Entropy notions for state estimation and model detection with finite-data-rate measurements

Daniel Liberzon and Sayan Mitra

***Abstract*— We study a notion of *estimation entropy* for continuous-time nonlinear systems, formulated in terms of the number of system trajectories that approximate all other trajectories up to an exponentially decaying error. We also consider an alternative definition of estimation entropy which uses approximating functions that are not necessarily trajectories of the system, and show that the two entropy notions are equivalent. We establish an upper bound on the estimation entropy in terms of the sum of the desired convergence rate and an upper bound on the matrix measure of the Jacobian, multiplied by the system dimension. We describe an iterative procedure that uses quantized and sampled state measurements to generate state estimates that converge to the true state at the desired exponential rate. The average bit rate utilized by this procedure matches the derived upper bound on the estimation entropy. We also show that no other algorithm of this type can perform the same estimation task with bit rates lower than the estimation entropy. Finally, we discuss an application of the estimation procedure in determining, from the quantized state measurements, which of two competing models of a dynamical system is the true model. We show that under a mild assumption of *exponential separation* of the candidate models, detection always happens in finite time.**

I. Introduction

Entropy is a fundamental notion in the theory of dynamical systems. Roughly speaking, it describes the rate at which the uncertainty about the system's state grows as time evolves. One can think of this alternatively as the exponential growth rate of the number of system trajectories distinguishable with finite precision, or in terms of the growth rate of the size of reachable sets. Different entropy definitions (notably, topological and measure-theoretic ones) and relationships between them are studied in detail in the book [1] and in many other sources, and continue to be a subject of active research in the dynamical systems community. The concept of entropy of course also plays a central role in thermodynamics and in information theory (see, e.g., in [2]).

In the context of control theory, if entropy describes the rate at which uncertainty is generated by the system (when no measurements are taken), then it should also correspond to the rate at which information about the system should be collected by the controller in order to induce a desired behavior (such as invariance or stabilization). This link has not escaped the control community, and suitable entropy definitions for control systems have been proposed and related to minimal data rates necessary for controlling the system over a communication channel. The first such result was obtained by Nair et al. in [3], where topological feedback entropy for discrete-time systems was defined in terms of cardinality of open covers in the state space. An alternative definition was proposed later by Colonius and Kawan in [4], who instead counted the number of "spanning" open-loop control functions. The paper [5] summarized the two notions and established an equivalence between them. Colonius subsequently extended the formulation of [4] from discrete-time to continuous-time dynamics and from invariance to exponential stabilization in [6]. The survey [7] provides a broader overview of control under data rate constraints.

In this work we are concerned with the problem of estimating the state of a continuous-time system when state measurements are transmitted via a limited-data-rate communication channel, which means that only quantized and sampled measurements of continuous signals are available to the estimator. Observability over finite-data-rate channels and its connection to topological entropy has been studied, most notably by Savkin [8]. Our point of departure in this paper is a synergy of ideas from Savkin [8] and Colonius [6]. As in [8], we focus on state estimation rather than control. However, we follow [6] in that we consider continuous-time dynamics and require that state estimates converge at a prescribed exponential rate. As a result, our definition of *estimation entropy* combines some features of the definitions used in [8] and [6]. We also consider an alternative definition of entropy which uses approximating functions that are not necessarily trajectories of the system. We show that, somewhat surprisingly, the two entropy notions turn out to be equivalent (Theorem 1). We proceed to establish an upper bound of $(M+\alpha)n/\ln 2$ for the estimation entropy of an n-dimensional nonlinear dynamical system whose Jacobian matrix has matrix measure bounded by M, when the desired exponential convergence rate of the estimate is α (Proposition 2). We note that the same estimation entropy notions for globally Lipschitz systems were introduced and studied in our recent paper [9]. When the system's right-hand side is differentiable and not just Lipschitz, our upper bound on the estimation entropy is sharper than the one given in [9] which has the Lipschitz constant in place of M.

We proceed to describe an iterative procedure that uses quantized and sampled state measurements to generate state estimates that converge to the true state at the desired exponential rate. The main idea in the algorithm, which borrows some elements from [10] and earlier work cited therein, is to exponentially increase the resolution of the quantizer while keeping the number of bits sent in each

The authors are with Coordinated Science Laboratory, University of Illinois at Urbana-Champaign. Email: {liberzon,mitras}@illinois.edu. D. Liberzon's research was supported in part by the NSF grants CNS-1217811 and ECCS-1231196. S. Mitra's research was supported in part by the NSF grants CCF-1422798 and CNS-1338726.

round constant. This is achieved by using the quantized state measurement at each round to compute a bounding box for the state of the system for the next round. Then, at the beginning of the next round, this bounding box is partitioned to make a new and more precise quantized measurement of the state. We show that the bounding box is exponentially shrinking in time at a rate α when the average bit rate utilized by this procedure matches the upper bound $(M+\alpha)n/\ln 2$ on the estimation entropy (Theorem 3 and Proposition 4). We also show that no other algorithm of this type can perform the same estimation task with bit rates lower than the estimation entropy (Proposition 5). In other words, the "efficiency gap" of our estimation procedure is at most as large as the gap between the estimation entropy of the dynamical system and the above upper bound on it.

In the last part of the paper, we briefly discuss an application of the estimation procedure in solving model detection problems. Suppose we are given two competing candidate models of a dynamical system and from the quantized state measurements we would like to determine which one is the true model. For example, the different models may arise from different parameter values or they could model "nominal" and "failure" operating modes of the system. This can be viewed as a variant of the standard system identification or model (in)validation problem (see, e.g., [11], [12]) except, unlike in classical results which rely on input/output data, here we use quantized state measurements and do not apply a probing input to the system. We show that under a mild assumption of *exponential separation* of the candidate models' trajectories, a modified version of our estimation procedure can always definitively detect the true model in finite time (Theorem 6).

II. Preliminaries

In this paper we work with the continuous-time system

$$\dot{x} = f(x), \qquad x(0) \in K \tag{1}$$

where $x \in \mathbb{R}^n$ is the state, $f : \mathbb{R}^n \to \mathbb{R}^n$ is a C^1 (continuously differentiable) function, and $K \subset \mathbb{R}^n$ is a known compact set of initial states. Let $\xi : K \times \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ denote the trajectories or solutions of (1), so that $\xi(x,t)$ is the solution from the initial state x evaluated at time t. We assume that these solutions are defined globally in time, i.e., the system (1) is forward complete.[1]

We denote by $|\cdot|$ some chosen norm in $\mathbb{R}^n$. In general definitions and results this norm can be arbitrary, but in specific quantized algorithm implementations we will find it convenient to use the ∞-norm $\|x\|_\infty := \max_{1\leq i\leq n}|x_i|$; in those places, the choice of the ∞-norm will be explicitly declared. For any $x \in \mathbb{R}^n$ and $\delta > 0$, $B(x,\delta) \subseteq \mathbb{R}^n$ is the closed ball of radius δ centered at x, that is, $B(x,\delta) = \{y \in \mathbb{R}^n : |x-y| \leq \delta\}$; for the ∞-norm this is a hypercube.

Let $\|\cdot\|$ be the induced matrix norm on $\mathbb{R}^{n\times n}$ corresponding to a chosen norm $|\cdot|$ on $\mathbb{R}^n$. Then the *matrix measure* $\mu : \mathbb{R}^{n\times n} \to \mathbb{R}$ is defined by $\mu(A) := \lim_{\varepsilon\to 0^+} \dfrac{\|I+\varepsilon A\|-1}{\varepsilon}$ (see, e.g., [13]). One of the basic properties of matrix measures is that for every matrix A we have

$$\mu(A) \leq \|A\| \tag{2}$$

and we note that the left-hand side of (2) may be negative while the right-hand side is always positive. The role that matrix measures will play in our analysis of the nonlinear system (1) is enabled by the following assumption, which we impose throughout the paper, and by the well-known fact stated in Lemma 1 below.

Assumption 1 The matrix measure of the Jacobian matrix of f is bounded: for some $\bar{\mu} \in \mathbb{R}$ we have

$$\mu\big(\partial f/\partial x\big)(x) \leq \bar{\mu} \qquad \forall\, x \in \mathbb{R}^n \tag{3}$$

Lemma 1 *Consider the system* (1) *satisfying Assumption 1. Then for every pair of initial states $x_1, x_2 \in \mathbb{R}^n$ the corresponding solutions of* (1) *satisfy $|\xi(x_1,t)-\xi(x_2,t)| \leq e^{\bar{\mu}t}|x_1-x_2|$ for all $t \geq 0$.*

From the proof of this result (see, e.g., [14], [15]) it can be seen that instead of requiring the bound (3) to hold globally over $\mathbb{R}^n$ it is enough to know that it holds for all points x reachable from K at some time $t \geq 0$, provided that K is a convex set. Moreover, if all solutions (1) starting from K remain in a bounded invariant set then a $\bar{\mu}$ with the indicated property always exists (by continuity of $\partial f/\partial x$).

By default, the base of all logarithms is 2. When we use the natural logarithm, we write $\ln$. For a bounded set $S \subseteq \mathbb{R}^n$ and $\delta > 0$, a *δ-cover* is a finite collection of points[2] $C = \{x_i\}$ such that $\cup_{x_i\in C} B(x_i,\delta) \supseteq S$. For a hyperrectangle $S \subseteq \mathbb{R}^n$ and $\delta > 0$, a *δ-grid* is a special type of δ-cover of S by hypercubes centered at points along axis-parallel planes that are 2δ apart. The boundaries of the δ-hypercubes centered at adjacent δ-grid points overlap. For a given set S, there are many possible ways of constructing specific δ-grids. We can choose any strategy for constructing them without changing the results in this paper. For example, we can construct a special grid on, say, the unit interval. Then, when working with a general interval I (a cross-section of S in any given dimension), we map I to the unit interval, mark the chosen grid on it, and then map it back to I. We denote the δ-grid on S by $grid(S,\delta)$.

III. Estimation entropy

In this section we review the notion of estimation entropy recently introduced in [9]. Let us select a number $\alpha \geq 0$ that defines a desired exponential convergence rate, and let $T > 0$ be a time horizon (which is initially fixed but ultimately approaches ∞). For each $\varepsilon > 0$, we say that a finite set of functions $\hat{X} = \{\hat{x}_1(\cdot),\ldots,\hat{x}_N(\cdot)\}$ from $[0,T]$ to $\mathbb{R}^n$ is *(T,ε,α,K)-approximating* if for every initial state $x \in K$ there exists some function $\hat{x}_i(\cdot) \in \hat{X}$ such that

$$|\xi(x,t)-\hat{x}_i(t)| < \varepsilon e^{-\alpha t} \qquad \forall\, t \in [0,T]. \tag{4}$$

[1] We will later impose a condition on the Jacobian of f guaranteeing that the distance between solutions of (1) grows at most exponentially, and this implies forward completeness.

[2] With a slight abuse of terminology, we take the elements of a cover to be the centers of the balls covering S and not the balls themselves.

Let $s_{\text{est}}(T,\varepsilon,\alpha,K)$ denote the minimal cardinality of such a (T,ε,α,K)-approximating set, and define *estimation entropy* as

$$h_{\text{est}}(\alpha,K) := \lim_{\varepsilon\searrow 0}\overline{\lim_{T\to\infty}}\,\frac{1}{T}\log s_{\text{est}}(T,\varepsilon,\alpha,K).$$

It is easy to see that instead of $\lim_{\varepsilon\searrow 0}$ we could equivalently write $\sup_{\varepsilon>0}$, because $s_{\text{est}}(T,\varepsilon,\alpha,K)$ grows as $\varepsilon\to 0$ for fixed T,α,K. Intuitively, since s_{est} corresponds to the minimal number of functions needed to approximate the state with desired accuracy, h_{est} is the average number of bits needed to identify these approximating functions. The inner $\overline{\lim}$ extracts the base-2 exponential growth rate of s_{est} with time and the outer limit gives the worst case over $\varepsilon>0$.

As a special case, further considered below, we can define the $\hat{x}_i(\cdot)$'s to be trajectories $\xi(x,\cdot)$ of the system from different initial states. Then, s_{est} corresponds to the number of quantization points needed to identify the initial states, and h_{est} gives a measure of the long-term bit rate needed for communicating sensor measurements to the estimator. We pursue this connection in more detail in Section V. We note that the norm in the above definition can be arbitrary.

A. Alternative entropy notion

In the above definition, the functions $\hat{x}_i(\cdot)$ are arbitrary functions of time and not necessarily trajectories of the system (1). If we insist on using system trajectories, then we obtain the following alternative definition: a finite set of points $S=\{x_1,\dots,x_N\}\subset K$ is (T,ε,α,K)*-spanning* if for every initial state $x\in K$ there exists some point $x_i\in S$ such that the corresponding solutions satisfy

$$|\xi(x,t)-\xi(x_i,t)|<\varepsilon e^{-\alpha t}\qquad\forall\, t\in[0,T]. \tag{5}$$

Letting $s^*_{\text{est}}(T,\varepsilon,\alpha,K)$ denote the minimal cardinality of such a (T,ε,α,K)-spanning set, we could define estimation entropy differently as

$$h^*_{\text{est}}(\alpha,K) := \lim_{\varepsilon\searrow 0}\overline{\lim_{T\to\infty}}\,\frac{1}{T}\log s^*_{\text{est}}(T,\varepsilon,\alpha,K).$$

Since every (T,ε,α,K)-spanning set gives rise to a (T,ε,α,K)-approximating set via $\hat{x}_i(t):=\xi(x_i,t)$, and since entropy is determined by the minimal cardinality of such a set, it is clear that $s_{\text{est}}(T,\varepsilon,\alpha,K)\le s^*_{\text{est}}(T,\varepsilon,\alpha,K)$ for all T,ε,α,K, and therefore $h_{\text{est}}(\alpha,K)\le h^*_{\text{est}}(\alpha,K)$ for all α,K. Interestingly, this last inequality is actually always equality. In other words, there is no advantage—as far as estimation entropy is concerned—in using any approximating functions (even possibly discontinuous ones) other than system trajectories.

Theorem 1 *For every $\alpha\ge 0$ and every compact set K we have $h_{est}(\alpha,K)=h^*_{est}(\alpha,K)$.*

This result was proved in [9]. By compactness of K and by the property of continuous dependence of solutions of (1) on initial conditions, for given ε,α,T there exists a $\delta>0$ such that (5) holds whenever x and x_i satisfy $|x-x_i|<\delta$. From this it immediately follows that $s^*_{\text{est}}(T,\varepsilon,\alpha,K)$, and hence also $s_{\text{est}}(T,\varepsilon,\alpha,K)$, is finite for every $\varepsilon>0$. This does not in principle preclude $h^*_{\text{est}}(\alpha,K)$ and $h_{\text{est}}(\alpha,K)$ from being infinite (the supremum over positive ε could still be ∞). However, we will see next that this does not happen if the system satisfies Assumption 1.

IV. Entropy bounds

In this section, we establish an upper bound on the estimation entropy of (1). This entropy bound is independent of the choice of the initial set K; without significant loss of generality, we assume in the sequel that K is a set of positive measure and "regular" shape, such as a hypercube, large enough to contain all initial conditions of interest.

The result given below relies on the global bound $\bar{\mu}$ on the matrix measure of the Jacobian of f provided by Assumption 1. While this assumption is restrictive, we note the following points. First, as we commented after Lemma 1, this can be replaced by a bound over the reachable set, which automatically exists if the reachable set is bounded. Second, we are not assuming that $\bar{\mu}<0$, i.e., the system need not be contractive. Finally, it is useful to compare the entropy bound given here to the one established in [9], which applies to globally Lipschitz (but not necessarily C^1) systems and looks similar but has the Lipschitz constant L of f in place of $\bar{\mu}$. When f is C^1, the bound derived here is sharper because the Lipschitz constant is equal to the induced norm of the Jacobian and so, in light of (2), we have $\bar{\mu}\le L$.

Proposition 2 *For the system* (1) *satisfying Assumption 1, the estimation entropy $h_{est}(\alpha,K)$ is finite and does not exceed $(M+\alpha)n/\ln 2$, where $M:=\max\{\bar{\mu},-\alpha\}$.*

The proof proceeds along the lines of the proof of Proposition 2 in [9] (see also [14] and the references therein for earlier results along similar lines).

Remark 1 In the case when (1) is a linear system

$$\dot{x}=Ax \tag{6}$$

the result of Proposition 2 can be sharpened. Namely, in this case one can show that the exact expression (not just an upper bound) for the estimation entropy is

$$1/(\ln 2)\sum_{\operatorname{Re}\lambda_i(A)>-\alpha}(\operatorname{Re}\lambda_i(A)+\alpha) \tag{7}$$

where $\operatorname{Re}\lambda_i(A)$ are the real parts of the eigenvalues of A. This follows from results that are essentially well known, although not well documented in the literature (especially for continuous-time systems); for discrete-time systems this is shown, e.g., in [8]. Namely, since the flow is $\xi(x,t)=e^{At}x$, the volume of the reachable set at time T from the initial set K is $\det(e^{AT})\mathrm{vol}(K)$ which by Liouville's trace formula equals $e^{(\mathrm{tr}A)T}\mathrm{vol}(K)$. The decaying factor $e^{-\alpha t}$ on the right-hand side of (4) can be canceled by multiplying by $e^{\alpha t}$ on both sides; the effect of doing this on the left-hand side is that of replacing solutions of $\dot{x}=Ax$ by solutions of $\dot{x}=(A+\alpha I)x$, and suitably modifying the approximating functions. Projecting onto the unstable subspace of $A+\alpha I$, we can refine the trace to be the sum of only unstable eigenvalues of this matrix. The number of approximating functions must be at least proportional to the above volume (since the ε-balls around their endpoints must have enough volume to cover the reachable set), and after taking the logarithm, dividing by T, and letting $T \to 0$ we obtain (7) as the lower bound. The upper bound is obtained by reducing A to Jordan normal form followed by an argument similar to the proof of Proposition 2 above applied to each Jordan block (with the corresponding eigenvalue replacing M), and ends up giving the same expression (7).

V. Estimation over infinite horizon

We will first describe a procedure for state estimation of the system (1) over infinite time horizon. Next, we will show that the output of this estimation procedure exponentially converges to the actual state of the system. Finally, we will give a bound on the bit rate sufficient to achieve this convergence.

A. Estimation procedure

From this point on in this section, we will discuss a specific estimation procedure based on quantized state measurements. The norm used here will be the infinity norm $\|\cdot\|_\infty$. Accordingly, the $B(x,\delta)$ balls will be the hypercubes and the grids will be sets of hypercubes. We will treat all previous definitions and results related to entropy in terms of the infinity norm.

The estimation procedure computes a function $v : [0,\infty) \to \mathbb{R}^n$ and an exponentially shrinking envelope around $v(t)$ such that the actual state of the system $\xi(x,t)$ is guaranteed to be within this envelope. It has several inputs: (1) a sampling period $T_p > 0$, (2) a desired exponential convergence rate $\alpha > 0$, (3) an initial set K and an initial partition size $d_0 > 0$, and (4) the constant M defined in Proposition 2, and (5) a subroutine for computing solutions of the differential equation (1). In this paper we do not distinguish between this subroutine for computing solutions and the actual solutions $\xi(\cdot,\cdot)$. The procedure works in rounds $i = 1, 2, \ldots$ and each round lasts T_p time units. In each round, a new state measurement q is obtained and the values of three state variables S, δ, C are updated. We denote these updated values in the i^{th} round as q_i, δ_i, S_i, and C_i. Roughly, $S_i \subseteq \mathbb{R}^n$ is a hypercubic over-approximation of the state estimate, δ_i is the radius of the set S_i, and C_i is a grid on S_i which defines the set of possible state measurements q_{i+1} for the next round. We think of the quantized state measurements q_i as being transmitted from the sensors to the estimator via a finite-data-rate communication channel, while the variables δ_i, S_i, and C_i are generated independently and synchronously on both sides of the channel.

The initial values of these state variables are: $\delta_0 = d_0$; S_0 is a hypercube with center, say x_c, and radius $r_c = \frac{\mathrm{diam}(K)}{2}$, such that $K \subseteq B(x_c, r_c)$; and $C_0 = grid(S_0, \delta_0 e^{-(M+\alpha)T_p})$. Recall the definition of a grid cover from Section II: C_0 is a specific collection of points in $\mathbb{R}^n$ such that $S_0 \subseteq \cup_{x\in C_0} B(x, \delta_0 e^{-(M+\alpha)T_p})$.

At the beginning of the i^{th} round, the algorithm takes as input (from the sensors) a measurement q_i of the current state of the system with respect to the cover C_{i-1} computed in the previous round. The measurement q_i is obtained by choosing a grid point $c \in C_{i-1}$ such that the corresponding $\delta_{i-1}e^{-(M+\alpha)T_p}$-ball $B(c, \delta_{i-1}e^{-(M+\alpha)T_p})$ contains the current state $\xi(x, iT_p)$ of the system. (If there are multiple grid points satisfying this condition—and this may happen as C_{i-1} is a cover with closed sets having overlapping boundaries—then one is chosen arbitrarily.) Using this measurement, the algorithm computes the following: (1) $v_i : [0,T_p] \to \mathbb{R}^n$, which is an approximation function for the state over the interval spanning this round, defined as the solution of the system (1) from q_i, (2) δ_i is updated as $e^{-\alpha T_p}\delta_{i-1}$, (3) $S_i \subseteq \mathbb{R}^n$ is an estimate of the state after T_p time, that is, at the beginning of round $i+1$, and (4) C_i is a $\delta_i e^{-(M+\alpha)T_p}$-grid on S_i. Specifically, S_i is computed by first evaluating the solution $v_i(T_p) = \xi(q_i, T_p)$ of the system starting from q_i after time T_p, and then constructing the hypercube $B(v_i(T_p), \delta_i)$. Note that the size of this hypercube decays geometrically at the rate $e^{-\alpha T_p}$ with each successive round. Recall Section II where we defined grids and discussed specific ways of constructing them; here the specific construction is less important than the fact that each C_i can be computed from q_i by translating and scaling C_{i-1}.

```
input: T_p, α, K, d_0, M, ξ(·,·)
i = 0;
δ_0 ← d_0;
S_0 ← B(x_c, r_c);   // x_c is the center of K
C_0 ← grid(S_0, δ_0 e^{-(M+α)T_p});
while (true)
// at i^th round, i > 0
  i++;
  input q_i ∈ C_{i-1};
  // measurement of current state
  v_i(·) ← ξ(q_i, ·)|[0, T_p];
  δ_i ← e^{-αT_p} δ_{i-1};
  S_i ← B(v_i(T_p), δ_i);
  C_i ← grid(S_i, δ_i e^{-(M+α)T_p});
  output S_i ⊆ R^n, C_i, v_i : [0, T_p] → R^n;
  wait(T_p);
```

Fig. 1. Estimation procedure.

Consider the beginning of the i^{th} round for some $i > 0$. From the algorithm it follows that if the current state x is contained in the estimate S_{i-1} computed in the last iteration, then the measurement q_i is one of the points in the cover C_{i-1} computed in the last iteration, and further, the error in the measurement $|q_i - x|$ is at most the precision of the cover which is $\delta_{i-1}e^{-(M+\alpha)T_p}$.

In order to analyze the accuracy of this estimation procedure, we define a piecewise continuous estimation function $v : [0,\infty) \to \mathbb{R}^n$ by $v(0) := v_1(0)$ and

$$v(t) = v_i(t-(i-1)T_p) \quad \forall\, t \in ((i-1)T_p, iT_p],\ i = 1, 2, \ldots$$

The next theorem, proved along the lines of [9, Theorem 3], establishes an exponentially decaying upper bound on the error between the actual state of the system and the computed approximating function.

Theorem 3 *For any choice of the parameters $\alpha, d_0, T_p > 0$, the procedure in Figure 1 has the following properties: for $i = 0, 1, 2, \dots$ and for any initial state $x \in K$,*

(a) $\xi(x,t) \in S_i$ *for each* $t = iT_p$, *and*
(b) $\|\xi(x,t) - v(t)\|_\infty \leq d_0 e^{-\alpha t} \quad \forall\, t \in [iT_p, (i+1)T_p)$.

B. Bit rate of estimation scheme and its relation to entropy

Now we estimate the communication bit rate needed by the estimation procedure in Figure 1. As the states S_{i-1} and C_{i-1} are maintained and updated by the algorithm in each round, the only information that is communicated from the system to the estimation procedure in each round is the measurement q_i. The number of bits needed for that is $\log(\#C_i)$, where $\#$ stands for the cardinality of a set. The long-term average bit rate of the algorithm is given by $b_r(\alpha, d_0, T_p) := \overline{\lim}_{j\to\infty} \dfrac{1}{jT_p} \sum_{i=1}^{j} \log(\#C_{i-1})$. We proceed to characterize this quantity from the description of the estimation procedure in Figure 1. We calculate $\#C_0 = \lceil \frac{\mathrm{diam}(K)}{2d_0 e^{-(M+\alpha)T_p}} \rceil^n$. For each successive iteration i, $\#C_i = \lceil \frac{\delta_i}{\delta_i e^{-(M+\alpha)T_p}} \rceil^n = \lceil e^{(M+\alpha)T_p} \rceil^n$. Thus, $b_r(\alpha, d_0, T_p) = \lim_{i\to\infty} \frac{1}{T_p} \log(\#C_i) = (M+\alpha)n/\ln 2$ is the bit rate utilized by the procedure; it is actually independent of d_0 and T_p. We state our conclusion as follows.

Proposition 4 *The average bit rate used by the estimation procedure in Figure 1 is $(M+\alpha)n/\ln 2$, where M is defined in Proposition 2.*

By Proposition 2, the bit rate $(M+\alpha)n/\ln 2$ used by the above algorithm is an upper bound on the entropy $h_{\text{est}}(\alpha, K)$. We now establish that no other similar algorithm can perform the same task with a bit rate lower than the entropy $h_{\text{est}}(\alpha, K)$. In other words, the "efficiency gap" of the algorithm is at most as large as the gap between the entropy and its upper bound known from Proposition 2. (Incidentally, combining this result with Proposition 4 we can arrive at an alternative proof of Proposition 2.)

In order to state this result, we need to formalize the class of algorithms to which it applies and to which the above algorithm also belongs. As before, assumed given are the system (1), the associated constant M and initial set K, as well as the desired estimation parameters d_0 (initial bound) and α (convergence rate). We also select the sampling period T_p, which we can think of as a design parameter in the algorithm. On the encoder side, at each step i (corresponding to time $t = (i-1)T_p$), a codeword q_i from a finite set (coding alphabet) C_i is generated based on the state behavior history up to this time. On the decoder side, using this codeword and the previously received codewords, an estimate $v(\cdot)$ of the state over the next sampling interval $((i-1)T_p, iT_p]$ is defined. Such encoding-decoding schemes are by now quite standard (cf. [8, Section 2] and the references therein).

The lower bound on the bit rate in terms of entropy is given below for an algorithm that uses a constant number of bits at each round; since in our estimation algorithm $\#C_0$ may be higher than $\#C_i$ for $i \geq 1$, we can think of this comparison as being valid in "steady state."

Proposition 5 *Consider an algorithm of the above type such that at each step i the set C_i has the same number of elements: $\#C_i = N \ \forall\, i$ (i.e., the coding alphabet is of fixed size). If this algorithm achieves the properties listed in Theorem 3 for an arbitrary $d_0 > 0$, then its bit rate cannot be smaller than $h_{\text{est}}(\alpha, K)$.*

The proof follows along the same lines as the proof of Statement 1 of Theorem III.1 in [8]. We note that the algorithm described in [8] performs a similar estimation task (with $\alpha = 0$ and in discrete time) and operates at an arbitrary bit rate above the entropy. However, that algorithm is quite abstract, since it relies on the existence of a suitable spanning set and performs block coding over a sufficiently large time window using sequences from this spanning set. By contrast, our algorithm given in Section V-A is constructive in that it utilizes a specific quantization procedure and works with an arbitrary fixed sampling period.

Remark 2 For the case of a linear system (6), the algorithm of Section V-A can be modified so that its average bit rate equals the entropy of the linear system given by the formula (7). This can be achieved by aligning the grids C_i used in the algorithm with eigenvectors of the matrix A and replacing the constant M with eigenvalues of A (i.e., using a different number of quantization points for each dimension). Constructions of this type for linear systems are well established in the literature; see, e.g., [16], [17].

VI. Model Detection

In this section we briefly discuss how the estimation algorithm of Figure 1 can be used to distinguish two system models, provided they are in some sense adequately different. Consider two continuous-time system models:

$$\dot{x} = f_1(x), \quad x \in \mathbb{R}^n, \tag{8}$$

$$\dot{x} = f_2(x), \quad x \in \mathbb{R}^n \tag{9}$$

where the initial state is in the known compact set $K \subset \mathbb{R}^n$ and f_1 and f_2 are C^1 functions satisfying Assumption 1, with respective constants M_1 and M_2 defined as in Proposition 2 (see also the comments immediately before that proposition). We denote the trajectories of the systems (8) and (9) by $\xi_1 : \mathbb{R}^n \times \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ and $\xi_2 : \mathbb{R}^n \times \mathbb{R}_{\geq 0} \to \mathbb{R}^n$, respectively. From runtime data, we are interested in distinguishing whether the true dynamics of the system is f_1 or f_2. For example, if f_1 and f_2 correspond to models with different sets of parameter values, then solutions to this problem could be used for model parameter identification. As another example application, consider a scenario where f_1

captures the nominal dynamics of the system and f_2 models a known aberration or failure mode. Then, solution to the above detection problem can be used for failure detection. It is straightforward to generalize the solution proposed below to handle multiple competing models.

For $M_s, T_s > 0$ we say that the two models are (M_s, T_s)-*exponentially separated* (locally) if there exists a constant $\varepsilon_{\min} > 0$ such that for any $\varepsilon \le \varepsilon_{\min}$, for any two states $x_1, x_2 \in \mathbb{R}^n$ with $|x_1 - x_2| \le \varepsilon$, we have $|\xi_1(x_1, T_s) - \xi_2(x_2, T_s)| > \varepsilon e^{M_s T_s}$. The exponential separation property can be shown to hold over a compact set if the vector fields of the two models are different at each point in this set; see also [9] for further discussion and numerical experiments.

In the above definition of exponential separation the norm can be arbitrary, but in the algorithm below we work with the infinity norm. With some modifications, the procedure in Figure 1 can detect models using observations. In Figure 2, we show the procedure for detecting models. First of all, before taking the measurement in each round (T_p time) it makes an additional check. If the current state is not in the set S_i (line 8) computed from the previous round, then the procedure immediately halts by detecting model 2. If the current state is in S_i, then it proceeds as before and records a measurement q_i of the current state as one of the points in the cover C_i. Secondly, the function v_i (line 13) is now computed as a solution $\xi_1(q_i, \cdot)$ of the system given by (8). Finally, in computing the radius of the elements in the cover C_i (line 16), the constant M_1 of the system (8) is used.

1 **input**: $T_p, \alpha, K, d_0, M_1, \xi_1(\cdot,\cdot)$
2 $i = 0$;
3 $\delta_0 \leftarrow d_0$;
4 $S_0 \leftarrow B(x_c, r_c)$;
5 $C_0 \leftarrow grid(S_0, \delta_0 e^{-(M_1+\alpha)T_p})$;
6 **while** (true) // *at* i^{th} *round*, $i > 0$
7 $i{+}{+}$;
8 **if** *current state* $\notin S_{i-1}$
9 **output** ''second model'';
10 **break**;
11 **else**
12 **input** $q_i \in C_{i-1}$;
13 $v_i(\cdot) \leftarrow \xi_1(q_i, \cdot)|[0, T_p]$;
14 $\delta_i \leftarrow e^{-\alpha T_p}\delta_{i-1}$;
15 $S_i \leftarrow B(v_i(T_p), \delta_i)$;
16 $C_i \leftarrow grid(S_i, \delta_i e^{-(M_1+\alpha)T_p})$;
17 **wait**(T_p);

Fig. 2. Procedure for detecting models.

Theorem 6 *Suppose that the true system model is either (8) or (9) and that the two models are (M_1, T_p)-exponentially separated. Then the procedure in Figure 2 outputs "second model" if and only if the true model is (9).*

The proof is very similar to the proof of [9, Theorem 6] with M_1 replacing the Lipschitz constant L_1.

Remark 3 The definition of exponential separation does not imply that the value of the upper bound $\varepsilon_{\min}$ is known, and short of that we cannot conclude for sure that the true model is the first model even if the state measurements conform with the constructed bound S_i in every round. However, if we know such an upper bound $\varepsilon_{\min}$ for which the models are (M_1, T_p)-exponentially separated, then the algorithm can be made to decisively halt with the output "first model" [9].

VII. Conclusions and Future Directions

We introduced two different notions of *estimation entropy* and established their equivalence. We derived an upper bound of $(M + \alpha)n/\ln 2$ for the estimation entropy of an n-dimensional nonlinear dynamical system whose Jacobian's matrix measure does not exceed M, when the desired exponential convergence rate of the estimate is α. We developed a procedure for generating exponentially converging state estimates using an average bit rate that matches this upper bound on the entropy, and showed that no other similar state estimation algorithm can work with bit rates lower than the entropy. Finally, we presented an application of the estimation procedure in picking out one from a pair of candidate models using measurement data.

References

[1] A. Katok and B. Hasselblatt, *Introduction to the Modern Theory of Dynamical Systems*. Cambridge University Press, 1995.

[2] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.

[3] G. N. Nair, R. J. Evans, I. M. Y. Mareels, and W. Moran, "Topological feedback entropy and nonlinear stabilization," *IEEE Trans. Automat. Control*, vol. 49, pp. 1585–1597, 2004.

[4] F. Colonius and C. Kawan, "Invariance entropy for control systems," *SIAM J. Control Optim.*, vol. 48, pp. 1701–1721, 2009.

[5] F. Colonius, C. Kawan, and G. Nair, "A note on topological feedback entropy and invariance entropy," *Systems Control Lett.*, vol. 62, pp. 377–381, 2013.

[6] F. Colonius, "Minimal bit rates and entropy for exponential stabilization," *SIAM J. Control Optim.*, vol. 50, pp. 2988–3010, 2012.

[7] G. N. Nair, F. Fagnani, S. Zampieri, and R. J. Evans, "Feedback control under data rate constraints: An overview," *Proc. IEEE*, vol. 95, pp. 108–137, 2007.

[8] A. V. Savkin, "Analysis and synthesis of networked control systems: Topological entropy, observability, robustness and optimal control," *Automatica*, vol. 42, pp. 51–62, 2006.

[9] D. Liberzon and S. Mitra, "Entropy and minimal data rates for state estimation and model detection," in *Proceedings of Hybrid Systems: Computation and Control (HSCC 2016)*, 2016, pp. 247–256.

[10] D. Liberzon and J. P. Hespanha, "Stabilization of nonlinear systems with limited information feedback," *IEEE Trans. Automat. Control*, vol. 50, pp. 910–915, 2005.

[11] P. A. Ioannou and J. Sun, *Robust Adaptive Control*. New Jersey: Prentice Hall, 1996.

[12] R. S. Smith and J. C. Doyle, "Model validation: a connection between robust control and identification," *IEEE Trans. Automat. Control*, vol. 37, pp. 942–952, 1992.

[13] M. Vidyasagar, *Nonlinear Systems Analysis*, 2nd ed. New Jersey: Prentice Hall, 1993.

[14] V. A. Boichenko and G. A. Leonov, "Lyapunov's direct method in estimates of topological entropy," *J. Math. Sci.*, vol. 91, pp. 3370–3379, 1998.

[15] E. D. Sontag, "Contractive systems with inputs," in *Perspectives in Mathematical System Theory, Control, and Signal Processing*, J. Willems, S. Hara, Y. Ohta, and H. Fujioka, Eds. Springer, 2010, pp. 217–228.

[16] J. P. Hespanha, A. Ortega, and L. Vasudevan, "Towards the control of linear systems with minimum bit-rate," in *Proc. 15th Int. Symp. on Mathematical Theory of Networks and Systems (MTNS)*, 2002.

[17] S. Tatikonda and S. K. Mitter, "Control under communication constraints," *IEEE Trans. Automat. Control*, vol. 49, pp. 1056–1068, 2004.