

Adaptive Magnetic Control using Stein Variational Gradient Descent computed Distribution of Object Parameters

Griffin F. Tabor^{1,2}, Tucker Hermans^{1,2,3}

I. INTRODUCTION

Traditional magnetic manipulation methods are severely limited in what they can manipulate. They rely on an object’s ferromagnetic properties, which are only present in a limited set of materials. While many engineering materials are not ferromagnetic, they are electrically conductive. It well understood that electrically conductive material exposed to a time-varying magnetic field results in eddy currents on the surface of the material[1]. The induced flow of electrons creates their own magnetic field, which in turn reacts with the original magnetic field. The interaction between these two fields induces force and torques on the conductive object.

We have recently begun investigating using eddy currents to enable contact free manipulation to clean up space debris. A recent study found that because of the current quantity of objects in orbit, “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” ([2]). These hazardous pieces of debris are largely composed of aluminum ([3]), which is not ferromagnetic, but is very conductive.

Towards this goal, in [4] and [5], we developed a model of the force and torque induced from rotating a dipole source near a solid conductive object. In [4] we showed that, leveraging this model and an array of fixed-in-position electromagnets producing rotating magnetic fields, it was possible to perform full six-degree-of-freedom manipulation of a conductive nonmagnetic sphere. Further, in [5] we showed that adaptive control let us use the spherical object model to control unknown spheres and even generalize to some nonspherical objects.

In this paper, we apply ideas from [6] to view online system identification as a variational inference problem. We approximately solve this variational inference problem using Stein Variational Gradient Descent [7], a recently proposed non-parametric Bayesian inference algorithm. We propose the following contributions to adaptive magnetic manipulation of nonmagnetic objects

- We propose an inverse dynamics controller to better take advantage of online dynamic parameter estimation
- We propose fitting a non-parametric distribution over possible object parameters instead of single best parameter
- We propose a controller that finds the best control in expectation under the object parameter distribution

II. REVIEW OF MODEL

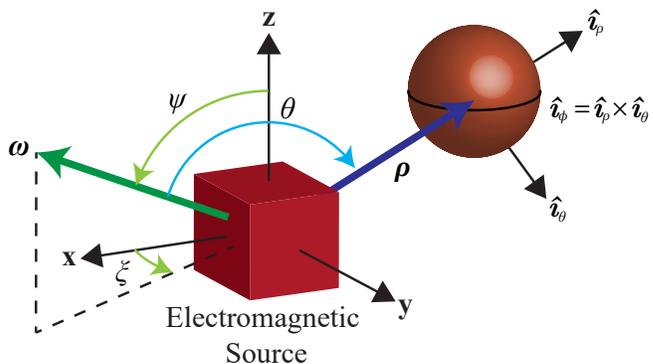


Fig. 1: Figure reproduced from [5]. Spherical coordinate systems describing the dipole rotation vector ω with respect to the world frame, and the conductive sphere with respect to ω .

The force torque model from [4] models two specific configurations of relative magnet rotation and object direction. This model can be seen in Eq. 1 with coefficients from Table. I to compute the nonzero elements of the force-torque vector. Note in the $\theta = 0$ configuration there is only force and torque in a single direction, and in $\theta = 90^\circ$ configuration there is a force component in two axes and only a torque component in the third. It is important to note the force-torques are expressed in a spherical frame shown in Fig. 1 uniquely defined by $\hat{\omega}$ and $\hat{\rho}$.

- σ : conductivity of the object in units S/m
- r : radius of the object in meters
- ρ : vector from the magnetic dipole source to the object in meters. Which can be expressed as unit vector direction $\hat{\rho}$ and distance ρ
- μ_0 known permeability of free space
- m : dipole strength in units A · m²

*This work was supported by the National Science Foundation under Grant 2149585.

¹University of Utah Robotics Center, Salt Lake City, UT, USA

²Kahlert School of Computing, University of Utah, Salt Lake City, UT, USA

³NVIDIA, Seattle, WA, USA

- $\boldsymbol{\omega}$: three vector rotational velocity of the dipole source in units 1/s. Which can be expressed as unit vector direction $\hat{\boldsymbol{\omega}}$ and frequency ω

TABLE I: Coefficients from [4] for model in Eq. (1). These values are the average of the FEA and physical values as suggested in [5]

Coefficients							
θ	f, τ	c_0	c_1	c_2	c_3	c_4	c_5
0°	f_ρ	448.5	2.88	-0.0989	-9.505	7	4
0°	τ_ρ	6870	3.175	-0.0988	-14.05	6	3
90°	f_ρ	274	2.90	-0.0995	-8.53	7	4
90°	f_ϕ	5955	3.47	-0.0996	-14.45	7	4
90°	τ_θ	8050	3.50	-0.0957	-15.35	6	3

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

In the world frame, the spherical frame is defined as

$$\begin{aligned} \hat{\mathbf{i}}_\rho &= \frac{\boldsymbol{\rho}}{\rho} \\ \hat{\mathbf{i}}_\phi &= \frac{\hat{\boldsymbol{\omega}} \times \boldsymbol{\rho}}{\|\hat{\boldsymbol{\omega}} \times \boldsymbol{\rho}\|} \\ \hat{\mathbf{i}}_\theta &= \hat{\mathbf{i}}_\phi \times \hat{\mathbf{i}}_\rho \end{aligned} \quad (2)$$

Further, arbitrary configurations can be computed using trigonometric interpolation between our known configurations as a function of θ the angle between $\hat{\boldsymbol{\omega}}$ and $\boldsymbol{\rho}$ [5].

$$f_\rho(\rho, \theta) = - \left(\frac{f_\rho(\rho, 90^\circ) - f_\rho(\rho, 0^\circ)}{2} \right) \cos(2\theta) + \left(\frac{f_\rho(\rho, 90^\circ) + f_\rho(\rho, 0^\circ)}{2} \right) \quad (3)$$

$$f_\theta(\rho, \theta) \approx 0 \quad (4)$$

$$f_\phi(\rho, \theta) = f_\phi(\rho, 90^\circ) \sin(\theta) \quad (5)$$

$$\tau_\rho(\rho, \theta) = \tau_\rho(\rho, 0^\circ) \cos(\theta) \quad (6)$$

$$\tau_\theta(\rho, \theta) = \tau_\theta(\rho, 90^\circ) \sin(\theta) \quad (7)$$

$$\tau_\phi(\rho, \theta) = 0 \quad (8)$$

We assume we have a workspace surrounded by electromagnets that can each produce a magnetic dipole in arbitrary directions. Instead of reasoning about the direction of a magnetic dipole we abstract away the current direction of the dipole and consider dipoles rotating about an axis $\hat{\boldsymbol{\omega}}$ with frequency ω . As in [5] and [4] we select a fixed frequency of 15 HZ and control the magnets by selecting m and $\hat{\boldsymbol{\omega}}$.

For simplicity let $w(\mathbf{x}, \boldsymbol{\lambda}, \{i, m, \psi, \xi\})$, be the force-torque in the world frame given object pose \mathbf{x} , object properties $\boldsymbol{\lambda}$ (radius and conductivity), and control parameters (magnet id, dipole strength and two angles to parameterize $\hat{\boldsymbol{\omega}}$). In [5] and [4] we proposed solving a constrained optimization problem to find the control actions that are closest to a desired force-torque, as the model has no closed form inverse. Importantly, as the model is only accurate for a single rotating dipole source and a manipulation environment requires multiple dipole sources around the workspace, this optimization also searches over which of the available sources to use.

$$\begin{aligned} \operatorname{argmin}_{i, m, \psi, \xi} & \left\| \begin{bmatrix} \mathbf{f}_{\text{des}} \\ \boldsymbol{\tau}_{\text{des}} \end{bmatrix} - \begin{bmatrix} \mathbf{f} \\ \boldsymbol{\tau} \end{bmatrix} \right\|_Q^2 \\ \text{s.t.} & \quad i \in \{1, \dots, n\} \\ & \quad m \in [0, m_{\max}] \\ & \quad \psi \in [0, \pi] \\ & \quad \xi \in [-\pi, \pi] \\ & \quad \mathbf{f}, \boldsymbol{\tau} = w(\mathbf{x}, \boldsymbol{\lambda}, \{i, m, \psi, \xi\}) \end{aligned} \quad (9)$$

This mixed integer optimization problem is solved with parallelized Newton's method solver, with two initializations for each of the n magnets.

III. STEIN VARIATIONAL GRADIENT DESCENT

Variational inference attempts to approximate unknown distributions $p(z)$ [8]. This is done by selecting a family of distributions Q and finding the member $q^*(z)$ that closest matches the target distribution $p(z)$. Formally, this is expressed as

$$q^*(z) = \operatorname{argmin}_{q(z) \in Q} f(q(z), p(z)) \quad (10)$$

Where $f(q, p)$ is some statistical divergence between two distributions, often Kullback-Leibler (KL) divergence [9] [10] [11]. We use $p(z)$ as shorthand for $p(z|O)$ where we have a set O of observations.

$$p(z|O) = \frac{1}{K} \bar{p} \quad (11)$$

$$\bar{p}(z) = p_0(z) \prod_{o \in O} p(o|z) \quad (12)$$

$$K = \int \bar{p}(z) dz \quad (13)$$

The main strength of the variational approach is that gradient based optimization scales to the large number of observations that are necessary to compute $p(z)$ in practice. One main weakness is that it depends upon a good choice of Q , if you pick a family of distributions with too much capacity you will have difficulty solving the optimization and a family with too little capacity won't approximate the true distribution well.

Stein Variational Gradient Descent (SVGD) uses a set of particles as an implicit distribution q . Instead of updating the underlying parameters of q with the gradient of the KL divergence, SVGD updates the particles directly. It defines a one-to-one distribution transformation T_ϵ such that $T_\epsilon(x) = x + \epsilon g(x)$ [7]. This transform is a single step in the iterative algorithm. The goal then is to find a transform that will bring q closer to p , and thus minimizes $KL(q_{[T]} || p)$.

Note, by definition, a transformation on q and an inverse transformation on p has the same effect on KL divergence.

$$KL(q_{[T]} || p) = KL(q || p_{[T^{-1}]}) \quad (14)$$

By taking the gradient of KL divergence with respect to epsilon, we have a function for finding how transformation

directions will affect KL divergence.

$$\nabla_{\varepsilon} \text{KL} \left(q || p_{[T-1]} \right) = -\mathbb{E}_{z \sim q} \left[\begin{array}{l} \nabla_{\varepsilon} \log p(z + g(z))^{\top} g(z) \\ + \text{trace} \left((I + \varepsilon g(z))^{-1} \cdot \nabla_z g(z) \right) \end{array} \right]$$

If we evaluate this gradient at $\varepsilon = 0$ we get $-\mathbb{E}_{z \sim q} [\nabla_{\varepsilon} \log p(z)^{\top} g(z) + \text{trace}(\nabla_z g(z))]$. It follows if we search for g in some functional family G that maximizes this quantity it will be the steepest descent direction, within G , to improve the KL divergence.

$$g^*(x) = \underset{g(z) \in G}{\text{argmax}} \quad -\mathbb{E}_{z \sim q} \left[\nabla_{\varepsilon} \log p(z)^{\top} g(z) + \text{trace}(\nabla_z g(z)) \right] \quad (15)$$

Liu et al. [12] shows if we select the family G to be functions within the unit ball of the reproducing kernel Hilbert Space, for a kernel function k , this optimization has a closed form solution.

$$g^* = \mathbb{E}_{z \sim q} [k(z) \nabla_z \log p(z) + \nabla_z k(z)] \quad (16)$$

SVGD approximates this expectation using a finite set of particles and updates them incrementally using the functional gradient g^* . This non-parametric form provides more capacity and flexibility than a fixed family of analytic distributions.

$$z_i^{l+1} \leftarrow z_i^l + \varepsilon g^*(z_i^l) \quad (17)$$

$$g^*(z) \approx \frac{1}{n} \sum_j^n \left[k(z_j^l, z) \nabla_{z_j^l} \log \left(p(z_j^l) \right) + \nabla_{z_j^l} k(z_j^l, z) \right] \quad (18)$$

There is an intuitive explanation for the terms in the update equation, the first term is the weighted average of the gradient of nearby particles and the second term is a repulsive term pushing particles away from each other.

IV. ADAPTIVE CONTROL

In [5] we demonstrated an adaptive controller able to manipulate unknown metal spheres. This controller can also generalize to nonspherical objects by approximating them online as a sphere with changing size and conductivity. The adaptive controller attempts to solve the following system identification optimization problem using a minibatch of the 25 most recent timesteps ([13]).

$$\mathcal{L}_a(\boldsymbol{\lambda}, b) = \left\| \left(\frac{\dot{\mathbf{x}}[b+1] - \dot{\mathbf{x}}[b]}{\delta t} \right) - M(\boldsymbol{\lambda})^{-1} w(\mathbf{x}[b], \boldsymbol{\lambda}, \boldsymbol{\eta}[b]) \right\|_{\mathcal{Q}}^2, \quad (19)$$

$$\boldsymbol{\lambda}^* = \underset{\boldsymbol{\lambda}}{\text{argmin}} \sum_{b=1}^B \mathcal{L}_a(\boldsymbol{\lambda}, b) \quad (20)$$

Where M is the mass matrix of a sphere based on current object parameters and a fixed density (that of copper). Each timestep the adaptive controller took a single momentum step, with backtracking line search, enforcing the parameters smoothly varied over time. In [5], we used the control optimization in Eq. 9 with the most up-to-date object parameters.

In this paper, we replace this force-torque optimization with an inverse dynamics control optimization. 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 Eq. 9:

$$\underset{i, m, \psi, \xi}{\text{argmin}} \quad \left\| \begin{bmatrix} \mathbf{a}_{\text{des}} \\ \boldsymbol{\alpha}_{\text{des}} \end{bmatrix} - M(\boldsymbol{\lambda})^{-1} \begin{bmatrix} \mathbf{f} \\ \boldsymbol{\tau} \end{bmatrix} \right\|_{\mathcal{Q}}^2 \quad (21)$$

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

This improved Adaptive controller will serve as our baseline method.

V. METHODS

In this section, we will present our updated control method. First, we will explain how we use SVGD to fit a distribution of possible object parameters instead of only trusting a single estimate. Second, we will show how we use the object parameter distribution to find the best control action in expectation.

A. Stein Variational

Similar to [6], we construct a likelihood function of our dynamic parameters conditioned on our observed motion. However, instead of using the previous timesteps distribution as a prior and only a single observation as the likelihood, we use a minibatch of the most recent B observations with a continuous uniform prior over the object properties. Let $O_{1:B}$ be the B most recent observed accelerations and control actions. Because we can not actually sample object parameters, we map values for $\boldsymbol{\lambda}$ into accelerations space using our dynamic equation and control parameters. We assume the observations for each timestep are only dependent on the object properties and that the true underlying dynamics function has additive Gaussian noise with zero mean.

$$l(\boldsymbol{\lambda} | O_{1:B}) = \prod_{b=1}^B p(O_b | \boldsymbol{\lambda}) \quad (22)$$

$$= \frac{1}{Z} \prod_{b=1}^B \exp \mathcal{L}_a(\boldsymbol{\lambda}, b) \quad (23)$$

Where \mathcal{Q} is the precision of the observation noise and Z is a normalization constant to keep the total probability mass equal to one. Note, if we attempted to find the MAP (Maximum A Posteriori) estimate for l by minimizing the negative log likelihood, we recover exactly our system identification loss plus a constant, $\log(l(\boldsymbol{\lambda}) | O_{1:B}) = \sum_{b=1}^B \mathcal{L}_a(\boldsymbol{\lambda}, b) - \log(Z)$. The constant is unimportant for the optimization.

Within the region of our continuous uniform prior $p(\boldsymbol{\lambda} | O_{1:B}) = l(\boldsymbol{\lambda} | O_{1:B})/a$ where $1/a$ is the prior density. We use SVGD to approximate this entire distribution. We never need to actually create p explicitly, as we only need the gradient of \mathcal{L}_a . We use k as the RBF kernel with the median

heuristic as chosen by [7].

$$k(z, z') = \exp\left(-\frac{1}{h}\|z - z'\|^2\right) \quad (24)$$

$$\nabla_z k(z, z') = -\frac{1}{h}\exp\left(-\frac{1}{h}\|z - z'\|^2\right)(z - z') \quad (25)$$

We sample our initial distribution from the continuous uniform prior distribution and use projection based bound constraints to keep every particle within the support of p .

B. Control in Expectation

We select the control action that is the best in expectation with respect to our distribution of object parameters.

$$\begin{aligned} \operatorname{argmin}_{i, m, \psi, \xi} \quad & \mathbb{E}_{\lambda \sim q} \left\| \begin{bmatrix} \mathbf{a}_{\text{des}} \\ \boldsymbol{\alpha}_{\text{des}} \end{bmatrix} - M(\boldsymbol{\lambda})^{-1} \begin{bmatrix} \mathbf{f} \\ \boldsymbol{\tau} \end{bmatrix} \right\|_{\mathbf{Q}}^2 \\ \text{s.t.} \quad & i \in \{1, \dots, n\}, m \in [0, m_{\text{max}}], \\ & \psi \in [0, \pi], \xi \in [-\pi, \pi], \\ & \mathbf{f}, \boldsymbol{\tau} = w(\mathbf{x}, \boldsymbol{\lambda}, \{i, m, \psi, \xi\}) \end{aligned} \quad (26)$$

Which given J particles in q we can approximate as

$$\begin{aligned} \operatorname{argmin}_{i, m, \psi, \xi} \quad & \frac{1}{J} \sum_j \left\| \begin{bmatrix} \mathbf{a}_{\text{des}} \\ \boldsymbol{\alpha}_{\text{des}} \end{bmatrix} - M(\boldsymbol{\lambda}_j)^{-1} \begin{bmatrix} \mathbf{f} \\ \boldsymbol{\tau} \end{bmatrix} \right\|_{\mathbf{Q}}^2 \\ \text{s.t.} \quad & i \in \{1, \dots, n\}, m \in [0, m_{\text{max}}], \\ & \psi \in [0, \pi], \xi \in [-\pi, \pi], \\ & \mathbf{f}, \boldsymbol{\tau} = w(\mathbf{x}, \boldsymbol{\lambda}_j, \{i, m, \psi, \xi\}) \end{aligned} \quad (27)$$

VI. EXPERIMENTS

To test our method compared to the baseline adaptive control method, we constructed a six DOF numerical simulator with a manipulation workspace surrounded by four magnets. We compared both methods' performance on a regulation task where a time parameterized cubic polynomial generated position targets, smoothly moving the target object from the initial pose to the origin. We use a PD controller to convert position and velocity errors in six DOF to desired accelerations.

Our baseline method uses Eq. 19 to iteratively find the single best object parameter to fit the 5 most recent observations, and Eq. 21 to solve for optimal control actions.

Our proposed method uses SVGD to iteratively find a distribution of object parameters to iteratively fit the 5 most recent observations and Eq. 21 to solve for optimal control actions in expectation under the object parameter distribution.

For both controllers, we take a single momentum step, with backtracking line search, using identical parameters to fit the object parameters every timestep. For the SVGD parameter adaptation, the backtracking cost is the sum over the negative log likelihood for all particles. For both controller's parameter adaptation, we project onto the bound constraints $r \in [0.001, 0.1]$ meters, $\sigma \in [2e6, 1.2e8]$ S/m.

While we can not truly simulate manipulation of nonspherical objects, our goal is to test the robustness of our methods against tasks that are reflective of what we might

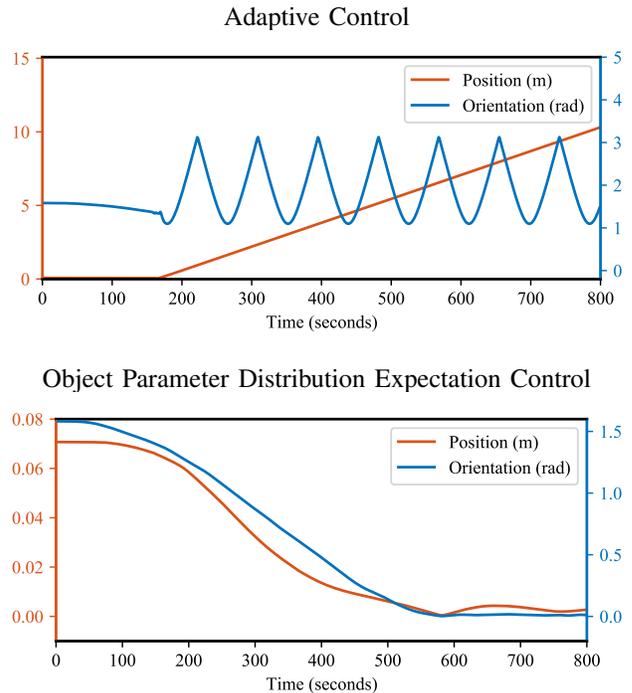


Fig. 2: Errors from six DOF regression task moving object from arbitrary pose to origin with stochastic dynamics. We see the adaptive control baseline loses control over the object at approximately 160 seconds into the trajectory while the proposed method still achieves the task.

expect when applying our spherical object dynamics model to nonspherical objects. We expect nonspherical objects can not be globally modeled by our model. However, because the force-torques produced are coupled, we expect the wrench for a given timestep to map well onto our model for some unknown object properties. To test our methods' robustness to time varying object properties, our simulator samples its internal radius and conductivity from independent Gaussian distributions with means 0.05m and $4e7$ S/m and standard deviations 0.009m and $9e6$ S/m respectively.

The baseline adaptive controller is initialized with the true mean of the simulator's object property distribution, while the SVGD adaptive controller is initialized uniformly over the possible parameter space.

VII. RESULTS AND CONCLUSIONS

We see in Fig. 2 that having a distribution of control parameters is much more robust to extreme stochastic dynamics functions. We believe the improved robustness to wildly changing dynamic parameters in simulation will extend to improved generalization to unmodeled objects in the future. It is also encouraging for the baseline method because we had to make the simulator distribution impressively large before the baseline was unstable. With a deterministic simulator or a much tighter simulator distribution, the results were indistinguishable.

REFERENCES

- [1] H. Hertz, *Miscellaneous papers*. London: Macmillan, 1896, ch. II. On induction in rotating spheres, english translation by D. E. Jones and G. A. Schott.
- [2] National Aeronautics and Space Administration, "Astromaterials Research & Exploration Science: Orbital Debris Program Office," <https://orbitaldebris.jsc.nasa.gov/>, 2021, accessed: 2021-07-08.
- [3] J. N. Opiela, "A study of the material density distribution of space debris," *Adv. Space Res.*, vol. 43, no. 7, pp. 1058–1064, 2009.
- [4] L. N. Pham, G. F. Tabor, A. Pourkand, J. L. B. Aman, T. Hermans, and J. J. Abbott, "Dexterous magnetic manipulation of conductive non-magnetic objects," *Nature*, vol. 598, no. 7881, pp. 439–443, 10 2021. [Online]. Available: <https://rdcu.be/czSK5>
- [5] G. F. Tabor, L. N. Pham, J. J. Abbott, and T. Hermans, "Adaptive Manipulation of Conductive, Nonmagnetic Objects via a Continuous Model of Magnetically Induced Force and Torque," in *Robotics: Science and Systems (RSS)*, 2022. [Online]. Available: <https://sites.google.com/gcloud.utah.edu/adaptiveeddycurrent/>
- [6] L. Barcelos, A. Lambert, R. Oliveira, P. Borges, B. Boots, and F. Ramos, "Dual online stein variational inference for control and dynamics," *CoRR*, vol. abs/2103.12890, 2021. [Online]. Available: <https://arxiv.org/abs/2103.12890>
- [7] Q. Liu and D. Wang, "Stein variational gradient descent: A general purpose bayesian inference algorithm," 2019.
- [8] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, "Variational inference: A review for statisticians," *Journal of the American Statistical Association*, vol. 112, no. 518, pp. 859–877, apr 2017. [Online]. Available: <https://doi.org/10.1080%2F01621459.2017.1285773>
- [9] S. Kullback and R. A. Leibler, "On information and sufficiency," *The annals of mathematical statistics*, vol. 22, no. 1, pp. 79–86, 1951.
- [10] M. Yi and S. Liu, "Sliced wasserstein variational inference," 2022.
- [11] T. Campbell and X. Li, "Universal boosting variational inference," 2019.
- [12] Q. Liu, J. D. Lee, and M. I. Jordan, "A kernelized stein discrepancy for goodness-of-fit tests and model evaluation," 2016.
- [13] C. G. Atkeson, C. H. An, and J. M. Hollerbach, "Estimation of Inertial Parameters of Manipulator Loads and Links," *Int. J. Robot. Res.*, vol. 5, no. 3, pp. 101–119, 1986.