

Empirical Individual State Observability

Benjamin Cellini, Burak Boyacıoğlu, and Floris van Breugel

Abstract—A dynamical system is observable if there is a one-to-one mapping from the system’s measured outputs and inputs to all of the system’s states. Analytical and empirical tools exist for quantifying the (full state) observability of linear and nonlinear systems; however, empirical tools for evaluating the observability of individual state variables are lacking. Here, a new empirical approach termed Empirical Individual State Observability (E-ISO) is developed to quantify the level of observability of individual state variables. E-ISO first builds an empirical observability matrix via simulation, then determines the subset of its rows required to estimate each state variable individually. We present a convex optimization approach to do this efficiently. Finally, (un)observability measures for these subsets are calculated to provide independent estimates of the observability of each state variable. Multiple example applications of E-ISO on linear and nonlinear systems are shown to be consistent with analytical results. Broadly, E-ISO will be an invaluable tool both for designing active sensing control laws or optimizing sensor placement to increase the observability of individual state variables for engineered systems, and analyzing the trajectory decisions made by organisms.

I. INTRODUCTION

State estimation is a critical component of many tasks involving dynamic systems. A prerequisite for accurate estimation is that a system’s states are observable, i.e. that there is a one-to-one mapping from a system’s measured outputs and inputs to its states. Most estimation methods (e.g. a Kalman filter) rely on all of a system’s states being observable and will otherwise fail in most cases [1]. Mathematical tools have been developed to assess the observability of systems and therefore help us understand how well an estimator will perform [2], [3], [4]. For nonlinear systems, where the observability may depend on the current states and control inputs, these tools can inform the design of control laws to achieve state trajectories and/or sensor locations or sensor configurations that improve the observability of the system.

However, certain control tasks only require estimating a single state variable (or a subset), instead of the full state vector. For instance, given a (nonlinear) system with unknown model parameters, it may be desirable to only estimate these parameters periodically, while ensuring continuous observability of other states. In some high-dimensional systems, such as fluid-structure interactions, it may only be necessary to estimate specific states for purposes of control [5]. Many

navigation tasks also only require partial observability to achieve a desired outcome. A classic example is proportional navigation, a guidance law to ensure a collision course that only requires estimating one state: absolute bearing angle [6]. Flying insects engaged in chemical plume tracking behaviors [7] may also prioritize estimating ambient wind direction over other states such as ground speed. Prior work has shown that the stereotyped zigzagging flight trajectories flying insects use may be tuned to enhance the observability of the high-priority state (ambient wind direction) [8], [9]. In each of these examples, it is critical to have tools for evaluating the observability of individual state variables, instead of the full state vector.

Although analytical tools exist for determining if a particular state variable of a dynamic system is observable [8], [10], or if a linear combination of states is functionally observable [11], there are no established methods to quantify the level of observability. Attempts to accomplish this task include using measures based on eigenvalues of the linear observability Gramian [12], the singular values of the observability matrix [13], [14], filter performances [15], and the estimation ambiguity [16]. These analytical methods become impractical when the dynamic model is missing or sometimes when the dynamics are nonlinear. As an alternative, empirical methods have been developed for quantifying a system’s observability. The advantage of empirical tools, such as the empirical observability Gramian [17], is that they do not rely on an analytical model of the system and only require the ability to simulate it. However, how to tease out the relative observability of each state variable remains unclear. To our knowledge, an empirical method to quantify the observability of individual state variables does not exist.

This work presents a derivative-free empirical framework for evaluating the observability of individual states: Empirical Individual State Observability (E-ISO). Whereas prior approaches primarily employ the empirical observability Gramian, E-ISO utilizes the empirical observability *matrix*. First, analytical observability tools are reviewed. Then, a framework for constructing an empirical observability matrix is presented. From here, E-ISO selects a subset of rows from the observability matrix that are necessary for reconstructing each state individually, and a convex optimization framework to efficiently perform this selection is presented. Lastly, measures of (un)observability are calculated for each subset of rows, yielding an independent measure of (un)observability for each state variable. By considering example linear and nonlinear systems, it is shown that E-ISO can facilitate both sensor selection and trajectory planning to increase the observability of specific state variables.

*This work was partially supported by funding from the Air Force Office of Scientific Research (FA9550-21-0122), the Sloan Foundation (FG-2020-13422), and the National Science Foundation AI Institute in Dynamic Systems (2112085).

The authors are with the Department of Mechanical Engineering at the University of Nevada, Reno, NV 89557 USA. bcellini@unr.edu, bboyacioglu@unr.edu, fvanbreugel@unr.edu

II. NONLINEAR OBSERVABILITY BACKGROUND

To relate our E-ISO method to existing tools, we begin with a brief review of analytical and empirical observability.

A. Analytical Observability: Review

Observability is a fundamental system property that characterizes the existence of an one-to-one (injective) mapping from measurements to state space with the knowledge of inputs. For nonlinear systems, we discuss weak observability, i.e. distinguishability of the unknown initial state in an open neighborhood based on finite-time measurements and input information [18].

Consider the continuous-time/discrete-time nonlinear time-invariant system dynamics,

$$\Sigma_c : \begin{array}{l} \dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t)) \\ \mathbf{y}(t) = \mathbf{h}(\mathbf{x}(t)), \end{array} / \Sigma_d : \begin{array}{l} \mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k) \\ \mathbf{y}_k = \mathbf{h}(\mathbf{x}_k), \end{array} \quad (1)$$

where states \mathbf{x} take values in a smooth n -dimensional state manifold \mathbf{X} , control inputs \mathbf{u} take values in a subset \mathbf{U} of an m -dimensional manifold \mathcal{U} , and outputs \mathbf{y} take values in \mathbb{R}^p . $\mathbf{f}_{\mathbf{u}} := \mathbf{f}(\cdot, \mathbf{u})$ is a smooth vector field for each $\mathbf{u} \in \mathbf{U}$, and $\mathbf{h} = [h_1 \ h_2 \ \dots \ h_p]^\top$ is the smooth output map of the system from \mathbf{X} to \mathbb{R}^p . Given some control \mathbf{u}^* , the first $(w-1)$ time-derivatives of the output for the continuous-time system with $w p \geq n$ are given by

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{y}' \\ \mathbf{y}'' \\ \vdots \\ \mathbf{y}^{(w-1)} \end{bmatrix} = \begin{bmatrix} \mathbf{h}(\mathbf{x}(t)) = L_{\mathbf{f}_{\mathbf{u}^*}}^0 \mathbf{h} \\ \mathbf{h}'(\mathbf{x}(t)) = L_{\mathbf{f}_{\mathbf{u}^*}}^1 \mathbf{h} \\ \mathbf{h}''(\mathbf{x}(t)) = L_{\mathbf{f}_{\mathbf{u}^*}}^2 \mathbf{h} \\ \vdots \\ \mathbf{h}^{(w-1)}(\mathbf{x}(t)) = L_{\mathbf{f}_{\mathbf{u}^*}}^{w-1} \mathbf{h} \end{bmatrix} := \mathcal{G}_c(w, \mathbf{x}(t), \mathbf{u}^*), \quad (2)$$

where the Lie derivative, $L_{\mathbf{f}_{\mathbf{u}^*}} \mathbf{h}$, denotes the derivative of \mathbf{h} with respect to \mathbf{x} on the vector field $\mathbf{f}_{\mathbf{u}^*}$, i.e. $L_{\mathbf{f}_{\mathbf{u}^*}} \mathbf{h} = \frac{\partial \mathbf{h}}{\partial \mathbf{x}} \mathbf{f}_{\mathbf{u}^*}$, and repeated Lie derivatives are calculated as $L_{\mathbf{f}_{\mathbf{u}^*}}^k \mathbf{h} = L_{\mathbf{f}_{\mathbf{u}^*}} L_{\mathbf{f}_{\mathbf{u}^*}}^{k-1} \mathbf{h}$. The invertibility of the mapping \mathcal{G}_c at a given state vector $\mathbf{x}_0 \in \mathbb{R}^n$ requires its Jacobian to have the same rank as the dimension of the state space at \mathbf{x}_0 , i.e. if the observability matrix $\mathcal{O}_{c,w} := d\mathcal{G}_c = \frac{\partial \mathcal{G}_c(w, \mathbf{x}, \mathbf{u}^*)}{\partial \mathbf{x}} \Big|_{\mathbf{x}=\mathbf{x}_0}$ is full column rank, Σ_c is observable [19].

Similarly, given an input sequence $\mathbf{u}^\dagger = (\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{w-1})$, w consecutive measurements from the discrete-time system dynamics would give

$$\begin{bmatrix} \mathbf{y}_k \\ \mathbf{y}_{k+1} \\ \mathbf{y}_{k+2} \\ \vdots \\ \mathbf{y}_{k+w-1} \end{bmatrix} = \begin{bmatrix} \mathbf{h}(\mathbf{x}_k) \\ \mathbf{h}(\mathbf{x}_{k+1}) = \mathbf{h} \circ \mathbf{f}_{\mathbf{u}_0}(\mathbf{x}_k) \\ \mathbf{h}(\mathbf{x}_{k+2}) = \mathbf{h} \circ \mathbf{f}_{\mathbf{u}_1}(\mathbf{x}_k) \circ \mathbf{f}_{\mathbf{u}_0}(\mathbf{x}_k) \\ \vdots \\ \mathbf{h}(\mathbf{x}_{k+w-1}) = \mathbf{h} \circ \mathbf{f}_{\mathbf{u}_{w-1}}(\mathbf{x}_k) \circ \dots \circ \mathbf{f}_{\mathbf{u}_0}(\mathbf{x}_k) \end{bmatrix} \quad (3)$$

$$:= \mathcal{G}_d(w, \mathbf{x}_k, \mathbf{u}^\dagger),$$

where \circ denotes function composition. If the mapping \mathcal{G}_d at \mathbf{x}_0 is invertible, then the discrete-time system is said to be w -step observable at \mathbf{x}_0 , that is, the observability matrix $\mathcal{O}_{d,w} := d\mathcal{G}_d = \frac{\partial \mathcal{G}_d(w, \mathbf{x}, \mathbf{u}^\dagger)}{\partial \mathbf{x}} \Big|_{\mathbf{x}=\mathbf{x}_0}$ being full column rank implies w -step observability of Σ_d [20].

One can also check the observability of a particular state variable by augmenting \mathcal{O}_c or \mathcal{O}_d with the basis vector corresponding to the state of interest $\mathbf{e}_j \in \mathbb{R}^n$ (e.g.

$\mathbf{e}_1 = [1 \ 0 \ 0]^\top$ for the first state of a three-state system) and checking if the rank changes. If $\text{rank}([\mathcal{O}^\top \ \mathbf{e}_j]) = \text{rank}(\mathcal{O}^\top)$, then the information required to obtain the state is already contained within \mathcal{O} , thus the j^{th} state is observable. If the rank does change, then new information about the state was added to \mathcal{O} , thus the state is unobservable [8].

B. Empirical Observability: Review

Although analytical observability tools are valuable for systems with a known model, analytically obtaining the observability matrix is not always possible due to requirements like differentiability. For such systems, the empirical observability Gramian was introduced [17]. Here, we show how to build an empirical observability matrix and relate it to the observability Gramian.

An empirical continuous-time observability matrix can be obtained by numerically computing $\mathcal{O}_{c,w}$. However, calculating the higher-order (time) derivatives that appear in the Jacobian of Eq. 2 using difference formulas can be unreliable [21]. Hence, we focus on building an empirical observability matrix from discrete-time measurements.

Let \mathbf{x}_0 be the initial state vector of interest of the observability analysis, and let \mathbf{u}^\dagger be a nominal input. To construct a w -step empirical observability matrix, we perturb each initial state variable in positive and negative directions with a perturbation amount ε , that is, we simulate the given system dynamics $2n$ times in total and define the perturbed system's output vectors at time k as:

$$\mathbf{y}_k^{\pm j}(\mathbf{x}_0, \mathbf{u}^\dagger, \varepsilon) = \mathbf{y}_k(\mathbf{x}_0 \pm \varepsilon \mathbf{e}_j, \mathbf{u}^\dagger). \quad (4)$$

Then the w -step empirical discrete-time observability matrix can be obtained as:

$$\mathcal{O}_{d,w,\varepsilon} = \frac{1}{2\varepsilon} \begin{bmatrix} \Delta \mathbf{y}_0^1 & \Delta \mathbf{y}_0^2 & \dots & \Delta \mathbf{y}_0^n \\ \Delta \mathbf{y}_1^1 & \Delta \mathbf{y}_1^2 & \dots & \Delta \mathbf{y}_1^n \\ \vdots & \vdots & \ddots & \vdots \\ \Delta \mathbf{y}_{w-1}^1 & \Delta \mathbf{y}_{w-1}^2 & \dots & \Delta \mathbf{y}_{w-1}^n \end{bmatrix}, \quad (5)$$

where $\Delta \mathbf{y}_k^j$'s are the differences between the output vectors at time k for the perturbed state j , $\mathbf{y}_k^{+j}(\cdot)$ and $\mathbf{y}_k^{-j}(\cdot)$, i.e.

$$\Delta \mathbf{y}_k^j(\mathbf{x}_0, \mathbf{u}^\dagger, \varepsilon) = \mathbf{y}_k^{+j}(\mathbf{x}_0, \mathbf{u}^\dagger, \varepsilon) - \mathbf{y}_k^{-j}(\mathbf{x}_0, \mathbf{u}^\dagger, \varepsilon). \quad (6)$$

For the purposes of evaluating observability along a state trajectory, it is convenient to construct the observability matrix in sliding windows assuming no noise (Fig. 1A).

Finally, the continuous-time observability Gramian for the time interval $[0, w\Delta t]$ is defined as [22]:

$$W_{\mathcal{O}_c}(0, w\Delta t) = \int_0^{w\Delta t} \partial_{\mathbf{x}_0} \mathbf{y}^\top(\tau) \partial_{\mathbf{x}_0} \mathbf{y}(\tau) d\tau, \quad (7)$$

and it can be shown that $\mathcal{O}_{d,w,\varepsilon}^\top \mathcal{O}_{d,w,\varepsilon} \Delta t$ with constant $w\Delta t$ would converge to $W_{\mathcal{O}_c}(0, w\Delta t)$ as the perturbation amount ε and the discretization time step size Δt go to zero, that is,

$$W_{\mathcal{O}_c}(0, w\Delta t) \approx \mathcal{O}_{d,w,\varepsilon}^\top \mathcal{O}_{d,w,\varepsilon} \Delta t, \quad (8)$$

for small $\varepsilon, \Delta t$. Hereafter, we simplify $\mathcal{O}_{d,\varepsilon}$ and $W_{\mathcal{O}_c}$ to \mathcal{O}_ε and $W_{\mathcal{O}}$, respectively, and we use a value of $\varepsilon = 10^{-3}$, considering that the standard size of each state variable is of order one.

C. Unobservability Measures: Review

Since observability is determined by the invertibility of $W_{\mathcal{O}}$, established measures for quantifying the level of a nonlinear system's unobservability include the reciprocal of the minimum eigenvalue of $W_{\mathcal{O}}$, $1/\lambda(W_{\mathcal{O}})$, and the condition number of the same matrix, $\kappa(W_{\mathcal{O}}) = \bar{\lambda}(W_{\mathcal{O}})/\lambda(W_{\mathcal{O}})$, and they are called the unobservability index and estimation condition number, respectively [23]. If $\mathcal{O}_{\varepsilon}$ has full column rank, then the singular values of $\mathcal{O}_{\varepsilon}$ approach the square root of the eigenvalues of $\frac{1}{\Delta t}W_{\mathcal{O}}$ for small $\varepsilon, \Delta t$. Since the method presented in this paper analyzes $\mathcal{O}_{\varepsilon}$, we apply these established measures to the singular values of $\mathcal{O}_{\varepsilon}$ and focus on $\kappa(\mathcal{O}_{\varepsilon}^{\top}\mathcal{O}_{\varepsilon}) = \kappa(\mathcal{O}_{\varepsilon})^2 = [\bar{\sigma}(\mathcal{O}_{\varepsilon})/\underline{\sigma}(\mathcal{O}_{\varepsilon})]^2$ for brevity.

III. MOTIVATING EXAMPLES

To illustrate the challenges associated quantifying observability for individual state variables, consider the following series of examples.

$${}_1\mathcal{O}_{\varepsilon} = \begin{bmatrix} 10^{-4} & 1 & 0 \\ 0 & 1 & 0 \end{bmatrix}, \quad {}_2\mathcal{O}_{\varepsilon} = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \end{bmatrix}.$$

In neither case is the system full rank. To assess the observability of the first state we could augment each system with the first state basis vector ($\mathbf{e}_1^{\top} = [1 \ 0 \ 0]$) and check to see if the rank has changed (as in [8]). In both cases the rank does not change, suggesting that the first state is observable, however, it is clearly *more* observable in ${}_2\mathcal{O}_{\varepsilon}$, and a quantification of this difference would be helpful.

Established approaches for quantifying the level of observability involve looking at the eigenvalues of the Gramian. To see the challenges associated with these methods, consider the following.

$${}_3\mathcal{O}_{\varepsilon} = \begin{bmatrix} 1 & 0 \\ 0 & 10 \end{bmatrix}, \quad {}_4\mathcal{O}_{\varepsilon} = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}, \quad {}_5\mathcal{O}_{\varepsilon} = \begin{bmatrix} 1 & 10^{-16} \\ 0 & 10^{-16} \end{bmatrix}.$$

Estimating the first state variable ($j = 1$) is equally well-posed for ${}_3\mathcal{O}_{\varepsilon}$ and ${}_4\mathcal{O}_{\varepsilon}$, and practically, ${}_5\mathcal{O}_{\varepsilon}$ too. However, the condition numbers of the Gramians are all different. Investigating the rows of the Gramians does not provide additional insight either, e.g. ${}_4W_{\mathcal{O}_{\varepsilon}} = \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix}$ hides the fact that the first state is *directly* observable. The second row of ${}_4W_{\mathcal{O}_{\varepsilon}}$ does not actually provide any information about the first state (unless there is a different measurement, or there are dynamics, that make it possible to decouple this combined measurement into its components). To summarize, each row of $\mathcal{O}_{\varepsilon}$ that does not contribute to closely reconstructing the basis vector \mathbf{e}_j can confound efforts to quantify the observability of the j^{th} state. Thus, we develop an approach for finding a small subset of $\mathcal{O}_{\varepsilon}$ that is sufficient for reconstructing \mathbf{e}_j , and which yields a small condition number. There are, however, likely cases where in practice (e.g. with noisy sensors) using more rows of $\mathcal{O}_{\varepsilon}$ (e.g. more sensors, or more measurements in time) will yield a better state estimate for the j^{th} state, despite corresponding to a larger condition number compared the small subset of rows that we select.

IV. EMPIRICAL INDIVIDUAL STATE OBSERVABILITY (E-ISO)

E-ISO provides measures of (un)observability for individual state variables by quantifying how well-posed the problem of estimating a single state variable is given measurements and inputs from a time window of length w . This process involves finding the combination of sensors and measurements, i.e. rows of $\mathcal{O}_{\varepsilon}$, that provide the “best” value for the chosen measure (e.g. the smallest condition number). Solving this problem with a brute force approach would be computationally intractable for large systems, as this is equivalent to finding the best possible combination of rows in $\mathcal{O}_{\varepsilon}$ for observing the state variable of interest. The number of combinations that would need to be evaluated for such an approach is given by

$$N = \sum_{r=1}^{pw} \binom{pw}{r} = \sum_{r=1}^{pw} \frac{(pw)!}{r!(pw-r)!}, \quad (9)$$

where pw is the number of rows in $\mathcal{O}_{\varepsilon}$. For perspective, in a system with four outputs and 25 simulation steps ($pw = 100$) there are $N > 10^{30}$ combinations.

This section details an efficient, but approximate, solution consisting of three steps. First, a sparse subset of rows is selected ($\mathcal{O}_{\varepsilon}^{\mathbf{e}_j}$) whose linear combination can reconstruct \mathbf{e}_j to within a user-specified tolerance (β). Second, (un)observability measures of the corresponding approximate observable subspace ($\hat{\mathcal{O}}_{\varepsilon}^{\mathbf{e}_j}$) are evaluated. Finally, since many unique subsets of rows of $\mathcal{O}_{\varepsilon}$ can be found to reconstruct \mathbf{e}_j , these unique subsets are sequentially gathered, in order of decreasing sparsity, into a collection ${}^i\mathcal{O}_{\varepsilon}^{\mathbf{e}_j}$ that increases in size with each iteration (i). For each iteration, (un)observability measures of the approximate observable subspace are calculated. The final measure that describes the approximate observability of the state variable of interest is the “best” of these. Pseudo-code is provided at the end of the section.

A. Selecting a sparse subset of the observability matrix

For a state to be observable given $\mathcal{O}_{\varepsilon}$, it must be possible to linearly combine the rows to reconstruct the basis state vector corresponding to the state variable of interest (\mathbf{e}_j for the j^{th} state variable). That is, there must be a vector \mathbf{v} such that

$$\mathbf{e}_j^{\top} = \mathbf{v}^{\top}\mathcal{O}_{\varepsilon}. \quad (10)$$

We define $\mathcal{O}_{\varepsilon}^{\mathbf{e}_j}$ as the subset of $\mathcal{O}_{\varepsilon}$ corresponding to non-zero elements in \mathbf{v} . To efficiently exclude as many rows from $\mathcal{O}_{\varepsilon}^{\mathbf{e}_j}$ as possible, we use established optimization tools [24], [25] to find \mathbf{v}_o that minimizes a constrained convex problem,

$$\mathbf{v}_o = \underset{\mathbf{v}}{\arg \min} \quad \|\mathbf{e}_j - \mathcal{O}_{\varepsilon}^{\top}\mathbf{v}\|_2 + \alpha\|\mathbf{v}\|_1, \quad (11)$$

$$\text{s.t.} \quad |(\mathbf{e}_j - \mathcal{O}_{\varepsilon}^{\top}\mathbf{v}^{\top})_s| \leq \beta, \quad s = \{1, 2, \dots, n\}$$

where α is a scalar hyper-parameter, and $\beta > 0$ is a tolerance on how closely each state element of \mathbf{e}_j must be reconstructed. The objective function consists of two terms: 1) the ℓ_2 -norm of the reconstruction error, and 2) the ℓ_1 -norm of the free variable \mathbf{v} . The first term drives \mathbf{v} to reconstruct

the basis vector \mathbf{e}_j , whereas the second term is a regularizer that promotes sparsity in \mathbf{v} . Furthermore, the ℓ_1 -norm penalty on \mathbf{v} also helps to prioritize the selection of rows of \mathcal{O}_ε containing large values, thereby choosing rows that are likely to increase the magnitude of the singular values of \mathcal{O}_ε . In practice, the ℓ_1 -norm does not always drive small elements of \mathbf{v} to zero. Thus, we add an extra step to eliminate small values in \mathbf{v}_0 by sequentially adding the largest elements until the tolerance set on \mathbf{e}_j is met. Figure 1C, gray shading, shows that this optimization selects a single row from \mathcal{O}_ε to estimate \mathbf{e}_1 . Rows highlighted in green and teal are discussed in Sec. IV-C. If no solution to the optimization can be found, barring computational idiosyncrasies of the selected solver, we conclude that the system is *approximately* unobservable.

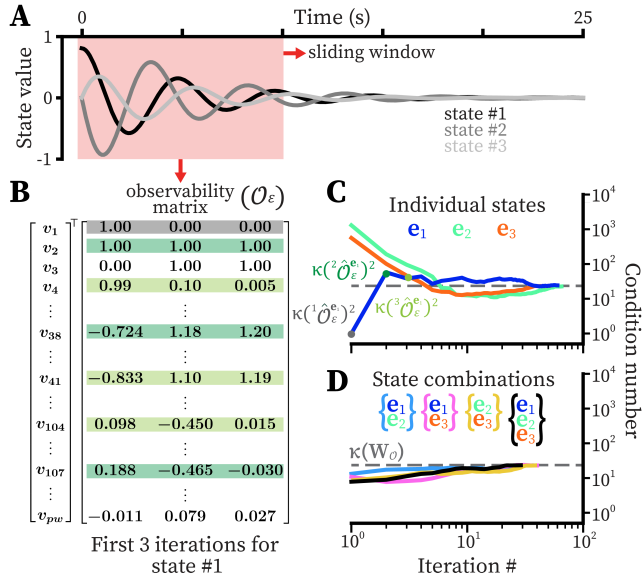


Fig. 1. Graphical illustration of E-ISO applied to a fully observable linear system with the dynamics $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} = \begin{bmatrix} 0 & 1 & 0 \\ -2 & 0 & 1 \\ 1 & 0 & 1 \end{bmatrix}\mathbf{x}$, $\mathbf{y} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \end{bmatrix}\mathbf{x}$. **A.** Simulation from which the empirical observability matrix (\mathcal{O}_ε) is constructed over time in a sliding window ($w = 100$). **B.** Example \mathcal{O}_ε (Eq. 5) and free parameter vector \mathbf{v} (Eq. 10) used for optimization. The rows of \mathcal{O}_ε selected from the first three optimization iterations of E-ISO are highlighted. **C.** The condition number of the observable subspace of each iteration of E-ISO for each individual state variable, which converge to the condition number of the observability Gramian (gray dashed line). **D.** Same as C, but for combinations of state variables. E-ISO parameters: $\alpha = 10^{-2}$, $\beta = 10^{-3}$, $\sigma_0 = 10^{-6}$.

B. Obtaining a quantitative (un)observability measure

To compute quantitative (un)observability measures for individual state variables, the singular values of $\mathcal{O}_\varepsilon^{\mathbf{e}_j}$ can be analyzed. Since this subset may not have full column rank, an approximate observable column-subspace of $\mathcal{O}_\varepsilon^{\mathbf{e}_j}$ is found by calculating a rank-truncated singular value decomposition given a user-specified threshold (σ_0). We define the projection of $\mathcal{O}_\varepsilon^{\mathbf{e}_j}$ onto the approximately observable subspace as $\hat{\mathcal{O}}_\varepsilon^{\mathbf{e}_j}$. Now, established (un)observability measures, such as the condition number, can be applied to $\hat{\mathcal{O}}_\varepsilon^{\mathbf{e}_j}$ to obtain a quantitative (un)observability measure for the state variable of interest (Fig. 1B–C, iteration # = 1).

C. Quantifying observability for iterated subsets of \mathcal{O}_ε

The optimization problem defined by Eq. 11 yields a single set of rows corresponding to an observable state variable, however, there may be multiple valid combinations of rows. To ensure that all relevant rows are accounted for, the optimization should ideally be iterated for every possible combination of rows in \mathcal{O}_ε , and then the subset of selected rows with the minimum condition number could be used to evaluate the observability of the state variable of interest. The first optimization of Eq. 11 yields ${}^1\mathcal{O}_\varepsilon^{\mathbf{e}_j}$ (Fig. 1B, gray row). The rows in ${}^1\mathcal{O}_\varepsilon^{\mathbf{e}_j}$ are then removed from \mathcal{O}_ε and the optimization is repeated to select new rows (Fig. 1B, green rows), which are added to ${}^1\mathcal{O}_\varepsilon^{\mathbf{e}_j}$ and the new collection is defined as ${}^2\mathcal{O}_\varepsilon^{\mathbf{e}_j}$. This process is repeated (e.g. Fig. 1B, teal rows) until the optimization fails to find an observable subset. At each iteration the condition number of ${}^i\hat{\mathcal{O}}_\varepsilon^{\mathbf{e}_j}$ is determined (Fig. 1C). To obtain the best single measure, the minimum condition number is selected:

$$\kappa_{\min} = \min\{\kappa({}^1\hat{\mathcal{O}}_\varepsilon^{\mathbf{e}_j})^2, \kappa({}^2\hat{\mathcal{O}}_\varepsilon^{\mathbf{e}_j})^2, \dots, \kappa({}^w\hat{\mathcal{O}}_\varepsilon^{\mathbf{e}_j})^2\}, \quad (12)$$

where $\kappa({}^i\hat{\mathcal{O}}_\varepsilon^{\mathbf{e}_j})$ is the condition number of the collection of rows selected from \mathcal{O}_ε after i iterations. For directly measurable states, the smallest κ will occur at the first iteration (Fig. 1C, \mathbf{e}_1); for states requiring an accumulation of sensor measurements (Fig. 1C, \mathbf{e}_2 and \mathbf{e}_3) the smallest κ typically occurs at some intermediate iteration number.

The E-ISO method can be extended to determine a single observability measure for any combination of z states by stacking z unique \mathbf{e}_j 's and defining \mathbf{v} as a $pw \times z$ matrix (Fig. 1D). For a fully observable linear system, the condition number $\kappa({}^i\hat{\mathcal{O}}_\varepsilon^{\mathbf{e}_j})^2$ for each individual state, or any combination of states, will converge to the condition number of the observability Gramian $\kappa(W_\mathcal{O})$ after many iterations (Fig. 1C–D), provided that every row of \mathcal{O}_ε is eventually selected. For some systems, our iterative algorithm will not, however, eventually select every row.

Pseudo-code for E-ISO, which we implemented in Python, is provided below.

Algorithm 1 E-ISO

Input:	$\mathcal{O}_\varepsilon, \mathbf{e}_j$	Parameters:	α, β, σ_0	Output:	κ_{\min}
1:	$\hat{\mathcal{O}}_\varepsilon^{\mathbf{e}_j} \leftarrow ()$, $\kappa \leftarrow ()$, $i \leftarrow 1$				▷ initialize variables
2:	$IsObservable \leftarrow \text{True}$				
3:	while $IsObservable$ do				
	▷ reconstruct state with convex optimization				
	▷ r is the collection of row indices used				
4:	$IsObservable, r \leftarrow \text{OPTIMIZE}(\mathcal{O}_\varepsilon, \mathbf{e}_j, \alpha, \beta)$				
5:	if $IsObservable$ then				
6:	$\hat{\mathcal{O}}_\varepsilon^{\mathbf{e}_j} \leftarrow [\hat{\mathcal{O}}_\varepsilon^{\mathbf{e}_j}, \mathcal{O}_\varepsilon(r, :)]$				▷ add rows to subset
7:	$\kappa(i) \leftarrow \text{CONDITIONNUMBER}(\hat{\mathcal{O}}_\varepsilon^{\mathbf{e}_j}, \sigma_0)$				
8:	else				
9:	if $i = 1$ then				▷ failed on 1st iteration
10:	$\kappa_{\min} \leftarrow \infty$				▷ condition # is undefined
11:	else				
12:	$\kappa_{\min} \leftarrow \min(\kappa)$				▷ minimum condition #
13:	$\mathcal{O}_\varepsilon(r, :) \leftarrow 0$				▷ set rows to zero
14:	$i \leftarrow i + 1$				▷ next iteration

V. APPLICATIONS

We apply the E-ISO approach to two examples: a linear system to highlight sensor selection applications, and a nonlinear biological system to highlight applications to trajectory planning for active sensing.

A. Sensor selection

The need for efficiently estimating states of high dimensional systems while minimizing the quantity of physical sensors has spurred the development of sparse sensing approaches [26], [27], but these methods focus on full state estimation. Here, E-ISO is applied to a simple discrete-time linear system with various output configurations to illustrate how sensors could be chosen to maximize the observability of an individual state variable of interest. The analysis shows that the condition number of the Gramian is correlated with the least observable state variable ($R^2 = 0.99$)—thus obscuring information about the more observable state variables—whereas E-ISO can resolve differences in observability between state variables (Fig. 2).




Fig. 2. The condition number of the observability Gramian (top) is generalized to individual state variables with E-ISO (middle). Different sensor sets are represented by C matrices (bottom) applied to a discrete-time linear system with dynamics $\mathbf{x}_{k+1} = A\mathbf{x}_k$, $\mathbf{y}_k = C\mathbf{x}_k$ where $A = \begin{bmatrix} 0.9952 & 0.095 & 0 \\ -0.095 & 0.9902 & 0 \\ 0 & 0 & 0.9948 \end{bmatrix}$. E-ISO parameters: $\alpha = 10^{-4}$, $\beta = 10^{-2}$, $\sigma_0 = 10^{-6}$ multiplied by maximum eigenvalue.

B. State trajectory planning for active sensing

For nonlinear systems, the observability of the state variables can depend on the current state, and some non-zero inputs may be required to guarantee observability [28]. To illustrate E-ISO's application to such *active sensing* objectives, consider the following system inspired by a flying insect such as a fruit fly [8]:

$$\dot{\mathbf{x}} = \begin{bmatrix} d \\ \dot{g} \\ \dot{w} \\ \dot{\phi} \\ \dot{\zeta} \end{bmatrix} = \begin{bmatrix} 0 \\ u_g \\ 0 \\ u_\phi \\ 0 \end{bmatrix}, \quad \mathbf{h}(\mathbf{x}(t)) = \begin{bmatrix} \phi \\ g/d \\ \gamma \end{bmatrix}. \quad (13)$$

States 1–3 represent magnitudes describing the fly's altitude d , ground speed g , and the ambient wind speed w . States 4–5 are angular quantities describing the fly's heading ϕ and the ambient wind direction ζ . For simplicity heading and course direction are defined to be equal. In this example, d , w , and ζ are constant, whereas ϕ and g are directly controlled with inputs u_ϕ and u_g . All dynamics (inertial, aerodynamics, etc.) are excluded to simplify the presentation of the observability analysis. The nonlinear outputs $\mathbf{h}(\mathbf{x}(t))$ of the system consist of ϕ measured directly, optic flow

approximated by g/d [29], and the air speed angle in the global frame:

$$\gamma = \arctan \left(\frac{-g \sin \phi + w \sin \zeta}{-g \cos \phi + w \cos \zeta} \right). \quad (14)$$

For a fly engaged in a chemical plume tracking behavior, estimating ζ is especially important [7]. Thus, ζ can be considered a high-priority estimate, whereas g, d, w are not needed for most plume tracking algorithms. E-ISO can reveal which inputs are needed to ensure observability of an individual state variable, such as ζ . The following E-ISO results confirm prior analytical work and its extension to the exact dynamics given above ([8], see also: <https://github.com/BenCellini/EISO>).

First, the observability of all the state variables was evaluated for a non-zero constant optic flow trajectory $\dot{g} = 0$ with zero inputs $u_g = u_\phi = 0$. Only ϕ , a direct measurement, is observable (Fig. 3A).

Prior work has shown that flies perform rapid turns called saccades at a rate of 0.5 Hz during flight [30], which could potentially improve their ability to estimate the ambient wind direction [8], [9]. In Eq. 13, saccades can be emulated by introducing a nonzero change in heading $\dot{\phi} \neq 0$ by setting $u_\phi \neq 0$. Applying E-ISO to the resulting trajectories reveals that ζ becomes observable, but the rest of the states (aside from ϕ) remain unobservable (Fig. 3B).

Prior work has also shown that with monocular optic flow measurements g/d , both d and g are only observable when a known translational acceleration $\dot{g} \neq 0$ is applied (i.e. $u_g \neq 0$ in Eq. 13) [29], [31]. Applying E-ISO to accelerating trajectories revealed that all the state variables become observable (Fig. 3C). Together, the E-ISO results suggest that flies could use two distinct active sensing strategies, turning and/or accelerating, to estimate ζ —but flies would require acceleration to observe the full state (this result does require any model parameters to be calibrated [8]).

E-ISO can also identify which sensor combinations are necessary to observe an individual state variable given some trajectory. E-ISO shows that the angular sensor set alone is enough to estimate ζ for turning trajectories, but the full sensor set is still required to estimate the full state when accelerating (Fig. 3D).

VI. DISCUSSION

Variations. Although the narrative and examples presented here focus on the discrete-time empirical observability matrix as the starting point, trivial modifications include starting from an analytically determined observability matrix, or a constructability matrix. Furthermore, although we focus on the condition number, alternative measures such as the unobservability index can be used instead. In fact, by working with \mathcal{O}_ε , instead of the Gramian $W_{\mathcal{O}}$, additional measures can be explored including: how many distinct sensors are necessary, how many measurements are needed, and the size of the time window required to capture the level of observability along state trajectories.

Limitations and practical recommendations. We recommend scaling state variable units such that all states are as

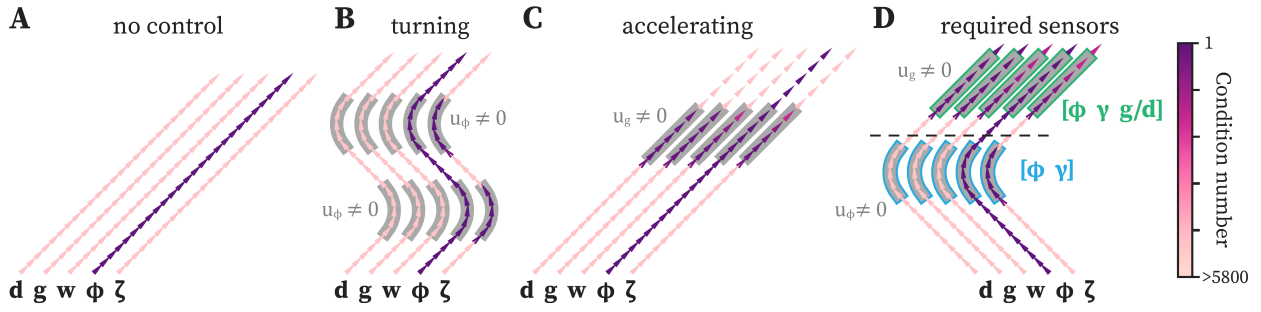


Fig. 3. E-ISO reveals flies must turn or accelerate to estimate wind direction. **A.** Simulated five-second trajectory from Eq. 13 with constant optic flow, no turning, and $\Delta t = 0.1$. For each individual state variable in Eq. 13, color shading indicates the observability level (Eq. 12) for sliding windows ($w = 3$). E-ISO parameters: $\alpha = 10^{-6}$, $\beta = 10^{-3}$, $\sigma_0 = 10^{-8}$. **B.** Same as A, but for a trajectory with turns. **C.** Same as A, for a trajectory with translational acceleration. **D.** Different sensors from Eq. 13 are needed to estimate ζ for different trajectories: only ϕ and γ are necessary when turning, but g/d is also required when accelerating.

comparable in magnitude over time as possible, and scaling outputs according to their expected noise levels. Normalizing the outputs is not recommended, as this would change the interpretation of the singular values of $\hat{\mathcal{O}}_{\varepsilon}^{e_j}$ with respect to the states themselves. E-ISO has three hyper-parameters that may require tuning, or a methodical sweep. When state values are scaled from 0.1 to 10, we recommend starting with: $\alpha = 10^{-6}$, $\beta = 10^{-3}$. For systems with a large number of states and/or measurements, we suggest increasing α and β to find sparse sets of $\mathcal{O}_{\varepsilon}$, at the expense of reconstruction tolerance. When choosing the singular value threshold σ_0 , for which to calculate the rank-truncated condition number, it is generally critical to pick a value that excludes any singular values that are not required to reconstruct the state variable of interest. As a test, a rank-truncated observability matrix can be constructed from a subset of the singular values of the full observability matrix $\mathcal{O}_{\varepsilon}$. This procedure can be iterated starting from the largest singular value until the state variable of interest can be reconstructed within the β tolerance from the rank-truncated observability matrix—and the last value of the smallest singular value required can be set as the σ_0 threshold. Without this routine, an arbitrary chosen σ_0 can lead to inaccurate condition numbers. When possible, for smaller systems, this method of choosing σ_0 combined with evaluating every possible combination of rows will yield the most accurate condition numbers for individual states variables. We provide an example of this implementation on the E-ISO GitHub. E-ISO can become computationally cumbersome for large systems. Rather than implement E-ISO in real-time for trajectory planning, we recommend using E-ISO to identify trajectory motifs a priori (e.g. turning or accelerating), and using these motifs to rank active sensing trajectories.

Applications. E-ISO provides a practical solution to the open problem of methodically discovering trajectories that guarantee observability of specific state variables (or parameters) in partially observable (nonlinear) systems [32]. E-ISO’s focus on individual state variables can serve as a practical generalization to full state methods such as empirical Gramian-based observability methods, and sparse sensor selection algorithms [26]. In particular, E-ISO goes beyond

established methods like Kalman canonical decomposition, which can be used to find the observable subspace, but does not single out individual state variables within that subspace. Beyond sensor selection and trajectory planning, E-ISO can also be used to curate data prior to building bespoke observers using data-hungry machine learning methods to limit extraneous/unobservable information. E-ISO results could also be incorporated into partial update Kalman filters that use observability measures to throttle state estimation [33]. Finally, E-ISO can be used as an elegant analysis and hypothesis generation tool for understanding active sensing in biological systems that may not be concerned with full state estimation.

ACKNOWLEDGMENT

The authors thank Richard Murray, Trevor Avant, Natalie Brace, Aditya Nair, Petros Voulgaris, and Noah Cowan for feedback on the manuscript.

REFERENCES

- [1] J. Li, H. Song, X. Wei, and J. Liang, “Result analysis of Kalman filter for unobservable systems,” in *Proc. of the Int. Conf. on Artificial Intelligence, Information Processing and Cloud Computing*. New York, NY: ACM, 2019, pp. 1–5.
- [2] S. Zeng, “Observability measures for nonlinear systems,” in *2018 IEEE Conf. on Decision and Control*. IEEE, 2018, pp. 4668–4673.
- [3] B. Southallzy, B. Buxtony, and J. Marchant, “Controllability and observability: Tools for Kalman filter design,” in *British Machine Vision Conf.*, vol. 98, 1998, pp. 164–173.
- [4] J. Mei, S. L. Brunton, and J. N. Kutz, “Mobile sensor path planning for Kalman filter spatiotemporal estimation,” 2022, *arXiv:2212.08280*.
- [5] M. K. Hickner, U. Fasel, A. G. Nair, B. W. Brunton, and S. L. Brunton, “Data-driven unsteady aeroelastic modeling for control,” *AIAA Journal*, vol. 61, no. 2, pp. 780–792, 2023.
- [6] S. A. Murtaugh and H. E. Criel, “Fundamentals of proportional navigation,” *IEEE Spectrum*, vol. 3, no. 12, pp. 75–85, 1966.
- [7] F. van Breugel and M. H. Dickinson, “Plume-tracking behavior of flying drosophila emerges from a set of distinct sensory-motor reflexes,” *Current Biology*, vol. 24, no. 3, pp. 274–286, 2014.
- [8] F. van Breugel, “A nonlinear observability analysis of ambient wind estimation with uncalibrated sensors, inspired by insect neural encoding,” in *60th IEEE Conf. on Decision and Control*, 2021, pp. 1399–1406.
- [9] F. van Breugel, R. Jewell, and J. Houle, “Active anemosensing hypothesis: how flying insects could estimate ambient wind direction through sensory integration and active movement,” *Journal of The Royal Society Interface*, vol. 19, no. 193, p. 20220258, 2022.

- [10] M. Anguelova, "Nonlinear observability and identifiability: General theory and a case study of a kinetic model for *S. cerevisiae*," Ph.D. dissertation, Chalmers University of Technology and Göteborg University, 2004.
- [11] T. L. Fernando, H. M. Trinh, and L. Jennings, "Functional observability and the design of minimum order linear functional observers," *IEEE Trans. on Automatic Control*, vol. 55, no. 5, pp. 1268–1273, 2010.
- [12] A. J. Marques, "On the relative observability of a linear system," Master's thesis, Naval Postgraduate School, 1986.
- [13] J. R. Yim, "Autonomous spacecraft orbit navigation," Ph.D. dissertation, Texas A&M University, 2002.
- [14] L. Rui, Q. Yong-yuan, and J. Ji-chao, "Observable degree analysis of SINS initial alignment based on singular value decomposition," in *IEEE Int. Symposium on Knowledge Acquisition and Modeling Workshop*, 2008, pp. 444–448.
- [15] S. Zhigang, M. Xiaochuan, L. Yu, and Y. Shefeng, "Adaptive hybrid kalman filter based on the degree of observability," in *34th Chinese Control Conf. IEEE*, 2015, pp. 4923–4927.
- [16] Q. Gong, W. Kang, C. Walton, I. Kaminer, and H. Park, "Partial observability analysis of an adversarial swarm model," *Journal of Guidance, Control, and Dynamics*, vol. 43, no. 2, pp. 250–261, 2020.
- [17] S. Lall, J. E. Marsden, and S. Glavaški, "Empirical model reduction of controlled nonlinear systems," *IFAC Proc. Volumes*, vol. 32, no. 2, pp. 2598–2603, 1999.
- [18] H. Nijmeijer and A. Van der Schaft, *Nonlinear dynamical control systems*. Springer, 1990.
- [19] E. D. Sontag, "A concept of local observability," *Systems & Control Letters*, vol. 5, no. 1, pp. 41–47, 1984.
- [20] P. Moraal and J. W. Grizzle, "Observer design for nonlinear systems with discrete-time measurements," *IEEE Trans. on Automatic Control*, vol. 40, no. 3, pp. 395–404, 1995.
- [21] F. van Breugel, J. N. Kutz, and B. W. Brunton, "Numerical differentiation of noisy data: A unifying multi-objective optimization framework," *IEEE Access*, vol. 8, pp. 196 865–196 877, 2020.
- [22] D. Georges, "Towards optimal architectures for hazard monitoring based on sensor networks and crowdsensing," *Journal of Integrated Disaster Risk Management*, vol. 10, no. 1, 2020.
- [23] A. J. Krener and K. Ide, "Measures of unobservability," in *Proc. of the 48th IEEE Conf. on Decision and Control held jointly with 28th Chinese Control Conf.*, 2009, pp. 6401–6406.
- [24] S. Diamond and S. Boyd, "CVXPY: A Python-embedded modeling language for convex optimization," *Journal of Machine Learning Research*, vol. 17, no. 83, pp. 1–5, 2016.
- [25] M. ApS, *The MOSEK optimizer API for Python. Version 8.1.*, 2018.
- [26] K. Manohar, B. W. Brunton, J. N. Kutz, and S. L. Brunton, "Data-driven sparse sensor placement for reconstruction: Demonstrating the benefits of exploiting known patterns," *IEEE Control Systems Magazine*, vol. 38, no. 3, pp. 63–86, 2018.
- [27] N. L. Brace, N. B. Andrews, J. Upsal, and K. A. Morgansen, "Sensor placement on a cantilever beam using observability Gramians," in *IEEE 61st Conf. on Decision and Control*, 2022, pp. 388–395.
- [28] A. Kunapreddy and N. J. Cowan, "Recovering observability via active sensing," *Proc. of the American Control Conf.*, pp. 2821–2826, 2018.
- [29] F. van Breugel, K. Morgansen, and M. H. Dickinson, "Monocular distance estimation from optic flow during active landing maneuvers," *Bioinspiration & Biomimetics*, vol. 9, no. 2, p. 025002, 2014.
- [30] B. Cellini and J.-M. Mongeau, "Hybrid visual control in fly flight: Insights into gaze shift via saccades," *Current Opinion in Insect Science*, vol. 42, pp. 23–31, 2020.
- [31] B. Lingenfelter, A. Nag, and F. van Breugel, "Insect inspired vision-based velocity estimation through spatial pooling of optic flow during linear motion," *Bioinspiration & Biomimetics*, 2021.
- [32] H. Mania, M. I. Jordan, and B. Recht, "Active learning for nonlinear system identification with guarantees," *Journal of Machine Learning Research*, vol. 23, pp. 32–1, 2022.
- [33] J. H. Ramos, D. W. Adams, K. M. Brink, and M. Majji, "Observability informed partial-update Schmidt Kalman filter," in *IEEE 24th Int. Conf. on Information Fusion*, 2021, pp. 1–8.