Evaluating the Pre-Dressing Step: Unfolding Medical Garments via Imitation Learning

David Blanco-Mulero, Julia Borras, Carme Torras

***Abstract*— Robotic-assisted dressing has the potential to significantly aid both patients as well as healthcare personnel, reducing the workload and improving the efficiency in clinical settings. While substantial progress has been made in robotic dressing assistance, prior works typically assume that garments are already unfolded and ready for use. However, in medical applications gowns and aprons are often stored in a folded configuration, requiring an additional unfolding step. In this paper, we introduce the *pre-dressing* step, the process of unfolding garments prior to assisted dressing. We leverage imitation learning for learning three manipulation primitives, including both high and low acceleration motions. In addition, we employ a visual classifier to categorise the garment state as closed, partly opened, and fully opened. We conduct an empirical evaluation of the learned manipulation primitives as well as their combinations. Our results show that highly dynamic motions are not effective for unfolding freshly unpacked garments, where the combination of motions can efficiently enhance the opening configuration.**

I. Introduction

Dressing assistance is a crucial task performed by caregivers to assist individuals in nursing facilities or hospitals. Beyond patient care, dressing is also crucial in pre-surgery rooms and intensive care units, where nurses assist each other dress up in protective gowns. Robotic-assisted dressing [1] has emerged as a promising solution to address the shortage of healthcare personnel while also reducing the burden of repetitive tasks on caregivers and nurses.

Robotic-assisted dressing presents several challenges. First of all, dressing is a task that entails the manipulation of a deformable object, which thus needs to deal with the infinite degrees of freedom of garments as well as self-occlusions [2], [3]. Secondly, dressing assistance requires manipulation close to a human, imposing harder restrictions in the acceleration forces of the motions used compared to other deformable object manipulation tasks [4]. Although recent works have demonstrated great progress in assisted dressing [5], [6], they commonly assume that the garment is already unfolded, ready for dressing. In reality, however, garments such as medical gowns and aprons are typically stored in a folded configuration, requiring an additional unfolding step before they can be used. For this reason, in this work we explicitly address what we denote as the *pre-dressing* step: unfolding an initially folded garment and prepare it for assisted dressing, see Fig. 1.

This work was financially supported by the European Union's Horizon Europe programme project SoftEnable (grant agreement No. 101070600).

The authors are with the Institut de Robòtica i Informàtica Industrial, CSIC-UPC, 08028, Barcelona, Spain. (e-mail: david.blanco.mulero@upc.edu)

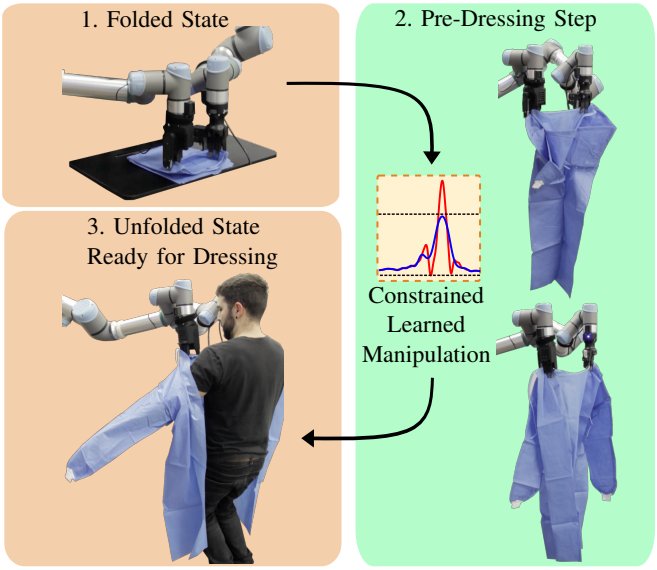

Fig. 1: The *pre-dressing* step is the transition from an initially folded garment to an unfolded garment prepared for assisted dressing. Here, we utilise imitation learning for learning manipulation primitives that perform the garment unfolding while adhering to the robot hardware limits.

While recent research on deformable object manipulation tackles the problem of garment unfolding [4], [7]–[9], these works consider fabrics that start in a crumpled configuration. To the best of our knowledge, this work is the first to address unfolding a garment starting from an initial folded configuration. In this work, we investigate the effectiveness of different manipulation primitives for unfolding a medical gown. To manipulate the garment we leverage the imitation learning framework of Dynamic Movement Primitives (DMPs) [10]. Then, to evaluate the success of the task, we propose to use a visual classifier to determine the opening state of the garment. Supplementary materials, as well as details of the DMP and visual classifier, are available on our website[1].

In summary, our contributions include:

- Introducing the problem of pre-dressing, focusing on the unfolding of a garment before dressing assistance.
- Presenting a solution for pre-dressing based on DMPs and a visual classifier to identify the opening state of the garment.
- Conducting an empirical evaluation of different manipulation primitives, comparing high-acceleration and low-acceleration manipulation actions to assess their effectiveness in unfolding garments.

[1] https://sites.google.com/view/pre-dressing

(a) Folded. (b) Closed with sleeves tangled. (c) Closed untangled. (d) Partly Opened. (e) Opened 1. (f) Opened 2.

Fig. 2: Example of a medical gown in: (a) folded configuration, and (b) closed configuration with the sleeves tangled after lifting the garment from a flat surface. The categories utilised for identifying the gown state and training a visual classifier are: (c) Closed, (d) Partly Opened, (e) Opened with sleeves forward, and (f) Opened with both sleeves hidden.

II. Related work

A. Unfolding Deformable Objects

The problem of unfolding a deformable object has been studied in the literature for objects such as cloths [9], garments [4], and bags [11], to name a few. Prior works on cloth and garment unfolding have primarily focused on the task of flattening [4], [7]–[9], [12]. In this task, given a garment in a crumpled state, the robot manipulator needs to find the necessary actions to flatten the object, maximising the area that it covers. Similarly, recent works on bag unfolding address the problem of maximising the volume of an initial crumpled bag [11], [13], [14]. These approaches share the common goal of moving the deformable object from an initially unstructured state to a more structured configuration. In contrast, this work tackles the opposite problem, unfolding a garment from an initial structured folded configuration.

B. Assisted Dressing

Assisted dressing has gained significant attention over the past years [1], [5], [6], [15], [16]. Here, we primarily focus on the recent advances related to dressing the upper body of a human, one important aspect in occupational therapy. Previous works have developed learning-based approaches for inserting the garment sleeves onto a person's arm [5], [6], assuming the garment is already grasped and ready for dressing. Alternatively, other works have also incorporated grasping the garment into their assisted dressing pipeline [1], [17]. As an example, Zhang et al. [17] proposed an edge tracing method for preparing the garment before dressing. However, the aforementioned works assume that the garment is already unfolded, facilitating the dressing step. In this work, we tackle this crucial step required prior to the dressing assistance step, evaluating the effectiveness of different manipulation primitives to unfold the garment.

III. Pre-Dressing Manipulation and Metrics

A. Task Definition

The objective of the pre-dressing task is to prepare a folded garment for dressing assistance. Hence, given a garment in a folded configuration placed on a flat surface, the robot needs to perform the necessary actions to unfold the garment (see Fig. 1). More specifically, in this paper we focus on the task of unfolding medical gowns. Unlike other garments such as t-shirts, a folded gown may exhibit sleeve intersections, adding complexity to its unfolding process. As shown in Fig. 2b, when the gown is grasped from a flat surface and the contact is removed, the overlapped fabric layers hinder the complete opening of the garment. In line with prior works on robotic-assisted dressing [5], [17], [18], we consider a bi-manual robot set-up scenario. In addition, we assume that the detection of the grasping location is solved, and perform a top grasp of a single layer of the garment in the collar.

B. Metrics for Evaluating Unfolding Success

Following the definition of bending levels proposed in the taxonomy in [19], the pre-dressing task requires to move from an structured (folded) configuration to an unstructured (opened) bending configuration, increasing the unstructuredness of the garment. The unstructured configuration is defined by the number of accessible keypoints. Here, once the garment has been lifted and the sleeves are untangled, losing visibility of the sleeves indicates that the opening state has improved. Hence, we utilise as keypoints the position of the sleeves.

To identify the state of the garment, rather than manually engineering a reward based on the position of the sleeves, we train a visual classifier to identify the opening state. We define three categories, illustrated in Fig. 2, as follows:

- **Closed:** the sleeves of the gown are either tangled and not visible or fully visible at the back of the gown, see Fig. 2b and Fig. 2c. Furthermore, the left and right back sides of the gown are tangled.
- **Partly Opened:** at least one of the sleeves faces forward. In addition, the back sides of the gown are untangled and slightly opened, see Fig. 2d.
- **Opened:** one or both sleeves are either partly visible or not visible. Here, both back sides of the gown are untangled and widely opened, see Fig. 2e and Fig. 2f.

Note that the main difference between the categories of partly opened and opened lies in the position of the sleeves and a wider opening of the back side of the gown.

The approach of recognising the state using a visual classifier circumvents the challenge of self-occlusion, which arises when using motion capture systems. Unlike other deformable objects such as bags [13], where positioning markers remain visible, markers placed on the sleeves of a gown may become obscured during unfolding. Another alternative would be to utilise the Chamfer Distance, which has been previously used for identifying the success in cloth unfolding tasks [7], [8]. However, this would require access to a canonical shape. In this task, there are thousands of valid configurations for pre-dressing, which would require a vast amount of canonical shapes to contrast against, undermining the potential use of such metric.

C. Manipulation Primitives for Pre-Dressing

To comprehensively study the complexity of the pre-dressing step, we devise three bi-manual manipulation primitives that target different unfolding dynamics of the gown. The manipulation primitives are as follows:

- **Fling**: similar to the primitives used in [4], [13], this dynamic motion entails rapidly accelerating the robot, leveraging inertia to push air into the gown and improve its opening state. The main rotation takes place around the Y-axis of the robot frame.
- **Shake**: rapid motion where the robot moves repeatedly back and forth while rotating the end-effector to loosen the garment folds. As in the previous motion, the main rotation is around the Y-axis of the robot frame.
- **Twist**: motion with lower accelerations where the robot moves twice forward while rotating the end-effector in opposite directions to create tension in the garment. In this motion the main rotation is around the Z-axis of the robot frame.

One solution for learning these primitives is utilising Reinforcement Learning and training a policy in simulation [4], [9]. However, due to the sim-to-real gap in garment simulation [20], as well as the challenge of simulating the aerodynamics when dynamically unfolding a deformable object [13], we decide to apply the imitation learning framework of DMPs [10] to circumvent these challenges.

The three motions are demonstrated by a human and captured using a motion capture system. Then, the demonstrations are pre-processed to reduce noise from the capture system and to filter out movements along non-essential axes. Additionally, the maximum distance between the end-effectors is constrained to prevent excessive tension that could tear apart the garment. Taking into account that a learned dynamic motion can exceed the hardware limitations of the manipulator, we follow the same approach as in [13] and use a DMP formulation that constrains the position, velocity and acceleration to those of the robot system [10].

In addition to the three learned manipulation primitives, we introduce a quasi-static primitive inspired by [13], designed to further refine the opening state of the garment. This motion slowly moves forward the garment, maintaining the distance between the end-effectors, and finally positioning the gown for handover to a human for the dressing step.

IV. Experiments

The goal of our experiments is to assess the requirements for succeeding in the pre-dressing step. To that aim, our experiments investigate: 1) which manipulation primitives are more suitable for pre-dressing, 2) the performance of combining motion primitives, and 3) the performance in previously opened and recently unpacked gowns.

A. Experimental Set-up

Our set-up consists of a dual-arm system of two 6-DoF UR5e robot arms, which are mounted in a platform resembling the pose of a human. The robot arms are equipped with a gripper designed for the EU project SoftEnable [21], which enables automatically grasping a single layer of the garment. Our set-up also includes a RealSense D435i camera for capturing the gown configuration and classifying its state. Additionally, a table is placed beneath the robots for grasping the flat gowns. The table is set at a height ensuring that the garment remains suspended in the air after being lifted.

Our experiments evaluate garments in two configurations: 1) medical gowns that have been previously opened, and 2) recently unpackaged medical gowns. For the previously opened gowns, we use the same gown, folding it in the same sequence in which the sleeves are tangled, ending after multiple folds in a square folded configuration. The unpacked gowns start from a randomly folded configuration as they come out of a sealed bag. Thus, some of these garments may exhibit more complex folds that require more effort to unfold. In the experiments, the initial configuration of the gown once lifted from the flat surface is random. Each experiment consists of up to 5 iterations to assess improvement in gown opening. For previously opened gowns we performed three trials per primitive. For packed gowns, due to time and resource constraints, we performed one trial per primitive or combination, except for primitive and quasi-static combination, in which we performed two trials.

For the constrained learned DMPs we set the same constraints as in [13], that is, 98% of the joint position, velocity and acceleration limits. Finally, for the visual classifier we utilise YOLOv11 [22] and train it using a dataset created from samples of the three categories presented in Section III-B, gathering data from manipulation performed by a human. Additional details can be found on our website[1].

B. Manipulation Primitives Comparison

We start by evaluating the three learned primitives without the quasi-static motion. The results are shown in Table I and Fig. 3. The percentage in Table I indicates the percentage of experiments in which the gown was classified as opened or as partly opened. First, we analyse the results of previously opened gowns. We can notice that all primitives achieve a large success rate, where the twist primitive performs best. Additionally, all the primitives manage to move the sleeves to the forward position of the gown. However, looking at Fig. 3, the only primitive that results in the sleeves not visible by the camera is the twist primitive. This is a result of the final position of the robot arms in the twist primitive, which

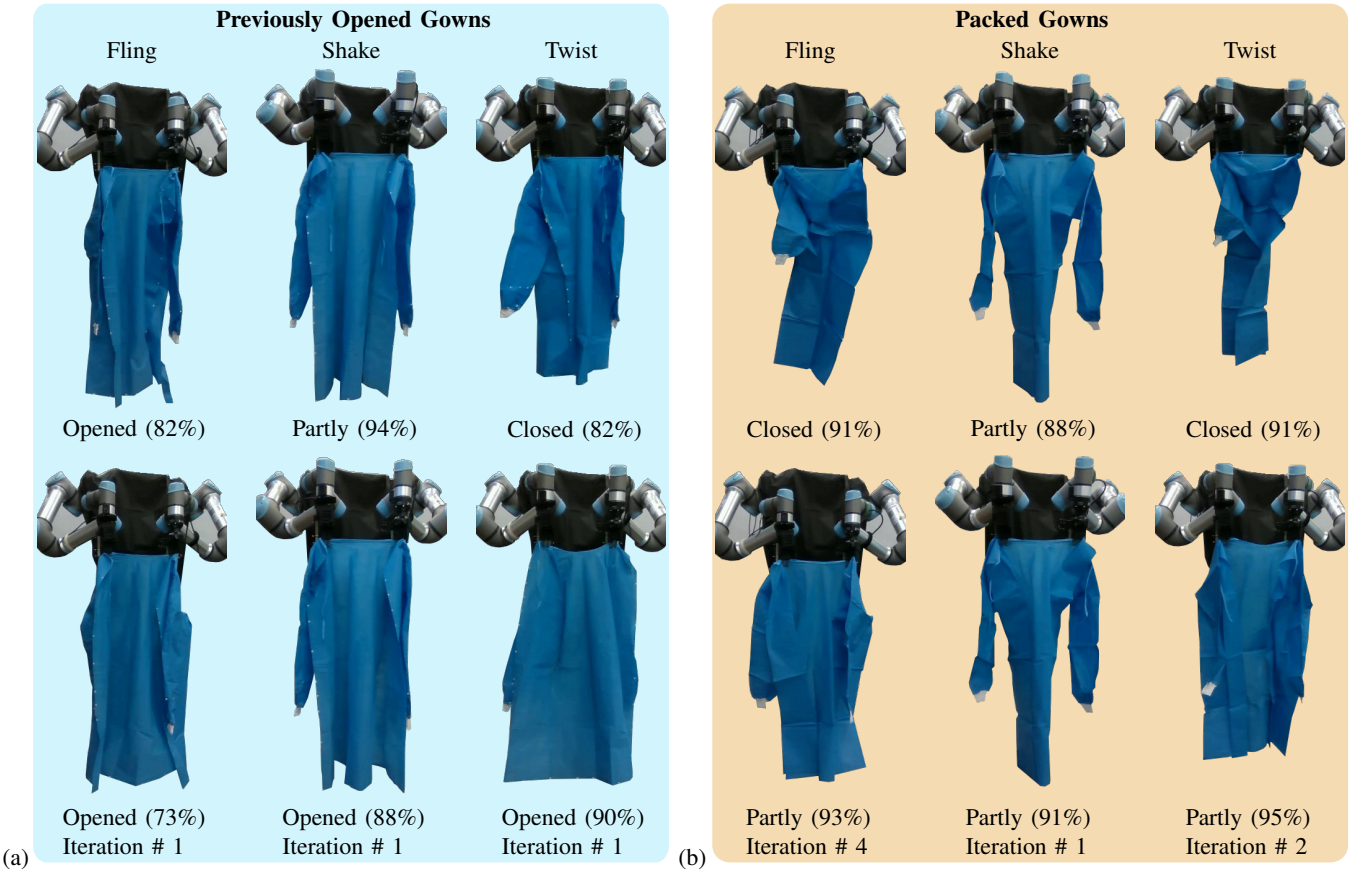

Fig. 3: Qualitative results for the three manipulation primitives in (a) previously opened gowns and (b) unpacked gowns. The first row shows the gown before execution, and the second row after reaching an opened or a partly opened state. We indicate the category predicted by the visual classifier and its confidence, as well as the number of iterations it took to reach either an opened state, or, if no further improvement occurred, the first iteration at which a partial state was reached.

rotates the grippers facilitating the opened state. Moreover, all primitives take on average only one iteration to achieve a fully opened state. In general, this indicates that both quasi-static and dynamic primitives are great candidates for unfolding a garment which has been previously opened.

Now, looking at the results for packed gowns we observe that the performance drastically drops. None of the primitives is able to reach the opened state. Furthermore, while all primitives are able to achieve a partly opened state, the number of iterations increases. It is important to note that although the shake primitive is able to partly open the gown within one iteration, the opening is significantly worse than the one achieved by the twist primitive, as shown in Fig. 3. By contrast, the width of the gown after applying either the fling or the twist is significantly larger, while requiring more iterations. These results highlight the need for a motion that refines the state of the gown further.

C. Combining Dynamic and Quasi-Static Motions

Next, we investigate the effectiveness of combining different motions for the unpacked gown case. Here, we evaluate two approaches: 1) combining two of the learned manipulation primitives, and 2) performing the quasi-static motion described in Section III-C after a learned manipulation primitive. The results are shown in Table II and Fig. 4.

Firstly, we analyse the results from the combination of the learned primitives. We select as starting manipulations

TABLE I: Quantitative results for previously opened and unpacked gowns using the three learned primitives. The results indicate the percentage of gowns classified as opened or partly opened according to the trained classifier, and if the arms are positioned Forward (Fwd.), which indicates that the gown is fully opened. The results also indicate the mean number of iterations (It.) that it takes to reach either the open or partly opened state of the gown.

		Opened	Partly Opened	Arms Fwd.	# It.
Prev. Opened	Fling	33.33%	66.67%	100%	1
	Shake	66.67%	33.33%	100%	1
	Twist	100%	0%	100%	1
Unpacked	Fling	0%	100%	0%	4
	Shake	0%	100%	0%	1
	Twist	0%	100%	0%	2

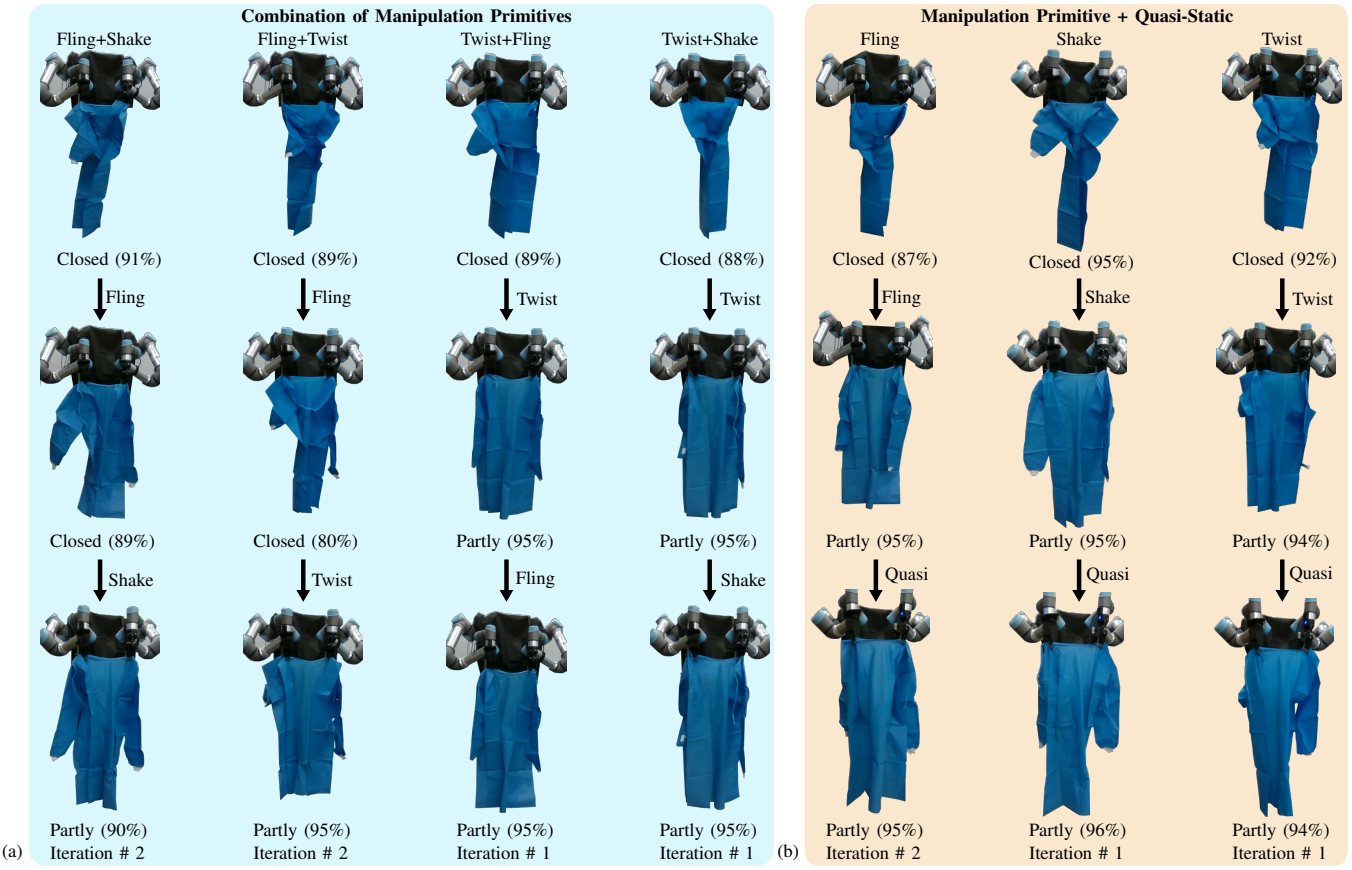

Fig. 4: Qualitative results for unfolding unpacked gowns combining the three manipulation primitives (a) with each other and (b) with a quasi-static motion. The first row shows the gown shape before execution, while the second and third rows show the resulting gown state after the corresponding manipulation indicated next to an arrow. We indicate the category predicted by the visual classifier and its confidence, as well as the iteration at which a partly opened state was reached.

the twist and fling primitives, where the first is selected due to the superior performance in the prior results, and the second to evaluate the performance of the fling dynamic motion followed by primitives with lower accelerations. In this experiment, an iteration involves the execution of two primitives. The combination of the fling with either the shake or twist motion requires an average of two iterations. This yields the same number of motions as the fling standalone motion, four in total. Secondly, we analyse the results of combining the learned primitives with the quasi-static motion. Compared to the results in Table I, the quasi-static motion helps to reduce the number of iterations required for partly opening the gown. These results showcase that combining different motion primitives helps to unfold the garment further than single primitives for unpacked gowns.

More generally, looking at Table II we can notice that none of the combined approaches is able to achieve a fully opened state. This results from the fact that the electrostatic forces between the layers of the unpacked garment are really high, and even after using dynamic manipulation primitives, the manipulation does not separate these layers. Nevertheless, looking at Fig. 4, the partly opened state is such that the back of the gown is opened enough for a human to insert the arms. Furthermore, the sleeves are positioned at the front, where all manipulations achieve 100% success, which enables further manipulation for the subsequent dressing step.

TABLE II: Quantitative results for unpacked gowns using combinations of the three learned primitives and a quasi-static motion. The results indicate the percentage of gowns classified as opened or partly opened according to the trained classifier, and if the arms are positioned Forward (Fwd.). The results also indicate the mean number of iterations (It.) that it takes to reach either the open or partly opened state.

	Opened	Partly Opened	Arms Fwd.	# It.
Fling + Shake	0%	100%	100%	2
Fling + Twist	0%	100%	100%	2
Twist + Fling	0%	100%	100%	1
Twist + Shake	0%	100%	100%	1
Fling + Quasi	0%	100%	100%	1.5
Shake + Quasi	0%	100%	100%	1
Twist + Quasi	0%	100%	100%	1

V. Discussion

Our results highlight the challenges of unfolding packed medical gowns, where electrostatic forces and tightly folded layers make it difficult to fully open the garment. While the dynamic motions helped to separate some layers, none of them were sufficient to achieve a fully unfolded state. Although in theory higher-velocity motions could overcome the electrostatic forces, these high-speed motions would be undesirable in healthcare environments, where safety in the human interactions plays a pivotal role. Nevertheless, the learned primitives are able to succeed in the opening of previously opened gowns. This suggests that these motions would suffice for unfolding other types of garments, such as t-shirts or trousers, which lack long sleeves that may become tangled when folded.

Our experiments combining multiple motions show that this combination is crucial for achieving a partly opened state. In certain settings, such as assisting nurses in dressing, the partly opened state can be enough for facilitating gown placement and reducing the nurses burden in repetitive tasks. Alternative strategies such as air-based manipulation [23] could separate the fabric layers by using airflow actions. However, this type of actions introduce a significant problem in healthcare settings, since airflow actions could interfere with sterile conditions. For this reason, although the combination of lower-acceleration motions such as the twist with quasi-static motions are not sufficient for reaching the gown opened configuration, they are an effective and safe approach for healthcare environments.

VI. Conclusions

In this work, we introduced the task of pre-dressing, the step of unfolding medical garments starting in a folded configuration before assisted dressing. We evaluated three manipulation primitives for pre-dressing: fling, shake, and twist, learned from human demonstrations through the DMP framework. To assess the garment opening, we introduced three categories: closed, partly opened, and opened, and trained a visual classifier for state recognition.

Our experiments evaluated the effectiveness of the learned primitives in unfolding medical gowns, both individually, in combination with each other, as well as with a quasi-static motion. Our results showed that individual primitives could effectively unfold previously opened gowns in few iterations, while they required more attempts to partly open unpacked gowns. Here, the combination of different primitives proved more efficient to reach the partly opened state. While our results showed that neither high- nor low-acceleration motions could break the electrostatic forces to unfold unpacked gowns, their combination was more effective for reaching a sufficient partly opened state for assisted dressing.

Future work could explore combining the dynamic and quasi-static motions presented here with visual-tactile methods or re-grasping strategies to enhance the unfolding success. Building on this, we aim to integrate the pre-dressing step into a robotic-assisted dressing pipeline, with the broader objective of supporting nurses in hospital settings.

References

[1] F. Zhang and Y. Demiris, "Learning garment manipulation policies toward robot-assisted dressing," *Science Robotics*, vol. 7, no. 65, 2022.

[2] J. Zhu, A. Cherubini, C. Dune, D. Navarro-Alarcon, F. Alambeigi, D. Berenson, F. Ficuciello, K. Harada, J. Kober, X. Li, J. Pan, W. Yuan, and M. Gienger, "Challenges and outlook in robotic manipulation of deformable objects," *IEEE Robot. & Automat. Mag.*, vol. 29, no. 3, pp. 67–77, 2022.

[3] A. Longhini, Y. Wang, I. Garcia-Camacho, D. Blanco-Mulero, M. Moletta, M. Welle, G. Alenyà, H. Yin, Z. Erickson, D. Held, J. Borràs, and D. Kragic, "Unfolding the literature: A review of robotic cloth manipulation," *Annual Review of Control, Robotics, and Autonomous Systems*, 2024.

[4] H. Ha and S. Song, "Flingbot: The unreasonable effectiveness of dynamic manipulation for cloth unfolding," in *Proc. of the 5th Conf. Robot Learn.*, vol. 164. PMLR, 08–11 Nov 2022, pp. 24–33.

[5] Z. Sun, Y. Wang, D. Held, and Z. Erickson, "Force-constrained visual policy: Safe robot-assisted dressing via multi-modal sensing," *IEEE Robot. Automat. Lett.*, vol. 9, no. 5, pp. 4178–4185, 2024.

[6] S. Kotsovolis and Y. Demiris, "Garment diffusion models for robot-assisted dressing," *IEEE Robot. Automat. Lett.*, vol. 10, no. 2, pp. 1217–1224, 2025.

[7] X. Lin, Y. Wang, Z. Huang, and D. Held, "Learning visible connectivity dynamics for cloth smoothing," in *Proc, of the 5th Conf. Robot Learn.*, ser. PMLR, vol. 164. PMLR, 08–11 Nov 2022, pp. 256–266.

[8] Z. Huang, X. Lin, and D. Held, "Mesh-based Dynamics with Occlusion Reasoning for Cloth Manipulation," in *Proc. of Robots: Sci. Syst. (RSS)*, New York City, NY, USA, June 2022.

[9] D. Blanco-Mulero, G. Alcan, F. J. Abu-Dakka, and V. Kyrki, "QDP: Learning to sequentially optimise quasi-static and dynamic manipulation primitives for robotic cloth manipulation," in *2023 IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023.

[10] A. Sidiropoulos, D. Papageorgiou, and Z. Doulgeri, "A novel framework for generalizing dynamic movement primitives under kinematic constraints," *Autonomous Robots*, vol. 47, no. 1, pp. 37–50, 2023.

[11] L. Y. Chen, B. Shi, D. Seita, R. Cheng, T. Kollar, D. Held, and K. Goldberg, "Autobag: Learning to open plastic bags and insert objects," in *2023 IEEE Int. Conf. Robot. Automat.*, 2023.

[12] N. Gu, R. He, and L. Yu, "Learning to unfold garment effectively into oriented direction," *IEEE Robot. Automat. Lett.*, vol. 9, no. 2, pp. 1051–1058, 2024.

[13] E. Hannus, T. N. Le, D. Blanco-Mulero, and V. Kyrki, "Dynamic manipulation of deformable objects using imitation learning with adaptation to hardware constraints," in *2024 IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2024.

[14] L. Y. Chen, B. Shi, R. Lin, D. Seita, A. Ahmad, R. Cheng, T. Kollar, D. Held, and K. Goldberg, "Bagging by learning to singulate layers using interactive perception," in *2023 IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023.

[15] Y. Wang, Z. Sun, Z. Erickson, and D. Held, "One Policy to Dress Them All: Learning to Dress People with Diverse Poses and Garments," in *Proc. of Robots: Sci. Syst.*, 2023.

[16] J. Zhu, M. Gienger, G. Franzese, and J. Kober, "Do you need a hand? – a bimanual robotic dressing assistance scheme," *IEEE Transactions on Robotics*, vol. 40, pp. 1906–1919, 2024.

[17] F. Zhang and Y. Demiris, "Visual-tactile learning of garment unfolding for robot-assisted dressing," *IEEE Robot. Automat. Lett.*, vol. 8, no. 9, pp. 5512–5519, 2023.

[18] H. Zhang, J. Ichnowski, D. Seita, J. Wang, H. Huang, and K. Goldberg, "Robots of the lost arc: Self-supervised learning to dynamically manipulate fixed-endpoint cables," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 4560–4567.

[19] D. Blanco-Mulero, Y. Dong, J. Borras, F. T. Pokorny, and C. Torras, "T-DOM: A taxonomy for robotic manipulation of deformable objects," *arXiv preprint arXiv:2412.20998*, 2024.

[20] D. Blanco-Mulero, O. Barbany, G. Alcan, A. Colomé, C. Torras, and V. Kyrki, "Benchmarking the sim-to-real gap in cloth manipulation," *IEEE Robot. Automat. Lett.*, pp. 1–8, 2024.

[21] S. Foix, J. Oriol, J. Borras, and P. Grosch, "Gripper and procedure for the separation of layers," 2024, filled EU patent ref EP24382974.4.

[22] G. Jocher, J. Qiu, and A. Chaurasia, "Ultralytics YOLOv11," Jan. 2023.

[23] Z. Xu, C. Chi, B. Burchfiel, E. Cousineau, S. Feng, and S. Song, "Dextairity: Deformable manipulation can be a breeze," in *Proc. of Robots: Sci. Syst. (RSS)*, 2022.