

ViTa-Zero: Zero-shot Visuotactile Object 6D Pose Estimation

Hongyu Li^{1,2†}, James Akl¹, Srinath Sridhar^{1,2}, Tye Brady¹, Taşkın Padır^{1,3‡}

Abstract—Object 6D pose estimation is a critical challenge in robotics, particularly for manipulation tasks. While prior research combining visual and tactile (visuotactile) information has shown promise, these approaches often struggle with generalization due to the limited availability of visuotactile data. In this paper, we introduce *ViTa-Zero*, a zero-shot visuotactile pose estimation framework. Our key innovation lies in leveraging a visual model as its backbone and performing feasibility checking and test-time optimization based on physical constraints derived from tactile and proprioceptive observations. Specifically, we model the gripper-object interaction as a spring-mass system, where tactile sensors induce attractive forces, and proprioception generates repulsive forces. We validate our framework through experiments on a real-world robot setup, demonstrating its effectiveness across representative visual backbones and manipulation scenarios, including grasping, object picking, and bimanual handover. Compared to the visual models, our approach overcomes some drastic failure modes while tracking the in-hand object pose. In our experiments, our approach shows an average increase of 55% in AUC of ADD-S and 60% in ADD, along with an 80% lower position error compared to FoundationPose.

I. INTRODUCTION

For robots to intelligently manipulate objects like humans, they must perceive the object’s state within their environment accurately. A key aspect of this perception is object 6D pose estimation, *i.e.*, estimating position and orientation in 3D space, which serves as a critical representation of object state [1]. Accurate pose estimation can improve the performance of state-based manipulation policies and lead to improved learning efficiency and success rate [2–8]. Prior research has focused on instance-level [9–12], category-level [13–15] and more recently novel object estimation [16–20]. Despite these advancements, state-based approaches still face significant challenges in practical, real-world applications compared to end-to-end visuomotor approaches [2, 3, 21, 22]. This is due to the challenges of pose estimation in real-world scenarios, particularly during contact-rich and in-hand manipulation tasks characterized by frequent occlusions and dynamic interactions [5, 23].

To better estimate the pose of in-hand objects, previous works have explored the combination of visual and tactile (visuotactile) sensory observations [23–31]. These works

[†]The work was done when Hongyu Li was an intern at Amazon. Corresponding to: hli230@cs.brown.edu

[‡]Taşkın Padır holds concurrent appointments as a Professor of Electrical and Computer Engineering at Northeastern University and as an Amazon Scholar. This paper describes work performed at Amazon and is not associated with Northeastern University.

¹Amazon Fulfillment Technologies & Robotics, Westborough, MA, 01581.

²Brown University, Providence, RI, 02912.

³Northeastern University, Boston, MA, 02120.

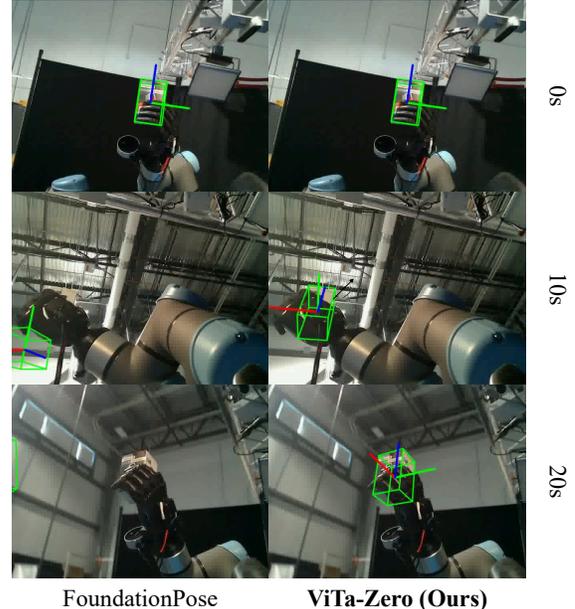


Fig. 1: FoundationPose [16] (left) fails due to errors that are acute (*e.g.*, occlusions) or accumulative (*e.g.*, noise) while tracking the in-hand object. Our approach (right) leverages tactile and proprioceptive observations for stable tracking.

have shown improvements in grasping and in-hand manipulation tasks, as tactile sensors can reveal occluded parts of objects to enhance pose estimation during manipulation. However, this introduces some practical challenges: (1) collecting a visuotactile dataset is arduous due to the fragility and diversity of tactile sensors, making it hard to scale in the real world [32]. While simulation offers a potential solution, the sim-to-real gap remains substantial [7, 25, 26, 33]; (2) visuotactile models often overfit to a single hardware setup and scenarios involving rich contact signals, such as in-hand object pose estimation, and struggle to perform on par with the visual counterparts for static and unoccluded scenes; and, (3) adopting the generalization techniques used for visual models for these visuotactile models is non-trivial due to the inclusion of tactile data and the non-uniformity of tactile sensors and grippers. Therefore, most models are sensor-specific and instance-level, limiting their practicality.

In this paper, we address these challenges by proposing **ViTa-Zero**, a framework that estimates and tracks the 6D pose of the novel object through the *zero-shot integration* of visual and tactile information. Our framework can leverage any capable visual model as its backbone, allowing it to

inherit the generalization capabilities of models trained on large-scale visual datasets. During our experiments, we found that visual models often fail under the adverse visual conditions faced during object manipulation (Fig. 1). To counter this, we introduce a feasibility check on the visual estimates based on three physical constraints: (1) *contact constraint*: the object must make contact with the tactile sensor if a positive contact signal is detected; (2) *penetration constraint*: at any time, the object cannot penetrate the robot (per its proprioception); and (3) *kinematic constraint*: the object’s motion must be physically feasible.

When visual estimation is deemed infeasible based on these constraints, we refine it through a test-time optimization process using tactile and proprioceptive information. Specifically, we model the gripper–object interaction as a spring–mass system, where the object is pulled by a virtual “attractive spring” towards the tactile sensors according to the tactile signals. Concurrently, a “repulsive spring” repels the object to avoid penetration with the robot. This optimization process allows for robust, real-time tracking even in dynamic, occluded scenarios.

To demonstrate effectiveness, we apply our framework onto two representative novel object visual estimation models [16, 17], conducting extensive experiments on a real-world robotic setup without collecting any tactile dataset for fine-tuning. Our experiments show that ViTa-Zero outperforms the underlying visual models by a large margin. Our contributions could be summarized as follows:

- 1) We present the first zero-shot visuotactile novel object pose estimation framework, greatly improving the practicality within the visuotactile domain.
- 2) We propose a test-time optimization approach with both efficacy and efficiency using tactile and proprioception to refine the visual estimation results.
- 3) We evaluate our approach in real-world robotic manipulation scenarios, demonstrating large performance improvements over the underlying visual models [16, 17].

II. RELATED WORK

In this section, we survey the representative works in pose estimation from two perspectives: vision-based and visuotactile-based approaches. We also briefly discuss their relevance and implications for manipulation tasks in robotics.

A. Vision-based Pose Estimation

Traditionally, pose estimation research has focused on instance-level estimation [9–12], restricting trained models to specific, known objects and therefore limiting their practicality in real-world applications. Recent works have expanded the scope to category-level [13–15] and even to novel object estimation [16–20], addressing the challenges associated with broader applications. Notably, FoundationPose [16] achieves state-of-the-art performance in novel object pose estimation, leveraging massively scaled synthetic datasets and photo-realistic rendering for effective sim-to-real transfer. However, contact-rich manipulation tasks often introduce adverse conditions (*e.g.*, occlusion, high dynamicity), which makes

visual approaches less effective. Furthermore, applying such scaling law to tactile sensing is non-trivial due to the non-uniformity of tactile sensors and the wider sim-to-real gap.

Pose estimation has been heavily used in state-based policies, *e.g.*, [6–8] perform in-hand object re-orientation task using visual pose estimation. Their experiments, however, typically depend on elaborate setups, such as cages or multi-view systems, and are generally limited to a single object, restricting their real-world applicability. Accordingly, state-based policies are more challenging to deploy in real-life environments compared to visuomotor approaches [2–4].

B. Visuotactile-based Pose Estimation

In contrast, visuotactile-based methods offer a different approach by integrating visual inputs with tactile sensing. Kuppaswamy et al. [34] perform tactile-based pose estimation using the Iterative Closest Point (ICP) algorithm with a (vision-based) soft bubble sensor that provides dense geometric data of the contact area. Dikhale et al. [25] propose to leverage both RGB-D and tactile information to enhance in-hand object pose estimation. Later works have explored various techniques, such as applying graph neural networks on taxel signals [27, 28], integrating tactile sensing with proprioception [27], using explicit shape completion module [24], and hyper-resolving the sparse taxel signals [31]. While these works have shown great promise, they often face challenges related to datasets and often overfit to a specific type of tactile sensor and gripper.

The work most closely related to ours is by Liu et al. [29], which also uses tactile signals to refine visual results. However, their method relies on tracking object velocity by training an ad-hoc model and calculating the optical flow of marker points [35] on two vision-based tactile sensors mounted on a parallel gripper. This approach presents challenges when applied to non-visual sensors, where optical flow is difficult to obtain. In contrast, our zero-shot framework leverages physical constraints and imposes fewer conditions on the type of tactile sensor and end-effectors used, making it more versatile for real-world applications.

III. PROBLEM DEFINITION

Given a rigid object \mathcal{O} , we estimate (in the camera’s coordinate frame) its pose $T = (R, \mathbf{t}) \in \mathbb{SE}(3)$, comprised of its orientation $R \in \mathbb{SO}(3)$ and position $\mathbf{t} \in \mathbb{R}^3$. Following the problem setup of novel object pose estimation [16, 17], we assume the 3D mesh $M_{\mathcal{O}}$ of the object \mathcal{O} is either given or reconstructed [36, 37]. The object \mathcal{O} is observed using an RGB-D sensor and visible for at least the initial frame.

Robot. The object \mathcal{O} is manipulated by the robot \mathcal{R} with rigid links and joints. We assume the robot kinematic model and the part meshes are given, *e.g.*, through Unified Robot Description Format (URDF), to enable an accurate calculation of the robot’s end-effector poses. The relative transform between the end-effectors and camera frames is obtained using hand-eye calibration [38].

Sensor. The robot \mathcal{R} is equipped with tactile sensors \mathcal{S} , and their installation positions are known. We assume that

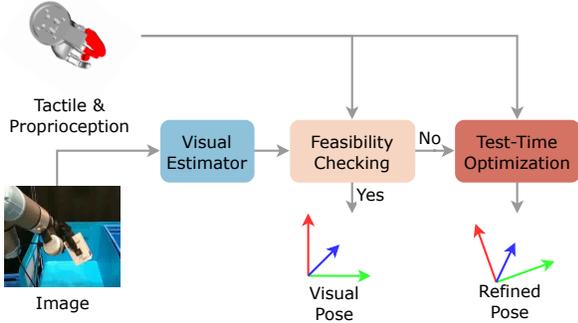


Fig. 2: **Overview of ViTa-Zero.** Red fingers in the robot hand model represent activated fingertip tactile sensors.

the sensors will only make contact with the object \mathcal{O} and no self-contact. Although we focus on taxel-based sensors in this paper, our framework can extend to other types of sensors, which we will discuss in the next section.

IV. METHODOLOGY

Our framework consists of three modules: visual estimation, feasibility checking, and tactile refinement. An overview of our framework is in Fig. 2. Initially, a visual model estimates the pose denoted as T . Then, we assess the feasibility of T using constraints derived from the tactile signals and proprioception. If T does not meet these constraints, we refine it through our test-time optimization algorithm using tactile and proprioceptive observations, yielding the final pose estimate, denoted as T^* .

A. Visual Estimation

We pass the visual observation into a visual model to obtain the initial pose estimate T . In this paper, we utilized MegaPose [17] for RGB input and FoundationPose [16] for RGB-D input. We follow the estimate-then-track procedure: we utilize the estimation model for the first frame and the pose tracking model for subsequent frames to estimate the relative pose between consecutive frames. Our experiments show that the visual models struggle to cope with severe occlusions and dynamic scenes due to the lack of visual cues, leading to significant failures and accumulative errors during manipulation tasks (Fig. 3). Therefore, refining visual results with tactile sensing is crucial for reliable manipulation.

B. Feasibility Checking

We introduce the geometric representations of each observation used for feasibility checking: tactile signals, the object, and the robot (proprioception).

Representation for the tactile signals. Prior studies have represented tactile signals as image [39–44] and raw values [45, 46]. However, these representations do not generalize well across different sensors, *e.g.* image-based methods are ineffective for taxel-based or F/T sensors. To improve generalization across different tactile sensors, we represent

the tactile signal as a point cloud P_S in the camera coordinate system, whose positions could be obtained through forward kinematics. For taxel-based sensors, such as those used in this paper, point values (in Boolean) are obtained by discretizing sensor outputs with a binary threshold [24, 25, 46, 47]. For vision-based sensors [35, 48–50], the point cloud could be obtained through depth estimation [23, 34, 51–54].

Representation for the object. Given the object mesh $M_{\mathcal{O}}$, we obtain the model point cloud $P_{\mathcal{O}}$ for contact constraint checking by randomly sampling points on the mesh. $P_{\mathcal{O}}$ is further downsampled using voxelization to create a voxel grid $V_{\mathcal{O}}$, allowing faster collision checks.

Representation for the robot. We use a mesh representation $M_{\mathcal{R}}$ for the robot \mathcal{R} , parameterized by joint angles, instead of directly using joint angles or end-effectors’ positions [27, 45, 46, 55] due to the latter’s specificity to robot embodiments and less direct physical interpretation. This mesh is converted into a downsampled point cloud $P_{\mathcal{R}}$, using the same downsampling techniques applied to the object mesh to accelerate feasibility checks.

1) *Contact constraint:* Each point in the tactile point cloud P_S must be in contact with some point on the object model point cloud. The minimum pairwise Euclidean distance between a point $\mathbf{p}_{\mathcal{O}}^i$ on the transformed object point cloud $P_{\mathcal{O}}$ and a point \mathbf{p}_S^j in the tactile point cloud P_S must satisfy:

$$\min_{i,j} \|T(\mathbf{p}_{\mathcal{O}}^i) - \mathbf{p}_S^j\| \leq \theta_c, \quad (1)$$

where θ_c is the contact distance threshold.

2) *Penetration constraint:* No point on the robot model should penetrate the object model and vice versa. This constraint is enforced by checking the overlap (intersection) between the transformed object voxel grid $T(V_{\mathcal{O}})$ and the robot model point cloud $P_{\mathcal{R}}$. We use the downsampled robot point cloud $P_{\mathcal{R}}$ as queries and calculate the proportion of the occupied voxels in the object model:

$$\frac{\text{card}[T(V_{\mathcal{O}}) \cap P_{\mathcal{R}}]}{\text{card}[T(V_{\mathcal{O}})]} \leq \theta_p, \quad (2)$$

where $\text{card}[\cdot]$ denotes set cardinality and θ_p is the allowable overlapping threshold to compensate for sensor noise and calibration errors.

3) *Kinematic constraint:* The object’s motion must be physically feasible. The following first-order differential constraint on the object’s position prevents unrealistic motion, such as pose flipping between frames for symmetrical objects, by ensuring the movement of the object complies with the actual dynamics:

$$\|P_n - P_{n-1}\| \leq \theta_d, \quad (3)$$

where $P_n \subset P_{\mathcal{O}}$ represents the contact patch on the object model for n^{th} frame.

C. Test-time Optimization

If the visual pose estimate T fails to meet any constraints, we refine it using a test-time optimization approach. We

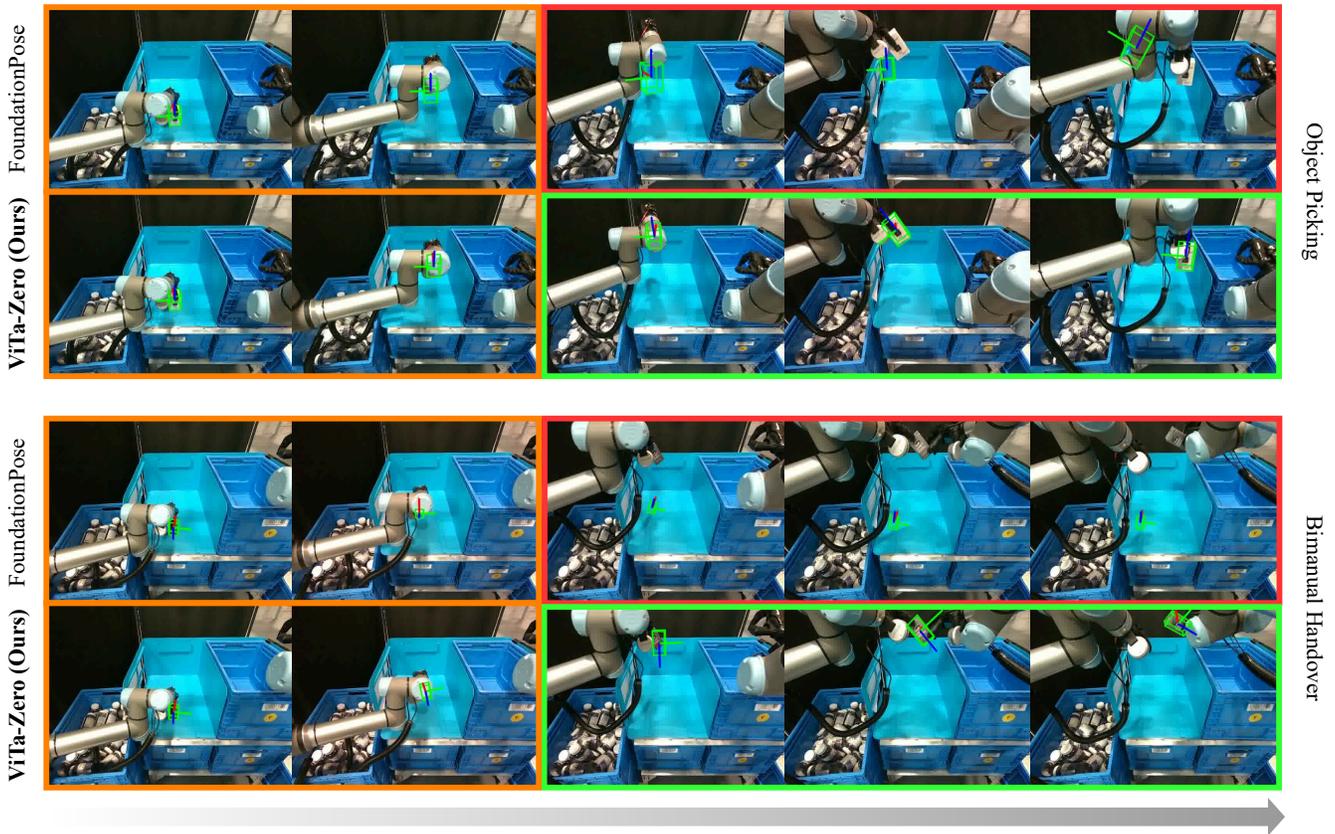


Fig. 3: **Qualitative results.** We demonstrate the performance during the object picking and bimanual handover tasks with “camera” and “eyedrop” objects. During these manipulation tasks, it is common to encounter scenarios where the object is **highly occluded while moving**, as illustrated in the figure. Visual approaches, like FoundationPose, can **lose tracking** and fail due to the absence of visual information. In contrast, our method utilizes additional tactile and proprioceptive feedback to **maintain object tracking**, ensuring robust performance. We note that **the wrist cameras were not used** in this study.

employ a spring–mass model comprising attractive and repulsive springs. The spring generates elastic potential energy in the form of $\frac{1}{2}kx^2$ (from Hooke’s Law), where x is the displacement of the spring. Unlike the previous works that apply spring–mass models for hand grasp synthesis [56–58], we repurpose it for object pose refinement based on tactile and proprioceptive feedback. The optimization objective is to *find a relative pose T_Δ that minimizes the total potential energy from both springs.*

The attractive spring pulls the object towards the tactile sensors having *in-contact* states and reaches its rest length when the object is in perfect contact with all such sensors. The potential energy of attractive spring is defined as $E_a = \frac{1}{2}k_a \min_{i,j} \|T_\Delta(\mathbf{p}_O^i) - \mathbf{p}_S^j\|^2$. The repulsive spring prevents penetration between the robot model P_R and the object model P_O . We define the potential energy as $E_r = \frac{1}{2}k_r \max(0, \gamma)^2$ where γ is the penetration distance defined as:

$$\gamma = \frac{1}{\text{card}(P_R)} \sum_i (\tilde{\mathbf{p}}_O^i - \mathbf{p}_R^i) \cdot \tilde{\mathbf{n}}_O^i. \quad (4)$$

Here, $\tilde{\mathbf{p}}_O^i = \underset{\mathbf{p} \in P_O}{\text{argmin}} \|\mathbf{p} - T_\Delta(\mathbf{p}_R^i)\|$ represents the nearest neighbor of the robot’s refined point i ; and $\tilde{\mathbf{n}}_O^i$ is its unit

normal on P_O pointing outward of the mesh surface. Hence, γ is the mean of signed distances along $\tilde{\mathbf{n}}_O^i$ between points P_R and their nearest neighbors in P_O .

We parameterize the relative pose T_Δ into relative rotation R_Δ (expressed in angle–axis form using the RoMa library [59]) and relative translation \mathbf{t}_Δ . We apply L_2 regularization on the parameters to stabilize the optimization. The refined pose T^* is obtained by:

$$\begin{aligned} T_\Delta &= \underset{R_\Delta, \mathbf{t}_\Delta}{\text{argmin}} (E_a + E_r + \lambda \mathcal{L}_2), \\ T^* &= T_\Delta \circ T, \end{aligned} \quad (5)$$

where \mathcal{L}_2 denotes the L_2 regularization loss. In the subsequent frame, we perform pose tracking on the T^* .

We implement our algorithm using PyTorch [60] and Adam [61] solver. To facilitate the optimization process, we initialize the T_Δ using the average translation changes of the activated taxels. This initialization strategy smoothes the optimization process and improves the overall performance.

V. EXPERIMENTS

We evaluate our approach on a real-world robot platform (Fig. 4), consisting of two Universal Robots UR5e robot arms and two PSYONIC Ability hands. The anthropomorphic

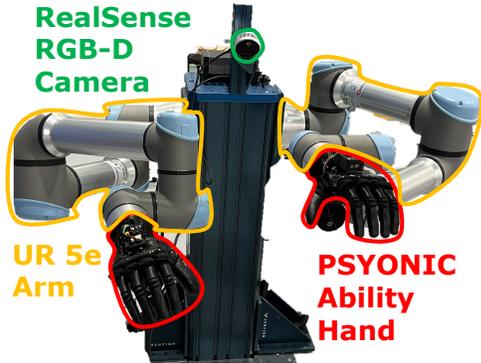


Fig. 4: **Our robot platform.** Our setup consists of two Universal Robots UR 5e arms and PSYONIC Ability hands.

hand has five fingers, each equipped with six FSR sensors on the fingertip to provide tactile sensing. More details regarding this platform are available in [45]. Visual sensing is provided by RealSense RGB-D cameras, specifically a combination of the D455 and L515 models. While our framework is designed to generalize across various embodiments and tactile sensors, this paper focuses on the Ability hand platform. We test on five daily objects (Fig. 5) purchased on Amazon.com, which could be easily held by the Ability hand and, to our best knowledge, are not part of the Objaverse [62] and GSO datasets [63]. Therefore, these five objects are considered novel for the following experiments.

In the feasibility checking, we set the contact threshold $\theta_c = 0.05$, the allowable penetration threshold $\theta_p = 0.008$, and the kinematic threshold $\theta_d = 0.03$. For the test-time optimization algorithm, the optimizer’s learning rate is set to 10^{-3} , and the number of iterations is empirically tuned to 10. The loss function parameters are set as follows: attractive energy weight $k_a = 1$, repulsive energy weight $k_r = 1000$, and L_2 regularization parameter $\lambda = 1000$.

In the following experiments, we compare our method with the state-of-the-art novel object pose estimation models: MegaPose (RGB) and FoundationPose (RGB-D)¹. We obtain the segmentation mask and bounding box needed for these models using off-the-shelf models [64, 65].

A. Qualitative Results

Figure 3 presents qualitative results for object picking and bimanual handover tasks. We observe that the visual model often loses tracking during manipulation (e.g., grasping) causing catastrophic failures. In contrast, our approach effectively leverages tactile and proprioception to refine the visual estimation result and maintain object tracking.

B. Quantitative Results

Due to the lack of public visuotactile object pose estimation datasets², we collect a dataset for testing purposes.

¹We implement FoundationPose using their released checkpoint at <https://github.com/NVlabs/FoundationPose>. Since the released checkpoint was trained on a slightly different dataset, it achieves similar but not the exact performance reported in their paper.

²[24, 25, 27, 31]’s datasets are proprietary, and [23, 26] are not released by the time we finish this paper.

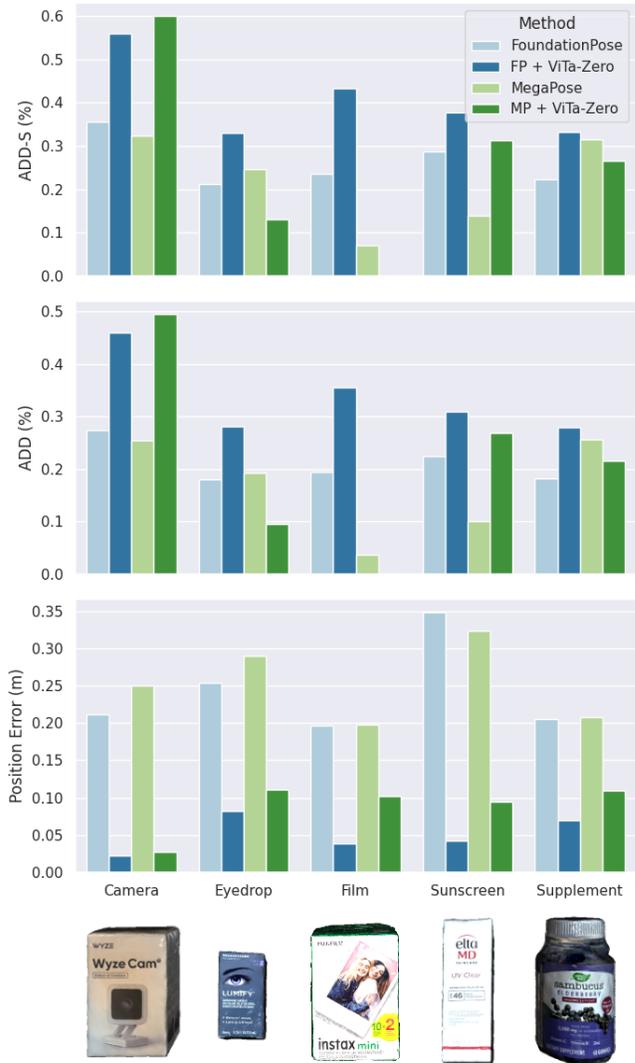


Fig. 5: **Comparison with FoundationPose (FP) and MegaPose (MP).** The performance is measured by the AUC of ADD-S and ADD (the higher, the better) and position error (PE) (the lower, the better).

We evaluate the performance using position error (PE) and the area under the curve (AUC) of ADD and ADD-S [9, 16, 24, 25]. PE is determined as the L_2 norm between the ground-truth translation and the estimated translation.

TABLE I: Quantitative comparison with FoundationPose (FP) and MegaPose (MP). The best results are bolded.

Method	ADD-S \uparrow	ADD \uparrow	PE \downarrow
FoundationPose	26.30	21.09	0.24
FP + ViTa-Zero	40.68	33.68	0.05
MegaPose	21.90	16.88	0.25
MP + ViTa-Zero	26.19	21.48	0.09

1) *Comparison with baselines:* We leverage FoundationPose [16] and MegaPose [17] as our backbones and compare the performance against them (Fig. 5 and Tab. I). Since MegaPose lacks depth information, *MegaPose* and *MP* +

ViTa-Zero have lower performance than *FoundationPose* and *FP + ViTa-Zero*. Nonetheless, our framework demonstrates a consistent performance improvement upon its visual backbone, achieving an average increase of 55% in the AUC of ADD-S and 60% in ADD, along with an 80% lower PE compared to *FoundationPose*. Despite these significant improvements, the error rate remains relatively high compared to other computer vision benchmarks [66], highlighting the inherent challenges of pose estimation during manipulation tasks.

TABLE II: Ablation studies on loss terms.

Ablations	ADD-S	ADD	PE
Without Attractive Loss	39.91	32.70	0.06
Without Penetration Loss	37.04	30.83	0.05
Without L_2 Loss	37.24	30.68	0.05
All Loss	40.68	33.68	0.05

2) *Ablation studies on each loss term:* We perform ablation studies to assess the impact of each loss term on our approach (Tab. II). When the attractive loss is removed, the refined pose remains static, disregarding tactile feedback. However, the presence of the penetration loss ensures that the pose is still “pushed” by the hand model as it moves. Without the penetration loss, the optimization often leads to solutions in which the object model penetrates the robot model. Without L_2 regularization loss, the optimization process becomes unstable and more challenging to converge.

Overall, our empirical results suggest that excluding any of these three loss terms lead to a decrease in performance, as measured by ADD and ADD-S metrics.

TABLE III: Ablation studies on refinement algorithm.

Method	ADD-S	ADD	PE
ICP	28.37	23.64	0.05
Ours	40.68	33.68	0.05

3) *Ablation studies on refinement algorithm:* A preliminary version of our test-time optimization algorithm can be viewed as performing point cloud registration between the tactile point cloud and the object surface patches [34]. In this study, we investigate the effect of using the ICP algorithm [67] as our refinement algorithm (Tab. III). Specifically, we calculate the T_Δ by registering the object point cloud P_O , serving as the source, with the tactile point cloud P_S , used as the target. In this paper, we employ the point-to-point ICP estimation method, which minimizes the Euclidean distance between corresponding points.

Our experiments demonstrate that the ICP algorithm effectively refines the visual pose T by aligning it with the nearest tactile points. However, due to the absence of penetration penalties, the solutions frequently exhibit unintended interpenetration with the robot model.

4) *Ablation studies on initialization for optimization:* In Table IV, we study the impact of removing the initialization for the test-time optimization algorithm. To recall, we initialize the optimized parameters T_Δ using the average

TABLE IV: Ablation studies on optimization initialization.

Method	ADD-S	ADD	PE
Without Init	39.83	35.85	0.06
With Init	40.68	33.68	0.05

translation of the activated taxels. Given the highly non-convex nature of our optimization problem, we find that this initialization strategy leads to more stable optimization and facilitates convergence.

TABLE V: Running time for each component.

Module	Time (s)	Hz
Pose tracking	0.011	90.9
Feasibility checking	0.006	166.7
Test-time optimization	0.051	19.6

5) *Running time:* Given our focus on manipulation applications, we evaluate the running time of each critical component in our framework. Our platform is equipped with an Intel i9-14900K CPU and an NVIDIA RTX 4090 GPU. For the pose tracking speed of *FoundationPose*, we set the refinement iterations as two. Overall, our framework could reach an online refresh rate of around 20 Hz, given ten iterations during test-time optimization.

VI. CONCLUSION

In this paper, we introduced a novel zero-shot visuotactile object 6D pose estimation framework. Our approach eliminates the need to collect a tactile dataset and demonstrates robust real-world performance on novel objects. The proposed test-time optimization algorithm leverages a spring-mass model to refine the infeasible visual estimations using tactile and proprioceptive observations. Experiment results show that our framework could consistently improve the performance of the visual models by a large margin. Notably, we achieved an average increase of 55% in the area under the curve (AUC) of ADD-S and 60% in ADD, along with an 80% lower position error on *FoundationPose*. We also conducted extensive ablation studies to evaluate the impact of various design choices, including different loss functions, refinement algorithms, and optimization initialization techniques. For future work, we plan to evaluate more object shapes and integrate our pose estimator into manipulation policies to assess its impact on the overall performance of manipulation tasks.

Limitations. We identify two limitations in our framework: i) Following the setting of novel object pose estimation, we assume the 3D CAD model of the object is either provided or can be reconstructed from reference views. However, this might be challenging for in-the-wild manipulation applications, where active visuotactile object reconstruction techniques [23] could be leveraged to mitigate this issue. ii) Our approach relies on the forward kinematics of the robot, which is available in most commercially available grippers. However, this might be non-trivial when dealing with deformable or compliant grippers [68]. Recent works [69, 70] offer potential solutions to address this challenge.

REFERENCES

- [1] O. Kroemer, S. Niekum, and G. Konidaris, “A Review of Robot Learning for Manipulation: Challenges, Representations, and Algorithms,” *Journal of Machine Learning Research*, vol. 22, no. 30, pp. 1–82, 2021.
- [2] T. Chen, M. Tippur, S. Wu, V. Kumar, E. Adelson, and P. Agrawal, “Visual dexterity: In-hand reorientation of novel and complex object shapes,” *Science Robotics*, vol. 8, no. 84, p. eadc9244, Nov. 2023, publisher: American Association for the Advancement of Science.
- [3] T. Chen, J. Xu, and P. Agrawal, “A System for General In-Hand Object Re-Orientation,” in *Proceedings of the 5th Conference on Robot Learning*. PMLR, Jan. 2022, pp. 297–307, ISSN: 2640-3498.
- [4] H. Qi, B. Yi, S. Suresh, M. Lambeta, Y. Ma, R. Calandra, and J. Malik, “General In-hand Object Rotation with Vision and Touch,” Aug. 2023.
- [5] t. Thalhammer, D. Bauer, P. Hönig, J.-B. Weibel, J. García-Rodríguez, and M. Vincze, “Challenges for Monocular 6D Object Pose Estimation in Robotics,” *IEEE Transactions on Robotics*, pp. 1–20, 2024, conference Name: IEEE Transactions on Robotics.
- [6] OpenAI, I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas, J. Schneider, N. Tezak, J. Tworek, P. Welinder, L. Weng, Q. Yuan, W. Zaremba, and L. Zhang, “Solving Rubik’s Cube with a Robot Hand,” Oct. 2019, arXiv:1910.07113 [cs, stat].
- [7] O. M. Andrychowicz, B. Baker, M. Chociej, R. Józefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, J. Schneider, S. Sidor, J. Tobin, P. Welinder, L. Weng, and W. Zaremba, “Learning dexterous in-hand manipulation,” *The International Journal of Robotics Research*, vol. 39, no. 1, pp. 3–20, Jan. 2020, publisher: SAGE Publications Ltd STM.
- [8] A. Handa, A. Allshire, V. Makoviychuk, A. Petrenko, R. Singh, J. Liu, D. Makoviichuk, K. Van Wyk, A. Zhurkevich, B. Sundaralingam, and Y. Narang, “DeXtreme: Transfer of Agile In-hand Manipulation from Simulation to Reality,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, May 2023, pp. 5977–5984.
- [9] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, “PoseCNN: A Convolutional Neural Network for 6D Object Pose Estimation in Cluttered Scenes,” vol. 14, Jun. 2018.
- [10] C. Wang, D. Xu, Y. Zhu, R. Martin-Martin, C. Lu, L. Fei-Fei, and S. Savarese, “DenseFusion: 6D Object Pose Estimation by Iterative Dense Fusion,” 2019, pp. 3343–3352.
- [11] K. Park, T. Patten, and M. Vincze, “Pix2Pose: Pixel-Wise Coordinate Regression of Objects for 6D Pose Estimation,” 2019, pp. 7668–7677.
- [12] Y. Li, Y. Mao, R. Bala, and S. Hadap, “MRC-Net: 6-DoF Pose Estimation with MultiScale Residual Correlation,” 2024, pp. 10476–10486.
- [13] H. Wang, S. Sridhar, J. Huang, J. Valentin, S. Song, and L. J. Guibas, “Normalized Object Coordinate Space for Category-Level 6D Object Pose and Size Estimation,” 2019, pp. 2642–2651.
- [14] D. Chen, J. Li, Z. Wang, and K. Xu, “Learning Canonical Shape Space for Category-Level 6D Object Pose and Size Estimation,” 2020, pp. 11973–11982.
- [15] T. Lee, J. Tremblay, V. Blukis, B. Wen, B.-U. Lee, I. Shin, S. Birchfield, I. S. Kweon, and K.-J. Yoon, “TTA-COPE: Test-Time Adaptation for Category-Level Object Pose Estimation,” 2023, pp. 21285–21295.
- [16] B. Wen, W. Yang, J. Kautz, and S. Birchfield, “FoundationPose: Unified 6D Pose Estimation and Tracking of Novel Objects,” 2024, pp. 17868–17879.
- [17] Y. Labbé, L. Manuelli, A. Mousavian, S. Tyree, S. Birchfield, J. Tremblay, J. Carpentier, M. Aubry, D. Fox, and J. Sivic, “MegaPose: 6D Pose Estimation of Novel Objects via Render & Compare,” Aug. 2022.
- [18] X. He, J. Sun, Y. Wang, D. Huang, H. Bao, and X. Zhou, “OnePose++: Keypoint-Free One-Shot Object Pose Estimation without CAD Models,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 35103–35115, Dec. 2022.
- [19] Y. Liu, Y. Wen, S. Peng, C. Lin, X. Long, T. Komura, and W. Wang, “Gen6D: Generalizable Model-Free 6-DoF Object Pose Estimation from RGB Images,” in *Computer Vision – ECCV 2022*, S. Avidan, G. Brostow, M. Cissé, G. M. Farinella, and T. Hassner, Eds. Cham: Springer Nature Switzerland, 2022, pp. 298–315.
- [20] J. Lin, L. Liu, D. Lu, and K. Jia, “SAM-6D: Segment Anything Model Meets Zero-Shot 6D Object Pose Estimation,” 2024, pp. 27906–27916.
- [21] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. C. Burchfiel, and S. Song, “Diffusion Policy: Visuomotor Policy Learning via Action Diffusion,” vol. 19, Jul. 2023.
- [22] Y. Ze, G. Zhang, K. Zhang, C. Hu, M. Wang, and H. Xu, “3D Diffusion Policy,” Mar. 2024, arXiv:2403.03954 [cs].
- [23] S. Suresh, H. Qi, T. Wu, T. Fan, L. Pineda, M. Lambeta, J. Malik, M. Kalakrishnan, R. Calandra, M. Kaess, J. Ortiz, and M. Mukadam, “Neural feels with neural fields: Visuo-tactile perception for in-hand manipulation,” Dec. 2023, arXiv:2312.13469 [cs].
- [24] H. Li, S. Dikhale, S. Iba, and N. Jamali, “ViHOPE: Visuotactile In-Hand Object 6D Pose Estimation With Shape Completion,” *IEEE Robotics and Automation Letters*, vol. 8, no. 11, pp. 6963–6970, Nov. 2023.
- [25] S. Dikhale, K. Patel, D. Dhingra, I. Naramura, A. Hayashi, S. Iba, and N. Jamali, “VisuoTactile 6D Pose Estimation of an In-Hand Object Using Vision and Tactile Sensor Data,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2148–2155, Apr. 2022.
- [26] Z. Wan, Y. Ling, S. Yi, L. Qi, W. W. Lee, M. Lu, S. Yang, X. Teng, P. Lu, X. Yang, M.-H. Yang, and H. Cheng, “VinT-6D: A Large-Scale Object-in-hand Dataset from Vision, Touch and Proprioception,” in *Proceedings of the 41st International Conference on Machine Learning*. PMLR, Jul. 2024, pp. 49921–49940, ISSN: 2640-3498.
- [27] A. Rezaeadeh, S. Dikhale, S. Iba, and N. Jamali, “Hierarchical Graph Neural Networks for Proprioceptive 6D Pose Estimation of In-hand Objects,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, May 2023, pp. 2884–2890.
- [28] Y. Tu, J. Jiang, S. Li, N. Hendrich, M. Li, and J. Zhang, “PoseFusion: Robust Object-in-Hand Pose Estimation with SelectLSTM,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2023, pp. 6839–6846, ISSN: 2153-0866.
- [29] Y. Liu, X. Xu, W. Chen, H. Yuan, H. Wang, J. Xu, R. Chen, and L. Yi, “Enhancing Generalizable 6D Pose Tracking of an In-Hand Object with Tactile Sensing,” *IEEE Robotics and Automation Letters*, vol. 9, no. 2, pp. 1106–1113, Feb. 2024, arXiv:2210.04026 [cs].
- [30] Y. Gao, S. Matsuoka, W. Wan, T. Kiyokawa, K. Koyama, and K. Harada, “In-Hand Pose Estimation Using Hand-Mounted RGB Cameras and Visuotactile Sensors,” *IEEE Access*, vol. 11, pp. 17218–17232, 2023, conference Name: IEEE Access.
- [31] H. Li, S. Dikhale, J. Cui, S. Iba, and N. Jamali, “HyperTactel: Hyper-Resolution for Tactel-Based Tactile Signals Through Contrastive Learning,” Aug. 2024, arXiv:2408.08312 [cs].
- [32] F. Yang, C. Ma, J. Zhang, J. Zhu, W. Yuan, and A. Owens, “Touch and Go: Learning from Human-Collected Vision and Touch,” Jun. 2022.
- [33] S. Wang, M. Lambeta, P.-W. Chou, and R. Calandra, “TACTO: A Fast, Flexible, and Open-Source Simulator for High-Resolution Vision-Based Tactile Sensors,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3930–3937, Apr. 2022, conference Name: IEEE Robotics and Automation Letters.
- [34] N. Kuppaswamy, A. Castro, C. Phillips-Grafflin, A. Alspach, and R. Tedrake, “Fast Model-Based Contact Patch and Pose Estimation for Highly Deformable Dense-Geometry Tactile Sensors,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1811–1818, Apr. 2020, conference Name: IEEE Robotics and Automation Letters.
- [35] I. H. Taylor, S. Dong, and A. Rodriguez, “GelSlim 3.0: High-Resolution Measurement of Shape, Force and Slip in a Compact Tactile-Sensing Finger,” in *2022 International Conference on Robotics and Automation (ICRA)*, May 2022, pp. 10781–10787.
- [36] B. Wen, J. Tremblay, V. Blukis, S. Tyree, T. Muller, A. Evans, D. Fox, J. Kautz, and S. Birchfield, “BundleSDF: Neural 6-DoF Tracking and 3D Reconstruction of Unknown Objects,” Mar. 2023, arXiv:2303.14158 [cs].
- [37] T. Müller, A. Evans, C. Schied, and A. Keller, “Instant Neural Graphics Primitives with a Multiresolution Hash Encoding,” *ACM Trans. Graph.*, vol. 41, no. 4, pp. 102:1–102:15, Jul. 2022, place: New York, NY, USA Publisher: ACM.
- [38] E. Marchand, F. Spindler, and F. Chaumette, “ViSP for visual servoing: a generic software platform with a wide class of robot control skills,” *IEEE Robotics & Automation Magazine*, vol. 12, no. 4, pp. 40–52, Dec. 2005, conference Name: IEEE Robotics & Automation Magazine.
- [39] M. B. Villalonga, A. Rodriguez, B. Lim, E. Valls, and T. Sechopoulos, “Tactile Object Pose Estimation from the First Touch with Geometric Contact Rendering,” in *Proceedings of the 2020 Conference on Robot Learning*. PMLR, Oct. 2021, pp. 1015–1029, ISSN: 2640-3498.
- [40] G. M. Caddeo, N. A. Piga, F. Bottarel, and L. Natale, “Collision-aware In-hand 6D Object Pose Estimation using Multiple Vision-based Tactile Sensors,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, May 2023, pp. 719–725.
- [41] I. Guzey, B. Evans, S. Chintala, and L. Pinto, “Dexterity from Touch:

- Self-Supervised Pre-Training of Tactile Representations with Robotic Play,” in *Proceedings of The 7th Conference on Robot Learning*, PMLR, Dec. 2023, pp. 3142–3166, iSSN: 2640-3498.
- [42] F. Yang, C. Feng, Z. Chen, H. Park, D. Wang, Y. Dou, Z. Zeng, X. Chen, R. Gangopadhyay, A. Owens, and A. Wong, “Binding Touch to Everything: Learning Unified Multimodal Tactile Representations,” Jan. 2024, arXiv:2401.18084 null.
- [43] Y. Dou, F. Yang, Y. Liu, A. Loquercio, and A. Owens, “Tactile-Augmented Radiance Fields,” May 2024, arXiv:2405.04534 [cs].
- [44] J. Zhao, Y. Ma, L. Wang, and E. H. Adelson, “Transferable Tactile Transformers for Representation Learning Across Diverse Sensors and Tasks,” Jun. 2024, arXiv:2406.13640 [cs].
- [45] T. Lin, Y. Zhang, Q. Li, H. Qi, B. Yi, S. Levine, and J. Malik, “Learning Visuotactile Skills with Two Multifingered Hands,” Apr. 2024, arXiv:2404.16823 [cs].
- [46] Z.-H. Yin, B. Huang, Y. Qin, Q. Chen, and X. Wang, “Rotating without Seeing: Towards In-hand Dexterity through Touch,” vol. 19, Jul. 2023.
- [47] Z. Xue, H. Zhang, J. Cheng, Z. He, Y. Ju, C. Lin, G. Zhang, and H. Xu, “ArrayBot: Reinforcement Learning for Generalizable Distributed Manipulation through Touch,” Jun. 2023, arXiv:2306.16857 [cs].
- [48] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer, D. Jayaraman, and R. Calandra, “DIGIT: A Novel Design for a Low-Cost Compact High-Resolution Tactile Sensor With Application to In-Hand Manipulation,” *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 3838–3845, Jul. 2020, conference Name: IEEE Robotics and Automation Letters.
- [49] W. Yuan, S. Dong, and E. H. Adelson, “GelSight: High-Resolution Robot Tactile Sensors for Estimating Geometry and Force,” *Sensors*, vol. 17, no. 12, p. 2762, Dec. 2017.
- [50] E. Donlon, S. Dong, M. Liu, J. Li, E. Adelson, and A. Rodriguez, “GelSlim: A High-Resolution, Compact, Robust, and Calibrated Tactile-sensing Finger,” May 2018, arXiv:1803.00628 [cs].
- [51] M. Bauza, O. Canal, and A. Rodriguez, “Tactile Mapping and Localization from High-Resolution Tactile Imprints,” in *2019 International Conference on Robotics and Automation (ICRA)*, May 2019, pp. 3811–3817, iSSN: 2577-087X.
- [52] S. Suresh, Z. Si, S. Anderson, M. Kaess, and M. Mukadam, “Midastouch: Monte-Carlo inference over distributions across sliding touch,” in *Proceedings of The 6th Conference on Robot Learning*. PMLR, Mar. 2023, pp. 319–331, iSSN: 2640-3498.
- [53] S. Suresh, Z. Si, J. G. Mangelson, W. Yuan, and M. Kaess, “ShapeMap 3-D: Efficient shape mapping through dense touch and vision,” in *2022 International Conference on Robotics and Automation (ICRA)*, May 2022, pp. 7073–7080.
- [54] R. Ambrus, V. Guizilini, N. Kuppaswamy, A. Beaulieu, A. Gaidon, and A. Alspach, “Monocular Depth Estimation for Soft Visuotactile Sensors,” in *2021 IEEE 4th International Conference on Soft Robotics (RoboSoft)*, Apr. 2021, pp. 643–649.
- [55] M. A. Lee, Y. Zhu, P. Zachares, M. Tan, K. Srinivasan, S. Savarese, L. Fei-Fei, A. Garg, and J. Bohg, “Making Sense of Vision and Touch: Learning Multimodal Representations for Contact-Rich Tasks,” *IEEE Transactions on Robotics*, vol. 36, no. 3, pp. 582–596, Jun. 2020, conference Name: IEEE Transactions on Robotics.
- [56] L. Yang, X. Zhan, K. Li, W. Xu, J. Li, and C. Lu, “CPF: Learning a Contact Potential Field to Model the Hand-Object Interaction,” in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, Oct. 2021, pp. 11 077–11 086, iSSN: 2380-7504.
- [57] L. Yang, X. Zhan, K. Li, W. Xu, J. Zhang, J. Li, and C. Lu, “Learning a Contact Potential Field for Modeling the Hand-Object Interaction,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 46, no. 8, pp. 5645–5662, Aug. 2024, conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- [58] S. Brahmhatt, A. Handa, J. Hays, and D. Fox, “ContactGrasp: Functional Multi-finger Grasp Synthesis from Contact,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Nov. 2019, pp. 2386–2393, iSSN: 2153-0866.
- [59] R. Brégier, “Deep Regression on Manifolds: A 3D Rotation Case Study,” in *2021 International Conference on 3D Vision (3DV)*, Dec. 2021, pp. 166–174, iSSN: 2475-7888.
- [60] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimeshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “PyTorch: An Imperative Style, High-Performance Deep Learning Library,” in *Advances in Neural Information Processing Systems*, vol. 32. Curran Associates, Inc., 2019.
- [61] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” Jan. 2017, arXiv:1412.6980 [cs].
- [62] M. Deitke, D. Schwenk, J. Salvador, L. Weihs, O. Michel, E. Vanderbilt, L. Schmidt, K. Ehsani, A. Kembhavi, and A. Farhadi, “Objaverse: A Universe of Annotated 3D Objects,” 2023, pp. 13 142–13 153.
- [63] L. Downs, A. Francis, N. Koenig, B. Kinman, R. Hickman, K. Reymann, T. B. McHugh, and V. Vanhoucke, “Google Scanned Objects: A High-Quality Dataset of 3D Scanned Household Items,” in *2022 International Conference on Robotics and Automation (ICRA)*, May 2022, pp. 2553–2560.
- [64] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollar, and R. Girshick, “Segment Anything,” 2023, pp. 4015–4026.
- [65] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, Q. Jiang, C. Li, J. Yang, H. Su, J. Zhu, and L. Zhang, “Grounding DINO: Marrying DINO with Grounded Pre-Training for Open-Set Object Detection,” Jul. 2024, arXiv:2303.05499 [cs].
- [66] T. Hodan, M. Sundermeyer, Y. Labbe, V. N. Nguyen, G. Wang, E. Brachmann, B. Drost, V. Lepetit, C. Rother, and J. Matas, “BOP Challenge 2023 on Detection Segmentation and Pose Estimation of Seen and Unseen Rigid Objects,” 2024, pp. 5610–5619.
- [67] P. J. Besl and N. D. McKay, “Method for registration of 3-D shapes,” in *Sensor Fusion IV: Control Paradigms and Data Structures*, vol. 1611. SPIE, Apr. 1992, pp. 586–606.
- [68] L. U. Odhner and A. M. Dollar, “Stable, open-loop precision manipulation with underactuated hands,” *The International Journal of Robotics Research*, vol. 34, no. 11, pp. 1347–1360, Sep. 2015, publisher: SAGE Publications Ltd STM.
- [69] B. Chen, R. Kwiatkowski, C. Vondrick, and H. Lipson, “Fully body visual self-modeling of robot morphologies,” *Science Robotics*, vol. 7, no. 68, p. eabn1944, Jul. 2022, publisher: American Association for the Advancement of Science.
- [70] S. L. Li, A. Zhang, B. Chen, H. Matusik, C. Liu, D. Rus, and V. Sitzmann, “Unifying 3D Representation and Control of Diverse Robots with a Single Camera,” Jul. 2024, arXiv:2407.08722 [cs].