbf{T} is close to \mathbf{T}_0 . Saying this, the theory for clustering in the lumpable case may not hold for an arbitrary lumpable \mathbf{T} , but only those close to Markov matrices with informative spectra.

5. Approximation Guarantees

With perturbation $\epsilon_{\mathbf{A}}$, $\epsilon_{\mathbf{B}}$, $\epsilon_{\mathbf{T}}$, the reduced $\hat{\Sigma}$ may not be equivalent to the original Σ as in Proposition 4. In this case, if certain approximation guarantees can be established, they can serve verification purposes such as safety (Julius and Pappas, 2009) and invariance (Soudjani and Abate, 2011) evaluations. In this section, we show that when Σ is autonomous and stable, the reduced system $\hat{\Sigma}$ can provably well approximate the original system Σ in terms of the trajectory difference $\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|$. Note that autonomous MJSs also include MJSs under mode-dependent state-feedback controller $\mathbf{K}_{1:s}$ such that the closed-loop dynamics $\mathbf{x}_{t+1} = (\mathbf{A}_{\omega_t} + \mathbf{B}_{\omega_t}\mathbf{K}_{\omega_t})\mathbf{x}_t$ is autonomous. Autonomous system reduction is also considered in Abate et al. (2011); Tkachev and Abate (2014); Bian and Abate (2017). Since we have shown in Theorem 5 and 7 that $\text{MR}(\hat{\Omega}_{1:r}) = 0$ when perturbation $\epsilon_{\mathbf{A}}$, $\epsilon_{\mathbf{B}}$, $\epsilon_{\mathbf{T}}$ are small, in the following, we assume $\Omega_{1:r} = \hat{\Omega}_{1:r}$ for simplicity.

The stability under consideration is the mean-square stability (MSS) that is typically studied in MJS Costa et al. (2006). We say Σ is mean-square stable (MSS), if there exists Σ_{∞} such that $\lim_{t \rightarrow \infty} \mathbb{E}[\mathbf{x}_t \mathbf{x}_t^T] = \Sigma_{\infty}$. Define the augmented state matrix $\mathcal{A} \in \mathbb{R}^{sn^2 \times sn^2}$ with its ij -th $n^2 \times n^2$ block given by $[\mathcal{A}]_{ij} := \mathbf{T}(j, i) \cdot \mathbf{A}_j \otimes \mathbf{A}_j$, and let ρ denote the spectral radius of \mathcal{A} . Then, under the autonomous case, i.e., $\mathbf{u}_t = 0$, Σ being MSS is equivalent to $\rho < 1$. Furthermore, we define $\tau := \sup_{k \in \mathbb{N}} \|\mathcal{A}^k\|/\rho^k$, which measures how fast $\|\mathcal{A}^k\|^{1/k}$ converges to ρ . We let $\bar{A} := \max_i \|\mathbf{A}_i\|$ and $\rho_0 := \frac{1+\rho}{2}$. Then we have the following bounds on the trajectory difference.

Theorem 8 Consider Σ that is autonomous and MSS. Assume $\hat{\Omega}_{1:r} = \Omega_{1:r}$ in Algorithm 1. Suppose Σ and $\hat{\Sigma}$ have the same initial states, i.e., $\mathbf{x}_0 = \hat{\mathbf{x}}_0$. Mode $\hat{\omega}_t$ of $\hat{\Sigma}$ is synchronous to ω_t of Σ , i.e., for all t , if $\omega_t \in \Omega_k$ then $\hat{\omega}_t = k$. Then, when $\epsilon_{\mathbf{A}} \leq \min\{\bar{A}, \frac{1-\rho_0}{6\sqrt{s}\tau\bar{A}}\}$, we have,

$$\mathbb{E}[\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|^2] \leq 12nst\rho_0^{t-1}\tau^2\bar{A}\|\mathbf{x}_0\|^2\epsilon_{\mathbf{A}}. \quad (8)$$

We require the synchrony between the mode sequences $\omega_{0:t-1}$ and $\hat{\omega}_{0:t-1}$ so that evaluating $\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|$ is meaningful as we discussed in Section 3.2. When ω_t is observed while \mathbf{x}_t is not, by forcing mode synchrony, we can estimate \mathbf{x}_t with $\hat{\mathbf{x}}_t$. Since $\rho < 1$ due to the MSS of Σ , we know $\rho_0 < 1$, which implies that the trajectory difference $\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|$ converges to 0 exponentially with t . Note that this also implies that $\hat{\Sigma}$ is MSS, as otherwise the difference would not converge. Proposition 4 provides a the sanity check for Theorem 8 that when $\epsilon_{\mathbf{A}} = 0$, we have $\mathbf{x}_t = \hat{\mathbf{x}}_t$.

6. Controller Design with Case Study on LQR

A typical class of controllers for MJS is known to be *mode-dependent*. A mode-dependent controller is essentially a collection of individual controllers such that each mode is associated with one, and the deployed controller switches with corresponding modes. With the reduced MJS $\hat{\Sigma}$, we can design mode-dependent controller $\hat{\mathbf{K}}_{1:r}$ for $\hat{\Sigma}$ and then associate every mode i in Σ with $\hat{\mathbf{K}}_k$ if $i \in \hat{\Omega}_k$. Since $\hat{\Sigma}$ has a smaller scale than Σ , the computation cost may be reduced. In the following, with the classical infinite time linear quadratic regulator (LQR) problem as an example, we showcase the implementation of this idea and provide its suboptimality guarantees.

In the infinite LQR problems, given positive definite cost matrices \mathbf{Q} and \mathbf{R} , we define quadratic cost $J_T = \mathbb{E} \left[\sum_{t=0}^{T-1} (\mathbf{x}_t^\top \mathbf{Q} \mathbf{x}_t + \mathbf{u}_t^\top \mathbf{R} \mathbf{u}_t) + \mathbf{x}_T^\top \mathbf{Q} \mathbf{x}_T \right]$. The goal is to design inputs to minimize the infinite time average cost $\limsup_{T \rightarrow \infty} \frac{1}{T} J_T$ under Σ . The optimal solution is well studied in [Costa et al. \(2006\)](#). Mode-dependent controllers $\mathbf{K}_{1:s}$ is obtained by solving s coupled discrete algebraic Riccati equations each of which is parameterized by $\mathbf{A}_i, \mathbf{B}_i, \mathbf{T}(i, \cdot), \mathbf{Q}, \mathbf{R}$ and resembles the Riccati equation used in LQR for LTI systems. Then, at time t , the optimal input is given by $\mathbf{u}_t = \mathbf{K}_{\omega_t} \mathbf{x}_t$.

To design controllers with the reduced $\hat{\Sigma}$, we can first compute controller $\hat{\mathbf{K}}_{1:r}$ by solving LQR problem with $\hat{\Sigma}$ as the dynamics. This requires solving r coupled Riccati equations each of which is parameterized by $\hat{\mathbf{A}}_i, \hat{\mathbf{B}}_i, \hat{\mathbf{T}}(i, \cdot), \mathbf{Q}, \mathbf{R}$. To apply $\hat{\mathbf{K}}_{1:r}$ to the original Σ , we simply let $\mathbf{u}_t = \hat{\mathbf{K}}_k \mathbf{x}_t$ if $\omega_t = k$. Since the number of coupled Riccati equations is the same as the number of modes, the computation cost for Σ is $\mathcal{O}(s)$ while only $\mathcal{O}(r)$ for $\hat{\Sigma}$, thus the saving is prominent when $r \ll s$. Additionally, we have the following suboptimality guarantees.

Theorem 9 *Assume system Σ is mean square stabilizable and has additive Gaussian process noise $\mathcal{N}(0, \sigma_w^2 \mathbf{I}_n)$ that is independent of the mode switching. Assume $\hat{\Omega}_{1:r} = \Omega_{1:r}$ in Algorithm 1. Let J^* and \hat{J} respectively denote the infinite time average cost incurred by the optimal controller $\mathbf{K}_{1:s}$ and controller $\hat{\mathbf{K}}_{1:r}$ (at time t , $\mathbf{u}_t = \hat{\mathbf{K}}_k \mathbf{x}_t$ if $\omega_t \in \hat{\Omega}_k$). Then, there exists constant $\bar{\epsilon}_{\mathbf{A}, \mathbf{B}}, \bar{\epsilon}_{\mathbf{T}}, C_{\mathbf{A}, \mathbf{B}}$, and $C_{\mathbf{T}}$, such that when $\max\{\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}\} \leq \bar{\epsilon}_{\mathbf{A}, \mathbf{B}}$ and $\epsilon_{\mathbf{T}} \leq \bar{\epsilon}_{\mathbf{T}}$, we have $\hat{J} - J^* \leq \sigma_w^2 (C_{\mathbf{A}, \mathbf{B}} \max\{\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}\} + C_{\mathbf{T}} \epsilon_{\mathbf{T}})^2$.*

Having additive noise means the MJS dynamics in (1) becomes $\mathbf{x}_{t+1} = \mathbf{A}_{\omega_t} \mathbf{x}_t + \mathbf{B}_{\omega_t} \mathbf{u}_t + \mathbf{w}_t$ where $\mathbf{w}_t \sim \mathcal{N}(0, \sigma_w^2 \mathbf{I}_n)$. Constants $\bar{\epsilon}_{\mathbf{A}, \mathbf{B}}, \bar{\epsilon}_{\mathbf{T}}, C_{\mathbf{A}, \mathbf{B}}$, and $C_{\mathbf{T}}$ only depend on the original Σ and cost matrices \mathbf{Q} and \mathbf{R} with the exact expressions provided in [Du et al. \(2022b\)](#).

7. Numerical Experiments

In this section, we present experiment results to evaluate the main results in the paper. System Σ is randomly generated with desired levels of perturbation $\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}, \epsilon_{\mathbf{T}}$ under the uniform partition $\Omega_{1:r}$ on $[s]$, i.e., $|\Omega_i| = |\Omega_j|$ for any $i, j \in [r]$. We first evaluate the clustering performance of Algorithm 1. Then, we evaluate the performance of LQR controller designed with the reduced MJS $\hat{\Sigma}$.

To evaluate Algorithm 1, we fix $n = 5, p = 3, r = 4$ and record the average clustering error $\text{CE} = \min_{h \in \mathcal{H}} \sum_{k \in [r]} |\{i : i \in \Omega_k, i \notin \hat{\Omega}_{h(k)}\}|$, i.e., the number of misclustered modes under the best cluster labeling, over 100 runs. Figure 2 presents the clustering performances under different number of modes s and perturbation $\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}$ and $\epsilon_{\mathbf{T}}$. In Figure 2(a), we set perturbation $\epsilon_{\mathbf{A}} = \epsilon_{\mathbf{B}}$ and hyper-parameters $\alpha_{\mathbf{A}} = \frac{1}{\max_i \|\mathbf{A}_i\|}, \alpha_{\mathbf{B}} = \frac{1}{\max_i \|\mathbf{B}_i\|}$, and $\alpha_{\mathbf{T}} = 0$ in Algorithm 1. The clustering error goes up with increasing s and $\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}$, and we can indeed see when the perturbation is small (< 0.5), there are no misclustered modes. The impact of $\epsilon_{\mathbf{T}}$ is shown in Figure 2(b) (lumpable case P1) and 2(c) (aggregatable case P2), where we set $\alpha_{\mathbf{A}} = \alpha_{\mathbf{B}} = 0, \alpha_{\mathbf{T}} = 1$.




Figure 2: Clustering performance vs (a) $\epsilon_A = \epsilon_B$; (b) ϵ_T (lumpable case P1); (c) ϵ_T (aggregatable case P2)

Next we implement the idea of designing LQR controllers for the original Σ through the reduced $\hat{\Sigma}$ as discussed in Section 6. We let $r = 3$, $n = 10$, $p = 5$, process noise variance $\sigma_w^2 = 0.1$, initial state $\mathbf{x}_0 = \mathbf{1}$, and perturbation $\epsilon_A = \epsilon_B = \epsilon_T$. In Figure 3(a), the finite time LQR with horizon $T = 500$ is considered, and the plot shows time-averaged sub-optimality $(\hat{J}_t - J_t^*)/t$ averaged over 4000 runs. At $t = 500$, since the behavior has reached stationary state, $(\hat{J}_t - J_t^*)/t$ can be viewed as the infinite time suboptimality $\hat{J} - J^*$ discussed in Theorem 9. We can easily see the trend that larger perturbations result in larger suboptimality. Figure 3(b) shows the time to compute controllers via Riccati iterations using Σ and $\hat{\Sigma}$. We see when s is large, Σ needs significantly more time than $\hat{\Sigma}$. Since the optimal infinite time LQR controllers also can be obtained through Riccati iterations until convergence, on each curve we use circles to mark the time needed when the controller difference between two adjacent iterations falls below 10^{-12} .




Figure 3: Suboptimality and computation time evaluation of LQR controllers designed with the reduced $\hat{\Sigma}$

8. Conclusion and Future Work

In this work, we propose a clustering-based method that reduces the mode complexity of an MJS. The reduced MJS provably approximates the original MJS in terms of trajectory difference and controller optimality. Several potential extensions include: (i) stronger approximation metrics that allow for inputs; (ii) stability of the reduced MJS; (iii) partial observation case, i.e., the state \mathbf{x}_t is observed through $\mathbf{y}_t = \mathbf{C}_{\omega_t} \mathbf{x}_t$ for some mode-dependent output matrices $\mathbf{C}_{1:s}$.

Acknowledgments

N. Ozay and Z. Du were supported by ONR grants N00014-18-1-2501 and N00014-21-1-2431, L. Balzano and Z. Du were supported by AFOSR YIP award FA9550-19-1-0026 and NSF CAREER award CCF-1845076, and L. Balzano was supported by NSF BIGDATA award IIS-1838179. We thank Dan Li for his suggestions on an earlier draft.

References

- Alessandro Abate. Approximation metrics based on probabilistic bisimulations for general state-space markov processes: a survey. *Electronic Notes in Theoretical Computer Science*, 297:3–25, 2013.
- Alessandro Abate, Joost-Pieter Katoen, John Lygeros, and Maria Prandini. Approximate model checking of stochastic hybrid systems. *European Journal of Control*, 16(6):624–641, 2010.
- Alessandro Abate, Alessandro D’Innocenzo, and Maria D Di Benedetto. Approximate abstractions of stochastic hybrid systems. *IEEE Transactions on Automatic Control*, 56(11):2688–2694, 2011.
- Alessandro Abate, Marta Kwiatkowska, Gethin Norman, and David Parker. Probabilistic model checking of labelled markov processes via finite approximate bisimulations. In *Horizons of the Mind. A Tribute to Prakash Panangaden*, pages 40–58. Springer, 2014.
- Rajeev Alur, Thomas A Henzinger, Gerardo Lafferriere, and George J Pappas. Discrete abstractions of hybrid systems. *Proceedings of the IEEE*, 88(7):971–984, 2000.
- Gaoang Bian and Alessandro Abate. On the relationship between bisimulation and trace equivalence in an approximate probabilistic context. In *International Conference on Foundations of Software Science and Computation Structures*, pages 321–337. Springer, 2017.
- Andreas Bittracher and Christof Schütte. A probabilistic algorithm for aggregating vastly under-sampled large markov chains. *Physica D: Nonlinear Phenomena*, 416:132799, 2021.
- Lars Blackmore, Stanislav Funiak, and Brian C Williams. Combining stochastic and greedy search in hybrid estimation. In *AAAI*, pages 282–287, 2005.
- Peter Buchholz. Exact and ordinary lumpability in finite markov chains. *Journal of applied probability*, 31(1):59–75, 1994.
- DO Cajueiro. *Stochastic optimal control of jumping Markov parameter processes with applications to finance*. PhD thesis, PhD thesis, 2002, Instituto Tecnológico de Aeronáutica-ITA, Brazil, 2002.
- Edmund M Clarke Jr, Orna Grumberg, Daniel Kroening, Doron Peled, and Helmut Veith. *Model checking*. MIT press, 2018.
- Oswaldo Luiz Valle Costa, Marcelo Dutra Fragoso, and Ricardo Paulino Marques. *Discrete-time Markov jump linear systems*. Springer, 2006.
- Josée Desharnais, Abbas Edalat, and Prakash Panangaden. Bisimulation for labelled markov processes. *Information and Computation*, 179(2):163–193, 2002.
- Josee Desharnais, Vineet Gupta, Radha Jagadeesan, and Prakash Panangaden. Metrics for labelled markov processes. *Theoretical computer science*, 318(3):323–354, 2004.
- Zhe Du, Necmiye Ozay, and Laura Balzano. Mode clustering for markov jump systems. *arXiv preprint arXiv:1910.02193*, 2019.

- Zhe Du, Necmiye Ozay, and Laura Balzano. Clustering-based mode reduction for markov jump systems, 2022a. doi: [10.7302/4257](https://doi.org/10.7302/4257).
- Zhe Du, Yahya Sattar, Davoud Ataee Tarzanagh, Laura Balzano, Samet Oymak, and Necmiye Ozay. Certainty equivalent quadratic control for markov jump systems. In *2022 American Control Conference (ACC) (forthcoming)*. IEEE, 2022b.
- Antoine Girard and George J Pappas. Approximation metrics for discrete and continuous systems. *IEEE Transactions on Automatic Control*, 52(5):782–798, 2007.
- Serkan Gugercin and Athanasios C Antoulas. A survey of model reduction by balanced truncation and some new results. *International Journal of Control*, 77(8):748–766, 2004.
- A Agung Julius and George J Pappas. Approximations of stochastic hybrid systems. *IEEE Transactions on Automatic Control*, 54(6):1193–1203, 2009.
- Georgios Kotsalis and Anders Rantzer. Balanced truncation for discrete time markov jump linear systems. *IEEE Transactions on Automatic Control*, 55(11):2606–2611, 2010.
- Robert P Kurshan. *Computer-aided verification of coordinating processes: the automata-theoretic approach*, volume 302. Princeton university press, 2014.
- Kim G Larsen and Arne Skou. Bisimulation through probabilistic testing. *Information and computation*, 94(1):1–28, 1991.
- Jing Lei, Alessandro Rinaldo, et al. Consistency of spectral clustering in stochastic block models. *The Annals of Statistics*, 43(1):215–237, 2015.
- KA Loparo and F Abdel-Malek. A probabilistic approach to dynamic power system security. *IEEE transactions on circuits and systems*, 37(6):787–798, 1990.
- Yuriy Zacchia Lun, Jack Wheatley, Alessandro D’Innocenzo, and Alessandro Abate. Approximate abstractions of markov chains with interval decision processes. *IFAC-PapersOnLine*, 51(16): 91–96, 2018.
- Oded Maler, Amir Pnueli, and Joseph Sifakis. On the synthesis of discrete controllers for timed systems. In *Annual symposium on theoretical aspects of computer science*, pages 229–242. Springer, 1995.
- Marina Meilă and Jianbo Shi. A random walks view of spectral segmentation. In *International Workshop on Artificial Intelligence and Statistics*, pages 203–208. PMLR, 2001.
- Yahya Sattar, Zhe Du, Davoud Ataee Tarzanagh, Laura Balzano, Necmiye Ozay, and Samet Oymak. Identification and adaptive control of markov jump systems: Sample complexity and regret bounds, 2021.
- Ying Shen, Zheng-Guang Wu, Peng Shi, and Choon Ki Ahn. Model reduction of markovian jump systems with uncertain probabilities. *IEEE Transactions on Automatic Control*, 65(1):382–388, 2019.

- Bruno Sinopoli, Luca Schenato, Massimo Franceschetti, Kameshwar Poolla, and Shankar Sastry. An lqg optimal linear controller for control systems with packet losses. In *Proceedings of the 44th IEEE Conference on Decision and Control*, pages 458–463. IEEE, 2005.
- Sadegh Esmail Zadeh Soudjani and Alessandro Abate. Adaptive gridding for abstraction and verification of stochastic hybrid systems. In *2011 Eighth International Conference on Quantitative Evaluation of SysTems*, pages 59–68. IEEE, 2011.
- Minhui Sun and James Lam. Model reduction of discrete markovian jump systems with time-weighted h2 performance. *International Journal of Robust and Nonlinear Control*, 26(3):401–425, 2016.
- Lars EO Svensson, Noah Williams, et al. Optimal monetary policy under uncertainty: a markov jump-linear-quadratic approach. *Federal Reserve Bank of St. Louis Review*, 90(4):275–293, 2008.
- Ilya Tkachev and Alessandro Abate. On approximation metrics for linear temporal model-checking of stochastic systems. In *Proceedings of the 17th international conference on Hybrid systems: computation and control*, pages 193–202, 2014.
- Thao HT Truong, Peter Seiler, and Lauren E Linderman. Analysis of networked structural control with packet loss. *IEEE Transactions on Control Systems Technology*, 2021.
- Valery Ugrinovskii and Hemanshu R Pota. Decentralized control of power systems via robust control of uncertain markov jump parameter systems. *International Journal of Control*, 78(9):662–677, 2005.
- Majid Zamani and Alessandro Abate. Approximately bisimilar symbolic models for randomly switched stochastic systems. *Systems & Control Letters*, 69:38–46, 2014.
- Anru Zhang and Mengdi Wang. Spectral state compression of markov processes. *IEEE transactions on information theory*, 66(5):3202–3231, 2019.
- Huiyan Zhang, Ligang Wu, Peng Shi, and Yuxin Zhao. Balanced truncation approach to model reduction of markovian jump time-varying delay systems. *Journal of the Franklin Institute*, 352(10):4205–4224, 2015a.
- Huiyan Zhang, Ligang Wu, Peng Shi, and Yuxin Zhao. Model reduction on markovian jump systems with partially unknown transition probabilities: balanced truncation approach. *IET Control Theory & Applications*, 9(9):1411–1421, 2015b.
- Liqian Zhang, Biao Huang, and James Lam. \mathcal{H}_∞ model reduction of markovian jump linear systems. *Systems & Control Letters*, 50(2):103–118, 2003.