Conditional Visual Servoing for Multi-Step Tasks

Sergio Izquierdo*, Max Argus*, Thomas Brox

Abstract— **Visual Servoing has been effectively used to move a robot into specific target locations or to track a recorded demonstration. It does not require manual programming, but it is typically limited to settings where one demonstration maps to one environment state. We propose a modular approach to extend visual servoing to scenarios with multiple demonstration sequences. We call this conditional servoing, as we choose the next demonstration conditioned on the observation of the robot. This method presents an appealing strategy to tackle multi-step problems, as individual demonstrations can be combined flexibly into a control policy. We propose different selection functions and compare them on a shape-sorting task in simulation. With the reprojection error yielding the best overall results, we implement this selection function on a real robot and show the efficacy of the proposed conditional servoing. For videos of our experiments, please check out our project page: https://lmb.informatik.uni-freiburg.de/projects/conditional_servoing/**

I. INTRODUCTION

Programming robots to complete tasks remains a challenging and time-consuming problem, which prevents robots from being used more widely. Servoing-from-Demonstration (SfD) approaches enable few-shot imitation from demonstration videos and, thus, are an attractive direction for robot programming. The key challenge is how to make SfD approaches flexible enough to automatically adapt to many variations of a task without the need to record a demonstration for each particular case.

Our work builds on FlowControl [1], which servos to the configuration specified in a demonstration frame, then iterating to the next demonstration frame once one is close enough. This process is repeated for all demonstration frames, enabling the replication of demonstration tasks. The algorithm requires RGB-D observations as well as foreground object segmentations for the demonstrations. Robustness to variations in appearance is achieved by establishing correspondence between observation and demonstration via learned optical flow.

FlowControl is bound to a single demonstration that must match the control problem. In this paper, we allow for a set of demonstrations to choose from. This enables the option to select the most suitable demonstration, conditioned on the initial demonstration frame, which is why we call the approach conditional servoing. Moreover, it enables the recombination of subtask demonstrations, similar to skills, to

* equal contribution authors

All authors are at the University of Freiburg and are members of Brain-Links-Brain-Tools. The corresponding author can be contacted via: argusm@cs.uni-freiburg.de
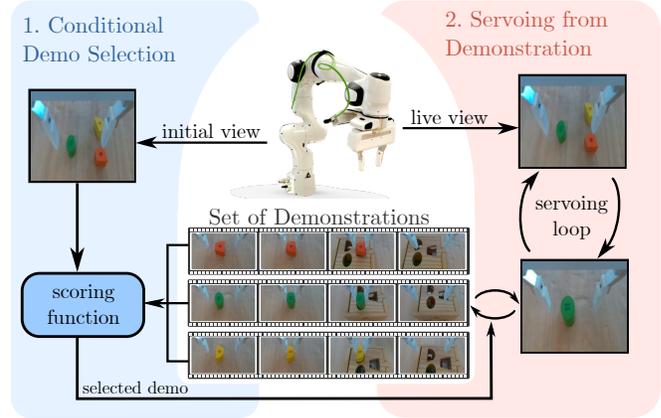
Fig. 1: In conditional servoing, the most suitable demonstration is selected. In our approach, this is done by evaluating a distance score between initial live views and initial demonstrations views, shown in the blue half. This demonstration is then used with a servoing algorithm designed to replicate demonstration sequences to complete tasks, shown in the red half. A more detailed description of the method is shown in Figure 2. The images are from an end-of-arm camera, which we use for robot control.

more complex policies and avoids the necessity of a complete demonstration for each possible permutation of the scene.

More precisely, the system is given a set of demonstrations of different subtasks, some of which may be unsuitable to follow, for instance, the object is not currently visible. In conditional servoing, the system must select the best demonstration to follow based on its current observation. To do so, we propose to perform the selection based on a matching score that assesses the similarity between the demonstration and the current observation

We evaluated a number of different matching scores in simulation and verified that the most promising scoring metric works well also on a real robot. Our experiments in simulation and on the real robot deal with a shape sorting task in which multiple blocks are present in the scene and must be inserted in the right opening of a sorter box. This task is challenging from the control point of view, due to the tight holes of the sorter box, and requires the selection of the right entry point. Conditional servoing increases the success rate of FlowControl on the real robot from 40% to 78%. Furthermore, conditional servoing significantly increases the chances of correctly recovering by re-grasping if an object is dropped by mistake. In addition, we perform a qualitative

evaluation on a number of related configurations.

II. RELATED WORK

Servoing: Visual Servoing [2], [3] is a well-established approach to vision-based robotics problems, in which a comparison between sensor observations and a target view, or demonstration, is used to generate control commands in a closed feedback loop. This allows the control to adapt to the surrounding environment as observed by the system. A large number of robotics problems can be approached with visual servoing techniques [4], examples of this include manufacturing [5] and surgery [6], among many others.

While we did not find any prior works featuring demonstration selection in visual servoing, there are some related concepts. In hybrid visual servoing, control laws are extended with additional constraints [7]. This allows the definition of different subtasks, e.g., obstacle avoidance, manipulation tasks, and joint-limits avoidance [8]. Subtasks are usually defined simultaneously and often pertain to some constraining of the robot state. Task sequencing in visual servoing describes a situation, where the control laws are changed over the course of iteration [9]. In contrast to these works, we use the term *tasks* to describe the problems in the environment that the robot should complete.

Our work extends FlowControl [1], which combines the use of learned optical flow with demonstration sequence tracking to replicate tasks. The use of learned optical flow for servoing is similar to [10], [11], [12], and the use of demonstration sequence tracking is similar to [13], [6]. For the sake of completeness, we briefly recapitulate how FlowControl works. It has two main parts, a frame alignment part and a sequence tracking part. The frame alignment part computes the transformation between the live state, as seen by the robot's camera, and the state shown in the selected demonstration frame. This transformation is used to move the robot accordingly. The method relies on optical flow estimates to obtain correspondences between both states. The correspondences are further filtered using a foreground segmentation (attention mask) of the demonstration before being converted into 3D points using the camera's depth measurements. Finally, a rigid 3D transformation is computed between the point clouds. After achieving a good alignment with the current frame, the method performs sequence tracking by iterating to the next demonstration frame, allowing for the replication of tasks. A visual explanation of the method is available on Youtube: `https://youtu.be/SSC8RyBGUug`.

Compositionality: Individual demonstrations can be seen as encoding different skills. A large number of robotics works investigate compositional skills. These include [14], associative skill memories [15], as well as the closely related hierarchical approaches [16]. More recent works have extended this line of work to the deep-learning domain [17], [18].

Frame Selection: Selecting a suitable demonstration sequence can be addressed as a registration problem, in which frames from demonstrations are registered to a live view. A large number of similar problems exist in computer vision. Image retrieval tries to find the most similar image from a database of candidates [19], [20], [21]. It is the basis of visual place recognition, where one tries to localize based on a query photograph [22], [23]. Some early works using visual servoing can be regarded as predecessors of SLAM [24]. More recent incarnations of similar ideas can be found, for example, in [25]. In the field of SLAM, similar problems are encountered in global localization [26] or re-localization [27], which is closely related to the problem of loop closure.

Registering images to a map is also encountered in Structure-from-Motion (SfM), in which large numbers of images are registered to a scene [28].

The registration methods described here differ substantially in the representations they compare. SLAM methods often compare sequences of observations to maps. We are interested in comparing single images to single images. Furthermore, we are interested in just the registration of the relevant foreground objects. Thus, these methods cannot be used in our approach out-of-the-box. A similar setting is encountered in the context of object pose estimation. There is also the problem of template-based object instance detection [29], [30], in which objects need to be localized within a scene.

III. METHOD

This work considers a setting where multiple demonstration sequences are available and a suitable demonstration must be chosen. We call this conditional servoing because our choice of sequence depends on observations of the current environment.

Our setting implies environments in which correspondences between the current state and some of the demonstrations may be poor or even nonexistent. Hence, we investigate selection mechanisms to retrieve the most suitable demonstration from a set of multiple demonstration sequence options. This allows following the selected demonstration and discarding unsuitable ones.

This selection is implemented as an outer loop around a servoing algorithm. The loop starts by selecting the best demonstration and replicating it. Once it is completed, the next demonstration can be selected and replicated. This allows addressing multi-step tasks. A visual depiction of our method can be seen in Figure 2.

In the following, we describe the setting formally. We want to select the best demonstration, D_{best}, from a set of demonstrations, $\mathcal{D} = \{D_1, D_2, ..., D_n\}$. Given a demonstration D_i, our selection is based on only d_i, its first frame, as well as the current observation of the robot, x. To do so, we employ a distance function to assign a score to each demonstration, based on the similarity to the current observation, $f(x, d)$

As we build on FlowControl as the Servoing-from-Demonstration method in the inner loop, we can start with evaluating the residual of the rigid-transformation estimation to assess the quality and reliability of the matching. For a
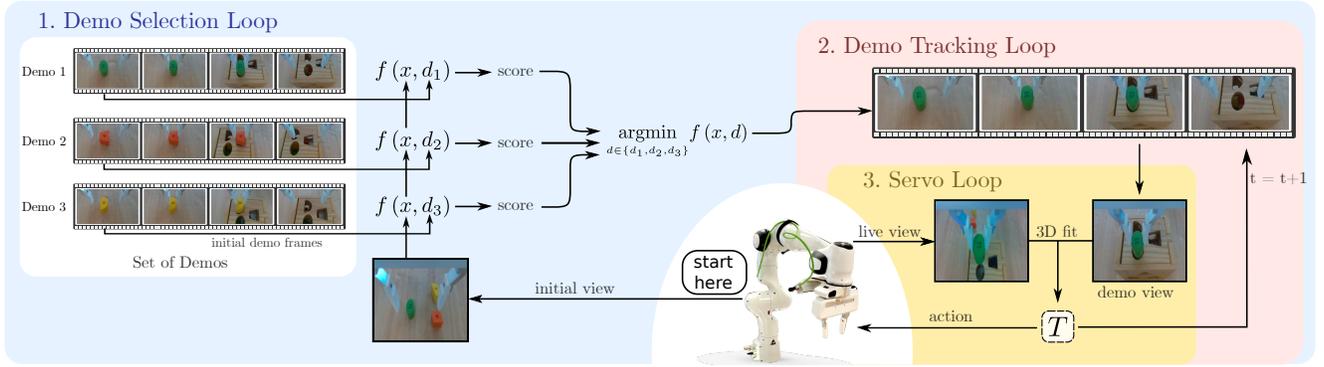
Fig. 2: The demonstration selection, shown in blue, is performed by taking the first frame from a demonstration sequence and comparing it to the initial live view. We then select the closest demonstration and continue with servoing from a demonstration sequence using FlowControl, shown in red and yellow. Once a demonstration sequence is done, demonstration selection can be repeated until the overall task is completed.

more detailed description of FlowControl, please see Section II.

We propose four different distances to assess the quality of a given demonstration:

Point-Quality Distance: We consider using the estimated 3D transformation between the live and the demo's initial point clouds, p^L, p^D. We transform the live point cloud into the demo's one according to the transformation retrieved by FlowControl, (R, t). After this, outliers that are more than 5mm apart are removed. Finally, the L2 norm between the two point clouds is computed to obtain the distance.

$$D_{\text{PQ}} = \frac{\sum_{i \in \Omega_p} \| (R \cdot p_i^D + t) - p_i^L \|}{| \Omega_p |} \tag{1}$$

Color-Quality Distance: Instead of comparing the 3D positions of the points, in this distance, we focus on the points' colors, $C(p)$. The L2 distance between the RGB values of the point cloud correspondences is calculated. This distance aims to detect situations where the point clouds obtain a good registration, but the colors of the object in the live state and in the considered demonstration do not match. As in the previous distance, possible outliers have been discarded, and are not considered in the distance.

$$D_{\text{CQ}} = \frac{\sum_{i \in \Omega_p} \| C^D(p_i^D) - C^L(p_i^L) \|}{| \Omega_p |} \tag{2}$$

Reprojection Distance: Similar to the color-quality distance, the reprojection distance evaluates the accuracy of the RGB correspondences on the basis of images. As point correspondences between the current observation and the demonstration frames are obtained with an optical flow estimation, we propose to evaluate the suitability of the demonstration by evaluating the quality of the flow. As we, naturally, lack the optical flow ground truth, we follow the approach of some self-supervised flow estimation methods: the use of photometric reconstruction as the evaluation function [31], [32].

The live view, I^L, is warped into the demo view, I^D, by using the optical flow, u, v, between the demo view and the live view. By computing the dissimilarity between the warped (reconstructed view) and the demo, we obtain a measurement of the fit of the flow. We use the L2 distance between the two, and only calculate the dissimilarity in the pixels selected by the demonstration's foreground mask, Ω_M, as other pixels do not need to be estimated correctly. An example of a reprojection image is shown in Figure 3.

In contrast to the previous functions, this one relies only on two-dimensional images. Moreover, it considers all pixels in the segmentation mask, i.e., it does not discard possible outliers. Outlier removal can have both benefits and drawbacks, while it can make the distances more robust it can also be problematic to apply in situations where the objects are not visible in the current state, which can result in incorrect transformation estimations.

$$D_{\text{RP}} = \frac{\sum_{x, y \in \Omega_M} \| I^L(x + u, y + v) - I^D(x, y) \|}{| \Omega_M |} \tag{3}$$

Learned distance: Trying to find interactions between the previously presented distances, we propose to train a classifier to predict the probability of success of a given pair (x, d). We train a small multilayer perceptron that takes the other distances as inputs and predicts the probability of success. The best demonstration is selected by taking the one with the highest estimated probability of success.

We train a small MLP with three hidden layers of $\theta = 64$, 64, and 32 neurons. The last layer has one neuron to predict the probability of success. We use ReLU activation in all intermediate layers and logistic in the last one. For the loss the binary cross-entropy is used and optimized with *Adam* [33] using its default recommended parameters of $\beta_1 = 0.9$ and $\beta_2 = 0.999$ as well as an initial learning rate of 10^{-5}. We used a batch size of 128 and trained for 2500 iterations or until the learning has converged.

$$D_{\text{MLP}} = 1 - P_{\text{MLP}}(y \mid x, d, \theta) \tag{4}$$

live view	demo view	re-projection
I^L	I^D	$I^L(x + u, y + v)$
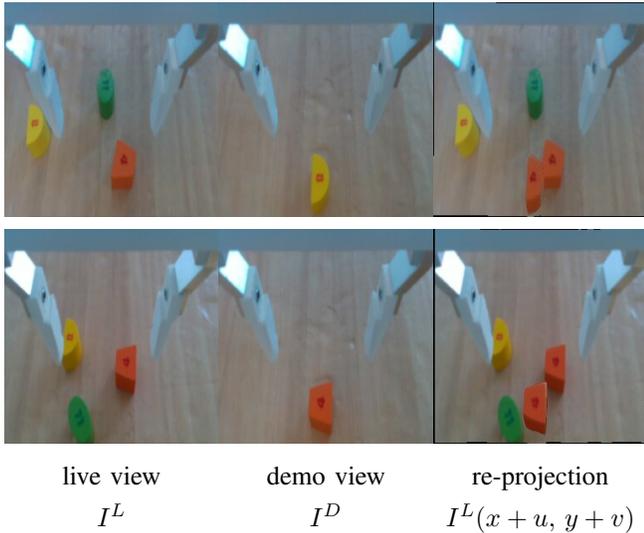
Fig. 3: Reprojection examples. Optical flow is computed from the demo view (center) to the live view (left). Using this flow, we warp the live view to be similar to the demonstration view. The figure shows a bad warping in the top row, where the flow matched the yellow object with the orange one, as well as a good warping in the bottom row.

IV. EXPERIMENTS

We use a shape sorting task to evaluate the conditional servoing. It is a popular multi-step task also used in other works [34], [35]. The robot has to grasp three pieces and introduce them precisely on a shape sorting cube, one after the other. To do so, the robot has to select the best demonstration to follow before grasping the next piece. Given the small tolerances, the task requires very precise grasping, positioning, and insertion of the objects. This is much easier if the optimal demonstration is selected.

A. Experimental Setup

We re-created the physical shape sorting task in the PyBullet physics simulator [36] to afford to collect a large number of executions, something prohibitively expensive in a real robot[1]. As shown in Figure 5, we aligned the simulation and the real robot setup, but the simulation uses the Kuka iiwa robot, whereas the physical robot was a Franka Emika Panda robot. It was equipped with an Intel RealSense D435 RGB-D camera mounted on the robot's end-effector, obtaining an eye-in-hand setup. For the gripper, two 3D printed fingers are attached to the robot's fingers. This setup is shown in Figure 4.

For the control, we restricted the end-effector's freedom such that it always faces downwards, limiting its orientation to the yaw rotation. We parametrize the control as in [1], with a 5DoF continuous action $a = [\Delta x, \Delta y, \Delta z, \Delta\theta, a_{\text{gripper}}]$

---

[1]Note that the simulation was only used for additional evaluation. In contrast to many contemporary robot learning methods, the simulation is not needed to learn the task. The real robot can be instructed by demonstration of the subtasks in the physical world.
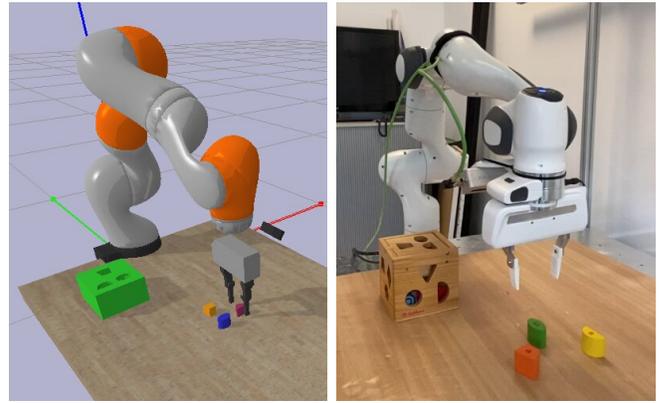
Fig. 4: Robot setup with end-of-arm camera in simulation (left) and in the real environment (right).
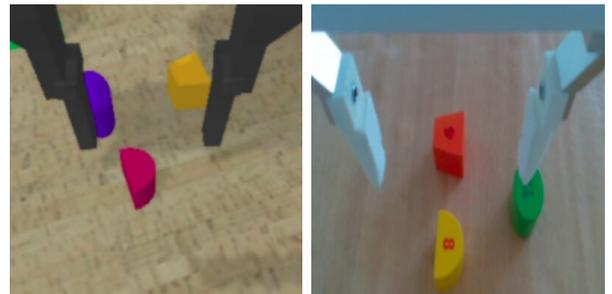
Fig. 5: Simulation environment (left) used to better perform large-scale quantitative evaluation of selection scoring functions compared to the real environment (right).

in the end-effector frame. $\Delta x, \Delta y, \Delta z$ specify a Cartesian offset for the desired end-effector position, $\Delta\theta$ defines the yaw rotation of the end effector, and a_{gripper} is the gripper action that is mapped to the binary command to open or close the fingers.

B. Simulation Experiments

First, we evaluated the scoring functions defined above. Because the simulation allows for effortless larger-scale experiments, we ran this evaluation in simulation. As the MLP distance function requires training, we recorded 500 executions with three objects each, giving a total of 1500 runs. Of these, we use 1000 samples to learn the distance and evaluate on the remaining samples as a test set. Details of the network's architecture and the training process are described in Section III.

The suitability of a demonstration can only be evaluated by tracking it with visual servoing to completion and seeing if the task is successfully completed. Thus, to evaluate a proposed scoring function, we run 100 simulations with each of the three objects, resulting in 300 samples, selecting the optimal demonstration according to the given scoring function. The starting arrangement of the objects in each of the 100 simulations is randomized, but the scoring functions were evaluated with the same set of initial configurations to enable a fair evaluation of their effectiveness.

TABLE I: Simulation success rates for different selection scores evaluated on the simulated shape sorting task.

| Score for argmin sampling | Success Rates [%] | |
of demonstration sequence	First Object	All Objects
Uniform Random	55	47
Point-Quality	69	60
Color-Quality	76	58
Reprojection	81	63
MLP	81	65

Fig. 6: Error analysis for simulation experiments. Error types were assessed manually by observing a set of 60 and 88 failure runs in FlowControl-C and FlowControl respectively. Our method effectively reduces the top types of error of the baseline method.

Table I shows the success rate obtained in simulation when following the proposed distances. We separate the results considering *all objects*, counting all recorded samples, and *first object*, only considering runs inserting the first of the three objects. This distinction shows the benefits of our conditional servoing in multimodal situations, where the flow estimation can match incorrect objects with the demonstrations. The point cloud based methods were also tested without outlier removal, which did not improve results.

All our proposed distances obtain a higher success rate than the uniform random baseline, where demonstrations are picked according to a uniform random distribution. Of these, the reprojection and MLP distance yield similar success rates for the *first object*. The MLP distance achieves slightly higher performance on the *all objects* set, as it was trained on this quantity. However, we still chose the reprojection distance as our preferred method, as it does not require collecting a dataset and training.

Although the color-quality and the point-quality distance may seem more robust, as they discard possible outliers, they achieve a lower success rate than the reprojection distance. We argue that incorrectly computed transformations in the servoing may lead these functions into inconsistent results when the fitting is based on incorrect flow estimates. On the other hand, the reprojection distance directly evaluates the quality of the flow prediction.

Figure 6 shows a detailed manual analysis of a set of 150 simulation runs. For these, we categorize the main cause of failures for the baseline FlowControl method, as well as our conditional servoing method, which uses reprojection distance as selection criterion. This is named FlowControl-C. The three main causes of error in FlowControl stem from wrong optical flow estimates. In *undecided flow* the flow estimation is unstable and jumps between different objects in each step. In *wrong object* cases, the robot grasps an incorrect shape due to a wrong projection, and finally, in *incorrect flow*, estimates are not accurate and, thus, the servoing does not converge or converges to an incorrect robot position. Including our conditional servoing significantly decreases these errors, showing one of the main benefits of our approach: an improvement in correspondence computation.

In an additional experiment, we show that the conditional servoing approach can also help when recovering from failure cases. To evaluate this in simulation we generate failure cases by opening the gripper according to a Bernoulli
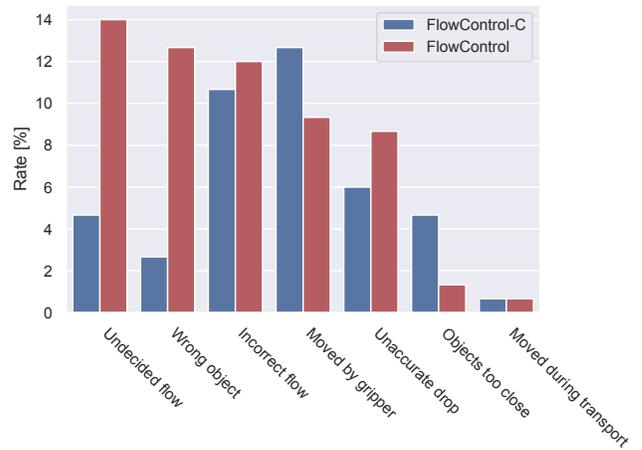
distribution with a probability of $p = 0.25$. To evaluate recovery we compare to the baseline of tracking the previously followed trajectory, which results in a success rate of 27%. This is clearly outperformed by the conditional selection of the trajectory which has a success rate of 59%.

C. Real Experiments

To test the effectiveness of our proposed conditional servoing we use a shape sorting task on the real robot. Based on the simulation experiments we selected the reprojection distance as the best overall selection score and subsequently carried out all our real robot experiments with this choice. These experiments are named FlowControl-C. An example of reprojected images is given in Figure 3. We then performed quantitative and qualitative experiments with our new method.

Success Rate: Similar to experiments in simulation we start by evaluating the success rate achieved on the shape sorting task when using the baseline FlowControl method compared to our FlowControl-C method with the reprojection distance. The shape sorting task is evaluated using the *first object*, as we wanted to test the method in multi-modal scenarios. With this setup, the flow estimation module is more likely to produce erroneous results due to multiple objects laying on the scene. An example of this possible failure case is shown in Figure 3 top row.

The results of this experiment can be seen in Table II. We obtain a similar improvement rate as in the simulation, confirming that a guided selection of the demonstration significantly increases the probability of success.

Scene Permutations: The quantitative experiments shown to this point improve performance through improved flow and registration. The conditional approach, however, also

TABLE II: Success rates for the shape insertion task evaluated in simulation and on the real robot. Compared to the baseline FlowControl method, success rates on both tasks increase when using our improved conditional servoing method called FlowControl-C.

	Simulation	Real Robot
FlowControl	55%	40%
FlowControl-C	81%	78%
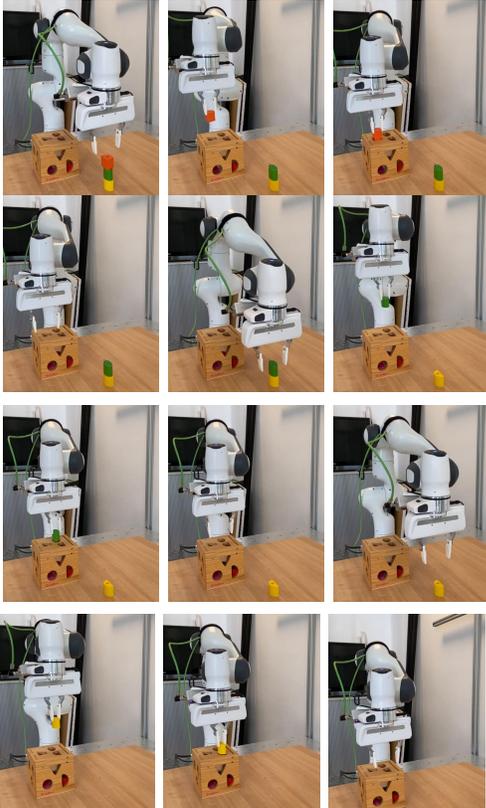
Fig. 7: Using a set of demonstrations allows tackling an initial configuration of stacked shapes. This shows that demonstrations allow addressing scenes that can be permuted geometrically, which would otherwise require an unfeasible number of complete demonstration sequences.

has other advantages: it is able to deal with geometric permutations of the scene. We show this experimentally by giving the robot a challenging scenario, in which the shapes are initially stacked in a tower. To solve this task without dividing demonstrations into subtasks one would theoretically need a demonstration for each order in which the cubes could be stacked. This is clearly infeasible for complex scenes as the number of required demonstrations to cope with all possible permutations increases rapidly. Examples from this experiment are shown in Figure 7.

Robustness: As an additional verification that our method does not lead to any unexpected failure cases, we evaluate a number of robustness tests. These introduce variations be-
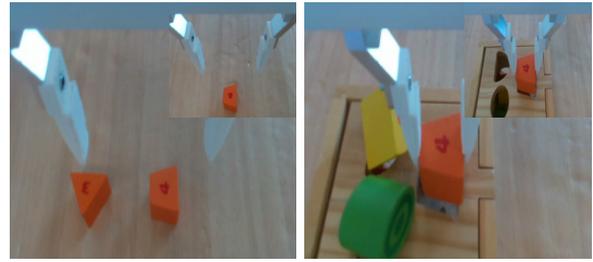
(a) Same Color Objects (b) Occlusions
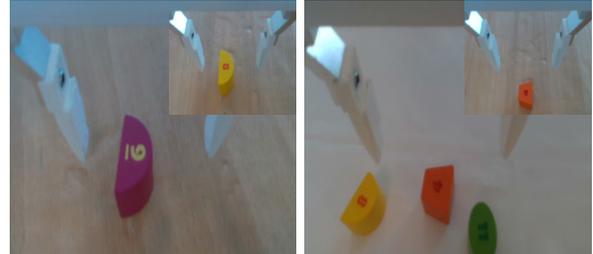
(c) Different Color Objects (d) New Backgrounds
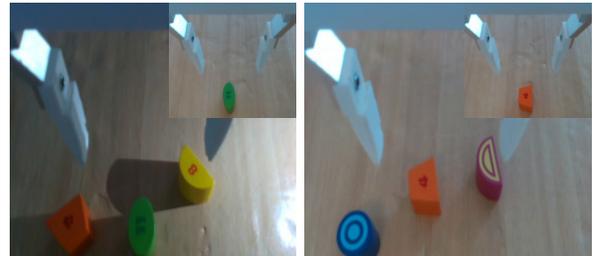
(e) Lighting (f) Unseen Objects

Fig. 8: Robustness to a number of nuisance variations results from the use of learned optical flow. Large images show live views and inset images show demonstrations. All examples from recording of successful insertions.

tween the live view and the demonstration view, for example by changing the lighting conditions or the backgrounds for the tasks. Example images from these experiments are shown in Figure 8.

V. CONCLUSIONS

We have presented conditional servoing as a concept – an outer loop around a Servoing-from-Demonstration loop. A practical implementation based on the FlowControl method has been provided. This implementation was tested both in simulation and on a real robot. In these experiments, we showed that reprojection error is a suitable scoring function for demonstration sequence selection. The method achieves high success rates on the challenging problem solving an insertion cube. In addition to these quantitative tests, we conducted qualitative tests showing that the system is able to handle scene permutations and works robustly despite the presence of nuisance variations in the scene.

REFERENCES

[1] M. Argus, L. Hermann, J. Long, and T. Brox, "Flowcontrol: Optical flow based visual servoing," in *Proceedings of the International Conference on Intelligent Robots and System (IROS)*, Las Vegas, Arizona, 2020. [Online]. Available: https://ras.papercept.net/proceedings/IROS20/634.pdf

[2] D. Kragic and H. I. Christensen, "Survey on visual servoing for manipulation," Computational Vision and Active Perception Laboratory, Tech. Rep., 2002.

[3] F. Chaumette and S. Hutchinson, "Visual servo control. i. basic approaches," *IEEE Robotics & Automation Magazine*, vol. 13, pp. 82–90, 2006.

[4] S. Hutchinson, G. D. Hager, and P. I. Corke, "A tutorial on visual servo control," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 5, pp. 651–670, Oct 1996.

[5] A. Kheddar, S. Caron, P. Gergondet, A. Comport, A. Tanguy, C. Ott, B. Henze, G. Mesesan, J. Englsberger, M. A. Roa, P.-B. Wieber, F. Chaumette, F. Spindler, G. Oriolo, L. Lanari, A. Escande, K. Chappellet, F. Kanehiro, and P. Rabate, "Humanoid robots in aircraft manufacturing," *IEEE Robotics and Automation Magazine*, vol. 26, no. 4, pp. 30–45, Dec. 2019. [Online]. Available: https://hal-lirmm.ccsd.cnrs.fr/lirmm-02303117

[6] B. Huang, M. Ye, S.-L. Lee, and G.-Z. Yang, "A vision-guided multi-robot cooperation framework for learning-by-demonstration and task reproduction," *International Conference on Intelligent Robots and Systems (IROS)*, pp. 4797–4804, 2017.

[7] F. Chaumette, S. Hutchinson, and P. Corke, "Visual Servoing," in *Handbook of Robotics, 2nd edition*, B. Siciliano and O. Khatib, Eds. Springer, 2016, pp. 841–866. [Online]. Available: https://hal.inria.fr/hal-01355384

[8] V. Lippiello, J. Cacace, A. Santamaria-Navarro, J. Andrade-Cetto, M. A. Trujillo, Y. R. Rodriguez Esteves, and A. Viguria, "Hybrid visual servoing with hierarchical task composition for aerial manipulation," *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 259–266, 2016.

[9] N. Mansard and F. Chaumette, "Tasks sequencing for visual servoing," *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566)*, vol. 1, pp. 992–997 vol.1, 2004.

[10] B. J. Nelson, N. Papanikolopoulos, and P. K. Khosla, "Visual servoing for robotic assembly," in *Visual Servoing-Real-Time Control of Robot Manipulators Based on Visual Sensory Feedback*, K. Hashimoto, Ed. River Edge, NJ: World Scientific Publishing Co. Pte. Ltd., 1993, pp. 139–164.

[11] T. Mitsuda, Y. Miyazaki, N. Maru, K. F. MacDorman, A. Nishikawa, and F. Miyazaki, "Visual servoing based on coarse optical flow," *IFAC Proceedings Volumes*, vol. 32, no. 2, pp. 539 – 544, 1999, 14th IFAC World Congress 1999. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1474667017560929

[12] W. Sepp, S. Fuchs, and G. Hirzinger, "Hierarchical featureless tracking for position-based 6-dof visual servoing," *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 4310–4315, 2006.

[13] P. Sermanet, C. Lynch, J. Hsu, and S. Levine, "Time-contrastive networks: Self-supervised learning from multi-view observation," *CoRR*, vol. abs/1704.06888, 2017. [Online]. Available: http://arxiv.org/abs/1704.06888

[14] O. Kroemer, S. Niekum, and G. Konidaris, "A review of robot learning for manipulation: Challenges, representations, and algorithms," *J. Mach. Learn. Res.*, vol. 22, pp. 30:1–30:82, 2021.

[15] P. Pastor, M. Kalakrishnan, L. Righetti, and S. Schaal, "Towards associative skill memories," *2012 12th IEEE-RAS International Conference on Humanoid Robots (Humanoids 2012)*, pp. 309–315, 2012.

[16] O. Kroemer, C. Daniel, G. Neumann, H. V. Hoof, and J. Peters, "Towards learning hierarchical skills for multi-phase manipulation tasks," *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1503–1510, 2015.

[17] T. Shu, C. Xiong, and R. Socher, "Hierarchical and interpretable skill acquisition in multi-task reinforcement learning," *ArXiv*, vol. abs/1712.07294, 2018.

[18] O. Mees, M. Merklinger, G. Kalweit, and W. Burgard, "Adversarial skill networks: Unsupervised robot skill learning from video," *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4188–4194, 2020.

[19] O. Chum, M. Perdoch, and J. Matas, "Geometric min-hashing: Finding a (thick) needle in a haystack," in *CVPR*, 2009.

[20] H. Jégou, M. Douze, C. Schmid, and P. Pérez, "Aggregating local descriptors into a compact image representation," *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 3304–3311, 2010.

[21] H. Jun, B. Ko, Y. Kim, I. Kim, and J. Kim, "Combination of multiple global descriptors for image retrieval," *ArXiv*, vol. abs/1903.10663, 2019.

[22] R. Arandjelović, P. Gronát, A. Torii, T. Pajdla, and J. Sivic, "Netvlad: Cnn architecture for weakly supervised place recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, pp. 1437–1451, 2018.

[23] T. Weyand, A. Araújo, B. Cao, and J. Sim, "Google landmarks dataset v2 – a large-scale benchmark for instance-level recognition and retrieval," *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2572–2581, 2020.

[24] C. Rasmussen and G. Hager, "Robot navigation using image sequences," in *AAAI/IAAI, Vol. 2*, 1996.

[25] J. Delfin, H. Becerra, and G. Arechavaleta, "Visual path following using a sequence of target images and smooth robot velocities for humanoid navigation," *2014 IEEE-RAS International Conference on Humanoid Robots*, pp. 354–359, 2014.

[26] H. Lim, S. N. Sinha, M. F. Cohen, and M. Uyttendaele, "Real-time image-based 6-dof localization in large-scale environments," *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1043–1050, 2012.

[27] B. Williams, G. S. W. Klein, and I. Reid, "Automatic relocalization and loop closing for real-time monocular slam," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, pp. 1699–1712, 2011.

[28] J. L. Schönberger and J.-M. Frahm, "Structure-from-motion revisited," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[29] J.-P. Mercier, M. Garon, P. Giguère, and J.-F. Lalonde, "Deep template-based object instance detection," *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1506–1515, 2021.

[30] M. Rad and V. Lepetit, "Bb8: A scalable, accurate, robust to partial occlusion method for predicting the 3d poses of challenging objects without using depth," *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 3848–3856, 2017.

[31] J. Y. Jason, A. W. Harley, and K. G. Derpanis, "Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness," in *European Conference on Computer Vision*. Springer, 2016, pp. 3–10.

[32] Z. Ren, J. Yan, B. Ni, B. Liu, X. Yang, and H. Zha, "Unsupervised deep learning for optical flow estimation," in *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.

[33] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[34] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *J. Mach. Learn. Res.*, vol. 17, pp. 39:1–39:40, 2015.

[35] J. Ibarz, J. Tan, C. Finn, M. Kalakrishnan, P. Pastor, and S. Levine, "How to train your robot with deep reinforcement learning: lessons we have learned," *The International Journal of Robotics Research*, vol. 40, pp. 698 – 721, 2021.

[36] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," http://pybullet.org, 2016–2021.