This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Suomalainen, Markku; Calinon, Sylvain; Pignat, Emmanuel; Kyrki, Ville
Improving dual-arm assembly by master-slave compliance

Published in:
Proceedings of the International Conference on Robotics and Automation, ICRA 2019

DOI:
[10.1109/ICRA.2019.8793977](https://doi.org/10.1109/ICRA.2019.8793977)

Published: 01/05/2019

Document Version
Peer reviewed version

Please cite the original version:
Suomalainen, M., Calinon, S., Pignat, E., & Kyrki, V. (2019). Improving dual-arm assembly by master-slave compliance. In *Proceedings of the International Conference on Robotics and Automation, ICRA 2019* (pp. 8676-8682). [8793977] (IEEE International Conference on Robotics and Automation). IEEE.
<https://doi.org/10.1109/ICRA.2019.8793977>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

© 2019 IEEE. This is the author's version of an article that has been published by IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Improving dual-arm assembly by master-slave compliance

Markku Suomalainen^{1,2}, Sylvain Calinon³, Emmanuel Pignat³ and Ville Kyrki²

Abstract—In this paper we show how different choices regarding compliance affect a dual-arm assembly task. In addition, we present how the compliance parameters can be learned from a human demonstration. Compliant motions can be used in assembly tasks to mitigate pose errors originating from, for example, inaccurate grasping. We present analytical background and accompanying experimental results on how to choose the center of compliance to enhance the convergence region of an alignment task. Then we present the possible ways of choosing the compliant axes for accomplishing alignment in a scenario where orientation error is present. We show that an earlier presented Learning from Demonstration method can be used to learn motion and compliance parameters of an impedance controller for both manipulators. The learning requires a human demonstration with a single teleoperated manipulator only, easing the execution of demonstration and enabling usage of manipulators at difficult locations as well. Finally, we experimentally verify our claim that having both manipulators compliant in both rotation and translation can accomplish the alignment task with less total joint motions and in shorter time than moving one manipulator only. In addition, we show that the learning method produces the parameters that achieve the best results in our experiments.

I. INTRODUCTION

When a human performs alignment motions as depicted in Fig. 1, often only one hand moves the piece actively and the other hand only complies with the motions to ease the alignment. Whenever there is not a fixture available to hold one of the workpieces, they must both be held to perform alignment. Therefore, to allow low-threshold automation of currently human-performed tasks, robots must be able to execute these tasks similarly to humans and learn the tasks efficiently [1].

When a dual-arm robot needs to grasp and assemble two workpieces, their relative position and orientation will differ in each consecutive trials. This issue is independent of the controller being centralized or decentralized because location of a grasp on a workpiece is always limited in precision. Pose errors during alignment caused by this uncertainty are difficult to mitigate with a vision system due to contact between the workpieces and occlusions. Another approach to mitigate these errors is leveraging compliance, i.e. take advantage of the workpieces geometry to slide them into contact and, furthermore, into alignment. For this kind of tasks, impedance control [3] is a useful tool due to its ability

Fig. 1: An illustration of a heavy-duty assembly task with two manipulators grasping workpieces and taking advantage of compliance (source: [2])

to perform both free-space and in-contact motions. Even though compliant motions for dual-arm manipulation can be achieved through pre-planned trajectories [4], this can be a strenuous task.

For efficiently teaching robots new tasks, Learning from Demonstration (LfD) is a well established paradigm [5] where the idea is that the user can simply show the robot the required task, which is often more efficient than designing trajectories. Teaching a dual-arm manipulator using LfD, however, poses certain challenges. Especially when using kinesthetic teaching, often perceived as the easiest methods by users [6], it can be difficult to simultaneously guide two manipulators due to friction at the joints, kinematic redundancy and inertia of the manipulator. Furthermore, kinesthetic teaching cannot be used with heavy robots (Fig. 1) or in dangerous environments, restricting the use of methods limited to kinesthetic teaching only.

We propose to use linear translational motions in one arm along with compliance in both arms to perform dual-arm alignment tasks under pose uncertainty. We first present an analysis of how rotational and translational compliances affect the alignment task. We also show how important the choice of Center of Compliance (CoC) is for a successful alignment. Then we present how this kind of motions for an impedance controller can be efficiently learned from human demonstrations with an LfD method presented in [7], [8]. Finally, we explain how compliances for both arms can be learned from a teleoperated demonstration where one arm is fixed, thus simplifying the demonstration. With the learned compliances we can effectively use contact forces to guide the workpieces during alignment.

This work was supported by Academy of Finland, decision 286580.

¹ M. Suomalainen is currently with University of Oulu, Finland markku.suomalainen@oulu.fi

²School of Electrical Engineering, Aalto University, Finland ville.kyrki@aalto.fi

³ Idiap Research Institute, Switzerland firstname.surname@idiap.ch

In this work we present the following contributions: 1) we present an analysis on the importance of choosing the CoC and compliant axes in an assembly task; 2) we show that the compliance parameters learned by the method in [8] can be used in both arms performing a dual-arm assembly task, even from demonstration with one arm only; and 3) we show through experiments that having compliance on both manipulators eases an assembly task by decreasing the time and joint motions required to complete the task.

II. RELATED WORK

There exist several works on leveraging geometry in single- and dual-arm assembly. Verscheure et al. [9] showed how probing motions can be used to deduce geometry of a task and Karayiannidis et al. [10] presented how to calibrate tools when facing pose uncertainties due to grasping. Almeida and Karayiannidis [11] performed dual-arm folding assembly by explicitly computing the contact point kinematics and developed a controller based on feedback linearisation. Whereas these papers consider similar problems as we do, the difference is the simplicity: in this paper we show how to maximize the power of simple linear motions with compliance and also how a lay user could teach them to a robot efficiently.

Learning from Demonstration for dual-arm robots has been a topic for research for more than ten years [12]. Silvrio et al. [13] showed how to learn the full pose for a task-parametrized dual-arm task. Recently, Batinica et al. [14] demonstrated how Compliant Movement Primitives [15] can be used to learn from human demonstration dual-arm motions requiring compliance by combining desired joint motion trajectories and corresponding joint torque signals. Whereas this method can learn more complex trajectories in free space, the tight temporal coupling between pose and torque trajectories makes it susceptible to pose errors. As the scope of this paper is in-contact alignment tasks, we prefer error tolerance over complexity of applicable motions and thus base our work on the method from [8].

Recently there has been an increasing interest in learning from teleoperated demonstrations. As shown by Fischer et al. [6], this requires more error tolerance from the learning method. Pervez et al. [16] overcame this problem by manually choosing one good demonstration as a baseline. Havoutis and Calinon [17] looked into mixed teleoperation, where the human would first guide the robot through teleoperation, but the robot could then finish the task autonomously. Thus, a method that can learn from a teleoperated demonstration will have a broader range of applications than learning from kinesthetic teaching only.

A strategy for dual-arm peg-in-hole assembly with a compliant manipulator was presented by Zhang et al. [18]. They proposed a hand-crafted two-phase method to overcome positional inaccuracies, creating a sequential search using feedback from a Force/Torque (F/T) sensor. Their observation from a human performing the task was that after initiating contact (i.e. their second phase, which corresponds to the work in this paper), the human mainly adjusts the

orientation of the workpieces. This enhances our intuition that for a dual-arm peg-in-hole task, actively moving both arms does not necessarily improve completion of the task compared to having a single-arm operation. However, we observe that a human performing the peg-in-hole task with two arms allows the passive arm to move in both orientation and translation according to the motions of the active arm moving the peg. This motivates our research on the properties of dual-arm alignment task with rotational and translational compliance on both arms.

III. METHOD

We consider a scenario where a *master* manipulator is moved actively along a linear trajectory in the Cartesian space. The master can also be compliant (defined as a stiffness of 0 in this paper) in rotation and/or translation, in which case environmental forces can displace it from the linear trajectory. The *slave* manipulator is commanded to stay stationary, but may be set compliant along chosen axes and thus move because of contact forces. We assume both manipulators are grasping their respective workpieces to be aligned and that the uncertainties regarding the location of the workpieces tooltips are low enough such that the workpieces can be brought into contact. In addition, we assume the CoC of each manipulator match the Tool Coordinate System (TCS) and can be set arbitrarily. We do not require the manipulators to be centrally controlled.

An illustration of a peg-in-hole scenario meeting the preceding requirements is shown in Fig. 2. In this scenario the task is to insert the peg into the hole by moving the right master manipulator along the direction of the tool, the z-axis in TCS as shown in Fig. 2.

Methods presented in this paper intend to make the convergence region of the manipulators, i.e. the margin of error from which the alignment is accomplished, as large as possible. In addition, we investigate whether there are additional benefits for using dual-arm compliant robots for assembly, such as smaller workspace requirements. In Section III-A we observe how different compliance parameters can be set to enlarge the convergence region and decrease the required joint motions. In Section III-B we shortly present the method from [8] for learning compliant behaviour from human demonstration and how it is applied to a dual-arm scenario. Finally, we also show how to control the manipulators during reproduction of the motion.

A. Mechanics of compliant dual-arm assembly

When performing dual-arm assembly with compliance, there are a number of possible failure types and variables that affect them. In [19] there is a thorough analysis on two possible insertion failures, jamming and wedging, which depend on the geometry and forces affecting the peg. We assume that our workpieces cannot be deformed, which is required in wedging, but jamming can occur with high friction and high error in orientation, as explained in [19]. However, when also the hole is compliant, the problem becomes different. If there is enough friction the peg can still

Fig. 2: Possible outcomes of dual-arm peg-in-hole with compliance. The right manipulator is the master, moving along the direction of the arrow, and the left one is the slave.

be jammed, but there is also a new failure mode, illustrated in Fig. 2c.

Considering the figure, when the peg is moved along the direction of the arrow and both master and slave are compliant, there are three possible outcomes: (a) jamming, which means that the peg does not move at all or very little, (b) the alignment is completed and (c) both workpieces start rotating in the wrong direction and the task fails. Besides the geometry of the workpieces, this behaviour depends on three factors: 1) center of compliance, 2) axes of compliance and 3) the orientation error between the peg and the hole. As factor 3) depends on the setup, it is considered only in a comparative manner in Section IV. In the next sections we consider the effect of the first two factors.

1) *Center of Compliance*: CoC is the point where the stiffness matrix is diagonal in coordinates whose origin is at the CoC [19]. This means that both the control commands and the axes of compliance are defined in a coordinate system at this point. Using Fig. 3, we investigate two possible choices: having the CoC at the wrist of the manipulator (blue manipulators) or having it at the tooltip (orange manipulators).

We first analyse the scenario where the slave manipulator is fixed (gray colour), i.e. the setting is similar to a single-arm peg-in-hole with the master manipulator moving the peg forward as explained earlier, while the master manipulator is rotationally compliant in either tooltip (Fig. 3b) or wrist (Fig. 3d) but there is no translational compliance. The key point is to observe that in the starting scenario (Fig. 3a), the origin of the TCS at tooltip is already in the hole and thus the rotational compliance is enough to align the master tool with the slave, whereas the TCS at the wrist would require compliance in translation along x-axis to align the tools. However, as there is no translational compliance in this scenario, the tool of the master can only rotate to the position shown in Fig. 3d. As this is not physically possible, the motion would result in jamming and a failed task.

The results are identical if the master is stiff (i.e. motion but no compliance) and the slave manipulator is set to rotationally compliant. Fig. 3c shows the slave manipulator rotating around the CoC at tooltip caused by the external

Fig. 3: Illustration of the relevance of choosing the CoC (●). The orange manipulators have the CoC at tooltip and the blue manipulators at the wrist. (a) shows the starting poses of manipulators of each case, along with the coordinate frame used throughout this paper. Figures on the left, (b) and (d), represent the end result of the master having rotational compliance at different CoCs. Figures on the right ((c), (e)) represent the end result of the slave having rotational compliance at different CoCs. The blue manipulators would not actually move due to physical constraints, but the figures show how they could rotate without the constraints.

forces of the master such that alignment is possible. Rotation around the wrist of the slave manipulator would result in the pose of Fig. 3e, and thus jamming and a failed task. This analysis is further examined in the experiments in Section IV. We note that with translational compliance the alignment would succeed with either location of the CoC, but the fact that rotational compliance at tooltip is more efficient is expected to have an effect on the results.

2) *Axes of compliance*: Starting again from the scenario of Fig. 3a, without any compliance it is impossible to perform alignment. As discussed in Section III-A.1, with a correctly chosen CoC, rotational compliance only is sufficient for completing the alignment. However, the results in [8] show that adding translational compliance increases the convergence region for a single-arm task. As our goal is to maximize the convergence region and consider methods that accomplish the alignment regardless of the choice of CoC, the option of single manipulator rotation-only compliant will not be considered further.

We consider the most prominent types of compliance for achieving the goal regardless of the choice of CoC to be:

- compliance in translation and rotation in the master only (e.g. single-arm peg-in-hole, the exact setting from peg-in-hole experiments in [8]),
- compliance in translation and rotation in the slave only,
- compliance in translation and rotation in both robots,
- compliance in rotation only in both robots.

There exist several other combinations in which the task can succeed as well, such as both translational and rotational

compliance in one manipulator and either translational or rotational compliance on the other. However, we will focus on the four itemized cases as they can be considered the corner cases, and thus we expect them to present the most interesting results.

The first case is essentially a single-arm assembly and serves as a baseline. In the second case we explore if there is difference whether the compliance is in the master or in the slave. In the third option, we expect to see an improvement over first and second option, since now both robots can contribute to the alignment. The fourth option is a theoretically interesting one. As mentioned, one robot being rotationally compliant only at the wrist is insufficient to complete the defined task, but with both robots being compliant in rotation the missing compliance in translation is compensated even if CoC is chosen at the wrist. In contrast, having only translational compliance in both robots is not enough to complete the motion in any CoC unless the rotations are already aligned. These cases are experimentally studied in Section IV.

B. Control and Learning

In this section we shortly present the learning and control methods from [8] that we use for first learning the controller parameters and then performing the actual alignment. For performing the alignment motion we define impedance controllers for both manipulators, separately for translations and rotations, as

$$\begin{aligned} \mathbf{F}_C &= K_f(\mathbf{x}^* - \mathbf{x}) + D_f \mathbf{v} \\ \mathbf{T}_C &= K_\theta(\boldsymbol{\beta}^* - \boldsymbol{\beta}) + D_\theta \boldsymbol{\omega} \end{aligned} \quad (1)$$

where $\mathbf{F}_C, \mathbf{T}_C$ are the commanded forces and torques, $\mathbf{x}^*$ is the desired position, $\mathbf{x}$ the current position, $\boldsymbol{\beta}^*$ the desired orientation, $\boldsymbol{\beta}$ the current orientation, K_f and K_θ stiffness matrices where the axes of compliance are defined and $D_f \mathbf{v}$ and $D_\theta \boldsymbol{\omega}$ linear damping terms.

The trajectory for the master manipulator is computed in a feed-forward manner as

$$\mathbf{x}_t^* = \mathbf{x}_{t-1}^* + \hat{\mathbf{v}}_d^* \nu \Delta t \quad (2)$$

where Δt is the sample time of the control loop, $\hat{\mathbf{v}}_d^*$ the desired direction and ν the translational velocity. The goal of the method from [8] is to learn parameters $\hat{\mathbf{v}}_d^*$, K_f and K_θ for both manipulators from one or more human demonstrations performed such that the slave manipulator is stiff.

The intuition for learning the desired direction $\hat{\mathbf{v}}_d^*$ stems from geometry: to slide the robot's end-effector along a surface, there is always a friction-dependent sector s of directions from which the robot can apply a force to accomplish the sliding. If this sector is calculated at intervals over a whole demonstration, the intersection of all sectors s_i would signify a direction which can lead the end-effector through the whole demonstrated motion either in free space or in contact. We call sector s a set of desired directions and it is visualized for a single time-instant in 2-D in Fig. 4.

When the manipulator is teleoperated, the sum of environmental forces $\mathbf{F}$ from Fig. 4 can be estimated with the

Fig. 4: Forces acting on the end-effector during a sliding motion. $\mathbf{F}$ is the force applied by the controller during a teleoperated demonstration, $\mathbf{F}_N$ the normal force, $\mathbf{F}_\mu$ the friction force, $\mathbf{v}_a$ the actual direction of motion and s the sector of desired directions.

desired Cartesian force of the teleoperated arm's controller, which is of similar form as (1). Therefore, as seen from Fig. 4, the sum of environmental forces for a tool sliding with constant speed can be written as

$$\mathbf{F} = \mathbf{F}_N + \mathbf{F}_\mu \quad (3)$$

where $\mathbf{F}_\mu = -|\mu \mathbf{F}_N| \hat{\mathbf{v}}_a$ is the force caused by Coulomb friction with μ being the friction coefficient and $\hat{\mathbf{v}}_a$ a unit vector representing the direction of motion, and $\mathbf{F}_N$ the normal force. Similarly, the environmental torque can be estimated from the torque that the controller applies, written as

$$\mathbf{T} = \mathbf{r} \times \mathbf{F}_\mu + \mathbf{r} \times \mathbf{F}_N \quad (4)$$

where $\mathbf{T}$ is the controller torque, $\mathbf{r}$ the position vector between the measurement point and location of contact.

In [8] demonstrations were performed by kinesthetic teaching, i.e. grabbing the robot and leading it through the motions. However, there are cases where kinesthetic teaching is not feasible, for example if the robot is operating in a confined or dangerous workspace. In [20] the authors showed that for a 2 degrees of freedom task the method from [8] can successfully learn compliant motions from teleoperated demonstrations, even though in general providing demonstrations by teleoperation is more difficult than by kinesthetic teaching [6]. In this paper we show that the method from [8] can be directly applied to teleoperating a robot manipulator by moving another manipulator, similarly as in [17].

The key assumption for detecting the axes of compliance is that if there is motion in other directions besides $\hat{\mathbf{v}}_d^*$, that motion must be caused by the environment and therefore compliance is required in the direction of that motion. To find the axes, we perform Principal Component Analysis (PCA) on the motion in other directions besides $\hat{\mathbf{v}}_d^*$ to find out where the most prominent motion besides $\hat{\mathbf{v}}_d^*$ has occurred. We compute likelihoods of how well each PCA vector fits the data. Based on the likelihoods, we use Bayesian Information Criterion (BIC) [21] to decide which of the PCA vectors need to be compliant, thus constructing K_f and K_θ .

Initially, these K_f and K_θ are learned for the master manipulator. We propose using the same matrices directly for the slave manipulator as well, defined in the slave's TCS. We can see from Figs. 3b and 3c that both manipulators need to rotate around the same axis according to their respective TCSs to succeed in alignment. Additionally, as

Fig. 5: The experimental setup: two Panda robots (peg attached to the left, master, and hole attached to the right, slave) used for performing the task, and the Baxter robot used for teaching by teleoperation.

the workpieces are essentially coupled when touching each other, their translational motion caused by compliance will also occur along the same direction. We argue that the difference in the translational axes of compliance caused by the misalignment of the workpieces is not significant enough to have a major effect on the task.

IV. EXPERIMENTS AND RESULTS

We performed experiments with a Rethink Robotics Baxter dual-arm manipulator and two Franka Emika Panda manipulators. We used a peg-in-hole setup made of PVC plastic, where the peg has radius 16.5 mm, length 80 mm and a rounded tip, and the clearance with the hole is 0.25 mm. The arrangement is depicted in Fig. 5. The part with hole was attached to the wrist of the right Panda robot (in the view of Fig. 5) and the peg was attached to the wrist of the left Panda robot, the master robot. Instead of grasping, the peg and hole were rigidly attached to the wrists of the robots to allow better measurements for evaluating the method.

For the demonstrations, we teleoperated one of the Panda arms with one of the Baxter arms, while the other Panda arm was set stiff. The teacher moved the Baxter arm kinesthetically by using a controller compensating the effect of gravity. The teleoperation was performed with relative motions in TCS, i.e. the Panda repeated in its own TCS the motions that were performed with the Baxter in the Baxter’s TCS. Motion from the master Panda was recorded and the environment wrenches were estimated based on the desired controller wrenches. With the method from [8] we learned the parameters required to perform the peg-in-hole alignment in the case of a stiff slave. The learned desired direction of translation $\hat{v}_d^*$ was along z-axis in TCS. No desired direction of orientation was required. One axis of compliance in translation detected, along the x axis. One axis of compliance in rotation detected, around the y axis.

The parameters were learned for a CoC at the wrist only. In [8] we already analysed the possibilities of learning in different coordinate systems. As these parameters are also what our reasoning in Section III-A suggests, for the sake of comparison we use them with all our experiments. However, to get a better idea of which of the parameters are most significant in terms of performance, we chose only some of the parameters for certain tests, but the individual axes were

Configuration	A	B	C	D	E
Max. orientation error (deg)	4	12	-	8	12

TABLE I: A table summing up the maximum orientation errors with which the alignment was successful. In case C the reproduction failed even with 4 degrees error.

learned from demonstrations. Thus, from now on when we write *compliance in rotation*, we mean compliance around y axis as learned, and similarly for translation. By varying the number of compliant axes, we studied the following five compliance configurations:

- A Compliance in translation and rotation in master’s wrist only, similar to teaching.
- B Compliance in translation and rotation in master’s tooltip only.
- C Compliance in translation and rotation in slave’s tooltip only.
- D Compliance in rotation only, both master and slave tooltip.
- E Compliance in rotation and translation in both master and slave tooltip.

A. Convergence region

With the term *convergence region* we mean the region from which reproduction with the same set of parameters is possible. More specifically, in these experiments we mean the maximum orientation error in the rotation around y axis in TCS from which the reproduction is successful. Even though the resulting values are specific for our setup, we can use this measure to compare configurations A-E. Table I shows the largest angles for which the alignment was successful with each set of parameters. During the experiment the orientation error was increased with 1 degree resolution from 4 degrees until the reproduction failed. Each experimental condition was repeated 5 times. In all cases all 5 repetitions were either successful or unsuccessful, which is why measurement uncertainty is not presented in Table I.

The first choice was the location of CoC, i.e. comparison between experiments A and B. It was observed that compliance in tooltip has a more significant effect on the convergence region, as expected after the analysis of Section III-A.1. After this result, we performed the rest of the experiments with CoC at tooltip.

In condition C, only the slave was set compliant, with the CoC at tooltip. However, with this setting even 4 degrees error prevented the motion. This is contradictory to our analysis in Section III-A.1. A probable explanation is the internal friction at the robot joints. Since there are already forces applied at the joints of the master, it requires only a small amount of additional force to cumulatively overcome the static friction. However, in the slave manipulator, overcoming the static friction would require higher forces, and thus in practice the compliance in the master facilitates the alignment more than compliance in the slave.

Finally we set both manipulators compliant at tooltip, with the parameters learned earlier used for the respective TCS of each manipulators. First we set both master and slave

manipulators compliant in rotations only (condition *D*), and then both manipulators to compliant in both translation and rotation (condition *E*). It can be observed that translational compliance is important in this setting, and setting both robots to compliant in both rotation and translation increased tolerance to errors in the starting orientation. However, having both manipulators compliant did not achieve larger region of convergence than having only master manipulator compliant (conditions *B* and *E*).

We conclude that to maximize robustness against positional errors in an alignment task, at least the master arm should be compliant in both translation and rotation at tooltip. The translational compliance seems to play a key role, since rotational compliances only in both master and slave did not ease the alignment as much, even though when the coordinate system is set to tooltip there is only rotational error, not translational. In the next section we investigate how the motions are affected when also the slave is compliant.

B. Motion analysis

We recorded the joint positions, forces and duration of the motion during reproduction. Each of these criteria can present certain advantages. If the workspace of the robot is limited, it would be advantageous that the robot can perform the motion in a small space. Minimizing the joint motions can also help keep the robot operational for longer time. Performing the motions faster would also be an advantage. Finally, keeping the forces lower decreases the probability of breaking the workpieces being aligned.

We observed, as was already shown by Zhang et al. [18], that having both arms compliant reduces the amount of forces applied to the end-effector. As they did a thorough numerical validation of this, we will not present these results from our experiments for brevity but concentrate on the metrics which they did not analyse.

We compared the metrics of conditions *B*, *D* and *E* at 8 degrees error since this was the maximum orientation error where the motion was successful in all these three conditions. From 5 repetitions for each condition, we computed the means of the following metrics, summed over all joints: total covered joint distance (rad), total covered joint distance weighed by link lengths, joint covered distance maximum (the most motion by a single joint) and movement duration (seconds). The results are shown in Table II. When applicable, the distances of the two arms are presented separately.

It can be seen that having compliance in both manipulators (conditions *D* and *E*) produces superior results across all measures. In addition, condition *E* outperformed condition *D*, showing that translational compliance plays a significant role in an alignment task. Furthermore, we empirically observed that having both manipulators compliant greatly eased the final push when the workpieces were nearly aligned.

C. Discussion

In the results of Table I the failure type that confined the convergence region was always jamming (Fig. 2a). In the initial trials we also observed failures according to Fig. 2c

		Configuration				
		B	D		E	
			M	S	M	S
Metric	<i>Joint dist</i>	1.45	0.64	0.35	0.41	0.18
	<i>W. joint dist</i>	1.18	0.53	0.23	0.36	0.15
	<i>Max. joint dist</i>	1.1	0.47	0.19	0.32	0.13
	<i>Time</i>	26.5	13.0		5.00	

TABLE II: Various metrics evaluated for a successful alignment task with 8 degrees rotational error. M stands for master manipulator and S for slave, W. means weighed with the link lengths and Max. the maximum norm

when the orientation error was considerably higher than the maximum orientation error in Table I. Thus we reason that failure type presented in Fig. 2c occurs either with high orientation error or low friction.

It should also be noted that the set of parameters learned from the human demonstration were the ones that produced best results in the experiments when applied to both manipulators. This experiment showed that the parameters can be learned even when the demonstrations are performed with a difficult teleoperation method. Finally, we also validated that with the same parameters the task can also be completed when the error angle is in the other direction (i.e. the master is sliding "up").

V. CONCLUSIONS AND FUTURE WORK

We presented an analysis on how to choose the center and axes of compliance for a dual-arm alignment task where the master moves along a linear trajectory and both arms can be set to compliant. In addition, we showed how to easily learn these parameters from a human demonstration performed with one teleoperated manipulator only. A slightly unexpected result was that adding compliance to the slave manipulator in addition to the master did not increase the convergence region. We suspect the reason to be similar that prevented alignment with compliance in slave manipulator only: due to the static joint frictions, the master manipulator must move first and the slave compliance will only have an effect later in the motion. We showed through experiments that slave compliance in addition to the master had other significant benefits, namely shorter joint motions and time.

In this paper we concentrated more on the rotational compliance, the main reason being that the assumed error was in rotation. However, the results showed that translational compliance is a major factor in easing the alignment, even though the analysis from Section III-A suggests that rotational compliance suffices. The exact mechanics for the role of translational compliance would require further mechanical analysis of the system.

An interesting concept would be to have the CoC of the slave match that exactly of the master—in other words, make the CoC change dynamically during the motion according to the pose of the master. Our experimental setup did not allow a dynamically changing CoC, but the experiment that produced the best results, having the CoC of both robots at the tooltip is an approximation of this.

REFERENCES

- [1] C. Smith, Y. Karayiannidis, L. Nalpantidis, X. Gratal, P. Qi, D. V. Dimarogonas, and D. Kragic, "Dual arm manipulation—a survey," *Robotics and Autonomous systems*, vol. 60, no. 10, pp. 1340–1353, 2012.
- [2] M. Suomalainen and V. Kyrki, "Learning compliant assembly motions from demonstration," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 871–876, IEEE, 2016.
- [3] N. Hogan, "Impedance control: An approach to manipulation: Part ii—implementation," *Journal of dynamic systems, measurement, and control*, vol. 107, no. 1, pp. 8–16, 1985.
- [4] F. Caccavale, C. Natale, B. Siciliano, and L. Villani, "Achieving a cooperative behavior in a dual-arm robot system via a modular control structure," *Journal of Robotic Systems*, vol. 18, no. 12, pp. 691–699, 2001.
- [5] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and autonomous systems*, vol. 57, no. 5, pp. 469–483, 2009.
- [6] K. Fischer, F. Kirstein, L. C. Jensen, N. Krüger, K. Kukliński, T. R. Savarimuthu, *et al.*, "A comparison of types of robot control for programming by demonstration," in *The Eleventh ACM/IEEE Int. Conf. on Human Robot Interaction*, pp. 213–220, IEEE Press, 2016.
- [7] M. Suomalainen and V. Kyrki, "A geometric approach for learning compliant motions from demonstration," in *Humanoid Robots (Humanoids), 2017 IEEE-RAS 17th International Conference on*, pp. 783–790, IEEE, 2017.
- [8] M. Suomalainen and V. Kyrki, "Learning 6-d compliant motion primitives from demonstration," *Autonomous Robots*, 2019, submitted. arXiv:1809.01561.
- [9] D. Verscheure, J. Swevers, H. Bruyninckx, and J. De Schutter, "On-line identification of contact dynamics in the presence of geometric uncertainties," in *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, pp. 851–856, IEEE, 2008.
- [10] Y. Karayiannidis, C. Smith, F. E. Vina, and D. Kragic, "Online contact point estimation for uncalibrated tool use," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pp. 2488–2494, IEEE, 2014.
- [11] D. Almeida and Y. Karayiannidis, "Folding assembly by means of dual-arm robotic manipulation," in *Robotics and Automation (ICRA), 2016 IEEE International Conference on*, pp. 3987–3993, IEEE, 2016.
- [12] R. Zöllner, T. Asfour, and R. Dillmann, "Programming by demonstration: dual-arm manipulation tasks for humanoid robots.," in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 479–484, 2004.
- [13] J. Silvério, L. Rozo, S. Calinon, and D. G. Caldwell, "Learning bimanual end-effector poses from demonstrations using task-parameterized dynamical systems," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 464–470, IEEE, 2015.
- [14] A. Batinica, B. Nemec, A. Ude, M. Raković, and A. Gams, "Compliant movement primitives in a bimanual setting," in *Humanoid Robotics (Humanoids), 2017 IEEE-RAS 17th International Conference on*, pp. 365–371, IEEE, 2017.
- [15] M. Deniša, A. Gams, A. Ude, and T. Petrič, "Learning compliant movement primitives through demonstration and statistical generalization," *IEEE/ASME transactions on mechatronics*, vol. 21, no. 5, pp. 2581–2594, 2016.
- [16] A. Pervez, A. Ali, J.-H. Ryu, and D. Lee, "Novel learning from demonstration approach for repetitive teleoperation tasks," in *IEEE World Haptics Conference (WHC)*, pp. 60–65, 2017.
- [17] I. Havoutis and S. Calinon, "Learning from demonstration for semi-autonomous teleoperation," *Autonomous Robots*, 2018.
- [18] X. Zhang, Y. Zheng, J. Ota, and Y. Huang, "Peg-in-hole assembly based on two-phase scheme and f/t sensor for dual-arm robot," *Sensors*, vol. 17, no. 9, p. 2004, 2017.
- [19] D. E. Whitney, "Quasi-static assembly of compliantly supported rigid parts," *Journal of Dynamic Systems, Measurement, and Control*, vol. 104, no. 1, pp. 65–77, 1982.
- [20] M. Suomalainen, J. Koivumäki, S. Lampinen, V. Kyrki, and J. Mattila, "Learning from demonstration for hydraulic manipulators," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, p. 3579–3586, 2018.
- [21] G. Schwarz *et al.*, "Estimating the dimension of a model," *The annals of statistics*, vol. 6, no. 2, pp. 461–464, 1978.