

Autonomous Robotic Assembly: From Part Singulation to Precise Assembly

Kei Ota^{1,2}, Devesh K. Jha^{1,2}, Siddarth Jain¹, Bill Yerazunis¹, Radu Corcodel¹, Yash Shukla¹,
Antonia Bronars¹, Diego Romeres¹

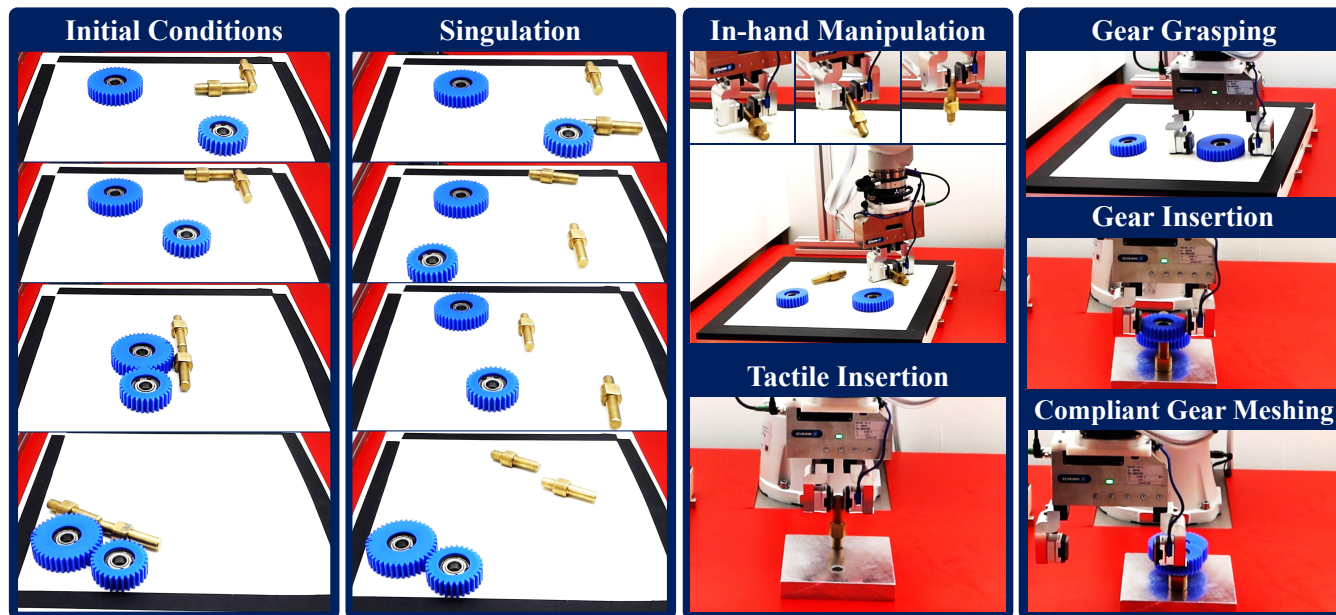


Fig. 1: **Autonomous Robotic Assembly:** We present an autonomous robotic assembly system that can assemble a gear box from any given initial condition, as shown in the first column. The assembly system can reason about grasp feasibility and slide selected objects out of a clutter to create grasp affordances for the assembly parts. Then it performs the pose manipulation and grasping required for the downstream assembly task. Finally, using various different controllers, it performs the required insertion and meshing of gears to assemble a functioning gear box. The proposed system works in a closed-loop fashion, where it can deliberate on the success and failure of individual steps and react accordingly.

Abstract—Imagine a robot that can assemble a functional product from the individual parts presented in any configuration to the robot. Designing such a robotic system is a complex problem which presents several open challenges. To bypass these challenges, the current generation of assembly systems is built with a lot of system integration effort to provide the structure and precision necessary for assembly. These systems are mostly responsible for part singulation, part kitting, and part detection, which is accomplished by intelligent system design. In this paper, we present autonomous assembly of a gear box with minimum requirements on structure. The assembly parts are randomly placed in a two-dimensional work environment for the robot. The proposed system makes use of several different manipulation skills such as sliding for grasping, in-hand manipulation, and insertion to assemble the gear box. All these tasks are run in a closed-loop fashion using vision, tactile, and Force-Torque (F/T) sensors. We perform extensive hardware experiments to show the robustness of the proposed methods as well as the overall system. See supplementary video at <https://www.youtube.com/watch?v=cZ9M1DQ230I>.

¹All authors are with Mitsubishi Electric Research Labs (MERL), Cambridge, MA, USA. {ota, jha, sjain, yerazunis, corcodel, shukla, bronars, romeres}@merl.com

²Authors with equal contribution.

I. INTRODUCTION

Designing robots that can achieve human-like dexterity and reasoning while performing complex, long-horizon tasks has been the long-standing goal of robotics. While robots have been getting very good at performing repetitive pick-and-place kind of operations, achieving reliable & high degrees of dexterity remains elusive. Consequently, robots require a lot of structure to perform tasks that require complex and long-horizon reasoning. Imagine a factory floor where assembly blueprints and parts are presented to a robot in arbitrary configurations, and the robot can autonomously assemble the desired product [1]. While this task is effortless for humans, it is quite challenging for robotic systems. In this paper, we study a simplified assembly problem to understand the challenges of the underlying long-horizon manipulation.

Assembly is arguably the single biggest application of robots in current society. However, this generally requires creating a complex system for singulation and manipulating the pose of parts. Robots are simply used as machines to

perform a very precise pick-and-place task. Robots with intelligent manipulation skills can simplify the design and operation of these systems on factory floors. However, manipulation is a complex problem from the aspects of planning, sensing, and control. Despite being an active area of research for decades, it still remains a widely open and challenging problem to all of robotics community [2]. Through a simplified example of the assembly of a gear box (see Fig. 1), we highlight the challenges associated with designing agents that can achieve high levels of dexterity and autonomy with simple hardware.

Robotic assembly is a challenging, long-horizon manipulation problem. Thus, performing autonomous assembly requires reasoning about various sub-goals given a particular instance of the environment. Apart from reasoning about a feasible plan, it also presents the challenge of achieving very high accuracy and precision to make a tight tolerance assembly feasible. Another major challenge is to allow robustness to various initial conditions so that the system can always achieve the desired goal regardless of the initial condition. To demonstrate all these challenges, we present the assembly of a gear box shown in Fig. 1. Under some mildly restrictive assumptions, we present solutions to the problems in order to create an autonomous agent that can perform assembly with a perfect success rate. We make use of three different kinds of sensors – an RGBD camera, vision-based Gelsight tactile sensors, and a six-axis F/T sensor mounted on the wrist of the manipulator.

As shown in Fig. 1, our system can perform part singulation, part manipulation, as well as assembly in a closed-loop fashion while using an array of controllers and perception methods for individual steps. In particular, this paper has the following contributions:

- 1) We present a benchmark assembly task which can be used to test the dexterity and precision of related manipulation techniques.
- 2) We present an autonomous assembly system which can combine various sensing modalities such as vision, touch, and F/T while running in closed-loop feedback. The proposed system is extensively tested on hardware and has been shown to achieve an almost perfect success rate.

To the best of our knowledge, an autonomous, multi-modal system that can perform high-precision multistep assembly task has not been presented earlier in the literature. We hope that the proposed system can be used as a benchmark in the near future for testing manipulation algorithms and methods.

II. RELATED WORK

Robotic manipulation primitives. Autonomous assembly has been one of the most challenging and widely studied problems in robotics in the past several decades [3], [4], [5], [6]. Several challenges related to assembly tasks have been widely studied in the literature, including grasping in clutter [7], [8], part singulation [9], pose estimation [10], [11], [12], [13], part insertion [14], etc. Similarly, in-hand manipulation and non-prehensile manipulation have also

been widely studied in the literature [15], [16], [17], [18], [19], [20]. However, all these work would require an external point of contact for manipulation. In-hand pose estimation using tactile sensors has also been explored recently [12], [13]. Similarly, insertion using force or tactile feedback has also been explored [14], [21]. However, these methods could struggle for cases with very tight tolerances required during assembly. Furthermore, these primitive skills have not been composed for long horizon tasks with sensory feedback.

Robotic assembly benchmarks. Long-horizon manipulation involves many challenges, including the need to dynamically switch skills based on observations and the ability to recover from failures. Recent robotic benchmarks, such as furniture assembly [3], [5], [22], irregular block stacking [23], and manufacturing tasks [5], [6], [22], provide reproducible environments to evaluate the manipulation performance. Our task is inspired by the NIST robotic manipulation benchmark [6], employing a simplified insertion tolerance of approximately 1 mm as opposed to the submillimeter tolerance (0.029 mm and 0.005 mm) of the standard board. This simplification was explored similarly in the NIST-i benchmark, but focused only on gear insertion [24]. Previous approaches to solve these assembly benchmarks typically omit tactile feedback, relying heavily on vision sensors, leading to issues with occlusion or kinematic constraints imposed by extra links for vision sensors mounted on robot links. Although tactile feedback has been explored for pick-and-place tasks [25], we believe that its integration with diverse visuotactile techniques for the design of autonomous assembly systems represents a significant and novel contribution to the existing literature.

III. PROBLEM STATEMENT

While the autonomous assembly problem in its entirety is a very complex problem, we make certain assumptions to solve it in a limited scope. This section briefly describes it.

A. Task Assumptions

The main objective of the paper is to evaluate the manipulation capability and robustness of the proposed system. In order to evaluate these, we make some simplifying assumptions to limit the scope of this study.

- 1) The positions of the holes are known and fixed.
- 2) The surface for the manipulation of parts and the base plate for assembly is perfectly leveled (see Fig. 1).
- 3) The assembly order of the parts is known.

Assumption (1) is not very restrictive. We can easily relax this by using a vision system to track holes, as shown earlier in our previous work [21]. The tray is kept leveled so that the parts do not slide under gravity. However, in the most general case, the manipulation primitives can be generalized to an inclined one. Assumption (3) can also be relaxed by using a Large Language Model [26] or a graph planner [1].

B. System Description

Robot platform. The MELFA RV-5AS-D Assista robot, a 6 DoF collaborative robot, is used in this study. Two

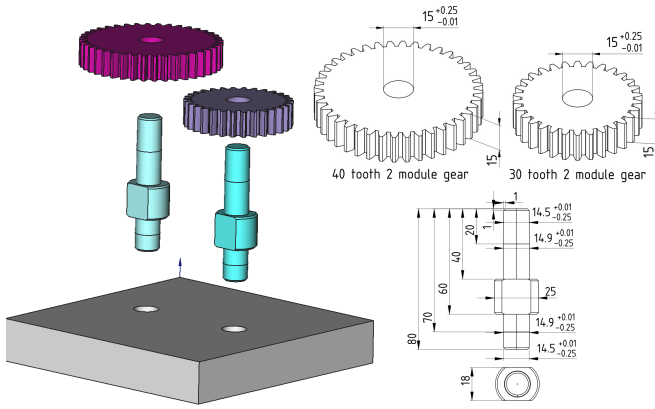


Fig. 2: This figure shows the assembly parts with two identical pegs and two gears. Accurate dimensions of the parts are provided up to machining tolerances. The holes in the base plate are 15 mm in diameter and 70 mm apart.

tactile sensors are mounted on the WSG-50 gripper, and a Force-Torque (F/T) sensor is mounted on the wrist (see the supplementary video for the system). The robot fingers also have a silicone gel padding at the bottom, which is used during contact interactions.

Tactile sensor. We use commercially available GelSight Mini tactile sensors, which provides 320×240 compressed RGB images at a rate of approximately 25 Hz, with a field of view of 18.6×14.3 millimeters.

Vision sensor. An Intel Realsense L515 camera is used as a vision sensor for the RGB-D input of the scene. This vision sensor is used to observe the table-top used for manipulation of assembly parts. The assembly location is known to the robot and is not observed by the vision sensor.

Assembly Parts. Figure 2 shows the assembly parts. The gear box consists of two pegs, a small and a large gear, and a baseplate that has two holes for the pegs. The dimensions of the individual parts are shown in Figure 2. The tolerance for the various insertion tasks is approximately 1 mm. While the tolerances are much higher than what is expected in assembled products [6], we believe the proposed approaches could be used for tighter assembly. We denote the parts by p_1, p_2, g_l, g_s for the pegs, the large gear, and the small gear, respectively.

Assembly Task. Imagine a task where assembly parts are presented in any given configuration in the robot workspace. The robot has to assemble a functional gear box by manipulating these parts. As shown in the first column of Fig. 1, we present the parts in any planar configuration to the robot. The task for the robot is to perform the required manipulation and assemble a functioning gear box using these parts. It is noted that the parts are always presented in a workspace that is visible to an overhead camera (the black boundaries in Fig. 1 show this workspace). As specified earlier, the location of the base plate where the assembly has to be performed is known to the robot.

What makes this task challenging? There are multiple tasks that an agent needs to reason about for the assembly

task. The first task is to reason about grasp affordances when the parts are presented in a clutter. In case the parts cannot be grasped directly, the robot has to reason about how it can move them apart or singulate them for grasping. The second challenge is for the robot to reason about how it should grasp the individual parts so that it is suitable for the desired downstream assembly task. Consequently, how should the robot achieve the precision required for assembly when the grasp pose is not fixed? Furthermore, how can the system be designed to achieve extremely robust performance? We address some of these challenges in this paper.

IV. SYSTEM DETAILS

A system-level flow chart of the designed assembly system is shown in Fig. 3. As could be seen, most steps use multi-modal feedback during execution of the task. We believe that this provides robustness to the system to perform the task reliably.

A. Vision System

Achieving reliable robotic grasping hinges on accurately determining the pose of target objects within the input data. Recent advancements in deep learning-based object pose estimation [10], [11], [27] have garnered attention for their ability to achieve centimeter-level accuracy across diverse datasets. Nevertheless, these methods are often insufficient for meeting the stringent requirements of precise robotic manipulation tasks at the sub-millimeter level. Moreover, in our scenario, assembly parts are initially placed in arbitrary configurations, involving multiple instances of the same object type.

In our settings, the workspace of the robot for autonomous assembly is constrained in that the parts lie on a planar workspace, and thus, the grasp synthesis can be constrained from one direction. Therefore, the synthesis of the grasp of each object can define a pose output, $g = (p, \phi, q)$, where $p = (x, y, z)$ is the position of the object, ϕ is the singular rotation angle around the vertical axis, and q represents the label of the object class. We adopt a hybrid approach that involves deep learning for part detection and segmentation, combined with analytical pose computation. The network architecture is built on a backbone convolutional neural network architecture for feature extraction. We use a feature pyramid network (FPN) based on ResNet-50 [28]. Mask-RCNN [29] is based on region proposals that are generated through a region proposal network. It adds a network head, a fully convolutional neural network, to produce the desired instance segmentation. Finally, mask and class predictions are decoupled; the mask network head predicts the mask independently from the network head predicting the class. Typically, this involves using a multitask loss function $L = L_{cls} + L_{bbox} + L_{mask}$. For training the network, we perform transfer learning from the MS COCO dataset pre-trained weights in a supervised manner. We capture 380 images of size 640×480 under different initial conditions and annotate the data to indicate segmentation pixels and class labels. A 5:1 ratio was used to divide the dataset into training and

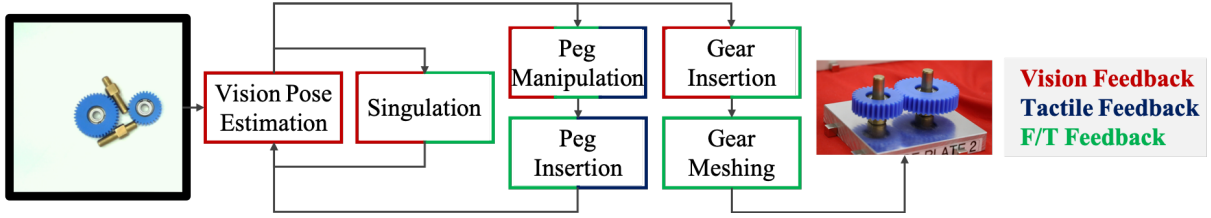


Fig. 3: System level overview of the assembly controller. The color codes indicate the feedback modality for the particular related operation. Multiple color blocks for the same operation indicate that multiple sensors are used for feedback and/or controller design for the particular operation. [Best seen in color].

validation sets. We selected the maximum iterations of 4000. We identify the resulting segmentation masks for parts with the network prediction. The detected segmentation masks are utilized with depth estimation to compute the corresponding registered point cloud data points. These clusters are then processed to perform analytic pose estimation [30] for grasping and part singulation. We analyze the principal components of a covariance matrix created from the nearest neighbours of the point \mathbf{p}_i in a local neighbourhood of size n . We denote the centroid $\bar{\mathbf{p}}$ of the neighbours of \mathbf{p}_i , given by Eq. (1). The covariance matrix $C \in \mathbb{R}^{3 \times 3}$ is computed as in Eq. (2) and the eigenvector problem is presented in Eq. (3). The eigenvalues $\lambda_j \in \mathbb{R}$ and the eigenvectors \vec{v}_j form an orthogonal frame, corresponding to the principal components of the point set in the neighbourhood n .

$$\bar{\mathbf{p}} = \frac{\sum_{k=1}^n \mathbf{p}_k}{n}, \quad (1)$$

$$C = \frac{1}{n} \sum_{i=1}^n (\mathbf{p}_i - \bar{\mathbf{p}})(\mathbf{p}_i - \bar{\mathbf{p}})^T, \quad (2)$$

$$C \cdot \vec{v}_l = \lambda_l \cdot \vec{v}_l, l \in \{0, 1, 2\}. \quad (3)$$

The eigenvalue λ_l quantifies the variation of \mathbf{p}_i along the direction of the corresponding eigenvector and the smallest eigenvalue λ_0 ($\lambda_0 \leq \lambda_1 \leq \lambda_2$) corresponds to the variation along the surface normal vector (i.e. the eigenvector \vec{v}_0) of the surface patch of neighbourhood size n . The object pose is then represented with $\bar{\mathbf{p}}$ and the singular rotation angle around the vertical axis ϕ computed from the angle between the minor eigenvector and the horizontal axis. The object class label q is obtained from the Mask-RCNN output.

B. Part Singulation

The assembly parts are placed in any initial configuration in a clutter to allow generalization. Consequently, the robot first has to reason about grasp affordance of the assembly parts. We design a model-based method using random-shooting to reason about part singulation.

Model-based Singulation using Random Shooting. Given the pose of all the objects from the vision system, we create the manipulation environment in simulation using the MuJoCo physics engine [31]. We use a simple random shooting method in this simulation environment to optimize the robot’s action using a suitable cost function. More concretely, we sample a fixed number of actions and choose the one that minimizes the cost function. The action is described

as $a = \{o, dx, dy\}$, where $o \in \{p_1, p_2, g_l, g_s\}$ is the object to interact with and dx, dy are the relative planar positions on the manipulation surface to which the selected object o is moved. We sample actions from a uniform distribution independently, specifically setting $dx, dy \sim U(-0.3, 0.3)$ [m] to sample a direction and a distance to move.

While pushing provides a good manipulation primitive to move the object on the table top surface, it is susceptible to uncertainty in friction and point of contact [32]. To allow better robustness, we use the sliding manipulation primitive to move objects on the table-top [33]. To imitate the sliding behavior in the simulation, we divide the direction dx, dy into smaller steps and check the collision during the move of the object. We then compute the cost of the sampled action using the following function:

$$c = -\lambda_1 \mathbb{I}_{\text{collision}} + \lambda_2 \mathbb{I}_{\text{gear}} - \lambda_3 d_{\text{obj}} - \lambda_4 d_{\text{center}}, \quad (4)$$

where $\mathbb{I}_{\text{collision}}$ and \mathbb{I}_{gear} are indicator functions of whether a collision between the selected object and other objects happens during sliding and whether a gear is selected. d_{obj} and d_{center} are the distance between the selected object and others before and after the sliding operation and the distance between the selected object from the center of the manipulation surface, respectively. The first term encourages optimization to avoid collisions, and the second one encourages the agent to interact with the gears since they have larger contact surfaces than pegs, resulting in fewer chances of losing contact during sliding on the real system. The third term encourages the selected object to have more affordance after the sliding operation. The final term avoids the part from dropping from the manipulation surface. We sample $N^{\text{RS}} = 100$ actions and choose the one with the minimum cost as the action to apply in the real system.

Singulation execution by sliding. We do not consider the frictional interaction for sliding in simulation to simplify the sim2real transfer. Rather, we design a reliable and compliant sliding controller that can execute sliding for simple linear paths (see Fig. 4 right). We use the object pose estimate from vision feedback to establish contact with the chosen object and then slide the object using a predefined normal contact force. The normal force is optimized on the real system to minimize slipping between the object and the manipulator end-effector. We terminate the singulation process if all pegs can be grasped without colliding with other objects. This is checked in the physics engine based on the vision feedback after every sliding action performed by the robot.

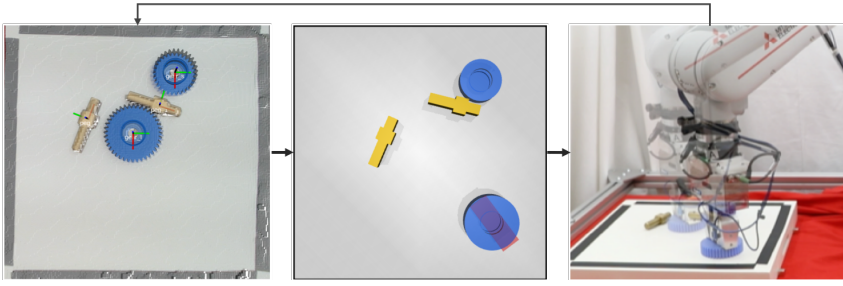


Fig. 4: Singulation procedure. Given part pose from vision module (left), we first reconstruct the manipulation environment in the MuJoCo physics engine and generate an action using the random shooting method (middle). We then apply the action in the real system by sliding the target object using a suitable impedance controller with the F/T sensor (right).

C. Peg Grasping and Manipulation

Once the parts are singulated, the objective is to grasp the singulated parts and perform the downstream assembly task. In order to assemble the gear box, the peg has to be inserted in the correct hole. Thus, apart from grasping the peg, it has to be re-oriented so that it could be inserted (see the third column of Fig. 1). There are several possibilities one can choose from. There is a wealth of work on using non-prehensile manipulation using extrinsic contacts that could be used to re-orient the peg [15], [16], [17]. However, non-prehensile manipulation is generally difficult to track online as one would need to estimate the slip, which requires a more careful design of tactile sensors and fingers.

The other option we investigate is in-hand manipulation [20], [18], as it could be tracked easily using tactile sensors co-located at the gripper fingers. We search for a grasp such that the peg re-orient in grasp as the peg is lifted from the table-top. In particular, we find a parameterized grasp denoted as $g_p = g_p(l_p, f_p)$, where l_p is the grasp location in the local peg frame and f_p is the grasping force that results in the desired peg re-orientation. This is shown in Fig. 5.

D. In-hand Pose Estimation

One of the biggest challenges with assembly is to localize the hole w.r.t. the peg after the robot has been able to grasp the peg. This presents a big challenge as with variation in grasp, achieving very tight localization between the peg and the hole could be challenging [14]. Previous approaches to performing tight tolerance insertion have tried to correct relative pose using force or tactile observations during contact formation with the environment (see [14], [21], [34]). However, allowing grasp to be a variable for the underlying problem makes learning challenging. An alternative approach could be to localize the peg in the grasp using a reliable in-hand pose estimation method. Vision-based tactile sensors co-located at the gripper fingers are very convenient for this task and have previously been explored for similar tasks [13], [12]. We train an in-hand pose estimation model that can predict the relative pose of the peg w.r.t. the local gripper frame. The model takes two tactile images obtained from

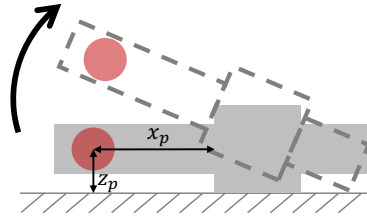


Fig. 5: For performing the peg manipulation, we design a grasp that allows in-hand rotation as the peg is grasped from the table-top. The grasp is defined as $g_p = g_p(l_p, f_p)$, where $l_p = (x_p, z_p)$ as shown in the figure.



Fig. 6: Example tactile images when grasping a peg, which is inputted to the in-hand pose estimation model to estimate the grasp error and correct it before insertion.

the tactile sensors attached to the two fingers $I^{\text{left}}, I^{\text{right}}$ (see Fig. 6), and estimates the pose transformation in the local gripper frame.

To train the model, we collect tactile images of the peg in the grasp by introducing displacement from a uniform distribution of $dx \sim U(-10, 10)$ [mm] in the local gripper frame. We then train a ResNet18 [28] model with regression loss to predict the displacement in the local gripper frame. Since the hole position is known in the local gripper frame, we can perform insertion once the peg is localized in the local gripper frame.

E. Gear Insertion

After peg insertion, the robot needs to insert the gears to assemble the gear box. Similar to the peg insertion problem, grasping the gears using vision-based localization introduces grasp uncertainty, and thus the robot needs to localize the gear in the local grasp (or gripper) frame. Even though the tactile sensors provide us with very high-resolution object features, tactile images do not provide enough information for in-hand localization of gears. This is because the gear teeth are identical and do not have any additional features to estimate the pose of the gear due to their limited field of view. Consequently, we make use of the force observations up on contact formation between the gear and the peg to localize the gear w.r.t. the peg. This is similar to the method presented earlier in [21]. Although one can train such a policy using tactile sensors (similar to [14]), we make use of six-axis force observations from the F / T sensor mounted on the robot wrist.

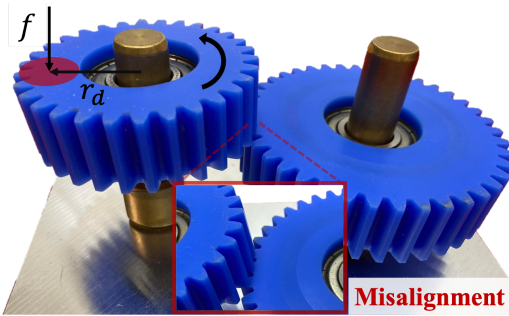


Fig. 7: This figure shows the misalignment between the gears during meshing and the controller designed to align the gear teeth by using a suitable compliance control input $u = u(f, r_d)$. This compliance controller is used to rotate the smaller gear while creating relative movement between the gears so that eventually the teeth align.

To train the gear insertion policy, we collect data by introducing known error between the gear center and the peg. We then measure the six-axis force observations up on contact formation between the peg and the gear with known pose error. The sensor can return measurements at 280 Hz. To collect training data, we add the known pose error between the gear and the peg, such as $dx \sim \{-4, 4\}$ [mm]. The robot then makes contact with the peg, and we collect data for $T = 3$ seconds for each insertion attempt. After collecting the raw signals, we compute the moving average over 70 time steps, since the raw F/T signals are noisy.

F. Gear Meshing

The tolerances of the parts are such that simply inserting the gears does not result in alignment of the gear teeth (see Fig. 7). Consequently, the robot has to perform fine manipulation so that the gear teeth could be meshed. We first investigated a force-feedback strategy to understand the alignment of the gear teeth. However, we observed that there was no reliable trend between the observed force signature and teeth alignment. Thus, we propose an open-loop compliant gear meshing method. In particular, we designed a compliant force controller $u = u(f, r_d)$ where f is the meshing force and r_d is the radial distance from the gear center at which the force is applied, as shown in Fig. 7. This compliant controller has the task of rotating the small gear while minimally disturbing the large gear. Note that the large gear will still rotate due to friction between the two gears. However, the compliant controller simply needs to create relative movement between the two gears. Due to the relative movement between the gears, the gear teeth eventually align.

V. RESULTS

We perform rigorous experiments to understand and address the following questions.

- What is the success rate of the individual system components? What are some of the observed failure cases?
- What is the success rate and performance of the overall system?

TABLE I: **Vision Pose Estimation.** Validation results for both detection and instance segmentation. APb and APm denotes detection and instance segmentation results.

APb	APb@0.5	APb@0.75	APm	APm@0.5	APm@0.75
93.32	99.50	99.08	93.37	99.50	99.08

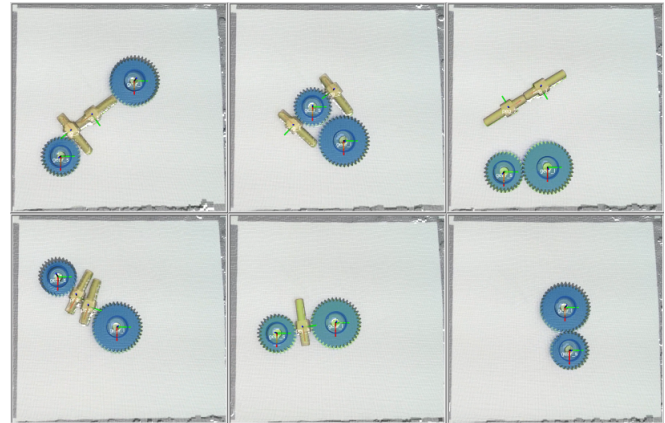


Fig. 8: **Vision Pose Estimation.** Qualitative results from the vision system. CAD model overlaid on point cloud for pose visualization.

In the rest of the paper, we present experiments to address these questions. We skip detailed analysis of compliance controller for peg manipulation and gear meshing for brevity and will be presented in an extended version of the paper.

A. Vision Pose Estimation

We report detection and instance segmentation results for validation (Table I). We use COCO evaluation metrics AP (averaged over IoU thresholds). Figure 8 shows a sample of qualitative results from the vision system. The efficacy of vision-based pose estimation is assessed on downstream tasks. Our findings indicate consistent success in effectively grasping and singulating assembly parts from vision across diverse test scenarios. The final precision required for insertion tasks was compensated by tactile and F/T sensor feedback, which is presented next.

B. Part Singulation

To evaluate the performance of part singulation, we compare the random shooting method with different numbers of samples, $N^{RS} \in \{1, 10, 100, 1000\}$. We set the maximum number of interactions to 10, and use this value if the algorithm fails to singulate, that is, the target object moves outside the manipulation surface. Table. II shows the success rate, number of interactions, and elapsed time to complete the singulation. As expected, the success rate and the number of interactions improve as the number of samples increases and converge with $N^{RS} = 100$. Since $N^{RS} = 1000$ takes approximately 9.3 times longer than $N^{RS} = 100$ with similar performance, we use $N^{RS} = 100$ for real experiments.

The model-based planner for singulation was implemented on the real system using a compliance controller designed

TABLE II: **Part singulation.** Success rate, number of interactions, and elapsed time required to singulate objects with different number of samples N^{RS} in simulation and the real system. The results are averaged over 100 runs on different initial conditions.

Number of samples N^{RS}	Simulation				Real
	1	10	100	1000	100
Success rate	0.19	0.80	1.00	1.00	1.00
Number of interactions	8.78	4.12	1.94	1.94	3.26
Elapsed time [secs]	0.05	0.21	1.38	12.87	-

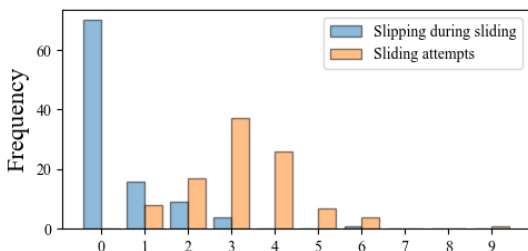


Fig. 9: **Part singulation.** The histogram shows the distribution of singulation attempts and slip frequency between the object and the manipulator over 100 real-system trials.

for sliding. We tested the execution on real systems for 100 initial conditions with a 100% success rate. The number of interactions required for singulation is presented in Fig. 9. We also observed slipping between the manipulator and objects during sliding, as shown in the histogram.

C. In-hand Pose Estimation

Next, we evaluate the performance of the in-hand pose estimation. We compare our method with a different number of training data (10, 100, 1000) to a pick-and-place policy, where the robot does not correct the grasping error and tries to insert the peg directly into the hole. To test the methods, we randomly added an error in grasp on the X -axis from a uniform distribution $dx \sim U(-10, 10)$ [mm].

The results in Table. III demonstrate that the baseline pick-and-place policy fails to insert the peg because the tolerance of the target hole is very tight (1 mm). On the contrary, our method with sufficient training data (more than 100) achieves a submillimeter precision in estimation, and enables the robot to insert the peg successfully. Data collection time for 100 data requires approximately 20 minutes.

D. Gear Insertion

Figure 10 shows the variation of the force data against the relative error between the peg and gear position during insertion. This data is then used to train a supervised gear insertion policy that can predict the expected error based on the force readings. We use LSTM [35] with two layers, each having 256 units, to capture patterns in time-series data. The supervised gear insertion predicts the direction of error between the peg and the gear in an iterative fashion. This supervised policy was tested for 100 times where it was able to achieve a success rate of 95% for gear insertion.

TABLE III: **In-hand pose estimation.** success rate and average error of baseline and our method with different number of training data (10, 100, 1000).

	Baseline	10	Ours 100	1000
Success rate	3/20	10/20	20/20	20/20
Average error [mm]	N/A	1.52	0.43	0.18

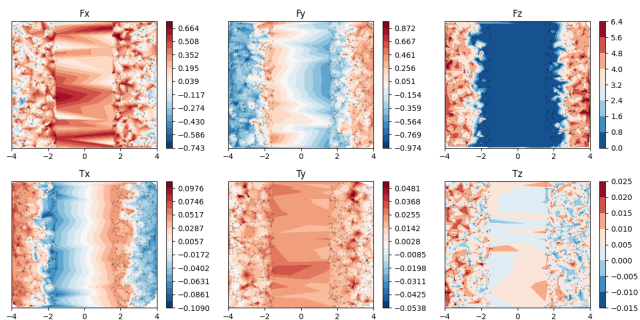


Fig. 10: **Gear insertion.** Force observations up on contact formation between the large gear and peg. We use these to learn the dependence between the alignment and the expected force readings. The horizontal axis is the error in X -axis.

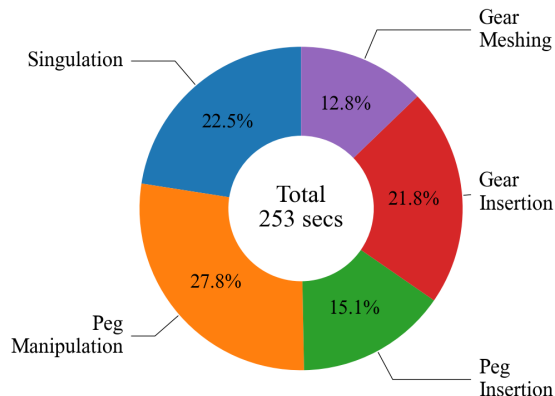


Fig. 11: **End-to-End evaluation.** Elapsed time for each operation averaged over 20 runs.

E. End-to-End Evaluation

We have tested the entire system under 225 different initial conditions (randomized by a diverse set of audience), and have been able to achieve a success rate of 99.11 % (223/225). In one of the failure cases, the robot could not singulate the parts. The other failure case was when the robot could not mesh the gear after 5 attempts. Fig. 11 shows the average elapsed time for each operation over 20 different runs. The majority of time is dedicated to manipulating the peg, demanding delicate adjustments and re-grasping (see supplementary video). Misalignment related to gear grasping necessitate intervention from the gear insertion policy trained on force signals. This tendency is evident in the duration of gear insertion. Notably, the timing results encompass the manipulator arm’s movement time.

VI. CONCLUSION AND DISCUSSION

Creating autonomous assembly agents can have a huge impact on the manufacturing industry. However, designing such agents presents myriads of challenges to achieve operational flexibility and robustness. In order to understand some of the challenges associated with autonomous assembly, we presented an autonomous robotic assembly system. The proposed assembly system can perform singulation and pose manipulation of assembly parts followed by high-precision assembly using vision as well as tactile feedback. The proposed system is extensively tested with various different initial configurations of the assembly parts to demonstrate the robustness of the proposed system.

We discovered several different failure cases during the development of the assembly system, which gave us useful insights. Since the singulation planning used a random shooting method, oftentimes it resulted in sub-optimal plans. The in-hand manipulation of the peg would fail after continued use of the Gelsight sensors. This was mainly because the gel pads of the sensors would become smooth, resulting in a reduction of the friction coefficient causing an undesirable slip of the peg during the in-hand manipulation. Similarly, the gear meshing was very susceptible to the amount of force applied during the meshing operation, requiring very fine manipulation. In the future, it would be desirable to have the capability of detecting these failures automatically and to have failure recovery modules.

Although we demonstrate the effectiveness of the proposed system with familiar components, there's a crucial need for robots to handle unfamiliar parts. Achieving this entails enabling robots to swiftly perceive and construct object models through interactive perception. This process mirrors how humans intuitively form object models during interactions. Nevertheless, comprehensively understanding the creation of object models tailored for object manipulation remains an area ripe for exploration.

REFERENCES

- [1] X. Zhu, D. K. Jha, D. Romeres, L. Sun, M. Tomizuka, and A. Cherian, "Multi-level reasoning for robotic assembly: From sequence inference to contact selection," *arXiv preprint arXiv:2312.10571*, 2023.
- [2] M. T. Mason, "Toward robotic manipulation," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, pp. 1–28, 2018.
- [3] R. A. Knepper, T. Layton, J. Romanishin, and D. Rus, "Ikeabot: An autonomous multi-robot coordinated furniture assembly system," in *Proc. of ICRA*. IEEE, 2013, pp. 855–862.
- [4] G. Thomas, M. Chien, A. Tamar, J. A. Ojea, and P. Abbeel, "Learning robotic assembly from cad," in *Proc. of ICRA*, 2018, pp. 3524–3531.
- [5] F. Suárez-Ruiz, X. Zhou, and Q.-C. Pham, "Can robots assemble an ikea chair?" *Science Robotics*, vol. 3, no. 17, p. eaat6385, 2018.
- [6] 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.
- [7] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg, "Learning ambidextrous robot grasping policies," *Science Robotics*, vol. 4, no. 26, p. eaau4984, 2019.
- [8] J. Lundell, F. Verdoja, and V. Kyrki, "Ddgc: Generative deep dexterous grasping in clutter," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 6899–6906, 2021.
- [9] A. Zeng, S. Song, S. Welker, J. Lee, A. Rodriguez, and T. Funkhouser, "Learning synergies between pushing and grasping with self-supervised deep reinforcement learning," in *Proc. of IROS*. IEEE, 2018, pp. 4238–4245.
- [10] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, "Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes," *RSS*, 2018.
- [11] J. Tremblay, T. To, B. Sundaralingam, Y. Xiang, D. Fox, and S. Birchfield, "Deep object pose estimation for semantic robotic grasping of household objects," *CoRL*, 2018.
- [12] M. Bauza, A. Bronars, and A. Rodriguez, "Tac2pose: Tactile object pose estimation from the first touch," *The International Journal of Robotics Research*, vol. 42, no. 13, pp. 1185–1209, 2023.
- [13] K. Ota, D. K. Jha, H.-Y. Tung, and J. B. Tenenbaum, "Tactile-filter: Interactive tactile perception for part mating," *RSS*, 2023.
- [14] S. Dong, D. K. Jha, D. Romeres, S. Kim, D. Nikovski, and A. Rodriguez, "Tactile-rl for insertion: Generalization to objects of unknown geometry," in *Proc. of ICRA*, 2021, pp. 6437–6443.
- [15] M. Zhang, D. K. Jha, A. U. Raghunathan, and K. Hauser, "Simultaneous trajectory optimization and contact selection for multi-modal manipulation planning," *RSS*, 2023.
- [16] Y. Shirai, D. K. Jha, A. U. Raghunathan, and D. Romeres, "Robust pivoting: Exploiting frictional stability using bilevel optimization," in *Proc. of ICRA*, 2022, pp. 992–998.
- [17] Y. Shirai, D. K. Jha, and A. U. Raghunathan, "Robust pivoting manipulation using contact implicit bilevel optimization," *arXiv preprint arXiv:2303.08965*, 2023.
- [18] N. Chavan-Dafle, R. Holladay, and A. Rodriguez, "Planar in-hand manipulation via motion cones," *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 163–182, 2020.
- [19] S. Kim, A. Bronars, P. Patre, and A. Rodriguez, "Texterity–tactile extrinsic dexterity: Simultaneous tactile estimation and control for extrinsic dexterity," *arXiv preprint arXiv:2403.00049*, 2024.
- [20] B. Liang, K. Ota, M. Tomizuka, and D. Jha, "Robust in-hand manipulation with extrinsic contacts," in *Proc. of ICRA*, 2024.
- [21] D. K. Jha, D. Romeres, S. Jain, W. Yerazunis, and D. Nikovski, "Design of adaptive compliance controllers for safe robotic assembly," in *2023 European Control Conference (ECC)*, 2023, pp. 1–8.
- [22] M. Heo, Y. Lee, D. Lee, and J. J. Lim, "Furniturebench: Reproducible real-world benchmark for long-horizon complex manipulation," in *RSS*, 2023.
- [23] A. X. Lee and et al., "Beyond pick-and-place: Tackling robotic stacking of diverse shapes," in *CoRL*, 2021.
- [24] K. Bousmalis, G. Vezzani, D. Rao, C. M. Devin, A. X. Lee, and et al., "Robocat: A self-improving generalist agent for robotic manipulation," *Transactions on Machine Learning Research*, 2024.
- [25] M. Bauza, A. Bronars, Y. Hou, I. Taylor, N. Chavan-Dafle, and A. Rodriguez, "simple: a visuotactile method learned in simulation to precisely pick, localize, regrasp, and place objects," *arXiv preprint arXiv:2307.13133*, 2023.
- [26] L. Sun, D. K. Jha, C. Hori, S. Jain, R. Corcodel, X. Zhu, M. Tomizuka, and D. Romeres, "Interactive planning using large language models for partially observable robotics tasks," *Proc. of ICRA*, 2024.
- [27] B. Wen, J. Tremblay, V. Blukis, S. Tyree, T. Müller, A. Evans, D. Fox, J. Kautz, and S. Birchfield, "Bundlesdf: Neural 6-dof tracking and 3d reconstruction of unknown objects," in *Proc. of CVPR*, 2023.
- [28] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. of CVPR*, 2016, pp. 770–778.
- [29] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *Proc. of CVPR*, 2017, pp. 2961–2969.
- [30] S. Jain and B. Argall, "Grasp detection for assistive robotic manipulation," in *Proc. of ICRA*. IEEE, 2016, pp. 2015–2021.
- [31] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *Proc. of IROS*, 2012, pp. 5026–5033.
- [32] Y. Shirai, D. K. Jha, A. U. Raghunathan, and D. Romeres, "Chance-constrained optimization for contact-rich systems using mixed integer programming," *Nonlinear Analysis: Hybrid Systems*, vol. 52, p. 101466, 2024.
- [33] X. Yi and N. Fazeli, "Precise object sliding with top contact via asymmetric dual limit surfaces," *RSS*, 2023.
- [34] K. Ota, D. K. Jha, K. M. Jatavallabhula, A. Kanazaki, and J. B. Tenenbaum, "Tactile estimation of extrinsic contact patch for stable placement," in *Proc. of ICRA*, 2024.
- [35] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.