

# Learning Fabric Manipulation in the Real World with Human Videos

Robert Lee<sup>1</sup>, Jad Abou-Chakra<sup>1</sup>, Fangyi Zhang<sup>1</sup>, Peter Corke<sup>1</sup>

**Abstract**—Fabric manipulation is a long-standing challenge in robotics due to the enormous state space and complex dynamics. Learning approaches stand out as promising for this domain as they allow us to learn behaviours directly from data. Most prior methods however rely heavily on simulation, which is still limited by the large sim-to-real gap of deformable objects or rely on large datasets. A promising alternative is to learn fabric manipulation directly from watching humans perform the task. In this work, we explore how demonstrations for fabric manipulation tasks can be collected directly by humans, providing an extremely natural and fast data collection pipeline. Then, using only a handful of such demonstrations, we show how a pick-and-place policy can be learned and deployed on a real robot, without any robot data collection at all. We demonstrate our approach on a fabric smoothing and folding task, showing that our policy can reliably reach folded states from crumpled initial configurations. Code, video and data are available on the project website: <https://sites.google.com/view/foldingbyhand>

**Index Terms**—Manipulation, Visual Learning, Perception for Grasping and Manipulation

## I. INTRODUCTION

Fabric manipulation involves complex dynamics and a high-dimensional state space, making it a challenging task for traditional robotic control methods. This complexity makes learning-based approaches more suitable. However, these often demand large datasets, human demonstrations, or simulations—each with their own drawbacks. Simulations, for instance, struggle to capture the real-world complexities of deformable objects. To address these challenges, we propose a learning approach that leverages a small set of human demonstrations to teach a robot how to manipulate fabric in the real world, thereby eliminating the sim-to-real gap and reducing data collection time.

Rather than relying on robot-generated data or extensive human annotation, our method learns from a few high-quality human demonstrations. We capture these through videos of humans performing the task, using off-the-shelf hand-tracking to extract pick-and-place actions, which serve as training data for the robot’s task execution policy. Our method aims to learn a policy for fabric manipulation directly from a small number of human actions, without requiring any robot time at all. Prior approaches using only real-world data have utilized random self-supervised robot data collection. However, this can lead to poor performance on long-horizon multi-step tasks [1], or can require large datasets of real-world data and/or additional manual human annotations to achieve good performance [2], [3].

<sup>1</sup>Authors are with the Queensland University of Technology (QUT) Centre for Robotics, 2 George Street, Brisbane City, 4000, Queensland, Australia. (email: r21.lee@hdr.qut.edu.au; peter.corke@qut.edu.au)

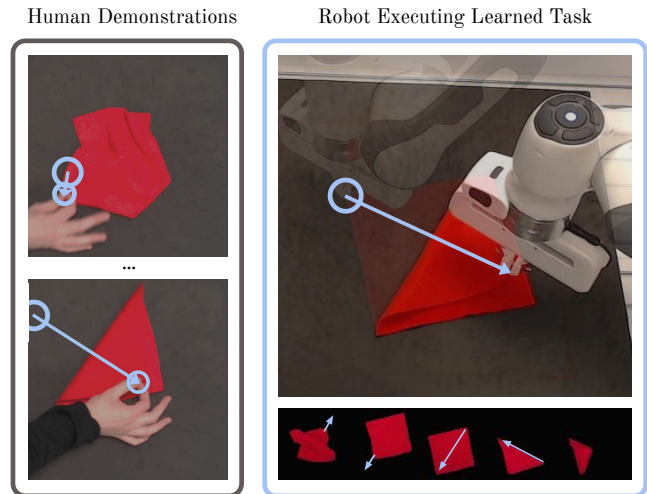


Fig. 1. Our policy learns to smooth and fold fabric from just 15 human demonstrations, enabling the robot to handle new configurations without prior robotic training.

Our sample-efficient architecture predicts place actions conditioned on pick locations, enabling focused learning of the placing component while selecting picks that yield the best placing outcome. Our work extends this concept to real-world tasks, learning from a limited amount of high-quality human demonstration data. Since humans fold cloth with ease, data collection is streamlined, eliminating the need for complex methods like kinesthetic guiding, engineered control interfaces, or VR. Our work is a step towards learning deformable object manipulation directly from human videos. This holds the potential to leverage vast, unlabeled online resources like YouTube tutorials (examples in Figure 2).

We demonstrate our approach on a challenging, long-horizon multi-step manipulation task, smoothing a cloth from a crumpled initial configuration before folding it to a goal configuration. While several prior works attempt smoothing and folding, they typically separate these tasks with some threshold for smoothness, typically learning each task with large amounts of (typically simulated) data [2], [3]. We show that using only 15 demonstrations, which takes approximately 15 minutes of human demonstration time, is adequate to learn an effective manipulation policy for the whole task. Furthermore, we demonstrate that our spatial action space, pick-conditioned policy approach is effective for learning fabric manipulation from a small number of demonstrations.

## II. RELATED WORK

### A. Deformable Object Manipulation

Manipulation of cloth has been a long-standing challenge for robotics [4]–[8]. Early approaches made use of geometric cues [9] and engineered features [10] for grasp detection for smoothing and folding.

However, recent progress in cloth manipulation has been impressive due to advancements in learning-based approaches [11]–[14]. In particular, spatial pick-and-place action spaces allow for the use of Fully Convolutional Networks for learning action affordance heatmaps [1] as well as dynamics models [15] to achieve sample efficiency. Such architectures have also been used to learn fabric smoothing via flinging [16], bimanual stretching [17] and pick-and-place [2], as well as one-step fabric folding policies [18].

Several approaches have achieved impressive results for reaching goals with dynamics models [15], [19], [20]. [18] learn a flow-based model for predicting single timestep folds, making use of simulated cloth ground-truth positions. Despite these successes, such approaches are usually limited by the requirement for large simulation datasets, while the sim-to-real gap is still a major issue with deformable objects. In contrast, our approach learns directly in the real world from a small number of demonstrations of humans performing the task.

[21] learns a pick-conditioned placing reinforcement learning value function for fabric smoothing, which shows improved sample efficiency. Their method achieves impressive performance by training a reinforcement learning policy on a large amount of simulated data, and then deploying the policy in the real world. Our method builds on the idea of learning pick-conditioned placing to create a policy based on a fully-convolutional network architecture for imitation learning with a small number of demonstrations.

[3] learns to fold T-shirts from a dataset of 4300 self-supervised and human-annotated observations, by training a number of smoothing primitives as well as a classifier for detecting when the shirt is smooth enough to begin following user-provided folding instructions. [2] learns T-shirt folding in the real world by combining imitation learning with analytic methods. Human demonstrations for folding a T-shirt were collected through a mouse-click interface to select actions for the robot to execute. The task is divided into smoothing and then folding, with separate models trained for both. In contrast, our approach learns a single policy that manipulates a cloth directly from a crumpled initial state to a folded state with demonstrations collected from human



Fig. 2. Motivation: demonstrations of cloth manipulation are widely available on the internet. We show several sample frames from YouTube videos.

hands.

### B. Learning from Human Videos

Recently, robotics research has made progress towards learning from watching humans perform tasks directly. [22] use pre-trained human hand and body pose estimation models to control a robot arm and hand in real time for a variety of manipulation tasks. They achieve this by training a re-targeting model to estimate corresponding robot hand and arm poses based on the observed human poses. In contrast, our work aims to learn a pick-and-place behaviour policy directly from human videos. Other work has investigated learning to cook from videos of humans, by detecting both objects and human grasps to create a sequence of robot actions to replicate the task [23]. By obtaining priors such as hand motion, human video demonstrations have also been used in combination with 1-2 hours of robot learning to perform a variety of mobile manipulation tasks [24]. Our work aims to learn difficult fabric manipulation behaviour directly from human videos without any robot training or data collection time.

## III. APPROACH

### A. Problem Definition

We consider the task of manipulating fabric by a single robot arm from some initial crumpled configuration, to a desired folded state. The robot observes the object resting on a flat workspace from an overhead camera, and performs pick-and-place actions, corresponding directly to locations in the image. This is referred to as a spatial action space [25], and is widely used in deformable object manipulation [1], [2], [15], [18]. Our objective is to find the optimal pick-and-place action for the task given an image observation. While prior approaches either perform smoothing alone [21], [26], [27]; assume a pre-folded state for folding [1], [18]; or combine separate smoothing and folding policies with a switching criteria [2], [3]; we aim to learn a single policy that can perform the full smoothing-folding task.

### B. Human Demonstrations

Fabric manipulation is a domain with complex dynamics and severe challenges with respect to long-term task horizons. Learning behaviours from self-supervision with no human guidance is challenging, as random actions will rarely bring the fabric to desirable configurations. Engineered biases, such as corner grasping, can assist with the data collection process [1], [18], but learning long-horizon tasks is still a challenge. A natural alternative is learning from demonstrations of humans manipulating fabric. However, demonstrating actions for fabric manipulation has typically involved user-interfaces for controlling the robot [2], [28]. Ideally, we would like to learn directly from videos of humans performing fabric manipulation behaviours directly with their own hands, as it is a very natural and efficient way to demonstrate behaviours.

To this end, we propose to employ a pre-trained hand-tracking system, which allows for accurate digit position




Fig. 3. We visualize the hand-tracking demonstration collection process. An observation is taken before the action, the hand is tracked with localization of thumb and forefingers. A grasp is detected when the fingers meet. The movement is tracked during the action, and a release is detected when the fingers separate. Finally, an observation is taken after the action.

estimation. We use Mediapipe Hands [29], a real-time hand tracking system trained to localize 21 hand coordinates from 30,000 annotated real-world hand images. The hand coordinates includes finger tip locations, which we utilize for our data-collection pipeline.

In this work, we use a pick-and-place action space for manipulation, and so we would like to retrieve the pick-and-place locations that the human selected during their manipulation of the cloth. We detect grasps by tracking when the distance between the index and thumb tip locations is below a threshold, and when this occurs, the average of the finger locations is recorded as the pick location. Similarly, the place location is recorded as the average of the thumb and index tip locations when the grasp is released, and the distance exceeds the grasp threshold.

The result is a tuple of the pick-and-place pixel coordinate locations. We record the action before and after the action is conducted, by detecting a period of minimal pixel change in the video stream. Thus, the human removes their hand from view, an image is detected, and the human performs an action. A dataset can then be naturally collected for fabric manipulation tasks.

### C. Learning to Imitate Pick-and-Place Actions




Fig. 4. A visualization of our pick-conditioned place pipeline. We concatenate the observation image with a candidate pick, taken from the cloth mask, represented as a 2D Gaussian image. The network is trained against the place heatmap label.

Our approach focuses on learning pick-and-place behaviors from limited human demonstration data, making efficient learning crucial to minimize data collection costs. In the context of deformable objects like cloth, understanding the interdependency between the 'pick' and 'place' actions is vital. For instance, in a smoothing task, multiple 'pick' points may seem viable, but each requires a specific 'place' point for proper smoothing. A mismatch between these points could inadvertently result in folding the cloth rather than smoothing it. Prior work [21] with pick-and-place action spaces for

deformable object manipulation has shown that learning only placing, conditioned on pick location, improves sample efficiency in reinforcement learning of fabric smoothing policies. This approach learns a continuous Q-network for placing conditioned on pick location, and then selects the pick location that maximizes the Q-value across picks. We propose to extend this approach for use with spatial action spaces for supervised imitation learning of pick-and-place actions, by learning to predict place heatmaps conditioned on pick location.

Our approach trains a fully-convolutional network conditioned on pick locations, specified by an image channel with a 2D Gaussian centered on the pick location, similar to the approach used for bimanual pick-and-place in [18]. We train the model to predict place locations via heatmap image labels created with 2D Gaussians centered on the place location. Thus, conditioned on a candidate pick location, we can estimate the corresponding place location by taking the argmax pixel location over the place heatmap.

Since this approach would only be effective for known pick locations, as we only have examples of ground truth pick-and-place pairs, we are unable to optimize for the pick location as the model has not been trained to lower probabilities for bad pick locations. To solve this, we simply create artificial negative samples by randomly selecting pick locations on the cloth and training against a heatmap of zeros. As such, the model can learn to output higher placing probabilities for correct pick locations, and lower probabilities for incorrect picks. Therefore, we can find the pick action that maximizes place probability by taking the argmax over possible pick locations. Since the approach in [21] trains a Q-function on rewards, this step of creating artificial negative samples is not required, as the possibility of low and high reward values makes this unnecessary for reinforcement learning. Thus, our approach allows this idea to be applied in supervised, imitation learning settings, with spatial action spaces.

Given a dataset of human demonstrations, we train our pick-conditioned place network to imitate the human behaviour. The network is trained to output place heatmaps to mimic the human action. The spatial action space of the fully-convolutional network lends itself well to modelling the multi-modal nature of human demonstrations for fabric manipulation.

#### D. Implementation Details

We implement our model as a fully-convolutional network in PyTorch. The input to the network is a  $64 \times 64$  observation image, concatenated with a 2D Gaussian pick image. The output is a  $64 \times 64$  place heatmap. Thus, there are  $64 \times 64$  possible place locations, while the pick locations are limited to the cloth mask. The network is trained with Binary Cross Entropy loss against the ground-truth place heatmaps, with a batch size of 16. Validation is performed every 500 training steps, by evaluating the mean squared error of predicted pick-and-place actions against the ground truth actions in a batch sampled from the validation set. We save the best model according to this validation error after 500,000 training steps.

During training, we evaluate our model by sub-sampling the cloth mask by a factor of 2 to reduce time and memory. At test time, we evaluate every possible pick location in the mask for full-resolution pick selection.

### IV. EXPERIMENTS

We demonstrate our approach on a real-world fabric manipulation task, folding a cloth in half twice into a small triangle, beginning from a crumpled initial configuration. Folding a square fabric into this shape, assuming a smooth initial state, has been shown in prior work [1], [18], [19]. To the best of our knowledge, we are the first to train a single policy that can reach a folded state from a crumpled state.

#### A. Experimental Setup

Our experimental setup consists of a Franka Emika Panda robot workstation, with a RealSense L515 camera. We make use of depth for observation, and infrared for cloth masking, while RGB is used for visualization only. The workspace is  $40 \times 40$ cm square of PVA foam. The fabric is a  $22 \times 22$ cm square of red polar fleece.

#### B. Data Collection

We use our hand-tracking-based approach to collect a small dataset of 15 demonstration episodes, each beginning from a crumpled initial state. We use 10 demonstrations for training and the remaining 5 for validation. An episode consists of the human demonstrating the given task with a single hand, first smoothing the cloth to a suitable stage to begin folding, and then folding the cloth in half twice to create a small square.

While both depth and RGB have been shown to be suitable for fabric manipulation tasks [1], [18], [19], we opt to use depth to allow for direct generalization to cloths of various appearances. Additionally, we mask the depth image to isolate the cloth, setting the background to 0. Masks are widely used in deformable object manipulation to limit actions to the object itself [1], [3], [16], [18]. Generally, background removal is a mature problem, which can be solved effectively using a variety of methods. However, we note that cloths of many appearances are easily isolated from the foam background using infrared images, which makes it a convenient choice for us, and allows the policy to work across fabrics of a variety of colors.

We augment the dataset by a factor of 20, by randomly applying rotation, flips, small changes in image scaling and Gaussian noise to the depth images. Additionally, we apply depth scaling augmentation between the factors 0.8 and 1.2 to add robustness to cloth thickness.

#### C. Robot Execution

We execute the trained policy directly on the real robot and demonstrate its performance for the given task. Selecting the pick-and-place locations in the image, we transform the pixel coordinates to the XY plane of the robot’s workspace via a simple linear transform. The robot then moves to a pose directly above the grasp pose, rotates its wrist so the grasp direction is perpendicular to the action vector to ensure consistency between actions, moves down until it senses a cartesian contact, and then grasps. We then execute the action by raising the gripper, moving in a straight line to the pose above the drop point, descending a small amount, and releasing the gripper. A grasping width heuristic is used, with a larger finger width when grasping within a small pixel threshold of the edge of the cloth mask or when the place location is on the mask, and a narrower pinch grasp otherwise, which allows the robot to grasp only the top layer of cloth. This facilitates smoothing and folding, as folding requires grasping all layers, and is typically performed from the edges inwards, while smoothing usually requires top-layer grasping only. We use a simple 3D-printed gripper with a sloped, pointed design that facilitates effective pinch grasping with minimal fabric contact, reducing the risk of bunching or crumpling. This design worked well on the tested fabrics. However, generalization to significantly different materials would likely necessitate modifications such as alternative gripper geometries, tactile sensing, or different friction characteristics. The STL file for our gripper is available on the project website.

#### D. Tasks

**Full Task:** To ensure repeatability and consistent evaluation, we adopt the predefined initial cloth configurations as suggested by [26]. Our work, however, extends the scope to include both smoothing and folding operations in a single, multi-step task. Each of the three initial crumpled configurations is tested five times, with the policies running for 15 timesteps to complete both smoothing and folding autonomously. This task is particularly challenging due to its long-horizon nature, where success in the final folded state is contingent upon precise intermediate smoothing actions.

**Smooth Initial Configuration:** We also introduce a simplified task that starts from a pre-smoothed cloth. This is designed to isolate the folding behavior and evaluate the baseline models’ performance. The policies run for 5 timesteps for this shorter task. The results underline the importance of conditioning the ‘place’ actions on ‘pick’ locations, an essential feature for tasks involving deformable objects like cloth.

#### E. Metrics

Defining suitable metrics for deformable object manipulation is challenging, as it is impossible to perfectly perceive




Fig. 5. Metrics plotted for a successful (top) and failed (bottom) episode from the same initial configuration, ending at the timestep with the best IOU.

the object state. As such, we report several metrics:

**Success:** Following prior work, we qualitatively assess if the folding episode was a success, as defining a precise quantitative termination condition for deformable object manipulation is difficult due to variations in fabric configurations. We consider an episode successful if any timestep within the fixed-duration run achieves the desired folded state (triangle with minor defects allowed). Results are available on the project website.

**Intersection Over Union (IOU):** We report the IOU of the fabric mask, after translating and rotating it to best match a mask of a cloth in the final folded state. For each episode, we take the highest IOU across all timesteps. We note that IOU is not a perfect metric, as it only considers the silhouette of the fabric, and not the internal layers or 3D shape.

**Intermediate Smoothing Coverage (ISC):** We report the ratio of fabric coverage (sum of mask pixels) as a fraction of the coverage of a smoothed cloth mask. We take the highest coverage score across all timesteps in the episode, which is a secondary metric for the task, indicating how well the fabric was smoothed during the episode.

#### F. Baselines

To evaluate the overall performance of our method, and to validate the effectiveness of our pick-conditioned place model, we compare it against both human performance and a baseline architecture.

**Human** We show the performance of a human using a single hand to solve the task. This represents an approximate upper bound for our method, which is trained to mimic human demonstrations.

**Pick+Place** We demonstrate a simple FCN network that outputs two heatmaps directly, one for pick-and-place, based on the pick-and-place policy from [2]. The network architecture is the same as ours, aside from the input and outputs. Notably, the place is not conditioned on the pick. Since the pick cannot be constrained to the fabric mask, as it might interfere with the pick-and-place correspondence, if the pick location is predicted off the fabric, we simply find the closest point on the fabric mask to the predicted pick location.

#### G. Results

In the robot evaluation experiments, we aim to assess whether the policy can learn to consistently reach a folded

configuration from initially crumpled states, from only a small number of human demonstrations collected via our approach. Additionally, we aim to evaluate the performance of our pick-conditioned place fully-convolutional model, compared to an alternative baseline model. We illustrate the performance visually in Figure 7, by showing the start and best timestep according to the IOU metric, for the best episode for each task. The results for the full task are shown in Table I.

Our proposed method was successful in 13 of the 15 runs, for an average of 4.334/5, or 86.7%. We find that the policy consistently performs to a high standard at smoothing, shown by the intermediate smoothing coverage (ISC) score, which approaches the human’s score, as well as having a high IOU score, just below that of the human. However, we note that for the single failure case, the policy chose to fold before the fabric was adequately smoothed. This is reflected by the average ISC score for that task being marginally lower than the other tasks. We illustrate this failure case, compared with a success case, in Figure 5. This shows the progression of the task metrics throughout the episode, showing how the ISC score progresses upwards until the smooth state is reached, followed by the IOU score progressing until the folded state is reached. The fail case folded too early in the smoothing phase, resulting in a somewhat crumpled final folded state.

We find that the baseline network architecture shows some

TABLE I  
EVALUATION RESULTS FROM CRUMPLED INITIAL CONFIGURATION.  
SUCCESS IS OUT OF 5 RUNS.




	Method	Success	IoU	ISC
	Human	5	<b>0.883</b>	<b>0.98</b>
	<i>PickToPlace (Ours)</i>	5	<b>0.843</b>	<b>0.965</b>
	Pick+Place	2	0.752	0.846
	Human	5	<b>0.876</b>	<b>0.98</b>
	<i>PickToPlace (Ours)</i>	4	<b>0.801</b>	<b>0.958</b>
	Pick+Place	0	0.776	0.851
	Human	5	<b>0.86</b>	<b>0.982</b>
	<i>PickToPlace (Ours)</i>	4	<b>0.811</b>	<b>0.978</b>
	Pick+Place	0	0.684	0.909
Mean	Human	5	<b>0.873</b>	<b>0.981</b>
	<i>PickToPlace (Ours)</i>	4.334	<b>0.818</b>	<b>0.967</b>
	Pick+Place	0.667	0.737	0.869

TABLE II  
EVALUATION RESULTS FROM SMOOTH INITIAL CONFIGURATION.  
SUCCESS IS OUT OF 5 RUNS.

Method	Success	IoU
Human	5	0.876
<i>PickToPlace (Ours)</i>	4	0.845
Pick+Place	1	0.622

smoothing and folding ability, successfully smoothing and folding the cloth from an initially crumpled state to the desired folded state twice out of the total 15 runs, with an average of 0.667/5 across tasks. While this success rate is low, we note the long-horizon nature of the task requires that consistently good actions are taken with small margin for failure. We hypothesize that the low success rate is caused by two main factors. Firstly, the network is learning both pick-and-place outputs, while our network is only learning place, conditioned on pick. Learning both pick-and-place is a more challenging learning problem, and likely requires more data to learn an adequate policy. Secondly, and more fundamental to the architecture’s capability, the place location is not conditioned on the pick location. Thus we observe several poor action choices that are caused by mismatching pick-and-place locations. This is confirmed by the initially smoothed task results.

The initially smoothed results, shown in Table II, confirm that even from a smooth initial state, the baseline policy is unable to consistently fold the cloth. We observe that the poor performance of this policy is directly caused by the inability to select pick-and-place actions that correspond correctly. When the cloth is in the smooth state, all corners are potentially good pick, and place, locations. However, given a particular corner is chosen, only the opposite corner is a good place location. The model will often choose a place location that does not correspond to the chosen pick, for example by folding from one corner to the nearest, rather than the opposite corner. We illustrate this in Figure 6. Conversely, since our architecture conditions on pick location to predict place, and then chooses the best pick-and-place action according to place probability, it only predicts a single place location that achieves the fold.

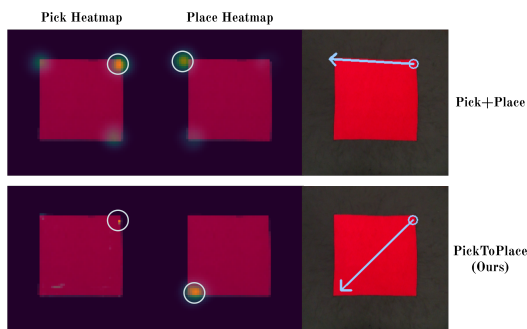


Fig. 6. A common failure case when predicting pick-and-place directly, without modeling the conditional relationship. Above: The output of the model trained to output both pick-and-place. Below: Our model trained to output place heatmaps conditioned on pick.




Fig. 7. Runs performed by our approach (above) and the Pick+Place baseline (below) from all initial configurations. Shown are the initial configurations and the best states from the best runs according to the IoU metric. Successes are outlined in green, while failures are outlined in red.

#### H. Generalization

Since our method is trained on depth images, the resulting model is invariant to visual properties. We explore this generalisation capability by running episodes on a variety of cloths shown in Figure 8. We find that in addition to generalizing to visual appearance, our model is also robust to a range of material properties, such as thickness, texture and even shape of the cloth. For these experiments, we vary only the grasping width of the robot fingers to handle grasping of different materials.



Fig. 8. Our method can generalize to fabrics of a variety of appearances and textures. We show the initial, smooth state before folding, and final configurations from successful folding episodes.

#### V. CONCLUSION

We present an approach for real-world, long-horizon fabric manipulation that leverages a small number of human demonstrations. Using off-the-shelf hand-tracking, we extract pick-and-place locations from videos. This data is used to train our sample-efficient pick-conditioned place heatmap prediction model. With only 15 demonstrations (approximately 15 minutes of collection time), we achieve reliable folding success rates exceeding 85%, outperforming alternative architectures. Our model generalizes to unseen fabrics with varying textures, colors, and shapes.

Our approach demonstrates the potential of learning fabric manipulation directly from human videos. While the current work requires suitable human demonstrations and is limited to settings where collecting this data is possible, our method’s sample efficiency and simple data collection procedure minimize this data burden. We envision scaling this approach to diverse tasks, materials, and unlabeled online video resources. Incorporating bimanual manipulation and addressing cloth occlusion will further expand the method’s applicability. Additionally, while we focus on a simple square cloth in a top-down setting, the core principles could be adapted to other top-down pick-and-place tasks. Extensive testing on broader applications is needed.

## REFERENCES

- [1] R. Lee, D. Ward, A. Cosgun, V. Dasagi, P. Corke, and J. Leitner, "Learning arbitrary-goal fabric folding with one hour of real robot experience," *arXiv preprint arXiv:2010.03209*, 2020.
- [2] R. Hoque, K. Shivakumar, S. Aeron, G. Deza, A. Ganapathi, A. Wong, J. Lee, A. Zeng, V. Vanhoucke, and K. Goldberg, "Learning to fold real garments with one arm: A case study in cloud-based robotics research," *arXiv preprint arXiv:2204.10297*, 2022.
- [3] Y. Avigal, L. Berscheid, T. Asfour, T. Kröger, and K. Goldberg, "Speedfolding: Learning efficient bimanual folding of garments," *arXiv preprint arXiv:2208.10552*, 2022.
- [4] Y. Li, X. Hu, D. Xu, Y. Yue, E. Grinspun, and P. K. Allen, "Multi-sensor surface analysis for robotic ironing," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 5670–5676.
- [5] —, "Multi-sensor surface analysis for robotic ironing," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 5670–5676.
- [6] F. Osawa, H. Seki, and Y. Kamiya, "Clothes folding task by tool-using robot," *Journal of Robotics and Mechatronics*, vol. 18, no. 5, pp. 618–625, 2006.
- [7] M. Cusumano-Towner, A. Singh, S. Miller, J. F. O'Brien, and P. Abbeel, "Bringing clothing into desired configurations with limited perception," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 3893–3900.
- [8] D. Seita, N. Jamali, M. Laskey, R. Berenstein, A. K. Tanwani, P. Baskaran, S. Iba, J. Canny, and K. Goldberg, "Deep Transfer Learning of Pick Points on Fabric for Robot Bed-Making," in *International Symposium on Robotics Research (ISRR)*, 2019.
- [9] J. Maitin-Shepard, M. Cusumano-Towner, J. Lei, and P. Abbeel, "Cloth grasp point detection based on multiple-view geometric cues with application to robotic towel folding," in *2010 IEEE International Conference on Robotics and Automation*. IEEE, 2010, pp. 2308–2315.
- [10] B. Willimon, S. Birchfield, and I. Walker, "Model for unfolding laundry using interactive perception," in *IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*, 2011, pp. 4871–4876.
- [11] R. Jangir, G. Alenya, and C. Torras, "Dynamic cloth manipulation with deep reinforcement learning," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 4630–4636.
- [12] J. Matas, S. James, and A. J. Davison, "Sim-to-real reinforcement learning for deformable object manipulation," in *Conf. on Robot Learning*, 2018, pp. 734–743.
- [13] Y. Tsurumine, Y. Cui, E. Uchibe, and T. Matsubara, "Deep reinforcement learning with smooth policy update: Application to robotic cloth manipulation," *Robotics and Autonomous Systems*, vol. 112, pp. 72–83, 2019.
- [14] J. Hietala, D. Blanco-Mulero, G. Alcan, and V. Kyrki, "Learning visual feedback control for dynamic cloth folding," 2022.
- [15] R. Lee, M. Hamaya, T. Murooka, Y. Ijiri, and P. Corke, "Sample-efficient learning of deformable linear object manipulation in the real world through self-supervision," *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 573–580, 2022.
- [16] H. Ha and S. Song, "Flingbot: The unreasonable effectiveness of dynamic manipulation for cloth unfolding," in *Proceedings of the 5th Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, A. Faust, D. Hsu, and G. Neumann, Eds., vol. 164. PMLR, 08–11 Nov 2022, pp. 24–33. [Online]. Available: <https://proceedings.mlr.press/v164/ha22a.html>
- [17] Z. Xu, C. Chi, B. Burchfiel, E. Cousineau, S. Feng, and S. Song, "Dexterity: Deformable manipulation can be a breeze," in *Proceedings of Robotics: Science and Systems (RSS)*, 2022.
- [18] T. Weng, S. Bajracharya, Y. Wang, K. Agrawal, and D. Held, "Fabricflownet: Bimanual cloth manipulation with a flow-based policy," in *Conference on Robot Learning*, 2021.
- [19] R. Hoque, D. Seita, A. Balakrishna, A. Ganapathi, A. K. Tanwani, N. Jamali, K. Yamane, S. Iba, and K. Goldberg, "Visuospatial foresight for multi-step, multi-task fabric manipulation," *arXiv preprint arXiv:2003.09044*, 2020.
- [20] W. Yan, A. Vangipuram, P. Abbeel, and L. Pinto, "Learning predictive representations for deformable objects using contrastive estimation," in *Conference on Robot Learning*. PMLR, 2021, pp. 564–574.
- [21] Y. Wu, W. Yan, T. Kurutach, L. Pinto, and P. Abbeel, "Learning to manipulate deformable objects without demonstrations," in *16th Robotics: Science and Systems, RSS 2020*. MIT Press Journals, 2020.
- [22] A. Sivakumar, K. Shaw, and D. Pathak, "Robotic telekinesis: learning a robotic hand imitator by watching humans on youtube," *arXiv preprint arXiv:2202.10448*, 2022.
- [23] Y. Yang, Y. Li, C. Fermuller, and Y. Aloimonos, "Robot learning manipulation action plans by" watching" unconstrained videos from the world wide web," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 29, no. 1, 2015.
- [24] S. Bahl, A. Gupta, and D. Pathak, "Human-to-robot imitation in the wild," 2022.
- [25] D. Wang, R. Walters, X. Zhu, and R. Platt, "Equivariant q-learning in spatial action spaces," in *Conference on Robot Learning*. PMLR, 2022, pp. 1713–1723.
- [26] S. Sharma, E. Novoseller, V. Viswanath, Z. Javed, R. Parikh, R. Hoque, A. Balakrishna, D. S. Brown, and K. Goldberg, "Learning switching criteria for sim2real transfer of robotic fabric manipulation policies," *arXiv preprint arXiv:2207.00911*, 2022.
- [27] 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.
- [28] B. Waymouth, A. Cosgun, R. Newbury, T. Tran, W. P. Chan, T. Drummond, and E. Croft, "Demonstrating cloth folding to robots: Design and evaluation of a 2d and a 3d user interface," in *2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*. IEEE, 2021, pp. 155–160.
- [29] F. Zhang, V. Bazarevsky, A. Vakunov, A. Tkachenka, G. Sung, C.-L. Chang, and M. Grundmann, "Mediapipe hands: On-device real-time hand tracking," *arXiv preprint arXiv:2006.10214*, 2020.