Check for updates

Article

The International Journal of Robotics Research
2025, Vol. 44(7) 1117–1137
© The Author(s) 2024
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/02783649241300167
journals.sagepub.com/home/ijr
Sage

Magnetic manipulation of unknown and complex conductive nonmagnetic objects with application in the remediation of space debris

Griffin F Tabor[1,2], Lan N Pham[1,3], Jake J Abbott[1,3] and Tucker Hermans[1,2,4]

Abstract

This article extends recent work in magnetic manipulation of conductive, nonmagnetic objects using rotating magnetic dipole fields. Eddy-current-based manipulation provides a contact-free way to manipulate metallic objects. We are particularly motivated by the large amount of aluminum in space debris. We previously demonstrated dexterous manipulation of solid spheres with all object parameters known a priori. This work expands the previous model, which contained three discrete modes, to a continuous model that covers all possible relative positions of the manipulated spherical object with respect to the magnetic field source. We further leverage this new model to examine manipulation of spherical objects with unknown physical parameters by applying techniques from the online-optimization and adaptive-control literature. Our experimental results validate our new dynamics model, showing that we get improved performance compared to the previously proposed model, while also solving a simpler optimization problem for control. We further demonstrate the first physical magnetic manipulation of aluminum spheres, as previous controllers were only physically validated on copper spheres. We show that our adaptive control framework can quickly acquire useful object parameters when weakly initialized. Finally, we demonstrate that the spherical-object model can be used as an approximate model for adaptive control of nonspherical objects by performing magnetic manipulation of a variety of objects for which a spherical model is not an obvious approximation.

Keywords

Space robotics and automation, dexterous manipulation, model learning for control, robust/adaptive control, optimization and optimal control

Received 21 June 2023; Revised 7 August 2024; Accepted 5 October 2024
Senior Editor: Manuel Catalano
Associate Editor: Hamid Sadeghian

1. Introduction

There have been significant advances on the topic of magnetic manipulation over the past two decades, the vast majority coming from the robotics community (Abbott et al., 2020). Researchers have developed methods to use sets of stationary electromagnets, or robot-controlled permanent magnets, to dexterously manipulate both tethered and untethered devices without any direct physical contact. However, the objects being manipulated have typically been primarily composed of ferromagnetic material (soft- or permanent-magnet). Traditional magnetic methods are severely limited in what they can manipulate, as they rely on the object's ferromagnetic properties, which are only present in a limited set of materials.

Many engineering materials, although not ferromagnetic, are electrically conductive, including aluminum, titanium, copper, and some stainless steels. It has long been known that when conductive objects are exposed to time-varying magnetic fields (as opposed to static magnetic fields), a flow of electrons known as eddy currents is induced in the material (Hertz et al., 1896). These eddy currents then interact with the applied magnetic field, inducing forces and torques on the conductive object. A common commercial application of this phenomenon is material separation in metal recycling plants (Smith et al., 2019).

[1]Robotics Center, University of Utah, Salt Lake City, UT, USA
[2]Kahlert School of Computing, University of Utah, Salt Lake City, UT, USA
[3]Department of Mechanical Engineering, University of Utah, Salt Lake City, UT, USA
[4]NVIDIA, Seattle, WA, USA

Corresponding author:
Griffin Tabor, The University of Utah, 201 Presidents' Cir, Salt Lake City, UT 84112, USA.
Email: Griffin.Tabor@utah.edu

The use of eddy-current-induced forces and/or torques for applications in space is a particularly promising and active area of research, as it enables contactless manipulation of objects. Traditional magnetic manipulation methods fall short because a large quantity of objects currently in space are composed of aluminum (Opiela, 2009), which is conductive but not ferromagnetic. Eddy-current-induced forces have been proposed as a method of traversing the exterior of the International Space Station (Reinhardt and Peck, 2016; Wilson et al., 2023; Wilson and Peck, 2020, 2021). We are particularly interested in contributing solutions to the problem of space debris (Liou and Johnson, 2016; Ma et al., 2023; Mark and Kamath, 2019; National Aeronautics and Space Administration, 2021; Papadopoulos et al., 2021; Shan et al., 2016). A study found that "even if no future launches occurred, collisions between existing satellites would increase the 10-cm and larger debris population faster than atmospheric drag would remove objects" (National Aeronautics and Space Administration, 2021). This will eventually lead to a phenomenon known as the Kessler Syndrome (Kessler et al., 2010), in which Earth's orbit becomes clogged with debris due to cascading collisions between objects, making it unusable. As such, there is a dire need for remediation strategies to remove or repair resident space objects in order to protect the fast-growing number of satellites that the world's population has grown to rely on (National Aeronautics and Space Administration, 2021). The majority of prior efforts using eddy currents have focused on detumbling satellites using static magnetic fields, both uniform and dipole (Liu et al., 2019; Nurge et al., 2018; Ortiz Gómez and Walker, 2015; Sugai et al., 2013). Eddy-current-based approaches do not have the risk of potentially destructive collision that is inherent to all other robotic approaches to detumbling, such as serial-link manipulators (Aghili, 2020; Vijayan et al., 2022), bio-inspired soft robots (Carambia et al., 2023; Frazelle et al., 2021), and tethered nets (Boonrath and Botta, 2024; Zeng et al., 2024). However, prior eddy-current-based efforts are not able to maintain the relative position of the object as it is detumbled, since torques and forces are not controlled independently, and objects tend to get pushed away as they are detumbled.

We recently showed that full six-degree-of-freedom (6-DOF) dexterous manipulation of conductive, nonmagnetic objects (specifically spheres) utilizing eddy currents is, in fact, possible (Pham et al., 2021). The method assumes that the object is surrounded (to some degree) by multiple static electromagnet field sources capable of generating continuously rotating magnetic dipole fields about arbitrary axes. It is noteworthy that we were able to achieve full 6-DOF manipulation of spherical objects, as 6-DOF manipulation of ferromagnetic objects has been shown to only be possible for complex geometries (Diller et al., 2016), with 5-DOF typical of most simple geometries, and only 3-DOF achievable for soft-magnetic spheres (Abbott et al., 2020). We have also shown that it is possible to create a tractor-beam-like effect to pull an object into a central manipulation workspace (Dalton et al., 2022).

The forces and torques induced on conductive, nonmagnetic spheres are small compared to those due to ferromagnetism, but they have the potential to be useful for applications in the microgravity environment of space, in which manipulation over relatively long time scales is feasible. We are particularly interested in application of detumbling a resident space object by inducing torque antiparallel with the object's angular velocity while maintaining its relative position via closed-loop position control, such that it can be subsequently manipulated by more traditional robotic means for maintenance or de-orbiting. We recently showed that a rotating dipole field is substantially more effective at braking a tumbling object than is a static dipole field (Allen et al., 2024). Surrounding the manipulated object with field sources could be achieved by a single robotic spacecraft with multiple arms, each equipped with a field source as its end-effector (Sugai et al., 2013). Alternatively, it could be achieved by multiple simple spacecraft that each contains a single field source; the field sources could even be used simultaneously to maintain the relative positions of the spacecraft as they perform the manipulation operation, to reduce the reliance on consumables (Abbott et al., 2017). The initial rendezvous with a resident space object is a solved problem (e.g., Volpe and Circi, 2019).

In this paper, which serves as an extension of our recent conference paper Tabor et al. (2022), we improve manipulation performance and vastly expand the class of objects that we can manipulate (Figure 1). We make five contributions

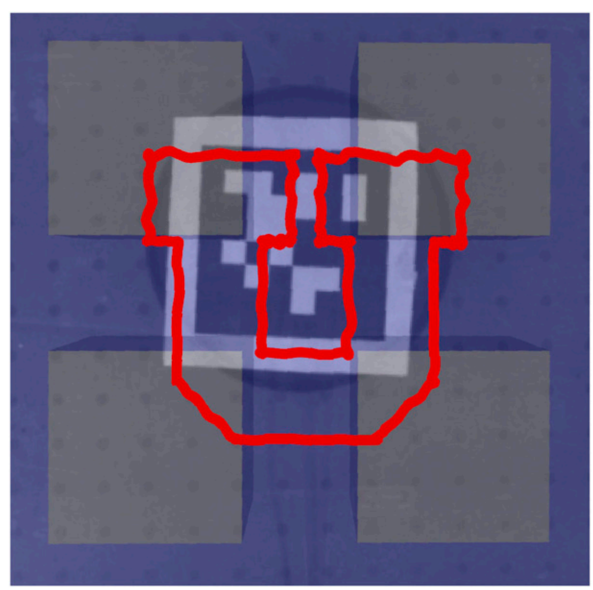

Figure 1. Example trajectory—the University of Utah "block U"—produced using the proposed adaptive controller (specifically, our inverse-dynamics controller with a prior/regularization term during system identification). The red line represents position over time (72 min) traced by the center of a copper cuboid, while holding a fixed orientation. In this planar simulation of microgravity, the conductive nonmagnetic object is placed in a plastic raft that floats with 3-DOF mobility on the surface of water, with four magnetic-dipole field sources placed beneath the water tank. The pose of the raft is tracked with a camera using a fiducial marker. The positions of the cube-shaped electromagnetic field sources are rendered in the image at true scale and with perspective; they are obstructed in the actual video.

relative to our earlier work Pham et al. (2021), and two contributions relative to Tabor et al. (2022), all motivated by manipulation of space debris:

1. In Pham et al. (2021), we modeled eddy-current-induced force–torque at three distinct canonical positions of a nonmagnetic, conductive sphere with respect to a rotating magnetic dipole: along the rotation axis of the rotating dipole (parallel and antiparallel) and orthogonal to the axis of rotation. We then used the canonical-position model in a manipulation framework, which forced the conductive sphere to be cast into one of the three canonical positions during actuation. Although this method was sufficient to enable 6-DOF manipulation (provided there were enough dipole-field sources), it unnecessarily constrained the dipole rotation axes that could be used, making the results suboptimal. Here, we provide a single, continuous model of force–torque across all positions of the conductive sphere relative to the rotating dipole. We then modify the manipulation framework to use the new continuous-position model, and show improved tracking performance to that of Pham et al. (2021), while solving a simpler optimization problem.
2. In Pham et al. (2021), we used a controller based on desired force–torque wrenches. Here, we additionally propose an inverse-dynamics controller that accounts for object mass to achieve desired acceleration, which results in improved performance.
3. In Pham et al. (2021), we assumed that object dynamics models were known. As a step toward manipulation of unknown space debris, we propose an approach to manipulate spheres with unknown physical parameters (i.e., radius and conductivity) through the use of adaptive control. We leverage the recently proposed view of adaptive control as online optimization (Ratliff et al., 2016). This enables us to more closely tie adaptive control to classical system identification (Atkeson et al., 1986), while also making use of exciting advances in online optimization (Hazan, 2016) such as solvers that are robust to noise while also handling constraints and injecting prior knowledge of system parameters.
4. In Pham et al. (2021), we only physically manipulated copper spheres. Here, we use the adaptive controller to enable the manipulation of aluminum spheres; aluminum is the most commonly used material in engineered space objects, but it is less electrically conductive than copper, and consequently has smaller induced forces and torques.
5. In Pham et al. (2021), we developed a model for induced force–torque on nonmagnetic, conductive spheres; this model was hypothesized to be a useful approximation for other geometries, but the model was only used to manipulate spheres. Here, we demonstrate that our adaptive controller can be used to manipulate nonspherical, nonmagnetic, conductive objects by locally approximating the dynamics using the model for spheres. We provide extensive testing and analysis to demonstrate our adaptive controller's ability to generalize to a variety of unknown objects.
6. In Tabor et al. (2022), we only explicitly considered compact solid objects such as cylinders and cuboids, for which a solid-sphere model is obviously a valid first-order approximation. Here, we consider a variety of additional interesting objects that are likely more representative of space debris, including: thin-walled objects; composite objects with elements that are either electrically connected or electrically isolated from each other; and objects containing small amounts of ferromagnetic material.
7. In Tabor et al. (2022), we only explicitly considered four field sources in the particular arrangement of our experimental testbed. Here, we consider how the number of field sources surrounding a workspace affects manipulability, and find that as few as two field sources enables full 6-DOF manipulation of objects in the central workspace. This result is particularly promising for the deployment of our method in space applications, in which fewer field sources are more practical.

In addition, Tabor et al. (2022) had some software bugs: the adaptive-control parameter optimization compared measured and expected accelerations in different frames, and the desired rotational velocity in the controller was always set to 0. These software bugs heavily reduced the performance of our controller, so the results in this article are markedly better.

The paper structure continues as follows. We review the existing state-of-the-art force–torque wrench model in Section 2.1, and we then describe our proposed continuous-in-position wrench model in Section 2.2. We discuss our manipulation framework in Section 3. We then detail our approach to object parameter optimization, both as system identification and adaptive control, in Section 4. We describe our experimental design and results in Section 5. We discuss remaining open problems in Section 6, before concluding in Section 7.

2. Model of induced force–torque wrench on a solid sphere

2.1. Review of prior model

In this section, we summarize the wrench (i.e., force and torque) model of Pham et al. (2021). However, we recast the model into spherical coordinates, which we have found enables an elegant way to extend the existing model to new, previously unmodeled locations.

The magnetic field source can be abstracted as a point dipole $\boldsymbol{m}$ (units $\mathrm{A{\cdot}m^2}$, with direction pointing from the south pole to the north pole) at its center of mass.

The center of the nonmagnetic conductive sphere is then described by a relative displacement vector $\boldsymbol{\rho}$, with both vectors expressed in a common frame of reference. Figure 2 shows the new spherical coordinate system, where any given position can be described by three coordinates with respect to the rotating magnetic dipole: a distance $\rho = \|\boldsymbol{\rho}\|$, a polar angle θ measured from the dipole's rotation vector $\boldsymbol{\omega}$, and an azimuthal angle ϕ measuring a right-handed rotation about $\boldsymbol{\omega}$. In this coordinate system, the three canonical positions from Pham et al. (2021) are described by $\theta = 0°$, $\theta = 90°$, and $\theta = 180°$.

The eddy-current-induced force f and torque τ was empirically modeled, using both finite-element analysis (FEA) and experiments, at the three canonical positions as a function of the magnetic dipole strength $m = \|\boldsymbol{m}\|$, the dipole rotation frequency $\omega = \|\boldsymbol{\omega}\|$ (units Hz), the radius r (units m) of the conductive sphere, the distance ρ, and the electrical conductivity σ (units S/m) of the conductive sphere:

$$f, \tau = \frac{\left(c_0 \sigma \mu_0 \omega r^2\right)^{c_1 \left(\sigma \mu_0 \omega r^2\right)^{c_2}} 10^{c_3} \left(\mu_0 m^2\right)}{\left(\frac{\rho}{r}\right)^{c_4} r^{c_5}} \tag{1}$$

where $\mu_0 = 4\pi \times 10^{-7}$ N·A^{-2} is the permeability of free space. The coefficients for the $\theta = 0°$ and $\theta = 90°$ positions (which is all that we will need going forward, due to the symmetry of $\theta = 0°$ and $\theta = 180°$) are provided in Table 1. Pham et al. (2021) recommended using both the experimentally derived coefficients and FEA-derived coefficients to bound the estimates on the resulting force and torque. The force–torque model is quasistatic, as it was empirically derived using a static conductive sphere.

The model in Eq. (1) is accurate in a "far-field" regime in which the center of the nonmagnetic sphere is approximately 1.5 sphere radii or farther away from the center of the magnetic field source ($\rho > 1.5r$). This is not a particularly restrictive assumption, considering that the theoretical lower limit on ρ is $\rho = r$ (for a point dipole) and any actual magnetic field source has its own finite dimensions. In the "near-field" regime ($\rho \leq 1.5r$), the model Eq. (1) underpredicts the force–torque magnitude. In practice, the near-field regime applies to scenarios in which the physical magnetic field source is close to a much larger nonmagnetic, conductive object.

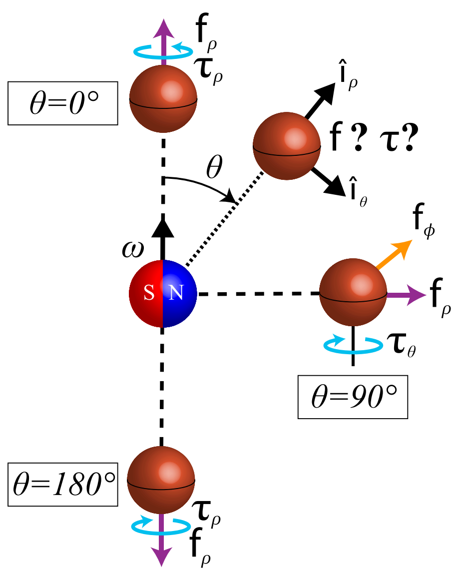

Figure 2. Eddy-current-induced forces and torques shown in a spherical coordinate system to describe arbitrary positions relative to a rotating dipole source. Note that the orthonormal basis for the coordinate system is defined such that $\hat{i}_\phi = \hat{i}_\rho \times \hat{i}_\theta$. The three canonical positions in Pham et al. (2021), and their respective forces and torques, are recast in the spherical coordinate system. The arrowhead on $\boldsymbol{\tau}_\rho$ at $\theta = 180°$ depicts the positive sign convention, which is opposite to the actual torque direction for the $\boldsymbol{\omega}$ shown. All other force/torque arrowheads depict both the positive sign convention and the actual force/torque direction for the $\boldsymbol{\omega}$ shown. In this image, $\boldsymbol{f}_\phi$ points into the page. The model makes no estimate of force/torque at other values of θ, such as those denoted by $\boldsymbol{f}$? and $\boldsymbol{\tau}$?.

2.2. New continuous-in-position model

We conducted new simulations of magnetically induced force and torque using Ansys multiphysics FEA, following the specifications provided in Pham et al. (2021). We placed the conductive sphere relative to the rotating dipole source from $\theta = 0°$ to $\theta = 180°$ at 15° increments, as shown in Figure 3(a) and (c). Our simulation had a dipole strength $m =$ 200 A·m^2, a dipole rotation frequency $\omega = 10$ Hz, a conductive-sphere radius $r = 50$ mm, a distance $\rho = 500$ mm, and conductive-sphere electrical conductivity of $\sigma = 5.8 \times 10^7$ S/m for copper.

The complete results of the FEA are shown in Figure 3, with the exception that components in the $\hat{i}_\phi$ direction are not depicted in Figure 3(a) and (c). From these results, it became evident that all six force and torque components can be expressed by simple trigonometric functions that provide a smooth transition between the modeled forces and torques

Table 1. Coefficients from Pham et al. (2021) for model in Eq. (1) for canonical positions, recast in spherical coordinates, obtained using both FEA simulations and experiments.

		Coefficients					
θ	f, τ	c_0	c_1	c_2	c_3	c_4	c_5
FEA simulations							
0°	f_ρ	430	2.95	−0.101	−9.26	7	4
0°	τ_ρ	6840	3.00	−0.0986	−13.2	6	3
90°	f_ρ	266	2.60	−0.101	−7.65	7	4
90°	f_ϕ	6040	3.45	−0.102	−14.3	7	4
90°	τ_θ	8100	3.60	−0.0985	−15.7	6	3
Experiments							
0°	f_ρ	467	2.81	−0.0969	−9.75	7	4
0°	τ_ρ	6900	3.35	−0.0990	−14.9	6	3
90°	f_ρ	282	3.20	−0.0980	−9.41	7	4
90°	f_ϕ	5870	3.49	−0.0973	−14.6	7	4
90°	τ_θ	8000	3.40	−0.0928	−15.0	6	3

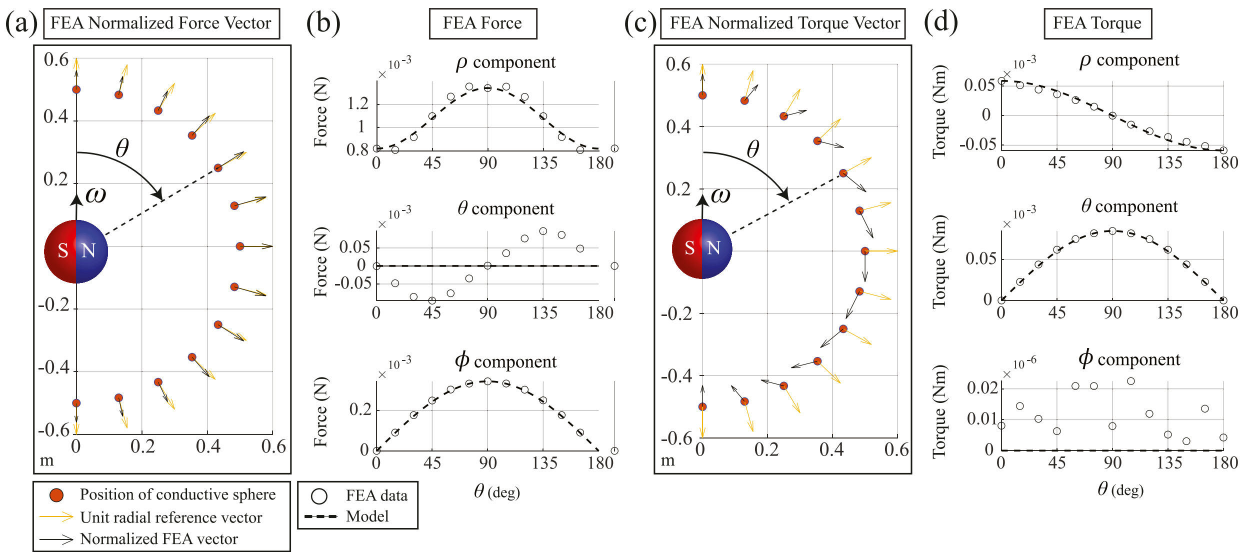

Figure 3. Complete results of FEA. (a) Force vectors in the $\hat{i}_\rho$-$\hat{i}_\theta$ plane for each conductive-sphere position, normalized by the value of f_ρ at $\theta = 90°$ (i.e., the maximum value). A $\hat{i}_\rho$ unit vector is shown in yellow for reference. (b) Three components of force vector as a function of θ, with trigonometric models (dashed lines) given by Eqs. (2)–(4), respectively. (c) Torque vectors in the $\hat{i}_\rho$-$\hat{i}_\theta$ plane for each conductive-sphere position, normalized by the value of τ_θ at $\theta = 90°$ (i.e., the maximum value). A $\hat{i}_\rho$ unit vector is shown in yellow for reference. (d) Three components of torque vector as a function of θ, with trigonometric models (dashed lines) given by Eqs. (5)–(7), respectively.

at the three canonical positions to arbitrary values of θ, and that also embody the symmetries that we would expect. The equations that describe the force and torque components in spherical coordinates—at arbitrary values of ρ and θ, and not requiring ϕ due to symmetry—which call the canonical-position model of Eq. (1), are as follows:

$$f_\rho(\rho,\theta) = -\left(\frac{f_\rho(\rho,90°) - f_\rho(\rho,0°)}{2}\right)\cos(2\theta) + \left(\frac{f_\rho(\rho,90°) + f_\rho(\rho,0°)}{2}\right) \tag{2}$$

$$f_\theta(\rho,\theta) \approx 0 \tag{3}$$

$$f_\phi(\rho,\theta) = f_\phi(\rho,90°)\ \sin(\theta) \tag{4}$$

$$\tau_\rho(\rho,\theta) = \tau_\rho(\rho,0°)\ \cos(\theta) \tag{5}$$

$$\tau_\theta(\rho,\theta) = \tau_\theta(\rho,90°)\ \sin(\theta) \tag{6}$$

$$\tau_\phi(\rho,\theta) = 0 \tag{7}$$

Our decision to model $\tau_\phi = 0$ comes from the observation that the simulation results appear to simply be numerical noise with no discernible pattern. The data produce none of the symmetry we would expect from a magnetic model and exhibit magnitudes three orders of magnitude smaller than the other torque components.

Our decision to model $f_\theta \approx 0$, even though Figure 3(b) suggests that f_θ is also described by a trigonometric function, is based on two considerations. First, it is evident from Figure 3(a) that forces in the $\hat{i}_\rho$-$\hat{i}_\theta$ plane are almost entirely in the $\hat{i}_\rho$ direction, such that ignoring the force component in the $\hat{i}_\theta$ direction will likely have a negligible impact on our ability to perform manipulation. Second, whereas the other five wrench components could be modeled as an interpolation between the values at the canonical positions, in the case of f_θ the value at the canonical positions is zero. Consequently, finding the value of f_θ at $\theta = 45°$ would require additional modeling, analogous to the efforts of Pham et al. (2021) that led to the model in Eq. (1). Since the negligibility of f_θ may be system/configuration dependent, modeling of f_θ may be justified in future work. For now, we will proceed with the assumption that closed-loop control will correct for any modeling deficiencies.

3. Manipulation framework

We now propose a control framework to perform dexterous manipulation with multiple dipole-field sources (at least partially) surrounding the conductive object.

3.1. Controller parameterization

We assume that each source is an electromagnet capable of dipole rotation about any axis (e.g., an Omnimagnet (Petruska and Abbott, 2014)). Both m and $\boldsymbol{\omega}$ can be controlled, but their maximum achievable magnitudes are coupled due to the low-pass-filtering effect of induction in the electromagnets; that is, an increase in ω results in a decrease in the maximum value of m that can be achieved before the amplifiers' voltage limits are reached. As in Pham et al. (2021), we have chosen to treat m and the direction of $\boldsymbol{\omega}$ (i.e., the unit vector $\hat{\boldsymbol{\omega}}$) as the control variables, and to use a constant rotation frequency ω, which simplifies our control problem but somewhat limits the peak tracking performance of the controller. We assume n electromagnets, with the ith electromagnet located at position $\mathscr{P}_i$. We assume a single conductive object at pose, $\boldsymbol{x}$ comprising a position $\mathscr{P}_c$ and orientation R_c (Lynch and Park, 2017). We can describe the conductive object by a displacement vector $\boldsymbol{\rho}_i = \mathscr{P}_c - \mathscr{P}_i$ with respect to each dipole source, where $\rho_i = \|\boldsymbol{\rho}_i\|$.

For each electromagnet, we parameterize the control variable $\hat{\boldsymbol{\omega}}$ with respect to the world frame (i.e., a common frame of reference) using spherical coordinates, with a polar angle ψ measured from the z-axis, and an azimuthal angle ξ measured about the z-axis and from the x-axis (see Figure 4). Given a pair (ψ, ξ), we can reconstruct

$$\hat{\boldsymbol{\omega}} = \begin{bmatrix} \sin(\psi)\cos(\xi) \\ \sin(\psi)\sin(\xi) \\ \cos(\psi) \end{bmatrix} \tag{8}$$

The angle θ can then be found using knowledge of $\hat{\boldsymbol{\omega}}$ and the $\boldsymbol{\rho}$ value for the electromagnet under consideration:

$$\theta = \operatorname{atan2}(\|\hat{\boldsymbol{\omega}} \times \boldsymbol{\rho}\|, \hat{\boldsymbol{\omega}} \cdot \boldsymbol{\rho}) \tag{9}$$

Let us first consider the special cases when $\theta = 0°$ or $\theta = 180°$, where only the radial force and torque components are non-zero. We construct a unit vector

$$\hat{\boldsymbol{\iota}}_\rho = \frac{\boldsymbol{\rho}}{\rho} \tag{10}$$

and then use Eq. (1) to solve for the induced force and torque on the conductive sphere:

$$\boldsymbol{f} = f_\rho(\rho, \theta)\hat{\boldsymbol{\iota}}_\rho \tag{11}$$

$$\boldsymbol{\tau} = \tau_\rho(\rho, \theta)\hat{\boldsymbol{\iota}}_\rho \tag{12}$$

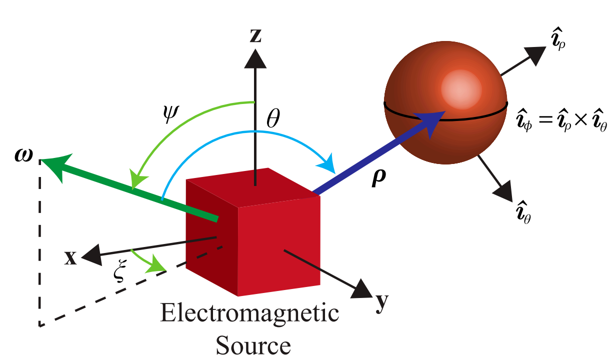

Figure 4. Spherical coordinate systems describing the dipole rotation vector $\boldsymbol{\omega}$ with respect to the world frame, and the conductive sphere with respect to $\boldsymbol{\omega}$ (as in Figure 2).

For all other values of θ, we can construct unit basis vectors that are compatible with the model of Section 2.2:

$$\hat{\boldsymbol{\iota}}_\phi = \frac{\hat{\boldsymbol{\omega}} \times \boldsymbol{\rho}}{\|\hat{\boldsymbol{\omega}} \times \boldsymbol{\rho}\|} \tag{13}$$

$$\hat{\boldsymbol{\iota}}_\theta = \hat{\boldsymbol{\iota}}_\phi \times \hat{\boldsymbol{\iota}}_\rho \tag{14}$$

where $\hat{\boldsymbol{\iota}}_\rho$ is calculated as in Eq. (10). The induced force and torque on the conductive sphere is then:

$$\boldsymbol{f} = f_\rho(\rho, \theta)\hat{\boldsymbol{\iota}}_\rho + f_\theta(\rho, \theta)\hat{\boldsymbol{\iota}}_\theta + f_\phi(\rho, \theta)\hat{\boldsymbol{\iota}}_\phi \tag{15}$$

$$\boldsymbol{\tau} = \tau_\rho(\rho, \theta)\hat{\boldsymbol{\iota}}_\rho + \tau_\theta(\rho, \theta)\hat{\boldsymbol{\iota}}_\theta \tag{16}$$

For ease of notation, we refer to this as our wrench model $\boldsymbol{f}, \boldsymbol{\tau} = w(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{\eta})$, where $\boldsymbol{\lambda}$ denotes a set of object parameters (e.g., sphere radius and conductivity) and $\boldsymbol{\eta} = \{i, m, \psi, \xi\}$ denotes the control parameters.

No closed-form inverse exists for the wrench model. Instead, for some instantaneous object pose and given set of object parameters, we solve a constrained optimization problem to effectively invert the model.

3.2. Wrench control policy

Given a desired wrench from some feedback controller, we select the dipole field source and associated dipole strength and axis of rotation that produces a wrench as close as possible. This is similar to the control algorithm used in Pham et al. (2021).

$$\begin{aligned} \underset{i, m, \psi, \xi}{\operatorname{argmin}} \quad & \left\| \begin{bmatrix} \boldsymbol{f}_{\text{des}} \\ \boldsymbol{\tau}_{\text{des}} \end{bmatrix} - \begin{bmatrix} \boldsymbol{f} \\ \boldsymbol{\tau} \end{bmatrix} \right\|_{\boldsymbol{Q}}^2 \\ \text{s.t.} \quad & i \in \{1, \ldots, n\}, \ m \in [0, m_{\max}], \\ & \psi \in [0, \pi], \ \xi \in [-\pi, \pi], \\ & \boldsymbol{f}, \boldsymbol{\tau} = w(\boldsymbol{x}, \boldsymbol{\lambda}, \{i, m, \psi, \xi\}) \end{aligned} \tag{17}$$

where the $\boldsymbol{Q}$-norm enables relative weighting between force and torque (which have different units), and for our system we set $m_{\max} = 40$ A·m^2. Note that reformatting $\hat{\boldsymbol{\omega}}$ as the pair (ψ, ξ) lets us construct the optimization without needing nonlinear constraints enforcing $\hat{\boldsymbol{\omega}}$ to be a unit vector.

We can efficiently find the optimal inputs using a parallelized (two initializations for each of the n electromagnets) Newton-method solver. We handle bound constraints through projection and use a backtracking line-search (Nocedal and Wright, 2006) to select step lengths.

3.3. Inverse-dynamics control policy

We additionally propose an inverse-dynamics controller. This controller uses the inverse of the object's mass matrix to solve for control actions given desired accelerations from some feedback controller, but otherwise using the same constraints and numerical techniques as before:

$$\underset{i,m,\psi,\xi}{\operatorname{argmin}} \quad \left\| \begin{bmatrix} \boldsymbol{a}_{\mathrm{des}} \\ \boldsymbol{\alpha}_{\mathrm{des}} \end{bmatrix} - M(\boldsymbol{\lambda})^{-1} \begin{bmatrix} \boldsymbol{f} \\ \boldsymbol{\tau} \end{bmatrix} \right\|_{\boldsymbol{Q}}^{2} \tag{18}$$

$$\begin{aligned} \text{s.t.} \quad & i \in \{1, \ldots, n\}, \ m \in [0, m_{\max}], \\ & \psi \in [0, \pi], \ \xi \in [-\pi, \pi], \\ & \boldsymbol{f}, \boldsymbol{\tau} = w(\boldsymbol{x}, \boldsymbol{\lambda}, \{i, m, \psi, \xi\}) \end{aligned}$$

3.4. Feedback controller

We assume we are given a time-varying wrench profile in the case of our forward-dynamics controller, or an acceleration profile in the case of our inverse-dynamics controller. Although not fundamental to our method, in this paper we generate a cubic spline between a series of waypoints and hand-tune the time allotted to each segment. We use simple proportional-derivative (PD) controllers to produce the requisite wrench or acceleration.

4. Adaptive control scheme

The model and control framework that we proposed in the previous sections expand the space of possible wrenches that we can induce compared to previous work, but they still require that the object being manipulated is a sphere of known properties. We aim to explore to what extent we can relax this assumption through the use of an adaptive control framework. Our hope is that, not only will this approach enable us to identify the parameters of spherical objects, but it will also enable manipulation of nonspherical objects by identifying online a spherical approximation that describes the observed behavior of the object.

We now formalize the system identification and adaptive control problem, which aims to find the optimal physical parameters of a sphere, $\boldsymbol{\lambda}^*$, given a discrete, time-varying sequence of object poses and input controls $\Omega = (\boldsymbol{x}[0], \boldsymbol{\eta}[0], \ldots, \boldsymbol{x}[K], \boldsymbol{\eta}[K])$ with time horizon K. Unlike the empirical model of wrenches induced in spheres, which was derived using wrench measurements, we now assume at the time of deployment that we only observe the object pose. Although many modalities (e.g., lidar, cameras and fiducial markers, RGB-D cameras) could be used to determine object pose in practice, we use visual observations in our experiments. We use an online smoothing formulation to track the object, deriving less noisy pose estimates as well as associated object velocities $\dot{\boldsymbol{x}}[k]$ (more details in Section 5). In our wrench model, the free model parameters $\boldsymbol{\lambda}$ are the spherical radius, r, and electrical conductivity, σ, which should both be constrained to be greater than zero. We can then estimate the mass matrix, $M(\boldsymbol{\lambda})$ as:

$$M(\boldsymbol{\lambda}) = \begin{bmatrix} \frac{4}{3}\pi r^3 \gamma \mathbb{I}_{3\times3} & 0_{3\times3} \\ 0_{3\times3} & \frac{8}{15}\pi r^5 \gamma \mathbb{I}_{3\times3} \end{bmatrix} \tag{19}$$

where we assume a solid sphere with given density γ.

We can connect these observed data with our magnetic wrench model, w, and thus our control inputs, by imposing a rigid-body motion model on the object dynamics. Given our target domain of space debris, our motion model assumes no friction and a simple linear mapping between the applied wrench and the resulting acceleration (i.e., Newton's second law). To solve for the parameters of our model as an optimization problem, we must define an associated loss (i.e., error) function over the observed data and dynamic object parameters. We consider two loss formulations. The first is the inverse-dynamics or acceleration-based loss formulation

$$\mathscr{L}_a(\boldsymbol{\lambda}, k) = \left\| \frac{\dot{\boldsymbol{x}}[k+1] - \dot{\boldsymbol{x}}[k]}{\delta t} - M(\boldsymbol{\lambda})^{-1} w(\boldsymbol{x}[k], \boldsymbol{\lambda}, \boldsymbol{\eta}[k]) \right\|_{\boldsymbol{Q}}^{2} \tag{20}$$

where δt is the controller's update period, and where the mass matrix enters via its inverse. The second is the forward-dynamics or wrench-based loss formulation,

$$\mathscr{L}_w(\boldsymbol{\lambda}, k) = \left\| M(\boldsymbol{\lambda}) \left(\frac{\dot{\boldsymbol{x}}[k+1] - \dot{\boldsymbol{x}}[k]}{\delta t} \right) - w(\boldsymbol{x}[k], \boldsymbol{\lambda}, \boldsymbol{\eta}[k]) \right\|_{\boldsymbol{Q}}^{2} \tag{21}$$

where the mass matrix enters directly. In both cases, we use finite differencing to estimate the acceleration from the observed object velocities. The weights used in the two $\boldsymbol{Q}$-norms would be different in general. Using either loss formulation, we can construct the batch system identification problem (Atkeson et al., 1986) as the following optimization:

$$\boldsymbol{\lambda}^* = \underset{\boldsymbol{\lambda}}{\operatorname{argmin}} \sum_{k=0}^{K-1} \mathscr{L}(\boldsymbol{\lambda}, k) \tag{22}$$

We investigate this batch formulation as a baseline in our experiments below. However, our primary interest lies in identifying the object parameters online. It is not obvious how to generate a safe set of controls to collect the data for system identification when the object properties are unknown. This motivates our adaptive control formulation.

In adaptive control, we leverage our model-based control to define the control signal, while updating the estimate of $\boldsymbol{\lambda}$ online based on our observations. Typically, in performing adaptive control, we would not fully solve this optimization at each step, but instead perform a single gradient step to update the parameters:

$$\boldsymbol{\lambda}[k+1] = \boldsymbol{\lambda}[k] - \alpha_k \nabla_{\boldsymbol{\lambda}} \mathscr{L}(\boldsymbol{\lambda}[k], k) \tag{23}$$

with some step length α_k (Siciliano et al., 2009; Slotine and Li, 1991). However, by framing the adaptive controller as an online optimization problem (Ratliff et al., 2016), we can use a broad set of tools in deciding on how to solve for the system parameters. In particular, we wish to explicitly model bound constraints on our object parameters and examine different solvers, including the momentum optimizer that has been shown to improve performance over gradient descent by smoothing out oscillations (Qian, 1999). We can additionally examine mini-batch formulations of the optimization, where we use the most recent k timesteps of pose and control data instead of the full batch as traditionally done in system identification or only a single step as traditionally done in adaptive control. In our experiments, we select our step length, α_k, online using the same backtracking line search used in our controller, and handle constraints using the same projection approach (Nocedal and Wright, 2006). We give further details of the design choices we examined for solving the adaptive control problem, including the performance of the different loss functions, in the following section.

5. Experimental validation

5.1. Microgravity simulators

We preformed physical experiments using the same system used in Pham et al. (2021) comprising four Omnimagnets (i.e., omnidirectional electromagnets) placed beneath a water tank (see Figure 5); the water's surface serves as a 3-DOF (2-DOF position + 1-DOF orientation) microgravity simulator. The Omnimagnets can each produce an approximate dipole source rotating about an arbitrary axis to match the fully continuous ω produced by our controller. Each Omnimagnet comprises three mutually orthogonal nested coils surrounding a ferromagnetic core. Each coil is driven by a current-drive amplifier (AMC16A8, Advanced Motion Control) with current and voltage limits of 8 A and 80 V, respectively. The amplifiers were connected in parallel to a shared power supply (PS16L80, Advanced Motion Control) with current and voltage limits of 10 A and 80 V, respectively. We used the linear approximation from Petruska and Abbott (2014) with α = 7.00 m^2 to map the desired dipole moment to the coil currents. In Pham et al. (2021), we found a frequency of 15 Hz and dipole strength of 40 A·m^2 maximized observed wrenches given our hardware limitations. We used a fixed 15 Hz for ω and set the maximum dipole strength to 40 A·m^2 for all experiments.

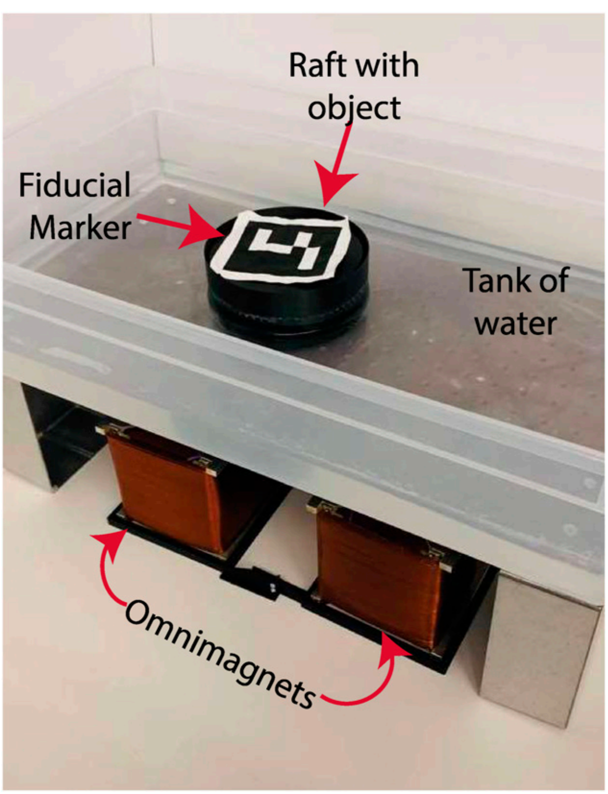

Figure 5. Physical microgravity simulation system, which uses four omnimagnets to perform 3-DOF (2-DOF position + 1-DOF orientation) manipulation experiments.

Figure 1 provides a top-down view of the environment. We placed a camera above the water tank to detect a fiducial marker rigidly attached to an object of interest. We fit independent cubic splines (Dierckx, 1995) online to the measured 2D position and orientation variables (x, y, τ). This allows us to decrease the influence of noise from the instantaneous marker locations and estimate the object velocity. We tuned the number of knot points for each spline so that the pose smoothly varied over the slow timescales we operate at. Buoyant objects can be placed in the water directly, whereas other objects are placed inside of a plastic raft that floats on the surface. For experiments with the object placed inside the plastic raft, we estimated the mass matrix of the raft as $M_r = \texttt{diag}$ (0.15, 0.15, 0.15, 3.7×10^{-4}, 3.7×10^{-4}, 6.7×10^{-4}), where the first three terms have units kg and the last three terms have units kg·m^2. We add the known raft mass matrix with the estimated mass of the object, $M_o(\boldsymbol{\lambda})$, in the dynamics model to remove its effects on control, $M(\boldsymbol{\lambda}) = M_r + M_o(\boldsymbol{\lambda})$. This additional mass enters into our system-identification problem and our inverse-dynamics control, but it does not enter into our wrench-based control.

Recently, in a related work, Dalton et al. (2022) showed that the drag created by our water-based microgravity simulator was not negligible, nor was it accurate to assume that drag is linear with respect to velocity (i.e., low Reynolds number, Stokes flow), even at the relatively slow velocities at which we are performing manipulation. Thus, we provide a simple drag model with linear and quadratic terms to the system identification on the physical experiments:

$$\begin{bmatrix} \boldsymbol{f}_{\text{drag}} \\ \boldsymbol{\tau}_{\text{drag}} \end{bmatrix} = -\begin{bmatrix} a\dot{\boldsymbol{x}}_{1:3} + c\|\dot{\boldsymbol{x}}_{1:3}\|\dot{\boldsymbol{x}}_{1:3} \\ b\dot{\boldsymbol{x}}_{4:6} + d\|\dot{\boldsymbol{x}}_{4:6}\|\dot{\boldsymbol{x}}_{4:6} \end{bmatrix} \tag{24}$$

Note that, although we construct this drag model as a full 6-DOF wrench, since we only observe velocities in 3-DOF (i.e., 2-DOF translational and 1-DOF rotational) we only have drag in 3-DOF. The four coefficients are a function of the shape of the raft and the combined mass of the raft and object (which affects where the raft sits in the water). For simplicity, we fit this model using a single object-manipulation trial and kept it fixed with $a = 8.12 \times 10^{-3}$ N·s·m^{-1}, $b = 1.13 \times 10^{-5}$ N·m·s·rad^{-1}, $c = 2.15$ N·s^2·m^{-2}, and $d = 1.90 \times 10^{-4}$ N·m·s^2·rad^{-2} for all experiments and objects.

For all physical experiments, we reproduced the trajectory-tracking experiments from Pham et al. (2021). This task requires the system to control the object of interest to track a 3-DOF planar Cartesian trajectory to draw a square in a counterclockwise direction, after starting from the center of the square and moving out to one corner, reorienting the object to point in the direction of motion each time it reaches a corner, and maintaining a fixed heading toward the next corner during motion along the edges.

We also implemented a numerical simulator, based on Jiang et al. (2017), to simulate 6-DOF manipulation that would be impractical in our physical simulation environment. For consistency, we impose the same limitations on our magnetic field sources as our physical experiments. With our numerical simulations, we can perform manipulation without any external wrenches on our system (e.g., gravity, drag).

5.2. *Quantifying force–torque model improvements*

To quantify the additional control authority offered by our new continuous-in-position model of Section 2.2 over our prior model of Section 2.1, we performed manipulation experiments in numerical simulation using each model. Each trial was initialized by sampling a 6-DOF pose from a uniform distribution. The object was then controlled to align with the workspace frame at the origin. For the new model, we use the control policy from Eq. (17), which we compare to the equivalent control policy from Pham et al. (2021). The underlying position controller producing the desired force–torque wrenches, and the 1000 random initial poses, were identical for both models. For both models, we consider the use of two field sources (the minimum required for manipulation) and four field sources in a tetrahedron arrangement (the minimum required to effectively cage the object, in an arrangement that optimally surrounds the object), with the sources equidistant from the center of the workspace. We define the settling time as the time at which the center of the object enters a spherical region that has a radius that is 1% of the initial error, and the orientation error also drops below 1% of its initial error, and then never respective errors never again become larger than these threshold values. The normalized path length is defined as the distance traveled by the object from the start of the trial to the settling time, normalized by the shortest possible distance (i.e., 99% of the initial error). This metric captures both non-straight paths and overshoot. In Figure 6, we visualize the normalized path length for both position and orientation using both control policies. We see the new model provides a statistically significant improvement (using a standard significance of $\alpha = 0.05$), taking a more direct path, which we use as proxy for control authority.

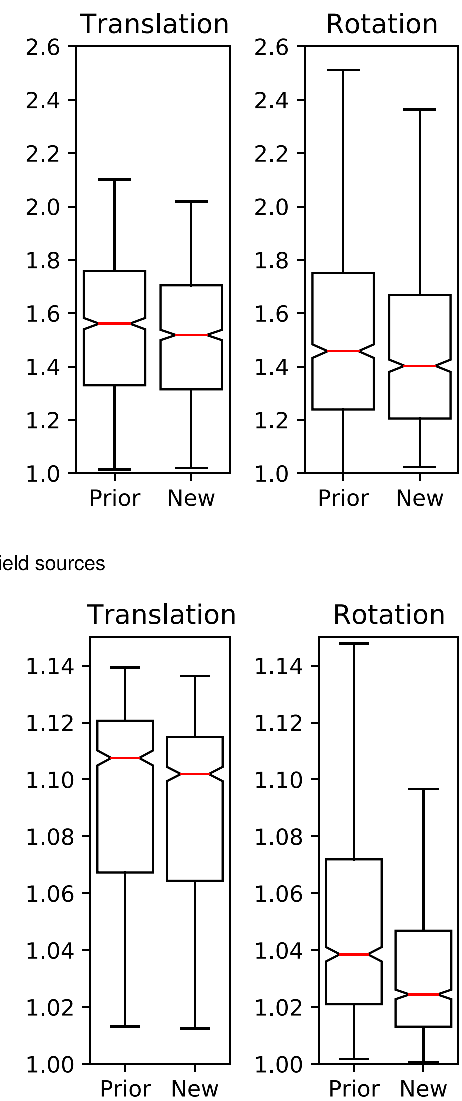

(a) Two field sources

(b) Four field sources

Figure 6. Notched-box-whisker plots showing the regularized path lengths from 1000 random initial poses to the workspace origin frame. A value of 1 corresponds to a path with the shortest possible distance. The notches indicate the 95% confidence interval on the median. Outliers have been omitted for clarity, but they maintain the same trends shown. (a) Two field sources. (b) Four field sources.

5.3. *Acceleration-based versus wrench-based loss for system identification*

We conducted a numerical simulation to examine optimization choices for system identification and adaptive control. Our simulation was in 6-DOF with six magnetic field sources, with the field sources equidistant from the center of the workspace and with a maximal-separation distribution (i.e., a tetrahedron). We used our proposed wrench-based controller shown in Eq. (17). Our primary objective was to examine which of our proposed loss functions ((20) or (21)) performed best when estimating the radius r and conductivity σ of the sphere being manipulated. To this end, we computed both losses with varying values of sphere radius, with known conductivity, in the batch system-identification setting. We visualize the log-loss for both functions in Figure 7. We see that the acceleration-based loss has a single minimum within the feasible range of radii, which coincides with the true radius. The wrench-based loss, on the other hand, has a local minimum at the true radius, but a global minimum at $r = 0$ as both the wrench and mass-matrix M decay to zero. With the wrench-based loss, there is a substantial region where the gradient points toward the global minimum at $r = 0$ instead of the true radius. As such, we elect to use the acceleration-based loss for all subsequent experiments to avoid the degenerate solution.

5.4. Robustness of parameter identification

To investigate the robustness of our system identification optimization with respect to initialization, we ran a square-trajectory physical manipulation trial with accurate physical parameters and the wrench-based controller. We randomly initialized our batch system identification optimization problem 100 times and optimized each of them until convergence, using Newton's method. The true object parameters were $r = 0.02$ m and $\sigma = 5.8 \times 10^7$ S/m, and the bound constraints on the object parameters were set to $r \in [0.001, 0.1]$ m and $\sigma \in [2 \times 10^6, 8 \times 10^7]$ S/m (this range of σ spans common metals). In Figure 8(a), we plot the evolution of each of the optimizations and see that all 100 particles reach the global minimum quickly; although it does not converge perfectly on the true parameters, it converges on parameters that have equivalent performance.

We note that, in our initial conference publication (Tabor et al., 2022), the adaptive system-identification problem was poorly scaled, so conductivity never meaningfully changed. Here, we optimized over radius directly while conductivity

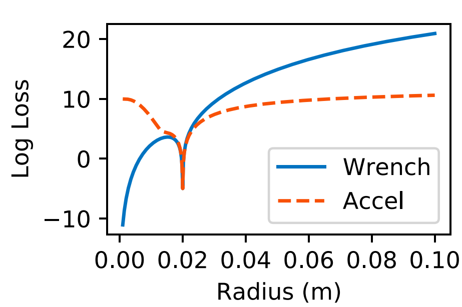

Figure 7. Natural logarithm of the wrench-space loss formulation $\mathscr{L}_w$ versus acceleration-space loss formulations $\mathscr{L}_a$, with a known conductivity and a true radius of 0.02 m.

was rescaled by 10^{-9} so the decision variables are the same order of magnitude.

In Figure 8(b), we consider a mini-batch (i.e., 25 timesteps) of the same data and we find that when using a small amount of data it is common for there to be local minima at a bound constraint. The heatmap suggests that there was not enough data to reject certain parameter regions (e.g., very low values of r). We observed a degenerate situation in our controller at the constraints; for example, with extremely low conductivity there is no gradient information for the controller optimization.

We can add an $L2$ regularization term to the system-identification loss formulation as $\mathscr{L}_a + \mathscr{L}_p$ to encode a prior preference for values in the middle of our parameters space (Murphy, 2021), where

$$\mathscr{L}_p(\boldsymbol{\lambda}) = \left\|\boldsymbol{\lambda}_p - \boldsymbol{\lambda}\right\|^2 \tag{25}$$

In Figure 8(c), we see the effect of this soft constraint, with results converging on the region known to have good performance from the complete data set.

5.5. Adaptive control of unknown spheres

In this section, we run adaptive-control experiments with unknown spheres using different permutations of our controller to select the best-performing version. For each permutation, we ran three trials with each of the copper and aluminum spheres shown in Figure 9, which each have a true radius of $r = 0.02$ m. We declare any trial where the raft runs into the wall a failure and prematurely halt the trial when it does so. We compare the resulting errors across the three controller permutations in Figure 10 for each object; we also compare to the two fixed-parameter controllers, which rely on imperfect models, manipulating the copper sphere. For the halted trials, we only compute the tracking error up to the time the trial was halted. In each adaptive trial, we examine optimizing over both the radius and electrical-conductivity parameters of the model. We initialized the parameters by setting the radius to $r = 0.01$ m, the electrical conductivity to the true value for copper, which is $\sigma = 5.80 \times 10^7$ S/m, and the density to that of copper, which is $\gamma = 8940$ kg/m^3. Even the aluminum sphere was set with the fixed density of copper and initialized with conductivity of copper. We conduct a single gradient update step at each iteration of the control loop, selecting the update step length using a backtracking line search. We elected to use the momentum strategy (Qian, 1999) given results reported by Ratliff et al. (2016). We used the same mini-batch size of 25 timesteps as we did in our offline experiments discussed earlier.

Let us compare our wrench-based controller with and without the added prior/regularization term. We find that including the prior/regularization term can improve the performance of adaptive control, reinforcing the results from our offline system-identification experiments of

Section 5.4. We show the adaptive parameters from the trials with the copper sphere in Figure 11, and we see that without the prior/regularization term the radius or conductivity were at the lower-bound constraints in all three trials shortly before running into a wall.

Note: We observed similar behavior when tuning our controller in our initial conference publication (Tabor et al., 2022), which lacked the prior/regularization term; without a more principled way to address the problem, we elected to use much tighter bounds on the adaptive parameters. Given the results described above, combined with the results of Section 5.4, we elect to include the prior/regularization term in all future experiments due to its lower errors and failure rate.

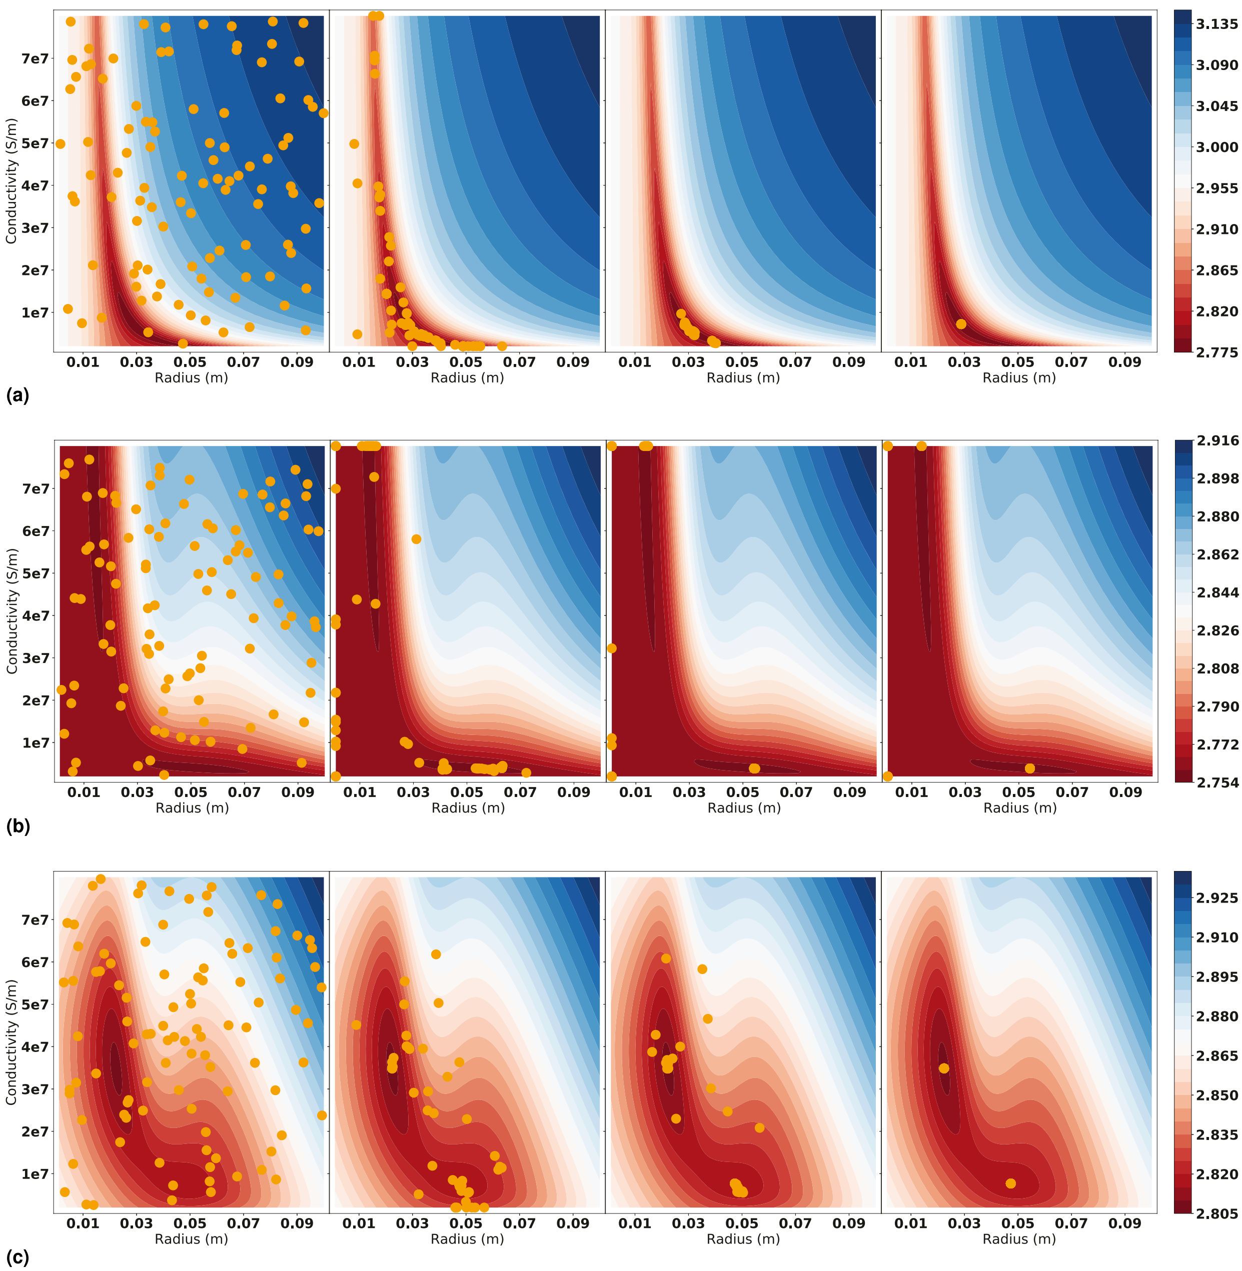

Figure 8. Offline system-identification results for a square-trajectory physical manipulation trial of a known object, after randomly initializing the problem 100 times and optimizing until convergence. The axes show object parameters and the heat map depicts the log of the log of the loss function. (a) Identification with data from the full trajectory, using loss function $\mathscr{L}_a$. (b) Identification with a mini-batch (25 timesteps) of the trajectory, using loss function $\mathscr{L}_a$. (c) Identification with the same mini-batch of data, using the loss function $\mathscr{L}_a + \mathscr{L}_p$ that incorporates a prior/regularization term that imposes a cost on parameters far from the center of the parameter space. The true object parameters are 0.02 m for radius and 5.8×10^7 S/m for conductivity.

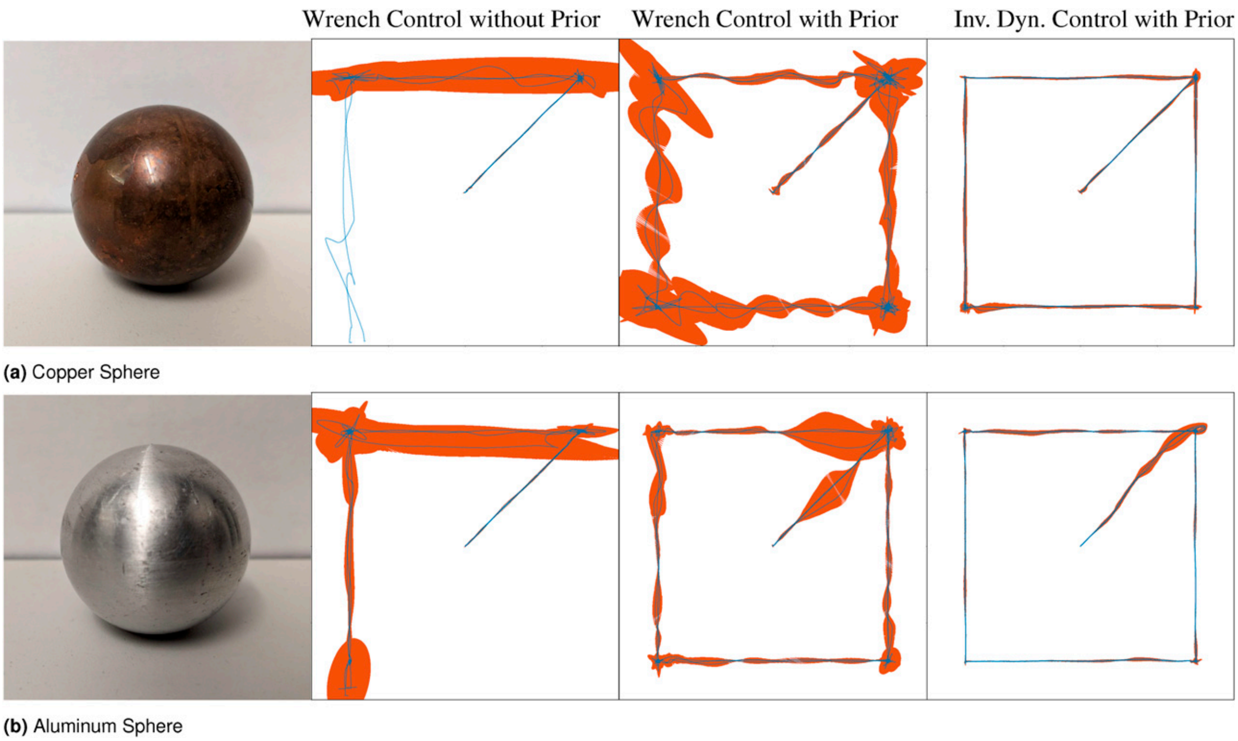

Figure 9. Adaptive control of spherical objects being manipulated along a square trajectory. The blue lines show individual trials and the red shading shows the 95% confidence path computed given three trials per object. We stop computing and visualizing the 95% confidence path after any of the trials collides with the edge of our water tank and is preemptively ended. (a) Copper sphere. (b) Aluminum sphere.

Next, let us compare our wrench-based controller with our inverse-dynamics controller, both with the prior/regularization term. In Figure 9, we see that our inverse-dynamics controller is markedly more repeatable, with lower position error. This lower position error is reinforced in Figure 10, and we see that orientation error is also lower. In Figure 11, we see that the improved control is associated with improved stability of the adaptive parameters. With the wrench-based controller, although the trials never fail, the control performance suffers to some degree as the object parameters vary dramatically. With the inverse-dynamics controller, both the estimated radius and estimated conductivity tend to vary less. We visualize characteristic trials in Figure 12.

We attribute the differences between the wrench-based controller and the inverse-dynamics controller to the inverse-dynamics controller being better aligned with the inverse-dynamics loss we use to fit the object parameters. We find the inverse-dynamics controller is more robust to the parameter fluctuations caused by the existence of multiple sets of parameters that would result in the same observations. A good model fit from our system-identification problem states that the object accelerates (which we can observe) like an object that had the estimated set of parameters would, not that the forces (which we cannot observe) are the same. As such, we elect to use the inverse-dynamics controller in all future experiments.

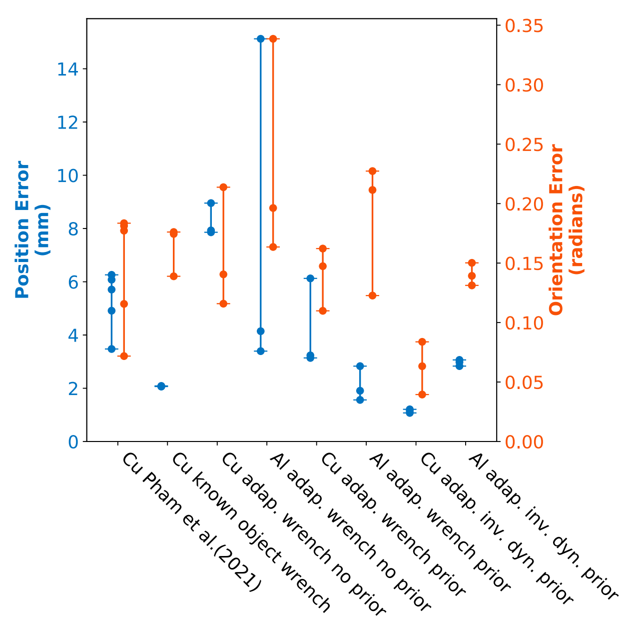

Figure 10. Tracking error for experiments manipulating a sphere. Points show average absolute error per trial, and bars show the range across all trials. Cu = copper; Al = aluminum. Results from "Pham et al. (2021)," which used the discrete wrench model of Section 2.1, and for "Cu known object wrench," which uses the continuous wrench model of Section 2.2, serve as a baseline with known objects.

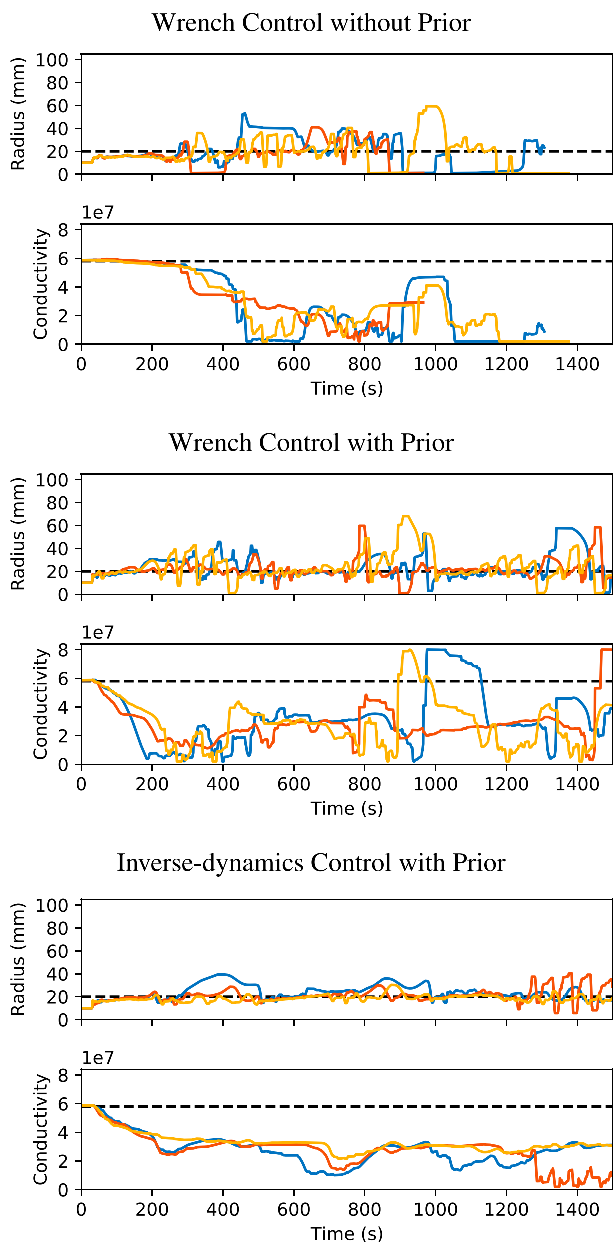

Figure 11. Adaptive parameters from three trials manipulating the copper sphere with each controller configuration. The black dashed line shows our estimate for the ground truth of the parameter. We visualize the first 1500 s from the 2400 s trials. Any trial that collides with the edge of our water tank is preemptively ended.

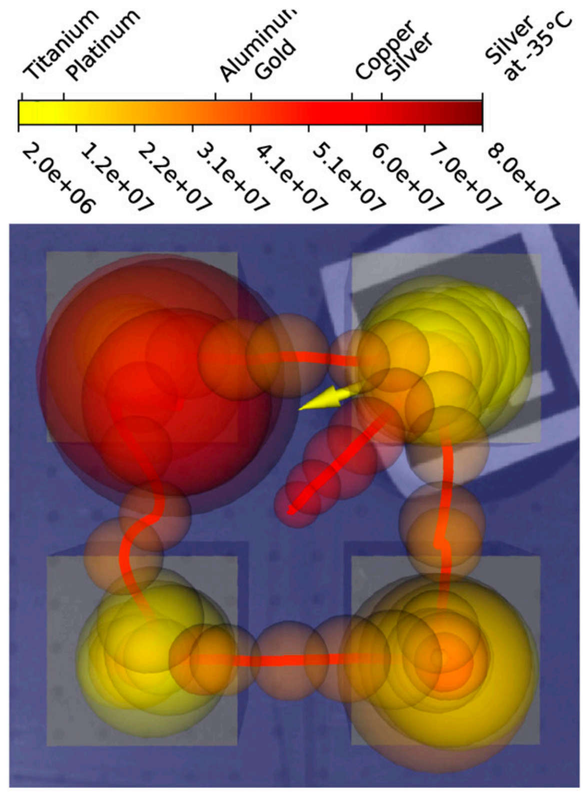

(a) Wrench-based control with prior

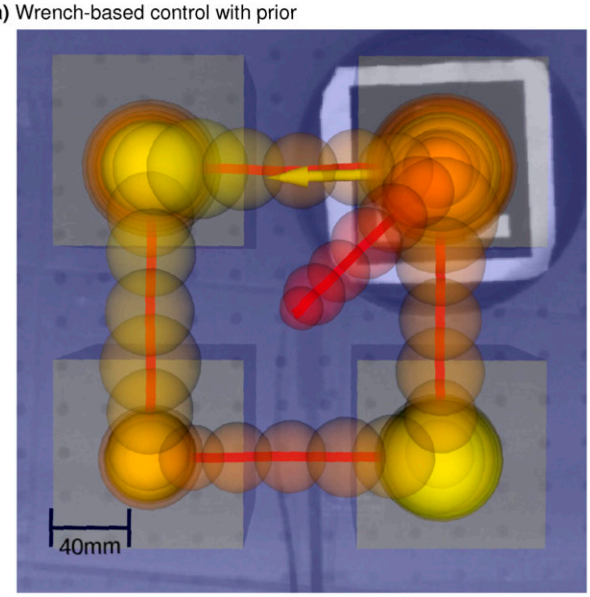

(b) Inverse-dynamics control with prior

Figure 12. Characteristic trials manipulating a copper sphere along a square trajectory using (a) our wrench-based controller and (b) our inverse-dynamics controller, both including a prior/regularization term. We visualize the radius and conductivity estimate at different points in the trajectory with periodically placed spheres. The conductivity-scale values are given at room temperature unless otherwise stated. The sphere's true diameter is 40 mm, indicated by the scale bar. See Supplemental Video for (b).

5.6. Sensitivity to the number of field sources

All of our physical experiments use four magnetic-dipole field sources, which are in one particular configuration, to perform 3-DOF manipulations tasks. We were also interested in how sensitive our adaptive controller's performance is to the number of field sources, and how it works with full

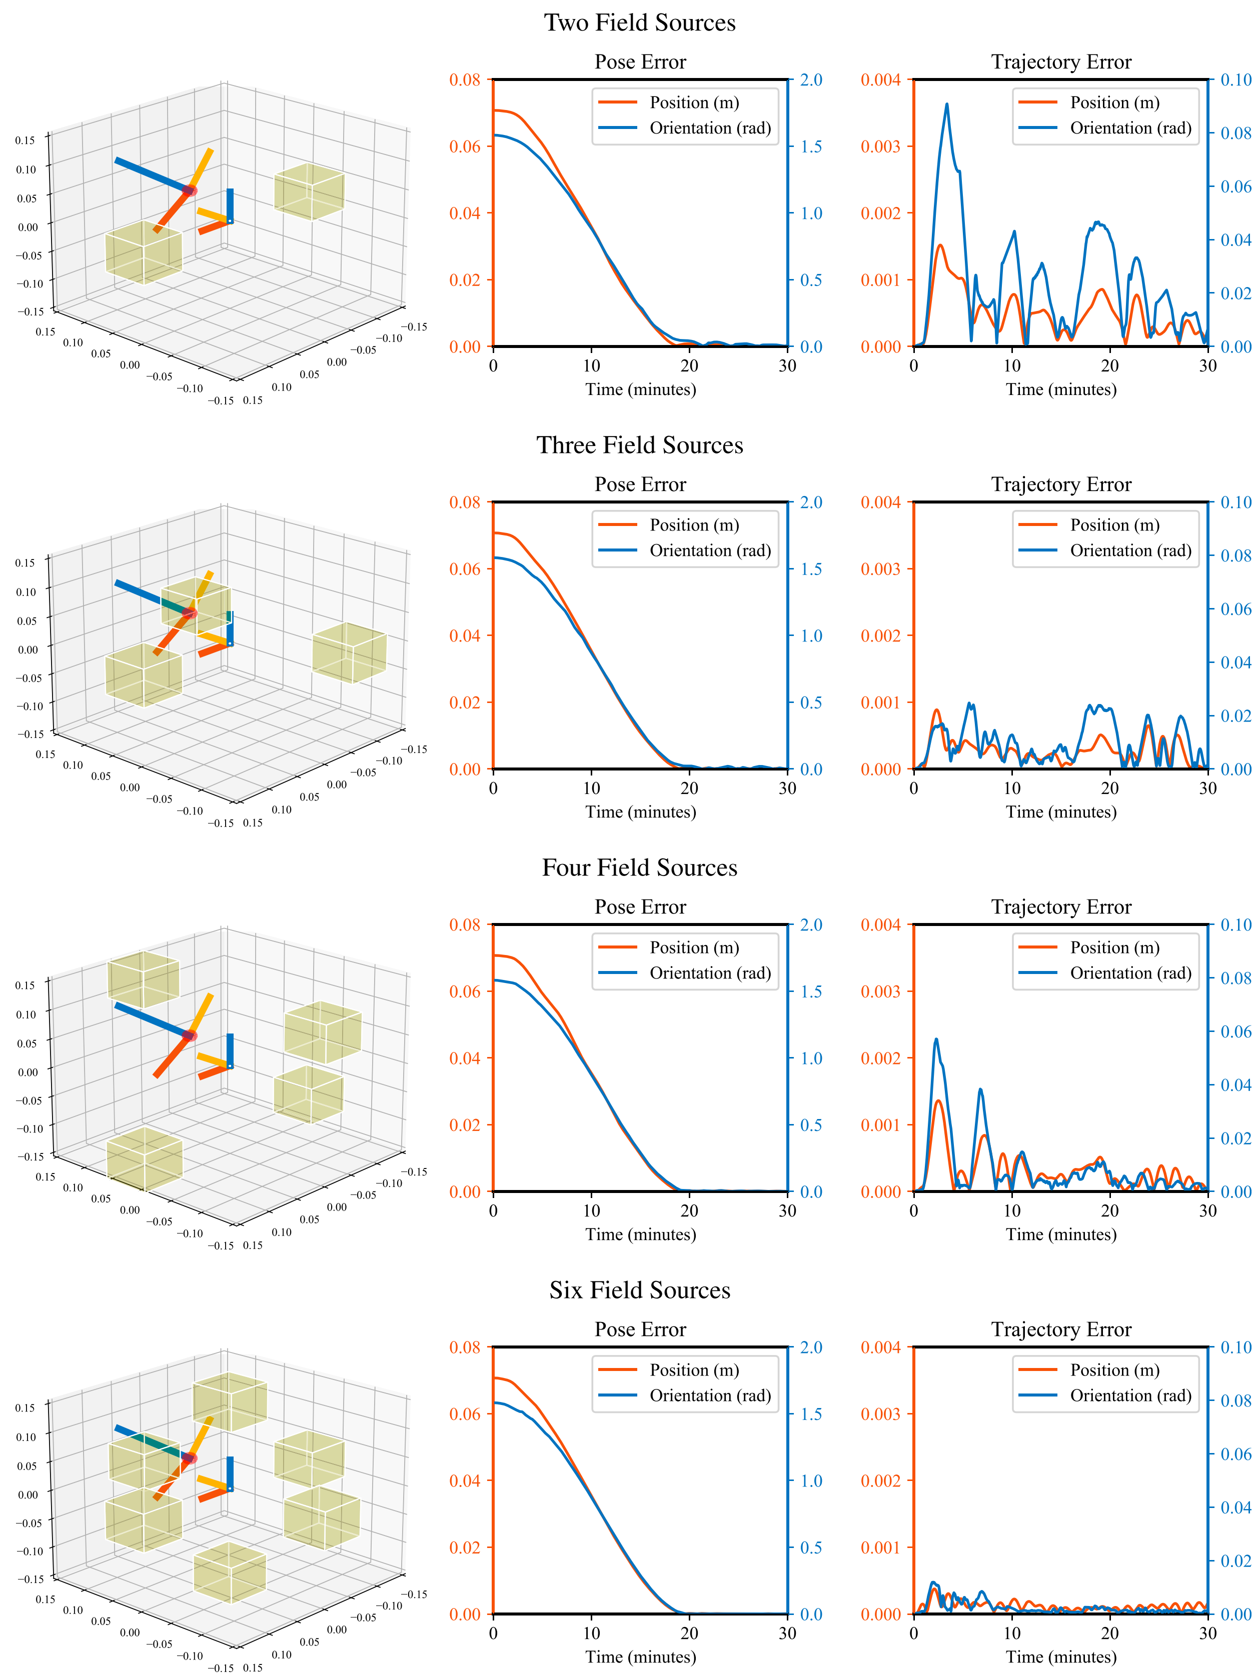

Figure 13. 6-DOF adaptive inverse-dynamics control, with a prior/regularization term, with varying number of magnetic-dipole field sources. The goal is to smoothly move from an arbitrary initial pose to align with a coordinate frame at the origin of the workspace (i.e., reject "pose error"). We also show "trajectory error" with respect to a simulated nominal trajectory in which we can directly induce the acceleration requested by our PD controller. See Supplemental Video.

6-DOF manipulation tasks. In Figure 13, we visualize the tracking performance of our adaptive inverse-dynamics controller with a prior/regularization term for an unknown copper sphere with a varying number of magnetic-dipole field sources surrounding the workspace, which we obtained using our numerical simulation environment. In each configuration, the field sources are equidistant from the center of the workspace and have a maximal-separation distribution. We compare the resulting performance to a nominal trajectory that was generated assuming that we can exactly induce the accelerations requested by our controller. We see that our ability to closely match an arbitrary trajectory increases with the number of field sources, with negligible error with six field sources. However, we also see that our controller is able to accomplish a 6-DOF manipulation task with as few as two field sources, with a performance that is ultimately very similar to that with six field sources, under the assumption that the object is near the center of the workspace.

5.7. Adaptive control of nonspherical objects

Our most challenging set of experiments test the hypothesis that the spherical model acts as a good first-order approximation for control of nonspherical objects. We use our adaptive inverse-dynamics controller, with a prior/regularization term, to manipulate nonspherical objects.

We conducted three square-trajectory manipulation trials on each of the eight nonspherical objects shown in Figure 14. The objects include: a solid copper cuboid; a solid aluminum cylinder with a small hole drilled through its axis; a pile of copper scrap objects, which are in contact and potentially electrically conducting; a complex aluminum structure, which is an Omnimagnet frame without the wire wrapped on it, with the entire structure electrically conducting; three pieces of extruded aluminum that are electrically isolated from each other; a five-sided thin-walled aluminum box that has been coated such that it is not electrically conducting with the water; the same aluminum box with a 12 g piece of (ferromagnetic) iron in the center; and the same aluminum box with a 40 g piece of iron in the center. We initialize the parameter estimates as in the previous experiment, with the density set to that of copper, the conductivity initialized to that of copper, and the radius initialized to $r = 0.01$ m. Because the aluminum box is larger than the plastic raft we used for our other experiments, the square trajectory was shrunken for all aluminum-box trials to give an appropriate buffer with the walls of our water tank.

We visualize tracking performance for all trials with these objects in Figure 14, with tracking error quantified in Figure 15. We find that tracking performance with the two simple solid objects is very similar to what we obtained with spheres, indicating that the spherical model is, in fact, a good approximation for these objects, as hypothesized. What is more surprising is that the tracking performance with the aluminum extrusions is nearly as good, even though it comprises three isolated conductive objects, whereas the model assumes a single spherical object in the center. Trials with the aluminum structure and the aluminum box are nearly as good after an initial learning period; this suggests that better initialization of object parameters could help improve performance, but how to better initialize them is left as an open problem. It is interesting to note that the addition of small amounts of ferromagnetic material does not have a major effect on tracking performance. We attribute this to our relatively high-frequency rotating magnetic fields generating net-zero periodic forces and torques on the ferromagnetic elements. We observed the worst performance with the copper scrap, and had to stop the trials prematurely when the raft collided with the wall of the water tank.

Characteristic trials of manipulation of the copper cuboid and the aluminum extrusions are shown in Figure 16, with the estimated object parameters visualized at different times across the trajectory. We see that the controller never settles on a single estimate of the object parameters, instead adapting the value continuously based on the locally observed motion.

In summary, even though our wrench model was built entirely from data derived from spheres, we achieve comparable tracking performance to that of the spheres when manipulating objects with significantly nonspherical geometry.

For a final experiment in this section, in Figure 17 we visualize the field source (i.e., Omnimagnet) that is used at each timestep along the trajectory. We can see that the algorithm predominately uses the two field sources that are nearest to the object, which is not surprising considering how rapidly induced forces and torques decrease with distance from a field source. However, occasionally the best choice is the field source that is farthest from the object, which is an indication that each field source is limited in its actuation authority, particularly in terms of directions of achievable wrenches.

5.8. Additional experimental analysis

Our final set of experiments attempt to address the shortcomings we observed in our manipulation experiments in Section 5.7, specifically, the relatively poor performance with the copper scrap and aluminum structure as well as the diminished performance during the beginning of the trials with the aluminum box.

The copper scrap failed in all three trials by hitting the edge of our tank. Upon further inspections, we found that the system was saturated at full power 95% of the time during the three trials. This is compared to the controller being saturated just 5% of the time during the three trials for manipulating the copper sphere in Figure 9. To address this problem, we performed an additional three trials with the copper scrap, eliminating the initial segment where the object is farthest from the field sources and doubling the length of time allotted for each segment. We obtained vastly improved performance, as shown in Figure 18. Interestingly, we also see that the adaptive control quickly finds useful parameters, but over time two of the trials switch to a different set of useful parameters. This improvement implies that the shortcoming was not in our adaptive

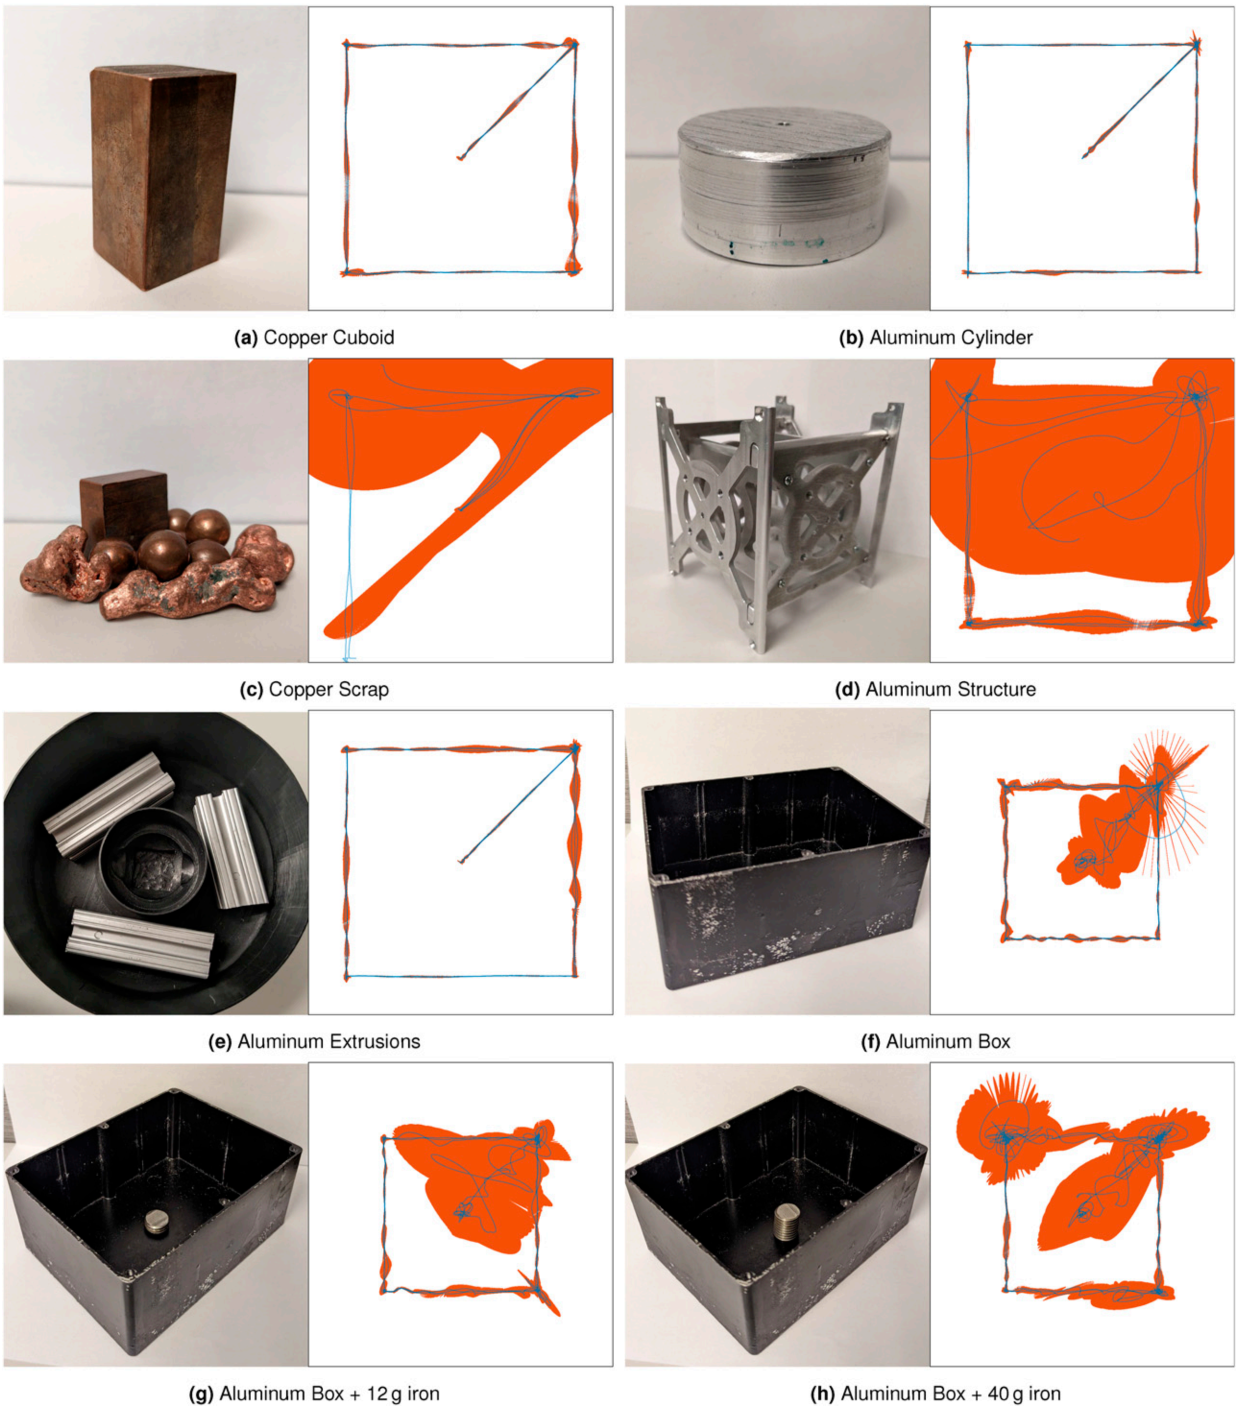

Figure 14. Adaptive inverse-dynamics control, with a prior/regularization term, on a variety of nonspherical objects being manipulated along a square trajectory. Blue lines show individual trials and the red shading shows the 95% confidence path computed given three trials per object. The copper scrap object failed all three trials by hitting the walls of the tank and was preemptively ended before completing the square trajectory. (a) Copper cuboid. (b) Aluminum cylinder. (c) Copper scrap. (d) Aluminum structure. (e) Aluminum extrusions. (f) Aluminum box. (g) Aluminum box + 12 g iron. (h) Aluminum box + 40 g iron.

controllers ability to extend to complex novel objects, but rather, that the given trajectory was not possible for our hardware system to achieve. In practice, this could be automatically accounted for by adjusting the time-parameterized trajectory online, either by simply slowing down when saturating or with full motion planning.

Next, we reconsider the aluminum structure. Interestingly, for this object the orientation tracking error is approximately 10 times higher than the other objects, whereas the position error is only double. During the three manipulation trials with the aluminum structure, the system was saturated at full power 75% of the time. Similar to the

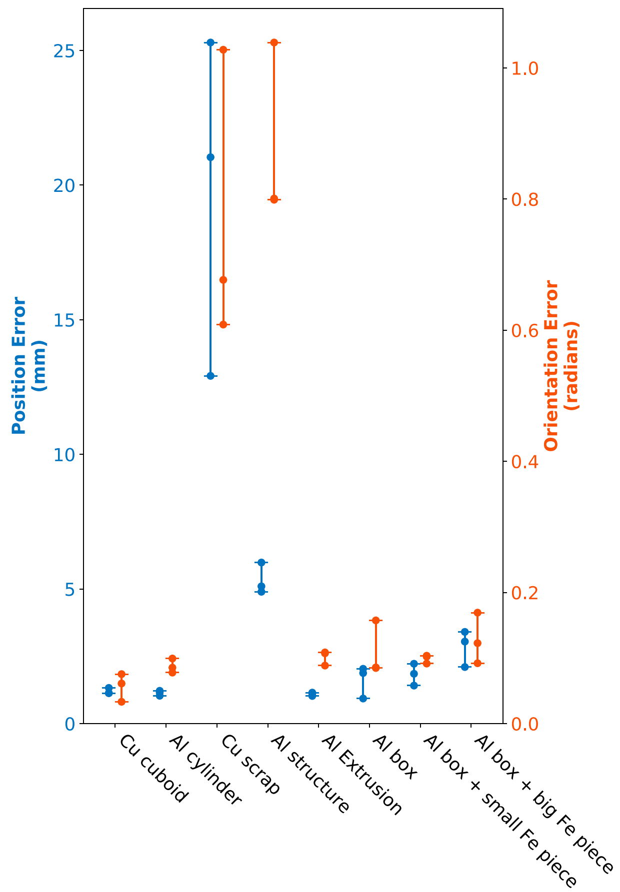

Figure 15. Tracking error for manipulating nonspherical objects along a square-trajectory. Points show average absolute error per trial, and bars show the range across all trials.

experiment described above with the copper scrap, we ran a trial manipulating the aluminum structure with the modified (slower) trajectory. In Figure 19—in which we use a different visualization where the object's orientation is encoded by its forward direction, indicated with a yellow arrow—we see that, much like with the copper scrap, the position accuracy improves and the orientation accuracy is now comparable to the other objects.

Finally, the aluminum-box manipulation trials have poor performance at the beginning of the trial when the object is in the middle of the workspace. We hypothesized that this is because our initialization is so poor for the object's parameters that the adaptation take a long time to find a good fit, and not because of control limits or the objects position in the state space. To test this hypothesis, we ran a trial of a more complex trajectory in the middle of our workspace where we initialize the object parameters to the median of the properties estimated during one of the previous trials. This trial is shown in Figure 20, and confirms our hypothesis. As expected, the worst tracking performance occurs when the system is initialized (see the wiggles in the "Block U"), with much lower error as the trajectory progresses, with the low-error portion comprising motions in every direction. If we directly compare with the similar

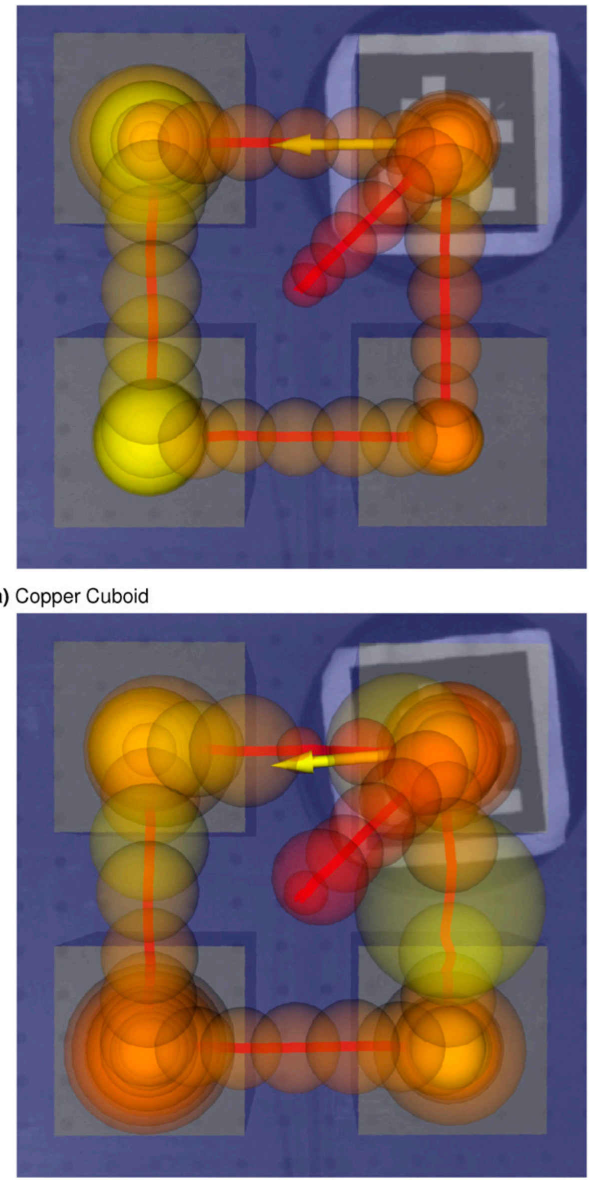

(a) Copper Cuboid

(b) Aluminum Extrusions

Figure 16. Characteristic trials manipulation (a) the copper cuboid and (b) the aluminum extrusions along a square trajectory. We visualize the radius and conductivity estimate at different points in the trajectory with periodically placed spheres. Note that no ground-truth radius exists for these nonspherical and anisotropic objects. Instead, the controller continuously adapts the radius and conductivity in order to provide a locally correct model of dynamics. See Supplemental Video.

copper-cuboid manipulation trial shown in Figure 1, we see quite comparable trajectory tracking.

In Figure 21, we quantitatively compare the tracking error in the new experiments from this section with the analogous original results of Section 5.7. We see drastically improved performance in all cases. In summary, to improve tracking performance, we should make our best attempt at initializing the object's parameters correctly, and we should slow down.

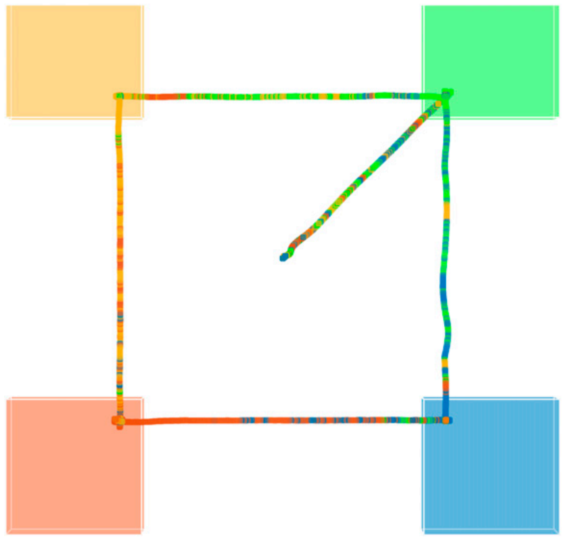

Figure 17. Characteristic trajectory manipulating the aluminum cylinder. At each timestep, we depict a color that matches the field source that is active at that instant.

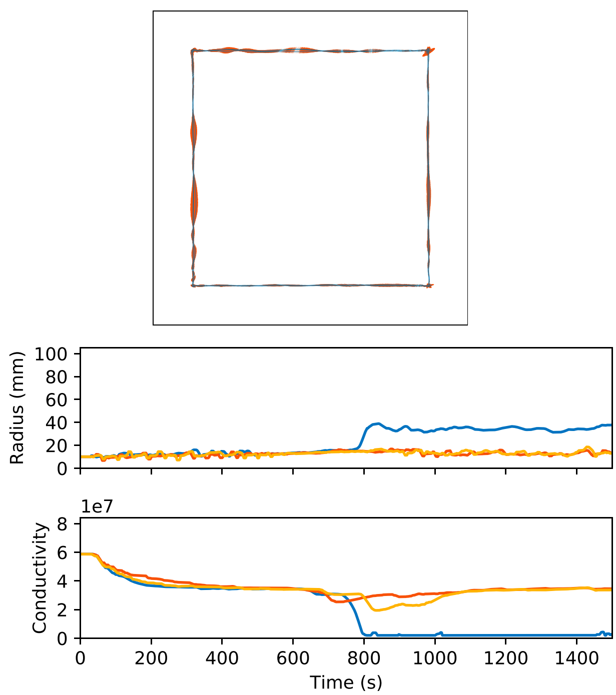

Figure 18. Three square-trajectory manipulation trials with the copper scrap, using the modified (i.e., slower) trajectory. See Supplemental Video.

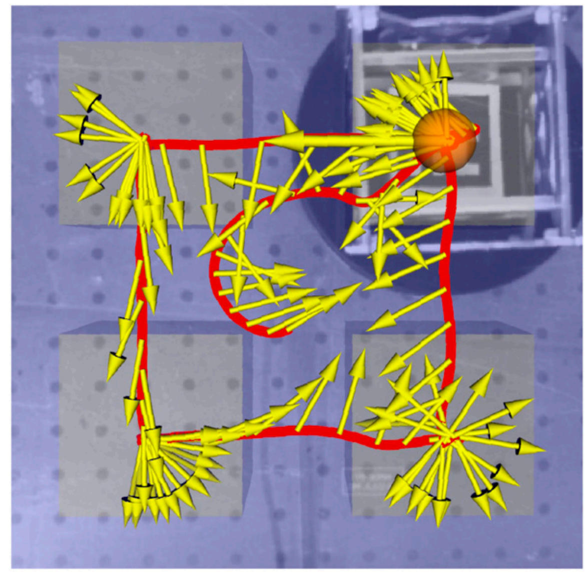

(a) Original Trajectory

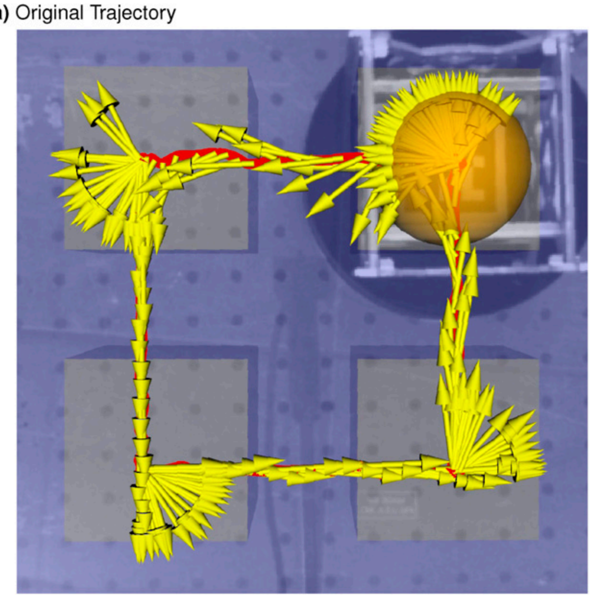

(b) Slower Trajectory

Figure 19. Trials manipulating the aluminum structure along a square trajectory. (a) Original trajectory of Figure 14(d), with alternative visualization. (b) Modified (i.e., slower) trajectory. See Supplemental Video.

6. Discussion

There are a number of open problems that would be interesting to explore in future work. In this work, as in all of our prior works, we only actuated one field source at a time, since our current models do not allow for superposition (we know that the induced wrenches from multiple field sources do not linearly superimpose). We expect that the proper application of superposition will lead to better-than-linear performance increases. Such a change in low-level field generation will not affect the methods described in this work.

In this work, we exclusively performed 6-DOF closed-loop pose-control tasks, but our wrench control policy is agnostic to how the desired force and torque are generated. One can apply our methodology in a space-debris detumbling task in which the desired force is generated by a 3-DOF closed-loop position regulator, and the desired torque is chosen to oppose, and thus reduce, the angular velocity of the object.

In this work, as in all of our prior dexterous-manipulation works, we only manipulate slowly moving objects, where

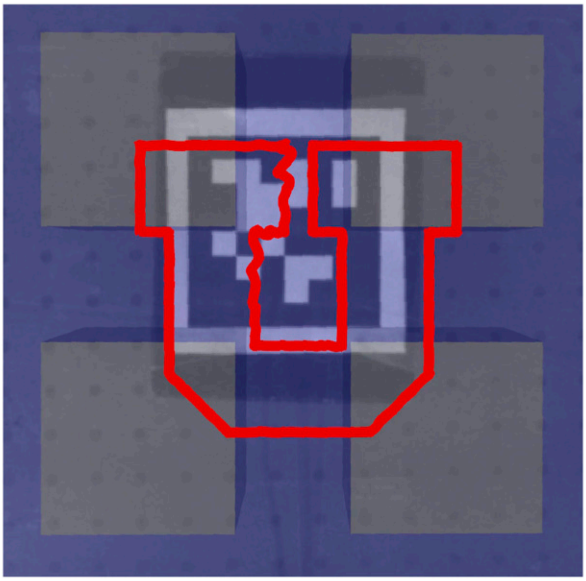

Figure 20. Trial manipulating the aluminum box along an alternative "Block U" trajectory, with strong initialization of object parameters.

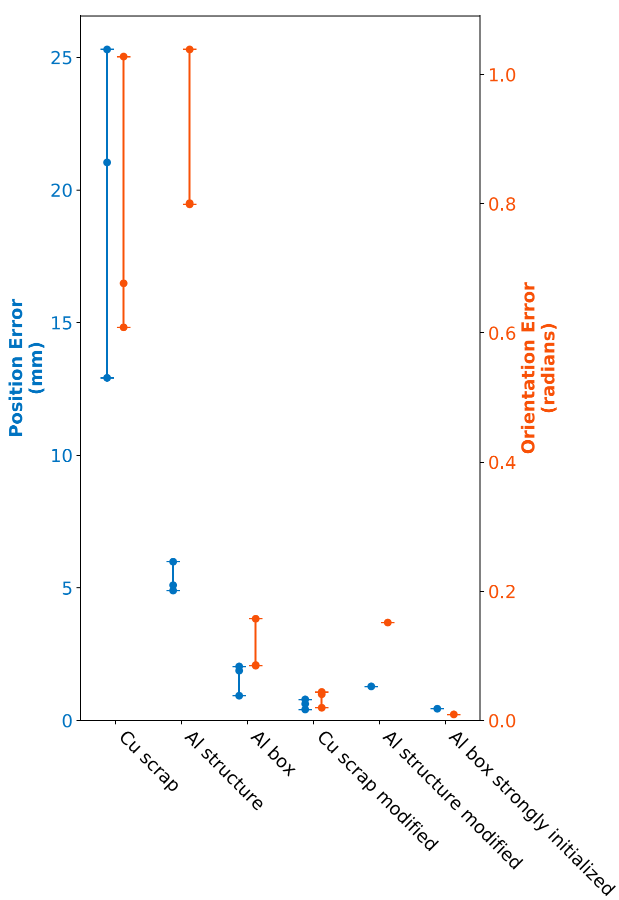

Figure 21. Tracking error for additional experiments, compared to the earlier results from Figure 15. Note: the copper-scrap and aluminum-structure trials have similar trajectories to the original, but the trial of the aluminum box with strong initialization is over a substantially different (and more complex) trajectory.

our quasistatic force–torque model is obviously applicable. There are rapidly rotating pieces of space debris where this might not be a valid assumption. Further research is needed on the change in induced force and torque due to object motion. We do know that object rotation will increase the induced braking torque compared to what is predicted by the quasistatic model (Allen et al., 2024).

In our targeted domain of space-debris remediation, the reaction force–torque on our field-generation system could be nontrivial to compensate for; accounting for these effects on our system is left as an open problem, but it will presumably require some combination of station keeping using the magnetic field sources (Abbott et al., 2017) and additional use of consumable gas thrusters.

In previous work, we showed that the model consistently under-predicated force–torque when an object's center was less than 1.5 times its radius away from the center of the dipole source, which we denoted the "near-field region." In this work, we build upon that model and never perform experiments in this near-field region where force–torques are dramatically higher. Additional modeling focused on larger objects might be necessary to accurately control in this region. This is likely worth doing, since it will result in more control authority once properly understood. Selecting optimal magnet configurations—whether fixed positions for a given task, or time-varying as part of the controller—will likely result in improved control authority as well.

In this work, we used a spherical-object model with a fixed density, with the result being that a change in the sphere's radius resulted in changes in the object's mass and moment of inertia that are coupled with the change in the induced wrench. We could decouple these effects by including the spherical-object's density as a third adaptive parameter, or by adaptively estimating the object's mass matrix directly.

In this work, we found that, for objects that are very distinct from a sphere, having a strong initialization of the spherical-object parameters can be very beneficial to manipulation quality. It may be worth pursing empirical models for some additional, nonspherical objects to this end. However, with so many potential geometries to consider, this may not be a good strategy. The modeling methods described here and in our prior work would be a reasonable starting point, with additional variables for relative orientations as other objects would lack as much symmetry as spheres. In addition to, or instead of, additional modeling, it may be worthwhile to use images (or other additional sensing modality) of objects to seed better initial object parameters and improve initial manipulation capabilities. Training such a system would likely require large amounts of manipulation data, likely provided by using our proposed system. Depending on the exact manipulation task, active learning, making control decisions specifically to help improve our understanding of the dynamics, might be helpful. This ties into ideas of persistent excitation.

Finally, in this work, as in our previous works on this topic, all of our physical manipulation experiments have been conducted using a water-based microgravity simulator with 3-DOF mobility, with 6-DOF experiments limited to our numerical microgravity simulator. Demonstrating physical 6-DOF manipulation is left as an open problem, but we have confidence that our force–torque models and manipulation methods are fundamentally 6-DOF in nature.

7. Conclusions

In this article, we expanded our previous model for the force–torque wrench induced on a conductive, nonmagnetic sphere by a rotating magnet dipole, which contained three discrete modes, to a continuous model that covers all possible relative positions of the sphere with respect to the rotating magnetic dipole. We leveraged this new model to examine manipulation of spherical objects with unknown physical parameters by applying techniques from the online-optimization and adaptive-control literature. Our experimental results validated our new dynamics model, showing that we get improved performance relative to the previous model, while also solving a simpler optimization problem for control. We further demonstrated the first physical magnetic manipulation of aluminum spheres, as previous controllers were only physically validated on copper spheres; aluminum is more relevant for the manipulation of engineered objects in space. We showed that our adaptive control framework can quickly acquire useful object parameters when weakly initialized. Finally, we demonstrated that the spherical-object model can be used as an approximate model for adaptive control of nonspherical objects by performing the first magnetic manipulation of a variety of nonspherical, nonmagnetic objects. We showed that multiple conductive objects—either rigidly connected to, or separated from, each other—can be approximated as a single sphere. Code, data, and videos associated with the experiments can be found at https://sites.google.com/gcloud.utah.edu/unknown-object-eddy-currents.

Acknowledgments

The authors would like to thank Rebecca Tabor for her technical assistance.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported in part by the National Science Foundation under Grants 1841845, 1846341, and 2149585. It was also supported in part by the AFRL/RVK under Contract No. FA9453-22-C-A043, which is a collaboration with Rogue Space Systems.

ORCID iDs

Griffin F Tabor https://orcid.org/0000-0003-3992-7391
Jake J Abbott https://orcid.org/0000-0002-0665-1934
Tucker Hermans https://orcid.org/0000-0003-2496-2768

Supplemental Material

Supplemental material for this article is available online.

References

Abbott JJ, Brink JB and Osting B (2017) Computing minimum-power dipole solutions for interdipole forces using nonlinear constrained optimization with application to electromagnetic formation flight. *IEEE Robotics and Automation Letters* 2(2): 1008–1014.

Abbott JJ, Diller E and Petruska AJ (2020) Magnetic methods in robotics. *Annual Review of Control, Robotics, and Autonomous Systems* 3: 57–90.

Aghili F (2020) Optimal trajectories and robot control for detumbling a non-cooperative satellite. *Journal of Guidance, Control, and Dynamics* 43(5): 981–988.

Allen TJ, Sperry AJ, Posselli NR, et al. (2024) Characterization of a rotating magnetic dipole field for contactless detumbling of space debris. In: International conference on space robotics, Luxembourg, 24–27 June 2024.

Atkeson CG, An CH and Hollerbach JM (1986) Estimation of inertial parameters of manipulator loads and links. *The International Journal of Robotics Research* 5(3): 101–119.

Boonrath A and Botta EM (2024) Validation of models for net deployment and capture simulation with experimental data. *Journal of Spacecraft and Rockets* 61(1): 181–199.

Carambia A, Frazell CG, Kapadia AD, et al. (2023) Non-cooperative space object capture and manipulation with soft robotics. In: IEEE aerospace conference, Big Sky, MT, 4–11 March 2023, 1–11.

Dalton DK, Tabor GF, Hermans T, et al. (2022) Attracting conductive nonmagnetic objects with rotating magnetic dipole fields. *IEEE Robotics and Automation Letters* 7(4): 11484–11491.

Dierckx P (1995) *Curve and Surface Fitting with Splines*. Oxford, UK: Clarendon Press.

Diller E, Giltinan J, Lum G, et al. (2016) Six-degree-of-freedom magnetic actuation for wireless microrobotics. *The International Journal of Robotics Research* 35(1–3): 114–128.

Frazelle C, Walker I, AlAttar A, et al. (2021) Kinematic-model-free control for space operations with continuum manipulators. In: IEEE aerospace conference, Big Sky, MT, 6–13 March 2021, 1–11.

Hazan E (2016) Introduction to Online Convex Optimization. *Found. Trends Optim.* 2(3–4): 157–325. Available at: https://doi.org/10.1561/2400000013.

Hertz H, Jones DE and Schott GA (1896) *Miscellaneous Papers, Chapter II. On Induction in Rotating Spheres*. London, UK: Macmillan.

Jiang L, Wang Y and Xu SH (2017) Integrated 6-DOF orbit-attitude dynamical modeling and control using geometric mechanics. *International Journal of Aerospace Engineering* 2017: 1–13.

Kessler DJ, Johnson NL, Liou JC, et al. (2010) The Kessler syndrome: implications to future space operations. *Advances in the Astronautical Sciences* 137(8): 1–15.

Liou JC and Johnson NL (2016) Risks in space from orbiting debris. *Science* 311: 340–341.

Liu X, Lu Y, Zhang Q, et al. (2019) An application of eddy current effect on the active detumble of uncontrolled satellite with tilt air gap. *IEEE Transactions on Magnetics* 55(12): 1–11.

Lynch KM and Park FC (2017) *Modern Robotics: Mechanics, Planning, and Control*. Cambridge, UK: Cambridge University Press.

Ma B, Jiang Z, Liu Y, et al. (2023) Advances in space robots for on-orbit servicing: a comprehensive review. *Advanced Intelligent Systems* 5(8): 2200397.

Mark CP and Kamath S (2019) Review of active space debris removal methods. *Space Policy* 47: 194–206.

Murphy KP (2021) *Machine Learning: A Probabilistic Perspective*. Cambridge, MA: MIT Press.

National Aeronautics and Space Administration (2021) *Astromaterials Research & Exploration Science: Orbital Debris Program Office*. Washington, DC: National Aeronautics and Space Administration. https://orbitaldebris.jsc.nasa.gov/ (Accessed: 2021-07-08).

Nocedal J and Wright SJ (2006) *Numerical Optimization*. 2nd edition. Berlin, Germany: Springer.

Nurge MA, Youngquist RC and Starr SO (2018) Drag and lift forces between a rotating conductive sphere and a cylindrical magnet. *American Journal of Physics* 86(6): 443–452.

Opiela JN (2009) A study of the material density distribution of space debris. *Advances in Space Research* 43(7): 1058–1064.

Ortiz Gómez N and Walker SJI (2015) Eddy currents applied to detumbling of space debris: analysis and validation of approximate proposed methods. *Acta Astronautica* 114: 34–53.

Papadopoulos E, Aghili F, Ma O, et al. (2021) Robotic manipulation and capture in space: a survey. *Frontiers in Robotics and AI* 8: 686723–686736.

Petruska AJ and Abbott JJ (2014) Omnimagnet: an omnidirectional electromagnet for controlled dipole-field generation. *IEEE Transactions on Magnetics* 50(7): 1–10.

Pham LN, Tabor GF, Pourkand A, et al. (2021) Dexterous magnetic manipulation of conductive non-magnetic objects. *Nature* 598(7881): 439–443.

Qian N (1999) On the momentum term in gradient descent learning algorithms. *Neural Networks* 12(1): 145–151.

Ratliff N, Meier F, Kappler D, et al. (2016) DOOMED: direct online optimization of modeling errors in dynamics. *Big Data* 4(4): 253–268.

Reinhardt BZ and Peck MA (2016) New electromagnetic actuator for on-orbit inspection. *Journal of Spacecraft and Rockets* 53(2): 241–248.

Shan M, Guo J and Gill E (2016) Review and comparison of active space debris capturing and removal methods. *Progress in Aerospace Sciences* 80: 18–32.

Siciliano B, Sciavicco L, Villani L, et al. (2009) *Robotics: Modelling, Planning and Control*. Berlin, Germany: Springer-Verlag.

Slotine J-JE and Li W (1991) *Applied Nonlinear Control*. Englewood Cliffs, NJ: Prentice Hall.

Smith YR, Nagel JR and Rajamani RK (2019) Eddy current separation for recovery of non-ferrous metallic particles: a comprehensive review. *Minerals Engineering* 133: 149–159.

Sugai F, Abiko S, Tsujita T, et al. (2013) Detumbling an uncontrolled satellite with contactless force by using an eddy current brake. In: IEEE/RSJ international conference on intelligent robots and systems, Tokyo, Japan, 3–7 November 2013, 783–788.

Tabor GF, Pham LN, Abbott JJ, et al. (2022) Adaptive manipulation of conductive, nonmagnetic objects via a continuous model of magnetically induced force and torque. In: Robotics: science and systems (RSS), New York, NY, 27 June–1 July 2022.

Vijayan R, De Stefano M and Ott C (2022) A detumbling strategy for an orbital manipulator in the post-grasp phase. In: IEEE international conference on robotics and automation, Philadelphia, PA, 23–27 May 2022, 4086–4092.

Volpe R and Circi C (2019) Optical-aided, autonomous and optimal space rendezvous with a noncooperative target. *Acta Astronautica* 157: 528–540.

Wilson KT and Peck MA (2020) Electromagnetic, free-flying mobility relative to a conductive body. In: AIAA scitech forum, Orlando, FL, 6–10 January 2020.

Wilson KT and Peck MA (2021) Electromagnetic actuation for non-contact, propellant-free, six-degree-of-freedom relative mobility. In AIAA scitech forum, Virtual Event, 11–15 January 2021.

Wilson KT, Bernal G and Peck MA (2023) Eddy-current actuator for attraction and repulsion of non-ferromagnetic, conductive spacecraft. *Journal of Spacecraft and Rockets* 60(4): 1063–1071.

Zeng C, Hecht GR, Chowdhury S, et al. (2024) Concurrent design optimization of tether-net system and actions for reliable space-debris capture. *Journal of Spacecraft and Rockets* 61(3): 773–783.