Deceptive Trajectory Imitation Using Affine Feedback

Melkior Ornik[1]

***Abstract*— In adversarial environments it is often beneficial for an agent to reach its objective while seeking to create the impression that it is progressing towards a different goal. This paper considers the setting where an agent seeks to tightly follow a public reference trajectory during the observation period, while actually using an affine feedback controller to ultimately guide the system towards a hidden objective. We pose the optimal synthesis of this affine controller as a nonlinear constrained optimization problem. Taking the sum of norms of trajectory deviations over time as the cost function, by using the power mean inequality we provide an approximation of the optimal controller as a solution of an ordinary least squares problem. We use a method inspired by Tikhonov regularization to ensure that the controlled trajectory converges to the intended objective. We illustrate our method on a variety of numerical examples, showing that the proposed method often generates trajectories that are nearly indistinguishable from the reference during the observation period, and identify some fundamental limits of trajectory imitation using affine feedback.**

I. Introduction

Deception — the art of making an adversary develop wrong beliefs to enable completion of a long-term goal — is of obvious interest to a variety of defense [4], bargaining [2], and cybersecurity scenarios [17]. In scenarios that involve a deceptive agent physically progressing towards its target, deception often relies on using the agent's position and trajectory to induce incorrect beliefs about its intentions [8], [20], [22], [23].

This paper considers a deception scenario where an agent seeks to follow a reference trajectory starting from a given state as closely as possible, up to a certain time, *while ultimately converging to its desired target*. Fig. 1 illustrates such an occurrence, motivated by prior investigations of a human recognition of robot objectives from partial data [8].

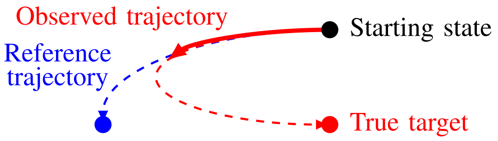

Fig. 1. An illustration of trajectory imitation. The agent's trajectory, in red, follows the reference trajectory (in blue) almost perfectly until a given time, and an observer might thus conclude that the agent is indeed progressing towards the blue target. However, the agent ultimately converges to a different target.

Existing efforts on deception in control often make highly simplifying assumptions on the agent dynamics [10], [23] or allow the use of complicated control laws [11], [19], which may render the problem of optimal trajectory imitation unrealistically effortless. In contrast, this paper recognizes that control systems are often, due to design and computational limitations, forced to use simple control laws even in complex environments [7]. Such a constraint is notably present in the related phenomenon of *motion camouflage*, where predator animals use a trajectory that imparts misleading beliefs about their motion from the point of view of their prey [12], [16]. While we recognize that an appropriate class of control laws may depend on the application domain, in this paper we consider the problem of optimal deceptive trajectory imitation by an agent capable only of using *time-invariant affine feedback control*.

In the framework of linear discrete-time systems, in Section II we formalize the problem of trajectory imitation as that of constrained optimization where the objective function is a metric of "partial reference tracking" — difference between the reference trajectory and the produced trajectory during the observation period — and the constraints ensure that the produced trajectory converges to a desired target. With an appropriate choice of metric, in Section III we show that this generally nonlinear problem may be approximated by a solution to a constrained ordinary least squares problem, and remove the constraints by pursuing a Tikhonov-like regularization method. Finally, in Section IV we illustrate the quality and fundamental limits of the obtained solutions through extensive numerical experimentation, showing that they outperform a naive controller designed to exactly follow a reference trajectory for a short amount of time.

Notation. Symbol $\mathbb{N}$ denotes the set of all positive integers. For $n \in \mathbb{N}$, $[n]$ denotes $\{1, \ldots, n\}$. We denote the set of all real $m \times n$ matrices by $\mathbb{R}^{m\times n}$. For a matrix $A \in \mathbb{R}^{n\times n}$, $\rho(A)$ denotes its *spectral radius*, i.e., the maximal magnitude of one of its eigenvalues. Symbol I denotes an identity square matrix whose size will be given in the subscript or will be clear from the context. We denote an ordered K-tuple $(x_1, \ldots, x_K)$, where $x_i \in \mathbb{R}^n$ for all i, by $x_{1:K}$. Unless stated otherwise, $\|\cdot\|$, without additional subscripts, denotes the usual vector and matrix 2-norms. To avoid confusion between a time index and a matrix transpose, we use M^τ to denote the transpose of a matrix M. For matrices M and N, $M \otimes N$ denotes their Kronecker product.

II. Problem Statement

We consider a scenario where a deceptive agent ("impostor") operates with linear dynamics

$$x_{t+1} = Ax_t + Bu_t, \tag{1}$$

[1] M. Ornik is with the Department of Aerospace Engineering and the Coordinated Science Laboratory, University of Illinois Urbana-Champaign, Urbana, IL 61801, USA. Email: mornik@illinois.edu

where $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{n \times m}$. The impostor seeks to deceive an observer who expects that the impostor is proceeding towards a *publicized target* x_{pub}. Namely, upon starting its mission at $x_0 \in \mathbb{R}^n$, the impostor publicizes both x_{pub} and its purported affine time-invariant control law

$$u_t^{ref} = K^{ref} x_t + L^{ref} \tag{2}$$

which would enable it to converge to x_{pub}. This control law generates a *reference trajectory* $x^{ref} = (x_0^{ref}, x_1^{ref}, \ldots)$ by applying control law (2) to (1) with $x_0^{ref} = x_0$.

To avoid technicalities and ease notation, we make the following assumption.

Assumption 1:

(a) System (1) is fully actuated, i.e., $\mathrm{Im}(B) = \mathbb{R}^n$.

(b) The publicized target is at the origin, i.e., $x_{pub} = 0$.

Removing Assumption 1(a) would not significantly change the discussion in the remainder of the paper, but would result in the requirement that some of the decision variables introduced in later optimization problems need to lie in particular affine spaces. Assumption 1(b) is trivially satisfied without loss of generality through coordinate translation.

Given that $x_t^{ref} \to x_{pub} = 0$, we know that 0 is an equilibrium of (1)–(2), so $BL^{ref} = 0$. Since the exact choice of L^{ref} does not affect x^{ref} as long as BL^{ref} remains the same, we take $L^{ref} = 0$. Thus, $u_t^{ref} = K^{ref} x_t$.

In order to imitate the reference trajectory while using a simple controller and achieving its ultimate goal, the impostor uses an affine control law

$$u_t^{imp} = K^{imp} x_t^{imp} + L^{imp} \tag{3}$$

so that the agent's true trajectory $x^{imp} = (x_0^{imp}, x_1^{imp}, \ldots)$ is obtained from applying control law (3) to dynamics (1) with $x_0^{imp} = x_0$.

To ensure the convergence of the impostor's trajectory x^{imp} to a target, we first assume that the impostor needs to choose K^{imp} and L^{imp} such that $\rho(A + BK^{imp}) < 1$. Indeed, x_t^{imp} will then converge to some state [9]. In order for the impostor to converge exactly to its desired target x_{tar}, it additionally needs to choose K^{imp} and L^{imp} such that $(A + BK^{imp})x_{tar} + BL^{imp} = x_{tar}$, i.e., $BL^{imp} = (I - A - BK^{imp})x_{tar}$ [9].

Remark 1: As the starting state x_0 is fixed, it may be the case that some of the eigenvalues of $A + BK^{imp}$ need not actually be in the unit circle to obtain convergence. However, we continue to require $\rho(A + BK^{imp}) < 1$ to avoid a technically cumbersome discussion.

We assume that the observer is able to see the impostor's partial trajectory $x_{0:T}^{imp} = (x_0^{imp}, x_1^{imp}, \ldots, x_T^{imp})$ until a given time T known to the impostor. Based solely on that trajectory, the observer seeks to distinguish between a well-meaning agent who follows the reference trajectory, and will thus presumably proceed towards the publicized target x_{pub}, and a malicious actor who does not. The observer is aware that its observations might not be perfect so it cannot exactly identify the applied control law. Instead, its belief that the agent is following the reference trajectory increases with a decrease in the distance between $x_{0:T}^{imp}$ and $x_{0:T}^{ref}$ in some appropriate metric. The impostor's goal is thus the following.

Problem 1: Let $T \in \mathbb{N}$, $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $K^{ref} \in \mathbb{R}^{m \times n}$, and $x_0, x_{tar} \in \mathbb{R}^n$. Let x^{ref} denote a solution to (1) with $u_t^{ref} = K^{ref} x_t^{ref}$ and $x_0^{ref} = x_0$. For $K^{imp} \in \mathbb{R}^{n \times n}$ and $L^{imp} \in \mathbb{R}^m$, let x^{imp} denote a solution to (1)–(3) with $x_0^{imp} = x_0$. Determine

$$\begin{aligned} \operatorname*{argmin}_{K^{imp}, L^{imp}} \quad & \|x_{0:T}^{err}\|_{\hat{1}} \\ \text{such that} \quad & \rho(A + BK^{imp}) < 1, \\ & BL^{imp} = (I - A - BK^{imp})x_{tar}, \end{aligned} \tag{4}$$

where $x^{err} = x^{imp} - x^{ref}$, and $\|x_{0:T}^{err}\|_{\hat{1}} = \sum_{t=0}^{T} \|x_t^{err}\|$, which we will refer to as the 1*-error norm*.

We acknowledge that the 1-error norm is not the only possibly intriguing metric, but note that many other metrics are equivalent to $\|x_{1:T}^{err}\|_{\hat{1}}$ in the sense of equivalence of norms. We omit their discussion because of page constraints.

Regardless of the chosen error norm, attempting to produce the optimal controller for the impostor by solving Problem 1 faces three challenges:

1° *The problem might not admit a solution*: set $\{K^{imp} \mid \rho(A + BK^{imp}) < 1\}$ is not compact. In practice, we will focus on finding an approximate solution to related optimization problems which penalize $\|x_{0:T}^{err}\|_{\hat{1}}$ and result in K^{imp} such that $\rho(A + BK^{imp}) < 1$.

2° Constraints on $\rho(A + BK^{imp})$ are possibly difficult to computationally tackle, given that they yield bounds on the roots of a polynomial, with little additional structure.

3° The dependence of the objective function $\|x_{0:T}^{err}\|_{\hat{1}}$ on the decision variables K^{imp} and L^{imp} seems at first glance to be complicated and unlikely to yield structures such as linearity or convexity.

We will attack 1° and 2° by slightly redefining the problem, thus yielding matrices K^{imp*} and L^{imp*} which satisfy the constraints in (4) while minimizing the weighted sum of $\|x_{0:T}^{err}\|_{\hat{1}}$ and an additional penalty related to $\rho(A + BK^{imp})$. Before we do so, we tackle 3° by providing an explicit expression for x_t^{err} dependent on K^{imp}. This expression will ultimately yield an approximation of Problem 1 by an ordinary least squares problem.

A. Closed-Form Error

Our approach to expressing x_t^{err} is motivated by a discrete interpretation of Grönwall's lemma [15]. Throughout the remainder of the paper, we will be denoting $A + BK^{ref}$ by M^{ref} and analogously $A + BK^{imp}$ by M^{imp}.

Proposition 1: Let $K^{ref}, K^{imp} \in \mathbb{R}^{m \times n}$, $L^{imp} \in \mathbb{R}^m$ and $x_0, x_{tar} \in \mathbb{R}^n$. Let x^{err} be defined as in Problem 1. Assume that $BL^{imp} = (I - M^{imp})x_{tar}$. Then,

$$x_t^{err} = x_{tar} - (M^{ref})^t x_0 - (M^{imp})^t (x_{tar} - x_0)$$

for all $t \in \mathbb{N} \cup \{0\}$.

Proof: We move by induction on t. For $t = 0$, $x_0^{err} = x_0^{imp} - x_0^{ref} = x_0 - x_0 = 0$, which matches the claim.

Assume now that the claim holds for some $t \in \mathbb{N} \cup \{0\}$. We will prove it for $t+1$.

By definitions of x^{ref}, x^{imp} and x^{err}, $x^{err}_{t+1} = x^{imp}_{t+1} - x^{ref}_{t+1} = (M^{imp})x^{imp}_t + BL^{imp} - (M^{ref})x^{ref}_t = (M^{imp})x^{err}_t + (M^{imp})x^{ref}_t + BL^{imp} - (M^{ref})x^{ref}_t = BL^{imp} + (M^{imp})x^{err}_t + (M^{imp} - M^{ref})x^{ref}_t$. Plugging in the claim for t into this equality, we obtain $x^{err}_{t+1} = BL^{imp} + (M^{imp})x_{tar} - (M^{imp})(M^{ref})^t x_0 - (M^{imp})^{t+1}(x_{tar} - x_0) + (M^{imp} - M^{ref})x^{ref}_t$. Since we assumed $BL^{imp} = (I - M^{imp})x_{tar}$, this equality can be expressed as $x^{err}_{t+1} = x_{tar} - (M^{imp})^{t+1}(x_{tar} - x_0) - (M^{imp})(M^{ref})^t x_0 + (M^{imp} - M^{ref})x^{ref}_t$.

Finally, noting that $(M^{ref})^t x_0 = x^{ref}_t$ from the definition of x^{ref}, we obtain $x^{err}_{t+1} = x_{tar} - (M^{imp})^{t+1}(x_{tar} - x_0) - (M^{ref})^{t+1}x_0$. ■

Because of Assumption 1(a), the correspondence $K \mapsto A + BK = M$ is surjective, i.e., for every $M \in \mathbb{R}^{n\times n}$, there exists at least one K such that $M = A + BK$, and Proposition 1 confirms that the particular choice of L^{imp} does not affect the objective function in (4), as long as $BL^{imp} = (I - M^{imp})x_{tar}$. Hence, Problem 1 reduces to

$$\begin{aligned} \underset{M^{imp} \in \mathbb{R}^{n\times n}}{\text{argmin}} \quad & \|x^{err}_{0:T}\|_{\hat{1}} \\ \text{such that} \quad & \rho(M^{imp}) < 1. \end{aligned} \tag{5}$$

Any solution M^{imp*} to (5) generates a solution to (4), by taking (K^{imp*}, L^{imp*}) such that $A + BK^{imp*} = M^{imp*}$ and $BL^{imp*} = (I - M^{imp*})x_{tar}$. We note that removing Assumption 1(a) would prevent Problem 1 from being written in the simple form (5). Instead, the decision variables would remain K^{imp} and L^{imp} and the problem would have an additional constraint $BL^{imp} = (I - A - BK^{imp})x_{tar}$.

We now proceed to further reframe problem (5) and propose its approximate solution.

III. Approximately Optimal Imitation

We begin our approach by noting that Proposition 1 offers a suggestive way of writing Problem 1. Namely, if we denote $v_t = x_{tar} - (M^{ref})^t x_0$, Problem 1 becomes

$$\begin{aligned} \underset{M^{imp} \in \mathbb{R}^{n\times n}}{\text{argmin}} \quad & \sum_{t=0}^{T} \|v_t - (M^{imp})^t v_0\| \\ \text{such that} \quad & \rho(M^{imp}) < 1. \end{aligned} \tag{6}$$

Although (6) is still a nonlinear optimization problem, its form allows us to approximate a solution while yielding theoretical guarantees on the quality of the approximation.

A. Power Mean Approximation

For a matrix M such that $\|M\| \leq 1$, let us define $h_t(M) = \|v_t - Mv_{t-1}\|$, for $t \in \mathbb{N}$. Then,

$$\begin{aligned} & \|v_t - M^t v_0\| \\ & = \|v_t - Mv_{t-1} + Mv_{t-1} - M^2 v_{t-2} + \ldots - M^t v_0\| \\ & \leq h_t(M) + \|M\| h_{t-1}(M) + \ldots + \|M\|^{t-1} h_1(M) \\ & \leq \sum_{k=1}^{t} h_k(M). \end{aligned}$$

Thus, *if we only consider matrices* M^{imp} such that $\|M^{imp}\| \leq 1$, the objective function of (6) is bounded from above by $\sum_{t=1}^{T}(T+1-t)h_t(M^{imp})$, which is, by the weighted power mean inequality [5], in turn bounded from above by

$$\sqrt{\frac{T(T+1)}{2} \sum_{t=1}^{T} (T+1-t)(h_t(M^{imp}))^2}. \tag{7}$$

We now use the standard approach of finding an approximate solution to an optimization problem by finding a solution to a problem generated by an upper bound of the original problem [1], [3], [24]. We thus seek to find

$$\begin{aligned} \underset{M^{imp} \in \mathbb{R}^{n\times n}}{\text{argmin}} \quad & \sum_{t=1}^{T} (T+1-t)\|v_t - (M^{imp})v_{t-1}\|^2 \\ \text{such that} \quad & \rho(M^{imp}) < 1. \end{aligned} \tag{8}$$

We note that (7) is only *guaranteed* to be an upper bound to the objective function in (6) on the domain given by $\|M^{imp}\| \leq 1$ and not all of $\rho(M^{imp}) < 1$. We continue considering problem (8) as a surrogate for (6), noting the connection between $\rho(M^{imp})$, $\|M^{imp}\|$, and the Frobenius norm $\|M^{imp}\|_F$ in Section III-B.

Problem (8) is now a constrained *ordinary least squares* problem. While the decision variable is a matrix, the objective function from (8) can be alternatively writen as

$$\sum_{t=1}^{T} (T+1-t)\|v_t - (v_{t-1}^{\tau} \otimes I_{n\times n})\text{vec}(M^{imp})\|^2,$$

where vec denotes matrix vectorization [26]. We now proceed to tackle the requirement that $\rho(M^{imp}) < 1$.

B. Tikhonov-Like Regularization

As mentioned, constraint $\rho(M^{imp}) < 1$ yields a non-convex and non-compact domain. We approach this challenge and provide an approximate solution by building on a classical method of *Tikhonov regularization* [13]. Tikhonov regularization is traditionally used to solve constrained least squares problems $\min_X \|A - BX\|^2$ where there exists a hard upper bound $\|X\| \leq \alpha$; it functions by instead minimizing an *unconstrained* least squares problem $\min_X \|A - BX\|^2 + \gamma\|X\|^2$. It is trivial to show that, for a large enough γ, the solution of the latter problem will certainly satisfy $\|X\| \leq \alpha$, as $\|X\|$ is afforded a sufficiently large weight in the objective function.

In the case of Problem 1 and subsequent relaxations, the constraint is on $\rho(M^{imp})$ instead of the norm of M^{imp}. However, $\rho(M) \leq \|M\|_F$ and $\|M\| \leq \|M\|_F$ for all matrices M [6], [14], where $\|\cdot\|_F$ stands for the Frobenius norm, i.e., $\|M\|_F = \|\text{vec}(M)\|$. Hence, as for standard Tikhonov regularization, for any large enough γ, a minimizer M^{imp*} of

$$\begin{aligned} & \sum_{t=1}^{T} (T+1-t)\|v_t - (v_{t-1}^{\tau} \otimes I_{n\times n})\text{vec}(M^{imp})\|^2 \\ & + \gamma\|\text{vec}(M^{imp})\|^2 \end{aligned} \tag{9}$$

will exist and satisfy $\|\text{vec}(M^{imp*})\| < 1$, and thus, $\rho(M^{imp*}) < 1$ as well as $\|M^{imp*}\| < 1$, required for the guaranteed upper bound established in Section III-A.

Naturally, in (9) we consider the smallest $\gamma \geq 0$ such that a solution to (9) satisfies $\rho(M^{imp}) < 1$. While possible analytical methods for choosing an appropriate γ are of interest to future work, in our numerical experiments we determine an appropriate γ by starting from $\gamma = 0$ and successively solving (9) while increasing γ until a solution that satisfies $\rho(M^{imp}) < 1$ is found.

Problem (9) is an ordinary least squares problem, with a solution that can be found analytically for every γ [13]. While it does not provide an exact solution to Problem 1 — nor does such a solution generally exist, as the problem is ill-posed — in practice we find that it often generates trajectories that imitate the reference trajectory well until time T, while *provably* converging towards the desired target. The following section will present relevant numerical examples and compare them to a naive approximate solution of (6) obtained by setting $Mv_{t-1} = v_t$ for all $t \in [T]$.

IV. Numerical Experiments

Having proposed a method for trajectory imitation, we now illustrate its effectiveness. We also provide a preliminary analysis on the quality of imitation given the length of the observation period, the initial state and target positions.

In the absence of previously established benchmarks, we compare the proposed methods against a method that generates a trajectory which exactly equals the reference trajectory for as long as possible. We first briefly introduce this method.

A. Naive Method of Imitation

From (6), it follows that $\|x^{err}_{0:T}\|_{\hat{1}} = 0$ if and only if $(M^{imp})^t v_0 = v_t$ for all $t \in [T]$. While such a matrix M^{imp} might not exist in general, and especially so with the constraint $\rho(M^{imp}) < 1$, the proposed naive method aims to satisfy

$$M^{imp} v_{t-1} = v_t \text{ for all } t \in [k], \tag{10}$$

for as large of a k as possible. While the exact maximal k depends on the linear independence of the subsets of $\{v_0, v_1, \ldots, v_k\}$, by definition of $v_t = x_{tar} - (M^{ref})^t x_0$ we generically [25] expect $k = n$, so we can define

$$M^{imp} = \begin{bmatrix} v_1 & v_2 & \cdots & v_n \end{bmatrix} \begin{bmatrix} v_0 & v_1 & \cdots & v_{n-1} \end{bmatrix}^{-1}. \tag{11}$$

Naturally, if $n < T$, (11) does not provide any guarantees for a low overall error $\|x^{err}_{1:T}\|_{\hat{1}}$. Additionally, (11) does not guarantee that $\rho(M^{imp}) < 1$. While the latter guarantee may be obtained by choosing a maximal $k \leq n$ such that there exists a matrix M^{imp} that satisfies (10) and $\rho(M^{imp}) < 1$, the determination of such a matrix — and even verification that it exists — may be difficult. We thus remain with (11) as the simple benchmark which generically produces a controller which sets $x^{err}_t = 0$ for $t \in [n]$, even if it may not guarantee convergence of x^{imp} to x_{tar}.

We now proceed to implement the proposed method for trajectory imitation using approximate minimization of the 1-error norm and compare it to the naive method (11).

B. Results

To run the experiments, we consider $M^{ref} = (I_{n\times n} + N_{n\times n})/(\rho(I_{n\times n}+N_{n\times n})+0.1)$, where N is a matrix whose each element is drawn randomly from a normal distribution with mean 0 and standard deviation 0.1. This form was chosen to ensure that the reference trajectory x^{ref} converges to 0, as well as to intuitively produce a "smoother looking" x^{ref}, as it may otherwise be difficult to visually judge the quality of the obtained trajectory imitation. The values of matrices A, B, and K^{ref} are themselves irrelevant in simulations as long as $A + BK^{ref} = M^{ref}$, although the knowledge of A and B is course necessary to compute the control law $u = K^{imp}x + L^{imp}$.

To compare the errors across methods and visualize the system state easily, we choose $n = 2$, initial state $x_0 = \begin{pmatrix} 1 & 5 \end{pmatrix}^\tau$, and target state $x_{tar} = \begin{pmatrix} 2 & 0 \end{pmatrix}^\tau$. The chosen values of x_0 and x_{tar} intuitively give the impostor a "decent chance" at emulating the system trajectory, but still make it a non-trivial task. We provide a more thorough discussion of the dependence of the error on x_0 and x_{tar} in Section IV-C.

We assume that the observer is able to see the first $T = 10$ states of the impostor's trajectory. This value of T is a trade-off that allows the observer to see some substantial segment of the impostor's trajectory, without showing it almost the entire path to the target. At the end of Section IV-B.2 we briefly comment on the quality of the trajectory imitation in dependence to the length of the observation period.

For the regularization in (9), γ starts from 0 and increases in steps of 0.1 until the found solution satisfies $\rho(M^{imp}) < 1$. Naturally, setting a lower γ step may produce better underapproximations, but increases the computational effort.

To provide a visual idea of the derived impostor trajectories, we begin by considering one randomly drawn M^{ref}.

1) Single Example: Fig. 2 illustrates the trajectories obtained using the proposed method from Section III and the naive method (11) for

$$M^{ref} = \begin{pmatrix} 0.8167 & -0.0520 \\ -0.0551 & 0.8926 \end{pmatrix},$$

with the above choice of parameters.

As seen in Fig. 2, the trajectory generated by (9) is nearly indistinguishable from the reference trajectory in the first $T = 10$ time steps. On the other hand, while the naive method produces a trajectory that exactly equals the reference for the first $n = 2$ time steps, it starts visibly diverging from it earlier than the trajectory generated by (9). Eventually, the trajectory x^{imp} generated by (9) converges to x_{tar} as guaranteed, and so does the naively generated trajectory in this case, while x^{ref} of course converges to 0.

Naturally, the visual appearance of x^{ref} and corresponding x^{imp} will generally differ depending on M^{ref}. We thus proceed to illustrate the quality of the proposed trajectory imitation on a larger set of samples.

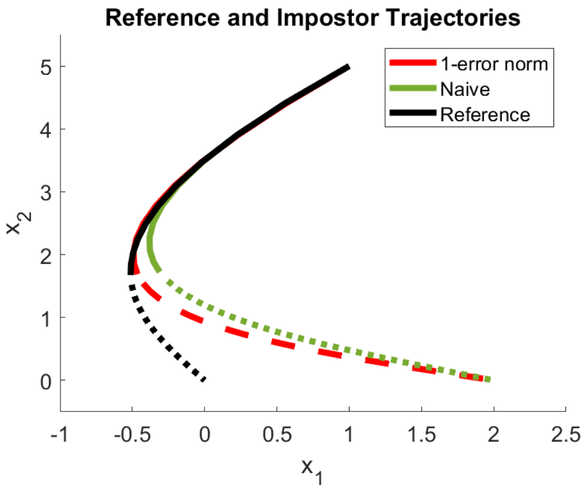

Fig. 2. Reference trajectory for a given M^{ref}, marked in black, and impostor trajectory generated by (9) marked in red. The benchmark naive impostor trajectory generated by (11) is marked in green. Full lines illustrate the trajectories until time $T = 10$, while dotted and dashed lines illustrate the trajectories after that time. In all figures, labels x_1 and x_2 denote the coordinates in the state space and not states at particular times.

2) Average Error Comparison: In this experiment, we implement the proposed method with 1000 matrices M^{ref}, randomly generated and with the same parameters as above.

While the proposed method is approximate and provides no hard guarantees of producing better results than the naive method (11), it yields a *lower* 1-*error norm in 97.6% of cases*, including the 24.7% of cases where the naive method does not yield M^{imp} which satisfies the constraint $\rho(M^{imp}) < 1$ at all. The proposed method, naturally, always satisfies $\rho(M^{imp}) < 1$.

Even though the proposed method only approximates the optimal trajectory imitation for 1-error norm, it also strongly outperforms the naive method when considering the 2-*error* and ∞-*error* norms defined by $\|x^{err}_{0:T}\|_{\hat{2}} = \sum_{i=0}^{T} \|x^{err}_t\|^2$ and $\|x^{err}_{0:T}\|_{\hat{\infty}} = \max_{0 \leq t \leq T} \|x^{err}_t\|$, respectively. Namely, produces a lower 2-error norm in 97.7% of cases, and a lower ∞-norm in 97.2% of cases.

A further slight increase in performance can be obtained by directly solving the nonlinear Problem 1 instead of its approximation (9). For instance, solving Problem 1 to minimize the 2-error norm yields a *constrained nonlinear least squares problem* in M^{imp}, which — after a Tikhonov-like regularization from Section III-B — may be tackled using a variety of prior work [18], [21], as well as general methods on nonlinear optimization. Doing so indeed reaps slight benefits in performance, producing lower 1-, 2-, and ∞-error norms than the naive method in $99.9\% - 100\%$ of the cases, and even lower error norms than the proposed method (9) in $97.5\% - 98.7\%$ of the cases. However, the computational price to pay for such an improvement may be vast. Preliminary experimental results, omitted from this discussion due to page limits, show that (9), implemented in MATLAB, produces results momentarily even for $n \approx 20$ and/or $T \approx 40$, and is faster than using MATLAB's native nonlinear optimization functionality to minimize the 2-error norm by, roughly, a factor of $100\times$.

Naturally, the error in the trajectory imitation depends on the length T of the observation period. For $T \leq n$, even the naive method will generically produce M^{imp} such that $\|x^{err}\|_{\hat{1}} = \|x^{err}\|_{\hat{2}} = \|x^{err}\|_{\hat{\infty}} = 0$, although it will not guarantee convergence to x_{tar}. On the other hand, $T \to +\infty$, none of the methods are able to provide a truly meaningful trajectory imitation, as the trajectories converge to different states. However, our preliminary analysis shows that the proposed method vastly outperforms the naive one across observation periods.

Having illustrated the quality of the proposed method for a particular initial state and target, we conclude this section by considering the success of trajectory imitation as a function of x_0 and x_{tar}.

C. Dependence on Initial State and Target Position

To describe the "difficulty of imitation" with respect to the initial state and the target, we first continue with the parameters from the beginning of Section IV-B, except we vary the initial state $x_0 \in S_0 = \{-3, -2.9, \ldots, 4\} \times \{-3, -2.9, \ldots, 8\}$, while keeping $x_{tar} = \begin{pmatrix} 2 & 0 \end{pmatrix}^\tau$. Fig. 3 displays the average 1-error norm with respect to x_0, obtained by choosing 100 matrices M^{ref} for each $x_0 \in S_0$, for a total of 788100 simulations. We always use method (9) to generate M^{imp}, with a step in γ equal to 1.

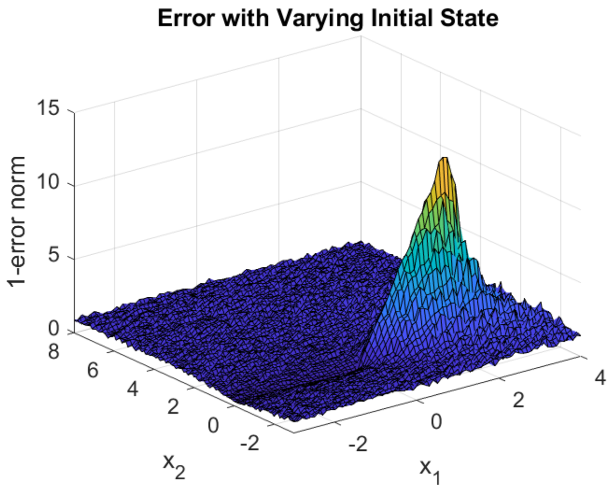

Fig. 3. The 1-error norm between x^{imp} generated by (9) and x^{ref}, with varying initial state x_0. Labels x_1 and x_2 denote the coordinates in the state space and not states at a particular time.

As Fig. 3 shows, the error increases as x_0 comes closer to the target state x_{tar}. Along with numerical reasons due to $v_0 \to 0$, we suggest that the reason for such a phenomenon may be that such parameters force the impostor to remain close to x_{tar} throughout its trajectory, while the reference trajectory moves away from $x_{tar} \approx x_0$ and towards the origin. Fig. 4 illustrates such a situation.

We now consider the case where the initial state x_0 is fixed to $\begin{pmatrix} 1 & 5 \end{pmatrix}^\tau$, while x_{tar} varies in $[-3, 4] \times [-3, 8]$. Fig. 5 displays the average error with respect to x_{tar}.

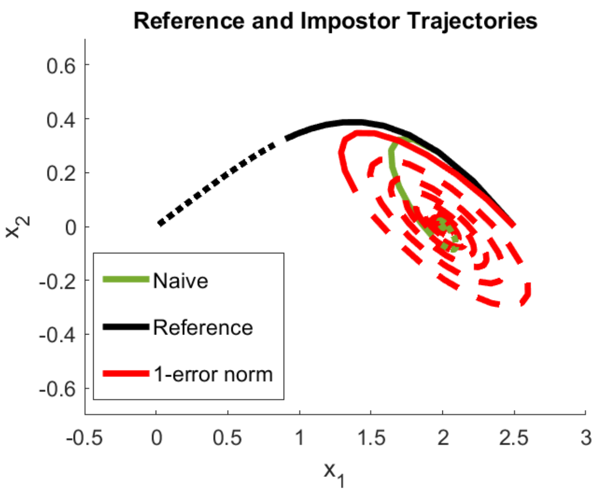

Fig. 4. Reference and impostor trajectories for a given M^{ref}, in the case where x_0 and x_{tar} are close, colored and with the notation as in Fig. 2.

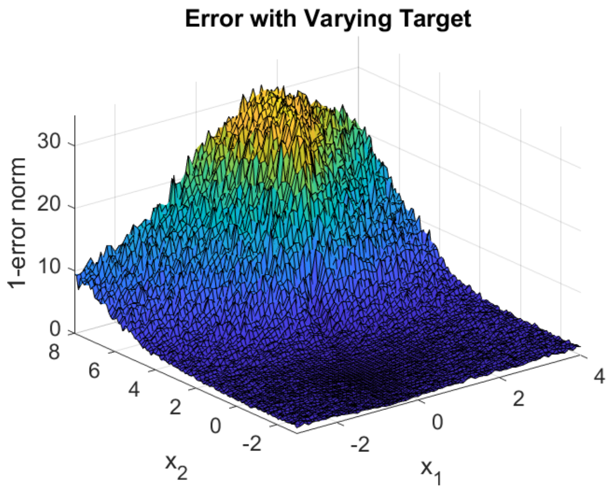

Fig. 5. The 1-error norm between x^{imp} generated by (9) and x^{ref}, with varying target x_{tar}. Labels x_1 and x_2 denote the coordinates in the state space and not states at a particular time.

Fig. 5 once again shows how heavily the relative positions — and not just distances — of the publicized target $x_{pub} = 0$, initial state x_0, and true target x_{tar} influence the impostor's success. In particular, x_{tar} that is chosen to the "north" of x_0 will often yield a much worse imitated trajectory using (9) than x_{tar} chosen to the "south". The problem of describing the dependence of the imitation error on target placement, and consequently the *problem of optimal target placement* — already identified in [22], but in a significantly simpler model — remain central issues for future work.

Acknowledgment

The author thanks Roy Dong for discussions on approximations of relevant optimization problems.

References

[1] Z.-Z. Bai, G. H. Golub, and M. K. Ng, "Hermitian and skew-Hermitian splitting methods for non-hermitian positive definite linear systems," *SIAM Journal on Matrix Analysis and Applications*, vol. 24, no. 3, pp. 603–626, 2003.

[2] M. A. Bhatt, T. Lohrenz, C. F. Camerer, and P. R. Montague, "Neural signatures of strategic types in a two-person bargaining game," *Proceedings of the National Academy of Sciences*, vol. 107, no. 46, pp. 19 720–19 725, 2010.

[3] L. Blackmore and B. Williams, "Finite horizon control design for optimal discrimination between several models," in *45th IEEE Conference on Decision and Control*, 2006, pp. 1147–1152.

[4] J. W. Caddell, *Deception 101 – Primer on Deception*. Strategic Studies Institute, 2004.

[5] Z. Cvetkovski, *Inequalities: theorems, techniques and selected problems*. Springer, 2012.

[6] N. A. Derzko and A. M. Pfeffer, "Bounds for the spectral radius of a matrix," *Mathematics of Computation*, vol. 19, no. 89, pp. 62–67, 1965.

[7] J. C. Doyle, B. A. Francis, and A. R. Tannenbaum, *Feedback Control Theory*. Macmillan, 1992.

[8] A. D. Dragan, R. M. Holladay, and S. S. Srinivasa, "An analysis of deceptive robot motion," in *Robotics: Science and Systems*, 2014.

[9] L. Farina and S. Rinaldi, *Positive linear systems: theory and applications*. John Wiley & Sons, 2000.

[10] W. Fu, J. Qin, W. X. Zheng, Y. Chen, and Y. Kang, "Resilient cooperative source seeking of double-integrator multi-robot systems under deception attacks," *IEEE Transactions on Industrial Electronics*, vol. 68, no. 5, pp. 4218–4227, 2020.

[11] Z. E. Fuchs, "Cooperative control strategies and deception in adversarial systems," Ph.D. dissertation, University of Florida, 2012.

[12] P. Glendinning, "The mathematics of motion camouflage," *Proceedings of the Royal Society B: Biological Sciences*, vol. 271, no. 1538, pp. 477–481, 2004.

[13] M. Gockenbach, *Linear Inverse Problems and Tikhonov Regularization*. Mathematical Association of America, 2016.

[14] G. H. Golub and C. F. Van Loan, *Matrix Computations*. Johns Hopkins University Press, 2013.

[15] J. M. Holte, "Discrete Gronwall lemma and applications," MAA-NCS North Central Section Meeting, University of North Dakota, 2009.

[16] E. W. Justh and P. S. Krishnaprasad, "Steering laws for motion camouflage," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 462, no. 2076, pp. 3629–3643, 2006.

[17] C. Kwon, W. Liu, and I. Hwang, "Security analysis for cyber-physical systems against stealthy deception attacks," in *American Control Conference*, 2013, pp. 3344–3349.

[18] J. B. Lasserre, K.-C. Toh, and S. Yang, "A bounded degree SOS hierarchy for polynomial optimization," *EURO Journal on Computational Optimization*, vol. 5, no. 1–2, pp. 87–117, 2017.

[19] M. Li, Y. Chen, and Y. Liu, "Sliding-mode secure control for jump cyber–physical systems with malicious attacks," *Journal of the Franklin Institute*, vol. 358, no. 7, pp. 3424–3440, 2021.

[20] P. Masters and S. Sardina, "Deceptive path-planning," in *26th International Joint Conference on Artificial Intelligence*, 2017, pp. 4368–4375.

[21] H. Mohammad, M. Y. Waziri, and S. A. Santos, "A brief survey of methods for solving nonlinear least-squares problems," *Numerical Algebra, Control & Optimization*, vol. 9, no. 1, pp. 1–13, 2019.

[22] M. Ornik, "Measuring target predictability for optimal environment design," in *59th IEEE Conference on Decision and Control*, 2020, pp. 5023–5028.

[23] M. Ornik and U. Topcu, "Deception in optimal control," in *2018 56th Annual Allerton Conference on Communication, Control, and Computing*, 2018, pp. 821–828.

[24] J. Roll, A. Nazin, and L. Ljung, "Local modelling of nonlinear dynamic systems using direct weight optimization," in *13th IFAC Symposium on System Identification*, vol. 36, no. 16, 2003, pp. 1513–1518.

[25] K. Tchoń, "On generic properties of linear systems: An overview," *Kybernetika*, vol. 19, no. 6, pp. 467–474, 1983.

[26] D. A. Turkington, *Generalized Vectorization, Cross-Products, and Matrix Calculus*. Cambridge University Press, 2013.