



HAL
open science

PAC-Bayesian theory for stochastic LTI systems

Deividias Eringis, John Leth, Zheng-Hua Tan, Rafal Wisniewski, Alireza Fakhrizadeh Esfahani, Mihaly Petreczky

► **To cite this version:**

Deividias Eringis, John Leth, Zheng-Hua Tan, Rafal Wisniewski, Alireza Fakhrizadeh Esfahani, et al.. PAC-Bayesian theory for stochastic LTI systems. IEEE Conference on Decision and Control, Dec 2021, Austin, United States. hal-03331763

HAL Id: hal-03331763

<https://hal.archives-ouvertes.fr/hal-03331763>

Submitted on 2 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

PAC-Bayesian theory for stochastic LTI systems

Deividas Eringis, John Leth and Zheng-Hua Tan, Rafal Wisniewski
Alireza Fakhrizadeh Esfahani, Mihaly Petreczky

Abstract—In this paper we derive a PAC-Bayesian error bound for autonomous stochastic LTI state-space models. The motivation for deriving such error bounds is that they will allow deriving similar error bounds for more general dynamical systems, including recurrent neural networks. In turn, PAC-Bayesian error bounds are known to be useful for analyzing machine learning algorithms and for deriving new ones.

I. INTRODUCTION

The goal of this paper to present a PAC-Bayesian error bound for learning autonomous stochastic linear time-invariant (LTI) state-space representations (LTI systems for short). Autonomous stochastic LTI systems are widely used to model time series, they correspond to ARMA models. They represent one of the simplest classes of dynamical systems, and their learning theory has a rich history. In particular, the learning of stochastic LTI systems has been the subject of system identification [14]. Despite the large body of literature for learning LTI models, there has been no result on PAC-Bayesian error bounds for LTI models.

In order to explain what PAC-Bayesian framework is, it is useful to recall from [14] the learning problem for LTI systems. Recall from [14] that stochastic LTI models can be viewed both as generators of an output process (its distribution), and as predictors which predict the current value of an output process based on its past values. These two views are equivalent. In this paper we will view dynamical systems as predictors, which use past outputs to predict future outputs. Informally, following [14] the learning problem for autonomous stochastic LTI models can be formulated as follows. Consider a process y which is generated by a black-box system. The goal is to find an LTI model which uses the past of the observed process y to predict the current value of y , and such that the prediction error is the smallest possible. The prediction error is measured by using a so called loss function. There are two types of prediction error to consider: the generalization error and the empirical loss. The generalization error is mathematical expectation of the difference between the current values of y and the one predicted by the model, if the model can use all the past observations of y to predict future ones. The generalization error is a number which tells us how well the model performs in average. In contrast, the average empirical error is a random variable, each realization of which

corresponds to the actual prediction error for a finite number of observations sampled from y .

As a rule, algorithms for learning LTI models [14] chose the model parameters in such a manner that the empirical error is small, when evaluated for the data available for learning. However, the fact that a certain model renders the empirical error small does not necessarily mean that the generalization error will be small too. In fact, one of the main theoretical challenge is to prove this implication for certain classes of learning problems and algorithms.

Informally, the immediate goal of PAC-Bayesian framework [9], [15], [16] is to derive an upper bound on the average generalization error, which involves the average empirical error of models and a number of tuning parameters. The averaging is done on the space of all possible models, using a probability distribution on the space of models. One of the tuning parameters is the prior distribution on the space of inputs, which represents our hypothesis on the model structure; We will refer to this inequality as the PAC-Bayesian inequality. Once a PAC-Bayesian inequality has been derived, learning algorithms can be formulated as follows.

- 1) Choose the value of the prior distribution and other tuning parameters.
- 2) Find probability distribution which minimizes the PAC-Bayesian upper bound for the average generalization error. Since this upper bound involves the empirical loss, we will tend to assign higher probability to models which fit well the training data. However, this will be tempered by the need to take into account the prior distribution and the tuning parameters. This step can be performed using sampled data. The thus obtained probability distribution is analogous to the posterior distribution in the Bayesian setting.
- 3) The final model can be obtained by either sampling randomly from the posterior distribution, choosing the element with the maximal density (if the posterior distribution is absolutely continuous), or by taking the mean of the posterior distribution.

Various choices of the prior distribution on the space of model and of the tuning parameters which enter the PAC-Bayesian inequality will lead to different learning algorithms. In fact, it can be shown [7], [9] that known learning algorithms can be derived from PAC-Bayesian inequality. The thus derived algorithms will automatically be equipped with an upper bound on the generalization error. In turn, we can use this upper bound for a systematic analysis of corresponding learning algorithm.

The PAC-Bayesian learning theory [15], [16]) is one of the several approaches for studying theoretical properties of ma-

Deividas Eringis, John Leth, Zheng-Hua Tan and Rafal Wisniewski are with Dept. of Electronic Systems, Aalborg University, Denmark, {der,jjl,zt,raf}@es.aau.dk.

Mihaly Petreczky and Alireza Fakhrizadeh Esfahani is with Laboratoire Signal et Automatique de Lille (CRISTAL) Lille, France, mihaly.petreczky@centralelille.fr, alireza.fakhrizadeh.esfahani@gmail.com

chine learning algorithms. This approach is subject of intensive research, see [9] for a recent survey. It has demonstrated its ability to provide computable generalization bounds on popular machine learning algorithms, such as neural networks [6] and SVMs [3]. Moreover, as its name suggests, PAC-Bayesian framework combines the classical Probably Approximately Correct theory based on a frequentist view with Bayesian inference, see [2], [7], [8], [19], [20] for a more detailed discussion on the topic.

The motivation for studying PAC-Bayesian error bounds for LTI models is as follows. Autonomous stochastic LTI models are a special case of stochastic LTI models with inputs. In fact, it is known [11], [13], [14] that learning stochastic LTI models with inputs can be decomposed into learning autonomous stochastic LTI models and learning deterministic LTI models with inputs. Hence, in order to derive PAC-Bayesian error bounds for stochastic LTI models with inputs the first step is to derive such error bounds for autonomous stochastic LTI models. In turn, stochastic LTI models with inputs are special cases of recurrent neural networks (RNN) and their various modifications, such as LSTM networks. Hence, deriving PAC-Bayesian error bounds for autonomous stochastic LTI models is the first step towards deriving PAC-Bayesian error bounds for stochastic LTI models with inputs, and then for RNNs and LSTMs. In turn, this will allow to understand the theoretical limitations of learning algorithms for RNNs.

Despite an impressive body of literature, PAC-Bayesian error bounds are not available for learning LTI systems. One of the reasons for this is that most of the existing literature dealt with bounded and Lipschitz loss functions [1], [6], [18], which is typical for classification problem, but not for regression problems. In [7] PAC-Bayesian bounds for linear regression with a quadratic loss function was developed, later this bound was improved and extended to non i.i.d. data in [17]. In particular, in [17] the derived PAC-Bayesian error bound for non i.i.d linear regression problem was applied to learning ARX models. Note that in [17] control inputs were allowed, while in the current paper we consider only autonomous systems. Since ARX models without inputs correspond to AR models, and the latter class represents a special case of autonomous stochastic LTI state-space representations, the present paper can be viewed as an extension of [17] to stochastic LTI state-space representations.

The structure of the paper is as follows. In Section II-A we present the formal framework for PAC-Bayesian learning of linear dynamical systems. In Section III we present the main result of the paper which is an analytic PAC-Bayesian error bound for stochastic LTI models. In Section IV we present a numerical example for illustrating the theoretical results. In Section V we present the conclusions and directions for future research.

II. PAC-BAYESIAN LEARNING FOR LINEAR DYNAMICAL SYSTEMS

The purpose of this section is to formulate the learning problem for stochastic dynamical systems and to present a

brief overview of PAC-Bayesian framework.

A. The learning problem

In order to formalize the problem, let us denote by $\mathcal{Y} = \mathbb{R}^p$ the set of output values. Moreover, we fix a probability space $(\Omega, \mathbf{P}, \mathbf{F})$, where \mathbf{F} is a σ -algebra on Ω and \mathbf{P} is a probability measure on \mathbf{F} , see for example [12] for the terminology. We use \mathbf{E} to denote the mathematical expectation with respect to the probability measure \mathbf{P} . Consider a stochastic process \mathbf{y} on this probability space, taking values in \mathcal{Y} , and with time axis \mathbb{Z} . That is, for any $t \in \mathbb{Z}$, $\mathbf{y}(t) : \Omega \rightarrow \mathcal{Y}$; $\omega \mapsto \mathbf{y}(t)(\omega) = y$ is a random variable on $(\Omega, \mathbf{P}, \mathbf{F})$. For technical reasons, we assume that \mathbf{y} is

- Stationary: its finite dimensional distributions are invariant under time displacement

$$\begin{aligned} \mathbf{P}(\mathbf{y}(t_1 + t) \leq y_1, \dots, \mathbf{y}(t_k + t) \leq y_k) \\ = \mathbf{P}(\mathbf{y}(t_1) \leq y_1, \dots, \mathbf{y}(t_k) \leq y_k) \end{aligned}$$

for any t, t_1, \dots, t_k and k .

- Square-integrable: $\mathbf{E}[\|\mathbf{y}(t)\|_2^2] < \infty$ for any t , with $\|\cdot\|_2$ denoting the Euclidean 2-norm.
- Zero mean: $\mathbf{E}[\mathbf{y}(t)] = 0$ for any t .

Recall that the first two assumptions imply constant expectation and that the covariance matrix $Cov(\mathbf{y}(t), \mathbf{y}(s)) = \mathbf{E}[(\mathbf{y}(t) - \mathbf{E}[\mathbf{y}(t)])(\mathbf{y}(s) - \mathbf{E}[\mathbf{y}(s)])^T]$ only depend on $t - s$.

The goal of learning is to estimate models for predicting \mathbf{y} base on a training set $\{\mathbf{y}(t)(\omega)\}_{t=0}^N$ formed by a finite portion of a sample path of \mathbf{y} .

Let $\mathcal{Y}^* = \bigcup_{k=1}^{\infty} \mathcal{Y}^k$ where here \bigcup denote disjoint union. By slight abuse of notation we simply write $y = (y_1, \dots, y_k)$ in place of the more correct $(y, k) = ((y_1, \dots, y_k), k)$ for an element in \mathcal{Y}^* . Moreover, whenever necessary we associate with \mathcal{Y}^k the topology generated by $\|\cdot\|_2$, and the Borel σ -algebra \mathcal{B}_k generated by the open sets of \mathcal{Y}^k .

In our setting, models will be a subset of measurable functions of the form

$$f : \mathcal{Y}^* \rightarrow \mathcal{Y} \quad (1)$$

In particular, $f^{-1}(B) \in \mathcal{B}_k$ for all k whenever $B \in \mathcal{B}_1$. For the easy of notation, for any $(y_1, \dots, y_k) \in \mathcal{Y}^k$, we denote $f((y_1, \dots, y_k))$ by $f(y_1, \dots, y_k)$.

In this paper, we will restrict attention to models which arise from stable transfer functions. More precisely, we will use the following definition.

Definition II.1 (Model). *A model is a function f of the form (1) such that there exists a sequence of $p \times p$ matrices $\{M_k\}_{k=1}^{\infty}$ such that the sequence is absolutely summable, i.e., $\sum_{k=0}^{\infty} \|M_k\|_2 < \infty$, and*

$$f(y_1, \dots, y_k) = \sum_{i=1}^k M_i y_i \quad (2)$$

Intuitively, a model f satisfying the conditions of Definition II.1 are input-output maps of linear systems defined by a stable transfer function $G(z) = \sum_{k=1}^{\infty} M_k z^{-(k-1)}$.

Intuitively, if f is applied to some tuple of past value of \mathbf{y} , then the resulting random process will be viewed as a prediction of the current value of \mathbf{y} based on those past values, that is,

Notation 1. In the sequel we use the following notation:

$$\hat{\mathbf{y}}_f(t | s) = f(\mathbf{y}(t-1), \dots, \mathbf{y}(s)).$$

The random variable $\hat{\mathbf{y}}_f(t | s)$ is interpreted as the prediction of $\mathbf{y}(t)$ generated by f , based on the past values $\mathbf{y}(t-1), \dots, \mathbf{y}(s)$, $s < t$. Clearly,

$$\hat{\mathbf{y}}_f(t | s) = \sum_{k=1}^{t-s} M_k \mathbf{y}(t-k) \quad (3)$$

Remark II.1. Notice that since \mathbf{y} is stationary, it follows that the probability distribution of $\hat{\mathbf{y}}_f(t | s)$ depends only on $t-s$

Naturally, we would like to find models which predict the output process well. The “quality” of the predictor f is usually assessed through a measurable loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow [0, \infty)$. The loss function evaluated at $\ell(\hat{\mathbf{y}}_f(t | s), \mathbf{y}(t))$ measures how well the predicted value $\hat{\mathbf{y}}_f(t | s)$ approximates the true value of \mathbf{y} at t : the smaller $\ell(\hat{\mathbf{y}}_f(t | s), \mathbf{y}(t))$ is, the smaller the prediction error is. Note that the use of the word smaller here needs to be quantified e.g., by \mathbf{E} , as $\ell(\hat{\mathbf{y}}_f(t | s), \mathbf{y}(t))$ is a random variable. Indeed, in the sequel we will refer to $\mathbf{E}[\ell(\hat{\mathbf{y}}_f(t | s), \mathbf{y}(t))]$ as the prediction error

Assumption II.1 (Quadratic loss function). *In the sequel, we will assume that the loss function is quadratic, i.e. $\ell(y, y') = \|y - y'\|_2^2 = (y - y')^T (y - y')$.*

Some of the subsequent discussion, especially Theorem II.1 can be extended to other loss functions, but for the purposes of the paper quadratic loss functions are sufficient.

Then the quantity

$$\mathbf{E}[\ell(\hat{\mathbf{y}}_f(t | s), \mathbf{y}(t))], \quad (4)$$

measures the mean difference between the actual process $\mathbf{y}(t)$ and the predicted value $\hat{\mathbf{y}}_f(t | s)$ based on $\{\mathbf{y}(\tau)\}_{\tau=s}^{t-1}$. However, the expectation (4) is not the most convenient measure of the predictive power of a model, as it depends on the prediction horizon $t-s$. In practice, the prediction horizon tends to increase with the increase of the number of available data points. For this reason, it is more convenient to consider the prediction error as the beginning of the prediction horizon goes to infinity.

Intuitively, we expect that the quality of the prediction will increase with the growth of the horizon $t-s$ used for prediction. In fact, we can state the following.

Lemma II.1 (Infinite horizon prediction, [10]). *The limit*

$$\hat{\mathbf{y}}_f(t) = \lim_{s \rightarrow -\infty} \hat{\mathbf{y}}_f(t | s) \quad (5)$$

exists in the mean square sense for all t , the process $\hat{\mathbf{y}}_f(t)$ is stationary,

$$\mathbf{E}[\ell(\hat{\mathbf{y}}_f(t), \mathbf{y}(t))] = \lim_{s \rightarrow -\infty} \mathbf{E}[\ell(\hat{\mathbf{y}}_f(t | s), \mathbf{y}(t))] \quad (6)$$

and $\mathbf{E}[\ell(\hat{\mathbf{y}}_f(t), \mathbf{y}(t))]$ does not depend on t .

This prompts us to introduce the following definition.

Definition II.2 (Generalization error of a model). *The quantity*

$$\mathcal{L}_{\mathbf{y}}^{\ell}(f) = \mathbf{E}[\ell(\hat{\mathbf{y}}_f(t), \mathbf{y}(t))] = \lim_{s \rightarrow -\infty} \mathbf{E}[\ell(\hat{\mathbf{y}}_f(t | s), \mathbf{y}(t))]$$

is called the generalization error of the model f when applied to process \mathbf{y} , or simply generalization error, when \mathbf{y} is clear from the context.

Intuitively, $\hat{\mathbf{y}}_f(t)$ can be interpreted as the prediction of $\mathbf{y}(t)$ generated by the model f based on all past values of \mathbf{y} . As stated in Lemma II.1 we consider the special case when $\hat{\mathbf{y}}_f(t)$ is the mean-square limit of $\hat{\mathbf{y}}_f(t | s)$ as $s \rightarrow -\infty$. Clearly, for large enough $t-s$, the prediction error $\mathbf{E}[\ell(\hat{\mathbf{y}}_f(t | s), \mathbf{y}(t))]$ is close to the generalization error.

The goal of learning is to find a model from a set of possible models with the smallest possible prediction error, using a finite portion of the sample path of \mathbf{y} . This can be achieved through minimizing a cost function which involves the so called empirical error. Assume that we would like to learn a model from the time series $S = \{\mathbf{y}(t)(\omega)\}_{t=0}^N$, for some $\omega \in \Omega$. The data S represents a finite portion of sample path $\{\mathbf{y}(t)(\omega)\}_{t \in \mathbb{Z}}$ of \mathbf{y} , and $N+1$ represents the number of data points used for the learning problem. Let us define first the concept of empirical (error) loss.

Definition II.3. *The empirical loss for a model f and process \mathbf{y} is defined by*

$$\hat{\mathcal{L}}_{\mathbf{y}, N}^{\ell}(f) = \frac{1}{N} \sum_{t=1}^N \ell(\hat{\mathbf{y}}_f(t | 0), \mathbf{y}(t)). \quad (7)$$

Note that $\hat{\mathcal{L}}_{\mathbf{y}, N}^{\ell}(f) : \Omega \rightarrow \mathbb{R}$ is a random variable, and for $\omega \in \Omega$, $\hat{\mathcal{L}}_{\mathbf{y}, N}^{\ell}(f)(\omega)$ corresponds to the average prediction error produced by f when applied to the samples $\mathbf{y}(t)(\omega)$ for $t = 0, \dots, N$ successively, that is, if the training data is $S = \{\mathbf{y}(t)(\omega)\}_{t=0}^{N-1}$, then

$$\begin{aligned} \hat{\mathcal{L}}_{\mathbf{y}, N}^{\ell}(f)(\omega) &= \frac{1}{N} \sum_{t=1}^N \ell(\hat{\mathbf{y}}_f(t | 0)(\omega), \mathbf{y}(t)(\omega)) \\ &= \frac{1}{N} \sum_{t=1}^N \ell(f(\mathbf{y}(t-1)(\omega), \dots, \mathbf{y}(0)(\omega)), \mathbf{y}(t)(\omega)) \end{aligned}$$

i.e. $\hat{\mathcal{L}}_{\mathbf{y}}^{\ell}(f)(\omega)$ depends only on the training data S .

Remark II.2 (Alternatives for defining the empirical loss and relationship with system identification). Note that the minimization of the empirical loss is quite a common method in system identification. However, for theoretical analysis usually not (7) is taken, but rather

$$V_{\mathbf{y}, N}(f) = \frac{1}{N} \sum_{t=1}^N \ell(\hat{\mathbf{y}}_f(t), \mathbf{y}(t)). \quad (8)$$

That is, instead of the prediction error using finite past $\{\mathbf{y}(s)\}_{s=0}^N$ they use the infinite past $\{\mathbf{y}(s)\}_{s=-\infty}^N$. As it was

pointed out above, $\lim_{s \rightarrow -\infty} (\hat{\mathbf{y}}_f(t | s) - \hat{\mathbf{y}}_f(t)) = 0$, where the limit is understood in the mean square sense. Using this observation and a Cesaro-mean like argument it can be shown that

Lemma II.2. $\lim_{N \rightarrow \infty} \hat{\mathcal{L}}_{\mathbf{y},N}^\ell(f) - V_{\mathbf{y},N}(f) = 0$, where the limit is understood in the mean sense.

In particular, as mean convergence implies convergence in probability, it is clear that for large enough N , $\hat{\mathcal{L}}_{\mathbf{y},N}^\ell(f)$, $V_{\mathbf{y},N}(f)$ will be close enough with a large probability. That is, for practical purposes minimizing $V_{\mathbf{y},N}(f)$ versus $\hat{\mathcal{L}}_{\mathbf{y},N}^\ell(f)$ will not make a big difference. The reason for using $\hat{\mathcal{L}}_{\mathbf{y},N}^\ell(f)$ instead of $V_{\mathbf{y},N}(f)$ is that we are interested in deriving PAC-Bayesian bounds, which are not asymptotic. This is in contrast to the usual approach of system identification, which focuses on asymptotic results, i.e., behavior of the algorithms as $N \rightarrow \infty$.

Depending on various choices of the cost function, which always involves the empirical loss, different learning algorithms can be derived. The challenge is to be able to characterize the generalization error of the results of these learning algorithms, and to understand how the choice of the learning algorithm influences this generalization error.

B. PAC-Bayesian framework

To this end, in the machine learning community the so called PAC-Bayesian framework was proposed. Next we present the PAC-Bayesian framework in detail. Let \mathcal{F} be a set of models, sometimes referred to as the set of hypothesis. Assume that \mathcal{F} can be parametrized by a parameter set Θ , i.e., there is a bijection $\Pi : \Theta \rightarrow \mathcal{F}$. For simplicity, denote $\Pi(\theta)$ by f_θ , $\theta \in \Theta$. That is, f_θ is the model which corresponds to the parameter θ . Assume that (Θ, B_Θ, m) is a measure space, where B_Θ is a σ -algebra on Θ and m is a measure on (Θ, B_Θ) . Furthermore, assume that the function $\theta \mapsto \mathcal{L}_{\mathbf{y}}^\ell(f_\theta)$ is measurable and integrable with respect to (Θ, B_Θ, m) . Assume that the function $\Theta \times \Omega \ni (\theta, \omega) \mapsto \hat{\mathcal{L}}_{\mathbf{y},N}^\ell(f_\theta)$ is measurable and integrable with respect to the joint product measure $m \times \mathbf{P}$. Let ρ be a probability density function on the measure space (Θ, B_Θ, m) , and $g : \mathcal{F} \rightarrow \mathbb{R}$ be a map such that $\Theta \ni \theta \mapsto g(f_\theta)$ is measurable and absolutely integrable. We then denote by $E_{f \sim \rho} g(f)$ the integral

$$E_{f \sim \rho} g(f) = \int_{\theta \in \Theta} \rho(\theta) g(f_\theta) dm(\theta) \quad (9)$$

Moreover, by a slight abuse of notation and terminology we sometimes write $\rho(f)$ for the value $\rho(\theta)$ whenever $f = f_\theta$ for $\theta \in \Theta$, and say that ρ is a density on \mathcal{F} . With this notation and terminology in mind, we can state the following theorem.

Theorem II.1. ([2] and [7]) For any two densities π and $\hat{\rho}$ on \mathcal{F} , any $\delta \in (0, 1]$, and a real number $\lambda > 0$,

$$\mathbf{P} \left(\left\{ \omega \in \Omega \mid E_{f \sim \hat{\rho}} \mathcal{L}_{\mathbf{y}}^\ell(f) \leq E_{f \sim \hat{\rho}} \hat{\mathcal{L}}_{\mathbf{y},N}^\ell(f)(\omega) + \frac{1}{\lambda} \left[KL(\hat{\rho} \parallel \pi) + \ln \frac{1}{\delta} + \Psi_{\ell,\pi}(\lambda, n) \right] \right\} \right) > 1 - \delta, \quad (10)$$

where $KL(\hat{\rho} \parallel \pi)$ is the Kullback-Leibler divergence:

$$KL(\hat{\rho} \parallel \pi) = E_{f \sim \hat{\rho}} \ln \frac{\hat{\rho}(f)}{\pi(f)}$$

and

$$\Psi_{\ell,\pi}(\lambda, N) = \ln E_{f \sim \pi} \mathbf{E}[e^{\lambda(\mathcal{L}_{\mathbf{y}}^\ell(f) - \hat{\mathcal{L}}_{\mathbf{y},N}^\ell(f))}]$$

Intuitively, in Theorem II.1, π plays the role of prior distribution density function and $\hat{\rho}$ plays the role of any candidate to posterior distribution on the space of models. The Kullback-Leibler divergence measures the distance between the prior and posterior densities. We expect the term $\Psi_{\ell,\pi}(\lambda, N)$ to converge to zero, as $N \rightarrow \infty$. In fact, under suitable assumptions on \mathbf{y} , this will indeed be the case.

The numbers λ and δ are tuning parameters. The parameter δ regulates the magnitude of the probability via $1 - \delta$ and influence the upper bound on the generalization error through a term inversely proportional to δ . This means that the smaller we would like to get the probability that the error bound holds higher, the error bound increases too. Finally, λ is a tuning parameter which tells us how much we care about the empirical loss being small, for details see below.

The way to use Theorem II.1 is as follows. Consider the sampled training data $S = \{\mathbf{y}(t)(\omega)\}_{t=0}^N$. Then, with probability at least $1 - \delta$, the inequality

$$E_{f \sim \hat{\rho}} \mathcal{L}_{\mathbf{y}}^\ell(f) \leq E_{f \sim \hat{\rho}} \hat{\mathcal{L}}_{\mathbf{y},N}^\ell(f)(\omega) + \frac{1}{\lambda} \left[KL(\hat{\rho} \parallel \pi) + \ln \frac{1}{\delta} + \Psi_{\ell,\pi}(\lambda, N) \right] \quad (11)$$

holds. If δ is small enough, it means that we should be unlucky for (11) not to hold. Our goal should be to find posterior distribution $\hat{\rho}$ such that the right-hand side of (11) is small. For such a density $\hat{\rho}$ the average generalization error $E_{f \sim \hat{\rho}} \mathcal{L}_{\mathbf{y}}^\ell(f)$ will be small. This means that under suitable statistical hypothesis, if we either

- randomly sample a model f_* from the probability distribution determined by $\hat{\rho}$, or
- we choose the most likely model $f_* = \operatorname{argmax}_{f \in \mathcal{F}} \hat{\rho}(f)$, or
- we choose f_* to be the mean or median of the probability distribution with the density $\hat{\rho}$

Then $\mathcal{L}_{\mathbf{y}}^\ell(f_*)$ will be small too (at least with a high probability). That is, we have found a model with a small generalization error.

Since the terms of $\Psi_{\ell,\pi}(\lambda, n)$ and $\ln \frac{1}{\delta}$ in (11) do not depend on the the posterior density $\hat{\rho}$, choosing a density $\hat{\rho}$ for which

the right-hand side of (11) is small is equivalent to choosing $\hat{\rho}$ so that

$$E_{f \sim \hat{\rho}} \hat{\mathcal{L}}_{\mathbf{y}, N}^{\ell}(f)(\omega) + \frac{1}{\lambda} KL(\hat{\rho} || \pi) \quad (12)$$

is small.

By tuning the parameters λ and the prior π the cost function (12) will be different, leading to different choices of $\hat{\rho}$ and model f_* , i.e., to different learning algorithms. The prior π encodes our hypothesis on the model structure. The parameter λ regulates the extent we care about fitting training data. For instance, small λ mean that we would like the posterior $\hat{\rho}$ to be close to the prior (we insist on a certain model structure) and we care less about how well the models fit the training data. In contrast, large λ means that we care less about the model having a certain structure, and more about it fitting the training data.

The parameter δ allows us to regulate the probability that the inequality (11) is not true, and hence the found model will have a potentially arbitrary generalization error. Clearly, the price we pay for making sure that (11) is true with a high probability is that the upper bound on the generalization error is large. Finally, the error term $\Psi_{\ell, \pi}(\lambda, N)$ depends on λ , the prior π and N , for large enough N we expect it to be small.

In fact, we can present an explicit formula for the density $\hat{\rho}$ which minimizes (12): the unique density $\hat{\rho}^*$ which achieves the minimum of (12) is the Gibbs posterior [2]:

$$\hat{\rho}^*(f) = \frac{1}{Z_{\mathbf{y}}(\omega)} \pi(f) e^{-\lambda \hat{\mathcal{L}}_{\mathbf{y}, N}^{\ell}(f)(\omega)}, \quad (13)$$

where

$$Z_{\mathbf{y}}(\omega) = E_{f \sim \pi} e^{-\lambda \hat{\mathcal{L}}_{\mathbf{y}, N}^{\ell}(f)(\omega)} \quad (14)$$

is the normalization term. Note that $\hat{\rho}^*(f)$ is a function of the empirical loss $\hat{\mathcal{L}}_{\mathbf{y}}^{\ell}(f)(\omega)$ evaluated for the training data $S = \{\mathbf{y}(t)(\omega)\}_{t=0}^N$.

If we choose f_* as the most likely model, i.e., $f_* = \operatorname{argmax}_{f \in \mathcal{F}} \hat{\rho}^*(f)$, then we are solving the optimization problem:

$$\min_{f \in \mathcal{F}} \left(\lambda \hat{\mathcal{L}}_{\mathbf{y}, N}^{\ell}(f)(\omega) - \ln \pi(f) \right), \quad (15)$$

That is, we choose f_* by minimizing a cost function which includes the empirical loss and a regularization term. The latter is just the prior distribution density. This is a further indication that a large variety of seemingly different learning algorithms can be viewed as special cases of the PAC-Bayesian framework.

Remark II.3 (Using infinite past filters). Theorem II.1 and the discussion above remains true, if we replace $\hat{\mathcal{L}}_{\mathbf{y}, N}^{\ell}(f)$ by the alternative definition $V_{\mathbf{y}, N}(f)$ from (8).

III. PAC-BAYESIAN FRAMEWORK FOR STOCHASTIC LTI SYSTEMS

In this paper we will derive a PAC-Bayesian bound for learning stochastic LTI systems without inputs. More precisely, we will assume that data generating process \mathbf{y} is the output of

a stochastic LTI system, and that the class of models is that of the class of stochastic LTI systems without inputs.

This latter point necessitates some explanation. As it was pointed out in the beginning of the paper, stochastic dynamical systems can be interpreted as generators of stochastic processes, parametrizations of distributions and predictors, and in a sense these three points of view are equivalent.

In system identification, the objective of learning is formulated as finding a model which approximately generates the same observed behavior as the underlying physical system. In Section II we formulate the learning problem as that of finding the best possible predictor for \mathbf{y} . In fact, it turns out that for the case of stochastic LTI systems, these two objectives are essentially equivalent, in the sense that there is a correspondence between best predictor of \mathbf{y} and stochastic LTI systems whose output is \mathbf{y} . We will come back to this point after having set up the formal mathematical framework.

A stochastic LTI is a system of the form

$$\begin{aligned} \mathbf{x}(t+1) &= A\mathbf{x}(t) + B\mathbf{v}(t) \\ \tilde{\mathbf{y}}(t) &= C\mathbf{x}(t) + D\mathbf{v}(t) \end{aligned} \quad (16)$$

where $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $C \in \mathbb{R}^{p \times n}$, $D \in \mathbb{R}^{p \times m}$ for $n \geq 0$, $m, p > 0$ and \mathbf{x} , $\tilde{\mathbf{y}}$ and \mathbf{v} are zero-mean square-integrable stationary stochastic process with values in \mathbb{R}^n , \mathbb{R}^p , and \mathbb{R}^m respectively. The processes \mathbf{x} , \mathbf{y} and \mathbf{v} are called state, output and noise process, respectively. Furthermore, we require that A is stable (all its eigenvalues are inside the open unit circle) and that for any $t, k \in \mathbb{Z}$, $k \geq 0$, $E[\mathbf{v}(t)\mathbf{v}^T(t-k-1)] = 0$, $E[\mathbf{v}(t)\mathbf{x}^T(t-k)] = 0$, i.e., $\mathbf{v}(t)$ is white noise and uncorrelated with $\mathbf{x}(t-k)$. We identify the system (16) with the tuple (A, B, C, D, \mathbf{v}) ; note that the state process \mathbf{x} is uniquely defined by the infinite sum $\mathbf{x}(t) = \sum_{k=1}^{\infty} A^{k-1} B \mathbf{v}(t-k)$. We say that the stochastic LTI system is a realization of the process \mathbf{y} , if its output $\tilde{\mathbf{y}}$ coincides with \mathbf{y} , i.e., $\mathbf{y}(t) = \tilde{\mathbf{y}}(t) = C\mathbf{x}(t) + D\mathbf{v}(t)$.

It is well-known [13], that if \mathbf{y} has a realization by a LTI system, then it has a realization by a minimal LTI system in forward innovation form, i.e., and LTI system (A, K, C, I, \mathbf{e}) such that I is the identity matrix, (A, K) is controllable, (C, A) is observable, and \mathbf{e} is the innovation process of \mathbf{y} .

To recall the notion of an innovation process, we need the following. Denote by \mathcal{H} the Hilbert space of zero-mean square-integrable real valued random variables, where the inner product between two random variables y, z is $E[yz]$. The Hilbert space generated by a set $U \subset \mathcal{H}$ is the smallest (w.r.t. set inclusion) closed subspace of \mathcal{H} which contains U . We denote by $E_l[z|U]$ the orthogonal projection of z onto U , see [13] for details. For a zero-mean square-integrable stochastic process $\mathbf{z}(t) \in \mathbb{R}^p$, we let $E_l[\mathbf{z}(t)|U] = [\hat{\mathbf{z}}_1(t), \dots, \hat{\mathbf{z}}_k(t)]^T$, where $\hat{\mathbf{z}}_i(t) = E_l[\mathbf{z}_i(t)|U]$, $i = 1, \dots, k$. That is, $E_l[\mathbf{z}(t)|U]$ is the random variable with values in \mathbb{R}^p obtained by projecting the coordinates of $\mathbf{z}(t)$ onto U . Accordingly, the orthogonality of a multidimensional random variable to a closed subspace in \mathcal{H} is meant element-wise.

For the stochastic process \mathbf{y} denote by $\mathcal{H}_{t-}^{\mathbf{y}}$ the closed subspace of \mathcal{H} generated by the set $\{\ell^T \mathbf{y}(s) \mid s \in \mathbb{Z}, s < t, \ell \in \mathbb{R}^p\}$ formed by the past values of \mathbf{y} . We call the process

$$\mathbf{e}(t) := \mathbf{y}(t) - E_l[\mathbf{y}(t) | \mathcal{H}_{t-}^{\mathbf{y}}], \quad \forall t \in \mathbb{Z}$$

the (forward) innovation process of \mathbf{y} , that is, the innovation process is the difference between \mathbf{y} and its orthogonal projection to its past. The variance $E[\mathbf{e}^T(t)\mathbf{e}(t)]$ gives us the minimal prediction error which is attainable by trying to predict $\mathbf{y}(t)$ as a linear function of the past values $\{\mathbf{y}(s)\}_{s < t}$.

Minimal realizations in forward innovation form of a process are unique up to a linear isomorphism. We will say that the process \mathbf{y} is coercive, if it has a minimal realization (A, K, C, I, \mathbf{e}) in forward innovation form such that $A - KC$ is stable. Since all minimal realization of \mathbf{y} in forward innovation form are isomorphic, the stability of $A - KC$ will hold for any minimal realization of \mathbf{y} .

Assumption III.1. *In the sequel we assume that \mathbf{y} can be realized by a stochastic LTI system and \mathbf{y} is coercive.*

Coercivity of \mathbf{y} implies that any minimal realization (A, K, C, I, \mathbf{e}) of \mathbf{y} in forward innovation form, defines a linear map $\{\mathbf{y}(s)\}_{s=-\infty}^{t-1} \mapsto \hat{\mathbf{y}}(t)$ as follows:

$$\mathbf{x}(t+1) = (A - KC)\mathbf{x}(t) + K\mathbf{y}(t), \quad \hat{\mathbf{y}}(t) = C\mathbf{x}(t). \quad (17)$$

In fact, $\hat{\mathbf{y}}(t)$ is the projection of $\mathbf{y}(t)$ onto the Hilbert-space generated by the past values $\{\mathbf{y}(s)\}_{s=-\infty}^{t-1}$ of \mathbf{y} , i.e., $\hat{\mathbf{y}}(t) = E_l[\mathbf{y}(t) | \mathcal{H}_{t-}^{\mathbf{y}}]$. Note that $\mathbf{y}(t) - \hat{\mathbf{y}}(t) = \mathbf{e}(t)$, i.e., $E[\|\mathbf{y}(t) - \hat{\mathbf{y}}(t)\|_2^2]$ is minimal among all linear mappings $\{\mathbf{y}(s)\}_{s=-\infty}^{t-1} \mapsto \hat{\mathbf{y}}(t)$, or all such mappings, if \mathbf{y} is Gaussian. That is, any minimal realization of \mathbf{y} in forward innovation form yields an optimal predictor.

The discussion above prompts us to define the set of models as follows. Let $\Sigma = (\hat{A}, \hat{K}, \hat{C})$ be a tuple of matrices such that \hat{A} and $\hat{A} - \hat{K}\hat{C}$ are stable. Define the map $f_{\Sigma} : \mathcal{Y}^* \rightarrow \mathcal{Y}$ as follows:

$$f_{\Sigma}(y_1, \dots, y_k) = \sum_{t=1}^k \hat{C}(\hat{A} - \hat{K}\hat{C})^{k-t} \hat{K}y_t \quad (18)$$

In other words, $\hat{\mathbf{y}}_{f_{\Sigma}}(t | s)$ is the output at time t of the following dynamical system:

$$\begin{aligned} \mathbf{x}_{\Sigma}(t+1 | s) &= (\hat{A} - \hat{K}\hat{C})\mathbf{x}_{\Sigma}(t | s) + \hat{K}\mathbf{y}(t), \quad \mathbf{x}_{\Sigma}(s | s) = 0 \\ \hat{\mathbf{y}}_{f_{\Sigma}}(t | s) &= \hat{C}\mathbf{x}_{\Sigma}(t | s) \end{aligned} \quad (19)$$

That is f_{Σ} is the input-output map of the deterministic LTI system $(\hat{A} - \hat{K}\hat{C}, \hat{K}, \hat{C})$.

It is clear that the function f_{Σ} is a model in the sense of Definition II.1, since f_{Σ} is generated by matrices $M_k = \hat{C}(\hat{A} - \hat{K}\hat{C})^{k-1} \hat{K}$ which are Markov parameters of the stable deterministic LTI system $((\hat{A} - \hat{K}\hat{C}), \hat{K}, \hat{C})$, and hence the sequence $\{M_k\}_{k=1}^{\infty}$ is absolutely summable.

In particular, the limit $\hat{\mathbf{y}}_{f_{\Sigma}}(t) = \lim_{s \rightarrow -\infty} \hat{\mathbf{y}}_{f_{\Sigma}}(t | s)$ exists in the mean square sense uniformly in t , and

$$\hat{\mathbf{y}}_{f_{\Sigma}}(t) = \sum_{k=1}^{\infty} \hat{C}(\hat{A} - \hat{K}\hat{C})^{k-1} \hat{K}\mathbf{y}(t-k), \quad (20)$$

in other words, $\hat{\mathbf{y}}_{f_{\Sigma}}(t)$ is the output of the following dynamical system

$$\begin{aligned} \mathbf{x}_{\Sigma}(t+1) &= (\hat{A} - \hat{K}\hat{C})\mathbf{x}_{\Sigma}(t) + \hat{K}\mathbf{y}(t) \\ \hat{\mathbf{y}}_{f_{\Sigma}}(t) &= \hat{C}\mathbf{x}_{\Sigma}(t) \end{aligned} \quad (21)$$

Recall from Section II that for the formulation of PAC-Bayesian framework we need to define a parametrized set of models. In order to do so, we remark that two different tuples Σ and Σ' may yield the same model $f_{\Sigma} = f_{\Sigma'}$. However, this source of non-uniqueness can be eliminated by using minimal tuple. Let us call the tuple $\Sigma = (\hat{A}, \hat{K}, \hat{C})$ minimal, if (\hat{A}, \hat{K}) is a controllable pair, (\hat{A}, \hat{C}) is an observable pair, and $\hat{A}, \hat{A} - \hat{K}\hat{C}$ are stable. From standard realization theory we can derive the following simple result:

Lemma III.1.

- If $\Sigma = (\hat{A}, \hat{K}, \hat{C})$ is a tuple such that \hat{A} and $\hat{A} - \hat{K}\hat{C}$ are stable, then there exists a minimal tuple Σ_m such that $f_{\Sigma} = f_{\Sigma_m}$.
- If $\Sigma = (\hat{A}, \hat{K}, \hat{C})$ and $\Sigma' = (\hat{A}', \hat{K}', \hat{C}')$ are two minimal tuples, then $f_{\Sigma} = f_{\Sigma'}$ implies that Σ and Σ' are isomorphic, i.e. there exists a nonsingular matrix T , such that $\hat{A}' = T\hat{A}T^{-1}$, $\hat{K}' = T\hat{K}$, and $\hat{C}' = \hat{C}T^{-1}$.

That is, it is enough to consider models which arise from minimal tuples. Moreover, we will consider parametrizations of models which arise via parametrizing minimal tuples and in which there are no two isomorphic minimal tuples.

More precisely, we consider a family of models \mathcal{F} which satisfies the following assumption.

Assumption III.2. *There exists a measure space (Θ, \mathcal{B}, m) and a map*

$$\Sigma : \Theta \rightarrow (\mathbb{R}^{n \times n} \times \mathbb{R}^{n \times p} \times \mathbb{R}^{p \times n})$$

such that the set of models is defined by

$$\mathcal{F} = \{f_{\Sigma(\theta)} \mid \theta \in \Theta\},$$

and the following holds.

- For any $\theta \in \Theta$, $\Sigma(\theta)$ is a minimal tuple,
- For any $\theta_1, \theta_2 \in \Theta$, the tuples $\Sigma(\theta_1)$ and $\Sigma(\theta_2)$ are not isomorphic, i.e., the map $\Theta \ni \theta \mapsto f_{\Sigma(\theta)} \in \mathcal{F}$ is a one-to-one function.
- There exists θ_0 such that $\Sigma(\theta_0) = (A_0, K_0, C_0)$ and $(A_0, K_0, C_0, I, \mathbf{e})$ is a minimal realization of \mathbf{y} in forward innovation form

The intuition behind these assumptions is as follows. The set \mathcal{F} of models under consideration arises by parametrizing a subset of tuples which then give rise to models as defined in (18). The first assumption says that we parametrize \mathcal{F} in such a manner that each parameter value corresponds to a minimal tuple. This is not a real constraints, as if for some parameter values the corresponding tuple of matrices is not a minimal tuple, then by Lemma III.1 we can replace it by a minimal tuple which gives rise to the same model. Hence, such a replacement will not change the set of models \mathcal{F} . The

second condition means that there are no two parameter values which give rise to the same input-output map. Finally, the third assumption says that \mathcal{F} contains a model arising from a realization of \mathbf{y} . The latter point is quite useful in terms of relating the learning problem formulated in Section II with the usual formulation of the system identification problem.

That is, the model which arises from the realization of \mathbf{y} is the model with the least generalization error. In fact,

Lemma III.2. *If*

$$\theta_* = \operatorname{argmin}_{\theta \in \Theta} \mathcal{L}_{\mathbf{y}}^{\ell}(f_{\theta}),$$

then $(A_0, K_0, C_0, I, \mathbf{e})$ is a minimal realization of \mathbf{y} , where $\Sigma(\theta_*) = (A_0, K_0, C_0)$.

Lemma III.2 means that trying to find a model \mathcal{F} based on sampled data is equivalent to finding the element of \mathcal{F} which corresponds to a realization of \mathbf{y} . In fact, even if we cannot find the model with the smallest generalization error exactly, this can still be interpreted as finding a stochastic LTI system output of which is close to \mathbf{y} . More precisely, if the parameter set Θ is a topological space and if $\Sigma(\theta)$ is continuous in θ and if we find $\theta \in \Theta$ such that θ is close to θ_0 , then the matrices of $\Sigma(\theta) = (\hat{A}, \hat{K}, \hat{C})$ will be close to $\Sigma(\theta_0) = (A_0, K_0, C_0)$. It then follows that the output of the stochastic LTI system $(\hat{A}, \hat{K}, \hat{C}, \mathbf{e})$ will be close to \mathbf{y} .

Now we are able to state the main result on PAC-Bayesian error bound for stochastic LTI systems.

Theorem III.1. *Assume that \mathbf{y} and \mathcal{F} satisfy Assumption III.1–III.2 and that \mathbf{y} is jointly Gaussian. For any two densities π and $\hat{\rho}$ on \mathcal{F} , $\delta \in (0, 1]$, and a real number $\lambda > 0$, (10) holds, where*

$$\Psi_{\ell, \pi}(\lambda, N) \leq u_{\ell, \pi}(\lambda, N) \quad (22)$$

with

$$u_{\ell, \pi}(\lambda, N) = \ln E_{\theta \sim \pi} \left[\frac{\exp(\lambda v(\theta))}{\left(1 + \frac{\lambda \rho_N(\theta)}{2}\right)^{\frac{pN}{2}}} \right], \quad (23)$$

and the number $v(\theta)$ and $\rho_N(\theta)$ are defined as follows:

$$v(\theta) = \mathcal{L}_{\mathbf{y}}^{\ell}(f_{\theta}) \quad (24)$$

and

$$\rho_N(\theta) = \min(\operatorname{eig}(\mathbf{E}[\mathbf{z}_{1:N}(f_{\theta})\mathbf{z}_{1:N}^T(f_{\theta})])) \quad (25)$$

with

$$\mathbf{z}_{1:N}(f_{\theta}) = \begin{bmatrix} (\mathbf{y}(1) - \hat{\mathbf{y}}_{f_{\theta}}(1 | 0)) \\ \vdots \\ (\mathbf{y}(N) - \hat{\mathbf{y}}_{f_{\theta}}(N | 0)) \end{bmatrix} \quad (26)$$

Proof Theorem III.1. The proof follows the same lines as that of [17, Theorem 2], for the sake of completeness we repeat the basic steps. From Theorem II.1 it follows that (11) holds with probability at least $1 - \delta$. Consider a random variable

$\mathbf{z}_i(f_{\theta}) = \mathbf{y}(i) - \hat{\mathbf{y}}_{f_{\theta}}(i | 0)$. Just like in the proof of [17, Theorem 2],

$$\begin{aligned} \Psi_{\ell, \pi}(\lambda, N) &= \ln E_{f \sim \pi} \mathbf{E}[e^{\lambda(\mathcal{L}_{\mathbf{y}}^{\ell}(f) - \hat{\mathcal{L}}_{\mathbf{y}, N}^{\ell}(f))}] = \\ &= \ln E_{f \sim \pi} e^{\lambda \mathcal{L}_{\mathbf{y}}^{\ell}(f)} \mathbf{E}\left[e^{-\lambda \hat{\mathcal{L}}_{\mathbf{y}, N}^{\ell}(f)}\right] = \\ &= \ln E_{f \sim \pi} e^{\lambda \mathcal{L}_{\mathbf{y}}^{\ell}(f)} \mathbf{E}\left[e^{-\frac{\lambda}{N} \sum_{i=1}^N \mathbf{z}_i^T(f_{\theta}) \mathbf{z}_i(f_{\theta})}\right] \end{aligned} \quad (27)$$

Since $\mathbf{z}_i(f_{\theta})$ is a linear combination of $\{\mathbf{y}(s)\}_{s=0}^i$, and \mathbf{y} is zero mean Gaussian process, it follows that $\mathbf{z}_i(f_{\theta})$ is zero mean Gaussian with covariance $Q_{\mathbf{z}_i(f_{\theta})} = \mathbf{E}[\mathbf{z}_i(f_{\theta})\mathbf{z}_i^T(f_{\theta})]$. Notice that the variable $\mathbf{z}_{1:N}(f_{\theta})$ from (26) is of the form

$$\mathbf{z}_{1:N}(f_{\theta}) = \begin{bmatrix} \mathbf{z}_1(f_{\theta}) \\ \vdots \\ \mathbf{z}_N(f_{\theta}) \end{bmatrix}. \text{ Define}$$

$$S(f_{\theta}) = Q_{\mathbf{z}_{1:N}(f_{\theta})}^{-1/2} \mathbf{z}_{1:N}(f_{\theta})$$

where $Q_{\mathbf{z}_{1:N}(f_{\theta})} = \mathbf{E}[\mathbf{z}_{1:N}(f_{\theta})\mathbf{z}_{1:N}^T(f_{\theta})]$ and let S_i be the i th entry of S , i.e., $S = [S_1 \ \dots \ S_{pN}]^T$. Then from (27) it follows that

$$\begin{aligned} \sum_{i=1}^N \mathbf{z}_i^T(f_{\theta}) \mathbf{z}_i(f_{\theta}) &= \mathbf{z}_{1:N}(f_{\theta})^T \mathbf{z}_{1:N}(f_{\theta}) = \\ &= \mathbf{z}_{1:N}(f_{\theta})^T Q_{\mathbf{z}_{1:N}(f_{\theta})}^{-1/2} Q_{\mathbf{z}_{1:N}(f_{\theta})} Q_{\mathbf{z}_{1:N}(f_{\theta})}^{-1/2} \mathbf{z}_{1:N}(f_{\theta}) = \\ &S^T(f_{\theta}) Q_{\mathbf{z}_{1:N}(f_{\theta})} S(f_{\theta}) \geq S^T(f_{\theta}) S(f_{\theta}) \rho_N(\theta) = \\ &\left(\sum_{i=1}^{pN} S_i(f_{\theta})^2\right) \rho_N(\theta) \end{aligned} \quad (28)$$

It then follows that

$$\begin{aligned} e^{-\frac{\lambda}{N} \sum_{i=1}^N \mathbf{z}_i^T(f_{\theta}) \mathbf{z}_i(f_{\theta})} &= e^{-\frac{\lambda}{N} \mathbf{z}_{1:N}^T(f_{\theta}) \mathbf{z}_{1:N}(f_{\theta})} \leq \\ &e^{-\frac{\lambda}{N} \rho_N(\theta) (\sum_{i=1}^{pN} S_i^2(f_{\theta}))} \end{aligned} \quad (29)$$

Notice now that $S(f_{\theta})$ is Gaussian and zero mean, with covariance

$$\begin{aligned} E[S(f_{\theta})S^T(f_{\theta})] &= Q_{\mathbf{z}_{1:N}(f_{\theta})}^{-1/2} E[\mathbf{z}_{1:N}(f_{\theta})\mathbf{z}_{1:N}^T(f_{\theta})] Q_{\mathbf{z}_{1:N}(f_{\theta})}^{-1/2} = \\ &= Q_{\mathbf{z}_{1:N}(f_{\theta})}^{-1/2} Q_{\mathbf{z}_{1:N}(f_{\theta})} Q_{\mathbf{z}_{1:N}(f_{\theta})}^{-1/2} = I \end{aligned}$$

That is, the random variables $S_i(f_{\theta})$ are normally distributed and $S_i(f_{\theta}), S_j(f_{\theta})$ are independent, and therefore $\sum_{i=1}^{pN} S_i^2(f_{\theta})$ has χ^2 distribution. Hence,

$$\mathbf{E}\left[e^{-\frac{\lambda \rho_N(\theta)}{N} (\sum_{i=1}^{pN} S_i^2(f_{\theta}))}\right] = \frac{1}{\left(1 + \frac{\lambda \rho_N(\theta)}{N}\right)^{\frac{pN}{2}}}$$

Combining this with (29) and (27) implies the statement of the theorem. \square

A. Discussion of the error bound

Below we will make a number of remarks concerning the error bound (22).

a) *Asymptotic convergence of $\Psi_{\ell,\pi}(\lambda, N)$:*

Lemma III.3. *For some subsequence N_k , $\lim_{k \rightarrow \infty} N_k = \infty$, $\lim_{k \rightarrow \infty} \Psi_{\ell,\pi}(\lambda, N_k) = 0$.*

This result indicates that $\Psi_{\ell,\pi}(\lambda, N)$ cannot converge to infinity or to a non-zero constant. This also indicates the possibility that $\Psi_{\ell,\pi}(\lambda, N)$ might converge to zero as $N \rightarrow \infty$.

b) *Behavior of the upper bound as $N \rightarrow \infty$, conservativity of the result:* Lemma (III.3) also suggests that the upper bound (22) is a conservative one, as the upper bound does not converge to 0. In fact, the following holds.

Lemma III.4. *The sequence $\rho_N(\theta)$ is decreasing and $\lim_{N \rightarrow \infty} \rho_N(\theta) = \rho^*(\theta) > \mu_* > 0$. In particular, if the spectral density of \mathbf{y} is denoted by $\Phi_{\mathbf{y}}$, then $\mu_* > 0$ can be taken such that $\Phi_{\mathbf{y}}(e^{i\theta}) > \mu_* I$ for all $\theta \in [-\pi, \pi]$. In particular, in this case $v(\theta) > p\mu_*$.*

Lemma III.4 says that $\rho_N(\theta)$ is bounded from below.

Lemma III.5.

$$\ln E_{\theta \sim \pi} [e^{\lambda(v(\theta) - p\rho_0)}] \leq \lim_{N \rightarrow \infty} u_{\ell,\pi}(\lambda, N) \leq \ln E_{\theta \sim \pi} [e^{\lambda(v(\theta) - p\mu_*)}],$$

where μ_* is as in Lemma III.4, and ρ_0 is the smallest eigenvalue of $\mathbf{E}[\mathbf{y}(t)\mathbf{y}^T(t)]$

Lemma III.5 together with Lemma III.3 indicates that the upper bound of Theorem III.1 might be quite conservative: while $\Psi_{\ell,\pi}(\lambda, N)$ converges to zero for a suitable strictly increasing sequence of integers N_k , the upper bound $u_{\ell,\pi}(\lambda, N)$ is bounded from below in the limit. Note that $\mathbf{y}(t)$ is an i.i.d process with variance σI_p , then $\rho_0 = \sigma = \mu_*$ and the upper and lower bounds of Lemma III.5 coincide.

The problem of finding a tight PAC-Bayesian upper bound for LTI systems seems to be a challenging problem. In fact, the first PAC-Bayesian upper bound for linear regression [7] with i.i.d. data had the same drawback, namely, the part of the upper bound on the term which corresponds to $\Psi_{\ell,\pi}(\lambda, N)$ did not converge to zero as $N \rightarrow \infty$. This drawback for eliminated in [17] for linear regression with i.i.d data, but not for linear regression with non i.i.d. data. In the light of this development, the conservativity of our error bound is not surprising.

Upper bound not involving N and expectation over the prior: Finally, we can formulate a non-asymptotic upper bound similar to that of Lemma III.5

Lemma III.6.

$$\begin{aligned} u_{\ell,\pi}(\lambda, N) &\leq \ln E_{\theta \sim \pi} \left[e^{\lambda(v(\theta) - \frac{p\mu_* p N * 0.5}{\lambda p \mu_* + 0.5 p N})} \right] \leq \\ &\leq E_{\theta \sim \pi} [e^{\lambda(v(\theta))}] - \frac{\lambda p \mu_* 0.5}{\lambda p \mu_* + 0.5} \end{aligned}$$

where μ_* is as in Lemma III.4

The result of Lemma III.5 – III.6 could be used to derive further PAC-Bayesian error bounds. In fact,

$$\lim_{N \rightarrow \infty} \ln E_{\theta \sim \pi} [e^{\lambda(v(\theta) - p\mu_*)}] = \ln E_{\theta \sim \pi} [e^{\lambda(v(\theta))}] - p\mu_*,$$

so any upper bound on $\ln E_{\theta \sim \pi} [e^{\lambda(v(\theta))}]$ yields an upper bound on the limit of $u_{\ell,\pi}(\lambda, N)$. Alternatively, we can also use the non-asymptotic bound from Lemma III.6. In particular, if

$$K_1 < v(\theta) < K_2$$

for every θ , then

$$K_1 - p\rho_0 \leq \lim_{N \rightarrow \infty} u_{\ell,\pi}(\lambda, N) \leq K_2 - p\mu_* \quad (30)$$

$$u_{\ell,\pi}(\lambda, N) \leq K_2 - \frac{0.5 p \mu_*}{p \mu_* + 0.5} \quad (31)$$

The latter bounds do not depend on the prior, but only the bounds K_1, K_2 and the constants $\rho_0 > \mu_*$, which, in turn, represent some mild assumptions on the distribution of \mathbf{y} . Furthermore, it is easy to see from the discussion in Section III-B that $v(\theta)$ is a continuous function of θ , hence, for a compact parameter set Θ the bounds K_1, K_2 will exist and they could be estimated based on the parametrization. That is, (30) could readily be applied in this case.

In turn, (30) together with Theorem III.1 and (10) yields bounds for the average generalization error of the posterior for a certain choice of N and λ .

c) *Using the alternative definition of empirical loss involving infinite past filters:* Recall from Remark II.2 the alternative definition of empirical loss, which uses the infinite past prediction $\hat{\mathbf{y}}_f(t)$, instead of the finite past prediction $\hat{\mathbf{y}}_f(t | 0)$. We could have stated Theorem II.1 with $V_{\mathbf{y},N}(f_\theta)$ instead of $\hat{\mathcal{L}}_{\mathbf{y},N}(f_\theta)$ and we could have stated Theorem III.1

too, with the difference that $\mathbf{z}_{1:N}(f_\theta) = \begin{bmatrix} (\mathbf{y}(1) - \hat{\mathbf{y}}_{f_\theta}(1)) \\ \vdots \\ (\mathbf{y}(N) - \hat{\mathbf{y}}_{f_\theta}(N)) \end{bmatrix}$.

In particular, with these changes in mind the discussion above still holds, but the proofs simplify a lot. In particular, $E[\mathbf{z}_{1:N}(f_\theta)\mathbf{z}_{1:N}(f_\theta)^T]$ has the Toeplitz-matrix structure, as $\mathbf{y}(t) - \hat{\mathbf{y}}_{f_\theta}(t)$ is a wide-sense stationary process.

d) *PAC-Bayesian error bounds and system identification:*

As it was pointed out in Subsection II, the derived error bound together with the Gibbs posterior distribution (13) leads to learning algorithms. In particular, choosing a model f which maximizes the likelihood of the Gibbs posterior amounts to minimizing the cost function (15). In particular, for $\lambda = 1$ and π being the uniform distribution, the cost function (15) becomes the empirical loss. The algorithm which calculates the model minimizing (15) then corresponds to prediction error minimization [14]. Note that in the theoretical analysis of the prediction error approach, usually the prediction error (8) using the infinite past is used. As it was pointed out in Remark II.2, for large enough N , the latter prediction error is close to the empirical loss with a high probability. In fact, in the practical implementation of prediction error method usually the empirical error (7) is used. The various other choices of λ and the prior π when minimizing (15) amount to minimizing the prediction error with a regularization term. Other methods for sampling from the posterior will lead to different system identification algorithms. Investigating the relationship between the thus arising algorithms and the existing system identification algorithms remains a topic of future research.

B. Computing $u_{\ell,\pi}(\lambda, N)$

The value of $u_{\ell,\pi}(\lambda, N)$ can be computed using sequential Monte-Carlo [4]. In details, sample a set of parameters $\Theta = \{\theta_i\}_{i=1}^{N_f}$, according to the prior distribution $\theta_i \sim \pi(\theta)$. Then for large enough N_f the right-hand side of (22) can be approximated by

$$\ln \left(\frac{1}{N_f} \sum_{i=1}^{N_f} \frac{\exp(\lambda v(\theta_i))}{\left(1 + \frac{\lambda \rho^*(\theta_i)}{\frac{N}{2}}\right)^{\frac{N}{2}}} \right) \quad (32)$$

For every θ_i drawn from the prior distribution $\pi(\theta)$, the quantities $v(\theta_i)$ and $\rho^*(\theta_i)$ can be calculated as follows. Given the system in forward innovation form (A_0, K_0, C_0) and a predictor f_{θ_i} with

$$\Sigma(\theta_i) = (\hat{A}, \hat{K}, \hat{C}) = (\hat{A}(\theta_i), \hat{K}(\theta_i), \hat{C}(\theta_i)).$$

Define a system $(A_e, K_e, C_e, I, \mathbf{e})$, whose output is the error $\mathbf{z}(t) = \mathbf{y}(t) - \hat{\mathbf{y}}_{f_{\theta_i}}(t | 0)$,

$$\begin{aligned} \underbrace{\begin{bmatrix} \mathbf{x}(t+1) \\ \hat{\mathbf{x}}(t+1) \end{bmatrix}}_{\tilde{\mathbf{x}}(t+1)} &= \underbrace{\begin{bmatrix} A_0 & 0 \\ \hat{K}C_0 & \hat{A} - \hat{K}\hat{C} \end{bmatrix}}_{A_e} \underbrace{\begin{bmatrix} \mathbf{x}(t) \\ \hat{\mathbf{x}}(t) \end{bmatrix}}_{\tilde{\mathbf{x}}(t)} + \underbrace{\begin{bmatrix} K_0 \\ \hat{K} \end{bmatrix}}_{K_e} \mathbf{e}(t), \\ \mathbf{z}(t) &= \underbrace{\begin{bmatrix} C_0 & -\hat{C} \end{bmatrix}}_{C_e} \begin{bmatrix} \mathbf{x}(t) \\ \hat{\mathbf{x}}(t) \end{bmatrix} + \mathbf{e}(t), \\ \begin{bmatrix} \mathbf{x}(0) \\ \hat{\mathbf{x}}(0) \end{bmatrix} &= \begin{bmatrix} \sum_{k=1}^{\infty} A_0^{k-1} K_0 \mathbf{e}(-k) \\ 0 \end{bmatrix}, \\ P_0 &= \begin{bmatrix} P_x & 0 \\ 0 & 0_{n \times n} \end{bmatrix}, \end{aligned} \quad (33)$$

where $P_x = E[\mathbf{x}(0)\mathbf{x}(0)^T]$ is the steady-state covariance, which satisfies the Lyapunov equation

$$P_x = A_0 P_x A_0^T + K_0 Q_e K_0^T, \quad (34)$$

with $Q_e = E[\mathbf{e}(t)\mathbf{e}(t)^T]$. In practice, the numerical value of the initial condition $\mathbf{x}(0)$ is chosen based on the particular application. Then $v(\theta)$ corresponds to the second moment of the random variable $\mathbf{z}(t)$

$$\begin{aligned} v(\theta) &= E \left[\sum \mathbf{z}_i^2 \right] = \text{trace}(E[\mathbf{z}(t)\mathbf{z}(t)^T]) = \\ &= \text{trace}(C_e P_e C_e^T + Q_e). \end{aligned} \quad (35)$$

The covariance matrix P_e is the solution to Lyapunov equation

$$P_e = A_e P_e A_e^T + K_e Q_e K_e^T. \quad (36)$$

Now, define the joint covariance matrix $Q_{\mathbf{z},N}$ of the random variable $\mathbf{z}_{1:N} = [\mathbf{z}(1), \dots, \mathbf{z}(N)]^T$ as

$$Q_{\mathbf{z},N} = E[\mathbf{z}_{1:N}\mathbf{z}_{1:N}^T], \quad (37)$$

i.e., the ij th $p \times p$ block matrix element $(Q_{\mathbf{z},N})_{ij}$ of $Q_{\mathbf{z},N}$ is $E[\mathbf{z}(i)\mathbf{z}(j)^T]$. Expanding (37) with

$$\mathbf{z}_{1:N} = \begin{bmatrix} C_e \\ C_e A_e \\ \vdots \\ C_e A_e^N \end{bmatrix} \tilde{\mathbf{x}}_e(0) + \begin{bmatrix} I & 0 & 0 & 0 \\ C_e A_e^0 K_e & I & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots \\ C_e A_e^{N-1} K_e & \dots & C_e A_e^0 K_e & I \end{bmatrix} \begin{bmatrix} \mathbf{e}(0) \\ \vdots \\ \mathbf{e}(N) \end{bmatrix}, \quad (38)$$

yields

$$(Q_{\mathbf{z},N})_{ij} = \begin{cases} C_e P_{i-1} C_e^T + Q_e & \text{if } i = j \\ (C_e P_{i-1} A_e^T + Q_e K_e^T)(A_e^{j-i-1})^T C_e^T & \text{if } i < j \\ C_e (A_e^{i-j-1})(A_e P_{j-1} C_e^T + K_e Q_e) & \text{if } i > j \end{cases} \quad (39)$$

with

$$P_{i+1} = A_e P_i A_e^T + K_e Q_e K_e^T. \quad (40)$$

The state covariances P_i , are computed under the assumption that the stochastic process $\mathbf{y}(t)$ has non-zero state and covariance at $t = 0$, encoding the infinite past.

IV. ILLUSTRATIVE EXAMPLE

Consider a stochastic process $\mathbf{y}(t)$, defined by the LTI system

$$\begin{aligned} \mathbf{x}(t+1) &= \begin{bmatrix} 0.5 & 1 \\ 0 & 0.5 \end{bmatrix} \mathbf{x}(t) + \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \mathbf{v}(t) \\ \mathbf{y}(t) &= \begin{bmatrix} 1 & 0 \end{bmatrix} \mathbf{x}(t) + \begin{bmatrix} 1 & 0 \end{bmatrix} \mathbf{v}(t) \\ \mathbf{v}(t) &\sim \mathcal{N}(0, 1) \end{aligned} \quad (41)$$

Given a realisation of $\mathbf{y}(t)$ of $N = 100$ samples, we wish to estimate a parameter θ of the system

$$\begin{aligned} \hat{\mathbf{x}}(t+1) &= \begin{bmatrix} 0.5 & 1 \\ 0 & 0.5 + \theta \end{bmatrix} \hat{\mathbf{x}}(t) + \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \mathbf{v}(t) \\ \hat{\mathbf{y}}(t) &= \begin{bmatrix} 1 & 0 \end{bmatrix} \hat{\mathbf{x}}(t) + \begin{bmatrix} 1 & 0 \end{bmatrix} \mathbf{v}(t) \\ \mathbf{v}(t) &\sim \mathcal{N}(0, 1) \end{aligned} \quad (42)$$

Remark For illustration purposes, θ represents the deviation from System (41), such that $\theta = 0$ would yield the smallest generalisation error.

A forward innovation form of a known LTI system in the form (16) can be found by solving algebraic Riccati equations, described in [13, Section 6.9]. Yielding a predictor

$$\begin{aligned} \mathbf{x}(t+1) &= \underbrace{\begin{bmatrix} 0.5 & 1 \\ 0 & 0.5 \end{bmatrix}}_{A_0} \mathbf{x}(t) + \underbrace{\begin{bmatrix} 1 \\ 0.12 \end{bmatrix}}_{K_0} \mathbf{e}(t) \\ \mathbf{y}(t) &= \underbrace{\begin{bmatrix} 1 & 0 \end{bmatrix}}_{C_0} \mathbf{x}(t) + \begin{bmatrix} 1 \end{bmatrix} \mathbf{e}(t) \end{aligned} \quad (43)$$

From System description (43) in forward innovation form, and parameterisation described in the example description (42), the hypothesis set can be defined as

$$\begin{aligned} \mathcal{F} &= \left\{ f_{\theta} = (A(\theta), K_0, C_0) \mid A(\theta) = A_0 + \begin{bmatrix} 0 & 0 \\ 0 & \theta \end{bmatrix}, \right. \\ &\quad \left. \theta \in \mathbb{R}, \rho(A(\theta)) < 1, \rho(A(\theta) - K_0 C_0) < 1 \right\} \quad (44) \end{aligned}$$

$$\det([K_0 \ A(\theta)K_0]) \neq 0, \det \left(\begin{bmatrix} C_0 \\ C_0 A(\theta) \end{bmatrix} \right) \neq 0 \}$$

where $\rho(\cdot)$ denotes the spectral radius.

Equipped with the hypothesis set (44), a realisation of stochastic process $S = \{\mathbf{y}(t)(\omega)\}_{t=0}^{100}$ of $N = 100$ samples of System (41), and a prior uniform distribution over the




Fig. 1. Empirical Loss over parameter set θ

hypothesis set as $\pi \sim \mathcal{U}(-1.5, 0.5)$. The posterior distribution $\hat{\rho}(\theta)$ and the upperbound on the expectation of generalised loss $E_{f \sim \hat{\rho}} \mathcal{L}_y^\ell(f)$, can be computed via sequential Monte-Carlo.

First sample a set of parameters

$$\Theta = \{\theta_i\}_{i=1}^{N_f}, \quad \theta_i \sim \pi \quad (45)$$

and construct a predictor $f_{\theta_i}, \forall \theta_i \in \Theta$ with

$$\Sigma(\theta_i) = (\hat{A}, \hat{K}, \hat{C}) = (A(\theta_i), K_0, C_0)$$

For each f_{θ_i} compute empirical predictions by iterating

$$\begin{aligned} \hat{x}(t+1) &= (\hat{A} - \hat{K}\hat{C})\hat{x}(t) + \hat{K}\mathbf{y}(t)(\omega), \\ \hat{y}(t) &= \hat{C}\hat{x}(t), \\ \hat{x}(0) &= 0. \end{aligned} \quad (46)$$

From empirical predictions compute empirical loss

$$\hat{\mathcal{L}}_{\mathbf{y}, N}^\ell(f_{\theta_i}) = \frac{1}{N} \sum_{t=1}^N (\mathbf{y}(t)(\omega) - \hat{\mathbf{y}}(t))^T (\mathbf{y}(t)(\omega) - \hat{\mathbf{y}}(t)). \forall i \quad (47)$$

Figure 1 shows empirical loss over the parameter set Θ . Note that $\theta = 0$ does not yield the smallest empirical loss, since we are using a finite number of samples. From (47) and prior distribution π , the Gibbs posterior distribution $\hat{\rho}$ (13) can be approximated, and is shown in Figure 2. In this example, the true system is known and therefore λ could be chosen to maximise the likelihood of $\theta = 0$. However, in an application, one would need to consider λ carefully.

Now, the computed posterior distribution can be used to calculate the expectation over empirical losses. Normally, one would have to repeat the sequential Monte-Carlo process, while sampling $\theta \sim \hat{\rho}$, and take the sample-mean as an approximation of the expectation. In this simplified example, the expectation can be approximated by approximating an integral

$$E_{f \sim \hat{\rho}} \mathcal{L}_{\mathbf{y}, N}^\ell(f_\theta) \approx \sum_{i=1}^{N_f-1} \hat{\mathcal{L}}_{\mathbf{y}, N}^\ell(f_{\theta_i}) \hat{\rho}(\theta_i) (\theta_{i+1} - \theta_i) \quad (48)$$

The KL divergence can similarly be approximated as

$$KL(\hat{\rho}||\pi) \approx \sum_{i=1}^{N_f-1} \ln \left(\frac{\hat{\rho}(\theta_i)}{\pi(\theta_i)} \right) \hat{\rho}(\theta_i) (\theta_{i+1} - \theta_i) \quad (49)$$




Fig. 2. Various Gibbs distributions with varying λ

Lastly $u_{\ell, \pi}(\lambda, N)$ can be computed according to Section III-B. Now we have everything necessary to compute the PAC-bayesian upper bound, the right hand side of (11). For $\lambda = 100$, the upper-bound ends up being 4.48, while $E_{f \sim \hat{\rho}} \mathcal{L}_y^\ell \approx 2.03$. Figure 3 showcases the upper bound for various number of training samples N and real values $\lambda > 0$. Recall, that in Gibbs distribution (13), λ plays a role of weighting the importance of data versus the importance of the prior distribution, larger λ will minimise $E_{f \sim \rho} \hat{\mathcal{L}}_y^\ell$, however it will increase $KL(\rho||\pi)$. Therefore, it is natural that there exists λ that compromises empirical data fit, with prior assumptions. For this specific example, realisation and prior, $\lambda = 5$ yields the smallest upper bound for all N .




Fig. 3. (a) shows the PAC-bayesian upper bound (right hand side of (11)) (b) shows the expectation of the generalisation error over posterior distribution, which depends on (N, λ) , the minimal value over all N and λ is denoted by *

V. CONCLUSIONS AND FUTURE WORK

In this paper we have formulated a non-trivial error PAC-Bayesian error bound for autonomous stochastic LTI models. This error bounds gives a non-asymptotic error bound on the generalization error. Future work will be directed towards extending this bound first to stochastic LTI models with inputs, and then to more complex models such as hybrid models and recurrent neural network. Another research direction would be to understand the effect of the choice of prior and the constant λ on the resulting learning algorithms and comparing it to various existing learning algorithms.

REFERENCES

- [1] P. Alquier, X Li, and O. Wintenberger. Prediction of time series by statistical learning: general losses and fast rates. *Dependence Modeling*, 1(2013):65–93, 2013.
- [2] P. Alquier, J. Ridgway, and N. Chopin. On the properties of variational approximations of Gibbs posteriors. *JMLR*, 17(239):1–41, 2016.
- [3] A. Ambroladze, E. Parrado-Hernández, and J. Shawe-Taylor. Tighter PAC-Bayes bounds. In *NIPS*, 2006.
- [4] A. Barbu and S. C. Zhu. *Monte Carlo Methods*. Springer Singapore, Singapore, 1st edition, 2020.
- [5] P. E. Caines. *Linear Stochastic Systems*. John Wiley and Sons, 1988.
- [6] G. K. Dziugaite and D. M. Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. In *UAI*. AUAI Press, 2017.
- [7] P. Germain, F. Bach, A. Lacoste, and S. Lacoste-Julien. Pac-bayesian theory meets bayesian inference. In *NIPS*, pages 1876–1884, 2016.
- [8] P. Grünwald. The safe Bayesian - learning the learning rate via the mixability gap. In *ALT*, 2012.
- [9] B. Guedj. A Primer on PAC-Bayesian Learning. *arXiv preprint arXiv:1901.05353*, 2019.
- [10] E. J. Hannan and M. Deistler. *The statistical theory of linear systems*. SIAM, 2012.
- [11] T. Katayama. *Subspace Methods for System Identification*. Springer-Verlag, 2005.
- [12] A. Klenke. *Probability Theory: A Comprehensive Course*. Springer International Publishing AG, Cham, 2020.
- [13] A. Lindquist and G. Picci. *Linear Stochastic Systems: A Geometric Approach to Modeling, Estimation and Identification*. Springer, 2015.
- [14] L. Ljung. *System Identification: Theory for the user (2nd Ed.)*. PTR Prentice Hall., Upper Saddle River, USA, 1999.
- [15] D. McAllester. Some PAC-Bayesian theorems. *Machine Learning*, 37(3):355–363, 1999.
- [16] D. McAllester. Simplified PAC-Bayesian margin bounds. In *COLT*, pages 203–215, 2003.
- [17] V. Shalaeva, A. Fakhrizadeh Esfahani, P. Germain, and M. Petreczky. Improved PAC-bayesian bounds for linear regression. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:5660–5667, Apr. 2020.
- [18] S. Shalev-Shwartz and S. Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- [19] R. Sheth and R. Khardon. Excess risk bounds for the bayes risk using variational inference in latent gaussian models. In *NIPS*, pages 5151–5161, 2017.
- [20] T. Zhang. Information-theoretic upper and lower bounds for statistical estimation. *IEEE Trans. Information Theory*, 52(4):1307–1321, 2006.

APPENDIX

Proof Lemma II.2. First we show that $\lim_{t \rightarrow \infty} (\hat{\mathbf{y}}_f(t | 0) - \hat{\mathbf{y}}_f(t)) = 0$ in the mean square sense. Indeed, for any $\tau \in \mathbb{Z}$ $\lim_{s \rightarrow -\infty} (\hat{\mathbf{y}}_f(\tau | s) - \hat{\mathbf{y}}_f(\tau)) = 0$ in the mean square sense. Notice that $(\hat{\mathbf{y}}_f(0 | s) - \hat{\mathbf{y}}_f(0)) = \sum_{k=1}^{\infty} M_{k-s} \mathbf{y}(s-k)$, where the infinite sum is absolutely convergent in the mean square sense. Define the sequence $b_N = \mathbf{E}[\|\sum_{k=1}^{\infty} M_{k+N} \mathbf{y}(-N-k)\|_2^2]$. Since $\mathbf{E}[\|\hat{\mathbf{y}}_f(0 | s) - \hat{\mathbf{y}}_f(0)\|_2^2] = b_{-s}$, it follows that the sequence b_N converges to zero as N converges to ∞ . Since

\mathbf{y} is stationary, it follows that for any $N \geq 0$, any $\tau \in \mathbb{Z}$ $b_N = \mathbf{E}[\|\sum_{k=1}^{\infty} M_{k+N} \mathbf{y}(-N-k)\|_2^2] = \mathbf{E}[\|\sum_{k=1}^{\infty} M_{k+N} \mathbf{y}(\tau-k-N)\|_2^2]$. In particular, by choosing $\tau = N$, it follows that $b_N = \mathbf{E}[\|\sum_{k=1}^{\infty} M_{k+N} \mathbf{y}(-k)\|_2^2]$. Finally, notice that $\hat{\mathbf{y}}_f(t | 0) - \hat{\mathbf{y}}_f(t) = \sum_{k=1}^N M_{k+t} \mathbf{y}(-k)$, and hence $\mathbf{E}[\|\hat{\mathbf{y}}_f(t | 0) - \hat{\mathbf{y}}_f(t)\|_2^2] = b_t$. Since $\lim_{t \rightarrow \infty} b_t = 0$, it follows that $\lim_{t \rightarrow \infty} (\hat{\mathbf{y}}_f(t | 0) - \hat{\mathbf{y}}_f(t)) = 0$ in the mean square sense, as claimed. Since $\lim_{t \rightarrow \infty} (\hat{\mathbf{y}}_f(t | 0) - \hat{\mathbf{y}}_f(t)) = 0$ in the mean square sense, $\lim_{t \rightarrow \infty} a_t = 0$, where $a_t = \mathbf{E}[\|\ell(\hat{\mathbf{y}}_f(t | 0), \mathbf{y}(t)) - \ell(\hat{\mathbf{y}}_f(t), \mathbf{y}(t))\|]$. It is easy to see that $\hat{\mathcal{L}}_{\mathbf{y}, N}(f) - V_{\mathbf{y}, N}(f) = \frac{1}{N} \sum_{t=1}^N (\ell(\hat{\mathbf{y}}_f(t | 0), \mathbf{y}(t)) - \ell(\hat{\mathbf{y}}_f(t), \mathbf{y}(t)))$, and hence $\mathbf{E}[\|\hat{\mathcal{L}}_{\mathbf{y}, N}(f) - V_{\mathbf{y}, N}(f)\|] \leq \frac{1}{N} \sum_{t=1}^N a_t$. Since a_t converges to 0 as $t \rightarrow \infty$, by the well-known property of Cesaro-means $\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N a_t = 0$. \square

Proof of Lemma III.1. Note that if $\Sigma = (\hat{A}, \hat{K}, \hat{C})$ is a minimal tuple, then $(\hat{A} - \hat{K}\hat{C}, \hat{K}, \hat{C})$ is a minimal deterministic LTI, since controllability of (\hat{A}, \hat{K}) implies that $(\hat{A} - \hat{K}\hat{C}, \hat{K})$ is controllable. Similarly, if Σ' is a minimal tuple, then $(\hat{A}' - \hat{K}'\hat{C}', \hat{K}', \hat{C}')$ is a minimal dimensional deterministic LTI system. Notice that f_{Σ} and $f_{\Sigma'}$ are the input-output maps of the deterministic LTI systems $S_1 = (\hat{A} - \hat{K}\hat{C}, \hat{K}, \hat{C})$ and $S_2 = (\hat{A}' - \hat{K}'\hat{C}', \hat{K}', \hat{C}')$ respectively. Hence, from standard realization theory it follows that as both S_1 and S_2 are minimal, $f_{\Sigma} = f_{\Sigma'}$ implies that there exists a nonsingular matrix T such that $\hat{A}' - \hat{K}'\hat{C}' = T(\hat{A} - \hat{K}\hat{C})T^{-1}$, $\hat{K}' = T\hat{K}$, and $\hat{C}' = \hat{C}T^{-1}$. From this the second statement of the lemma follows.

In order to prove the first part, consider the the deterministic LTI system $S_1 = (\hat{A} - \hat{K}\hat{C}, \hat{K}, \hat{C})$ and apply to it the standard minimization algorithm. Let the matrices of the resulting system be $S_2 = (\hat{F}_m, \hat{K}_m, \hat{C}_m)$ and define $\hat{A}_m = \hat{F}_m + \hat{K}_m\hat{C}_m$. Notice that \hat{F}_m is stable, since $\hat{A} - \hat{K}\hat{C}$ was stable and minimization preserves stability. It is not difficult to see that minimization preserves stability of \hat{A} too, i.e., \hat{A}_m will be stable too. Hence, the tuple $\Sigma_m = (\hat{A}_m, \hat{K}_m, \hat{C}_m)$ is a minimal tuple. Hence, since the input-output maps of S_1 and S_2 coincide, and the input-output map of the former is f_{Σ} and of the later is f_{Σ_m} , and hence $f_{\Sigma} = f_{\Sigma_m}$. \square

Proof of Lemma III.2. Indeed, assume that $\Sigma(\theta_0) = (A_0, K_0, C_0)$ is a minimal tuple such that $(A_0, K_0, C_0, I, \mathbf{e})$ is a minimal realization of \mathbf{y} . It is not difficult to see that in this case, $\hat{\mathbf{y}}_{f_{\Sigma_0}}(t) = E_l[\mathbf{y}(t) | \mathcal{H}_{t-}^{\mathbf{y}}]$, i.e., $\hat{\mathbf{y}}_{f_{\Sigma_0}}(t)$ is the orthogonal projection of $\mathbf{y}(t)$ onto its own past. Hence, $\mathbf{e}(t) = \hat{\mathbf{y}}_{f_{\Sigma_0}}(t) - \mathbf{y}(t)$, $\ell(\hat{\mathbf{y}}_{f_{\Sigma_0}}(t), \mathbf{y}(t)) = \|\hat{\mathbf{y}}_{f_{\Sigma_0}}(t) - \mathbf{y}(t)\|_2^2$, and therefore

$$E[\ell(\hat{\mathbf{y}}_{f_{\Sigma_0}}(t), \mathbf{y}(t))] = E[\|\mathbf{e}(t)\|_2^2] \quad (50)$$

It is easy to see that for any other element of $f \in \mathcal{F}$, the entries of $\hat{\mathbf{y}}_f(t | s)$ and $\hat{\mathbf{y}}_f(t)$ belong to the closed subspace $\mathcal{H}_{t-}^{\mathbf{y}}$ generated by the past of \mathbf{y} . By the properties of orthogonal projections, it then follows that

$$\begin{aligned} \mathcal{L}^{\ell}(f) &= E[\ell(\hat{\mathbf{y}}_f(t), \mathbf{y}(t))] = \\ &= E[\|\hat{\mathbf{y}}_f(t) - \mathbf{y}(t)\|_2^2] \leq E[\|\mathbf{e}(t)\|_2^2] = \\ &= E[\ell(\hat{\mathbf{y}}_{f_{\Sigma(\theta_0)}}(t), \mathbf{y}(t))] = \mathcal{L}^{\ell}(f_{\Sigma(\theta_0)}) \end{aligned} \quad (51)$$

and the equality holds only if $\hat{\mathbf{y}}_f(t) = \hat{\mathbf{y}}_{f_{\Sigma(\theta_0)}}(t)$. The latter means that $\mathbf{y}(t) - \hat{\mathbf{y}}_f(t) = \mathbf{e}(t)$. This means that if $f = f_{\Sigma}$ for a minimal tuple $\Sigma = (\hat{A}, \hat{K}, \hat{C})$, then by using (21) and substituting $\mathbf{y}(t) - \hat{\mathbf{C}}\mathbf{x}_{\Sigma}(t) = \mathbf{y}(t) - \hat{\mathbf{y}}_f(t) = \mathbf{e}(t)$ into the first equation of (21), it follows that

$$\mathbf{x}_{\Sigma}(t+1) = \hat{A}\mathbf{x}_{\Sigma}(t) + \hat{K}\mathbf{e}(t)$$

$$\mathbf{y}(t) = \hat{C}\mathbf{x}_{\Sigma}(t) + (\mathbf{y}(t) - \hat{C}\mathbf{x}_{\Sigma}(t)) = \hat{C}\mathbf{x}_{\Sigma}(t) + \mathbf{e}(t).$$

That is, $(\hat{A}, \hat{K}, \hat{C}, I, \mathbf{e})$ is a realization of \mathbf{y} in forward innovation form. Since (\hat{A}, \hat{K}) is controllable and (\hat{C}, \hat{A}) is observable, $(\hat{A}, \hat{K}, \hat{C}, I, \mathbf{e})$ is a minimal realization of \mathbf{y} in forward innovation form. Hence, it is isomorphic to $(A_0, K_0, C_0, I, \mathbf{e})$, that is, there exists a nonsingular matrix T , such that $\hat{A} = TA_0T^{-1}$, $\hat{K} = TK_0$ and $\hat{C} = C_0T^{-1}$, i.e., Σ and $\Sigma(\theta_0)$ are isomorphic. Since $\Sigma = \Sigma(\theta)$ for some $\theta \in \Theta$, by Assumption III.2, $\theta = \theta_0$. \square

Proof of Lemma III.3. Under Assumption III.1, by [14, Theorem 2.3] the alternative empirical loss converges to the generalization error, i.e. $\lim_{N \rightarrow \infty} V_{\mathbf{y}, N}(f) = \mathcal{L}_{\mathbf{y}}(f)$ with probability 1. Recall that by Lemma II.2 $\lim_{N \rightarrow \infty} V_{\mathbf{y}, N}(f) - \hat{\mathcal{L}}_{\mathbf{y}, N}(f) = 0$ in the mean sense, hence there exists a subsequence N_k , $\lim_{k \rightarrow \infty} N_k = \infty$ such that $\lim_{k \rightarrow \infty} V_{\mathbf{y}, N_k}(f) - \hat{\mathcal{L}}_{\mathbf{y}, N_k}(f) = 0$ with probability 1. Hence, $\lim_{k \rightarrow \infty} \hat{\mathcal{L}}_{\mathbf{y}, N_k}(f) = \mathcal{L}_{\mathbf{y}}(f)$ with probability 1. Using dominated convergence theorem the statement of the lemma follows. \square

Proof of Lemma III.4. Consider the matrix $Q_{\mathbf{z}, N} = \mathbf{E}[\mathbf{z}_{1:N}(f_{\theta})\mathbf{z}_{1:N}^T(f_{\theta})]$. It is easy to see that $Q_{\mathbf{z}, N}$ is the left-upper block of the matrix $Q_{\mathbf{z}, N+1}$. In particular, if $v \neq 0$ is an eigenvector of $Q_{\mathbf{z}, N}$ which corresponds to eigenvalue $\rho_N(\theta)$, then $\hat{v} = (v^T, 0^T)$, $\hat{v}^T Q_{\mathbf{z}, N+1} \hat{v} = v^T Q_N v = \rho_N(\theta) v^T v > \rho_{N+1}(\theta) v^T v$, where we used the fact $\rho_{N+1}(\theta) w^T w \leq w^T Q_{N+1} w$, due to the well-known property of smallest eigenvalue of a symmetric semi-definite matrix. Since $v \neq 0$, it follows that $v^T v > 0$, and hence $\rho_N(\theta) > \rho_{N+1}(\theta)$. Since $\rho_N(\theta)$ is decreasing, it follows that $\lim_{N \rightarrow \infty} \rho_N(\theta) = \rho^*(\theta) = \inf_N \rho_N(\theta)$. From [5, Section 2.2, eq. (2.1)] it follows $Q_{\mathbf{z}, N} = A_N Q_{\mathbf{y}, N} A_N^T$, where $(Q_{\mathbf{y}, N})_{i,j} = E[\mathbf{y}(i)\mathbf{y}^T(j)]$

$$\text{and } A_N = \begin{bmatrix} I & 0 & 0 & \cdots & 0 \\ -M_1 & I & 0 \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ -M_{N-1} & M_{N-2} & -M_{N-3} & \cdots & I \end{bmatrix}, \text{ where}$$

$M_i = C(A - KC)^{i-1}K$, $\theta = (A, K, C)$. In particular, for any $v \neq 0$, $v^T Q_{\mathbf{z}, N} v \geq \mu_N \|A_N v\|^2 \geq \mu_N \|A_N^{-1}\|_2 v^T v$, where μ_N is the minimal eigenvalue of $Q_{\mathbf{y}, N}$. Note that A_N is trivially

$$\text{invertable, and } A_N^{-1} = \begin{bmatrix} I & 0 & 0 & \cdots & 0 \\ \tilde{M}_1 & I & 0 \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \tilde{M}_{N-1} & \tilde{M}_{N-2} & \tilde{M}_{N-3} & \cdots & I \end{bmatrix},$$

where $\tilde{M}_i = CA^{i-1}K$. It then follows that $\|A_N^{-1}\|_2 > 1$ and hence $v^T Q_{\mathbf{z}, N} v > \mu_N v^T v$, and therefore $\rho_N(\theta) > \mu_N$. Finally, since \mathbf{y} is coercive, by [11, proof of Theorem 7.3],

it follows that there exists $\mu_* > 0$ such that the spectral density $\Phi_{\mathbf{y}}$ of \mathbf{y} satisfies $\Phi_{\mathbf{y}}(e^{i\theta}) > \mu_* I$ for all $\theta \in [-\pi, \pi]$, and $\mu_N > \mu_* > 0$. Finally, notice that $v(\theta) = \mathbf{E}[\|\mathbf{y}(t) - \hat{\mathbf{y}}_{f_{\theta}}(t)\|_2^2] = \text{trace} \mathbf{E}[(\mathbf{y}(t) - \hat{\mathbf{y}}_{f_{\theta}}(t))(\mathbf{y}(t) - \hat{\mathbf{y}}_{f_{\theta}}(t))^T]$, and $\mathbf{y}(t) - \hat{\mathbf{y}}_{f_{\theta}}(t)$ is the result of applying the linear system $(A - KC, K, -C, I)$ to θ , where $\theta = (A, K, C)$. In particular, if $H(z) = I - C(zI - A - KC)^{-1}K$ is the transfer function of that system, it follows by [5, Theorem 1.2] that $H(z)\Phi_{\mathbf{y}}(z)H^*(z)$ is the spectral density of $\mathbf{y}(t) - \hat{\mathbf{y}}_{f_{\theta}}(t)$, and $\mathbf{E}[(\mathbf{y}(t) - \hat{\mathbf{y}}_{f_{\theta}}(t))(\mathbf{y}(t) - \hat{\mathbf{y}}_{f_{\theta}}(t))^T] = \frac{1}{2\pi} \int_{-\pi}^{\pi} H(e^{i\theta})\Phi_{\mathbf{y}}(e^{i\theta})H^*(e^{i\theta})d\theta$, hence $v(\theta) = \frac{1}{2\pi} \int_{-\pi}^{\pi} \text{trace} H(e^{i\theta})\Phi_{\mathbf{y}}(e^{i\theta})H^*(e^{i\theta})d\theta \geq \mu_* \frac{1}{2\pi} \int_{-\pi}^{\pi} \text{trace} H(e^{i\theta})H^*(e^{i\theta})d\theta \geq p\mu_*$. Here we used the fact that by Parseval's equality $\frac{1}{2\pi} \int_{-\pi}^{\pi} \text{trace} H(e^{i\theta})H^*(e^{i\theta})d\theta = \|I_p\|_F^2 + \sum_{k=0}^{\infty} \|C(A - KC)^k K\|_F^2 > p\mu$, as $H(e^{i\theta}) = I_p + \sum_{k=0}^{\infty} C(A - KC)^k K e^{ik\theta}$ and $\|\cdot\|_F$ denotes the Frobenius norm. \square

Proof of Lemma III.5. For every $\theta \in \Theta$, $(1 + \frac{\lambda\rho_N(\theta)}{N/2})^{\frac{Np}{2}} \geq (1 + \frac{\lambda\mu_*}{N/2})^{\frac{Np}{2}}$, and $(1 + \frac{\lambda\rho_N(\theta)}{N/2})^{\frac{Np}{2}} \leq (1 + \frac{\lambda\rho_0}{N/2})^{\frac{Np}{2}}$ where we used the fact that $\rho_0 = \rho_1(\theta)$ and that by Lemma III.4 $\rho_N(\theta)$ is decreasing. Hence, $\frac{\exp(\lambda v(\theta))}{(1 + \frac{\lambda\rho_N(\theta)}{N/2})^{\frac{Np}{2}}} \leq \frac{\exp(\lambda v(\theta))}{(1 + \frac{\lambda\mu_*}{N/2})^{\frac{Np}{2}}}$, and $\frac{\exp(\lambda v(\theta))}{(1 + \frac{\lambda\rho_N(\theta)}{N/2})^{\frac{Np}{2}}} \geq \frac{\exp(\lambda v(\theta))}{(1 + \frac{\lambda\rho_0}{N/2})^{\frac{Np}{2}}}$. Notice that $\lim_{N \rightarrow \infty} (1 + \frac{\lambda\mu_*}{N/2})^{\frac{Np}{2}} = \lim_{N \rightarrow \infty} (1 + \frac{\lambda p \mu_*}{N/2})^{\frac{Np}{2}} e^{\lambda p \mu_*}$ and $\lim_{N \rightarrow \infty} (1 + \frac{\lambda\rho_0}{N/2})^{\frac{Np}{2}} = \lim_{N \rightarrow \infty} (1 + \frac{\lambda p \rho_0}{N/2})^{\frac{Np}{2}} e^{\lambda p \rho_0}$ and $\frac{\exp(\lambda v(\theta))}{(1 + \frac{\lambda p x}{N/2})^{\frac{Np}{2}}} \leq \exp(\lambda v(\theta))$ for $x \in \{\rho_0, \mu_*, \rho_N(\theta) \mid N \geq 1\}$. The statement of the lemma follows now using dominated convergence theorem. \square

Proof of Lemma III.6. The first equation follows by using the inequality

$$\left(1 + \frac{x}{y}\right)^y > \exp\left(\frac{xy}{x+y}\right) \text{ for } x, y > 0$$

with $x = p\lambda\rho_N(\theta)$ and $y = \frac{pN}{2}$, the second one by noticing that $\frac{xy}{x+y}$ is monotonically increasing in y , hence, $\frac{p\mu_* pN * 0.5}{\lambda p \mu_* + 0.5 p N} > \frac{p\mu_* * 0.5}{\lambda p \mu_* + 0.5}$. \square