2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

On Concentration Bounds for Bayesian Identification of Linear Non-Gaussian Systems

Yeoneung Kim Gihun Kim Insoon Yang

***Abstract*— We adopt a Bayesian perspective to identify the unknown parameters of linear stochastic systems with possibly non-Gaussian disturbance distributions. The key idea of our algorithm is to alternately execute L randomly selected linear state-feedback controllers and keep track of a maximum a posteriori estimator. The proposed algorithm asymptotically achieves the concentration of posterior distributions around the true system parameters. We also derive probabilistic bounds for the concentration based on the classical results regarding the asymptotic properties of posterior distributions. An empirical demonstration is provided as well.**

I. Introduction

System identification is a fundamental problem in the analysis and design of control systems [1]. Various classical methods have been extensively studied and have contributed to many practical domains including engineering, economics, finance, and biology. With advances in the intersection of machine learning and control theory, new perspectives, analyses, and algorithmic ideas have been emerging in terms of concentration bounds and sample complexity, among others.

A rich body of literature focuses on identifying a pair of unknown parameters, (A, B), of stochastic linear time-invariant (LTI) systems $x_{t+1} = Ax_t + Bu_t + w_t$. Most of the recent works consider the least square estimation exploiting its various benefits in theoretical analyses [2]–[10].

For a marginally stable uncontrolled system ($B = 0$ and $\rho(A) \leq 1$ where ρ denotes the spectral radius), [2] provides a non-asymptotic high-probability upper bound for the error using a single trajectory. Moreover, [3] carries out provable guarantees for both controlled and uncontrolled systems. Other notable works are [4], [5], which show that it is effective to estimate the system parameter first and then utilize the standard bootstrap argument; however, their methods are limited to the Gaussian noise case.

There are several works [6], [7], [9], [10] in which an extension to non-Gaussian noise is considered using a single trajectory. In [6], uncontrolled systems with sub-Weibull disturbance distributions are considered, and probabilistic guarantees are obtained using the standard mixing time argument [11]. However, the theoretical guarantees are valid only when the geometric multiplicity of eigenvalues is greater than unity is one. In parallel, [7] provides sharp error bounds for the case $1 - C/T \leq \rho(A) \leq 1 + C/T$ based on Mendelson's small-ball. A subsequent work, [10], handles error bounds for systems satisfying $\rho(A) \geq 1 + C/T$ under isotropic sub-Gaussian disturbance. However, the concentration bounds derived in the aforementioned literature require the evaluation of an unverifiable criterion or the knowledge of true system parameters. In [7], the concentration bound requires that $(k, \Gamma_{\text{sb}}, p)$-small ball condition. We also note that the bound provided in [10] relies on the true system parameters. Furthermore, it also has been acknowledged in the same paper that least squares estimation is statistically inconsistent under certain conditions.

Attempts have also been made to derive sample complexity using multiple trajectories. Notably, [4] presents concentration bound for least squares estimation using multiple full trajectories based on the bootstrap argument. The result holds regardless of the stabilizability but explicit dependence on the time horizon and rollouts is unavailable. More recently, [12] provides concentration bounds for least squares estimation using full trajectories under a certain assumption on the system parameters, k-controllability.

In this paper, we consider a fairly general class of LTI systems with inputs ($B \neq 0$) and non-Gaussian disturbances. From the Bayesian perspective, we propose a simple posterior update algorithm and analyze its asymptotic concentration properties that are valid regardless of the stabilizability of (A, B). Bayesian system identification also has an extensive body of literature (see e.g., [1], [13]–[15] and the references therein). The works most relevant to ours are [13], [16], where the posterior mean is used as an estimator with data from a very short period. They introduce a novel Markov Chain Monte Carlo (MCMC) algorithm and provide its convergence rate. However, MCMC methods usually require considerable computational resources; only a one-dimensional experiment is conducted using a simple piecewise-constant control input in [13].

Departing from these approaches, we devise a simple novel algorithm to obtain a maximum a posteriori (MAP) estimator, which can be computed efficiently for unimodal posterior distributions. The key idea of our algorithm is to alternatively execute a certain number of randomly selected linear state-feedback controllers satisfying a mild excitation condition, and update the posterior of unknown parameters (A, B). Leveraging the classical results on the asymptotic properties of posterior distributions [17], [18], we establish asymptotic concentration bounds for the MAP estimator around the true system parameter in terms of the time horizon

This work was supported in part by the National Research Foundation of Korea funded by MSIT (2020R1C1C1009766, RS-2023-00219980) and the Information and Communications Technology Planning and Evaluation funded by MSIT (2020-0-00857, 2022-0-00480).

Y. Kim is with the Department of Applied Artificial Intelligence, Seoul National University of Science and Technology, Seoul 01811, South Korea. G. Kim, and I. Yang are with the Department of Electrical and Computer Engineering, ASRI, Seoul National University, Seoul 08826, South Korea, insoonyang@snu.ac.kr

under the mild structural assumptions on the disturbance distribution, namely, the log-concavity of the probability density function (pdf). In this case, computational tractability is enhanced since the posterior pdf is unimodal. The performance of our method is numerically demonstrated through problems with non-Gaussian disturbance distributions in both stabilizable and unstabilizable cases.

II. Preliminaries

A. Linear Systems with Non-Gaussian Disturbances

Consider a linear stochastic system of the form

$$x_{t+1} = Ax_t + Bu_t + w_t, \quad t = 1, 2, \ldots, \tag{1}$$

where $x_t \in \mathbb{R}^n$ is the system input, and $u_t \in \mathbb{R}^m$ is the control input. The disturbance $w_t \in \mathbb{R}^n$ is an independent and identically distributed (i.i.d.) zero-mean random vector with covariance matrix W. Throughout the paper, the $d \times d$ identity matrix is denoted by I_d.

Assumption 1: For every $t = 1, 2, \ldots$, the random disturbance vector w_t satisfies the following properties:

1) The probability density function (pdf) of noise $p_w(\cdot)$ is known and twice differentiable. Additionally, p_w is strongly log-concave, i.e.,
$$-\nabla^2_{w_t} \log p_w(w_t) \succeq \underline{m} I_n$$
for some $\underline{m} > 0$;
2) $\mathbb{E}[w_t] = 0$ and $\mathbb{E}[w_t w_t^\top] = W$, where W is positive definite.

Remark 1: This assumption allows us to consider more general classes of disturbances beyond Gaussian. One example of noise that we use for the experiment is the Gaussian mixture distribution. Another class is an asymmetric noise whose last coordinate is assumed to be asymmetric while the others follow Gaussian. Please see Section IV for details.

Let $d := n + m$ and $\Theta \in \mathbb{R}^{d \times n}$ be the system parameter matrix defined by

$$\Theta := \begin{bmatrix} \Theta(1) & \cdots & \Theta(n) \end{bmatrix} := \begin{bmatrix} A & B \end{bmatrix}^\top,$$

where $\Theta(i) \in \mathbb{R}^d$ denotes the ith column of Θ. We also let

$$\theta := \mathrm{vec}(\Theta) := (\Theta(1), \Theta(2), \ldots, \Theta(n)) \in \mathbb{R}^{dn}$$

denote the vectorized version of Θ and often refer to θ as the parameter vector.

Let $h_t := (x_1, u_1, \ldots, x_{t-1}, u_{t-1}, x_t)$ be the *history* of observations made up to time t, and let H_t denote the collection of such histories at stage t.

B. Posterior Update

In this paper, we consider the Bayesian setting, where the prior of the true system parameter θ_* is given. The posterior update process is described using Bayes' rule inductively. Let the state-input pair be denoted by

$$z_t := (x_t, u_t) \in \mathbb{R}^d.$$

Then, the linear system (1) is expressed as $x_{t+1} - \Theta^\top z_t = w_t \sim p_w$, and therefore,

$$p(x_{t+1}|z_t, \theta) = p_w(x_{t+1} - \Theta^\top z_t),$$

which is strongly log-concave in θ under Assumption 1. On the other hand, we have

$$\begin{aligned} p(\theta|h_{t+1})p(x_{t+1}|z_t, h_t) &= p(\theta|x_{t+1}, z_t, h_t)p(x_{t+1}|z_t, h_t) \\ &= p(x_{t+1}|z_t, h_t, \theta)p(\theta|z_t, h_t) \\ &= p(x_{t+1}|z_t, \theta)p(\theta|h_t), \end{aligned}$$

where the second equality follows from Bayes' rule. Thus, the posterior at stage t is given by

$$\begin{aligned} p(\theta|h_{t+1}) &\propto p(x_{t+1}|z_t, \theta)p(\theta|h_t) \\ &= p_w(x_{t+1} - \Theta^\top z_t)p(\theta|h_t). \end{aligned} \tag{2}$$

Hence, if $p(\theta|h_t)$ is strongly log-concave, then so is $p(\theta|h_{t+1})$. Using mathematical induction, it is straightforward to show that the posterior pdf at any stage is strongly log-concave since strong log-concavity is preserved through the posterior update.

III. System Identification via Posterior Update

In this section, we propose an algorithm for estimating the posterior pdf of θ_* by applying only L randomly selected linear state-feedback controllers. Randomizing controllers seems to provide useful information for estimating system parameters, as discussed in the literature (e.g., [4], [6], [7]). Departing from the existing approaches, we simply adopt L randomly generated control gain matrices satisfying a mild excitation condition. Specifically, given $L \in \mathbb{N}$ which is determined by the dimension of the state and control input m and n, we introduce L control gain matrices $K_1, ..., K_L \in \mathbb{R}^{n \times m}$ and apply state-feedback controllers $u_t = K_i x_t$ when $t \equiv i \pmod{L}$ up to time T to collect the sequence of state-input pairs $(z_s)_{s=1}^T$ for N independent rollouts, which is denoted by $(z_s^{(\ell)})_{s=1}^T$ for $\ell = 1, \ldots, N$. We emphasize that the initial state x_1 is always set to 0 after each rollout. In the last part of this section, we rigorously show that the posterior distribution obtained using the collected data is asymptotically concentrated around the true system parameter θ_*.

A. Main Algorithm

Our algorithm is designed to compute the posterior pdf of θ that concentrates around $\theta = \theta_*$, given h_t. One can then estimate θ_* as $\arg\min_\theta\{-\log p(\cdot|h_t)\} = \arg\min_\theta\{U_t(\cdot)\}$, where U_t is the potential of the posterior pdf at time t, that is, $p(\theta|h_t) \propto e^{-U_t(\theta)}$.

We begin by specifying the control gain matrices used in our algorithm. The matrices are chosen to satisfy the following mild assumption need for excitation:

Assumption 2: For any $m, n \in \mathbb{N}$, let L be the smallest integer such that $m + n \leq Ln$. We assume that

$$\lambda_{\min}\left(\sum_{s=1}^{L} \begin{bmatrix} I_n \\ K_s \end{bmatrix} W \begin{bmatrix} I_n \\ K_s \end{bmatrix}^\top \right) > 0.$$

Algorithm 1 System identification via posterior update

Input: $\lambda > 0$, $N \in \mathbb{N}$, and L control gain matrices $K_1, K_2, \ldots, K_L$ satisfying Assumption 2;
$t \leftarrow 1$, $x_1 \leftarrow 0$, $\mathcal{D} \leftarrow \emptyset$;
for $\ell = 1, 2, \ldots, N$ **do**
    for $t = 1, 2, \ldots, T-1$ **do**
        if $t \equiv i \pmod{L}$ **then**
            Execute $u_t^{(\ell)} = K_i x_t^{(\ell)}$;
        end if
        Observe new state $x_{t+1}^{(\ell)}$;
        Update $\mathcal{D} \leftarrow \mathcal{D} \cup \{(z_t^{(\ell)}, x_{t+1}^{(\ell)})\}$;
        $t \leftarrow t + 1$;
    end for
end for
Compute $\tilde{U}$ as (3);
$\tilde{\theta} = \arg\min \tilde{U}(\cdot)$

Remark 2: It is intuitively clear that this condition is satisfied with a high probability if the L matrices (or one of the L matrices) are randomly selected. We empirically verified the condition for $n = m = 3$ and 5 by simply choosing $K_1 = -\frac{3}{2} I_n$ and K_2 to be a random matrix whose entries are sampled from the standard Gaussian distribution. Among 10,000 pairs of randomly generated matrices, none violate the condition.

Our Bayesian system identification algorithm is presented in Algorithm 1. Here, we simply assume that the prior distribution of the true system parameter follows $\mathcal{N}(0, \frac{1}{\lambda} I_{dn})$ for some $\lambda > 0$. Applying the control gain matrices K_i alternatively up to T, the data are collected. We then invoke Bayes' rule to obtain the potential of the posterior distribution of θ with N rollouts as follows:

$$\tilde{U}(\theta) = \frac{\lambda}{2}|\theta|^2 - \sum_{\ell=1}^{N} \sum_{s=1}^{T-1} \log p_w(x_{s+1}^{(\ell)} - \Theta^\top z_s^{(\ell)}), \qquad (3)$$

which consists of the prior and likelihood parts. Note that $\tilde{U}$ is strongly convex for any $\lambda > 0$ since p_w is log-concave. It is worth noting that $\tilde{\theta} \in \arg\min_\theta \tilde{U}(\cdot)$ is indeed an MAP estimator, and it can be computed using existing convex optimization algorithms.

To analyze the potential, we introduce the following useful quantity:

$$\hat{P}_t^{(\ell)} := \sum_{s=1}^{t-1} \mathrm{blkdiag}\{z_s^{(\ell)} z_s^{(\ell)\top}\}_{i=1}^{n}, \qquad (4)$$

where $\mathrm{blkdiag}\{A_i\}_{i=1}^n \in \mathbb{R}^{dn \times dn}$ denotes the block diagonal matrix consisting of $A_i \in \mathbb{R}^{d \times d}$ for $i = 1, \ldots, n$ in the diagonal. The next lemma describes the geometric relationship between $\hat{P}_t^{(\ell)}$ and the potential for the likelihood

$$\hat{U}_t^{(\ell)}(\theta) := -\sum_{s=1}^{t-1} \log p_w(x_{s+1}^{(\ell)} - \Theta^\top z_s^{(\ell)}). \qquad (5)$$

Lemma 1: Let $\ell \in \{1, \ldots, N\}$. Under Assumption 1, the following inequality holds

$$\nabla^2 \hat{U}_t^{(\ell)}(\theta) \succeq \underline{m} \hat{P}_t^{(\ell)}$$

for any t and any θ, where $\underline{m}$ is the constant specified in Assumption 1.

Proof: To avoid clutter, we suppress the superscript ℓ. By direct calculation, the following equality holds:

$$\begin{aligned}
&\nabla_\theta^2 \log p_w(x_{s+1} - \Theta^\top z_s) \\
&\quad = \nabla_{w_s}^2 \log p_w(x_{s+1} - \Theta^\top z_s) \otimes z_s z_s^\top,
\end{aligned}$$

where $\otimes$ denotes the Kronecker product. Then, the Hessian of $\hat{U}_t$ can be written as

$$\nabla_\theta^2 \hat{U}_t = -\sum_{s=1}^{t-1} \nabla_{w_s}^2 \log p_w(x_{s+1} - \Theta^\top z_s) \otimes z_s z_s^\top.$$

By Assumption 1, for any state-input pair $z_s = (x_s, u_s)$, we have

$$\begin{aligned}
&- \nabla_{w_s}^2 \log p_w(x_{s+1} - \Theta^\top z_s) \otimes z_s z_s^\top \\
&\quad \succeq \underline{m} \mathrm{blkdiag}(\{z_s z_s^\top\}_{i=1}^n),
\end{aligned}$$

and

$$\nabla_\theta^2 \hat{U}_t \succeq \underline{m} \sum_{s=1}^{t-1} \mathrm{blkdiag}(\{z_s z_s^\top\}_{i=1}^n).$$

Therefore, the result follows. ■

Using this lemma, we further compare the Hessian of the expected potential $U := \mathbb{E}[\hat{U}_T^{(\ell)}(\theta)]$ and $P := \mathbb{E}[\hat{P}_T^{(\ell)}]$ as

$$\begin{aligned}
\nabla_\theta^2 U(\theta) &= \nabla_\theta^2 \mathbb{E}[\hat{U}_T^{(\ell)}(\theta)] \\
&= \mathbb{E}[\nabla_\theta^2 \hat{U}_T^{(\ell)}(\theta)] \succeq \underline{m} \mathbb{E}[\hat{P}_T^{(\ell)}] = \underline{m} P.
\end{aligned} \qquad (6)$$

Now we verify that $\lambda_{\min}(\nabla_\theta^2 U(\theta))$ grows at least linearly in the time horizon T.

Theorem 1: Suppose that $T > L$. Let $(z_s^{(\ell)})_{s=1}^T$ be the data collected from Algorithm 1 and $P := \mathbb{E}[\hat{P}_T^{(\ell)}]$. Then, we have

$$\lambda_{\min}(\nabla_\theta^2 U(\theta)) \geq \lambda_{\min}(P) > \left(\frac{T}{L} - 1\right) \tilde{\lambda}_{\min}$$

for any $T > 1$, where $\tilde{\lambda}_{\min}$ represents the minimum eigenvalue of

$$\sum_{i=1}^{L} \begin{bmatrix} I_n \\ K_i \end{bmatrix} W \begin{bmatrix} I_n \\ K_i \end{bmatrix}^\top.$$

Proof: We again suppress the superscript ℓ for simplicity. For each $t \geq 1$, we have that $\mathbb{E}[x_t x_t^\top] = W$ where the expectation is taken with respect to $w_1, \ldots, w_t$.

By the definition of z_t, we have

$$\begin{aligned}
\sum_{s=1}^{T-1} \mathbb{E}[z_s z_s^\top] &\succeq \left(\frac{T}{L} - 1\right) \sum_{s=1}^{L} \begin{bmatrix} I_n \\ K_s \end{bmatrix} \mathbb{E}[x_i x_i^\top] \begin{bmatrix} I_n \\ K_s \end{bmatrix}^\top \\
&\succeq \left(\frac{T}{L} - 1\right) \sum_{s=1}^{L} \begin{bmatrix} I_n \\ K_s \end{bmatrix} W \begin{bmatrix} I_n \\ K_s \end{bmatrix}^\top,
\end{aligned}$$

and the result follows. ■

It follows from Theorem 1 that $\lambda_{\min}(\nabla_\theta^2(U(\theta)) \to \infty$ as $T \to \infty$, which can be understood in the following way. The growth of $\lambda_{\min}(\nabla_\theta^2 U(\theta))$ implies that the curvature of the potential $U(\theta)$ increases, which allows us to measure the concentration rate in terms of T.

B. Concentration Property

The classical result by Doob [17] asserts that the posterior distributions converge weakly (convergence in distribution) to the Dirac measure centered at the true parameter as the number of data N grows to infinity. The quantitative property of the potential investigated in the previous section allows us to analyze further beyond the weak convergence. Our focus is to reveal how the regularity of the potential contributes to establishing the probabilistic guarantee of system identification in the Bayesian framework.

In this section, we impose the following assumption on the $\{P_\theta\}$ which represents the natural parametric family of probability distributions for the system parameter θ in our problem.[1]

Assumption 3: The parameter family $\{P_\theta\}$ is identifiable, i.e., $\theta_1 \neq \theta_2$ implies $P_{\theta_1} \neq P_{\theta_2}$.

We denote the i.i.d samples from P_{θ_*} by $(X_1, ..., X_N)$, where X_ℓ represents a single trajectory $(x_s^{(\ell)}, z_s^{(\ell)})_{s=1}^T$. Throughout the paper, we let $B_r(\theta) = \{x : |x - \theta| \leq r\}$ and denote the complement of this ball by $B_r^c(\theta)$.

According to the Bernstein–von Mises theorem [18], the posterior distribution asymptotically follows the Gaussian distribution centered around the maximum likelihood estimator (MLE) denoted by $\hat{\theta}$, that is,

$$\hat{\theta} \in \arg\max_\theta p(X_1, ..., X_N|\theta).$$

Lemma 2 (Bernstein Von Mises Theorem [18]): Under Assumption 3, we have

$$\sqrt{N}(\theta - \hat{\theta})|(X_1, ..., X_N) \xrightarrow{d} \mathcal{N}(0, (\nabla_\theta^2 U(\theta_*))^{-1}),$$

where θ follows the posterior distribution associated with the data $(X_1, ..., X_N)$.

Recall that $\tilde{\theta}$ is the MAP estimator, obtained as

$$\tilde{\theta} \in \arg\min_\theta \tilde{U}(\theta)$$

in Algorithm 1. When the potential is strongly convex, the following inequality holds.

Lemma 3: [Lemma 10 in [19]] For a random variable $\theta \in \mathbb{R}^{dn}$ whose potential is a strongly convex function $U(\theta)$ satisfying $\lambda_{\min}(\nabla_\theta^2 U(\theta)) \geq \ell$, we have

$$\mathbb{E}[|\theta - \tilde{\theta}|^p] \leq 5^p \left(\frac{dnp}{\ell}\right)^{p/2}, \quad p > 0,$$

where the expectation is taken with respect to P_θ.

Finally, we claim the following result on the concentration of $\tilde{\theta}$ around θ_* as the number of rollouts N grows to infinity.

[1]Precisely, for a single trajectory $(x_s, z_s)_{s=1}^T$ generated by the algorithm, the parametric family of probability distributions is given by $P_\theta \sim p(\theta)\Pi_{t=1}^{T-1} p_w(x_{t+1} - \Theta^\top z_t)$.

Theorem 2: Suppose that Assumptions 1–3 hold. Fix arbitrary $\delta > 0$ and $r > 0$ and let $T_0 := \eta^{-1}(\frac{\delta}{2C})$, where $C = \frac{2^{-dn/2}}{\Gamma(dn/2+1)}$ and $\eta(T) := \int_{\sqrt{T/C_1}}^\infty e^{-r^2} r^{dn-1} \mathrm{d}r$ with constants $C_1 = \frac{L(L+1)}{\underline{m}\tilde{\lambda}_{\min}}$ and $\tilde{\lambda}_{\min}$ from Theorem 1. Then, for any $T \geq T_0$ in Algorithm 1, we have

$$\limsup_{N\to\infty} \Pr(\theta \in B_r^c(\theta_*)) \leq \delta, \tag{7}$$

and

$$\limsup_{N\to\infty} \Pr(\tilde{\theta} \in B_{2r}^c(\theta_*)) \leq \delta, \tag{8}$$

where θ follows the posterior distribution obtained in Algorithm 1.

Proof: Fix $\delta > 0$ and $r > 0$. First, to show the concentration bound (7), we observe that

$$\Pr(|\theta - \theta_*| \geq r) \leq \Pr\Big(|\theta - \hat{\theta}| \geq \frac{r}{2}\Big) + \Pr\Big(|\hat{\theta} - \theta_*| \geq \frac{r}{2}\Big).$$

In what follows, we bound the two terms on the right-hand side separately.

By Lemma 2, we have

$$\lim_{N\to\infty} \Pr(\theta \in B_{1/\sqrt{N}}^c(\hat{\theta})|X) = \Pr(|Y| \geq 1),$$

where $Y \sim \mathcal{N}(0, (\nabla_\theta^2 U(\theta_*))^{-1})$. Note that Y follows a zero-mean multivariate Gaussian distribution with covariance $(\nabla_\theta^2 U(\theta_*))^{-1}$. By the choice of C_1 and Theorem 1, we have

$$\lambda_{\min}(\nabla_\theta^2 U(\theta_*)) \geq \underline{m}\Big(\frac{T}{L} - 1\Big)\tilde{\lambda}_{\min} \geq \frac{T}{C_1}.$$

Thus, we deduce that

$$\begin{aligned}
&\Pr(|Y| \geq 1) \\
&= (2\pi)^{-dn/2} |\det(\nabla_\theta^2 U(\theta_*))|^{-1/2} \int_{|y|\geq 1} e^{-y^\top \nabla_\theta^2 U(\theta_*) y} \mathrm{d}y \\
&\leq (2\pi)^{-dn/2} \int_{\lambda_{\max}(\nabla_\theta^2 U(\theta_*)^{-1})|z|^2 \geq 1} e^{-|z|^2} \mathrm{d}z \\
&\leq (2\pi)^{-dn/2} \int_{|z|\geq\sqrt{T/C_1}} e^{-|z|^2} \mathrm{d}z \\
&\leq \frac{2^{-dn/2}}{\Gamma(dn/2+1)} \int_{\sqrt{T/C_1}}^\infty e^{-r^2} r^{dn-1} \mathrm{d}r.
\end{aligned}$$

By the definition of C and $\eta(T)$, the inequality above can be expressed as

$$\lim_{N\to\infty} \Pr(\theta \in B_{1/\sqrt{N}}^c(\hat{\theta})|X) \leq C\eta(T).$$

Hence,

$$\begin{aligned}
&\limsup_{N\to\infty} \Pr(\theta \in B_{r/2}^c(\hat{\theta})) \\
&= \limsup_{N\to\infty} \int \Pr(\theta \in B_{r/2}^c(\hat{\theta})|X) \mathrm{d}P_X \\
&\leq \limsup_{N\to\infty} \int \Pr(\theta \in B_{1/\sqrt{N}}^c(\hat{\theta})|X) \mathrm{d}P_X \\
&\leq C\eta(T).
\end{aligned}$$

Let $T_0 = \eta^{-1}(\frac{\delta}{2C})$. Then, for any $T \geq T_0$, we have

$$C\Big(1 - \mathrm{erf}\big(\sqrt{T/C_1}\big)\Big) \leq \frac{\delta}{2},$$

which implies that

$$\limsup_{N\to\infty} \Pr\Big(|\theta - \hat{\theta}| \geq \frac{r}{2}\Big) \leq \frac{\delta}{2}. \tag{9}$$

Now, recall the asymptotic property of MLE, i.e.,

$$\sqrt{N}(\hat{\theta} - \theta_*) \xrightarrow{d} \mathcal{N}(0, \nabla_\theta^2 U(\theta_*)^{-1}) \quad \text{as } N \to \infty.$$

Following the same argument as above, we have

$$\begin{aligned}\limsup_{N\to\infty} \Pr\Big(|\hat{\theta} - \theta_*| \geq \frac{r}{2}\Big) &\leq \lim_{N\to\infty} \Pr\Big(|\hat{\theta} - \theta_*| \geq \frac{1}{\sqrt{N}}\Big) \\ &= \Pr(|Y| \geq 1) \leq \frac{\delta}{2}\end{aligned}$$

for $T \geq T_0$. Combining this bound with (9) yields the first concentration bound (7).

To show the second concentration bound (8), we first observe that

$$\Pr(|\tilde{\theta} - \theta_*| \geq 2r) \leq \Pr(|\theta - \tilde{\theta}| \geq r) + \Pr(|\theta - \theta_*| \geq r).$$

The first term on the right-hand side is bounded as

$$\begin{aligned}\Pr(|\theta - \tilde{\theta}| \geq r) &= \int \Pr(|\theta - \tilde{\theta}| \geq r \mid X) \mathrm{d}P_X \\ &\leq \frac{1}{r^2} \int \mathbb{E}[|\theta - \tilde{\theta}|^2 \mid X] \mathrm{d}P_X \\ &\leq \frac{C_0}{r^2} \int \frac{1}{\lambda + \underline{m}\lambda_{\min}(\sum_{\ell=1}^N \hat{P}_T^{(\ell)})} \mathrm{d}P_X \\ &= \frac{C_0}{r^2} \int \frac{1}{\lambda + \underline{m}N\lambda_{\min}(\frac{\sum_{\ell=1}^N \hat{P}_T^{(\ell)}}{N})} \mathrm{d}P_X\end{aligned}$$

for some positive constant C_0, where the first inequality holds due to Markov's inequality and the second inequality follows from Lemma 3 with the property $\lambda_{\min}(\nabla_\theta^2 \tilde{U}(\theta)) \geq \lambda + \underline{m}\lambda_{\min}(\sum_{\ell=1}^T \hat{P}_t^{(\ell)})$. Thus, by the strong law of large numbers, we have

$$\limsup_{N\to\infty} \Pr(|\theta - \tilde{\theta}| \geq r) = 0.$$

Combining this with the first concentration bound (7) yields $\lim_{N\to\infty} \Pr(|\tilde{\theta} - \theta_*| \geq 2r) \leq \delta$ for $T \geq T_0$. ■

Example 1: For $t = 1, \ldots, T-1$, $\ell \in \{1, \ldots, N\}$, and $\lambda = 0$, consider the system $x_{t+1}^{(\ell)} = a_* x_t^{(\ell)} + u_t^{(\ell)} + w_t^{(\ell)}$, where $x_t^{(\ell)}, u_t^{(\ell)} \in \mathbb{R}$ and $w_t^{(\ell)} \sim \mathcal{N}(0,1)$. Choosing $u_t^{(\ell)} = x_t^{(\ell)}$, the MAP estimator a is written as $a = a_* + \frac{\sum_{\ell=1}^N \sum_{s=1}^{T-1} x_s^{(\ell)} w_s^{(\ell)}}{\sum_{\ell=1}^N \sum_{s=1}^{T-1} x_s^{(\ell)2}}$. By Markov's inequality,

$$\Pr(|a - a_*| > r) \leq \frac{1}{r^2} \mathbb{E}\Bigg[\frac{\sum_{\ell=1}^N \sum_{s=1}^{T-1} x_s^{(\ell)} w_s^{(\ell)}}{\sum_{\ell=1}^N \sum_{s=1}^{T-1} x_s^{(\ell)2}}\Bigg].$$

Roughly, the denominator scales in $O(N)$ while the numerator scales in $O(\sqrt{N})$ with high probability as shown in [5]. Our result provides a probabilistic bound of this complicated term from the Bayesian perspective.

The technique used in our analysis has distinguishing features compared to other works. In particular, [9], [10], [20] leverage martingale properties to handle the correlation between states and obtain a finite time sample complexity using a single trajectory. Our result is weaker in the sense that we derive asymptotic sample complexity by taking the number of rollouts N to infinity. However, we provide a simple and elegant approach from a Bayesian perspective by focusing solely on the geometric property of the potential, namely $\nabla_\theta^2 U(\theta) \geq \big(\frac{T}{L} - 1\big)\tilde{\lambda}_{\min}$.

IV. Experiments

We test the performance of our algorithm using three types of disturbance distributions: (i) Gaussian, (ii) Gaussian mixture, and (iii) asymmetric.[2] Furthermore, we consider both stabilizable and unstabilizable cases.[3] We compare the performance of our algorithm with that of [4] which takes the ordinary least squares estimation (LSE) approach using the full trajectory of N rollouts.[4]

The prior distribution of the true parameter is assumed to follow the standard Gaussian distribution and therefore $\lambda = 1$. The dimension of the state and input spaces is chosen as $n = m = 3$ and 5. We repeated the experiment 10 times for each case. Additional problem data and our implementation of the method are available online.[5] We use $L = 2$ control gain matrices, $K_1 = -\frac{3}{2} I_n$ and K_2 whose components are independently sampled from the standard Gaussian distribution.

1) *Stabilizable Case* : In the first set of experiments, the true system is chosen to be stabilizable and we set $N = 200$. The results are shown in Fig. 1 (a)–(c). For all three types of disturbances, our algorithm shows a comparable performance to the LSE method [4] in both three- and five-dimensional cases.
2) *Unstabilizable Case* : In the second set of experiments, we identify the parameters of unstabilizable systems with $N = 100$. As shown in Fig. 1 (d)–(f), our MAP method still performs well in both three- and five-dimensional problems even though only two different controllers are applied alternatively.

In conclusion, we observe the decay of the relative error as T increases, which is consistent with Theorem 2. Our MAP method supports that using only a limited number of state-feedback controllers is as efficient as the standard LSE method with random control inputs.

[2]For the Gaussian case, we use the standard Gaussian distribution. For the Gaussian mixture case, the pdf of w_t is chosen as $p(w_t) = \frac{1}{2\sqrt{2\pi}}(e^{-|w_t - a|^2/2} + e^{-|w_t + a|^2/2})$, where $a = \frac{1}{4}\mathbf{1}_m$. For the asymmetric case, we let all entries of w_t except for the last element $w_t[n]$ be independent and follow the standard Gaussian distribution. For the last element, we set $-\frac{\partial^2 \log p(w_t)}{\partial w_t[n]^2} = k$ if $w_t[n] < \alpha$, $-\frac{\partial^2 \log p(w_t)}{\partial w_t[n]^2} = \frac{K-k}{\beta-\alpha} w_t[n] + k - \frac{(K-k)\alpha}{\beta-\alpha}$ if $\alpha \leq w_t[n] < \beta$, and $-\frac{\partial^2 \log p(w_t)}{\partial w_t[n]^2} = K$ if $\beta \leq w_t[n]$, where $\alpha = 0$, $\beta = 40$, $k = 1$ and $K = 10$.

[3]For the former case, we randomly choose the system parameter (A, B) which is stabilizable. For the latter, we consider the case $\operatorname{rank}\begin{bmatrix} A - \lambda I & B \end{bmatrix} < n$ for some $\lambda > 0$.

[4]Precisely, for the trajectories $(z_s^{(\ell)})_{s=1}^T$ obtained via executing random controllers $u_t^{(\ell)} \sim \mathcal{N}(0, I_m)$, one can achieve the ordinary least squares estimation as $\tilde{\theta} \in \arg\min_\theta \sum_{\ell=1}^N \sum_{s=1}^{T-1} |x_{s+1}^{(\ell)} - \Theta^\top z_s^{(\ell)}|^2$.

[5]https://github.com/yeoneung/bayesian_system_id

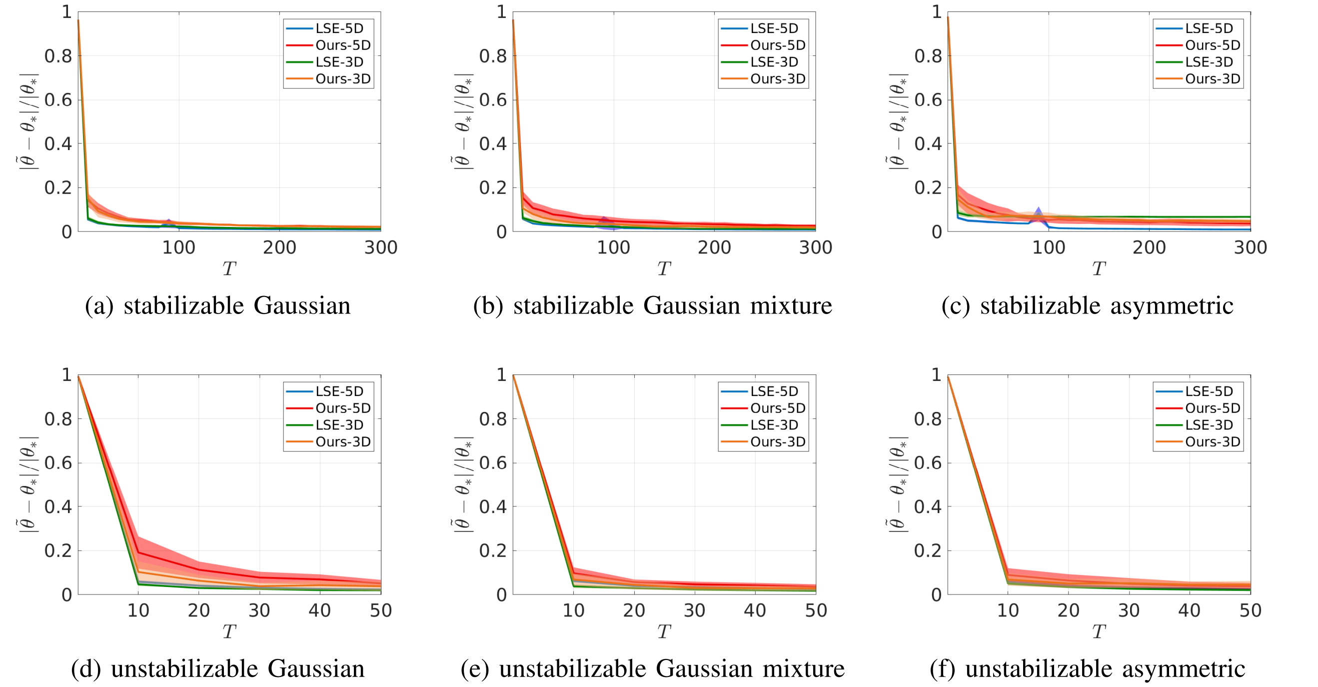

(a) stabilizable Gaussian (b) stabilizable Gaussian mixture (c) stabilizable asymmetric

(d) unstabilizable Gaussian (e) unstabilizable Gaussian mixture (f) unstabilizable asymmetric

Fig. 1: Relative estimation error $|\tilde{\theta}-\theta_*|/|\theta_*|$ under various disturbance distributions for stabilizable θ_* (top) and unstabilizable θ_* (bottom).

V. Conclusions

We have proposed a Bayesian approach to identifying unknown parameters of linear non-Gaussian systems via alternative executions of L randomly selected linear state-feedback controllers. The performance of our MAP estimator is characterized through rigorous concentration analyses that hold regardless of stabilizability. We believe that this work can be extended in some promising future research directions, including (i) concentration bounds for similar Bayesian system identification using a single trajectory, and (ii) understanding the geometry of random control gain matrices that lead to efficient system identification.

References

[1] L. Ljung, *System Identification*. Springer, 1998.
[2] A. Tsiamis and G. J. Pappas, "Finite sample analysis of stochastic system identification," in *Proc. IEEE Conf. Decis. Control*. IEEE, 2019, pp. 3648–3654.
[3] S. Oymak and N. Ozay, "Non-asymptotic identification of LTI systems from a single trajectory," in *Proc. Am. Control Conf*, 2019, pp. 5655–5661.
[4] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu, "On the sample complexity of the linear quadratic regulator," *Found. Comput. Math.*, vol. 20, no. 4, pp. 633–679, 2020.
[5] N. Matni and S. Tu, "A tutorial on concentration bounds for system identification," in *Proc. IEEE Conf. Decis. Control*. IEEE, 2019, pp. 3741–3749.
[6] M. K. S. Faradonbeh, A. Tewari, and G. Michailidis, "Finite time identification in unstable linear systems," *Automatica*, vol. 96, pp. 342–353, 2018.
[7] M. Simchowitz, H. Mania, S. Tu, M. I. Jordan, and B. Recht, "Learning without mixing: Towards a sharp analysis of linear system identification," in *COLT*. PMLR, 2018, pp. 439–473.
[8] Y. Jedra and A. Proutiere, "Sample complexity lower bounds for linear system identification," in *Proc. IEEE Conf. Decis. Control*. IEEE, 2019, pp. 2676–2681.
[9] ——, "Finite-time identification of linear systems: Fundamental limits and optimal algorithms," *IEEE Trans. Autom. Control*, 2022.
[10] T. Sarkar and A. Rakhlin, "Near optimal finite time identification of arbitrary linear dynamical systems," in *ICML*. PMLR, 2019, pp. 5610–5618.
[11] B. Yu, "Rates of convergence for empirical processes of stationary mixing sequences," *The Annals of Probability*, pp. 94–116, 1994.
[12] S. Tu, R. Frostig, and M. Soltanolkotabi, "Learning from many trajectories," *arXiv preprint arXiv:2203.17193*, 2022.
[13] B. Ninness, S. Henriksen, and T. Brinsmead, "System identification via a computational Bayesian approach," in *Proc. IEEE Conf. Decis. Control*, vol. 2. IEEE, 2002, pp. 1820–1825.
[14] E. J. Hannan and M. Deistler, *The statistical theory of linear systems*. SIAM, 2012.
[15] V. Peterka, "Bayesian approach to system identification," in *Trends and Progress in System Identification*. Elsevier, 1981, pp. 239–304.
[16] B. Ninness and S. Henriksen, "Bayesian system identification via Markov chain Monte Carlo techniques," *Automatica*, vol. 46, no. 1, pp. 40–51, 2010.
[17] J. L. Doob, "Application of the theory of martingales," *Le calcul des probabilites et ses applications*, pp. 23–27, 1949.
[18] A. W. Van der Vaart, *Asymptotic Statistics*. Cambridge University Press, 2000, vol. 3.
[19] E. Mazumdar, A. Pacchiano, Y.-A. Ma, P. L. Bartlett, and M. I. Jordan, "On Thompson sampling with Langevin algorithms," *arXiv preprint arXiv:2002.10002*, 2020.
[20] Y. Jedra and A. Proutiere, "Finite-time identification of stable linear systems optimality of the least-squares estimator," in *Proc. IEEE Conf. Decis. Control*. IEEE, 2020, pp. 996–1001.