

Towards Generalized Robot Assembly through Compliance-Enabled Contact Formations

Andrew S. Morgan, Quentin Bateux, Mei Hao, and Aaron M. Dollar .

Abstract—Contact can be conceptualized as a set of constraints imposed on two bodies that are interacting with one another in some way. The nature of a contact, whether a point, line, or surface, dictates how these bodies are able to move with respect to one another given a force, and a set of contacts can provide either partial or full constraint on a body’s motion. Decades of work have explored how to explicitly estimate the location of a contact and its dynamics, e.g., frictional properties, but investigated methods have been computationally expensive and there often exists significant uncertainty in the final calculation. This has affected further advancements in contact-rich tasks that are seemingly simple to humans, such as generalized peg-in-hole insertions. In this work, instead of explicitly estimating the individual contact dynamics between an object and its hole, we approach this problem by investigating compliance-enabled contact formations. More formally, contact formations are defined according to the constraints imposed on an object’s available degrees-of-freedom. Rather than estimating individual contact positions, we abstract out this calculation to an implicit representation, allowing the robot to either acquire, maintain, or release constraints on the object during the insertion process, by monitoring forces enacted on the end effector through time. Using a compliant robot, our method is desirable in that we are able to complete industry-relevant insertion tasks of tolerances $<0.25\text{mm}$ without prior knowledge of the exact hole location or its orientation. We showcase our method on more generalized insertion tasks, such as commercially available non-cylindrical objects and open world plug tasks.

I. INTRODUCTION

Robots interact with the world through physical contact. Contact has been studied in the literature for decades and can be conceptualized in a number of ways, but in its most fundamental sense, contacts are interaction-based constraints imposed on the relative motion between objects. As humans, we often rely on contact to change the state of our environment, whether it be removing keys from a bag or cleaning a dirty dish. While seemingly simple, contact is quite complex in that it requires accurate parameter estimation in order to sufficiently predict its current state, i.e., static, stick-slip, or slip conditions. Previous works have tried to estimate various properties of contact during manipulation, e.g., the location of the contact [1], frictional and curvature properties [2], [3], slip conditions [4], modes [5], etc., but have witnessed many bottlenecks.

Fundamentally, contacts define constraints and constraints define the available degrees-of-freedom (DOF) of an object.

This work was supported by the United States National Science Foundation under Grants IIS-1752134, IIS-1900681 & NRI-1734190.

Authors are with the Department of Mechanical Engineering and Materials Science, Yale University, USA. (email: {andrew.morgan, quentin.bateux, mei.hao, aaron.dollar}@yale.edu)

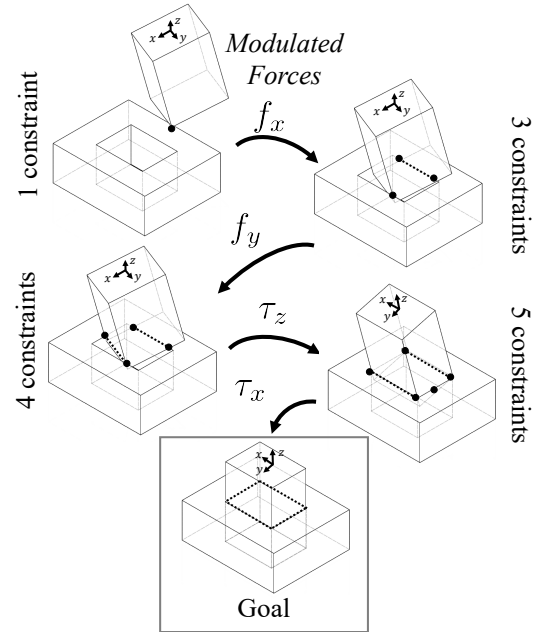


Fig. 1. Object insertion can be conceptualized as the continual addition and modulation through time of an object’s constrained degrees-of-freedom. By continually modulating forces once constraints are detected, tight tolerance insertion can be achieved without *a priori* knowledge of the object geometry or exact hole pose for convex objects.

In free space, an object is able to move in 6-DOF, but a non-cylindrical peg inserted into a slot is only able to translate in 1-DOF. When there exists some degree of uncertainty in the estimation of contact, the calculated constraints can be non-representative of the system’s true state. For instance, let’s assume there exists little uncertainty in the estimation of point contact locations. Here, we can claim that the restricted motion of an object with two point contacts may be similar or even identical in nature to the constraints imposed on an object with 100 edge point contacts on or near a single axis. In a practical sense, this distinction is insignificant in the free DOFs of the object, and merely adds needless computation. Now, let’s be more practical and assume there does exist some degree of uncertainty in our estimation of contact locations, even in a single point out of the 100. This calculation may estimate additional constraints placed on the object which are not physically valid. Thus, instead of explicitly defining the state of every contact between a robot and its environment, which is quite computationally expensive and is gravely subject to uncertainty, we in this work abstract out individual contact properties and are conversely interested in contact formations, or more generally, the constraints placed on an object’s motion when in contact with its environment.

Contact formations can be conceptualized as a grouping of contact sets that constrain an object in similar ways [6]. For instance, various contact scenarios could constrain either some translational DOFs, e.g., sphere in tube, or all translational DOFs, e.g. sphere in case. A set of all contact groupings that constrain an object in a similar way would be considered in a contact formation. Traditionally, acquiring and maintaining contact formations on a real robot platform was difficult due to an inability to practically modulate forces between two bodies over time. Within the past decade, a great deal of work has addressed the concept of dealing with uncertainty in robot mechanics, specifically, the ability to adapt to external forces as to continually maintain contact with an object. Compliant manipulation frameworks, either in the form of active [7] or passive architectures [8], [9], enable this ability – allowing a system to kinematically adapt to uncertain environments and maintain desired forces. This introduces the idea of *compliance-enabled contact formations*, which we leverage as to maintain a desired type of contact during the task.

Combining these ideas, this work investigates the applicability of controlling tight tolerance insertion tasks via the chaining of contact formations. Specifically, we leverage a control approach that modulates forces in controlled directions, and by continually adding forces along specific axes until contact, i.e., constraints, we transition between different contact formations until the task is successfully completed (Fig. 1). This approach utilizes external contact as a means of constraining the object’s potential motion [10], which in turn limits uncertainty of the object’s current configuration. While utilizing a compliant robot, this system aids in both acquiring and maintaining a desired contact configuration. Moreover, we show how in-hole jamming is limited due to this compliance, functioning similarly in concept to a remote center of compliance (RCC) device [11], [12]. Retrofitted with a 6-axis force/torque sensor at the end effector and an in-hand camera, we monitor the state of the object and servo to its desired goal forces and axial alignment. We can thus summarize our contributions:

- 1) Our algorithm is an object-agnostic procedure that does not require *a priori* knowledge of the object geometry, exact hole position, or exact hole orientation. To our knowledge, this is one of few works that can insert an object when axial rotation is largely unknown ($> 40^\circ$).
- 2) We do not require complex analytical or computationally expensive learned models of contact dynamics; our method relies solely on the idea of validating constraints via axial force application through time.
- 3) This method is conceptually straightforward, as it solely relies on following a trajectory in Contact Formation-space, and can thus be controlled physically in real-time.
- 4) We justify the true utility of using compliant systems for tight tolerance, contact-rich tasks by quantifying reconfigurability. Specifically, we quantify slip conditions for our underactuated hand, and show how the hand reorganizes its contacts when excessive forces are applied, which helps eliminate in-hole jamming.

II. RELATED WORK

Generalized robot assembly has been investigated for decades and the development of a practical, holistic solution has faced many challenges [13]. Elucidating this grand challenge can be found in previous works with: standard cylinders [7], [14], [15], multiple-peg objects [16], [17], soft or compliant cylinders [18], industrial inserts [19], and standard open world objects [8], [9], [20]–[22]. In an attempt to extend previous work, we outline a method that underscores the idea of *generalization*, in that our method provides a practical solution to insertion when there exists positional and orientational uncertainty of the hole pose, with minimal knowledge of the object geometry.

Peg-in-hole Assembly: Approaches to solving generic assembly tasks can be divided into two different categories: analytical approaches and learning-based approaches. In the former, approaches utilize contact models to reason about their state conditions [17], [18], controlling the manipulator based on force/torque sensor readings [14]. Model-based approaches using vision and rigid systems have been difficult to implement, as uncertainty in the manipulator’s pose can create dangerously large forces [23]. Compliance has become a key component to overcoming this issue, both in software-based [7] and hardware-based [8], [9] architectures. Many works have disregarded the use of a dexterous robot hand, directly attaching the object to the manipulator. Two previous works utilized a dexterous hand [24], [25], where both found it advantageously extended the robot’s workspace. In both cases, *a priori* knowledge of the hole pose was available.

Learning-based solutions help deal with robot sensor and model uncertainty, and typically require physical environment exploration or some form of human-in-the-loop demonstrations [15]. Reinforcement learning and self-supervised learning [26], [27] have been popular approaches, allowing the robot to interact with its environment and sufficiently explore desired regions [19], [20], [28]. Notably, learning manipulation policies is both, time consuming and data intensive, which increases the chance of damaging the robot.

While attempting to maintain a model-free nature but also limiting the need for expensive exploration, we are interested in forming a solution that can achieve reliable insertions of tight tolerance by solely using force data and vision feedback, and without *a priori* knowledge of the hole pose.

Contact Formations: *Contact formations* “provide a qualitative description of how 2 or more objects make contact with one another (e.g., vertex to surface, edge to edge)” [6]. This formulation is advantageous, in that it implicitly defines the contacts and thus constraints imposed on an object’s motion. Contact formations have been computed in different ways: from using CAD models [29] to probabilistic frameworks based on interaction [30]. From knowing these contact formations, other works have utilized them in planning [31] and control frameworks [32]. Fundamentally, the ability to group together different combinations of contacts that define the same or mere similar constraints, can be a powerful tool for defining more capable robot manipulation capabilities.




Fig. 2. The Contact Formation-space (CF-space) of an insertion task can be conceptualized as an explicit organization of different contact formations that share borders according to constraint similarity and possible transitions. Transitions between contact formations are represented when constraints are added or removed to the state of an object. By controlling paths through a CF-space from light (few constraints) to dark (more constraints), the object follows a progression towards insertion. **Note:** This depiction of contact types is not exhaustive and other intermediate formations may be possible.

III. METHODOLOGY

We are interested in solving a generalized peg-in-hole problem by leveraging *compliance-enabled contact formations* – the concept that object constraints can be more easily acquired, broken, and remade when operating within a compliant robot’s “reconfiguration range” [33], [34]. Conceptually, compliance allows the robot to convert a traditionally difficult force control problem into a velocity control problem, as there now exists an ability to “take up the slack” in control uncertainty. A F/T sensor is leveraged to modulate forces closed-loop. This is particularly advantageous for maintaining a contact formation, as now we can more easily ensure the system maintains constraints while operating along a desired path.

Fundamentally, contact formations (CF) represent groupings of discrete positional and physical relationships. Solving the insertion problem can thus be framed as a traversal from a starting, CF_0 , toward a final, CF_n , transitioning through intermediate CFs, i.e., $CF_{path} = \{CF_0, CF_1, \dots, CF_n\}$ (Fig. 2). In this work, we demonstrate the efficacy of this approach by following a fixed path in CF-space.

Formally, transitioning from one CF, CF_n , to a new CF, CF_{n+1} , along a trajectory in CF-space is comparable to adding or removing constraints on the object’s available DOF. Thus, the role of compliance serves to ensure that a single desired CF is maintained while manipulation occurs via force modulation, and transitions are detected when the *motion of the object is impeded by the environment*. Once a variational change in F/T signals is felt by the robot, a contact formation change is detected, and the system has transitioned in CF-space. Conceptually, *move along a single axis until no longer possible, and while maintaining that force, move orthogonally until another constraint is added*.

The theory of controlling compliance-enabled contact formations is that constraints can be continually added until the task of insertion is complete. Notably, any DOF that is not currently constrained can be controlled to perform other tasks without modifying the current CF of the system. For example,

an edge contact between an object and the hole’s surface allows the robot to explore via sliding until a lateral force is detected and thus a new CF is achieved. It is important to note that under this additive force contact process, explicitly estimating the current CF is possible if the starting CF is known and the transitions can be detected via F/T signals.

A. Assumptions and Prior Knowledge

By relying on F/T sensing for environment perception, our algorithm leverages the following priors:

- The object starts in a stable, centered, and upright grasp. This grasp is expected to be maintained.
- CF_0 has 0 constraints, i.e. 6 free DOFs starting in space.
- The hole is somewhere within a known 2D workspace boundary in the XY-plane (e.g., 10cm exploration size).
- Minimal knowledge of the object’s face type is known, but assumed to be convex or near-convex.
- There exists a hole with low contact friction and positive tolerance matching the peg within the workspace.
- Within the workspace, no objects other than the hole, plane, and manipulated object are present. This ensures the CF transitions match the preset CF-space path.

B. CF-space Insertion Algorithm

Our approach follows a desired path in CF-space, as defined through a total of 7 steps. For ease of notation, we set the object frame as an orthogonal frame with the x -axis pointing toward the direction of the object motion on the hole surface plane, and the z -axis pointing downward. For each step, we provide a corresponding implicit CF control target. Although we define a set of values for force modulation, this is for clarity and is in practice robot-specific (see Sec. IV-A).

Let’s assume we can detect the object forces, $F = \{f_x, f_y, f_z\}$ and torques, $T = \{\tau_x, \tau_y, \tau_z\}$, during insertion. Our goal is for the robot to traverse through a desired trajectory of CFs $\{CF_0, \dots, CF_5\}$ (Fig. 1). We outline the algorithm in pseudocode below, and provide a breakdown on a real robotic system in Sec. IV-B.

Insertion Pseudocode

1) Reaching the hole plane

Current constraints: None

Target constraint: $(+f_z, 1.5N)$

Implicit CF targets: (CF_1) point or edge on face

Additional motions:

- Lateral exploration within the workspace area, randomly selecting x/y direction within workspace limits.
- In-hand manipulation to tilt the object toward the direction of the lateral motion, to ensure an edge/face contact, and avoid a face/face contact.

2) Searching for the hole

Current constraints: $(+f_z, 1.5N)$

Target constraint: $(+f_x, 0.7N)$

Implicit CF target: (CF_2) 3-point contact with hole

Additional motions: Lateral exploration as previous step.

3) Wedging

Current constraints: $(+f_z, 1.5N), (+f_x, 0.7N)$

Target constraint: $(+f_y, 0.7N)$

Implicit CF target: (CF_3) 4-point contact with hole

4) Rotational alignment of peg and hole

Current constraints: $(+f_z, 1.5N), (+f_x, 0.7N), (+f_y, 0.7N)$

Target constraint: $(+\tau_z, 0.1N/m)$

Implicit CF target: (CF_4) hinge-type contact

Note: This step applies only to non-cylindrical objects.

5) Correcting upward tilt

Current constraints: $(+f_z, 1.5N)$

Target constraint: $(+f_x, 0N), (+f_y, 0N)$

Implicit CF target: (CF_5) Prismatic joint-type contact

Additional motion: Rotation around x - and y -axes to minimize the accumulated angle between the object and the hand due to fingertip slip during the previous steps. This rotation is performed both with in-hand manipulation and arm motions.

Note: The lateral forces are now minimized to avoid jamming the object. This also helps centering the object in the hole if the peg rotation is not perfectly centered on the object's center of mass. The angle information is provided by extracting the 3D pose of a marker placed on the surface of the object as seen from the palm camera.

6) Inserting peg

Current constraints: $(+f_z, 1.5N), (+f_x, 0N), (+f_y, 0N)$

Exit condition: When the fingertips start touching the hole surface or the object hits the hole bottom, we switch to disengagement. This can be detected as a sharp increase in f_z .

Note: We have already reached the CF state that enables peg insertion, thus the system maintains it while the final free DOF, i.e., the z -axis translation, is controlled to perform the final insertion motion.

7) Detaching hand grip and retracting arm

Action: The hand opens and the arm returns to its origin position in an open loop motion.




Fig. 3. A Yale OpenHand Model O and a 6-axis force/torque sensor are affixed to the end of Barrett WAM manipulator. Inside of the palm of the hand, an in-hand camera setup is fabricated as to monitor the state of the object during manipulation via an AprilTag.

IV. EXPERIMENTS

We evaluate our algorithm on a low-impedance manipulator and a compliant, underactuated hand (Fig. 3). The manipulator, a 7-DOF Barrett WAM, utilizes the *RRTCconnect* algorithm via OMPL [35] for global planning and is controlled locally via a Jacobian-based velocity controller. The manipulator is imprecise due to an inaccurate internal model of its true system dynamics, which further challenges our algorithm's robustness. The end effector, an adapted Yale OpenHand Model O [36], is a passively adaptive hand consisting of three actuators and six total joints. The hand is not equipped with tactile sensors or joint encoders, as to reduce weight and cost, making estimation of the true system state difficult during manipulation [37]. Modifications to the readily available open source design include rounded fingertips, as to facilitate in-hand manipulation, and bearings within the joints. An in-hand manipulation controller is devised, as in [38], and is utilized for fine motor control of object orientation up to $\pm 20^\circ$. Connecting the hand to the arm is a 6-axis ATI force/torque sensor sampled at 30Hz. Finally, a camera is fabricated into the palm of the hand as to enable the monitoring of object poses during manipulation via AprilTags.

The algorithm is validated through a variety of experiments. In our first experiment, we develop a linear pusher, comprised of a leadscrew and a stepper motor, to displace objects along different axes of the gripper's workspace. During this test, we measure the forces exerted on the object by the pusher and note the amount of force the gripper can resist before fingertip slip occurs. Thereafter, we test our algorithm with tight tolerance tasks ($< 0.25\text{mm}$), objects commercially available in a children's insertion toy, and two tasks within the NIST Assembly Task Board #1 [39] (Fig. 4).

A. Quantifying Axial Compliance

First, we are interested in quantifying the compliance of our system. Conceptually, a purely compliant system with fixed contact could comply indefinitely until a hard stop is reached, but practically for grippers, there exists a limit in which the forces applied to an object can be resisted by fingertip contacts. This information is valuable to quantify as it enables

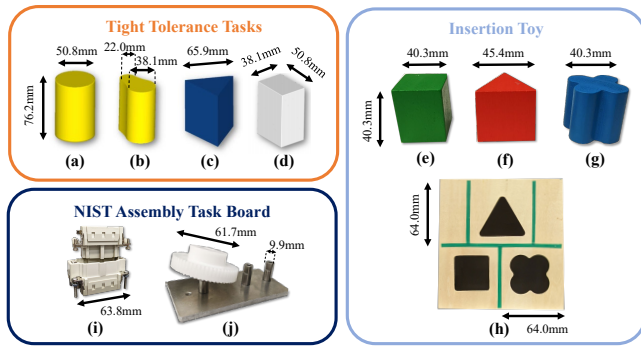


Fig. 4. Tested objects can be classified into three categories: Tight Tolerance Tasks, Insertion Toy, and NIST Assembly Task Board. Objects are referenced according to their face geometries: (a) circle, (b) pear, (c) large triangle, (d) rectangle, (e) cube, (f) small triangle, and (g) clove. Subfigure (h) illustrates the insertion toy’s hole layout with designated search spaces in green. NIST objects (i) and (j) are referred to as the plug and gear, respectively.

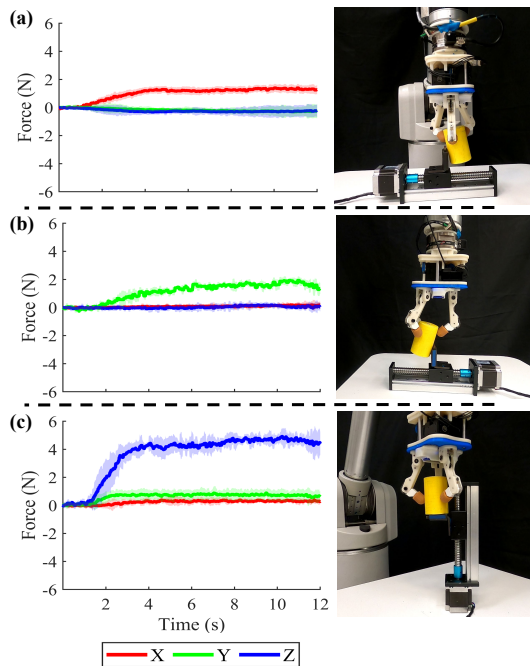


Fig. 5. An object is pushed along the (a) x -, (b) y -, and (c) z -axes with a linear pusher to evaluate the *force plateau* along each dimension, or more specifically, the amount of force the hand can resist before slip occurs.

the system to predict when slippage may occur, which can in turn be used to inform contact formation switching.

Given the linear pusher operating at a velocity of 3mm/s, we push the (a) circle grasped by the manipulator along the x -, y -, and z -axes for a duration of 12 seconds and over six trials. After each trial, the system is systematically reset. We record the forces measured by the F/T sensor via the pusher during each trial, and plot the mean and standard deviation for each evaluation (Fig. 5). Note that for each linear push, a *force plateau* is distinguishable, at $\sim 1.5\text{N}$ for the x - and y -axes (Fig. 5(a)(b)) and at $\sim 4\text{N}$ for the z -axis (Fig. 5(c)). These plateaus correspond to the hand’s ability to resist forces in corresponding directions, or more physically, the point at which static friction of the contact is overcome

and a new hand-object configuration is realized.

This data is valuable in that we are able to quantify the stable contact operating region given an external force. Once a force level is exceeded, sliding occurs and the hand-object state reconfigures until stability is once again realized. We can use these force plateaus for defining modulating forces for our algorithm in Sec. IV-B.

TABLE I
METRICS FOR OBJECT INSERTION EXPERIMENTS

| Obj. | Tol. (mm) | Explore (s) | Insert (s) | Offset ($^{\circ}$) | |
|------|-------------|-------------|------------|-----------------------|------|
| (a) | circle | 0.25 | 37.2 | 28.0 | *n/a |
| (b) | pear | 0.25 | 31.7 | 94.7 | 41.1 |
| (c) | l. triangle | 0.25 | 34.1 | 65.8 | 29.8 |
| (d) | rectangle | 0.25 | 44.2 | 48.1 | 27.1 |
| (e) | cube | 3.0 | 27.3 | 30.8 | 29.2 |
| (f) | s. triangle | 2.1 | 45.1 | 21.2 | 25.0 |
| (g) | clove | 2.6 | 63.2 | 19.4 | 23.7 |
| (i) | plug | 0.9 | 29.8 | 31.0 | 36.2 |
| (j) | gear | 0.1 | 27.9 | 21.4 | *n/a |

*not applicable for pegs with circular faces

B. Tight Tolerance Insertion

Rectangle Case Study: With knowledge of the *force plateaus*, we utilize this information to acquire and maintain contact formations. Leveraging our algorithm (Sec. III), we set target modulating forces to 0.7N, 0.7N, and 1.5N along the x -, y -, and z -axes, respectively, which is $\sim 40\text{-}50\%$ of the maximum force before slippage occurs.

In our first experiment, we attempt to insert the tight tolerance (d) rectangle with an initial 27.1° axial offset along the object’s z -axis (Fig. 6(a)). Notably, the exact location of the hole is unknown, and the system is given a search space of $8\text{cm} \times 8\text{cm}$. The process begins by first finding a downward force, i.e., the first constraint, and choosing an exploration direction (in this case $+f_x$) within the search space until an additional constraint is detected (transition 3). The additional force spike on the x -axis signals the hole’s perimeter and a force transition begins. This process is evidenced in Fig. 6(b), where forces are continually added and modulated around setpoints for states 1-7 (Sec. III). Note that although compliant, the WAM robot has difficulty modulating velocities, and thus forces, due to controllability. Moreover, forces in Fig. 6(b) are represented in the world frame to represent rotations during steps 4-5 and are smoothed for clarity, so force spikes are not represented.

All Objects: Generalization is underscored by our ability to insert objects of varying convex, or near-convex, geometries – specifically pegs with circle-, pear-, and triangle-shaped faces (Fig. 4). The result of our evaluations with tolerances, exploration times, insertion times, and offset degrees is presented in Table I (a-d). During evaluation, we noted that the (a) circle was the easiest, as it did not require z -axis offset control. The other objects (b-d), were more difficult and posed various challenges. First, the initial offset was different for each, ranging from 27.1° to 41.1° . Objects (c) and (d) had sharp edges which encouraged jamming, whereas object (b) presented difficulty due to its non-convexity.

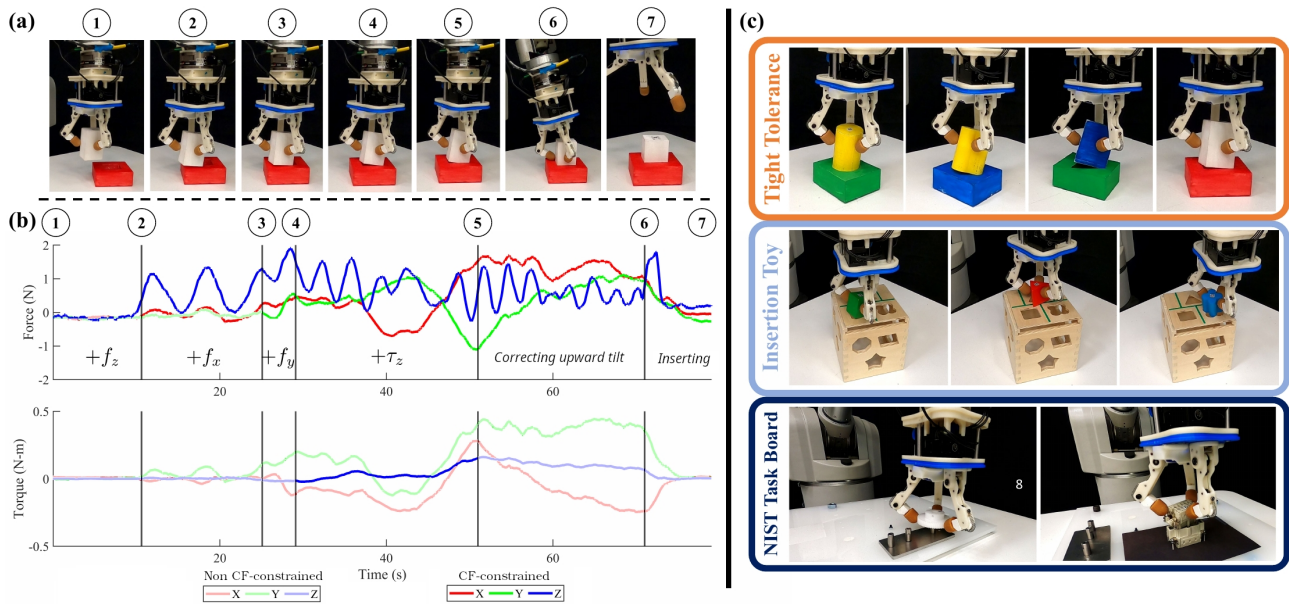


Fig. 6. The progression of insertion is depicted in (a), where an object goes from free space (0 constraints) to inserted (5 constraints) into its goal configuration. (b) During this process, forces are modulated and added through different steps in the insertion task (forces are smoothed and placed in the world frame for clarity). (c) We evaluate this algorithm with tight tolerance tasks, child insertion toys, and objects from the NIST Task Board.

Overcoming these challenges, we were able to complete insertions successfully with objects of $<0.25\text{mm}$ tolerance and without prior knowledge of the hole pose (Fig. 6(c)).

C. Open World Insertion Tasks

Beyond our tight tolerance evaluations, we were interested in applicability of our method to open world tasks. Our first experiment is with a commercially available children’s toy consisting of different object geometries and with a hole tolerance of approximately 2.5mm . Similar to the previous experiments, the pose of the hole is unknown, and the search area is now confined to a $6.4\text{cm} \times 6.4\text{cm}$ space (denoted in green in Fig. 4(h)). Here, we attempt to challenge the exploration component of our algorithm, ensuring the object started at the edge of the search space. All insertions were successfully completed, with the longest exploration phase of 63.2 seconds for the (g) clove. Interestingly, the non-convexity of the clove did not complicate the insertion process as much as originally believed, as the round edges of the object helped limit jamming from occurring.

Our final experiment evaluated plug and gear insertion from the NIST Assembly Task Board (Fig. 6(c)). The gear task is interesting in that the peg-in-hole paradigm is transformed instead into a hole-on-peg schema. This was not a problem for our system, as the search pattern was instead completed on the bottom of the hole, i.e., the gear, instead of by using a peg. Similarly, the plug insertion with a relaxed plunger spring was completed with ease, which started with an z – axis offset of 36.2° (Table I). These tests underscore the practicality and generalizability of our method, which is further showcased in the supplementary video.

V. DISCUSSIONS AND FUTURE WORK

This work presents a method that leverages *compliance-enabled contact formations* as a step towards generalized,

tight tolerance insertion for robots. Our algorithm is simple, yet mechanically grounded, and exploits the concept of reducing uncertainty through additive contact constraints, i.e., manipulation funnels. Notably, in this work we do not need to utilize costly learning frameworks or system-specific, idealized analytical models – this method is effective yet did not require a single equation – to describe in detail.

The authors are excited about this preliminary exploration, as our experiments illustrate validity of our approach for future applications. As an attempt to not overclaim contributions, the authors want to be forthcoming on known limitations that will warrant future investigation:

- 1) We cannot claim theoretical guarantees that an insertion will always be successful. In simulation, we verified that convex objects of non-negative tolerances should always succeed, but guarantees are not as clear for all non-convex circumstances.
- 2) Hole geometry requires a low-frictional “platform” for object exploration in order to find constraints. If this platform has variable friction, contact states may be transitioned prematurely and cause failure.
- 3) Negative tolerance insertion would be unlikely due to the maximum forces the system can apply (Sec. IV-A). A redesign of the end effector to apply a greater amount of force ($\sim 5\text{-}10\text{N}$) would be beneficial for varied tasks.

Overall, the authors believe compliance will continue to prove invaluable for future advancements in robot manipulation. By investigating how to build more capable end effectors and by developing robust control strategies for non-convex object insertion, our method should prove to extend to a vast array of everyday insertion tasks for service robots of the future. Please find a complete overview of our motivation, approach, and experiments in the supplementary video.

REFERENCES

- [1] J. Bimbo, A. S. Morgan, and A. M. Dollar, "Force-based simultaneous mapping and object reconstruction for robotic manipulation," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4749–4756, 2022.
- [2] J. C. Trinkle, J.-S. Pang, S. Sudarsky, and G. Lo, "On dynamic multi-rigid-body contact problems with coulomb friction," *Journal of Applied Mathematics and Mechanics*, vol. 77, no. 4, pp. 267–279, 1997.
- [3] D. J. Montana, "Contact Stability for Two-Fingered Grasps," *IEEE Transactions on Robotics and Automation*, vol. 8, no. 4, pp. 421–430, 1992.
- [4] D. L. Brock, "Enhancing the dexterity of a robot hand using controlled slip," in *Proceedings 1988 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 249–251.
- [5] 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.
- [6] M. Skubic and R. Volz, "Identifying contact formations from sensory patterns and its applicability to robot programming by demonstration," in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems. IROS '96*, vol. 2, 1996, pp. 458–464 vol.2.
- [7] H. Park, J. Park, D. Lee, J. Park, M. Baeg, and J. Bae, "Compliance-Based Robotic Peg-in-Hole Assembly Strategy Without Force Feedback," *IEEE Transactions on Industrial Electronics*, vol. 64, no. 8, pp. 6299–6309, 2017.
- [8] M. P. Polverini, A. M. Zanchettin, S. Castello, and P. Rocco, "Sensorless and constraint based peg-in-hole task execution with a dual-arm robot," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 415–420.
- [9] C. Choi, J. Del Preto, and D. Rus, "Using Vision for Pre- and Post-grasping Object Localization for Soft Hands," in *2016 International Symposium on Experimental Robotics*, D. Kulić, Y. Nakamura, O. Khatib, and G. Venture, Eds. Cham: Springer International Publishing, 2017, pp. 601–612.
- [10] N. Chavan-Dafle, A. Rodriguez, R. Paolini, B. Tang, S. Srinivasa, M. Erdmann, M. T. Mason, I. Lundberg, H. Staab, and T. Fuhlbrigge, "Extrinsic Dexterity: In-Hand Manipulation with External Forces," in *Proceedings of (ICRA) International Conference on Robotics and Automation*, May 2014, pp. 1578 – 1585.
- [11] Y. Xu and R. Paul, "A robot compliant wrist system for automated assembly," in *Proceedings., IEEE International Conference on Robotics and Automation*, 1990, pp. 1750–1755 vol.3.
- [12] *Remote Center of Compliance Reconsidered*, ser. International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, vol. Volume 2A: 24th Biennial Mechanisms Conference, 08 1996.
- [13] C. C. Kemp, A. Edsinger, and E. Torres-Jara, "Challenges for robot manipulation in human environments [Grand Challenges of Robotics]," *IEEE Robotics & Automation Magazine*, vol. 14, no. 1, pp. 20–29, 2007.
- [14] 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.
- [15] T. Tang, H. Lin, Y. Zhao, Y. Fan, W. Chen, and M. Tomizuka, "Teach industrial robots peg-hole-insertion by human demonstration," in *2016 IEEE International Conference on Advanced Intelligent Mechatronics (AIM)*, 2016, pp. 488–494.
- [16] J. Xu, Z. Hou, W. Wang, B. Xu, K. Zhang, and K. Chen, "Feedback Deep Deterministic Policy Gradient With Fuzzy Reward for Robotic Multiple Peg-in-Hole Assembly Tasks," *IEEE Transactions on Industrial Informatics*, vol. 15, no. 3, pp. 1658–1667, 2019.
- [17] 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.
- [18] 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.
- [19] G. Schoettler, A. Nair, J. A. Ojea, S. Levine, and E. Solowjow, "Meta-Reinforcement Learning for Robotic Industrial Insertion Tasks," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 9728–9735.
- [20] S. Levine, N. Wagener, and P. Abbeel, "Learning contact-rich manipulation skills with guided policy search," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 156–163.
- [21] W.-C. Chang, "Robotic assembly of smartphone back shells with eye-in-hand visual servoing," *Robotics and Computer-Integrated Manufacturing*, vol. 50, pp. 102–113, 2018.
- [22] Y. She, S. Wang, S. Dong, N. Sunil, A. Rodriguez, and E. Adelson, "Cable Manipulation with a Tactile-Reactive Gripper," in *Robotics: Science and Systems*, 2020.
- [23] J. Xu, Z. Hou, Z. Liu, and H. Qiao, "Compare Contact Model-based Control and Contact Model-free Learning: A Survey of Robotic Peg-in-hole Assembly Strategies," 2019.
- [24] K. Van Wyk, M. Culleton, J. Falco, and K. Kelly, "Comparative Peg-in-Hole Testing of a Force-Based Manipulation Controlled Robotic Hand," *IEEE Transactions on Robotics*, vol. 34, no. 2, pp. 542–549, 2018.
- [25] 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.
- [26] M. A. Lee, Y. Zhu, K. Srinivasan, P. Shah, S. Savarese, L. Fei-Fei, A. Garg, and J. Bohg, "Making Sense of Vision and Touch: Self-Supervised Learning of Multimodal Representations for Contact-Rich Tasks," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 8943–8950.
- [27] K. Zakka, A. Zeng, J. Lee, and S. Song, "Form2Fit: Learning Shape Priors for Generalizable Assembly from Disassembly," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 9404–9410.
- [28] C. C. Beltran-Hernandez, D. Petit, I. G. Ramirez-Alpizar, and K. Harada, "Variable Compliance Control for Robotic Peg-in-Hole Assembly: A Deep-Reinforcement-Learning Approach," *Applied Sciences*, vol. 10, no. 19, p. 6923, Oct 2020. [Online]. Available: <http://dx.doi.org/10.3390/app10196923>
- [29] U. Thomas and F. M. Wahl, *Assembly Planning and Task Planning — Two Prerequisites for Automated Robot Programming*, 2011, pp. 333–354.
- [30] A. Farahat, B. Graves, and J. Trinkle, "Identifying contact formations in the presence of uncertainty," in *Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human Robot Interaction and Cooperative Robots*, vol. 3, 1995, pp. 59–64 vol.3.
- [31] X. Cheng, E. Huang, Y. Hou, and M. T. Mason, "Contact mode guided sampling-based planning for quasistatic dexterous manipulation in 2d," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 6520–6526.
- [32] S. Cabras, M. E. Castellanos, and E. Staffetti, "Contact-state classification in human-demonstrated robot compliant motion tasks using the boosting algorithm," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 40, no. 5, pp. 1372–1386, 2010.
- [33] K. Hang, A. S. Morgan, and A. M. Dollar, "Pre-grasp sliding manipulation of thin objects using soft, compliant, or underactuated hands," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 662–669, 2019.
- [34] A. S. Morgan, K. Hang, B. Wen, K. Bekris, and A. M. Dollar, "Complex in-hand manipulation via compliance-enabled finger gaiting and multimodal planning," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4821–4828, 2022.
- [35] I. A. Sucan, M. Moll, and L. E. Kavraki, "The open motion planning library," *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, 2012.
- [36] R. Ma and A. Dollar, "Yale OpenHand Project: Optimizing Open-Source Hand Designs for Ease of Fabrication and Adoption," *IEEE Robotics Automation Magazine*, vol. 24, no. 1, pp. 32–40, 2017.
- [37] K. Hang, W. G. Bircher, A. S. Morgan, and A. M. Dollar, "Hand-object configuration estimation using particle filters for dexterous in-hand manipulation," *The International Journal of Robotics Research*, vol. 39, no. 14, pp. 1760–1774, 2020.
- [38] 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.
- [39] K. Kimble, K. Van Wyk, J. Falco, E. Messina, Y. Sun, M. Shibata, W. Uemura, and Y. Yokokohji, "Benchmarking protocols for evaluating small parts robotic assembly systems," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 883–889, 2020.