

# Bimanual rope manipulation skill synthesis through context dependent correction policy learning from human demonstration

Baturhan Akbulut<sup>1,\*</sup>, Tuba Girgin<sup>1,\*</sup>, Arash Mehrabi<sup>2</sup>, Minoru Asada<sup>3,4</sup>, Emre Ugur<sup>1</sup>, Erhan Oztop<sup>2,3</sup>

**Abstract**—Learning from demonstration (LfD) with behavior cloning is attractive for its simplicity; however, compounding errors in long and complex skills can be a hindrance. Considering a target skill as a sequence of motor primitives is helpful in this respect. Then the requirement that a motor primitive ends in a state that allows the successful execution of the subsequent primitive must be met. In this study, we focus on this problem by proposing to learn an explicit correction policy when the expected transition state between primitives is not achieved. The correction policy is learned via behavior cloning by the use of Conditional Neural Motor Primitives (CNMPs) that can generate correction trajectories in a context-dependent way. The advantage of the proposed system over learning the complete task as a single action is shown with a table-top setup in simulation, where an object has to be pushed through a corridor in two steps. Then, the applicability of the proposed method to bi-manual knotting in the real world is shown by equipping an upper-body humanoid robot with the skill of making knots over a bar in 3D space.

## I. INTRODUCTION

Learning from demonstration (LfD) [1] is an effective robot learning framework, where a robot is taught a desired skill through human demonstration. The demonstration may be kinaesthetic, visual, or through teleoperation [2]. The latter places the human operator in the control loop. Thus human motor learning capacity enables teaching more complex tasks to the robot [3]. The learning can be based on policy cloning [4], i.e., learning a state-to-action mapping, or a simple motor-tape approach [5], i.e., playing back the stored motor commands in an open-loop fashion. For complex and long tasks, both methods suffer from the so-called covariate shift problem [6]. When the system enters an off-learned distribution state, it can not recover. Even worse, it can diverge to states that are even further away. Assume a robot executing a rather long trajectory in order to make a knot using its two grippers. Even small mistakes that are made during the initial parts of the action, such as holding the rope from a slightly different position or making slight errors in the movement trajectory that cause the gripper to miss the rope, would result in failure of task execution. A remedy is to consider the actions as the composition of simpler action primitives less prone to covariate shift, allowing corrective actions to be inserted between primitives to ensure the



Fig. 1: Bi-manual knotting with Torobo in the real world.

successful execution of the subsequent primitive. Corrective action policies can be learned by the robot itself or again taught by a human demonstrator. In this study, we explore the latter approach, first in simulation, then show that it can be applied to real hardware by synthesizing a complex rope manipulation skill on Torobo (Tokyo Robotic Inc.), an upper-body humanoid robot (Fig 1).

In summary, we propose a method that enables the robot to recover from failure states that might be encountered during critical points in complex and long action executions. To do this, the action is segmented based on a set of identified critical points. During these critical points, the human operator demonstrates the movements that need to be taken to make the task execution succeed, and the demonstration is learned by our system as a function of robot-environment context. The corrective actions that might require very different trajectories are encoded in a single corrective movement primitive. Therefore, when encountered with a failure context, a corrective action is automatically generated to ensure successful task execution.

Overall, the main contributions of our paper are (1) reducing the human effort for data generation for LfD by focusing the learning on correction policies and (2) showing that the approach is general by realizing it on real hardware for a non-trivial task. In particular, with this method, we obtain a dexterous rope manipulation skill in 3D space, superseding earlier rope knotting skills with real robots.

## II. RELATED WORK

There has been a significant amount of research on robotic manipulation of deformable objects. Nair et al. [7] combined self-supervised learning with imitation for rope manipulation, where an inverse dynamics model predicts the action given the images of the current and next states, and human demonstration is used for high-level planning. The

\* Equal contribution, alphabetical order. <sup>1</sup> Baturhan Akbulut, Tuba Girgin and Emre Ugur are with Department of Computer Engineering, Bogazici University, Istanbul, Turkey. <sup>2</sup> Arash Mehrabi and Erhan Oztop are with Department of Computer Engineering, Ozyegin University, Istanbul, Turkey. <sup>3</sup> Erhan Oztop and Minoru Asada are with Symbiotic Intelligent Systems Research Center, Institute for Open and Transdisciplinary, Research Initiatives, Osaka University, Japan. <sup>4</sup> Minoru Asada is also affiliated with International Professional University of Technology in Osaka, Osaka, Japan.

inverse dynamics model works well for small deformations. However, it lacks the ability to construct a trajectory to reach to goal states. Wang et al. [8] extended [7] by replacing human demonstration with a Casual InfoGAN (CIGAN) model. However, high success rates could not be reached in generating plausible plans.

Yan et al. [9] developed a visual model-based learning framework that jointly learns the underlying visual latent representations and the dynamics models for deformable objects. The latent dynamics model was learned in a simulation environment, then transferred to a real PR2 robot. In another work, Yan et al. [10] proposed a model that learns the dynamics of linear deformable objects. In our approach, we do not learn the dynamics of the rope but directly build a correction policy which is applicable to the rope configurations faced during corrections. Chang et al. [6] proposed the network MILO where offline RL achieves learning without needing online demonstrations. Reichlin et al. [11] suggested a network to combine it with the offline RL setup to overcome the distribution-shift problem. The proposed recovery network encodes observations and creates a latent representation. The agent corrects its motion using the encoded information and approximating the training data density in latent space. However, the recovery network affects the predicted motion in each step, while our proposed method only needs correction in critical points, simplifying its logic.

Multiple other approaches have been tried for deformable object manipulation. Notably, [12] realized a model-free visual reinforcement learning algorithm for fabric and rope manipulation and a max-margin learning algorithm for rope untying. Sundaresan et al. [13] used interpretable deep visual representations for rope. They learned point-pair correspondences between initial and goal rope configurations in simulation from synthetic depth data, enabling them to plan actions to reach the goal configuration. Lee et al. [14] proposed a unified optimization method for an LfD method, which transfers a demonstration for the training scene to a test scene, and Seita et al. [15] used Transporter Networks [16] to obtain vision-based manipulation of deformable objects. Seita et al. [17] used deep imitation learning for fabric smoothing. They trained policies in simulation using dataset aggregation (DAGger) with an algorithmic supervisor, given images of a fabric sample as input. The difference between our algorithm and DAGger is that our algorithm only asks the expert for the correct action when it reaches a critical state, while DAGger needs continuous access to the supervisor's policy.

Conditional Neural Movement Primitives (CNMPs) [18] is an LfD method that can encode multiple modes of operations from multiple demonstrations and generate sensorimotor trajectories when conditioned at arbitrary time points. Compared to other LfD methods such as Dynamic Movement Primitives [19], [20], [21], Probabilistic Movement Primitives [22], [23], Compliant Parametric DMPs [24], [25] or statistical models [26], CNMP has the advantage of providing a single framework that can learn distributions of com-




Fig. 2: Illustration of the proposed concept with a single critical point. Demonstrated trajectory (dashed gray curve) passes through a critical region C, which depends on the task. For the robot execution, a correction movement is inserted to counteract the possible case of execution missing C, thereby splitting the execution into three motor primitives (MP): pre-MP, cor-MP, and post-MP. post-MP must be learned context-dependent to recover a wide range of failure cases. This is achieved by obtaining multiple correction demonstrations through teleoperation.

plex movement trajectories conditioned on high-dimensional task parameters by representing non-linear relations between the task parameters and movement trajectories from a few demonstrations. CNMPs have been shown effective in a wide range of learning problems such as learning manipulation [18] and locomotion skills [27], imitation learning [28], and adapting learned skills beyond demonstrations [29], [30]. However, correction policy learning with CNMPs has not been addressed before.

Overall, unlike previous works, our method learns to recover from failure states by forming an explicit correction policy based on expert correction demonstrations.

### III. METHOD

Complex tasks, e.g., rope manipulation, are often linked through critical points significant for successfully executing the task, where critical points are considered to be certain parts of the execution trajectory that may lead to failure. Furthermore, long and complex tasks are fragile to the accumulation of errors. Therefore, we propose to segment the trajectory by using the critical points as segment boundaries and insert corrective motor primitives to ensure the subsequent segment preconditions are met. Fig. 2 illustrates a trajectory with two such segments encoded as Pre-Movement (pre-MP) and Post-Movement (post-MP) primitives, and a Corrective Movement primitive (cor-MP) inserted, which enables the robot to recover from different failure situations.

In this work, we assume the critical points are known and focus on cor-MP learning through human demonstration. Since the correction would depend on the robot and environment configuration, i.e., the context reached after the execution of the preceding segment, we learn a context-dependent correction policy. A context is a vector extracted from sensors available to the robot that captures the task-relevant environment-robot relations. By imposing such a

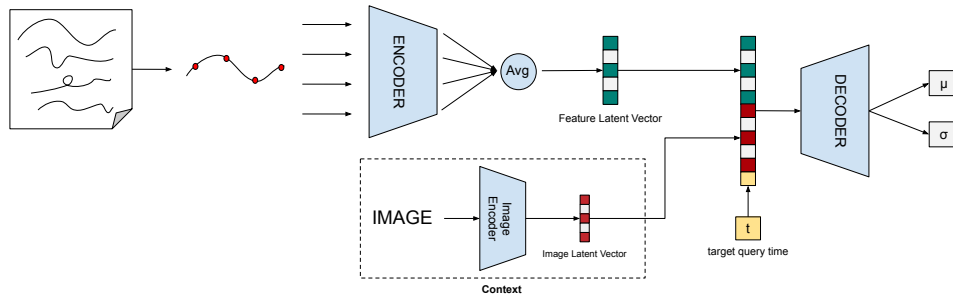


Fig. 3: Conditional Neural Movement Primitive architecture (CNMP) with a context input is shown. In the illustration, the context is provided by an autoencoder. See Section III for more details.

task decomposition, rather than learning a monolithic task policy from demonstration, we aim to improve performance and robustness while achieving a general correction policy for the critical point under consideration. To achieve a similar performance from a monolithic behavior cloning system, one would need to ask for multiple demonstrations for the full action; whereas in our proposal, the corrections for the critical points have to be demonstrated multiple times while the main task execution can be demonstrated a few times (typically once).

#### A. General Framework

Our model is built on top of Conditional Neural Movement Primitives (CNMPs) architecture [18]. CNMPs is a deep neural architecture capable of representing sensorimotor data conditioned on the given set of observation points as a latent vector, which then can be used to unfold the trajectory based on queried time points. Thus, when coupled with a feedback controller, it can act as a drive for the movements of a robot. In our work, for critical segment learning, we feed CNMPs with the context information so that the correction is conditioned on the context. A CNMPs is composed of two networks: an encoder and a decoder, as shown in Fig. 3. The encoder part learns to generate a latent representation (i.e., latent vector) from the samples taken from an input trajectory. By taking the mean of the representation for each sample, a general representation is obtained that is sufficient to reconstruct the trajectories learned. Since training and querying are done with a context input appended to the latent representation (Fig. 3), the generated trajectories become context-dependent. Thus, for the corrective motor primitives (cor-MP), CNMPs learn to generate different motor patterns based on the context. Priming the CNMPs with a context not only enforces the selection of the right correction policy but also allows generalization to unseen contexts, thereby bringing robustness to overall task execution.

In this study, we test our method on two tasks realized in simulated and real robots. Both tasks involve a single critical region. Therefore, we can unambiguously refer to pre-critical motor primitive (pre-MP), corrective motor primitive (cor-MP), and post-critical motor primitive (post-MP) in our descriptions (see Fig 2). Further details are explained comprehensively in Section IV. For the simulation task, human demonstration is emulated through a deterministic




Fig. 4: Simulation environment prepared in Gazebo.




Fig. 5: A typical action execution to accomplish the goal.

algorithm. For the hardware task, human demonstration is obtained through a teleoperation system using two haptic devices (GeoMagic Inc.).

#### B. Learning the Context

The context vector provides information about the current robot state and environment relation. As such, it can be a manually designed feature extractor, or it can be learned in an end-to-end fashion through an auto-encoder (AE). To test the latter approach, for the simulation experiment, we designed an AE with 2D convolutional layers of kernel sizes 3, strides [2, 1, 2, 1, 2], and filter sizes [32, 32, 64, 64, 64] in the encoder block. The output of the encoder is flattened and fed into a fully connected layer to get the latent representation of the image. The decoder mimics the encoder's operations in reverse and is used only in training the AE. In the simulation experiments, the encoder receives an image, and the resulting latent vector is used as the context. The AE is trained on a data set consisting of 400 images. 360 of them are used for training and the remaining 40 for testing. The images are taken after the first push and correction actions (described in Section IV-A). The AE is trained for 70 epochs using Adam optimizer [31] with a learning rate of  $3e-4$ . The size of the latent space representation of images was set to 8 for the experiments reported in the current paper.

## IV. EXPERIMENTS AND RESULTS

### A. Push object through corridor with AE-based context

We use the simulated upper body humanoid robot Torobo (Tokyo Robotics Inc.) in Gazebo simulation environment (Fig. 4) in order to teach the robot to push a target object through a corridor in two steps: First, the robot needs to align the object with the corridor, and then push it through, which is a simple task the robot can perform. However, to test our method, we perturb the position of the corridor stochastically in a 6 cm range, thereby deliberately creating a critical region for the task execution that would require a corrective step. Following the logic of the proposed model, the robot asks demonstrations for those critical points that are encountered after the first push. As stated in Section III, the human demonstration is emulated by an algorithmic logic that guides the robot to do the 'right' corrective action. While the correction is being carried out, the data is collected as the corrective action demonstration data and used for training cor-MP neural network.

The characteristics of the corrective motor primitives (cor-MP) differ according to the position of the corridor. When the corridor moves away from the robot, the robot needs to push the object to align it with the corridor. On the other hand, if the corridor comes closer, the robot needs to pull the object so as to align it with the corridor. A typical task execution with a pull-type correction is shown in Fig. 5.

Two hundred trajectories are collected together with environment information in total to create a data set. Each data unit consists of the Cartesian position of the gripper, the corridor location, and the target object. In addition, for each trajectory, two images are captured from the camera of the Torobo Robot. The first image is captured after the execution of the pre-critical movement primitive (pre-MP). The second image is captured after the execution of cor-MP. These images are used to obtain autoencoder-based context for the experiments that involve learned contexts (Section IV.B.2). Typical images captured from the camera of the robot head camera are shown in Fig. 6.

As mentioned in Section III, the task execution is composed of the executions of three sequential motor primitives. Thus our system is required to produce three trajectories, namely pre-MP network, cor-MP network, post-MP network to successfully complete the task. The trajectories are encoded in task space in this experiment. Particularly, the context is used by the post-MP network to make the final push since after cor-MP network correction, the location of the object would be, in general, changed. Two simulation experiments are conducted, one with 'perfect context' and one with the 'AE-based' context. For the perfect context, the corridor location and the object location after the correction are used. To be concrete, cor-MP network uses the position of the corridor on the posterior-anterior axis as the context,




Fig. 6: Camera images of the scene, showing the configurations before (left) and after (right) the correction action.

and post-MP network uses the position of the target object on the same axis. The context values are normalized before being used in cor-MP network and post-MP network.

For the AE-based context-based experiments, the captured images from the camera are fed into an autoencoder (AE) to extract the information related to the target object and the corridor positions. Then, the encoded images (i.e., latent vectors) are given as contexts to the cor-MP network and post-MP network, respectively. Table I shows the learning parameters used for training. To reiterate, the CNMP networks are trained by using (normalized) time and x, y, and z positions of the gripper collected during demonstrations for both perfect and AE-based context experiments. The CNMP weights are trained by using Adam optimizer[31]. During the training, 100 of 200 trajectories are used for training our model for each context type. The remaining ones are used for validation. Each model received 10M weight updates.

### B. Results of Push through corridor task Experiment (Simulation) Experiments

The trained pre-MP, cor-MP, and post-MP networks are integrated into the simulation environment to assess the performance of the proposed system in driving the Torobo Robot for successfully pushing the target object through the corridor. For execution, the predicted pre-MP Cartesian trajectory is mapped to a joint trajectory through inverse kinematics, which is fed to the robot system as desired joint angles at a nominal rate. The robot controller then tracks the obtained joint trajectory and ends up in the critical region. The robot observes the scene, obtains the latent vector from the trained auto-encoder (in the case of AE-based context) or logs the position of the corridor (in the perfect-context case), and uses it as the context for the cor-MP network. By taking the last gripper position as the initial observation for cor-MP network, cor-MP trajectory is predicted. Since the cor-MP network has learned corrections from demonstrations, the generated trajectories are expected to correspond to push, pull, or no action (as these have been strategies used in the demonstrations). After executing cor-MP, the post-MP trajectory is predicted according to the observed context and the position of the end effector. After the execution, the trial is counted as a success if the target center of the object passes the corridor's midpoint without touching the corridor's entrance. To assess the performance of the system, 25 trials (test data set) with unseen contexts are executed with the predicted pre-MP, cor-MP, and post-MP for both perfect context and AE-based context experiments.

#### 1) Results of CNMPs with Perfect Context (PC-CNMPs):

In order to successfully complete the task, the robot was required to make 9 push and 10 pull correction actions in the

TABLE I: Simulation Training Parameters

Network Type	Learning Rates	Encoder Layer Count	Decoder Layer Count	Layer Size
All CNMPs	0.0001	3	4	128

TABLE II: The action decision and task success of our proposed architecture with perfect context (PC) and autoencoder features (AE) are provided. For comparison, the results obtained from monolithic networks (Mono) is also provided.

Expected Correction	Action Choice			Decision Match			Task Success		
	PC	AE	Mono	PC	AE	Mono	PC	AE	Mono
Push	9 Push	8 Push, 1 No Correction	9 Push	100%	89%	100%	100%	100%	44%
No Correction	6 No Correction	6 No Correction	6 No Correction	100%	100%	100%	83%	100%	67%
Pull	10 Pull	9 Pull, 1 No Correction	10 Pull	100%	90%	100%	100%	100%	90%

TABLE III: Prediction Errors (cm) of MP trajectories.

Network Type	cor-MP		post-MP	
	mean	max	mean	max
PC	0.80	1.45	1.69	3.19
AE	0.80	0.99	2.64	4.83
B	1.30	1.92	5.75	7.24

25 cases tested. The remaining 6 cases required no correction (as the object was aligned with the corridor to a tolerable level). The PC-CNMPs results show that the network was able to predict an appropriate correction trajectory when faced with previously unseen contexts. This shows that the cor-MP network can learn to generate correct decisions and correct movement amounts. As a result of these corrective actions, the robot was able to move the object to a position to align it with the corridor. The only task failure was detected during a push type of correction generated by the cor-MP network, and the failure was due to the object slightly contacting the corridor at the entrance.

The mean and maximum error values for cor-MP and post-MP are shown in Table III. The error for each motion is calculated by taking the average of the Euclidean distances between the predicted action points and the ground truths. The mean error is 0.80 cm for cor-MP and 1.69 cm for post-MP. post-MP error is higher than the previous one because the desired motion for post-MP is more complex than the previous one. These errors cause small noise in the end position of the target object after each motion. Because these errors are small enough, 24 trials out of 25 yielded successful task executions. Predictions, robot executions, and the ground-truths for two different action types for cor-MP are compared in Fig. 7. As seen in the Figure, the predictions overlap the ground truths. Therefore, the simulated robot was able to execute the predicted motions accordingly.

2) *Results of CNMPs with AE (AE-CNMPs)*: As shown in Table II, when AE features were used as context by our system, in 24 trials out of 25, the robot was observed to accomplish the task by pushing the target object in front of the narrow corridor, correcting its position, and pushing it through the narrow corridor to the goal position. When the execution of each predicted motion is examined, the cor-MP network was mostly able to choose the right action type given the context. When we examine the predictions that perform differently from the test data set used to assess the performance of the predictions it is observed that the predicted trajectories by the cor-MP network slightly shifted in comparison to demonstrated trajectories. Since the captured image at boundary cases between push and no-correction, as well as pull and no-correction are similar, the AE representation for those cases is slightly smeared. Therefore, the cor-MP network may choose to generate

different action types around the decision boundaries. Still, because the correction amount required around the decision boundaries is small compared to the width of the corridor, the different decisions at these boundaries have not affected the success of task execution. As a result, the success of the system with AE context has been found to be 100%.

The mean cor-MP error is similar to that of perfect context version of the system. The mean error resulting from post-MP network prediction is 2.64 cm, which is higher than the one found for the perfect context experiments. This might be due to the fact that the AE-based context version of the system made different choices when doing the correction, leaving the target object in a different location experienced during the demonstrations, which is acted upon by post-MP.

3) *Comparison with Monolithic Baseline CNMP*: To assess the benefits of using segment action representation with explicit corrective motor primitives for the critical regions, we implement a monolithic baseline CNMP network. Given a fixed context, the full trajectory can be learned with the Vanilla CNMPs. However, since the position of the corridor changes trial-to-trial, without knowing any information about the context, the learning system cannot decide which trajectory to generate. Therefore, to be fair, we provide the corridor location to the Vanilla CNMPs and ask it to learn the demonstrated full trajectories (which include the proper corrections). The number of parameters of Vanilla CNMPs is adjusted to be the sum of the number of parameters of the pre-MP, cor-MP, and post-MP networks. The comparison of the Vanilla CNMPs network generated actions to those of PC-CNMPs and AE-CNMPs are shown in Table II with the Mono label. As shown, the Vanilla CNMPs network can generate the right correction action; however, when executed, only 17 trials out of 25 succeed in completing the task. This is due to the errors stemming from learning a longer trajectory, albeit the Vanilla CNMPs network uses an equal number of parameters as of the sum of the PC-CNMPs and AE-CNMPs networks. The mean and max error values for the Vanilla CNMPs network are also observed to be higher than the PC-CNMPs and AE-CNMPs networks. The results show that the proposed method outperforms the baseline with the total number of learnable parameters equalized.




Fig. 7: Comparison of ground-truth, prediction, and robot execution for different cor-MPs.




Fig. 8: Knot over bar execution illustrated with Torobo robot.

### C. Knotting with the real robot

For the hardware implementation, we chose bi-manual rope knotting over a bar, as it demonstrates the applicability of the proposed method for a range of complex tasks. With the implemented teleoperation setup (Section III), the users can control the pose of the individual grippers attached to the arms of the robot Torobo robot. With this, we obtain basic knotting demonstration on the robot by an expert demonstrator. Learning from demonstration with the Vanilla CNMPs yields knotting behavior; however small variations such as how the gripper holds the rope ends, the length of the rope between the grippers, and invisible twists in the rope cause the robot to miss the critical ‘pass through the loop and pull the free end’ (see Fig. 8) part of the knotting action. So, analogous to the simulation experiments, at the point just before the gripper would enclose (see Fig. 8-C), we collected multiple corrective demonstrations from the user to synthesize a corrective motor primitive to be inserted at this critical point with the aim of increasing the success of knotting. The actions before and after this point ,i.e., (pre-MP, post-MP) can be learned easily by individual CNMPs. For the corrective part, we have collected 28 corrective action demonstrations, which involve posterior-to-anterior, anterior-to-posterior, lateral corrections and their combinations. A few demonstrations without any corrections were also included. All the demonstrations were learned with a single corrective CNMP (cor-MP network) as in the simulations giving us a corrective motor primitive (cor-MP). The correction was given as relative Cartesian positions with respect to the gripper position. Thus the correction could be applied in any place in the workspace of the robot. For the Cartesian to joint space transformation we used inverse kinematics with fixed gripper orientation (as attained after the execution of pre-MP). As the context information, the cor-MP network received the relative position of the rope end with respect to the gripper mid-point. The rope end position is extracted with a straightforward color-segmentation and clustering pipeline on the RGB-D image captured by the camera (Intel SR300) mounted on the head of Torobo. The cor-MP network could learn the demonstrations well, thus near exact context (i.e., rope-gripper configuration) at the critical point successfully generated corrective movements. More critical test was to see the performance on unseen contexts. For that, we asked the robot to execute pre-MP and perform the correction by using the trajectory generated by the cor-MP network, which was followed by the execution of the post-MP, that included the enclose command of the gripper. Thus if the correction was successful, the knot operation would be completed successfully. In those trials, if the reached critical points




Fig. 9: Real robot knotting performance. Knot correction demonstrations (gray) with predictions generated (green) and the corresponding contexts (left panel).

did not require a correction, we moved the rope outside the fingers to see whether the required correction can be made. The experimental results are provided in Fig. 9. The gripper-robot configurations used in training are shown with gray, and the test configurations with successful correction and task execution are shown with green. In addition, some indicative failures are also shown (yellow and red). Overall, we can see that the correction skill was acquired by the robot and generalization to unseen contexts could be handled. Most of the failures were due to kinematic constraints (i.e., out of reach) and others are near misses (the rope is contacted by the grippers but the rope was not straddled).

### V. DISCUSSION AND FUTURE WORK

In this paper, we have proposed a novel LfD approach which can be applied to deformable object manipulation. By segmenting a given task into motor primitives and inserting explicit correction policies between the segments, we show that the potential failures that would be faced with a monolithic approach are eliminated, and accuracy and generalization of the task execution can be increased, while at the same human effort for demonstration is reduced.

The success of the proposed method relies on the knowledge on which regions are likely to need correction, i.e. critical for task success. Finding these regions automatically without excessive number of full task demonstrations would make the proposed method much more potent. Also, it would be interesting to see how other LfD methods would compare with CNMPs in learning context based correction policies.

### ACKNOWLEDGMENT

This research was supported by the International Joint Research Promotion Program, Osaka University under the project ‘‘Developmentally and biologically realistic modeling of perspective invariant action understanding’’, the Japan Society for the Promotion of Science, Grant-in-Aid for Scientific Research (project no 22H03670), the project JPNP16007 commissioned by the New Energy and Industrial Technology Development Organization (NEDO), and by the BAGEP Award of the Science Academy.

## REFERENCES

- [1] S. Schaal, "Learning from demonstration," *Advances in neural information processing systems*, vol. 9, 1996.
- [2] N. Amirshirzad, A. Kumru, and E. Oztop, "Human adaptation to human–robot shared control," *IEEE Transactions on Human-Machine Systems*, vol. 49, no. 2, pp. 126–136, 2019.
- [3] J. Babic, J. G. Hale, and E. Oztop, "Human sensorimotor learning for humanoid robot skill synthesis," *Adaptive Behavior*, vol. 19, no. 4, pp. 250–263, 2011. [Online]. Available: [Go to ISI://000294157400002](https://doi.org/10.1002/294157400002)
- [4] D. A. Pomerleau, "Alvinn: An autonomous land vehicle in a neural network," in *Advances in Neural Information Processing Systems*, D. Touretzky, Ed., vol. 1. Morgan-Kaufmann, 1988.
- [5] C. Atkeson, J. Hale, F. Pollick, M. Riley, S. Kotosaka, S. Schaul, T. Shibata, G. Tevatia, A. Ude, S. Vijayakumar, E. Kawato, and M. Kawato, "Using humanoid robots to study human behavior," *IEEE Intelligent Systems and their Applications*, vol. 15, no. 4, pp. 46–56, 2000.
- [6] J. Chang, M. Uehara, D. Sreenivas, R. Kidambi, and W. Sun, "Mitigating covariate shift in imitation learning via offline data with partial coverage," in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34. Curran Associates, Inc., 2021, pp. 965–979.
- [7] A. Nair, D. Chen, P. Agrawal, P. Isola, P. Abbeel, J. Malik, and S. Levine, "Combining self-supervised learning and imitation for vision-based rope manipulation," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 2146–2153.
- [8] A. Wang, T. Kurutach, K. Liu, P. Abbeel, and A. Tamar, "Learning robotic manipulation through visual planning and acting," 2019. [Online]. Available: <https://arxiv.org/abs/1905.04411>
- [9] W. Yan, A. Vangipuram, P. Abbeel, and L. Pinto, "Learning predictive representations for deformable objects using contrastive estimation," 2020. [Online]. Available: <https://arxiv.org/abs/2003.05436>
- [10] M. Yan, Y. Zhu, N. Jin, and J. Bohg, "Self-supervised learning of state estimation for manipulating deformable linear objects," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2372–2379, 2020.
- [11] A. Reichlin, G. L. Marchetti, H. Yin, A. Ghadirzadeh, and D. Kragic, "Back to the manifold: Recovering from out-of-distribution states," 2022. [Online]. Available: <https://arxiv.org/abs/2207.08673>
- [12] W. H. Lui and A. Saxena, "Tangled: Learning to untangle ropes with rgb-d perception," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2013, pp. 837–844.
- [13] P. Sundaresan, J. Grannen, B. Thananjeyan, A. Balakrishna, M. Laskey, K. Stone, J. E. Gonzalez, and K. Goldberg, "Learning rope manipulation policies using dense object descriptors trained on synthetic depth data," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 9411–9418.
- [14] A. X. Lee, S. H. Huang, D. Hadfield-Menell, E. Tzeng, and P. Abbeel, "Unifying scene registration and trajectory optimization for learning from demonstrations with application to manipulation of deformable objects," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014, pp. 4402–4407.
- [15] D. Seita, P. Florence, J. Tompson, E. Coumans, V. Sindhwani, K. Goldberg, and A. Zeng, "Learning to rearrange deformable cables, fabrics, and bags with goal-conditioned transporter networks," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 4568–4575.
- [16] A. Zeng, P. Florence, J. Tompson, S. Welker, J. Chien, M. Attarian, T. Armstrong, I. Krasin, D. Duong, A. Wahid, V. Sindhwani, and J. Lee, "Transporter networks: Rearranging the visual world for robotic manipulation," 2020.
- [17] D. Seita, A. Ganapathi, R. Hoque, M. Hwang, E. Cen, A. K. Tanwani, A. Balakrishna, B. Thananjeyan, J. Ichnowski, N. Jamali, K. Yamane, S. Iba, J. Canny, and K. Goldberg, "Deep Imitation Learning of Sequential Fabric Smoothing From an Algorithmic Supervisor," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [18] M. Y. Seker, M. Imre, J. H. Piater, and E. Ugur, "Conditional neural movement primitives," in *Robotics: Science and Systems*, vol. 10, 2019.
- [19] J. M. Rožanec and B. Nemeč, "Neural dynamic movement primitives—a survey," *arXiv preprint arXiv:2208.01903*, 2022.
- [20] P. Pastor, H. Hoffmann, T. Asfour, and S. Schaal, "Learning and generalization of motor skills by learning from demonstration," in *2009 IEEE International Conference on Robotics and Automation*. IEEE, 2009, pp. 763–768.
- [21] A. Ude, A. Gams, T. Asfour, and J. Morimoto, "Task-specific generalization of discrete and periodic dynamic movement primitives," *IEEE Transactions on Robotics*, vol. 26, no. 5, pp. 800–815, 2010.
- [22] A. Paraschos, C. Daniel, J. R. Peters, and G. Neumann, "Probabilistic movement primitives," *Advances in neural information processing systems*, vol. 26, 2013.
- [23] A. Paraschos, E. Rueckert, J. Peters, and G. Neumann, "Probabilistic movement primitives under unknown system dynamics," *Advanced Robotics*, vol. 32, no. 6, pp. 297–310, 2018.
- [24] E. Ugur and H. Girgin, "Compliant parametric dynamic movement primitives," *Robotica*, vol. 38, no. 3, pp. 457–474, 2020.
- [25] B. Nemeč, N. Likar, A. Gams, and A. Ude, "Human robot cooperation with compliance adaptation along the motion trajectory," *Autonomous robots*, vol. 42, pp. 1023–1035, 2018.
- [26] S. Calinon, "A tutorial on task-parameterized movement learning and retrieval," *Intelligent service robotics*, vol. 9, pp. 1–29, 2016.
- [27] Y. Yildirim and E. Ugur, "Learning social navigation from demonstrations with conditional neural processes," *Interaction Studies*, 2022.
- [28] M. Y. Seker, A. Ahmetoglu, Y. Nagai, M. Asada, E. Oztop, and E. Ugur, "Imitation and mirror systems in robots through deep modality blending networks," *Neural Networks*, vol. 146, pp. 22–35, 2022.
- [29] M. Akbulut, E. Oztop, M. Y. Seker, X. Hh, A. Tekden, and E. Ugur, "Acnmp: Skill transfer and task extrapolation through learning from demonstration and reinforcement learning via representation sharing," in *Conference on Robot Learning*. PMLR, 2021, pp. 1896–1907.
- [30] M. T. Akbulut, U. Bozdogan, A. Tekden, and E. Ugur, "Reward conditioned neural movement primitives for population-based variational policy optimization," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 10 808–10 814.
- [31] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.