High-Accuracy Injection Using a Mobile Manipulation Robot for Chemistry Lab Automation

Angelos Angelopoulos[1], Matthew Verber[2], Collin McKinney[2], James Cahoon[2], Ron Alterovitz[1]

***Abstract*— Lab automation has the potential to accelerate scientific progress in the natural sciences, allowing tedious experiments that would require many hours of human time to be automated, enabling higher accuracy, efficiency, and repeatability. Mobile manipulation robots have the potential to work in chemistry labs designed for humans to complete tasks for which setting up customized factory-scale automation is premature or infeasible. We present a new method to enable a mobile manipulation robot to automate injections, a common task in chemistry labs when using equipment such as gas chromatographs (GCs) for analyzing the contents of a sample mixture. This task is challenging for a mobile manipulation robot due to the need to navigate to the equipment in the lab and then achieve millimeter-scale accuracy required for the syringe positioning. Our approach leverages deep learning to create a model capable of localizing the syringe with high accuracy using cameras mounted on the chemistry equipment, and then uses a visual servoing approach based on the syringe's needle localization to achieve the injection. We demonstrate that our approach is robust to uncertainty in navigation as well as uncertainty in the grasping position and orientation of the syringe, achieving errors sufficiently small to enable the mobile manipulation robot to automate injections in real chemistry equipment.**

I. Introduction

Automation in chemistry laboratories has the potential to accelerate scientific research and discovery within the natural sciences. Automating experiments can reduce the need for tedious human labor, enable experiments to be conducted around the clock, and enhance the accuracy, efficiency, and repeatability of experimental procedures. Currently, chemists manually perform many tasks, ranging from fetching and storing reagents to operating various synthesis and measurement instruments. Current automation solutions vary from very expensive facility-wide systems to specialized tabletop assemblies, but the latter typically necessitate intricate tube networks, fixed instrument couplings, or are confined to handling materials and equipment within a limited proximity. These solutions are generally restricted to specific tasks and cannot automate the broad spectrum of lab-wide tasks customarily performed by chemists when conducting basic chemistry research. Mobile manipulation robots could offer a solution, performing many tasks traditionally executed by humans, and thus, safely increasing productivity within human-designed chemistry labs.

*This work was supported by the Creativity Hub at UNC-Chapel Hill.

[1]Angelos Angelopoulos and Ron Alterovitz are with the Department of Computer Science, University of North Carolina at Chapel Hill, NC 27599, USA {aangelos,ron}@cs.unc.edu

[2]Matthew Verber, Collin McKinney, and James Cahoon are with the Department of Chemistry, University of North Carolina at Chapel Hill, NC 27599, USA {mverber,collin,jfcahoon}@unc.edu

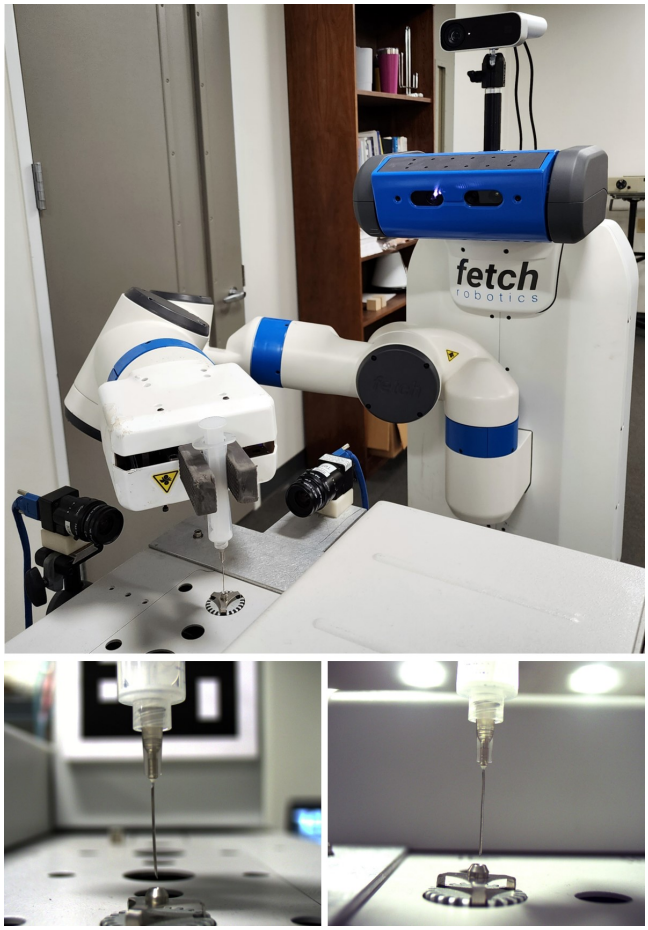

Fig. 1. A mobile manipulation robot (top) uses our method to autonomously perform an injection in a gas chromatograph (GC), a common instrument in chemistry labs for analyzing the content of a sample mixture. The task requires high accuracy, as can be seen from the 2 camera views (bottom).

Automating tasks in a chemistry lab using mobile manipulation robots is challenging. These robots must operate within human-centric environments while safely navigating and performing high-precision manipulation tasks. Chemistry labs typically have limited space that impedes navigation and cluttered tabletops that increase the difficulty of object manipulation. Additionally, the presence of fragile items, hazardous chemicals, and costly equipment demands careful robotic navigation and manipulation to prevent damage and ensure safety. After navigating several meters to a benchtop or specialized instrument, the robot's task frequently necessitates millimeter-level accuracy for successful and safe execution.

In this paper, we introduce a novel method enabling a mobile manipulation robot to automate injections — a common task in chemistry labs. For example, injection is integral for utilizing chemical analysis equipment such as a gas chromatograph (GC). A key component of chemistry research with a global market of $2.8 billion [1], a GC analyzes the chemical makeup of gas or liquid sample mixtures via separation on a stationary phase column, inducing varied residence times per molecule. For gas chromatography, an injection task requires the robot to transport a sample-loaded syringe to a GC, accurately align it with a millimeter-scale injection port (as shown in Figure 1), carefully perform the insertion, dispense the syringe's contents, and retract the syringe. Injection tasks are common in chemistry and extend beyond gas chromatography, including dispensing or withdrawing gases or liquids to/from specialized equipment and containers (e.g., vials) sealed with septums.

The automated injection problem is challenging for a mobile manipulation robot. The task requires navigating to the equipment and then achieving millimeter-level accuracy to align a syringe of submillimeter radius with a millimeter-scale injection target - in our case with a 1.5 mm radius. There are many aspects that complicate the problem, such as camera extrinsic calibration errors, joint encoder errors, and control and planning errors, which are exacerbated when using a low-cost, compact mobile manipulation robot as required by the application. Perception is a particularly significant challenge since these small objects are very difficult to detect from the robot's perspective. Furthermore, they are reflective, which makes perception even more difficult. There is very little room for error due to the small scale of the injection target; a failure due to missing the target may result in damage to the syringe's needle, leakage of the syringe's contents, damage to the environment, and delays to the experiment.

Our approach leverages a deep learning model capable of detecting the syringe's needle with high accuracy using machine vision cameras mounted on the chemistry equipment, combined with visual servoing, to achieve the injection. The robot first navigates to the chemistry equipment with the injection port and then uses an AprilTag [2] to position its end effector within the port's vicinity. The robot then aligns the needle with the injection port using visual servoing that leverages the deep learning model for detecting the needle in the equipment-mounted camera images and iteratively reduces the alignment error between the needle and the injection port. The deep learning model can detect the needle under different orientations and even when it is blurry. Once the robot aligns the needle, it inserts the needle into the injection port.

We demonstrate our method with the Fetch Mobile Manipulator Robot and two Toshiba TeliCam machine vision cameras. Our results show that our approach is robust to uncertainty in navigation as well as uncertainty in the grasping position and orientation of the syringe, achieving errors sufficiently small to enable the mobile manipulation robot to automate injections in real chemistry equipment.

II. Related Work

Whereas chemistry automation has traditionally relied on systems such as bespoke tabletop machines [3], [4], [5], [6], [7], fixed robotic arms [8], [9], [10], cartesian robots [11], and flow chemistry systems [12], [13]; a new wave of automation seeks to bring mobile, general-purpose robots to the lab [14], [15], [16], [17].

We first cover some literature on fixed robot arms. Coley et al. [18] used a robotic arm in a flow chemistry workcell to swap and reorder flow components, such as reactors and tubing, based on the desired reagent to be synthesized. Sparkes et al. [19] set up two robot arms in two separate workcells to interact with chemistry instruments. Lim et al. [20] employed an arm installed next to chemistry equipment, including a magnetic stirrer and a GC, to automate organic chemistry reactions. Notably, this study used a commercial auto-injector for performing GC injections. However, auto-injectors are fixed to the machine and do not work if the syringe requires its contents to be retrieved from elsewhere, which our mobile automated injection solution allows. Fixed arms can be useful for bespoke tabletop setups, but using multiple arms can be expensive and reduce flexibility and scalability.

Mobile robots address many of the problems of fixed robot arms by offering general, reusable platforms that can be used to automate tasks across the chemistry lab. Fakhruldeen et al. [16] used a KUKA KMR mobile robot to transport samples as well as load/unload chemistry equipment. A Franka Emika Panda arm was also used to automate a tabletop workcell consisting of a hot plate, a balance, and a stirrer. Kleine-Wechelmann et al. [15] created a mobile robot capable of handling sample trays. The robot could transport trays to and from storage, as well as load/unload them from chemistry instruments. The robot had integrated storage for several trays on its frame, enabling quicker transportation.

A notable project is by Burger et al. [14], who used a mobile robot and chemistry instruments to automatically execute 688 experiments for finding more active photocatalysts. The robot transported samples and loaded/unloaded them to/from instruments. The system used Bayesian optimization and identified photocatalysts six times more active than the base mixtures. The robot and various instruments were connected using the Robot Operating System (ROS) [21], which was used by a process management system for coordinating the experiments. The robot was part of a larger automation infrastructure, with the entire system taking about 2 years to develop.

While mobile robots have successfully been used to interact with chemistry equipment by loading/unloading samples, they have not been used for more complex tasks. This paper is a step towards automating more difficult lab tasks with a mobile manipulation robot, starting with automated injection, and shows that mobile manipulation robots can be used to automate chemistry procedures that would otherwise require human intervention.

III. Method

We tackle the problem of automated injection with a mobile manipulation robot. We begin by formalizing the problem and then describe our system components and method for automating this task.

A. Problem Definition

Mobile manipulation robots have the potential to automate tasks in existing chemistry labs because they can navigate around the lab, transport samples, interact with chemistry instruments, and perform a variety of manipulation tasks that would otherwise require human intervention. A single mobile manipulation robot can be used for many different tasks at many stations within the lab and can transport items from one station to another. Here we address the problem of enabling such a robot to autonomously perform high-accuracy injections into a target, a common task required when using a variety of chemistry instruments, such as GCs.

We use a mobile manipulation robot with a differential drive base and a 7-DoF arm mounted on a liftable torso. The differential drive allows the robot to rotate in place and more easily navigate in a cluttered environment. The DoF of the arm and torso are sufficient to enable obstacle avoidance while performing tasks using tabletop equipment. We define the base configuration as $\mathbf{q_b} = (x, y, \theta)$, where x and y are the coordinates of the base and θ is the base's heading. The arm and torso can be combined into one configuration as $\mathbf{q_a} = (\theta_1 ... \theta_7, \tau)$, where $\theta_1, ..., \theta_7$ are the joint angles of the arm and τ is the vertical position of the liftable torso.

The injection task involves two objects: an injection target, whose pose along its center we denote by $\mathbf{t} \in SE(3)$, and a syringe, whose pose along its center we denote by $\mathbf{s} \in SE(3)$. The target is typically composed of a circular port inside of which is a rubber septum which the syringe's needle can pierce through. The septum automatically re-seals after the needle is withdrawn. Some chemistry equipment include a conical structure on the injection port to help guide the needle into the septum, but this conical structure is typically limited in size to avoid deflecting the needle too much such that it sustains damage. We define t_r as the radius around $\mathbf{t}$ that defines a circular injection surface perpendicular to the Z axis, including the septum's surface area and the conical guide when present. We denote (t_x, t_y) as the center point of the injection target $\mathbf{t}$. The injection target lies flat, so its orientation components (roll, pitch, yaw) are zero. In our experiments, we use a Varian 450 GC that has an injection port with $t_r = 1.50$ mm.

The syringe has a needle of cylindrical shape, with a radius of s_r. In our experiments, $s_r = 0.35$ mm. We annotate two keypoints on the needle, $\mathbf{k_1}$, the base of the needle, and $\mathbf{k_2}$, the tip of the needle. The syringe, depending on its properties, can bend to a certain degree, allowing the target's cone to function as a guide such that any point on the target within the radius t_r is valid for the needle injection. We emphasize that the tip of the needle is fragile and can easily be damaged if it contacts a hard surface, so there are limits to the size and slope of the guide. Both the injection target and the needle are usually metallic and reflective.

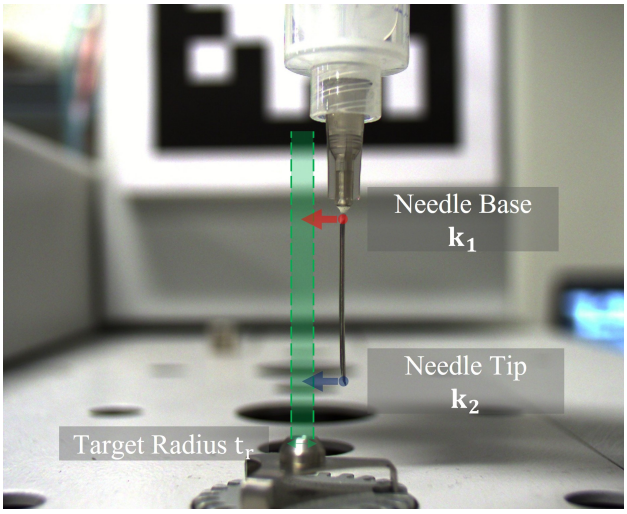

Fig. 2. Illustration of the automated injection problem, which requires aligning the robot-grasped syringe (which includes the syringe's needle) and the target with high accuracy.

The automated injection problem can be formally defined as the task of aligning the syringe $\mathbf{s}$ grasped by the robot's end effector with the injection target $\mathbf{t}$. Alignment is successful when $\mathbf{s}$ is in the circle (t_x, t_y) with radius t_r. Injection not only requires alignment of x and y but also that s_z is some ε above t_z so that when the syringe is lowered it is inserted into the target. Aside from translational alignment, orientational alignment is also necessary. The Z axes of $\mathbf{s}$ and $\mathbf{t}$ must ideally be anti-parallel ($\uparrow\downarrow$) to each other, where $\overrightarrow{s_z}$ is the axis pointing to the direction of the needle and $\overrightarrow{t_z}$ is the axis pointing upwards from the target. This ensures that the two are rotated in a way that would allow insertion of the needle into the target by moving down along the Z axis. Initially, we assume random configurations $\mathbf{q_b}$ and $\mathbf{q_a}$, random $\mathbf{s}$, and a fixed $\mathbf{t}$.

Accurate alignment of the needle is required before making contact with the injection port. Figure 2 illustrates the problem of aligning the needle with the target and equation 1 summarizes the injection problem, where the alignment function f controls the robot by computing and executing a set of motions in a closed loop, until the set of constraints to enable accurate alignment and successful injection are met. We note that due to the flexibility and fragility of the needle, as well as the force required to pierce the rubber septum inside the injection port, approaches that rely on force feedback will not work and can be dangerous. Another challenge is that the pose of $\mathbf{s}$ and $\mathbf{t}$ is unknown and must be estimated.

$$\mathbf{s}' = f(\mathbf{s}, \mathbf{t}, \mathbf{q_b}, \mathbf{q_a}) \,|\, (s_x - t_x)^2 + (s_y - t_y)^2 \leq t_r^2 \\ \wedge s_z = t_z - \varepsilon \wedge \overrightarrow{s_z} \uparrow\downarrow \overrightarrow{t_z} \tag{1}$$

B. Method Overview

Our method uses two machine vision cameras installed on the XZ and YZ planes of **t**. We trained a deep learning model to detect the base and tip keypoints of the needle, and we use these detections and a visual servoing controller to accurately move **s** above **t**, satisfying the constraints in equation 1. We also use an AprilTag to first position **s** near **t** such that the needle is visible by the cameras. The cameras are used to detect the needle on two planes using the deep learning model and the displacements are used to align the needle to the target in both planes simultaneously. This alignment is performed in a loop of needle detection and visual servoing until alignment is achieved. We show the workspace of the task in Figure 3.

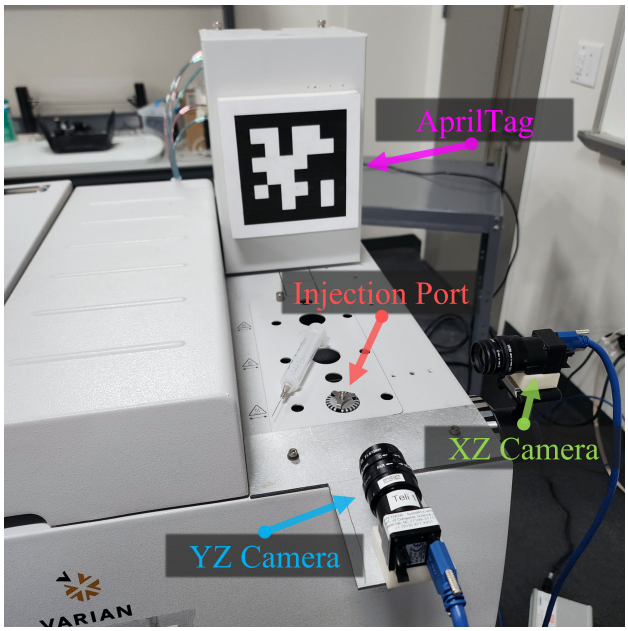

Fig. 3. Setup of the workspace for the injection task.

The robot we used was a Fetch Mobile Manipulator. It features a liftable torso, a rotating head with a built-in RGBD sensor, and a 7-DoF robotic arm mounted on top a differential drive base. Due to the low RGB image quality of the built-in RGBD sensor, we mounted an additional Azure Kinect RGBD sensor with up to 4K RGB resolution on top of the head. For the machine vision cameras, we used two Toshiba BU505MCF USB3 TeliCams with Kowa LM12JCM 2/3" 12mm/F1.4 lenses.

C. High-Accuracy Needle Detection

The dimensions of the task objects makes egocentric perception very difficult. They are also metallic and reflective, material properties that are widely known to complicate computer vision. Furthermore, common perception errors such as intrinsic and extrinsic calibration errors can have significant negative impact due to the millimeter scale on which the problem must be solved. Consequently, perception proved to be especially challenging for this problem, and robust perception was key to successful injection.

Accurately detecting the needle **s** and the target **t** is crucial to successful task completion. We need the position and orientation of **s** on the XZ and YZ planes. Unfortunately, due to the small scale of **s**, an RGBD sensor did not suffice. Consequently, we mounted two Toshiba TeliCam machine vision cameras facing **t** on two planes, providing both horizontal and vertical views. Since the cameras are fixed on the machine, we calibrated the location of **t** using its pixel position in the camera frames. The positioning of these cameras does not require strict precision (e.g., they do not need to perfectly face the XZ and YZ planes of the target) as the alignment method uses multiple iterations and can negate small positioning errors over time at the cost of longer alignment convergence. We note that augmentation of the environment (e.g., adding sensors like cameras) is acceptable in a chemistry lab, as it is a controlled environment.

To detect the needle using the two cameras, we trained a Keypoint-RCNN model (based on Mask-RCNN [22]) with a Feature Pyramid Network and a ResNet50 backbone pre-trained on MS-COCO. We collected and labeled a dataset of 1432 images of the needle and the target using images from the two cameras as well as a third top-down camera that was not used in the method. In each image we labeled the base and tip keypoints of the needle. The model also needed bounding boxes, which we did not label manually but computed using the keypoints. We used an augmentation pipeline with random horizontal and vertical flipping, random brightness and contrast, and random hue, saturation, and value. The input image resolution was 512x384 pixels. We trained the model for 20 epochs using the AdamW optimizer, step scheduling for the learning rate, and warmup for 1 epoch, using an 80/20 random data split.

D. High-Accuracy Needle Alignment

Accurate alignment of the needle, as shown in Figure 4, required controlling the various sources of uncertainty that make planning and control challenging in this scenario, particularly due to the requirement of using a mobile robot. Uncertainties include joint encoder errors, controller errors, and jerk and instability of the mobile robot's frame. For example, even small arm velocities can make the mobile robot's frame wobble, which can cause the needle to be thrown out of alignment as it is being lowered to the target.

We used *MoveIt* [23] for trajectory planning and control of the robot to approach the injection target and prepare for injection. For alignment, we interfaced with the Fetch robot's arm IK velocity controller and implemented a set of servoing routines to move the end effector to a desired pose - in this case the pose that aligns the end effector with the target.

Using the arm servoing routines, we implemented pose-based and image-based visual servoing using the AprilTag and the needle keypoints detected from the two cameras. The servoing routine can align the end effector's pose $\mathbf{e} \in SE(3)$ with the desired needle alignment pose until a distance threshold ε is met. For translational alignment, we computed the linear displacement d_l of $\mathbf{e}$ from the target pose in every translational axis. The linear velocity is $v_l = d_l / p_l$,

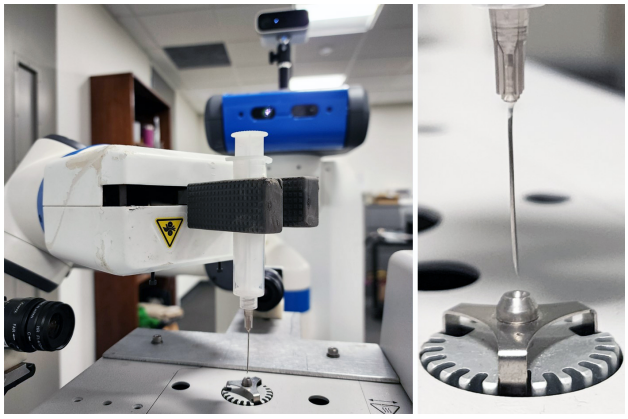

Fig. 4. Depiction of the robot aligning the syringe's needle using our method.

where p_l is a damping factor to slow down the motion. The angular displacement d_a is the minimum signed Euler angle difference for each rotational axis and so the angular velocity is $v_a = d_a/p_a$, where p_a is another damping factor. The velocities are recomputed in a loop until the Euclidean distance of $\mathbf{e}$ from the target pose for both translation and orientation meet distance thresholds ε_1 and ε_2 for each set of axes, respectively.

We use the visual servoing to align $\mathbf{s}$ with $\mathbf{t}$. We first detect the needle keypoint sets $\mathbf{k_h} = (k_{h,1}, k_{h,2})$ and $\mathbf{k_v} = (k_{v,1}, k_{v,2})$ from the two camera frames $\mathbf{F_h}$ and $\mathbf{F_v}$ using the deep learning model described in Sec. III-C. The keypoint sets $\mathbf{k_h}$ and $\mathbf{k_v}$ are both composed of two (x, y) image points $\mathbf{k_1}$ and $\mathbf{k_2}$, where $\mathbf{k_1}$ is the keypoint for the base of the needle and $\mathbf{k_2}$ is the keypoint for the tip of the needle. The goal is to move $\mathbf{s}$ so that it is inside the radius t_r. In order to determine if $\mathbf{s}$ is inside t_r, we precompute the pixel bounds $\mathbf{b_h}$ and $\mathbf{b_v}$ of $\mathbf{t}$ for $\mathbf{F_h}$ and $\mathbf{F_v}$, and then check if $\mathbf{k_2}$ for both $\mathbf{k_h}$ and $\mathbf{k_v}$ is in the bounds of $\mathbf{t}$ as seen in equation 2.

$$(b_{h,min} \leq k_{h,2} \leq b_{h,max}) \wedge (b_{v,min} \leq k_{v,2} \leq b_{v,max}) \quad (2)$$

To determine how much we need to move $\mathbf{s}$, we compute the pixel difference of $\mathbf{k_2}$ from the center of $\mathbf{b_h}$ and $\mathbf{b_v}$ and convert it to millimeters using the formula $d_{mm} = d_{px}/F_{ppmm}$, where d_{mm} is the distance in millimeters, d_{px} is the distance in pixels, and F_{ppmm} is the pixels-per-millimeter approximated by measuring the dimensions of $\mathbf{t}$ in both $\mathbf{F_h}$ and $\mathbf{F_v}$. The distance d_{mm} is computed for each camera frame, and we need to move the end effector so that this distance is minimized. To achieve this, we compute $\mathbf{k_h}$ and $\mathbf{k_v}$ and move the end effector using the servoing routines explained in the previous paragraph so that d_{mm} is minimized for both planes simultaneously. We do this in a loop until $\mathbf{k_2}$ for each $\mathbf{k_h}$ and $\mathbf{k_v}$ is within $\mathbf{b_h}$ and $\mathbf{b_v}$, respectively. When this is achieved, $\mathbf{s}$ is aligned with $\mathbf{t}$, and $\mathbf{s}$ can be safely inserted in $\mathbf{t}$ by lowering the end effector along the Z axis.

E. Implementation Details

The solution is implemented with ROS and we used three custom nodes: (1) an *injection controller* node, (2) a *needle detection* node, and (3) a *TeliCam streaming* node. The injection controller node is responsible for controlling the robot throughout the task. The needle detection node provides a service that returns the needle keypoints for a pair of cameras. Finally, the TeliCam streaming node interfaces with the Toshiba TeliCams and provides images. We also used two third-party nodes, including a node that interfaces with the Azure Kinect and a node for AprilTag detection [2].

In order to interface with the Toshiba TeliCams, we implemented a high-level C++ driver for USB3 Toshiba TeliCams, which we have open-sourced [24], and used it to create a ROS node. The driver and node make it easy to interface with multiple cameras at once and define parameters such as exposure time, gamma, white balancing, and decimation for every camera. The software allows the cameras to be used in continuous streaming mode or snapshot mode to capture images on demand.

IV. Evaluation

To evaluate our method, we measured the success rate, the distance of the syringe's needle tip from the target after alignment, how well the method performs under uncertainty, and the execution time. Performance under uncertainty was especially important, because uncertainty in grasping and navigation can make alignment of the needle more difficult. The sources of uncertainty we look at include the grasped syringe position and orientation as well as navigation uncertainty. An effective injection automation method needs to minimize variance in the distance of the needle tip from the target center, i.e., it must have high repeatability.

A. Comparison Methods

We compare our approach with a simpler method, which we label "AT", that only used an AprilTag for aligning the syringe needle. The method involves attaching an AprilTag close to the target $\mathbf{t}$ and defining a fixed transformation offset for where $\mathbf{t}$ is. We used the visual servoing routine described in Sec. III-D to align the needle with the target. The servoing target pose was the injection target $\mathbf{t}$, based on the AprilTag pose. We tried this method with multiple resolutions of the head-mounted Azure Kinect camera, including 1080p and 2160p, and noticed improved accuracy with higher resolution. We identify the automated injection methods using only the AprilTag with the 1080p and 2160p camera resolutions as AT 1080p and AT 2160p, respectively.

We compare three methods: our method, AT 2160p, and AT 1080p. While our method also utilizes AprilTag detection for the first stage of the algorithm, we use the two machine vision cameras and deep learning for needle alignment. For consistency, we used 2160p resolution for the head-mounted camera in our method, but the resolution has negligible effect on the method's performance as it primarily relies on the two cameras mounted on the instrument.

B. Experimental Setup

We conducted a total of 63 trials, 21 for each method. In each trial, we introduced uncertainty by adjusting the

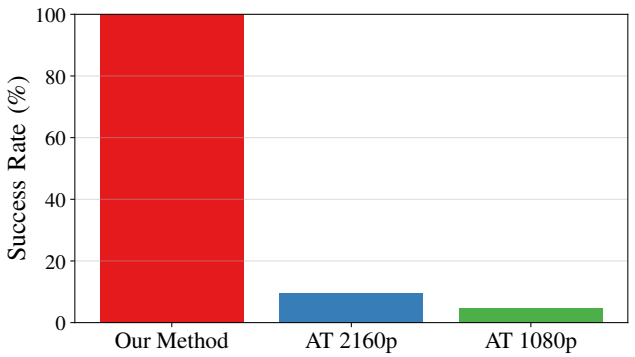

Fig. 5. Success rate of each method. Our method achieved a 100% success rate while the other two methods achieved less than a 10% success rate.

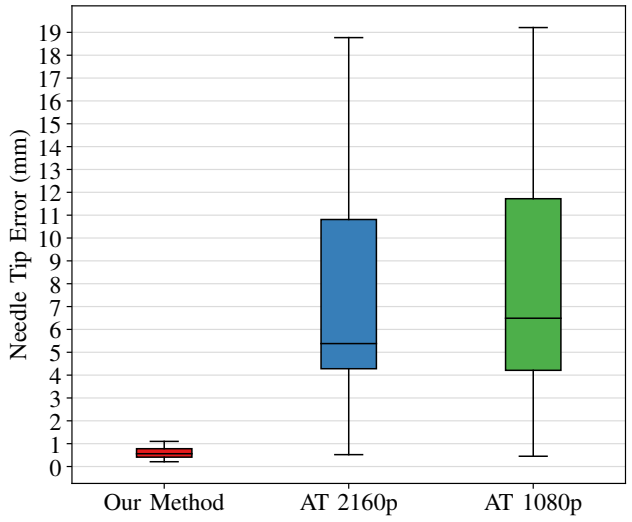

Fig. 6. Accuracy of each method based on a total of 63 trials, where lower needle tip error indicates higher accuracy.

yaw of the mobile base, its distance from the target, and the position and orientation of the syringe in the gripper. The base yaw was varied in the range of +/-15 degrees, where 0 degrees indicates that the base is directly facing the target. We measured the distance of the base from the target based on the distance of the camera, which was between 87.8 and 95 cm, constrained by the robot's arm reach and the GC. Since the syringe may not be grasped identically each time, we varied its position and orientation in the gripper. The syringe's position was uniformly sampled within a range of +/-5 mm from the center of the gripper, while the syringe's orientation was uniformly sampled within +/-9 degrees from perfectly vertical (90 degrees). For the first five trials, the syringe's position and orientation were fixed (i.e., it was kept straight and centered in the gripper) to establish a comparative baseline. We assessed the accuracy of each method by measuring the needle tip error (i.e., the XY (planar) distance of the needle's tip from the target after the method finished executing), with lower distance indicating higher accuracy. A trial was deemed successful if the needle tip error was ≤ 1.50mm.

C. Experimental Results

Our method achieved a 100% success rate, as shown in Figure 5, whereas the AT 2160p and AT 1080p methods achieved 9.52% and 4.76% success rates, respectively. Based on average accuracy, shown in Figure 6, our method was 11.44x more accurate than AT 2160p and 12.35x more accurate than AT 1080p. In its best trial, our method was 41.24x more accurate than AT 2160p and 46.52x more accurate than AT 1080p. The standard deviation of our method was more than 19 times less than that of the AT 2160p method and nearly 20 times less than that of the AT 1080p method. The end result was that our method achieved much lower distance and spread of the needle from the target.

Our method is robust to various sources of uncertainty, such as differences in grasped syringe position, grasped syringe orientation, and mobile base yaw. As shown in Figure 7, our method achieves high accuracy for both low and high uncertainty, whereas the AT methods degrade under higher uncertainty for syringe grasping orientation and position as well as mobile base yaw. We define low and high uncertainty based on the deviation for each source in regards to its range. Specifically, we define low uncertainty as indicating the deviation from the desired value being $\leq 30\%$ of the uncertainty range specified in Sec. IV-B (e.g., grasping position and orientation is close to centered and vertical, mobile base yaw is close to facing the target), whereas high uncertainty indicates deviation $> 30\%$. A deviation in base yaw changes the rotation of the head-mounted camera, causing the AprilTag to deviate from the image center, which increases the pose estimation error [25]. As for the mobile base distance, there is no known ideal value from which to measure deviations due to uncertainty, and varying distance had no significant effect on our method since all distances allowed the robot's arm to reach the target.

We achieve high success rate in the presence of uncertainty by accurately perceiving the position and orientation of the syringe's needle, which enables successful alignment regardless of how the robot grasps the syringe. Our method also utilizes fixed cameras near the injection target to perceive the needle and the injection target. This allows our method to be extended to different types of targets and needles, where only the pixel boundary of the target needs to be adjusted, while the needle's keypoints remain in the same relative locations. In contrast, the AT methods lack perception of the needle and precise knowledge of the target, leading to uncertainties in needle positioning and orientation that cannot be handled. Furthermore, errors in the AprilTag pose estimation have a minimal impact on our method, as it is only used to bring the needle in the vicinity of the target.

The average task completion time for our method was 57 seconds, whereas for the AT 2160p and AT 1080p methods the average time was 33 and 31 seconds. On average, our method was 75% slower than the AT 2160p method and 83% slower than the AT 1080p method. We note, however, that the runtime cost our method is not the primary cause of the

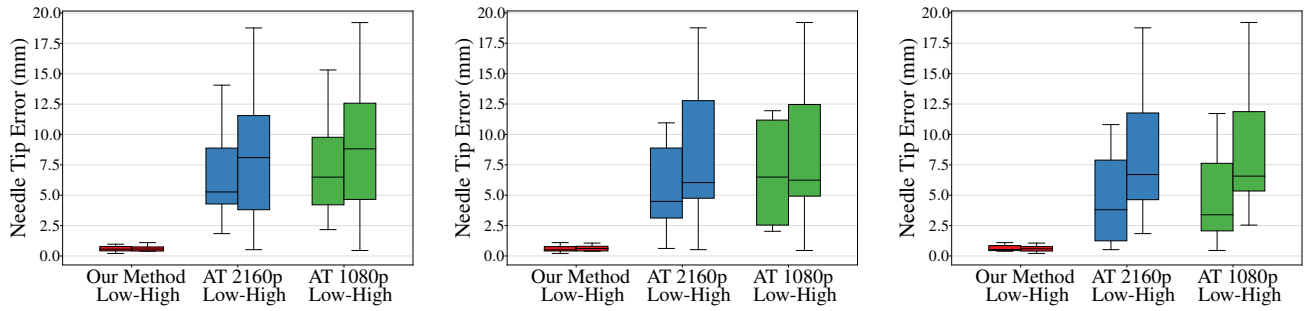

(a) Syringe Grasping Orientation (b) Syringe Grasping Position (c) Mobile Base Yaw

Fig. 7. Accuracy of each method grouped by source of uncertainty. Our method achieves lower needle tip error compared to the AT methods, and this effect is magnified when uncertainty increases in syringe grasping orientation and position and mobile base yaw.

slowdown, but rather low velocity during the alignment phase in order to keep the robot stable. The speed of the method can be significantly increased by using a higher-quality, more stable robot.

D. Needle Detection

The needle detection model demonstrated high accuracy with a validation L1 loss of 0.75 for the base keypoint $\mathbf{k_1}$ and 1.64 for the tip keypoint $\mathbf{k_2}$, based on 286 test images. Considering that the input image dimensions were 512x384 and that the needle is only aligned horizontally in the image, this translated to an error of at most 0.15% of the image width for $\mathbf{k_1}$ and 0.32% for $\mathbf{k_2}$. The loss plots for $\mathbf{k_1}$ and $\mathbf{k_2}$ are depicted in Figure 8.

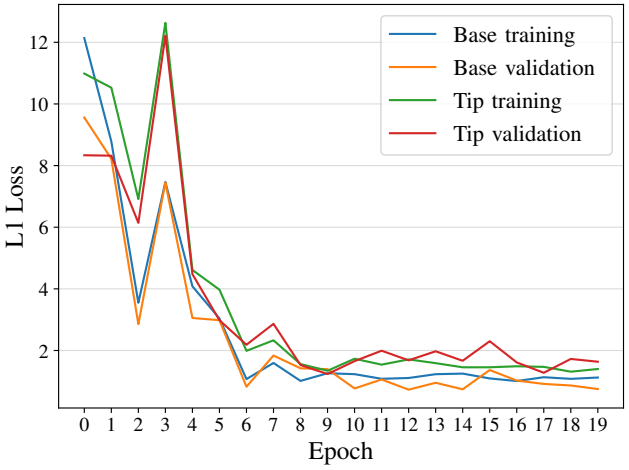

Fig. 8. Training and validation loss curves for the base and tip keypoints.

In general, higher accuracy was achieved with the base keypoint than the tip keypoint. We hypothesize that the higher accuracy for $\mathbf{k_1}$ is likely because of its clearer structure due to the prominent syringe features above it. The model is also robust to blur when the needle is out of focus. It was critical for the model to be robust to blur, since the cameras have a close focal distance. Overall, the model can detect the needle despite challenges such as

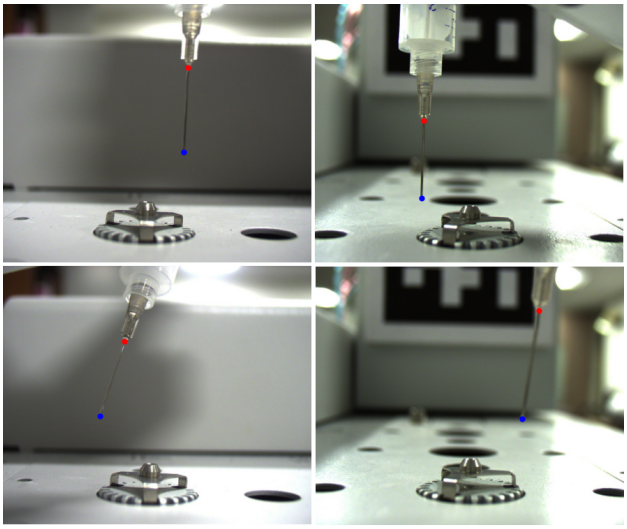

Fig. 9. Sample images with the needle keypoints detected by the model. The red keypoint is $\mathbf{k_1}$ and the blue keypoint is $\mathbf{k_2}$. The images in the left column are from the camera on the YZ plane, and the images on the right column are from the camera on the XZ plane.

orientation variations and blurriness. Figure 9 shows some sample images and the detected keypoints.

V. Conclusions

We presented a method for automating high-accuracy injections using a mobile manipulation robot for chemistry lab automation. Our method allows this critical step for chemical analysis to be automated with a mobile manipulation robot, unlocking greater automation scale. Our method can align a syringe to a target (e.g., a GC injection port) with submillimeter accuracy, while being robust to grasping errors and mobile base navigation errors. This is achieved by detecting the syringe's needle using a deep learning model from two machine vision cameras and performing image-based visual servoing to move the end effector such that the needle is aligned with the target. In addition, we use an AprilTag for initial positioning of the end effector. We compared our method to simpler methods that only use an AprilTag from which the pose of the target is calibrated. We showed that our method achieved a 100% success rate and that it was

on average over 11 times more accurate, while being robust to sources of uncertainty such as syringe grasping error and mobile base navigation error. In future work, we plan to continue development of the mobile manipulation robot by automating additional tasks critical for the acceleration of experiments in chemistry labs.

REFERENCES

[1] Global Industry Analysts Inc, "A $3.4 Billion Global Opportunity for Gas Chromatography Systems by 2026 - New Research from StrategyR." [Online]. Available: https://www.prnewswire.com/news-releases/a-3-4-billion-global-opportunity-for-gas-chromatography-systems-by-2026---new-research-from-strategyr-301505733.html

[2] J. Wang and E. Olson, "AprilTag 2: Efficient and robust fiducial detection," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2016, pp. 4193–4198, iSSN: 2153-0866.

[3] S. Asche, G. J. T. Cooper, G. Keenan, C. Mathis, and L. Cronin, "A robotic prebiotic chemist probes long term reactions of complexifying mixtures," *Nature Communications*, vol. 12, no. 1, p. 3547, June 2021, number: 1 Publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/s41467-021-23828-z

[4] I. Schmid and J. Aschoff, "A scalable software framework for data integration in bioprocess development," *Engineering in Life Sciences*, vol. 17, no. 11, pp. 1159–1165, 2017. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/elsc.201600008

[5] V. Dragone, V. Sans, A. B. Henson, J. M. Granda, and L. Cronin, "An autonomous organic reaction search engine for chemical reactivity," *Nature Communications*, vol. 8, no. 1, p. 15733, June 2017, number: 1 Publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/ncomms15733

[6] M. Christensen, L. P. E. Yunker, F. Adedeji, F. Häse, L. M. Roch, T. Gensch, G. dos Passos Gomes, T. Zepel, M. S. Sigman, A. Aspuru-Guzik, and J. E. Hein, "Data-science driven autonomous process optimization," *Communications Chemistry*, vol. 4, no. 1, pp. 1–12, Aug. 2021, number: 1 Publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/s42004-021-00550-x

[7] N. Hartrampf, A. Saebi, M. Poskus, Z. P. Gates, A. J. Callahan, A. E. Cowfer, S. Hanna, S. Antilla, C. K. Schissel, A. J. Quartararo, X. Ye, A. J. Mijalis, M. D. Simon, A. Loas, S. Liu, C. Jessen, T. E. Nielsen, and B. L. Pentelute, "Synthesis of proteins by automated flow chemistry," *Science*, vol. 368, no. 6494, pp. 980–987, May 2020, publisher: American Association for the Advancement of Science. [Online]. Available: https://www.science.org/doi/10.1126/science.abb2491

[8] J. Li, S. G. Ballmer, E. P. Gillis, S. Fujii, M. J. Schmidt, A. M. E. Palazzolo, J. W. Lehmann, G. F. Morehouse, and M. D. Burke, "Synthesis of many different types of organic small molecules using one automated process," *Science*, vol. 347, no. 6227, pp. 1221–1226, Mar. 2015, publisher: American Association for the Advancement of Science. [Online]. Available: https://www.science.org/doi/10.1126/science.aaa5414

[9] J. Grizou, L. J. Points, A. Sharma, and L. Cronin, "A curious formulation robot enables the discovery of a novel protocell behavior," *Science Advances*, vol. 6, no. 5, p. eaay4237, Jan. 2020, publisher: American Association for the Advancement of Science. [Online]. Available: https://www.science.org/doi/10.1126/sciadv.aay4237

[10] P. Courtney, M. S. Beck, and W. J. Martin, "A vision guided life-science laboratory robot," *Measurement Science and Technology*, vol. 2, no. 2, pp. 97–101, Feb. 1991, publisher: IOP Publishing. [Online]. Available: https://doi.org/10.1088/0957-0233/2/2/001

[11] F. Kong, L. Yuan, Y. F. Zheng, and W. Chen, "Automatic Liquid Handling for Life Science: A Critical Review of the Current State of the Art," *Journal of Laboratory Automation*, vol. 17, no. 3, pp. 169–185, June 2012, publisher: SAGE Publications Inc. [Online]. Available: https://doi.org/10.1177/2211068211435302

[12] K. F. Jensen, "Flow chemistry—Microreaction technology comes of age," *AIChE Journal*, vol. 63, no. 3, pp. 858–869, 2017. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/aic.15642

[13] J. Wegner, S. Ceylan, and A. Kirschning, "Flow Chemistry – A Key Enabling Technology for (Multistep) Organic Synthesis," *Advanced Synthesis & Catalysis*, vol. 354, no. 1, pp. 17–57, 2012. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/adsc.201100584

[14] B. Burger, P. M. Maffettone, V. V. Gusev, C. M. Aitchison, Y. Bai, X. Wang, X. Li, B. M. Alston, B. Li, R. Clowes, N. Rankin, B. Harris, R. S. Sprick, and A. I. Cooper, "A mobile robotic chemist," *Nature*, vol. 583, no. 7815, pp. 237–241, July 2020, number: 7815 Publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/s41586-020-2442-2

[15] S. Kleine-Wechelmann, K. Bastiaanse, M. Freundel, and C. Becker-Asano, "Designing the mobile robot Kevin for a life science laboratory," in *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, Aug. 2022, pp. 870–875, iSSN: 1944-9437.

[16] H. Fakhruldeen, G. Pizzuto, J. Glowacki, and A. I. Cooper, "ARChemist: Autonomous Robotic Chemistry System Architecture," Apr. 2022, arXiv:2204.13571 [cs]. [Online]. Available: http://arxiv.org/abs/2204.13571

[17] Q. Zhu, F. Zhang, Y. Huang, H. Xiao, L. Zhao, X. Zhang, T. Song, X. Tang, X. Li, G. He, B. Chong, J. Zhou, Y. Zhang, B. Zhang, J. Cao, M. Luo, S. Wang, G. Ye, W. Zhang, X. Chen, S. Cong, D. Zhou, H. Li, J. Li, G. Zou, W. Shang, J. Jiang, and Y. Luo, "An all-round AI-Chemist with scientific mind," *National Science Review*, p. nwac190, Sept. 2022. [Online]. Available: https://doi.org/10.1093/nsr/nwac190

[18] C. W. Coley, D. A. Thomas, J. A. M. Lummiss, J. N. Jaworski, C. P. Breen, V. Schultz, T. Hart, J. S. Fishman, L. Rogers, H. Gao, R. W. Hicklin, P. P. Plehiers, J. Byington, J. S. Piotti, W. H. Green, A. J. Hart, T. F. Jamison, and K. F. Jensen, "A robotic platform for flow synthesis of organic compounds informed by AI planning," *Science*, vol. 365, no. 6453, p. eaax1566, Aug. 2019, publisher: American Association for the Advancement of Science. [Online]. Available: https://www.science.org/doi/10.1126/science.aax1566

[19] A. Sparkes, W. Aubrey, E. Byrne, A. Clare, M. N. Khan, M. Liakata, M. Markham, J. Rowland, L. N. Soldatova, K. E. Whelan, M. Young, and R. D. King, "Towards Robot Scientists for autonomous scientific discovery," *Automated Experimentation*, vol. 2, no. 1, p. 1, Jan. 2010. [Online]. Available: https://doi.org/10.1186/1759-4499-2-1

[20] J. X.-Y. Lim, D. Leow, Q.-C. Pham, and C.-H. Tan, "Development of a Robotic System for Automatic Organic Chemistry Synthesis," *IEEE Transactions on Automation Science and Engineering*, vol. 18, no. 4, pp. 2185–2190, Oct. 2021, conference Name: IEEE Transactions on Automation Science and Engineering.

[21] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, A. Y. Ng, and others, "ROS: an open-source Robot Operating System," in *ICRA workshop on open source software*, vol. 3. Kobe, Japan, 2009, p. 5, issue: 3.2.

[22] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," Jan. 2018, arXiv:1703.06870 [cs]. [Online]. Available: http://arxiv.org/abs/1703.06870

[23] D. Coleman, I. Sucan, S. Chitta, and N. Correll, "Reducing the Barrier to Entry of Complex Robotic Software: a MoveIt! Case Study," Apr. 2014, arXiv:1404.3785 [cs]. [Online]. Available: http://arxiv.org/abs/1404.3785

[24] A. Angelopoulos, "UNC-Robotics/telicam_driver: High-level driver for Toshiba USB3 TeliCams." [Online]. Available: https://github.com/UNC-Robotics/telicam_driver

[25] S. M. Abbas, S. Aslam, K. Berns, and A. Muhammad, "Analysis and Improvements in AprilTag Based State Estimation," *Sensors*, vol. 19, no. 24, p. 5480, Jan. 2019, number: 24 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: https://www.mdpi.com/1424-8220/19/24/5480