

Robust Body Exposure (RoBE): A Graph-based Dynamics Modeling Approach to Manipulating Blankets over People

Kavya Puthuveetil¹, Sasha Wald¹, Atharva Pusalkar¹, Pratyusha Karnati², and Zackory Erickson¹

Abstract—Robotic caregivers could potentially improve the quality of life of many who require physical assistance. However, in order to assist individuals who are lying in bed, robots must be capable of dealing with a significant obstacle: the blanket or sheet that will almost always cover the person’s body. We propose a method for targeted bedding manipulation over people lying supine in bed where we first learn a model of the cloth’s dynamics. Then, we optimize over this model to uncover a given target limb using information about human body shape and pose that only needs to be provided at run-time. We show how this approach enables greater robustness to variation relative to geometric and reinforcement learning baselines via a number of generalization evaluations in simulation and in the real world. We further evaluate our approach in a human study with 12 participants where we demonstrate that a mobile manipulator can adapt to real variation in human body shape, size, pose, and blanket configuration to uncover target body parts without exposing the rest of the body. Source code and supplementary materials are available online³.

Index Terms—Physically Assistive Devices, Physical Human-Robot Interaction, Model Learning for Control

I. INTRODUCTION

FOR individuals who have substantial mobility impairments, robots that can provide physical assistance with activities of daily living (ADLs) hold significant promise to improve their independence, agency, and quality of life. Many people require assistance while in bed, often because their disease or injury prevents them from getting out of bed independently. In order to provide care in bed, the first, natural step is to manipulate and uncover blankets that generally will be covering the person’s body, which human caregivers are able to do with ease. It is therefore critical for robots to be able to manipulate blankets to uncover parts of the body and enable other assistive tasks to be performed. Robotic caregivers should be able to ensure the privacy and comfort of the person they are assisting by exposing individual limbs or

Manuscript received: March, 21, 2023; Revised June, 13, 2023; Accepted July, 7, 2023.

This paper was recommended for publication by Editor Markus Vincze upon evaluation of the Associate Editor and Reviewers’ comments. This work was supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE2140739 and by the NVIDIA Academic Hardware Grant.

¹ Kavya Puthuveetil, Sasha Wald, Atharva Pusalkar, and Zackory Erickson are with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA kavya@cmu.edu

² Pratyusha Karnati is with Google X, Everyday Robots, CA, USA

³ <https://rchi-lab.github.io/robust-body-exposure/>

Digital Object Identifier (DOI): see top of this page.

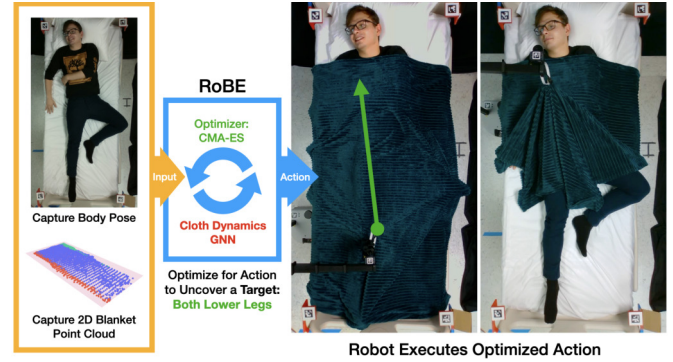


Fig. 1. Using RoBE to expose both lower legs of a participant in the real world using a Stretch RE1 mobile manipulator.

small sections of the body necessary for the completion of the task, just as human caregivers are trained to do [1], instead of removing the blanket entirely.

Previous work has considered reinforcement learning (RL) and self-supervised learning approaches for targeted bedding manipulation around people [2]. However, these model-free methods make several assumptions that limit their practicality and performance. For example, the robot must have a separate model-free bedding manipulation policy for each body part that we aim to uncover. On top of this, any changes to the action space based on the workspace of the robot requires each policy to be retained, which is an expensive process. Furthermore, prior approaches do not include any explicit representation of the cloth so variation in initial cloth state, or in other environmental factors such as human body shape, size, or pose, can degrade performance.

In this work, we propose a new method for bedding manipulation, which we call Robust Body Exposure (RoBE), that explicitly reasons about the blanket over the human body, allowing for greater flexibility and improved generalization to variation in the environment. Using data collected in simulation, we first learn a cloth dynamics model that takes in graphs of voxelized point clouds as input. Then, we plan over this model using a representation of the human’s pose and body shape in order to uncover a given target, both of which only need to be provided at run-time.

We show the robustness of RoBE when compared to geometric and reinforcement learning (RL) baselines using a number of in-distribution and generalization evaluations. We further compare our method to these baselines in the real world, deploying all approaches on a real robot to manipulate blankets over a manikin in a hospital bed. Lastly, we conduct

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

a human study with 12 participants to demonstrate RoBE’s ability to adapt to both a robot’s physical operating constraints and to challenging variations in real human body shape and pose, thus enabling us to uncover a variety of body targets in the real world without additional training on real-world data. For example, Fig. 1 shows the robot executing a trajectory to uncover both of the lower legs.

Through this work, we make the following contributions:

- We introduce a data-efficient and robust approach for targeted bedding manipulation around people that leverages graph-based cloth dynamics modeling.
- Through a series of evaluations, we quantitatively show that our method is more generalizable than existing reinforcement learning and geometric baselines, and can be easily transferred to real world robots.
- We present a human study using a mobile manipulator to uncover real people lying in bed. In this study, we demonstrate the effectiveness of our simulation-trained models in manipulating bedding in the real world, subject to real variation of human body shape, body size, and pose, as well as blanket configuration.

II. RELATED WORK

A. Data-Driven Cloth Manipulation

Early development of vision-based cloth manipulation strategies focused on leveraging traditional perception-based algorithms to identify cloth features, like corners, wrinkles, and hems, to inform manipulation [3], [4], [5], [6], [7], [8]. Since then, a number of data-driven approaches have emerged to improve generalization to different cloth shapes and configurations. These learning-based methods can be broadly defined as either model-free or model-based.

Model-free policies are often trained using imitation learning [9], [10], reinforcement learning [11], [12], [13], [2], or learning a value function [14] to complete smoothing or folding tasks. While these policies can achieve reasonable task performance on conditions in-distribution, they can struggle to generalize to unseen variation in cloth state, etc.

By contrast, model-based approaches are instead centered on learning a dynamics model of the cloth that can then be used for planning. Some previous work has attempted to learn visual cloth dynamics from RGB-D images [15], [16]. Contrastive predictive modeling has also been explored for learning a latent representation of cloth dynamics [17]. Recent advances have implemented particle or mesh-based dynamics models, which, unlike the previous approaches, explicitly reason about the topology of the cloth [18], [19], [20], [21], [22]. While model predictive control (MPC) is used to plan over dynamics models in many of the above approaches, Huang et al. demonstrate how test-time optimization can be used to plan over a mesh-based dynamics model [22], a principle that we use in our approach.

B. Cloth Manipulation for Caregiving In and Around Bed

Despite the great need for robots that can provide assistance to individuals while lying down in bed, previous work in this area is sparse. Kapusta et al. developed an interface

that enabled a user with quadriplegia to teleoperate a robot to manipulate a blanket down from the user’s knees to his feet [23]. While making an empty bed is not a physically assistive task, it is still a common home task that people with disabilities may require assistance with. Seita et al. tackle this problem with a self-supervised learning approach that uses depth images to select pick-and-place actions to pull a blanket over a bed [8]. In our own previous work, we began investigating how blankets could be manipulated over a person lying in bed to uncover target parts of their body via RL and self-supervised learning [2]. Seeking to address the shortcomings of our previous approach, including its limited capacity to generalize to a variety of new conditions and expensive training process, we introduce a new formulation of the bedding manipulation task that optimizes over a dynamics model to find bedding manipulation trajectories.

III. MODEL-BASED BEDDING MANIPULATION

In this section, we describe our dynamics modeling approach, including our action space, graph representation of cloth, and methodology for collecting data from simulation for model training. We then establish an objective function and discuss how it is used while optimizing over a trained dynamics model to complete the bedding manipulation task. Lastly, we present our set-up for evaluations in the real world. Our approach is summarized in Fig. 2.

A. Dynamics Modeling

We model cloth interactions for the bedding manipulation task as 2D, linear pick and place trajectories. These robot actions \mathbf{a} are 4D vectors containing the 2D position of the grasp (pick) location, $\mathbf{a}_g = (a_{g,x}, a_{g,y})$, and release (place) location $\mathbf{a}_r = (a_{r,x}, a_{r,y})$. The action space A , where a given action $\mathbf{a} = (a_{g,x}, a_{g,y}, a_{r,x}, a_{r,y}) \in A$, is constrained to a bounded planar region corresponding to the top surface of a standard hospital bed such that $a_{g,x}, a_{r,x} \in (-0.44 \text{ m}, 0.44 \text{ m})$ and $a_{g,y}, a_{r,y} \in (-1.05 \text{ m}, 1.05 \text{ m})$.

We represent the state of the cloth as a graph $G = (V, E)$ where the nodes $V = \{v_i\}_{i=1\dots N}$ represent particles that compose the cloth and edges $e_{ij} \in E$ represent the connectivity of two nearby particles v_i and v_j on the cloth. For every cloth particle $\mathbf{p}_i \in P'$ in the voxelized point cloud P' we create a node $v_i \in V$ that encodes the particle’s 2D position, $v_i = \mathbf{p}_i$. A particular edge, $e_{ij} \in E$, exists between given nodes v_i and v_j if the Euclidean distance between nodes is less than some distance threshold, $d: \|n_i - n_j\|_2 < d$. We also provide a candidate action \mathbf{a} used to manipulate the cloth as a global vector, which is appended to each node. We encode a 2D projection of a cloth particle’s position in the graph instead of its 3D position to ensure depth invariance of models trained using this graph representation. In an evaluation presented on our webpage, we find that both 2D and 3D cloth dynamics models perform similarly when used to manipulate bedding.

Graph neural networks (GNNs) have been well demonstrated for effectively learning cloth dynamics to predict how a robot’s actions will perturb a cloth state, particularly using the standard Graph Network-based Simulators (GNS) architecture [19], [21], [24]. In this work, we build a GNS-based

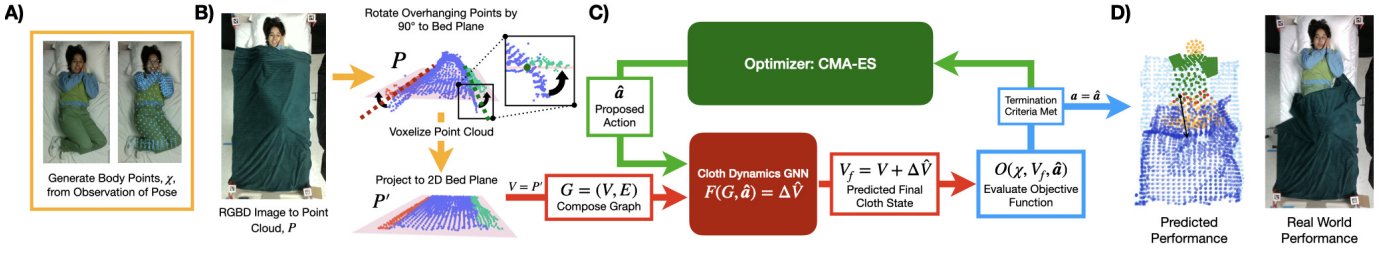


Fig. 2. Overview of our approach: A) We generate a set of body points from an observation of a person’s pose. B) We cover them with a blanket, then capture and process an initial cloth point cloud for composition of a graph. C) Given an input graph, we run an optimization loop over the dynamics model to find an action that uncovers a target body part. D) We generate an initial prediction of performance before executing the best action in the real world.

GNN with 4 processing layers that acts as a dynamics model $F(G, \mathbf{a}) = \Delta \hat{V}$ that, given a cloth graph G and a candidate action \mathbf{a} , outputs the predicted displacement of all nodes in the graph after execution of the action $\Delta \hat{V}$.

B. Dynamics Model Training in Simulation

We train our GNN in simulation with the PyBullet physics engine [25], which models the physics of planar deformables like cloth as a mass-spring system [26]. We interface with PyBullet via Assistive Gym [27], a simulation framework for physically assistive robotics. We utilize the same bedding manipulation environment presented in our prior work [2], in which we allow a capsulized human model of fixed body size to settle on a hospital bed in a supine pose with some variation in its joints positions. A blanket, represented by a $1.25\text{m} \times 1.7\text{m}$ grid mesh with mass, friction, spring damping, and spring elasticity tuned to match a real-world blanket, is then dropped onto the bed from a fixed position so that most of the human body, with the exception of the head and neck, will be covered. We treat the vertices of the ground truth blanket mesh as points in a point cloud P .

To interact with the cloth given an action \mathbf{a} , we first anchor a spherical manipulator, representing a robot’s end effector, to the closest cloth point p^* to the grasp point \mathbf{a}_g . We find the cloth point p^* via $p^* = \operatorname{argmin}_{p' \in P \odot (1, 1, 0)} \|p' - \mathbf{a}_g\|_2$, where p' represents the 2D position of a cloth point in point cloud P after being projected to a 2D plane, defined by the top surface of the hospital bed, via the Hadamard product, $P \odot (1, 1, 0)$. The manipulator then lifts the cloth off of the human by translating upwards to a fixed height of 40cm to avoid collision with the human body while executing an action. This fixed lift distance corresponds to the maximum lift height of the robot used in our real-world evaluations, described in Section III-E. After lifting the cloth, the manipulator moves linearly between the grasp location, \mathbf{a}_g , to the release location, \mathbf{a}_r , before finally releasing the manipulator-cloth anchor to drop the cloth. An example of a single simulation rollout, from environment setup to trajectory execution is pictured in Fig. 3.

To collect training data for our dynamics model, we uniformly sample linear trajectories from the action space A to perturb the cloth in a given simulation rollout. If a randomly sampled action contains a grasp point \mathbf{a}_g that is not over the cloth (i.e., there are no cloth points p within $\lambda = 2.8\text{cm}$ of \mathbf{a}_g such that $\exists p' \in P \odot (1, 1, 0) : \|\mathbf{a}_g - p'\|_2 < \lambda$), we re-sample actions until the grasp point is valid. We collect a dataset of size 10,000 (P, \mathbf{a}, P_f) pairs where P is the initial cloth point




Fig. 3. A simulation rollout: 1) We initialize a human model above the bed. 2) After the human model settles on the bed, we initialize a blanket above them. 3) We drop the blanket onto the human model and the end effector picks up the blanket at the grasp location. 4) The blanket is lifted upwards before being moved to the release location.

cloud, \mathbf{a} is the action, and P_f is the final cloth point cloud after the action is executed.

For each initial cloth state and action pair in the random interaction dataset, we compose a graph $G = (V, E)$ as described in Section III-A, where the nodes V correspond to points in a processed cloth point cloud P' . To compute P' from a given raw cloth point cloud P , we first rotate blanket points that hang over the side of the bed up to the bed plane. Rotating overhanging points, as visualized in Fig. 2-B, allows us to more accurately retain the full geometry of the blanket in the graph, which only encodes the 2D position of any particular cloth point. Our methodology for applying this rotation is described in detail on our webpage. After rotating the overhanging cloth points, we downsample the point cloud using a 5 cm centroid voxel grid filter [28]. Lastly, we project all of the cloth points to the 2D bed plane, yielding the processed point cloud P' .

We train a dynamics model $F(G, \mathbf{a}) = \Delta \hat{V}$, based on the GNS architecture, that given an input graph G that represents a cloth’s initial state and a proposed action \mathbf{a} , predicts the displacement of the cloth $\Delta \hat{V}$ after execution of a cloth manipulation trajectory. We train our models over 250 epochs on 10,000 graphs—enough training data for performance to saturate—using the Adam optimizer with a mean squared error loss function, batch size of 100, and learning rate of $1\text{e-}4$. We evaluate our dynamics models when used in the RoBE framework for the bedding manipulation task in Section IV.

C. Determining Covered Body Parts

To assess whether areas on the body are covered or uncovered, we represent the human body using a set of 384 discrete, uniformly spaced, points χ . To generate this set of points, we capture a 28D observation of the human’s pose, \mathbf{s} , defined by the 2D positions of 14 different joints. We also use a set of nine body shape measurements that includes the radius of several body segments. The 14 joint positions and the

nine body shape measurements we use are pictured in Fig. 4. We uniformly distribute points across the area of each body segment, as calculated using their length and radius/width, to build our representation of the human.

From a set of human body points χ , we can determine whether a given body point $x \in \chi$ is covered or uncovered by the final state of the cloth $V_f = V + \Delta V$ via the function:

$$C(x, V_f) = \begin{cases} 0 & \text{if } \exists v_f \in V_f : \|x - v_f\|_2 < \lambda \\ 1 & \text{otherwise} \end{cases} \quad (1)$$

which states that a body point $x \in \chi$ is considered covered ($C(x, V_f) = 0$) if there are any cloth points within $\lambda = 2.8\text{cm}$ of it, otherwise it is considered uncovered ($C(x, V_f) = 1$). The threshold λ was tuned manually in previous work [2].

D. Optimizing over Dynamics Models

Following training, we use our dynamics models to predict how the state of a cloth will change when manipulated using a proposed action. Compared to physics simulation, forward passes of our GNNs can be run within a fraction of a second. This enables them to be used in optimizing for actions that generate favorable cloth states in close to real time. Favorable cloth states uncover a target body part without exposing the rest of the body. We define ten discrete targets: *the right arm, right lower legs, right legs, left arm, left lower legs, left legs, both lower legs, upper body, lower body, and entire body*. The set of all body points χ can be further divided by whether they are target body points χ_t , non-target body points χ_n , or points on the head χ_h , such that $\chi_t, \chi_n, \chi_h \subseteq \chi$

In order to optimize for actions that uncover the body in a targeted way, we define an objective function:

$$O(\chi, V_f, \mathbf{a}) = \alpha O_t(\chi_t, V_f) + \phi O_n(\chi, V_f) + O_h(\chi_h, V_f) + O_d(\mathbf{a})$$

These terms represent, in order, a reward for uncovering target points and penalties for uncovering non-target points, covering the head, or choosing excessively long actions. $\alpha = 100$ and $\phi = 100$ are scalar weights that can be used to modify the reward for uncovering target points and penalty for uncovering non-target points, respectively.

The reward for uncovering target body points is given by: $O_t(\chi_t, V_f) = -\alpha (\rho_t / |\chi_t|)$ where $\rho_t = \sum_{x \in \chi_t} C(x, V_f)$ counts target points uncovered and $|\chi_t|$ is the cardinality of set χ_t .

The penalty for uncovering non-target body points is:

$$O_n(\chi, V_f) = \sum_{x \in \chi_n} C(x, V_f) * \Gamma(x, \chi_t) * \Xi(\chi, \chi_t) \quad (2)$$

where $\Gamma(x, \chi_t) = \text{argmin}_{x^* \in \chi_t} \|x^* - x\|_2$ is a scaling function that adjusts the penalty weight of a given non-target point $x \in \chi_n$ based on its distance to the closest target point $x^* \in \chi_t$. This term operates on the intuition that uncovering non-target points close to the target limb may be inevitable in some poses. We reason that uncovering such points should not incur as much cost as uncovering those that are farther away from the target (i.e. are theoretically easier to avoid uncovering). We also account for the increased difficulty of uncovering smaller targets with the scaling term $\Xi(\chi, \chi_t) = \frac{|\chi_t|}{0.05 * (|\chi_n| + |\chi_h|)}$, which further adjusts the penalty for uncovering a non-target body




Fig. 4. Capturing human body shape parameters and pose in the real world: 1) Measuring each of the nine body segments, 2) Adjusting the estimated position of the 14 joints in observation \mathbf{s} to better match the ground truth pose, 3) visualization of all the virtual points generated along the body.

point based on the relative size of the target, defined as the ratio between the number of target points and the number of non-target and head points.

We also include a heavy penalty for covering parts of the head, $O_h(\chi_h, V_f) = 2\alpha (\rho_h / |\chi_h|)$ where $\rho_h = |\chi_h| - \sum_{p \in \chi_h} C(p, V_f)$, since doing so is generally against human preferences. Lastly, $O_d(\mathbf{a}) = 150$ is a penalty applied when the Euclidean distance between grasp and release points, $\|\mathbf{a}_r - \mathbf{a}_g\|_2$, is greater than 1.5m. Otherwise, $O_d(\mathbf{a}) = 0$.

With a trained dynamics model and an objective function, we can introduce an optimizer, in our case covariance matrix adaptation evolution strategy (CMA-ES) [29], to identify optimal bedding manipulation trajectories to uncover people in bed in real time. Given an initial processed cloth point cloud P' , on each iteration of CMA-ES, a graph G is constructed with a new sampled action, $\hat{\mathbf{a}}$, and then fed into the dynamics model to predict the displacement of the cloth $\Delta \hat{V}$. We then compute the final cloth state $V_f = V + \Delta \hat{V}$, construct the set of all points along the body χ from an observation of human pose, define $\chi_t, \chi_n, \chi_h \subseteq \chi$ for a specific target limb, and finally evaluate objective function $O(\chi, V_f, \hat{\mathbf{a}})$ before re-sampling a new set of actions based on the cost value. We allow CMA-ES to optimize over the action space A for 38 iterations using a population size of 8 (304 total function evaluations), or until $O(\chi, V_f, \mathbf{a}) \geq 95$. Optimization hyperparameters are further described on our webpage.

Compared to other bedding manipulation methods, our strategy of optimizing over a dynamics model allows us to uncover new targets, modify objectives, or change the importance of sub-objectives, all in real time without needing to retrain. For example, we can, without retraining, restrict the optimizer search space to a given robot's workspace, which we demonstrate in Section III-E, or modify weights α and ϕ in $O(\chi, V_f, \mathbf{a})$ to prioritize uncovering the entirety of a target with less concern for uncovering non-target points.

E. Real World Implementation

We implemented RoBE using a Stretch RE1 mobile manipulator [30] to demonstrate its efficacy in the real world.

In order to capture a complete point cloud of the blanket, including points hanging over the sides of the bed, we use three cameras affixed above, and on the right and left sides of, the bed. After capturing point clouds from each camera



Fig. 5. Comparison of the action space used in simulation, A , and the action space used in the real world implementation, A_{RW} .

and applying color segmentation and position-based cropping on the point clouds to isolate the points on the blanket, we perform Iterative Closest Point (ICP) registration over 5e6 iterations to reconstruct a single point cloud of the blanket.

Before manually covering the human with the blanket, we capture an RGB image using a camera mounted above the center of the bed. We feed this image into BlazePose, as implemented in Google’s MediaPipe Pose [31], to capture an estimate of the human’s joint positions. Although we currently rely on pose estimation methods that require an RGB-based, unobstructed observation of the human body, future work could instead use occlusion-invariant methods that use images from a bed pressure mat [32] or from an depth camera [33]. Such methods may also help alleviate concerns associated with the use of RGB cameras in private areas. Using a custom GUI, we manually adjust pose estimates from BlazePose for better alignment with the body and also take body shape measurements, thus allowing us to generate the set of body points χ . Fig. 4 depicts this process.

Using fiducial tags fixed to each of the corners of the bed, we find the center of the real-world hospital bed, allowing us to transform the point cloud of the blanket and the human pose estimate to a planar global coordinate system comparable to that used in simulation. We apply the same point cloud preprocessing steps as in simulation, including rotation along bed edges, voxelization, and 2D projection. This allows us to construct an input graph to the dynamics model and optimize the objective function, as described in Section III-D, to find an action that uncovers a specific limb without fine-tuning. The Stretch robot’s maximum arm extension (52 cm) does not allow it to reach across a hospital bed (88 cm wide), however, our approach can easily accommodate the robot’s physical limitations by simply restricting the search space for CMA-ES to the robot’s workspace, which does not require any retraining. We also restrict the search space to prevent the robot from grasping near participants’ head or neck, which could cause discomfort. The resulting action space A_{RW} used in the real-world experiment is defined such that $a_{g,x}, a_{r,x} \in (0 \text{ m}, 0.44 \text{ m})$, $a_{g,y} \in (-0.525 \text{ m}, 1.05 \text{ m})$, and $a_{r,y} \in (-1.05 \text{ m}, 1.05 \text{ m})$, as shown in Fig. 5. To account for the robot’s limited reach, we only attempt to uncover seven of the ten targets enumerated in Section III-D in the real world, excluding the left arm, lower leg, and leg.

We provide computed actions to the robot for execution. At the beginning of a given trial, the robot uses its pan-tilt camera to observe the fiducial markers on the corners of the bed and localize itself to the real-world global coordinate system that the grasp and release points are defined with respect to. The




Fig. 6. Executing RoBE to uncover a participant’s right arm: 1) The initial pose of the participant and action to be executed. 2) The robot grasps the cloth at the grasp locations, 3) lifts the cloth, 4) linearly moves the cloth to the release location, 5) and drops the cloth. 6) The final state of the cloth and the target and non-target points that have been uncovered.

robot moves its end effector to the grasp location with its gripper closed, then lowers the gripper until it contacts the blanket, as detected by effort sensing in the robot’s actuators, to complete the grasp. It lifts up by 3cm to open its gripper, lowers to the cloth to close its gripper on the blanket, and then raises its end effector to 40cm above the bed before moving linearly to the release location to drop the blanket. Lastly, we capture an image from the above-bed camera to use in evaluation of the robot’s performance. This sequence of trial events is shown in Fig. 6. Completion of a real-world trial, from capturing pose to executing an action, takes 4-5 min, of which the robot motion takes ~ 1 min.

IV. EVALUATIONS

A. Evaluation Metrics

We perform experiments both in simulation and in the real world to assess the efficacy of our method. We define and compute the following metrics to assess task performance:

- True positive (TP) = ρ_t
- False negative (FN) = $|\chi_t| - \rho_t$
- False positive (FP) = $\gamma(\rho_n)$
- F-Score (F_1) = $\frac{TP}{TP+0.5(FP+FN)}$

TP is the number of target body points that are correctly uncovered while FN is the number of target points that remained covered. FP is a weighted representation of the number of non-target points that are uncovered, $\gamma(\rho_n)$:

$$\gamma(\rho_n) = \sum_{i=1}^{\rho_n} \begin{cases} w * i & \text{if } w * i \leq 1 \\ 1 & \text{otherwise} \end{cases}$$

where $\rho_n = \sum_{p \in \chi_n} C(p, V_f)$ is the total number of non-target points uncovered and $w = \frac{|\chi_t|}{|\chi|}$ is a weighting factor representing the size of a target limb relative to the size of the entire body. Similar to Equation 2, $\gamma(\rho_n)$ operates on the intuition that, when evaluating a bedding manipulation outcome, people are willing to accept that some non-target body points may inevitably be uncovered, like when uncovering regions of the body that are spatially close to other regions (e.g. uncovering an arm that is resting on the stomach).

Finally, F_1 is a weighted measure of the robot’s accuracy at performing bedding manipulation, where $F_1 \in [0, 1]$.

B. Baselines

We compare RoBE to two baselines, including a geometric approach and a reinforcement learning method:

1) *Geometric Baseline*: This approach is founded on two assumptions that: 1) there exists a geometric relationship between the cloth and human pose that can be leveraged to uncover a blanket from part of the body and, 2) the cloth behaves like paper—pick and place actions can produce straight-line or accordion folds. To compute actions that “fold” the cloth over the body to uncover a target limb, we apply three strategies depending on the target to be uncovered; *Arms*: grasps the edge of the cloth, pulls towards the midline of the body on a path normal to the length of the arm; *Lower Body Targets*: grasps along the midpoint of the target, pulls towards the head on a path parallel to the limb; *Upper Body and Entire Body*: grasps between the shoulders, pulls toward the feet on a path that bisects the body. We provide formal definitions of how each strategy computes these pick and place trajectories on the project webpage.

2) *Reinforcement Learning*: For this baseline, we use proximal policy optimization (PPO), a model-free deep RL method, to train policies to uncover a single target limb as demonstrated in prior work [2]. We train a single policy for each of the ten target limbs, using $-O(\chi, V_f, a)$ as a reward function and the 28D representation of human pose \mathbf{s} as the observation, similar to prior work [2]. Running PPO in the real world required the expensive process of training a second set of policies for each target limb, this time using the action space A_{RW} .

We evaluate all methods both in simulation, described in Section IV-C, and in the real world with a medical manikin, described in Section IV-D.

C. Evaluating Generalization in Simulation

We conduct all simulation evaluations over 500 rollouts where both the environment’s state and the target limb are randomized over the following generalization conditions:

1) *Training (Training Dist.)*: The model training conditions (Section III-B), including some small human pose variation with fixed body size and initial blanket configuration.

2) *Human Body Shape and Size Variation (Body Dist.)*: For this condition, we introduce large variation to the human body shape and size by fitting the capsulized human model to a SMPL-X body mesh [34], defined using 10 uniformly sampled body shape parameters $\beta \in U(0,4)$, before dropping the model onto the hospital bed as described in [2].

3) *High Human Pose Variation (Pose Dist.)*: We apply additional pose variation beyond that used in the training distribution by adding a δ to the joint angles of both of the capsulized model’s elbows and knees. For each of these four joints, this δ value is uniformly sampled from the range $(\theta_{min} + 1 \text{ rad}, \theta_{max} - 1 \text{ rad})$, where θ_{min} and θ_{max} are the limits of a given joint’s range of movement [27].

4) *Initial Blanket Variation (Blanket Dist.)*: We modify the initial blanket configuration by introducing variation uniformly sampled between $\pm 2\text{cm}$ to its x position, $[-25, 5]$ cm to its y position, and ± 45 degrees to its yaw θ_z orientation before dropping the blanket onto the human.

TABLE I
GENERALIZATION PERFORMANCE (F_1) IN SIMULATION AND IN A REAL-WORLD MANIKIN STUDY

Evaluation Conditions	Approaches		
	RoBE	PPO	Geometric
Training Dist.	0.77 ± 0.23	0.25 ± 0.40	0.41 ± 0.23
Body Dist.	0.72 ± 0.25	0.31 ± 0.40	0.43 ± 0.23
Pose Dist.	0.67 ± 0.29	0.26 ± 0.38	0.39 ± 0.26
Blanket Dist.	0.72 ± 0.25	0.27 ± 0.39	0.39 ± 0.25
Combo. Dist.	0.65 ± 0.30	0.33 ± 0.38	0.41 ± 0.31
Manikin Study	0.81 ± 0.12	0.23 ± 0.32	0.63 ± 0.16

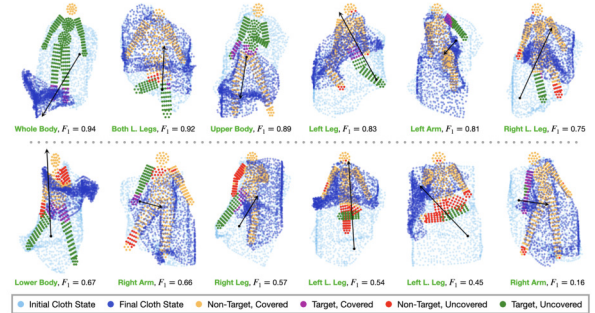


Fig. 7. Representative rollouts of our approach when evaluated on the combination distribution.

5) *Combination (Combo. Dist.)*: We use all variations from the Body, Pose, and Blanket Dists., resulting in challenging environmental states relative to the Training Dist.

Performance of our approach and both baselines on each generalization condition is summarized in Table I. We find that across all generalizations, our method achieves the highest F-scores, followed by the geometric approach, and PPO respectively. The poor performance of PPO policies can be attributed to them often selecting trajectories that induce little to no displacement of the cloth (e.g. not grasping the cloth at all, picking up and releasing the cloth in effectively the same place). The geometric approach does slightly better, but its predefined strategies limit its ability to generalize to a larger variety of poses and blanket configurations.

By contrast, our approach leverages a learned model that is close to the cloth’s true dynamics when optimizing for actions. RoBE’s F-scores vary from 0.65-0.72 on generalization conditions compared to 0.77 on the training distribution. Fig. 7 shows representative samples of how RoBE performs on the combination distribution and provides context to how F-score varies depending on the manipulation outcome. We observe decreased dynamics model accuracy when the cloth’s dynamics are more complex (Pose, Combo. Dists.), thus degrading task performance. Future work may consider how reasoning about the underlying connectivity of the blanket can improve performance in these cases [21], [22]. As shown in Fig. 7, we also note that some target limbs are near impossible to uncover in isolation due to the human body shape, pose, and initial blanket state.

D. Manikin Study

We compare our approach to baselines in the real world by deploying the robot to uncover target limbs of a medical

TABLE II
PERFORMANCE (F_1) PER LIMB IN THE REAL WORLD AND SIMULATION

Target Body Part	Human Study		Simulation
	Predefined	Randomized	Combo Dist.
Right Arm	0.77±0.29	0.70±0.26	0.50±0.32
Right Lower Leg	0.95±0.05	0.51±0.33	0.57±0.29
Upper Body	0.88±0.05	0.89±0.12	0.87±0.12
Right Leg	—	0.51±0.24	0.52±0.22
Both Lower Legs	—	0.67±0.27	0.77±0.18
Lower Body	—	0.61±0.28	0.84±0.10
Entire Body	—	0.87±0.05	0.96±0.09
Overall (Mean)	—	0.68±0.27	0.65±0.30

manikin lying in the hospital bed. In the study, we evaluate all three methods for uncovering each of the seven target limbs on a set of three poses (3 methods × 3 poses × 7 targets = 63 total trials). Table I summarizes the performance of each approach in the real world over the manikin. These results are generally consistent with what we observed in simulation, with RoBE ($F_1 = 0.81$) outperforming the geometric approach by 0.18 and PPO by 0.58.

E. Human Study

To evaluate the performance of our method in the real world, we run a human study (Carnegie Mellon University IRB Approval under 2022.00000351) with 12 participants (six female) with informed consent. For each participant, we demonstrate ten bedding manipulation trials (120 trials in total across participants). The ten trials are executed in two phases. The first phase consists of three trials where the participant is asked to assume a predefined pose, shown in the leftmost column of Fig. 8, and the robot attempts to uncover a predefined target body segment. By holding the pose and target constant, we examine the performance of our method given variation in body shape and size alone. The remaining seven trials are a part of Phase 2, in which we ask participants to assume any supine pose of their choosing and we randomly select a target body part to uncover. This second phase allows us to evaluate the robustness of our approach to a large set of unscripted resting poses. Since our approach operates under the assumption that the human’s pose is static, we ask participants to stay still during a given trial.

The mean F-score achieved by our method across target limbs in each phase of the human study is summarized in Table II. For comparison, we provide the detailed generalization results of our approach in simulation using the Combo. Dist. A representative selection of trials from Phase 1 and Phase 2 are visualized in Fig. 8 and Fig. 9, respectively.

In the Phase 1 trials, where both pose and target limbs were predefined, we observe strong performance when uncovering the right lower leg and upper body, with high mean F-scores (≥ 0.88) and low standard deviations (0.05), indicating that we can reliably uncover these targets despite variation in human body shape and size. We see lower performance ($F_1 = 0.77$) in uncovering the right arm. The increased standard deviation relative to the other two targets suggests that the right arm is a more challenging target to uncover for some body sizes.




Fig. 8. Representative performance from the first phase of the human study where both the target limb and participant’s pose are predefined. Images of the predefined poses shown to participants are in the leftmost column. Results are ordered in decreasing F-score.

Yet, overall we observe that the generalization of our approach across real variation in body shape and size is promising.

In Phase 2, we evaluated our approach on a variety of challenging scenarios where participants selected their pose and we randomized the target. We observed high performance consistent with simulation for most body targets.

Of the failure cases we observed, we found that many were less due to errors in our model and more attributable to the robot failing to complete a given trajectory. Topographical features created by the body that are normal to the bed’s surface, such as upturned feet or knees bent upwards, could generate high frictional forces as the robot attempts to drag the blanket over them. For longer trajectories, these forces can become so high that they can keep the robot from advancing forward or can pull the blanket out of the robot’s gripper. A failure case demonstrating an example of this phenomenon is shown the bottom right in Fig. 9.

Although using a robot that can pull at higher forces could potentially resolve this problem, there may be safety concerns with doing so. Instead, future work may consider incorporating feedback from estimated force and contact maps, as has been previously demonstrated for robot-assisted dressing [35], [36], [24], to modify or inform bedding manipulation trajectories. Aside from some limitations with our mobile manipulator, the results achieved in the human study with our current approach remain promising towards achieving robust bedding manipulation in the real world.

V. CONCLUSION

In this work, we leverage gradient-free optimization over simulation-trained graph-based cloth dynamics models to-

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.



Fig. 9. Example images from the second phase of the study where participants chose their own supine poses and the target limb was randomized. Results are ordered in decreasing F-score.

form targeted bedding manipulation over people. One of several benefits of this strategy is the capability to change the objective function and goals in real time without retraining. Compared to geometric and reinforcement learning baselines, our approach is robust and generalizable, as demonstrated by a battery of evaluations in simulation and in the real world with a medical manikin. In a human study, we also demonstrate RoBE's efficacy when performing targeted bedding manipulation over real people. We observed that our approach had promising performance uncovering targeted body regions across different human body sizes, variations in human pose, and different blanket configurations.

REFERENCES

- [1] P. Lynn and C. Lynn, *Hygiene: Giving a Bed Bath*, 3rd ed. Wolters Kluwer/Lippincott Williams & Wilkins Health, 2011, p. 316–325.
- [2] K. Puthuveetil, C. C. Kemp, and Z. Erickson, "Bodies uncovered: Learning to manipulate real blankets around people via physics simulations," *IEEE Robotics and Automation Letters*, 2022.
- [3] L. Sun, G. Aragon-Camarasa, S. Rogers, and J. P. Siebert, "Accurate garment surface analysis using an active stereo robot head with application to dual-arm flattening," in *IEEE ICRA*, 2015.
- [4] A. Ramisa, G. Alenyà, F. Moreno-Noguer, and C. Torras, "Using depth and appearance features for informed robot grasping of highly wrinkled clothes," in *IEEE ICRA*, 2012.
- [5] K. Yamazaki, "Grasping point selection on an item of crumpled clothing based on relational shape description," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014.
- [6] H. Yuba, S. Arnold, and K. Yamazaki, "Unfolding of a rectangular cloth from unarranged starting shapes by a dual-armed robot with a mechanism for managing recognition error and uncertainty," *Advanced Robotics*, vol. 31, no. 10, pp. 544–556, 2017.
- [7] J. Qian, T. Weng, L. Zhang, B. Okorn, and D. Held, "Cloth region segmentation for robust grasp selection," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020.
- [8] D. Seita, N. Jamali, M. Laskey, A. K. Tanwani, R. Berenstein, 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] D. Seita, A. Ganapathi, R. Hoque, M. Hwang, E. Cen, A. K. Tanwani, A. Balakrishna, B. Thananjeyan, J. Ichnowski, N. Jamali, et al., "Deep imitation learning of sequential fabric smoothing from an algorithmic supervisor," in *IEEE/RSJ IROS*, 2020.
- [10] A. Ganapathi, P. Sundaresan, B. Thananjeyan, A. Balakrishna, D. Seita, J. Grannen, M. Hwang, R. Hoque, J. E. Gonzalez, N. Jamali, et al., "Learning dense visual correspondences in simulation to smooth and fold real fabrics," in *IEEE ICRA*, 2021.
- [11] J. Matas, S. James, and A. J. Davison, "Sim-to-real reinforcement learning for deformable object manipulation," in *Conference on Robot Learning*. PMLR, 2018, pp. 734–743.
- [12] Y. Wu, W. Yan, T. Kurutach, L. Pinto, and P. Abbeel, "Learning to manipulate deformable objects without demonstrations," in *RSS*, 2020.
- [13] R. Lee, D. Ward, V. Dasagi, A. Cosgun, J. Leitner, and P. Corke, "Learning arbitrary-goal fabric folding with one hour of real robot experience," in *Conference on Robot Learning*, 2021.
- [14] H. Ha and S. Song, "Flingbot: The unreasonable effectiveness of dynamic manipulation for cloth unfolding," in *Conference on Robot Learning*. PMLR, 2022, pp. 24–33.
- [15] F. Ebert, C. Finn, S. Dasari, A. Xie, A. Lee, and S. Levine, "Visual foresight: Model-based deep reinforcement learning for vision-based robotic control," *arXiv preprint arXiv:1812.00568*, 2018.
- [16] 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.
- [17] 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.
- [18] Y. Li, J. Wu, R. Tedrake, J. B. Tenenbaum, and A. Torralba, "Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids," in *ICLR*, 2019.
- [19] A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and P. Battaglia, "Learning to simulate complex physics with graph networks," in *Proceedings of the 37th ICML*, 2020.
- [20] T. Pfaff, M. Fortunato, A. Sanchez-Gonzalez, and P. W. Battaglia, "Learning mesh-based simulation with graph networks," *arXiv preprint arXiv:2010.03409*, 2020.
- [21] X. Lin, Y. Wang, Z. Huang, and D. Held, "Learning visible connectivity dynamics for cloth smoothing," in *CoRL*, 2021.
- [22] Z. Huang, X. Lin, and D. Held, "Mesh-based dynamics with occlusion reasoning for cloth manipulation," in *RSS*, 2022.
- [23] A. S. Kapusta, P. M. Grice, H. M. Clever, Y. Chitalia, D. Park, and C. C. Kemp, "A system for bedside assistance that integrates a robotic bed and a mobile manipulator," *Plos one*, vol. 14, no. 10, 2019.
- [24] Y. Wang, D. Held, and Z. Erickson, "Visual haptic reasoning: Estimating contact forces by observing deformable object interactions," *IEEE Robotics and Automation Letters*, 2022.
- [25] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," 2016.
- [26] J. Stępień, "Physics-based animation of articulated rigid body systems for virtual environments," Ph.D. dissertation, Silesian University of Technology, 2013.
- [27] Z. Erickson, V. Gangaram, A. Kapusta, C. K. Liu, and C. C. Kemp, "Assistive gym: A physics simulation framework for assistive robotics," in *IEEE ICRA*, 2020.
- [28] R. B. Rusu and S. Cousins, "3d is here: Point cloud library (pcl)," in *IEEE international conference on robotics and automation*, 2011.
- [29] N. Hansen, S. D. Müller, and P. Koumoutsakos, "Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (cma-es)," *Evolutionary computation*, 2003.
- [30] C. C. Kemp, A. Edsinger, H. M. Clever, and B. Matulevich, "The design of stretch: A compact, lightweight mobile manipulator for indoor human environments," in *IEEE ICRA*, 2022.
- [31] V. Bazarevsky, I. Grishchenko, K. Raveendran, T. Zhu, F. Zhang, and M. Grundmann, "Blazepose: On-device real-time body pose tracking," in *Computer Vision and Pattern Recognition*, ser. Workshop on Computer Vision for Augmented and Virtual Reality, 2020.
- [32] H. M. Clever, A. Kapusta, D. Park, Z. Erickson, Y. Chitalia, and C. C. Kemp, "3d human pose estimation on a configurable bed from a pressure image," in *IEEE/RSJ IROS*, 2018.
- [33] H. M. Clever, P. Grady, G. Turk, and C. C. Kemp, "Bodypressure-inferred body pose and contact pressure from a depth image," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [34] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. Osman, D. Tzionas, and M. J. Black, "Expressive body capture: 3d hands, face, and body from a single image," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [35] Z. Erickson, A. Clegg, W. Yu, G. Turk, C. K. Liu, and C. C. Kemp, "What does the person feel? learning to infer applied forces during robot-assisted dressing," in *IEEE ICRA*, 2017.
- [36] Z. Erickson, H. M. Clever, G. Turk, C. K. Liu, and C. C. Kemp, "Deep haptic model predictive control for robot-assisted dressing," in *IEEE International Conference on Robotics and Automation*, 2018.