Learning to Visually Observe, Plan, and Control Compliant Robot In-Hand Manipulation

Andrew S. Morgan

I. Motivation

The dexterous capabilities afforded to humans by our hands is unparalleled to that of other species, allowing us to complete an array of daily manipulation tasks with ease [1]. As an illustration, let's examine the task of inserting a key into a lock. First, the hand must successfully acquire a stable grasp on the key, properly selecting contacts and maintaining them online via proprioceptive and tactile feedback. Once stable, the key is reconfigured within-hand via coordinated finger motions–balancing forces while making and breaking contact during a finger gaiting process. Upon reorientation, the hand-object system then evaluates how forces must be applied by the key into the lock tumbler for successful insertion and rotation.

This task decomposition serves as a single elaboration of the complex manipulation tasks humans mindlessly complete daily; seamlessly combining sensing, planning, and control. Numerous other tasks can be similarly decomposed, e.g. washing dishes [2], inserting batteries [3], or preparing food [4]. In emphasizing this narrative, service robots of the future must be able to complete a similar array of "everyday" tasks as that of a human, which has been unrealized by robots to date. In this research statement, I will elaborate on what we feel is a promising approach towards establishing these capabilities for hands and the role we believe compliance will play in this process.

Research Overview: *Complex robots complicate control; keep designs simple and exploit emergent behavior.* To this end, we propose to investigate how simple hands with compliant behavior can be leveraged to complete complex manipulation tasks. In our work, we (1) demonstrate the extent by which we can observe, plan, and control parasitic object motion [5] for in-hand manipulation solely through vision, i.e. without tactile sensors or joint encoders, and (2) propose fast, online multi-modal planning and control approaches for fixed-contact and fluid-contact scenarios with online adaptation for recovery.

II. Background

The fingers on hands can be conceptualized as a group of serial-link manipulators that must work in unison with one another; acquiring and modulating forces to maintain a stable grasp. Traditional works investigated how to do this with rigid, high-DOF end effectors [6]. However, these systems quickly became overly-complicated as they relied on tactile sensors, joint encoders, and advanced control/modeling architectures. Often, their utility in real-world scenarios were plagued by modeling inaccuracies, ultimately leading to task failure.

This complexity can be alleviated by thoughtful design. Within the past decade, soft, compliant, and underactuated hands have garnered much support within the robotics community, as these mechanisms are able to "absorb the slack" in modeling inaccuracies [7, 8, 9]. This classification of hands has shown to be particularly beneficial for grasping and online learning as they require less planning, less control, and little to no sensing in a more inexpensive package [10]. More explicitly, the kinematic reconfigurability of these mechanisms have turned the traditional *position* + *force* balancing problem into either a *position* or *force* control problem. While beneficial for grasping, in-hand manipulation, however, has seen little to no progress in the previous decade.

Estimated models, either analytical or learned [11], can be better leveraged when feedback is available. Closing the control loop through an external source, i.e. vision, can provide valuable information without complicating hand design. To this end, individual properties of the system, e.g. object poses, joint angles, contact phenomena, etc., are tracked while the passive adaptability of the hand innately modulates forces appropriately [12]. Formally, [13] has shown that visually extracted, mechanics-based features can sufficiently define the state of a hand-object system and we use this finding to support our approach with this sensing modality.

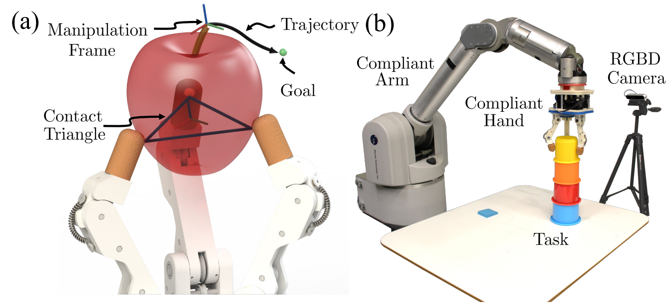

Fig. 1: (a) In-hand models that control virtual frames along an arbitrary trajectory can facilitate (b) completion of complex tasks such as cup stacking.

III. Extending Dexterous Capabilities of Hands

1) Fixed-contact manipulation: The traditional linear relationship between actuator motion and object motion, i.e. the Hand-Object Jacobian [14], does not generally exist for compliant hands [15]. Thus, as an initial investigation [16], we were interested in developing a representation that estimates this non-linear, compliant response. The model formulated, which utilizes geometric features extracted through vision, is generalized in terms of the estimated internal *energy* of the system, as determined by the compliant joint springs. It solves an optimization problem for the equilibrated next-state joint configuration, given: (i) an object-agnostic contact triangle

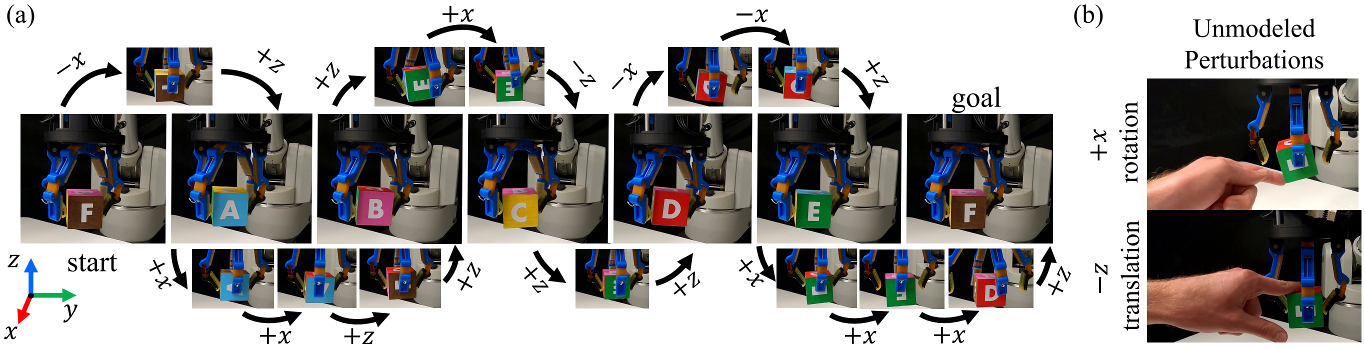

Fig. 2: (a) By visually tracking the cube online, our multi-modal planning and control algorithm is able to successfully achieve any object pose in $SO(3)$, (b) even with undesired perturbations. To the best of the author's knowledge, this is the first work in the literature to accomplish such capabilities against gravity and without a support surface. The success of our method can be largely attributed to the "inflated" modal transfer regions with online feedback from vision.

(Fig. 1.a), (ii) the current joint configuration, and, (iii) the actuation input. The principle of energy minimization supports our underlying theory for this method, where the system wants to lie in its lowest energy configuration–an idea similar to manipulation funnels in the literature [8, 17].

Our formulation serves as an estimated state transition model of the system, which provides for basic offline understanding of the hand's response to input. We then extend its utility in an online Model Predictive Control (MPC) framework that continually re-evaluates the transition model via a learned network. Compared to other online MPC methods in the literature, e.g. [18, 19, 20], ours was more computationally efficient and was shown to sufficiently adhere to the constraints of our task for real-world scenarios.

In-hand manipulation is particularly beneficial for assembly and insertion tasks, reducing system energy that would otherwise require whole-arm motion. Thus, we utilize our approach to assist in tight-tolerance (<0.25mm) and open-world insertion tasks (Fig. 1.b) [21]. Traditionally, precision manipulators equipped with expensive force/torque sensors mostly dominated the literature, e.g. [22, 23, 24]–representing complicated approaches to a seemingly simple problem. In our extension, we showed that by combining our MPC control framework, an external object tracker [25, 26], and the hand's passively compliant nature, we were able to account for any pose uncertainty between the object and the hole during insertion, providing a much simpler approach compared to those previously in the literature. Experiments demonstrated numerous successful insertions; from tight, industry-relevant tolerances of 5 different object geometries to delicate stacking, packing, and plug insertion open-world scenarios.

2) Fluid-contact manipulation: We can use the basis of our models in fixed-contact scenarios to extend to conditions where the contacts are always transitioning. Fundamentally, in-hand manipulation is often simplified by "fixing" contacts to an object, constraining the system to a single hand-object configuration manifold. This restraint limits the object's potential workspace and constrains it according to the kinematics of the hand. A way to alleviate such restraint is to continually change contact locations, either through coordinated slip or through finger gaiting. Unlike slip, finger gaiting can be a quasi-static process, which in turn helps enforce stability. Guided by this realization, we were interested in formulating an online multi-modal planning solution that can reorient objects in $SO(3)$ within-hand and without a support surface, unlike [27, 8, 28].

To accomplish this research goal [29], we devised two main manipulation modes via our energy formulation: x-axis rotation and z-axis rotation. The method was implemented on a four-fingered hand and was comprised of a bidirectional and multi-modal planning solution that solved for contact transfer regions between the modes. Compared to previous multi-modal planners, ours was unique in that compliance afforded an "inflated" transfer submanifold region between the modes. Ultimately, this realization permitted significant advancements in reliability to what has been previously possible in the literature. The final method is showcased by completing full $SO(3)$ rotations of a cube against gravity and with external, i.e. human, perturbations (Fig. 2). Results are presented in [29].

IV. Discussion and Future Work

Our results using simple, compliant hands support the aforementioned sentiment– *complex robots complicate control; keep designs simple and exploit emergent behavior*. We show that by (i) estimating a forward motion model of the hand and (ii) visually closing the control loop with a fast, online solution, we were able to extend in-hand capabilities well beyond what has been previously possible. Our general approach is advantageous as we can utilize simple, compliant, and inexpensive robot hands to study the fundamental building blocks of dexterous in-hand manipulation.

While the results we describe are promising, there are several future avenues to investigate. First, theoretical bounds by which compliant modal transfer regions for planning can be "inflated" beg several theoretical formulations, rather than empirical characterizations. Next, models describing compliance can be realized as differentiable funnels of system energy. Further formalizing this idea and understanding specific funnel properties would be an impactful contribution. And, finally, a promising direction would include studying energy models for compliant whole-hand caging manipulation [30].

Overall, we believe the culmination of these works elicits a promising approach to in-hand manipulation. While we have surely not reached human-level performance, the capabilities we have been able to demonstrate with simple hands continues to look promising for general-purpose robots of the future.

REFERENCES

[1] A. Bicchi, "Hands for dexterous manipulation and robust grasping: A difficult road toward simplicity," *IEEE Transactions on robotics and automation*, vol. 16, no. 6, pp. 652–662, 2000.

[2] W. Ruotolo, R. Thomasson, J. Herrera, A. Gruebele, and M. Cutkosky, "Distal Hyperextension Is Handy: High Range of Motion in Cluttered Environments," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 921–928, 2020.

[3] C. H. Kim and J. Seo, "Shallow-Depth Insertion: Peg in Shallow Hole Through Robotic In-Hand Manipulation," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 383–390, 2019.

[4] T. Bhattacharjee, G. Lee, H. Song, and S. S. Srinivasa, "Towards Robotic Feeding: Role of Haptics in Fork-Based Food Manipulation," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1485–1492, 2019.

[5] M. Liarokapis and A. Dollar, "Post-Contact, In-Hand Object Motion Compensation With Adaptive Hands," *IEEE Transactions on Automation Science and Engineering*, vol. PP, 10 2016.

[6] J. Kerr and B. Roth, "Analysis of multifingered hands," *The International Journal of Robotics Research*, vol. 4, no. 4, pp. 3–17, 1986.

[7] A. M. Dollar and R. D. Howe, "The Highly Adaptive SDM Hand: Design and Performance Evaluation," *The International Journal of Robotics Research*, vol. 29, no. 5, pp. 585–597, 2010.

[8] A. Bhatt, A. Sieler, S. Puhlmann, and O. Brock, "Surprisingly Robust In-Hand Manipulation: An Empirical Study," July 2021.

[9] M. G. Catalano, G. Grioli, E. Farnioli, A. Serio, C. Piazza, and A. Bicchi, "Adaptive synergies for the design and control of the Pisa/IIT SoftHand," *The International Journal of Robotics Research*, vol. 33, no. 5, pp. 768–782, 2014.

[10] A. S. Morgan, D. Nandha, G. Chalvatzaki, C. D'Eramo, A. M. Dollar, and J. Peters, "Model Predictive Actor-Critic: Accelerating Robot Skill Acquisition with Deep Reinforcement Learning," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 6672–6678.

[11] A. Sintov, A. S. Morgan, A. Kimmel, A. M. Dollar, K. E. Bekris, and A. Boularias, "Learning a State Transition Model of an Underactuated Adaptive Hand," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1287–1294, 2019.

[12] K. Hang, W. G. Bircher, A. S. Morgan, and A. M. Dollar, "Manipulation for self-Identification, and self-Identification for better manipulation," *Science Robotics*, vol. 6, no. 54, p. eabe1321, 2021. [Online]. Available: https://www.science.org/doi/abs/10.1126/scirobotics.abe1321

[13] A. S. Morgan, W. G. Bircher, and A. M. Dollar, "Towards Generalized Manipulation Learning Through Grasp Mechanics-Based Features and Self-Supervision," *IEEE Transactions on Robotics*, vol. 37, no. 5, pp. 1553–1569, 2021.

[14] R. M. Murray, Z. Li, and S. S. Sastry, *A mathematical introduction to robotic manipulation*. CRC press, 2017.

[15] L. U. Odhner and A. M. Dollar, "Dexterous manipulation with underactuated elastic hands," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 5254–5260.

[16] A. S. Morgan, K. Hang, and A. M. Dollar, "Object-Agnostic Dexterous Manipulation of Partially Constrained Trajectories," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5494–5501, 2020.

[17] M. Mason, "The mechanics of manipulation," in *Proceedings. 1985 IEEE International Conference on Robotics and Automation*, vol. 2. IEEE, 1985, pp. 544–548.

[18] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou, "Information theoretic MPC for model-based reinforcement learning," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 1714–1721.

[19] A. Lambert, F. Ramos, B. Boots, D. Fox, and A. Fishman, "Stein Variational Model Predictive Control," vol. 155, pp. 1278–1297, 16–18 Nov 2021.

[20] H. Bharadhwaj, K. Xie, and F. Shkurti, "Model-predictive control via cross-entropy and gradient-based optimization," in *Learning for Dynamics and Control*. PMLR, 2020, pp. 277–286.

[21] A. S. Morgan, B. Wen, J. Liang, A. Boularias, A. M. Dollar, and K. Bekris, "Vision-driven Compliant Manipulation for Reliable; High-Precision Assembly Tasks," in *Proceedings of Robotics: Science and Systems*, Virtual, July 2021.

[22] T. Tang, H. Lin, Yu Zhao, Wenjie Chen, and M. Tomizuka, "Autonomous alignment of peg and hole by force/torque measurement for robotic assembly," in *2016 IEEE International Conference on Automation Science and Engineering (CASE)*, 2016, pp. 162–167.

[23] K. Zhang, J. Xu, H. Chen, J. Zhao, and K. Chen, "Jamming Analysis and Force Control for Flexible Dual Peg-in-Hole Assembly," *IEEE Transactions on Industrial Electronics*, vol. 66, no. 3, pp. 1930–1939, 2019.

[24] Y. Fie and X. Zhao, "An Assembly Process Modeling and Analysis for Robotic Multiple Peg-in-hole," *Journal of Intelligent and Robotic Systems*, vol. 36, pp. 175–189, 2003.

[25] B. Wen, C. Mitash, B. Ren, and K. E. Bekris, "se(3)-TrackNet: Data-driven 6D Pose Tracking by Calibrating Image Residuals in Synthetic Domains," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 10 367–10 373.

[26] B. Wen, C. Mitash, S. Soorian, A. Kimmel, A. Sintov, and K. E. Bekris, "Robust, occlusion-aware pose estimation for objects grasped by adaptive hands," in *2020 IEEE International Conference on Robotics and Automation*

(ICRA). IEEE, 2020, pp. 6210–6217.
[27] O. M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray *et al.*, “Learning dexterous in-hand manipulation,” *The International Journal of Robotics Research*, vol. 39, no. 1, pp. 3–20, 2020.
[28] S. Abondance, C. B. Teeple, and R. J. Wood, “A Dexterous Soft Robotic Hand for Delicate In-Hand Manipulation,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5502–5509, 2020.
[29] A. S. Morgan, K. Hang, B. Wen, K. Bekris, and A. M. Dollar, “Complex In-Hand Manipulation Via Compliance-Enabled Finger Gaiting and Multi-Modal Planning,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4821–4828, 2022.
[30] W. G. Bircher, A. S. Morgan, and A. M. Dollar, “Complex manipulation with a simple robotic hand through contact breaking and caging,” *Science Robotics*, vol. 6, no. 54, p. eabd2666, 2021. [Online]. Available: https://www.science.org/doi/abs/10.1126/scirobotics.abd2666