BAYESIAN OPTIMIZATION FOR INFORMATIVE EXCITATION DESIGN FOR ESTIMATION OF INERTIA PARAMETERS IN SPACECRAFT RENDEZVOUS & DOCKING

Himadri Basu* **and Ricardo G. Sanfelice†**

This work proposes a Bayesian learning–based framework for designing informative torque inputs to estimate the inertial properties of a tumbling rigid body. Random or unstructured torque pulses may fail to sufficiently excite the system dynamics, leading to poor observability and incomplete or normalized parameter estimates. To address this limitation, we introduce a Bayesian optimization–based excitation design that selects sparse torque inputs online to enhance parameter observability. By actively learning the torque pulses that are most informative, the proposed approach enables efficient estimation of the moment of inertia using state-of-the-art estimation algorithms.

INTRODUCTION

Space exploration and on-orbit servicing missions require technologies for rendezvous, capture, and docking with objects in orbit. These missions include refueling, repairing, deorbiting malfunctioning satellites, and removing orbital debris. While several missions have demonstrated the feasibility of capturing cooperative targets with known properties, interacting with non-cooperative targets—such as tumbling satellites, debris, or defunct spacecraft—remains significantly more challenging and risky. Such targets are typically non-communicative and their physical properties, including mass, center of mass, and moments of inertia, are often unknown. Vision-based and active sensing technologies, such as LIDAR and time-of-flight cameras, can provide relative pose information [1], but they do not directly yield inertial properties, which are essential for predicting and controlling the target's motion.

Accurate knowledge of a target's mass and moments of inertia is critical for safe rendezvous, capture, and post-capture stabilization. For instance, the motion of a potential grasping point on a tumbling target must be propagated to plan a safe capture maneuver. After grasping, the combined dynamics of the servicing spacecraft and target must be stabilized, which is not possible without reliable estimates of the target's inertial parameters [2, 3]. This requirement is particularly evident in digital-twin-enabled on-orbit servicing architectures, where accurate inertial knowledge directly impacts state estimation, information sharing, and control coordination across multiple spacecraft subsystems [4]. In practice, because non-cooperative targets cannot be commanded or instrumented, their inertial properties must be inferred from external observations of their motion. Consequently, inertial parameter estimation becomes a prerequisite for both guidance and control during proximity operations and docking with unknown targets.

*Postdoctoral Researcher, Department of Electrical and Computer Engineering, University of California Santa Cruz, Santa Cruz, CA 95060, USA.

†Professor, Department of Electrical and Computer Engineering, University of California Santa Cruz, Santa Cruz, CA 95060, USA.

Several methods have been proposed to estimate the inertial properties of uncooperative space objects. Benninghoff et. al. [5] investigated the estimation of mass, center of mass, and moments of inertia for a tumbling target under zero external excitation. In the absence of applied torques, only a scaled or normalized version of the inertia tensor can be identified, limiting the estimation to relative parameters rather than true numerical values. Meng et al. [6] later introduced continuous excitation over a short time horizon to identify all inertial parameters of a non-cooperative object. While this approach improves identifiability, the use of random excitation may be impractical in realistic scenarios due to actuation limits, energy constraints, and operational safety considerations.

More recently, Boyacioglu et al. [7] analyzed nonlinear inertial parameter estimation for a two-dimensional rigid body using numerical observability tools. They showed that a nonzero net force is necessary to identify all inertial parameters and that the excitation profile strongly influences parameter observability. A similar observability analysis for inertial parameter estimation of a nonlinear rigid-body was conducted for planar multi-link systems with thrusters in [8]. Despite these advances, existing approaches face practical limitations. Methods based on continuous or arbitrary excitation can conflict with thruster constraints, safety requirements, and mission objectives. As a result, excitation may be insufficient to reliably estimate all inertia components, particularly in the presence of limited actuation authority and short observation windows.

Motivated by these challenges, this work focuses on the design of informative torque inputs to enable efficient and reliable estimation of inertial parameters for tumbling, uncooperative targets. Instead of relying on random or continuous excitation, we adopt a Bayesian learning–based optimization framework that selects sparse torque pulses online to maximize the observability of the target's inertial properties. By actively learning which inputs are most informative, the proposed approach enables accurate estimation of the moment of inertia using standard estimation algorithms such as the Extended Kalman Filter or Unscented Kalman Filter, while respecting practical actuation and safety constraints.

SYSTEM MODELS

Consider the rigid body rotational dynamics of a spacecraft expressed in a body-fixed frame $\mathcal{C}$ relative to an inertial frame $\mathcal{I}$, given by

$$J\dot{\omega}_{\mathcal{C}/\mathcal{I}} = -\omega_{\mathcal{C}/\mathcal{I}} \times \left(J\omega_{\mathcal{C}/\mathcal{I}}\right) + \tau, \tag{1}$$

where $\omega_{\mathcal{C}/\mathcal{I}} \in \mathbb{R}^3$ denotes the angular velocity of the body frame $\mathcal{C}$ relative to the inertial frame $\mathcal{I}$, $\tau \in \mathcal{U} \subset \mathbb{R}^3$ is the applied torque on the spacecraft, with $\mathcal{U}$ denoting the set of admissible control torques, and

$$J = \begin{bmatrix} J_{xx} & J_{xy} & J_{xz} \\ J_{xy} & J_{yy} & J_{yz} \\ J_{xz} & J_{yz} & J_{zz} \end{bmatrix} \tag{2}$$

is the symmetric positive definite inertia matrix of the rigid body. Let the available angular velocity measurements be given by

$$y_\omega = \omega_{\mathcal{C}/\mathcal{I}} + v_\omega, \quad v_\omega \sim \mathcal{N}(0, \sigma_\omega^2 I_3), \tag{3}$$

where v_ω represents zero-mean Gaussian measurement noise with standard deviation σ_ω. Let $q = \begin{bmatrix} q_0 & q_v^\top \end{bmatrix}^\top$ denote the unit quaternion representing the attitude of the body, where $q_0 \in \mathbb{R}$ and $q_v \in \mathbb{R}^3$ are the scalar and vector parts, respectively, satisfying $q_0^2 + \|q_v\|^2 = 1$. The corresponding

rotation matrix is denoted by $R(q) \in \text{SO(3)}$, where $\text{SO(3)} = \{R \in \mathbb{R}^{3\times3} \mid R^\top R = I_3,\ \det(R) = 1\}$.

Let $y_a \in \mathbb{R}^3$ and $y_m \in \mathbb{R}^3$ denote the body-frame measurements of the gravitational acceleration and the geomagnetic field, respectively, given by

$$y_a = R^\top(q)g + v_a, \qquad v_a \sim \mathcal{N}(0, \Sigma_a), \tag{4}$$

$$y_m = R^\top(q)m_e + v_m, \qquad v_m \sim \mathcal{N}(0, \Sigma_m), \tag{5}$$

where $g \in \mathbb{R}^3$ is the gravity vector and $m_e \in \mathbb{R}^3$ is the geomagnetic field vector, both expressed in the inertial frame. The noise processes v_a and v_m are assumed to be zero-mean Gaussian with covariances Σ_a and Σ_m, respectively. Using state-of-the-art state estimation techniques, such as the extended Kalman filter (EKF), the attitude quaternion q and the angular velocity $\omega_{\mathcal{C}/\mathcal{I}}$ can be estimated from the noisy measurements (y_ω, y_a, y_m).

Starting from the rotational dynamics (1), with $\omega \coloneqq \omega_{\mathcal{C}/\mathcal{I}} \in \mathbb{R}^3$ denoting the angular velocity of $\mathcal{C}$ relative to $\mathcal{I}$ expressed in $\mathcal{C}$, the equation can be rewritten as

$$\tau = J\dot{\omega} + [\omega]^\times J\omega, \tag{6}$$

where $[\omega]^\times \in \mathbb{R}^{3\times3}$ denotes the skew-symmetric matrix associated with $\omega = (\omega_x, \omega_y, \omega_z) \in \mathbb{R}^3$, defined such that $[\omega]^\times x = \omega \times x$ for all $x \in \mathbb{R}^3$, i.e.,

$$[\omega]^\times \coloneqq \begin{bmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & -\omega_x \\ -\omega_y & \omega_x & 0 \end{bmatrix}. \tag{7}$$

From (6), it is clear that the regression structure depends on the angular velocity ω and angular acceleration $\dot{\omega}$. In practice, ω is obtained from noisy gyroscope measurements, and $\dot{\omega}$ is not directly measured. Therefore, state-of-the-art Kalman filtering techniques (e.g., an EKF) are employed to obtain estimates $\hat{\omega} \in \mathbb{R}^3$ and $\dot{\hat{\omega}}\mathbb{R}^3$ from these noisy signals, suitable for subsequent parameter identification. Let

$$\theta := (J_{xx}, J_{xy}, J_{xz}, J_{yy}, J_{yz}, J_{zz})^\top \in \mathbb{R}^6. \tag{8}$$

be the vector of unknown inertia parameters. Since the inertia matrix J is symmetric and parameterized by θ, each component of τ is linear in the six independent inertia parameters. Collecting terms from (6) yields the linear-in-parameters regression form

$$\Psi(\hat{\omega}, \dot{\hat{\omega}})\,\theta = \tau, \tag{9}$$

where the regressor matrix $\Psi : \mathbb{R}^3 \times \mathbb{R}^3 \to \mathbb{R}^{3\times6}$ is defined as

$$\Psi(\hat{\omega}, \dot{\hat{\omega}}) \coloneqq \begin{bmatrix} \dot{\hat{\omega}}_x & \dot{\hat{\omega}}_y - \hat{\omega}_x\hat{\omega}_z & \dot{\hat{\omega}}_z + \hat{\omega}_x\hat{\omega}_y & -\hat{\omega}_y\hat{\omega}_z & \hat{\omega}_y^2 - \hat{\omega}_z^2 & \hat{\omega}_y\hat{\omega}_x \\ \hat{\omega}_x\hat{\omega}_z & \dot{\hat{\omega}}_x + \hat{\omega}_y\hat{\omega}_z & \hat{\omega}_z^2 - \hat{\omega}_x^2 & \dot{\hat{\omega}}_y & \dot{\hat{\omega}}_z - \hat{\omega}_x\hat{\omega}_y & -\hat{\omega}_x\hat{\omega}_z \\ 0 & \hat{\omega}_x^2 - \hat{\omega}_y^2 & \dot{\hat{\omega}}_x - \hat{\omega}_y\hat{\omega}_z & \hat{\omega}_x\hat{\omega}_y & \dot{\hat{\omega}}_y + \hat{\omega}_x\hat{\omega}_z & \dot{\hat{\omega}}_z \end{bmatrix}. \tag{10}$$

with the estimated angular velocity $\hat{\omega} = (\hat{\omega}_x, \hat{\omega}_y, \hat{\omega}_z) \in \mathbb{R}^3$ and angular acceleration $\dot{\hat{\omega}} = (\dot{\hat{\omega}}_x, \dot{\hat{\omega}}_y, \dot{\hat{\omega}}_z) \in \mathbb{R}^3$.

Let $\{t_k\}_{k\in\mathbb{N}}$ denote a sequence of discrete sampling instants. Define the stacked measurement vector $y_k \coloneqq (y_{\omega,k}, y_{a,k}, y_{m,k})$, where $y_{\omega,k}, y_{a,k}$ and $y_{m,k}$ are obtained from (3)–(5). Then, the state estimates $\hat{x}_k \coloneqq (q_k, \hat{\omega}_k, \theta_k) \in \mathbb{R}^7$ can be obtained via an EKF recursion approach as

$$(\hat{x}_k, P_k) = \mathcal{F}\big(\hat{x}_{k-1}, P_{k-1}, y_k, \tau_{k-1}\big), \tag{11}$$

where $\mathcal{F} : \mathbb{R}^7 \times \mathbb{R}^{7\times7} \times \mathbb{R}^9 \times \mathcal{U} \longrightarrow \mathbb{R}^3 \times \mathbb{R}^{3\times3}$ is a state-estimation map induced by the EKF, the vectors $\hat{q}_k$, $\hat{\omega}_k$, $\hat{\theta}_k$ and P_k denote the estimated attitude, angular velocity, inertia parameter and the associated estimation error covariance. The covariance of the inertia-parameter estimate is denoted by $P_{\theta,k}$ and is obtained as the principal submatrix of P_k corresponding to the inertia-parameter components $\hat{\theta}$ of $\hat{x}_k$. The inertia matrix estimate $\hat{J}_k$ is reconstructed from $\hat{\theta}_k$ using (8). The angular acceleration estimate is then obtained from the rotational dynamics evaluated at the current estimates,

$$\dot{\hat{\omega}}_k = \hat{J}_k^{-1}\Big(\tau_k - [\hat{\omega}_k]^\times \hat{J}_k \hat{\omega}_k\Big). \tag{12}$$

The regression model (9) shows that identifiability of the inertia parameter vector θ is governed by the information content of $\Psi(\hat{\omega}, \dot{\hat{\omega}})$. To quantify this information, we define the instantaneous information (Gram) matrix at discrete time instant t_k as

$$G_k \coloneqq \Psi_k^\top \Psi_k \in \mathbb{R}^{6\times6}, \qquad \Psi_k \coloneqq \Psi\big(\hat{\omega}(t_k), \dot{\hat{\omega}}(t_k)\big), \tag{13}$$

where $\{t_k\}_{k\in\mathbb{N}}$ denotes the discrete measurement time sequence. The matrix G_k characterizes the information contribution of the k-th measurement sample to inertia parameter identifiability. Given $\varepsilon > 0$ and over a finite window of N_T samples, the accumulated information matrix is defined as

$$\bar{G}(k_0; N_T) \coloneqq \sum_{k=k_0}^{k_0+N_T-1} G_k \; + \; \varepsilon I_6. \tag{14}$$

When $\bar{G}(k_0; N_T)$ is ill-conditioned or rank-deficient, the solution $\hat{\theta}$ to (9) becomes highly sensitive to noise and modeling errors, leading to poor identifiability and large estimation variance.

Objective and Problem Statement. The objective of this paper is to design torque inputs that maximize the information content of the regressor and thereby improve observability and identifiability of the inertia parameters θ. We use the minimum eigenvalue of the accumulated information matrix (13) as a scalar informativeness metric, namely,

$$\alpha(k_0; N_T) \coloneqq \lambda_{\min}\big(\bar{G}(k_0; N_T)\big), \qquad \mathcal{J}(k_0; N_T) \coloneqq -\log\big(\alpha + \varepsilon_\lambda\big), \tag{15}$$

where $\varepsilon_\lambda > 0$ prevents numerical singularities. Larger values of α (equivalently, smaller values of $\mathcal{J}$) correspond to a stronger excitation and improved parameter observability. The accumulated information matrix $\bar{G}(k_0; N_T)$ depends on the applied torque sequence $\{\tau_k\}_{k=k_0}^{k_0+N_T-1}$ through the regressor Ψ_k defined in (10), which in turn depends on the estimated state vectors $\hat{\omega}_k$ and $\dot{\hat{\omega}}_k$.

Given the current state estimate at time k_0, a finite design horizon N_T, and an admissible input set $\mathcal{U}$ encoding actuator and energy constraints (e.g., $|\tau_k|_\infty \le \tau_{\max}$ for all $k \in \mathbb{N}$), then, the proposed informative excitation design problem is formulated as the following constrained optimization problem:

$$\max_{\{\tau_k\}_{k=k_0}^{k_0+N_T-1} \subset \mathcal{U}} \lambda_{\min}\Big(\bar{G}(k_0; N_T)\Big). \tag{16}$$

Equivalently, using (15), the optimization in (16) can be written in minimization form as

$$\min_{\{\tau_k\}\in\mathcal{U}} \mathcal{J}(k_0; N_T) = \min_{\{\tau_k\}\in\mathcal{U}} -\log\Big(\lambda_{\min}(\bar{G}(k_0; N_T)) + \varepsilon_\lambda\Big). \tag{17}$$

In contrast to standard excitation strategies based on multisine or chirp signals [7], we restrict attention to sparse, safety-compatible episodic inputs and optimize their parameters online to directly maximize system observability.

EPISODIC EXPERIMENT DESIGN

Standard approaches for informative excitation often prescribe structured continuous inputs (e.g., multisine or chirp profiles) as in [7]. In our setting, continuous excitation is undesirable due to actuation constraints, power limitations, and safety considerations. Instead, we employ a sparse episodic excitation policy in which the commanded torque is identically zero for most of the time and is activated only during designated excitation episodes. Each episode i consists of a finite sequence of short torque bursts indexed by $j \in \{1, \dots, N\}$. Let $\mathcal{B}_{i,j}$ denote the time interval corresponding to burst j within episode i, each of duration T^i_{on}. During each burst interval $\mathcal{B}_{i,j}$, the torque is held constant and equal to

$$\tau^i = a^i\big(\hat{\omega}^i \times v^i\big) + b^i v^i, \qquad \tau^i \in \mathcal{U} \coloneqq \{\tau \in \mathbb{R}^3 : |\tau|_\infty \le \tau_{\max}\}, \tag{18}$$

while outside the union of burst intervals $\bigcup_j \mathcal{B}_{i,j}$ the torque is set to zero.

The burst parameters $(a^i, b^i, T^i_{\text{on}}, v^i)$ are selected at the beginning of episode i using the current state estimates obtained from the EKF recursion (11). In particular, $v^i \in \mathbb{R}^3$ is chosen to align with the dominant eigen-direction of the inertia-parameter covariance P^i_θ, namely the eigenvector corresponding to the largest eigenvalue of P^i_θ. The scalars a^i, b^i, and T^i_{on} are design variables subject to actuator, power, and mission constraints. The time gaps between consecutive bursts within episode i are selected from a prescribed admissible interval $[T_{g,\min}, T_{g,\max}]$, where $T_{g,\max} > T_{g,\min} > 0$ are known design constants, i.e., the inter-burst dwell time $\Delta T_{i,j}$ satisfies

$$T_{g,\min} \le \Delta T_{i,j} \le T_{g,\max}.$$

These bounds enforce dwell-time, thermal, and energy constraints and are treated as hard operational limits. Consequently, the resulting control signal τ is piecewise constant and sparse in time.

The structure of (18) provides targeted excitation both transverse to the current rotational motion and along the dominant uncertainty direction of the inertia parameters. The cross term $\hat{\omega}^i \times v^i$ is orthogonal to $\hat{\omega}^i$ and induces multi-axis rotational motion, thereby increasing variation in the regressor $\Psi(\hat{\omega}^i, \dot{\hat{\omega}}^i)$ in (10). The second term $b^i v^i$ injects excitation aligned with the dominant eigenvector of P^i_θ, targeting the largest component of inertia-parameter uncertainty.

For a given episode parameter vector $\xi^i := \big(a^i, b^i, T^i_{\text{on}}\big) \in \mathcal{X}$, where $\mathcal{X} \subset \mathbb{R}^3$ denotes the admissible set of episode parameters, we apply (18) over the burst intervals $\mathcal{B}_{i,j}$ of episode i, propagate the rigid-body dynamics together with the estimator (11), construct Ψ_k and $\bar{G}(k_0; N_T)$ via (13)–(14), and evaluate the scalar cost $\mathcal{J}$ defined in (15). This procedure defines a black-box mapping

$$\mathcal{J}^i = \mathcal{J}(\xi^i), \tag{19}$$

which is generally nonconvex and nondifferentiable due to actuator saturation, episodic scheduling, and the recursive nature of the state estimation.

Bayesian optimization formulation. Given actuator limits $\tau_{\max} > 0$, burst-duration bounds $0 < T_{\min} \le T_{\max}$ defining the admissible interval $T_{\text{on}} \in [T_{\min}, T_{\max}]$, and the torque structure (18), we seek episode parameters that minimize the informativeness cost:

$$\xi^\star \in \arg\min_{\xi \in \mathcal{X}} \mathcal{J}(\xi), \qquad \mathcal{X} := \Big\{(a, b, T_{\text{on}}) \in \mathbb{R}^3 : |\tau|_\infty \le \tau_{\max},\ T_{\min} \le T_{\text{on}} \le T_{\max}\Big\} \tag{20}$$

where $\mathcal{X}$ denotes the admissible design space of episode parameters satisfying actuator and duration constraints, and the objective function $\mathcal{J}$ represents the scalar informativeness cost defined in (15). Bayesian optimization (BO) is employed to solve (20) by constructing a probabilistic surrogate model for the unknown objective function $\mathcal{J}$. The surrogate is chosen as a Gaussian process (GP).

To solve (20), we perform a sequence of experiments (episodes). At episode j, the parameter ξ^j is applied and the corresponding objective value $\mathcal{J}(\xi^j)$ is observed. After i completed episodes, the collected dataset is $\mathcal{D}_i := \{(\xi^j, \mathcal{J}(\xi^j))\}_{j=1}^i$ where $\mathcal{J}(\xi^j)$ denotes the observed objective value corresponding to the parameter choice ξ^j. Conditioned on the dataset $\mathcal{D}_i$, the posterior distribution of the function $\mathcal{J}$ induces

$$\mathcal{J}(\xi) \mid \mathcal{D}_i \sim \mathcal{N}\big(\mu_i(\xi),\, \sigma_i^2(\xi)\big) \tag{21}$$

for each $\xi \in \mathcal{X}$, a predictive normal distribution with mean $\mu_i(\xi)$ and variance $\sigma_i^2(\xi)$.

Based on the posterior surrogate model of the objective, we determine where to evaluate the function next in order to reduce the performance cost. The next episode parameter ξ^{i+1} is selected by maximizing an acquisition function $\mathcal{A}_i : \mathcal{X} \to \mathbb{R}$ that balances exploitation (small $\mu_i(\xi)$) and exploration (large $\sigma_i(\xi)$):

$$\xi^{i+1} \in \arg\max_{\xi \in \mathcal{X}} \mathcal{A}_i(\xi). \tag{22}$$

In this work, we employ the probability of improvement (PI) acquisition function, which selects query points that have high posterior probability of improving upon the best objective value observed so far:

$$\mathcal{A}_i(\xi) = \Phi\big(Z_i(\xi)\big), \qquad Z_i(\xi) \coloneqq \frac{\mathcal{J}_{\min}^i - \kappa - \mu_i(\xi)}{\sigma_i(\xi)}, \qquad \mathcal{J}_{\min}^i := \min_{1 \le j \le i} \mathcal{J}(\xi^j), \tag{23}$$

where Φ denotes the standard normal cumulative distribution function (CDF), and $\kappa > 0$ is a desired exploration parameter. Since the standard normal CDF Φ is monotonically increasing, maximizing $\mathcal{A}_i(\xi)$ is equivalent to maximizing $Z_i(\xi)$. This yields the following two cases:

- $Z_i(\xi) > 0$ (Exploitation): Here $\mu_i(\xi) < \mathcal{J}_{\min}^i - \kappa$, so the posterior mean already predicts an improvement over the current best $\mathcal{J}_{\min}^i$. Moreover,

  $$\frac{\partial Z_i(\xi)}{\partial \mu_i(\xi)} = -\frac{1}{\sigma_i(\xi)} < 0 \quad \text{for all } \xi \neq \xi^j, \quad j \in \{1, 2, \cdots, i\}$$

  so decreasing $\mu_i(\xi)$ always increases $Z_i(\xi)$, and hence increases $\Phi(Z_i(\xi))$, regardless of the value of $\sigma_i(\xi)$. Therefore, PI favors sampling in regions with low predicted objective value, corresponding to exploitation.

- $Z_i(\xi) < 0$ (Exploration): Here $Z_i(\xi) = \frac{\mathcal{J}_{\min}^i - \kappa - \mu_i(\xi)}{\sigma_i(\xi)} < 0$ which implies $\mu_i(\xi) > \mathcal{J}_{\min}^i - \kappa$, i.e., the posterior mean does not predict improvement. Differentiating $Z_i(\xi)$ with respect to $\sigma_i(\xi)$ gives

  $$\frac{\partial Z_i(\xi)}{\partial \sigma_i(\xi)} = -\frac{Z_i(\xi)}{\sigma_i^2(\xi)} > 0 \quad \text{for all } \xi \neq \xi^j, \quad j \in \{1, 2, \cdots, i\}$$

Hence, increasing $\sigma_i(\xi)$ strictly increases $Z_i(\xi)$, and therefore increases $\Phi(Z_i(\xi))$. Consequently, when the predicted mean does not indicate improvement, PI favors points with larger posterior uncertainty, promoting exploration of less-certain regions of $\mathcal{X}$.

Algorithm 1 BO–EKF Episodic Excitation (Evaluation Budget M)

Require: Evaluation budget $M \in \mathbb{N}$; feasible set $\mathcal{X} \subset \mathbb{R}^3$ for $\xi = (a, b, T_{\mathrm{on}})$ satisfying $|\tau|_\infty \leq \tau_{\max}$, $T_{\mathrm{on}} \in [T_{\min}, T_{\max}]$; window length $N_T \in \mathbb{N}$; regularization $\varepsilon_\lambda > 0$; PI exploration parameter $\kappa > 0$.

Initialize EKF state and covariance according to (11): $(\hat{x}_{k_0}, P_{k_0})$
Initialize BO dataset $\mathcal{D}_0 = \emptyset$, episode counter $i \leftarrow 0$
while $i < M$ **do**
  $i \leftarrow i + 1$
  Extract P_{θ,k_0} as the principal submatrix of P_{k_0} corresponding to $\hat{\theta}_{k_0}$ in $\hat{x}_{k_0}$
  Select excitation direction v^i aligned with the dominant eigenvector of P_{θ,k_0}
  if $i = 1$ **then**
    Choose initial design $\xi^1 \in \mathcal{X}$ (e.g., space-filling)
  else
    Fit GP surrogate on $\mathcal{D}_{i-1}$
    $\xi^i \in \arg\max_{\xi \in \mathcal{X}} \mathcal{A}_{i-1}(\xi)$ defined in (23)
  end if
  Apply episodic burst command via (18). For each burst interval $\mathcal{B}_{i,j} \subset \{k_0, \ldots, k_0 + N_T - 1\}$ of duration T^i_{on},

$$\tau^i_k = a^i(\hat{\omega}_k \times v^i_k) + b^i v^i_k, \qquad |\tau^i_k|_\infty \leq \tau_{\max}, \qquad \forall k \in \mathcal{B}_{i,j}.$$

  for $k = k_0, \ldots, k_0 + N_T - 1$ **do**
    Propagate dynamics (1) with input τ_k
    Update EKF estimates $(\hat{x}_{k+1}, P_{k+1})$ via (11)
    Construct Ψ_k via (10) and compute G_k via (13)
  end for
  Accumulate $\bar{G}(k_0; N_T)$ via (14)
  Compute $\mathcal{J}(\xi^i) = -\log\big(\lambda_{\min}(\bar{G}(k_0; N_T)) + \varepsilon_\lambda\big)$
  $\mathcal{D}_i \leftarrow \mathcal{D}_{i-1} \cup \{(\xi^i, \mathcal{J}(\xi^i))\}$
  Update window start index $k_0 \leftarrow k_0 + N_T$
end while
Return: $\xi^\star \in \arg\min_{(\xi,\mathcal{J}) \in \mathcal{D}_M} \mathcal{J}$

The selected parameter $\xi^{i+1} = (a^{i+1}, b^{i+1}, T^{i+1}_{\mathrm{on}})$ determines the next pulse via (18). The system is excited for duration T^{i+1}_{on}, and the corresponding objective value $\mathcal{J}(\xi^{i+1})$ is computed from (15) and appended to the dataset $\mathcal{D}_i$ to update the GP posterior. During each episode, the EKF recursion (11) provides updated state estimates $(\hat{\omega}_k, \dot{\hat{\omega}}_k)$, which determine the regressor Ψ_k and the instantaneous Gramian G_k in (13). These quantities accumulate over the window to form $\bar{G}(k_0; N_T)$ in (14), from which the informativeness cost $\mathcal{J}$ in (15) is evaluated.

The BO procedure is repeated for a maximum of M objective evaluations, where M denotes a

prescribed evaluation budget reflecting computational and operational constraints. Iterating the BO steps (20)–(23) over $i = 1, \ldots, M$ generates a sequence of sparse torque bursts whose parameters are selected to increase $\lambda_{\min}(\bar{G}(k_0; N_T))$, thereby improving the observability and identifiability of the inertia parameters θ under practical actuation constraints.

The resulting BO–EKF episodic excitation strategy is summarized in Algorithm 1. To ensure robust parameter adaptation, the inertia-parameter process noise is initialized with a nonzero covariance and gradually reduced over time [6]. This prevents premature convergence to inaccurate inertia values when the initial parameter guess is poor, while allowing the estimates to stabilize as informative excitation data accumulate.

SIMULATION RESULTS

We consider a tumbling rigid body with initial angular velocity

$$\omega_0 = \begin{bmatrix} 0.5 & 10 & 0.45 \end{bmatrix}^\top \text{ rad/s},$$

and initial attitude corresponding to Euler angles (degrees) $[\phi, \theta, \psi] = [0.92,\ 0,\ 0.5]$. The true inertia matrix is

$$J_{\text{true}} = \begin{bmatrix} 100 & 6 & -4 \\ 6 & 80 & 3 \\ -4 & 3 & 70 \end{bmatrix},$$

which is assumed unknown to the estimator. The simulation integration time step is $\Delta t = 0.01\,\text{s}$. Measurements (3)–(5) are available intermittently with sampling intervals uniformly drawn from $[0.10,\ 0.20]\,\text{s}$ and are corrupted by additive noise with standard deviations $\sigma_\omega = 3 \times 10^{-2}\,\text{rad/s}$ for angular velocity and $\sigma_q = 10^{-2}$ for quaternion components (followed by quaternion normalization). The EKF is initialized with $\theta_0 = \begin{bmatrix} 85 & 135 & 210 & 0 & 0 & 0 \end{bmatrix}^\top$.

To excite the rotational dynamics and improve identifiability, the torque pulse is parameterized according to (18). The pulse duration satisfies $T_{\text{on}} \in [T_{\min}, T_{\max}] = [0.05,\ 1.0]\ \text{s}$. Between consecutive bursts, coast intervals are inserted. The inter-burst dwell time is selected from the prescribed admissible interval $\Delta T_j \in [T_{\text{g,min}}, T_{g,\max}] = [5,\ 20]\ \text{s}$, during which the torque is identically zero. These bounds enforce actuation, thermal, and energy constraints. The number of excitation pulses applied per episode is $N = 4$.

The BO decision vector is $\xi = (a, b, T_{\text{on}}) \in \mathcal{X}$. The search bounds for BO parameters are derived from actuator saturation and the initial angular-rate magnitude. Define $\omega_{\text{ref}} = \max(|\omega_0|,\ 10^{-3})$. Then $a \in \left[0,\ \frac{\tau_{\max}}{\omega_{\text{ref}}}\right]$, $b \in [0,\ \tau_{\max}]$. For this example, we select $\tau_{\max} = 5$. BO is run for $M = 100$ objective evaluations. For each candidate ξ, the objective returned to BO is the scalar cost $\mathcal{J}(\xi) = -\log\big(\alpha(\xi) + \varepsilon_\lambda\big)$, constructed from a sliding-window information matrix. At each time index k, the instantaneous information contribution is $G_k = \Psi_k^\top \Psi_k \in \mathbb{R}^{6\times 6}$, where $\Psi_k = \Psi(\hat{\omega}_k, \dot{\hat{\omega}}_k)$. Over a window of length $N_T = \max\big(5,\ \lfloor \frac{T_{\text{on}}}{\Delta t} \rceil\big)$, starting at index s, we define

$$\bar{G}(s) = \sum_{k=s}^{s+N_T-1} G_k + \varepsilon I_6, \qquad \lambda_{\min}(s) = \lambda_{\min}\big(\bar{G}(s)\big).$$

Let $\mathcal{S}$ denote the set of valid window indices. We aggregate the worst-case excitation using the soft-min operator

$$\alpha(\xi) = -\frac{1}{\beta} \log\left(\frac{1}{|\mathcal{S}|} \sum_{s\in\mathcal{S}} \exp\big(-\beta\, \lambda_{\min}(s)\big) \right). \tag{24}$$

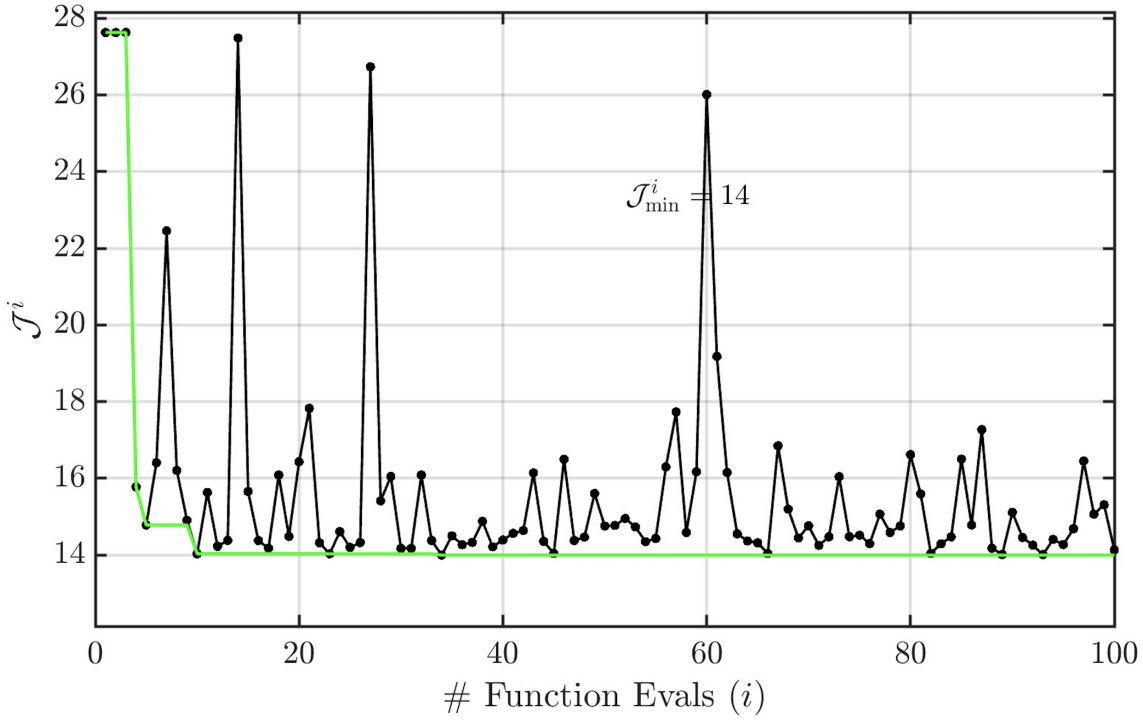

Figure 1 BO cost with repeated function evaluations.

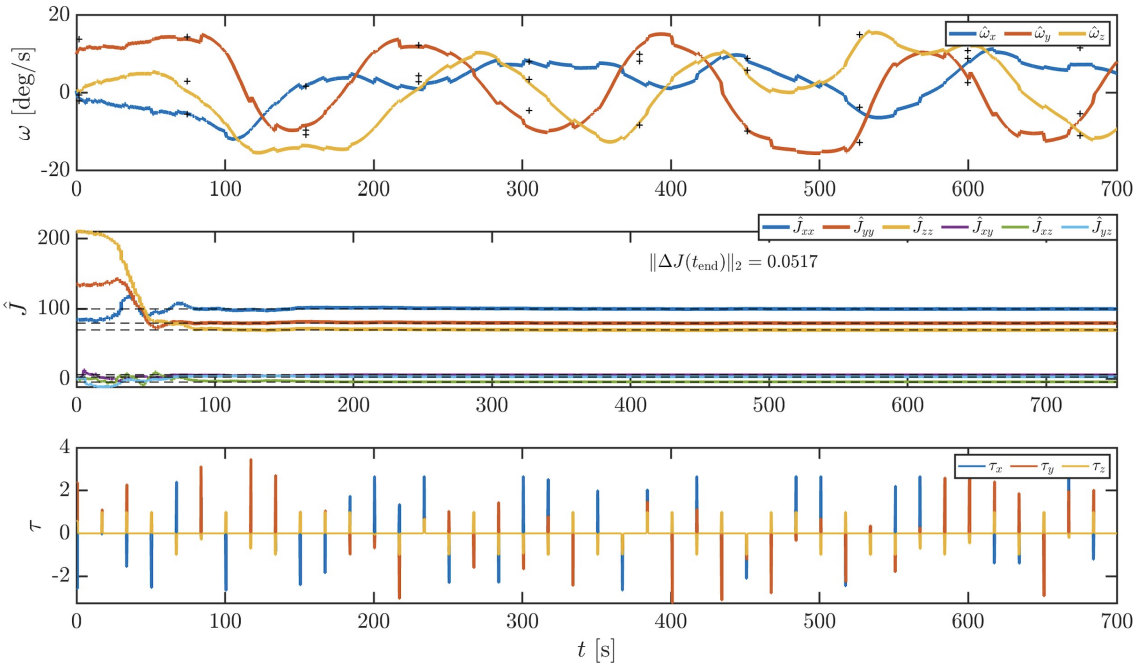

Figure 2 EKF estimates of $\hat{\omega}$, and $\hat{J}$ under the designed torque pulses τ.

Minimizing $\mathcal{J}(\xi)$ therefore promotes trajectories that maximize the (soft) worst-case information content and improve inertia identifiability.

Figure 1 depicts the BO objective values $\mathcal{J}(\xi^i)$ obtained at each function evaluation i (black curve). The green curve shows the running best value $\mathcal{J}^i_{\min} := \min_{1\le j\le i} \mathcal{J}(\xi^j)$. The initial evaluations rapidly reduce the cost, after which the running best stabilizes, indicating that BO has reached a near-optimal region of the search space. The occasional spikes in $\mathcal{J}(\xi^i)$ reflect stochasticity in the objective evaluation (e.g., intermittent noisy measurements and estimator variability), while the monotone decrease of $\mathcal{J}^i_{\min}$ confirms progress toward improved excitation design. With an evaluation budget of $M = 100$, as seen from Figure 1, BO converged to the parameter vector $\xi^\star = (3.76,\ 3.87,\ 0.99\,)$, achieving a minimum observed objective value $\mathcal{J}(\xi^\star) = 14$.

Finally, the optimized parameter vector $\xi^\star$ is validated in a long-horizon simulation of duration $T_{\text{val}} = 1000$ s. The torque is applied as a sequence of constant-amplitude bursts of duration $T_{\text{on}} = 0.99$ s, separated by coast intervals of approximately 90–100 s, during which $\tau(t) = 0$. Figure 2 reports the corresponding time histories of the EKF angular-velocity estimate $\hat{\omega}$, the inertia estimates $\hat{\theta}$, and the applied torque. The EKF is tuned with an initial covariance $P_0 = \text{blkdiag}\big(10^{-2}I_7,\ 15000\,I_6\big)$, measurement-noise covariance $R = \text{blkdiag}\big((3{\times}10^{-2})^2 I_3,\ (10^{-2})^2 I_4\big)$, and a process-noise model $Q_k = \alpha_k \text{blkdiag}(\bar{Q},\ Q_\theta)$ with $\bar{Q} = \text{blkdiag}(10^{-3}I_3,\ 10^{-4}I_4)$ and $Q_\theta = 5 \times 10^{-2} I_6$. To avoid premature locking to an inaccurate inertia model, the scalar factor α_k is scheduled to decrease over time [6],

$$\alpha_k = 0.999\,\gamma^k + q_{\min}, \qquad \gamma = 0.99, \qquad q_{\min} > 0,$$

so that the filter initially allows faster parameter adaptation and subsequently reduces the parameter drift as informative data accumulate. The convergence of $\hat{\omega}$ ensures accurate reconstruction of the regressor Ψ_k, which in turn enables the windowed information matrix to accumulate sufficient excitation for parameter identifiability. The rapid reduction and subsequent stabilization of $\hat{J}$ confirms that the BO-designed bursts provide persistently informative trajectories.

If the input is insufficiently exciting, the information content of the measurements is inadequate, and the inertia parameters are unidentifiable, in which case the inertia estimates do not converge to their true values, as seen from Figure 3.

CONCLUSION

This paper proposed a Bayesian optimization–based framework for designing informative, episodic torque inputs to estimate the inertial parameters of a tumbling, non-cooperative rigid body. By exploiting the linear-in-parameters structure of the rotational dynamics, we introduced an informativeness metric based on a windowed information (Gram) matrix and formulated excitation design as a constrained optimization problem. Sparse torque pulses were parameterized in a safety-compatible form and selected online using BO, while an EKF provided joint state and inertia estimates. Numerical results show that the resulting BO–EKF loop improves observability and enhances inertia identifiability under practical actuation limits. However, informative torque design alone is not sufficient to guarantee fast or robust convergence when the true inertia is unknown. A key objective is to minimize the number of excitation episodes needed for the estimator to converge with high confidence to a small neighborhood of the true parameters. This motivates extending the BO design beyond (A, B, T_{on}) to include filter-related parameters and additional cost terms that penalize posterior inertia uncertainty, enabling a joint input–estimator co-design that trades off informativeness, estimator confidence, and actuation effort.

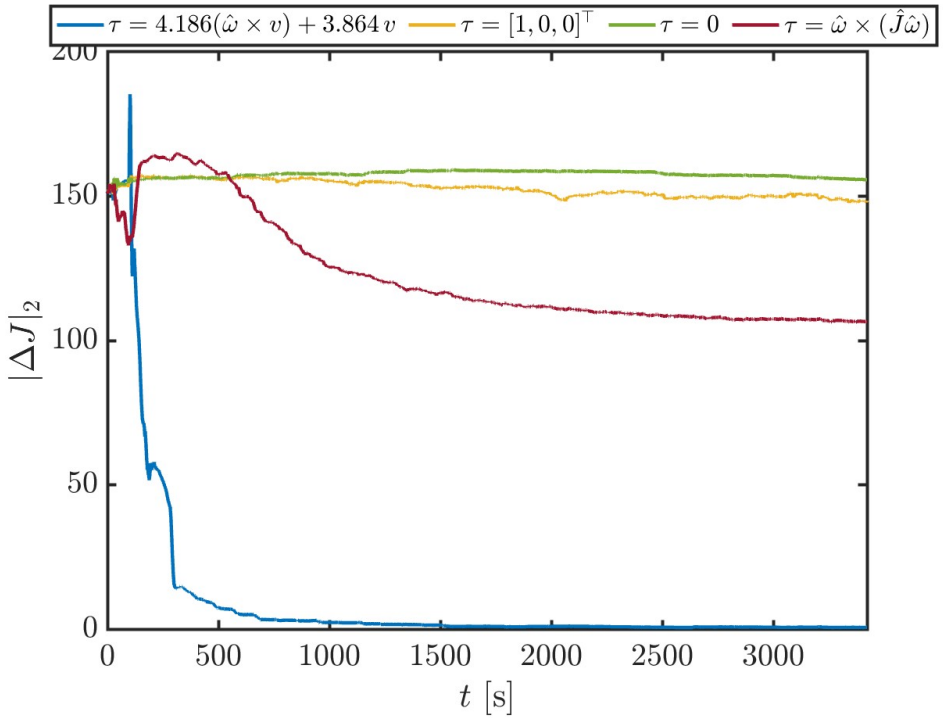

Figure 3 Inertia estimates under informative and noninformative excitation inputs.

ACKNOWLEDGMENTS

Regarding this work, we would like to express our sincere gratitude to Dr. Yasar Yanik and Prof. Daniele Venturi from the Department of Applied Mathematics at the University of California, Santa Cruz (UCSC) for their insightful discussions. We also sincerely acknowledge the support of our funding sponsors. This research was partially supported by NSF Grants Nos. CNS-2039054 and CNS-2111688; AFOSR Grants Nos. FA9550-19-1-0169, FA9550-20-1-0238, FA9550-23-1-0145, FA9550-23-1-0313, FA9550-23-1-0678; AFRL Grants Nos. FA8651-22-1-0017, FA8651-23-1-0004; ARO Grant No. W911NF-20-1-0253; and DoD Grant No. W911NF-23-1-0158.

REFERENCES

[1] A. Nocerino, R. Opromolla, G. Fasano, and M. Grassi, "LIDAR-Based Multi-Step Approach for Relative State and Inertia Parameters Determination of an Uncooperative Target," *Acta Astronautica*, Vol. 181, 2021, pp. 662–678.

[2] H. Basu, M. C. Fernandez, R. G. Sanfelice, and I. Kolmanovsky, "Hybrid Model Predictive Control Approach for Spacecraft Proximity Maneuvering and Docking Accounting for Collisions," *Proceedings of the 2025 American Control Conference (ACC)*, Denver, CO, USA, 2025.

[3] H. Basu, P. Jirwankar, R. G. Sanfelice, M. Castroviejo-Fernandez, and I. Kolmanovsky, "A Hybrid Model Predictive Control Framework for Docking & Stabilization of Composite Rigid Spacecraft Dynamics," *AIAA SciTech 2026 Forum*, Orlando, FL, USA, 2026. AIAA 2026-2855.

[4] S. H.-Garcia, M. Kapteyn, K. E. Willcox, M. Tezzele, M. C.-Fernandez, T. Kim, M. Ambrosino, I. Kolmanovsky, H. Basu, P. Jirwankar, and R. Sanfelice, "Digital-Twin-Enabled Multi-Spacecraft On-Orbit Operations," *AIAA SCITECH 2025 Forum*, Orlando, FL, AIAA, 2025.

[5] H. Benninghoff and T. Boge, "Rendezvous Involving a Non-Cooperative, Tumbling Target: Estimation of Moments of Inertia and Center of Mass of an Unknown Target," *Proceedings of the 25th International Symposium on Space Flight Dynamics*, Munich, Germany, German Aerospace Center (DLR), 2015.

[6] Q. Meng, J. Liang, and O. Ma, "Identification of All the Inertial Parameters of a Non-Cooperative Object in Orbit," *Aerospace Science and Technology*, Vol. 91, 2019, pp. 571–582.

[7] B. Boyacioglu, D. Sandursky, and K. A. Morgansen, "Nonlinear Estimation of Rigid Body Inertial Parameters," *AIAA SciTech 2023 Forum*, National Harbor, MD & Online, 2023. AIAA 2023-1809.

[8] N. B. Andrews and K. A. Morgansen, "Modeling, Observability, and Inertial Parameter Estimation of a Planar Multi-Link System with Thrusters*," *Proceedings of the 2025 IEEE 64th Conference on Decision and Control (CDC)*, Rio de Janeiro, Brazil, 2025, pp. 3244–3251.