



Personalized collaborative plans for robot-assisted dressing via optimization and simulation

Ariel Kapusta¹ · Zackory Erickson¹ · Henry M. Clever¹ · Wenhao Yu¹ · C. Karen Liu¹ · Greg Turk¹ · Charles C. Kemp¹

Received: 22 July 2018 / Accepted: 16 May 2019 / Published online: 26 June 2019
© The Author(s) 2019

Abstract

Robots could be a valuable tool for helping with dressing but determining how a robot and a person with disabilities can collaborate to complete the task is challenging. We present task optimization of robot-assisted dressing (TOORAD), a method for generating a plan that consists of actions for both the robot and the person. TOORAD uses a multilevel optimization framework with heterogeneous simulations. The simulations model the physical interactions between the garment and the person being dressed, as well as the geometry and kinematics of the robot, human, and environment. Notably, the models for the human are personalized for an individual's geometry and physical capabilities. TOORAD searches over a constrained action space that interleaves the motions of the person and the robot with the person remaining still when the robot moves and vice versa. In order to adapt to real-world variation, TOORAD incorporates a measure of robot dexterity in its optimization, and the robot senses the person's body with a capacitive sensor to adapt its planned end effector trajectories. To evaluate TOORAD and gain insight into robot-assisted dressing, we conducted a study with six participants with physical disabilities who have difficulty dressing themselves. In the first session, we created models of the participants and surveyed their needs, capabilities, and views on robot-assisted dressing. TOORAD then found personalized plans and generated instructional visualizations for four of the participants, who returned for a second session during which they successfully put on both sleeves of a hospital gown with assistance from the robot. Overall, our work demonstrates the feasibility of generating personalized plans for robot-assisted dressing via optimization and physics-based simulation.

Keywords Assistive robotics · Robot-assisted dressing · Manipulation

1 Introduction

Robotic assistance with activities of daily living (ADLs) (Wiener et al. 1990) could increase independence for people with disabilities. This may improve quality of life and help address societal challenges, such as aging populations, high healthcare costs, and shortages of healthcare workers found in the United States and other countries (Institute of Medicine 2008; Goodin 2003). A number of specially designed assistive devices exist to help people maintain their independence. However, many current assistive devices, such as those for dressing (e.g., reachers, dressing sticks, long-handled shoe-horns, and sock aids), provide limited support and rely on the user having substantial cognitive, perceptual, and motor capabilities (Canadian Partnership for Stroke Recovery 2015; The Wright Stuff Inc 2017). Robots could potentially serve as more versatile assistive devices.

This work was supported in part by the National Institute on Disability, Independent Living, and Rehabilitation Research (NIDILRR), Grant 90RE5016-01-00 via RERC TechSAge, and by NSF Awards IIS-1514258 and IIS-1150157. Dr. Kemp is a cofounder, a board member, an equity holder, and the CTO of Hello Robot, Inc., which is developing products related to this research. This research could affect his personal financial status. The terms of this arrangement have been reviewed and approved by Georgia Tech in accordance with its conflict of interest policies.

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s10514-019-09865-0>) contains supplementary material, which is available to authorized users.

✉ Ariel Kapusta
akapusta@gatech.edu

¹ Georgia Institute of Technology, Atlanta, USA



Fig. 1 Through our system implementation of TOORAD, four participants with disabilities were able to dress themselves with assistance from the robot. (Left-to-right) prior to dressing the left arm; end of

dressing the left arm; prior to dressing the right arm; end of dressing the right arm. After successful dressing, both of the participant's arms are in the hospital gown

Specialized robots are commercially available for a variety of ADLs (Topping and Smith 1998; Topping 1999), such as desktop feeding devices for feeding tasks, but robotic assistance for dressing remains in early stages of research. Studies suggest a need for robot-assisted dressing, with more older adults receiving assistance with dressing and bathing/showering than other ADLs, and over 80% of people in skilled nursing facilities requiring dressing assistance (Mitzner et al. 2013).

Dressing tasks are complicated, involving complex physical interactions between garments and the person's body. Many ADLs primarily involve moving a small tool with respect to the person's body. In contrast, dressing frequently involves extended flexible objects that must be maneuvered to surround the body. Earlier actions, such as a person's hand entering a sleeve, have important implications for subsequent actions, and can restrict the motion of the robot. In addition, the robotic workspace associated with dressing assistance can be large, making careful positioning of a fixed robot's base or use of a mobile base advantageous.

Determining how the robot and person can collaborate to complete the task is challenging. The collaboration involves two agents physically interacting, each of which typically has a high number of degrees of freedom (DoF). Disabilities can make both self and robot-assisted dressing more difficult. For example, involuntary movement, rapid fatigue, limited range of motion, and decreased dexterity are a few of the conditions we observed in our population of participants (Sect. 5.2). Personalization for individuals could be valuable due to the great diversity of disabilities. Additionally, other assistive tools the person may use, such as a wheelchair, can impede the robot's motion.

We present task optimization of robot-assisted dressing (TOORAD), a method for finding collaborative plans for a person and robot that will likely result in successful dressing. TOORAD makes use of geometric, kinematic, and physics simulation of the person, robot, and garment in its optimization. It uses personalized models of the geometry and

physical capabilities of the intended user. These models consider what the person is capable of doing, instead of what he/she typically does, with the notion that the person can learn to use the robot. With this approach, TOORAD is able to explore a wide range of actions for dressing in simulation, some of which might be challenging to test in the real world. Using a general-purpose mobile manipulator can mitigate some of the challenges in dressing by allowing the robot to move around to access different areas around the body, as explored by Kapusta and Kemp (2019). TOORAD optimizes the robot's base pose to improve the robot's ability to adapt to unexpected changes. Additionally, our method provides computer-generated instructions for the user receiving assistance from the robot. We have used TOORAD to optimize the actions of a person and a PR2 robot (a mobile manipulator made by Willow Garage) to collaborate in pulling two sleeves of a hospital gown onto the person's body.

We conducted a study with six people with disabilities who have difficulty dressing themselves to learn more about the habits, needs, and capabilities of some members of this population, as well as their views on robot-assisted dressing. Section 5.2 summarizes their responses and Sect. 6 provides a discussion. The results from this study may help guide future research in robot-assisted dressing.

We then evaluated TOORAD by having it generate collaborative plans to assist four of the six human participants and having a PR2 robot use these plans to assist the participants. The system was able to successfully assist all four participants in pulling on both sleeves of a hospital gown. Two of the four participants dressed one sleeve independently and received assistance from the robot with only the second sleeve, and two participants received assistance with both sleeves. Figure 1 shows one of the participants receiving dressing assistance from the robot. Our results provide evidence that TOORAD can be used to plan actions that will result in a robot and person with disabilities collaborating successfully to complete a dressing task.

The main contributions we present in this paper follow:

- A method (TOORAD) for how a robot and person's actions can be jointly optimized for dressing tasks.
- An implementation of our method applied to a PR2 assisting several participants with pulling sleeves of a hospital gown onto their arms.
- An evaluation of our implementation with participants with disabilities.
- A survey of people with disabilities on their needs and capabilities to inform this and other works.

2 Related work

Our approach uses offline optimization and simulation to generate personalized collaborative plans for robot-assisted dressing. In this way, our approach optimizes for an individual using computational models. Using simulation enables the system to discover novel plans and to evaluate candidate solutions that might be unsafe for a human. We also evaluated our system with participants who have disabilities. In this section, we compare and contrast our work with prior research, although many of these methods from the literature could potentially be complementary to our approach.

2.1 Learning for robot-assisted dressing

A number of efforts have investigated learning for robot-assisted dressing.

2.1.1 Learning from demonstration

Pignat and Calinon (2017) and Canal et al. (2018) use learning from demonstration (LfD) to have a human teach the robot how to perform dressing tasks. Canal et al. (2018) combines motion primitives with higher-level symbolic task planning to represent dressing tasks and demonstrate their method for putting a shoe onto a person's foot.

A potential advantage of LfD is that individuals can directly teach the robot, which gives them the opportunity to communicate their preferences. However, this approach requires that the person demonstrate a feasible plan for collaboration. This can be a challenge for some dressing tasks due to the complex and unintuitive ways in which some robots move, disabilities alter people's motions, articles of clothing behave, and robots, people, and clothing interact. Optimization via simulation has the potential to discover collaborative plans that would be difficult for a user to create.

2.1.2 Reinforcement learning

Tamei et al. (2011) proposed a method using reinforcement learning (RL) for a robot to learn trajectories to dress a mannequin in a t-shirt, focusing on topological relationships

between the mannequin and the shirt. Their experiments were initiated with the mannequin's arms in the shirt sleeves. Koganti et al. (2016) presented methods to improve RL for a similar shirt dressing task through dimensionality reduction via Bayesian nonparametric latent space learning. They also describe the possibility of using the low dimensional representation as a tool with which to demonstrate tasks for learning from demonstration. Twardon and Ritter (2018) presented a method to learn a policy to pull a knit hat on a mannequin head using policy search over a head-centric policy-space.

As with many efforts that have applied RL directly to real robots without simulation, this work requires that a reasonable initial policy be provided in some form, such as through a demonstration. As such, it can be beneficial for improving a policy through experience, but it is not well-suited to discovering a novel policy.

Applying RL in simulation in order to discover novel policies would be more closely related to our approach. Recently, Clegg et al. (2018) applied deep RL to synthetic character dressing. Their work included an example involving a freely floating simulated gripper providing assistance with a sleeve of a hospital gown to an able-bodied virtual character. However, RL for robot-assisted dressing in simulation has not been demonstrated with a full robot simulation or simulated disabilities, nor have policies been transferred to the real world.

2.2 Perceiving and modeling the user for robot-assisted dressing

Yamazaki et al. (2016) presented a system for a robot to assist users with pulling on trousers. Their work focused on using vision to estimate the state of the garment, of the person's body, and of the dressing task. They also used optical flow and force sensing to detect failures of the task and initiate recovery behaviors. Their system generated a fixed number of waypoints for a humanoid robot's end effectors along the outsides of the person's visually segmented legs. The robot performs a pre-defined set of subtasks that rely on cooperation from the person. The evaluation involved an able-bodied participant, a mannequin, and three different pairs of loose-fitting trousers made with different materials and visual patterns. A trial was deemed successful if the humanoid robot pulled the waist of the trousers above the knees of the participant or mannequin.

Gao et al. (2015) presented work on user modeling during robot-assisted dressing. They identify the pose of the user, model the movement space of the upper body joints, and select where a Baxter robot should place the openings of a sleeveless jacket. More recently, Gao et al. (2016) introduced a stochastic path optimization method to optimize a dressing trajectory for a Baxter robot using the estimated human pose

and haptic data from the sensor attached at the robot's end effector. Their method updates the dressing trajectory over several attempts of the task. For these two papers, a Baxter robot helped able-bodied participants put on a sleeveless jacket.

Zhang et al. (2017) presented a method that uses observed poses of a person's body to develop a model of a person's movement capabilities. Their control system has the robot follow waypoints attached to estimates of the person's arm pose, while also moving based on force-feedback to avoid applying high forces. In an evaluation with four able-bodied participants with simulated impairments, a Baxter robot successfully moved a sleeve of a sleeveless jacket onto one of the arms of the participants. Their use of waypoints is similar to the use of waypoints in our system.

Erickson et al. (2018a) presented a model predictive controller that can move a hospital gown sleeve onto a person's arm based on measurements from a 6-axis force-torque sensor. The controller used neural networks trained in a physics simulator that predicted the forces that would be applied to a person's body due to the robot's selected action. The controller was evaluated with a PR2 robot that moved a hospital gown sleeve onto the arms of able-bodied participants

For these papers, the placement of the robot relative to the person was fixed and selected by the researchers. In contrast, our system plans where the robot should be located relative to the person. This can be valuable, since success can be sensitive to the relative location of the robot. The systems in these papers used vision and force sensing to actively monitor and adapt to the person. Our method generates collaborative plans that rely on the human to be cooperative but are resilient to deviations from the plan. In our work, vision only serves to observe a fiducial marker, so that the models of the wheelchair and person can be aligned with the real wheelchair. Capacitive sensing serves as the primary way our robot reacts to the person's movements and adapts its plan, and force sensing detects potentially unsafe situations.

2.3 User-interaction for robot-assisted dressing

Jevtić et al. (2018) presented a system that provides assistance with putting on a slip-on shoe. The system provided interfaces with which the user could direct the robot with gestures and voice commands. The user could direct the robot to pick up, maneuver, and put on a shoe. The robot also integrated a number of perceptual and control processes to grasp the shoe and move it to an appropriate pose for the person's foot.

This work highlights the potential for richer interactions to enhance robot-assisted dressing. For our system, the human primarily influences the robot's actions via the capacitive sensor. First, our robot does not begin to move until the person moves his or her hand underneath the capacitive sensor. Sec-

ond, the robot moves to keep a constant distance between the capacitive sensor and the person's body. Other forms of interaction for dressing tasks like donning a gown merit further investigation.

2.4 Planning for robot-assisted dressing

Klee et al. (2015) presented a system that generates collaborative plans for a Baxter robot to assist a person with donning a hat. The system makes requests of the person when the robot is unable to succeed on its own. It also models the user's capabilities and generates new plans based on observations of the person's motions. The dressing task is represented as a sequence of goal poses with respect to the user's body. This results in a robotic system that interactively attempts to achieve a joint goal with a person.

Interactively working with a person to find ways to collaborate could be beneficial in a number of dressing contexts. However, as with other approaches, the applicability to higher complexity dressing tasks is unclear. Greedy approaches to interactively discovering strategies for collaboration might miss non-local solutions with beneficial properties. Using higher fidelity models in the form of simulation presents the opportunity for the robotic system to discover novel collaborative plans for more complex tasks, although the computational requirements might not permit interactive rates in the near term.

Umali and Berenson (2017) presented a system that generates plans for a Baxter robot to collaborate with a person to assist with the removal of personal protective equipment (PPE). This is perhaps the most similar work to ours, since it uses trajectory optimization to generate plans, includes a term in the objective function that rewards robot dexterity, uses subtasks, and plans for the robot and the human to move at different times.

However, the system they present does not use physics-based simulation, and instead uses geometric models with conservative bounds that approximate the constraints imposed by the PPE physics. Because of this coarse modeling and potentially other factors, their system requires a human operator to select a trajectory to use for a subtask out of a set of candidates. In addition, their system does not consider placement of the robot relative to the person, expects an initial demonstration for processing, does not consider disabilities or a wheelchair, and is focused on doffing instead of donning clothing.

2.5 Planning for mobile manipulators

A key distinction between our system and previous systems for robot-assisted dressing is our use of a mobile manipulator and the generation of plans that include optimized poses for the base and spine of the mobile manipulator. A com-

mon method is to select these poses using inverse-kinematics (IK) solvers (Diankov 2010; Beeson and Ames 2015; Kumar et al. 2010; Smits 2006). IK solvers typically seek a single joint configuration of the robot, although often many solutions exist or no solutions exist. The entire kinematic chain from end effector to the robot's base location can be solved using IK (Gienger et al. 2005; Grey et al. 2016). Alternatively, sampling-based methods can be used to find robot base poses that have valid IK solutions, often as part of motion planning (Elbanhawi and Simic 2014; Stilman and Kuffner 2005; Lindemann and LaValle 2005; Garrett et al. 2015; Diankov et al. 2008).

By relying solely on the existence of IK solutions to ensure that the robot can reach the goals, these methods are dependent on accurate models. Many of these methods are fast, but task execution may fail if there is modeling or state estimation error. Like these methods, TOORAD uses a sampling-based search in its optimization of the robot's configuration. However, there are often many robot configurations with valid IK solutions to all goals, and they cannot be distinguished using only IK. TOORAD uses measures of the robot's dexterity based on task-centric manipulability from Kapusta and Kemp (2019) to differentiate between those configurations. In that work, we showed that higher task-centric manipulability is correlated with improved performance for configurations that have collision-free IK solutions to all goals in a task.

2.6 Integrating task and motion planning

Our planning method relates to work that integrates task planning with motion planning, combining geometric and symbolic representations (Dogar and Srinivasa 2012). For example, Kaelbling and Lozano-Pérez (2013) uses a hierarchical planning architecture to layer symbolic planning on top of motion planning with uncertainty for generic mobile manipulation problems. We formulate our problem as a hierarchical optimization (Anandalingam and Friesz 1992; Migdalas et al. 2013), but our formulation is analogous to that of hierarchical planning.

3 Task optimization of robot-assisted dressing (TOORAD)

With our method, task optimization of robot-assisted dressing (TOORAD), we seek to find a sequence of actions for both the assistive mobile manipulator and the person receiving assistance that are likely to result in successful dressing. Key features of TOORAD are its simulation of cloth-human physics, its simulation of human and robot kinematics, its representations of robot dexterity, its planning of multiple actions for a task, and its planning of actions for both the robot and person using an optimization-based approach. TOORAD

is an offline process and it is suitable for situations when the garment and person can be modeled beforehand, when we would like to personalize the actions for the person's capabilities, and when we would like to configure the robot such that the task is successful despite variations between models and reality.

We first define the problem addressed by TOORAD and important assumptions it makes, followed by the details of the optimization. We provide examples from our implementation of TOORAD used in our evaluation, which is for the task of pulling the two sleeves of a hospital gown onto a person's arms.

3.1 Problem definition

TOORAD aims to find, within the space of the sequence of all actions that the person being assisted and the assistive robot can perform, \mathcal{U} , "What is a sequence of actions that will result in successful dressing?" Notably, TOORAD considers what humans *are capable of doing*, rather than what humans *typically do*. This choice follows from the notion that humans are currently more adaptable than robots.

With this approach, we avoid the challenges in modeling actions a person might take in given circumstances. Instead, we model the person's physical and kinematic capabilities. By taking this approach, we can take advantage of the robot's strengths and capabilities that may differ from a human assistant. For example, we can equip the robot with sensors, such as a capacitive distance sensor, that can estimate the distance to the person's arm in a way a human caregiver may find difficult. Additionally, instead of basing the robot actions on those from human caregivers, TOORAD explores strategies for dressing in a simulated environment that might be specific to robotic assistance. Some solutions can be difficult for caregivers or robotics experts to identify without computer assistance due to the high number of degrees of freedom, solutions that may be distinct from human practice, and the complexity of human disabilities that can vary greatly between individuals. The non-anthropomorphic kinematics of the robot presents challenges for approaches based on learning from demonstration or other forms of user control.

We can formulate the problem as optimizing the sequence of actions for the human and robot, $A_{h,r}$, such that

$$\arg \max_{A_{h,r} \in \mathcal{U}} R_d(A_{h,r}) \quad (1)$$

where \mathcal{U} is the domain of feasible robot and human actions and R_d is an objective function for dressing that we define in Eq. 3. We show the problem formulation in Table 1.

A challenge to solving this problem is that the space of all possible human and robot actions is large. To achieve

Table 1 The problem formulation for TOORAD as a single large optimization

$$\arg \max_{\mathbf{A}_{h,r} \in \mathcal{U}^c} R_d(\mathbf{A}_{h,r}) \quad (2)$$

$$R_d(\mathbf{A}_{h,r}) = -\psi C_n(N) + \frac{1}{N} \sum_{i=1}^N \left(-\zeta C_t(\mathbf{c}_{h,i}) - \eta C_s(\mathbf{c}_{h,i}, \mathbf{t}_{r,i}) + \gamma (\alpha R_{reach}(\mathbf{c}_{h,i}, \mathbf{c}_{r,i}, \mathbf{t}_{r,i}) + \beta R_{manip}(\mathbf{c}_{h,i}, \mathbf{c}_{r,i}, \mathbf{t}_{r,i})) \right) \quad (3)$$

Subject to:

- Collision constraints
- Garment stretching constraints
- Range-of-motion constraints of the person

Where:

- $\mathbf{A}_{h,r} = \{\{\mathbf{c}_{h,1}, \mathbf{c}_{r,1}, \mathbf{t}_{r,1}\}, \dots, \{\mathbf{c}_{h,N}, \mathbf{c}_{r,N}, \mathbf{t}_{r,N}\}\}$
- N is the number of subtasks, $N = \text{length}(\mathbf{A}_{h,r})$
- C_n is a cost on the number of subtasks
- C_t is a cost on torque experienced at the person's shoulder
- C_s is a cost on stretching the garment
- R_{reach} is a reward for how much of the trajectory the robot can reach
- R_{manip} is a reward for the dexterity of the robot arm along the trajectory
- $\psi, \zeta, \eta, \gamma, \alpha, \beta$ are weights for the terms in the objective function

tractability, we apply constraints on the action space that structure how interactions between the person and robot will take place. We constrain this space by limiting the robot's end effector to linear trajectories and alternating between the actions of the person and robot. Additionally, the person and robot base hold still as the robot moves the garment onto the person's body. This constrained search space, \mathcal{U}^c , simplifies the actions of the robot and the person. These simplified actions are also straightforward to communicate to people who will be using the robot for assistance.

We limit the sequence of actions to this constrained space, such that $\mathbf{A}_{h,r} \in \mathcal{U}^c$, and modify the optimization as shown in Eq. 2. $\mathbf{A}_{h,r}$ can then be defined as a sequence of paired actions by the human and robot,

$$\mathbf{A}_{h,r} = \{\{\mathbf{a}_{h,1}, \mathbf{a}_{r,1}\}, \dots, \{\mathbf{a}_{h,N}, \mathbf{a}_{r,N}\}\}, \quad (4)$$

where $\mathbf{a}_{h,i}$ and $\mathbf{a}_{r,i}$ are actions for the robot and human, respectively, for subtask i . Each pair of actions is performed in sequence. We call each pair of actions a subtask. N is the length of the sequence of actions, the number of subtasks for the dressing task. The actions of the person and robot are taken in that order, first the person, then the robot.

TOORAD does not automatically segment the task into subtasks. Instead, we provide a small set of candidate subtask sequences to the optimization. For example, a candidate sequence we used in our evaluation in Sect. 5.2 was to first dress the entire left arm and then dress the entire right arm. Other candidate sequences used included splitting each arm into two subtasks, first dressing the forearm and then dressing

the upper arm. Each subtask defines a kinematic trajectory for the robot's end effector, which holds the garment.

We narrow our definitions of the human and robot actions, based on our limited action space. The assumptions we use to define these actions are described in Sect. 3.2. We model the human's actions as only the pose he/she holds, and we ignore the movement of that pose. Therefore, we define the human's actions $\mathbf{a}_{h,i}$ as

$$\mathbf{a}_{h,i} = \mathbf{c}_{h,i}, \quad (5)$$

where $\mathbf{c}_{h,i}$ is a configuration of the human body that the person holds while the robot performs its action for the subtask, i . In our evaluation, the robot pulls sleeves onto a person's arms with each subtask using $\mathbf{c}_{h,i} \in \mathbb{R}^4$ (3 DoF at the shoulder, 1 DoF at the elbow) for the relevant arm.

We model the robot's actions for a subtask i as the pose of the robot's base, $\mathbf{c}_{r,i}$, and the trajectory of the robot's end effector, $\mathbf{t}_{r,i}$. We ignore the movement of the robot to achieve the base pose. Therefore, we define the robot's actions, $\mathbf{a}_{r,i}$, as

$$\mathbf{a}_{r,i} = \{\mathbf{c}_{r,i}, \mathbf{t}_{r,i}\}. \quad (6)$$

We hence refer to $\mathbf{c}_{r,i}$ as the robot configuration, which is a more general term for robot base pose as it may include other degrees of freedom such as the robot's spine height. In our evaluation, we used $\mathbf{c}_{r,i} \in \mathbb{R}^4$ (3 DoF for the robot base pose, 1 DoF for the robot spine height). The robot action is taken in sequence: first attaining the robot configuration and then executing the end effector trajectory.

We can substitute these definitions of the human and robot actions into Eq. 4, to obtain

$$A_{h,r} = \{\{c_{h,1}, c_{r,1}, t_{r,1}\}, \dots, \{c_{h,N}, c_{r,N}, t_{r,N}\}\}. \quad (7)$$

As before, the actions within each subtask occur in order: first the person attains his or her configuration, then the robot attains its configuration, and then the robot executes its end effector trajectory.

The goal of TOORAD is to find $A_{h,r}$, as defined in Eq. 7, that maximizes the value of the dressing objective function, R_d , as defined in Eq. 3. We discuss how we formulate the optimization architecture to maximize R_d , in Sect. 3.3. We define and discuss the details of the R_d , in Sects. 3.6, 3.7, and 3.8.

3.2 Assumptions

We make five key simplifying assumptions in addressing the problem described above. This list is not exhaustive, but highlights assumptions important to understanding the work. First, we assume that the robot is able to grasp and re-grasp the cloth, which we achieve during experiments through human intervention and special tools. We do not explore dexterous manipulation and grasping of cloth by the robot, as this problem has its own challenges and is being addressed by others (Fahantidis et al. 1997; Osawa et al. 2007; Maitin-Shepard et al. 2010; Miller et al. 2012; Bersch et al. 2011; Berenson 2013). Second, we assume that the participant receiving assistance is collaborative. That is, the participant will move his/her body to the extent of his/her ability in support of the task. Third, we assume the robot can estimate the pose of the participant's body prior to dressing. Fourth, we assume that the person and robot can achieve the desired configurations with negligible effect on the task. Fifth, we assume that human is able to hold his or her pose for the entire duration of the subtask.

3.3 Optimization architecture

We organize the optimization into three levels, shown in Table 2. We use layers of optimization (Sergeyev and Grishagin 2001; Anandalingam and Friesz 1992; Migdalas et al. 2013), to reduce computational cost. This structure reduces computation time compared to a single-layer, joint optimization of all parameters, by decreasing the times higher levels of optimization must be run, and more rapidly pruning areas of the search space. This architecture takes advantage of dependencies between the human configuration, robot configuration, and trajectory that we intend to optimize. The robot configuration and the trajectory of the robot's end effector depend on the configuration of the human. Both human and robot configurations and the robot trajectory depend on

the subtask being optimized. The top-level optimization is used to plan the optimal sequence of subtasks. The mid-level optimization is used to plan the optimal human configuration. The lower-level optimization is used to plan the optimal robot configuration.

An implication of this structure is that the nested optimizations heavily influence the computational requirements. Instead of optimizing the trajectory of the robot's end effector in a low-level optimization, we use previously selected trajectory policies for each subtask. The policy returns a kinematic trajectory of the robot's end effector given the human configuration. Trajectory optimization for dressing is complicated, as it involves complex physical interactions between the garment and the person's body, which can be computationally expensive to simulate. We trade accuracy in exchange for a policy that can rapidly be applied to many human configurations; the trajectory policy only gives a coarse approximation of the optimal trajectory.

3.3.1 Optimization algorithms

We provide TOORAD a small set of candidate subtask sequences to optimize, and we use an exhaustive search to consider all options in the set. In this case, exhaustive search is reasonable because the search space is small, with only a few candidate subtask sequences that are chosen manually. It is an open question how to better create candidate subtasks and search through them. The potential to develop more efficient computational methods for this type of problem is an interesting direction for future research but may not be strictly necessary to benefit people with disabilities.

The space of the human and robot optimization objective functions can be highly nonlinear and challenging to search, and the parameter space is large. These objective functions do not have an analytical gradient and estimating their gradients can be computationally expensive. Joint limits, model collisions, and cloth physics are constraints in the objective function that are highly non-convex and non-differentiable. We perform the optimizations of the human and robot configurations for each subtask using covariance matrix adaptation evolution strategy (CMA-ES) (CMA 2018; Hansen 2006), which works well for derivative-free, simulation-based, local optimization. We have observed that CMA-ES often performs better for local optimization than global optimization, and that starting with a good initialization often improves its performance.

3.4 Selecting candidate trajectory policies

We manually defined trajectory policies for each dressing subtask. A rectangular tool holding the sleeve pulls the sleeve first along the forearm, then along the upper arm, and finally moves to the top of the shoulder. Each linear

Table 2 The optimization performed by TOORAD is split into three levels

$\arg \max_{\pi} -\psi C_n(N) + \frac{1}{N} \sum_{i=1}^N (F(\pi_i))$	Top-level optimization to plan subtasks	(8)
<p>Where:</p> $F(\pi_i) = \max_{c_{h,i}} (-\zeta C_t(c_{h,i}) - \eta C_s(c_{h,i}, \pi_i(c_{h,i})) + \gamma R_r(c_{h,i}, \pi_i(c_{h,i})))$	Mid-level optimization of human configuration	(9)
$R_r(c_{h,i}, \pi_i(c_{h,i})) = \max_{c_{r,i}} (\alpha R_{reach}(c_{h,i}, c_{r,i}, \pi_i(c_{h,i})) + \beta R_{manip}(c_{h,i}, c_{r,i}, \pi_i(c_{h,i})))$	Lower-level optimization of robot configuration	(10)
<p>And where:</p> <ul style="list-style-type: none"> – $\pi = \{\pi_1, \dots, \pi_N\}$ – $t_{r,i} = \pi_i(c_{h,i})$ – $N = \text{length}(\pi)$ <p>All subject to:</p> <ul style="list-style-type: none"> – Collision constraints – Garment stretching constraints – Range-of-motion constraints of the person – C_n is a cost on the number of subtasks – C_t is a cost on torque experienced at the person’s shoulder – C_s is a cost on stretching the garment – $c_{h,i} \in \mathcal{H}$ – $\pi \in \Pi$ 		

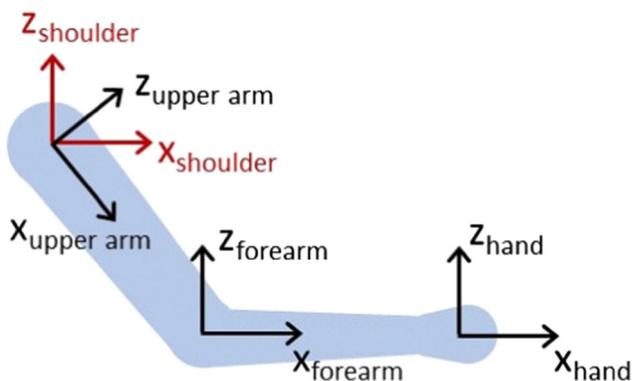


Fig. 2 The axes used in defining the trajectory policy overlaid on a diagram of an arm

trajectory is defined with respect to a coordinate frame at the base of the link being dressed, with its X-axis along the axis of the link and the Y-axis parallel to the ground plane. Figure 2 shows these axes overlaid on a diagram of the arm. The waypoints written as (x, y, z) in meters, were $(0.1, 0.0, 0.1)$, $(-0.03, 0.0, 0.1)$, $(-0.05, 0.0, 0.1)$, and $(0.0, 0.0, 0.1)$, with respect to the hand, forearm, upper arm, and shoulder, respectively. The policy is fixed for each subtask. Policies for different subtasks are created using these waypoints. For example, dressing the whole arm would consist of moving through all four waypoints and dressing the forearm would consist of moving through the hand and forearm waypoints.

Using a simulator of cloth-person physics, we verified that this policy succeeds in simulation for many configurations of the arm. For example, Fig. 3 shows the results in simulation of attempting different trajectories for pulling a sleeve onto a person’s forearm for different angles of the forearm. The figure also shows what a successful outcome looks like in the simulator as well as what it looks like when the sleeve catches on the arm. The center of mass of the successful trials for angles between 30° and -30° is consistently near 10 cm above the axis of the forearm. This result supports our selected policy. When the forearm deviates further from horizontal, deformation of the sleeve’s opening becomes problematic.

The policy π_i outputs a kinematic trajectory of the robot’s end effector, $t_{r,i}$, for pulling the sleeve onto the person’s body for subtask i for any human configuration in \mathcal{H}_i , where \mathcal{H}_i is the set of human configurations, $c_{h,i}$, where the trajectory policy succeeds. We write the trajectory as

$$t_{r,i} = \pi_i(c_{h,i}) \quad \forall c_{h,i} \in \mathcal{H}_i. \tag{11}$$

TOORAD uses a fixed-radius neighbors model¹ to quickly estimate if a human configuration, $c_{h,i}$ lies within \mathcal{H}_i .

¹ From <http://scikit-learn.org>.

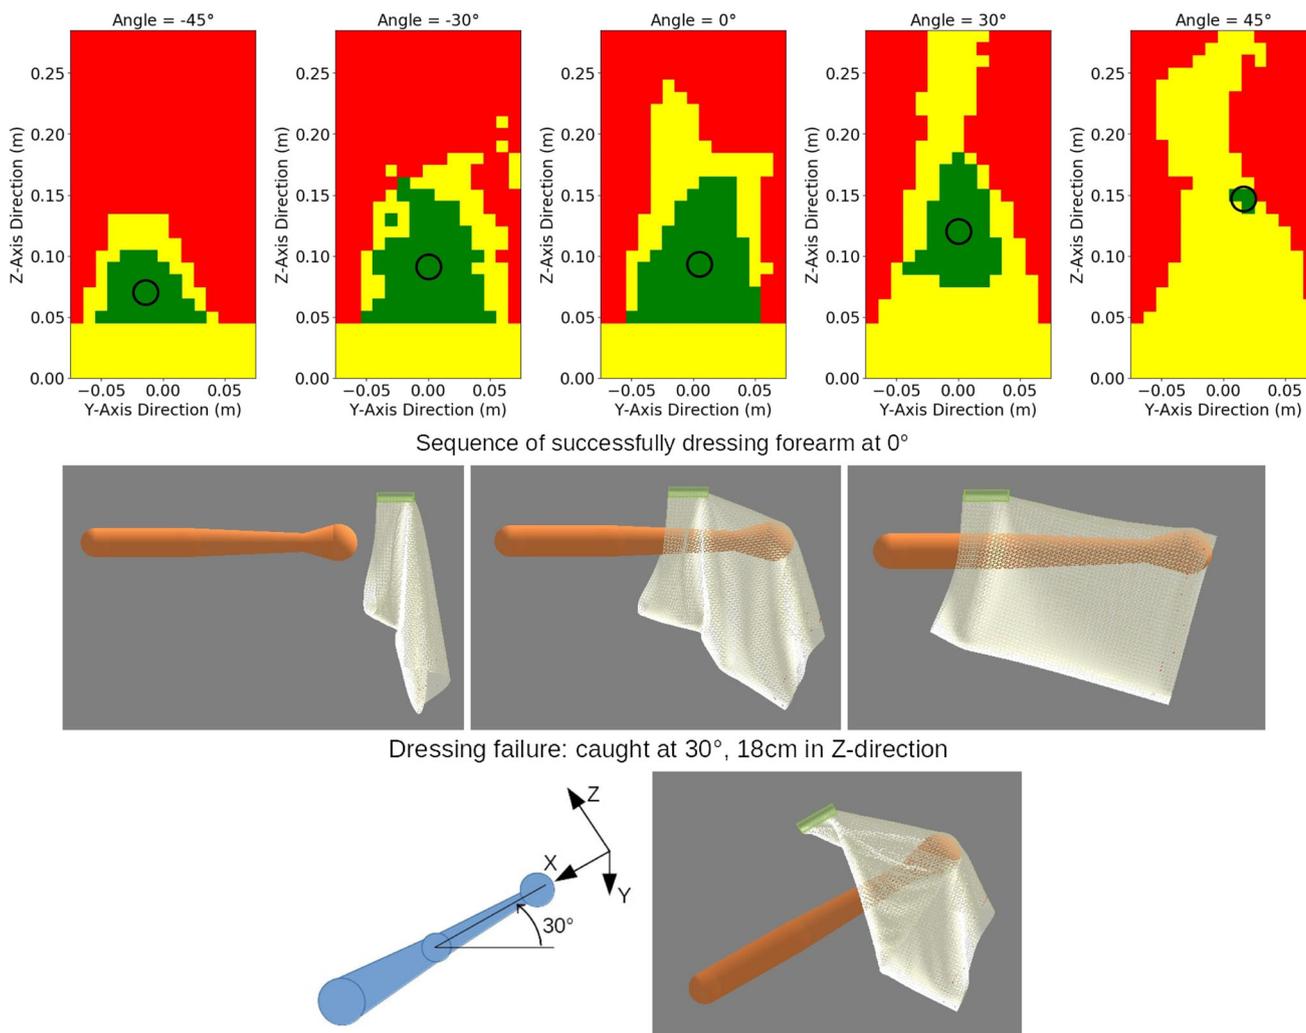


Fig. 3 (Top) plots showing the outcomes when pulling the sleeve onto a forearm for different poses of the arm. The sleeve is pulled by a tool along the axis of the forearm with varying start positions with respect to the arm. Green represents the forearm successfully going into the sleeve, yellow represents the arm getting caught on the sleeve, and red represents the arm missing the opening of the sleeve. The circle represents the centroid of the green area. For Z of ≤ 0.05 m, the tool holding the gown collides with the person’s arm. (Middle) a sequence of images

showing the sleeve successfully being pulled onto the forearm with the arm at 0° from horizontal and the tool moving 10 cm above the axis of the arm. The tool holding the sleeve is colored green. (Bottom-left) a diagram showing the forearm at 30° and the axes of the trajectory. X is along the axis of the forearm, Y is out of the plane, and Z is orthogonal. (Bottom-right) a view in simulation of the sleeve caught on the fist with the arm at 30° from horizontal with the tool moving at 18 cm in the Z direction and 0 cm in the Y direction (Color figure online)

3.4.1 Determining the human configuration space of the policy

Using both geometric and physics simulation, TOORAD estimates the space, \mathcal{H}_i , of human configurations, $c_{h,i}$, where the trajectory policy succeeds. It verifies that the following criteria hold true:

- The person’s body is not in self-collision.
- The robot’s tool does not collide with the person.
- The trajectory is successful in the cloth-person physics simulator.

If all criteria are satisfied, then we consider $c_{h,i} \in \mathcal{H}_i$. Otherwise, $c_{h,i} \notin \mathcal{H}_i$. The simulation used to determine the human configuration space where the policy is successful are shown in Fig. 4. The space \mathcal{H}_i is estimated once for each subtask using a generic human model based on a 50 percentile male from Tilley (2002) and a generic cylindrical sleeve of similar dimension to a hospital gown sleeve.

3.5 Constraints

The optimization of the human and robot configurations uses personalized models of the person, wheelchair (when

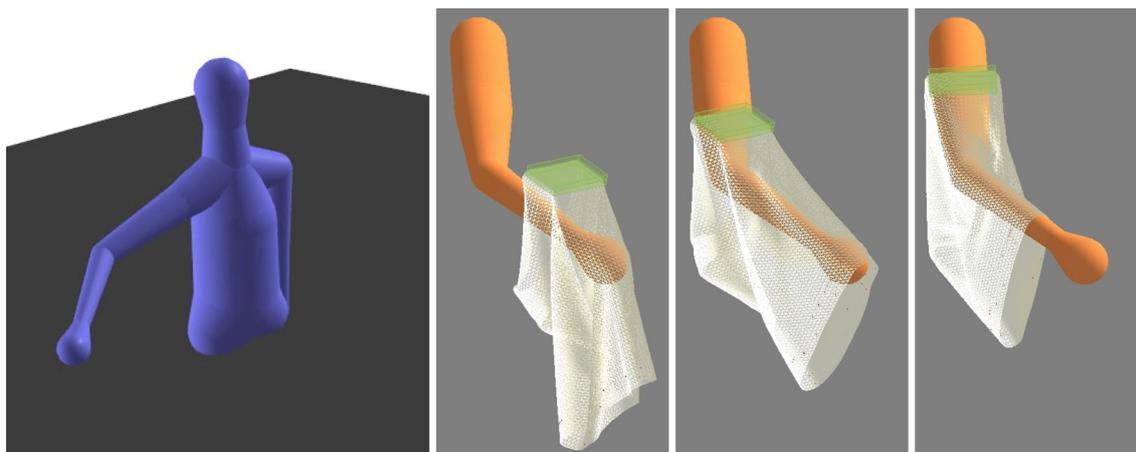


Fig. 4 Simulation is used to verify that our selected policy for the trajectory of the garment is reasonable and to find the space of human configurations, \mathcal{H} , where the policy succeeds. The simulators we used are described in Sect. 4.1. (Left) the DART simulator is used to verify collision constraints. Here we visualize the human model (without

legs) in the configuration being evaluated. (Right) the PhysX simulator used to simulate cloth physics when pulling the sleeve onto an arm in isolation in the same configuration (from a different perspective). The tool holding the sleeve is in green

applicable), and garment. Using a geometric simulator and a cloth-person physics simulator along with these models, TOORAD enforces six constraints in its optimization. These constraints are:

- Stretching limits of the garment.
- The person’s range of motion.
- The person’s body is not in self-collision.
- The robot does not collide with the person, wheelchair or garment.
- The garment does not experience interference from the wheelchair.
- $\mathbf{c}_{h,i} \in \mathcal{H}$

As the dressing task proceeds and subtasks are completed, TOORAD adds stretching constraints. These constraints are based on measurements of the real garment. For example, the maximum stretch between the two shoulders of the hospital gown we used in our evaluation was 0.5 m. Once the person’s left arm is in the sleeve of the gown, we assumed the left sleeve would remain in place. Therefore, for the subtasks that pull on the right sleeve, the right sleeve can move at most 0.5 m from the top of the person’s left shoulder. These stretching constraints model dependencies between subtasks.

We modeled a person’s range of motion by placing limits on the range of joint angles on a joint-by-joint basis using a 3-axis Euler-angle-joint at the shoulder and a single axis joint at the elbow. Some disabilities were easier to model with other constraints. For example, a range of motion constraint we used for one participant was that the participant’s upper arm could not be raised above parallel to the ground. We person-

alized TOORAD’s modeled constraints for the participants’ motor impairments described in Sect. 5.2.

3.6 Subtask optimization

For the candidate sequences of subtasks we provide, TOORAD performs the optimization in Eq. 8. The optimization is given a set Π of sequences of trajectory policies $\boldsymbol{\pi} = \{\pi_1, \dots, \pi_N\}$. The top-level optimization uses an exhaustive search to plan the best sequence of trajectory policies. We define C_n as a cost for each additional subtask,

$$C_n(N) = N. \quad (12)$$

The function F is the objective function for the mid-level optimization of the human configuration shown in Eq. 9. We manually selected $\psi = 0.1$ to be small, $\sim 100\times$ smaller than the values of F for human configurations that result in successful dressing.

3.7 Human optimization

Our formulation of F allows TOORAD to simultaneously consider the comfort of the person, the capabilities of the person, the kinematics of the robot, and the physics involved in manipulating garments onto a persons’ body. We formulate F in Eq. 9 as a linear combination of a few terms that we will define.

C_t is the cost on torque due to gravity experienced at the person’s shoulder in the configuration $\mathbf{c}_{h,i}$, C_s is the cost on exceeding a soft constraint on stretching the garment, and R_r is the score of how well the robot can execute the trajec-

tory (see Sect. 3.8). C_t is normalized to 0–1. C_s has units of distance with a maximum value of 0.04 m. R_r is unitless and ranges from 0 to 11. With these ranges, we manually selected 0.5, 5.0, and 1.0 for ζ , η , and γ , respectively. These gains emphasize functionality of the configuration, with more emphasis placed on the value of R_r while also considering the other costs.

3.7.1 Cost on torque

To encourage human arm configurations that are more comfortable for the person, we have included a cost on the torque at the person’s shoulder due to gravity. This cost ignores torques due to external forces, such as from the garment, robot, or wheelchair. We divide the torque at the current configuration by the maximum possible torque, to create a normalized torque cost. The maximum torque is when the arm is straight, parallel to the ground. This torque cost is independent of the weight and scale of the person’s arm. We use the body mass values for an average male.² The cost on torque follows as

$$C_t(\mathbf{c}_{h,i}) = \frac{\text{torque}(\mathbf{c}_{h,i})}{\text{maximum possible torque}} \quad (13)$$

3.7.2 Cost on stretching the garment

In addition to the hard constraint on stretching the garment described in Sect. 3.5, we added a soft boundary cost to encourage human configurations for which the garment is stretched slightly less. For some small distance before the hard constraint, the objective function receives a penalty defined as

$$C_s(\mathbf{c}_{h,i}, \pi_i(\mathbf{c}_{h,i})) = d_{\text{exceeded}}, \quad (14)$$

where d_{exceeded} is the maximum amount by which the soft boundary has been exceeded for all points along the kinematic trajectory of the robot’s end effector. This value comes from the kinematics simulator based on the human configuration and the end effector trajectory from the trajectory policy. For our hospital gown, based on its physical limits, we set a hard constraint on the maximum distance between the two shoulders of the gown to 0.5 m. We set a soft constraint of 0.04 m, resulting in a maximum value for C_s of 0.04.

3.8 Robot optimization

The reward function R_r estimates how well the robot can execute the trajectory for the given human configuration. TOORAD calculates it using the lower-level optimization of

the robot configuration, $\mathbf{c}_{r,i}$, shown in Eq. 10. This optimization is based on Kapusta and Kemp (2019) with modifications that are relevant to dressing tasks. This objective function uses two measures that we have developed to estimate how well the robot will be able to perform the dressing task: R_{reach} , which is based on TC-reachability, and R_{manip} , which is based on TC-manipulability (from Kapusta and Kemp (2019)). These two terms are related to task-centric (TC) robot dexterity.

Robot Dexterity Measures

TC-reachability and TC-manipulability are measures that consider the average values of reachability and manipulability across all goals in a task without order. Because dressing has a specific order that the robot should move to goals in its trajectory, we have added a graph-based search to the calculation of these terms to approximate this order and structure. Goal poses for the robot are determined from sampling the trajectory of the robot’s end effector. For each goal pose, TOORAD uses OpenRAVE’s *ikfast* (OpenRAVE 2018), to find a sample of collision-free IK solutions. These IK solutions are found for the PR2’s 7-DoF arm by creating an analytical IK solver for discretized values of one of its joints. We used the robot’s forearm roll joint discretized at 0.1 radian intervals. An IK sample is obtained for each discretized value of the forearm roll joint.

From these IK solutions for each goal pose, TOORAD creates a layered directed acyclic graph where a layer is created for each goal pose and each node is a joint configuration of the robot arm. Nodes in a layer are connected to nodes in the next layer if the maximum difference between each robot joint is less than a threshold (we used 40°). This threshold is used to predict that a path exists between the two joint configurations without the time cost of running a motion planner. To facilitate use of a standard graph search algorithm, we add a *start* node connected to each node in the first layer (IK solutions for the first goal pose in the trajectory) and an *end* node connected to each node in the last layer. Once it creates the graph, TOORAD performs a uniform cost search (Dijkstra 1959) to find a path with maximal JLWKI.

Joint-limit-weighted Kinematic Isotropy (JLWKI)

JLWKI is presented in Kapusta and Kemp (2019) and is based on kinematic isotropy, $\Delta(\mathbf{q})$, from Kim and Khosla (1991), shown in Eq. 15.

$$\Delta(\mathbf{q}) = \frac{\sqrt[a]{\det(\mathbf{J}(\mathbf{q})\mathbf{J}(\mathbf{q})^T)}}{\left(\frac{\text{trace}(\mathbf{J}(\mathbf{q})\mathbf{J}(\mathbf{q})^T)}{a}\right)} \quad (15)$$

Geometrically, kinematic isotropy is proportional to the volume of the manipulability ellipsoid of the manipulator, which is the volume of Cartesian space moved by the end effector for a unit ball of movement by the arm’s joints. It is based on manipulability from Yoshikawa (1984), with a

² From <https://msis.jsc.nasa.gov/sections/section03.htm>.

modification to remove order dependency and scale dependency. This metric can be useful when assessing kinematic dexterity between different configurations of the same robot. The values of kinematic isotropy are always in the range of 0 to 1 so they can be more directly compared across robot platforms.

JLWKI modifies kinematic isotropy to consider joint limits by scaling the manipulator’s Jacobian by an $n \times n$ diagonal joint-limit-weighting matrix T , defined in Eq. 16, where n is the number of joints of the manipulator.

$$T(q) = \begin{bmatrix} t_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & t_n \end{bmatrix} \tag{16}$$

t_i in T is defined as

$$t_i = 1 - \phi^\kappa \tag{17}$$

where

$$\kappa = \frac{q_i^r - |q_i^r - q_i + q_i^-|}{\lambda q_i^r} + 1 \tag{18}$$

and

$$q_i^r = \frac{1}{2}(q_i^+ - q_i^-). \tag{19}$$

We set $t_i = 1$ for infinite roll joints. The variable ϕ is a scalar that determines the maximum penalty incurred when joint q_i approaches its maximum and minimum joint limits, q_i^+ or q_i^- , and λ determines the shape of the penalty function. We used a value of 0.5 for ϕ and 0.05 for λ . This weighting function and the values for ϕ and λ were selected to halve the value of the kinematic isotropy at joint limits, have little effect in the center of the joint range, to begin exponentially penalizing joint values beyond 75% of the range, and to operate as a function of the percentage of the joint range. JLWKI is then defined as

$$JLWKI(q) = \frac{\sqrt[n]{\det(J(q)T(q)J(q)^T)}}{(\frac{1}{n})\text{trace}(J(q)T(q)J(q)^T)}. \tag{20}$$

Scoring Metrics

The graph-based search returns either a path from *start* to *end* or the longest path achievable towards *end*. The metric, R_{reach} , is related to the percentage of goal poses along the returned path to which the robot can find a collision-free IK solution from robot configuration, $c_{r,i}$. R_{reach} is defined as

$$R_{reach}(c_{h,i}, c_{r,i}, \pi_i(c_{h,i})) = \left(\frac{N_p}{N_{total}} \right), \tag{21}$$

where N_p is the number of nodes in the path and N_{total} is the total number of goal poses. R_{manip} is then defined as

$$R_{manip}(c_{h,i}, c_{r,i}, \pi_i(c_{h,i})) = \left(\frac{1}{N_{total}} \right) \sum_{i=1}^{N_p} JLWKI(p_i), \tag{22}$$

where p_i is the IK solution for the i th goal pose.

4 TOORAD implementation details

4.1 Simulators

TOORAD uses two simulators, one for simulation of human-cloth physics, and a second for simulation of human-wheelchair-robot-garment geometries and robot kinematics. We personalized the human and wheelchair models for each participant, and we customized the simulated garments based on the hospital gown dimensions.

We used the cloth simulator in Nvidia’s PhysX (PhysX 2011) to simulate human-cloth physics due to its efficiency and robustness at handling large contact forces (e.g., when the sleeve opening is caught on the hand). PhysX is based on position-based dynamics (PBD), which directly calculates position changes instead of through force integration. By avoiding solving and integrating forces and instead modifying position directly, PBD can be more stable, controllable, and efficient than alternative methods. We modified PhysX to add additional functionality and improved accuracy of friction handling, as described in the work by Yu et al. (2017). We manually selected the parameters of the simulator so the modeled fabric would behave similarly to the real-world hospital gown.

We modeled the person in PhysX using capsules, or pairs of spheres connected by a conical frustum, for simplicity: these are primitives in both PhysX and DART. We modeled the robot’s tool holding the garment as a parallel jaw gripper that holds the shoulder region of the gown. We did not model any other part of the robot. This setup corresponds to the tool we used with the real-world PR2. Figure 4 (right) shows the PhysX simulation environment with a sleeve being pulled onto an arm.

TOORAD simulates the human-wheelchair-robot-garment geometries in the DART (DART 2018) simulation environment. In this environment, TOORAD performs collision detection using Bullet (Coumans and Bai 2016–2018) and uses OpenRAVE’s ikfast (OpenRAVE 2018) for the robot’s inverse-kinematics. We set up this second simulation environment because DART has native support for robot and human models. DART uses the same human model as in PhysX, and we also added a personalized model of the

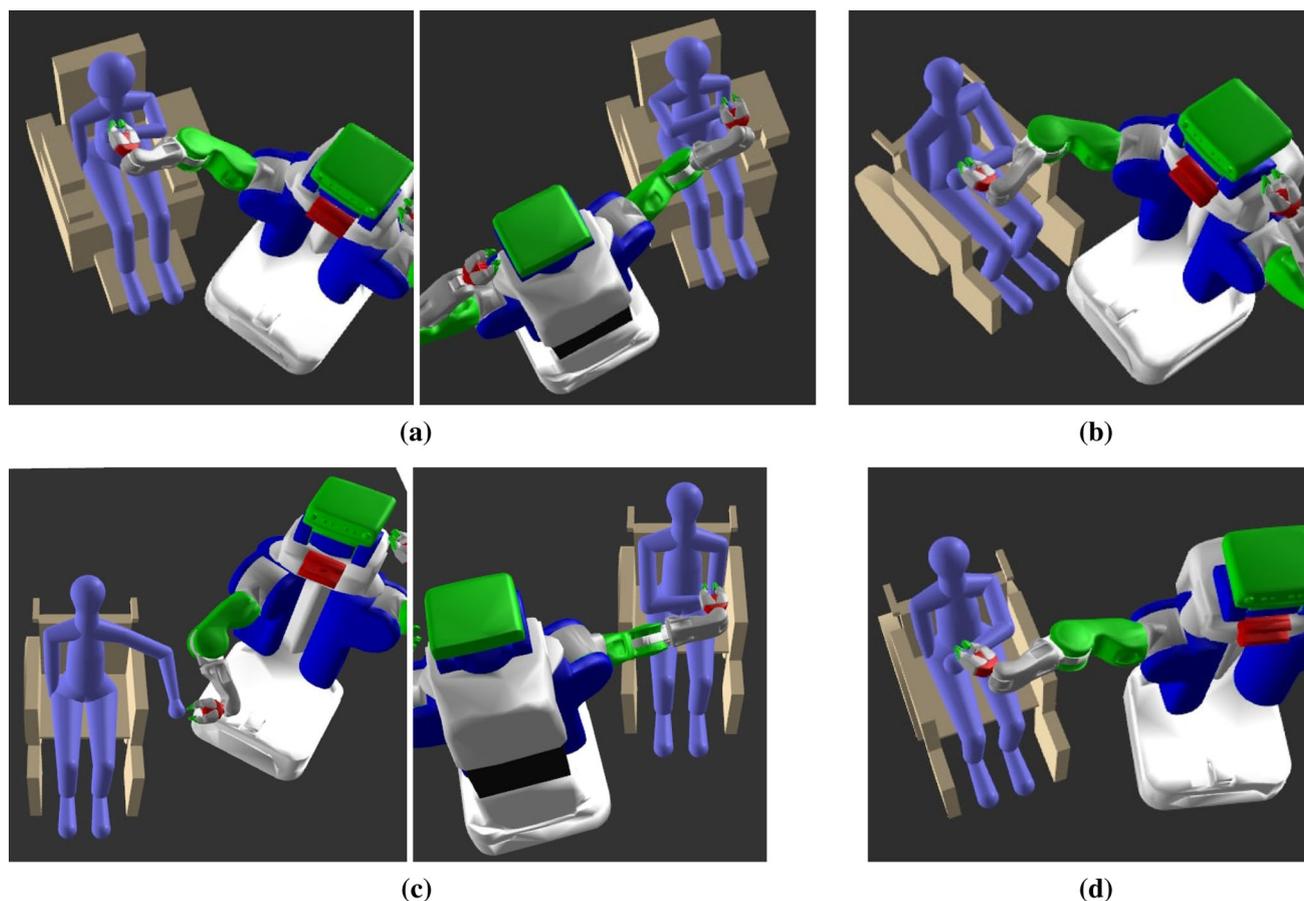


Fig. 5 The initial robot and human configurations planned by TOORAD for the four participants who interacted with the robot. The models have been personalized for the participants. The visualization comes from the DART simulator. The solutions found by TOORAD

were shown to participants to provide instruction on their expected arm poses. The configurations are grouped by participant. **b, d** show the configurations for the participants who received assistance with only one sleeve

person's wheelchair within DART, and a coarse geometric representation of the garment. Figure 5 shows this simulation environment (the garment model is not shown).

4.2 Practical additions to the optimization

In our implementation of TOORAD, we used several methods to speed up computation and to improve results. As mentioned in Sect. 3.3, we observed that the optimization algorithm we used, CMA-ES, was sensitive to its initialization.

We found three candidate human configurations with which to initialize the optimization of the human configuration performed by CMA-ES. To find these candidate configurations we first used a grid-search with a resolution of 2.5° for the 4 DoF of the human arm to find feasible human configurations: where $c_{h,i} \in \mathcal{H}$ and where constraints described in Sect. 3.4.1 were satisfied. We then used a K-means clustering algorithm to choose the three candidate

human configurations from these feasible configurations. We selected $K = 3$ manually based on the number of regions of feasible human configuration we often observed. We used L_2 distance on the 4 DoF configuration space of the human's arm as the distance metric for the K-means algorithm. We then ran CMA-ES three times, once with each of the three candidates, with a population size of 20, and 50 iterations and selected the best of the three results.

We manually selected five candidate robot configurations around the person to initialize the optimization of the robot configuration performed by CMA-ES. These candidate robot configurations were to the front and sides of the person, with the robot's arm facing the person. In TOORAD, the robot configuration optimization takes place in a lower level, so it is run many times. To decrease computation cost of this lower-level optimization, CMA-ES was limited to 20 iterations with a population of 20. To compensate for this coarser optimization, we added a step after the main optimization in which TOORAD more finely optimizes the robot confi-

uration. This step used the fixed human configuration from the main optimization and used the optimized robot configuration from the main optimization to initialize CMA-ES. This final step used smaller step-sizes and 80 iterations for CMA-ES, but otherwise used the same framework.

The grid search took ~ 30 min on 18-cores of a 64-bit, 14.04 Ubuntu operating system with 60GB of RAM and a 2.9GHz Intel Xeon E5-2666 v3. The main optimization took ~ 10 hours on the same machine. The fine-tuning of the robot configuration took ~ 10 min on a single core of the same machine.

We do not consider this computation time to be a limitation as it is an offline optimization problem that only needs to be performed a single time for a new user. For a device that would potentially be used daily by an individual, one overnight optimization should not be a practical concern. In addition, use of additional computers or optimization of the code could reduce the time required.

5 Evaluation

We used TOORAD to plan the actions for a PR2 robot (a mobile manipulator made by Willow Garage) and a participant with disabilities to collaborate in a dressing task. The participants self-identified as requiring assistance with dressing. The dressing task we chose was putting on the two sleeves of a hospital gown. We applied TOORAD to dressing a hospital gown, as it is a garment that tends to be both uniform and ubiquitous in hospital settings. Additionally, the actions involved in dressing a hospital gown, pulling sleeves onto the arms, is representative of many other dressing tasks for which a tube of fabric is pulled onto part of the body (e.g., shirt sleeves, pant legs).

5.1 System implementation

We implemented a system to test TOORAD using a PR2 robot. All TOORAD source code is openly available.³ We limited the PR2 to use only its right arm. TOORAD found, for each subtask of the dressing task, a configuration for the participant's arm, a pose of the PR2's base, a height of the PR2's spine, a kinematic trajectory for the PR2's end effector, and joint configurations for the PR2's arm for samples along the end effector's trajectory. In this case, $c_{h,i} \in \mathbb{R}^4$ (3 DoF at the shoulder, 1 DoF at the elbow), and $c_{r,i} \in \mathbb{R}^4$ (3 DoF in the robot base pose, 1 DoF in the robot spine height). We evaluated if, using our system, a participant and the PR2 could collaborate to successfully complete the dressing task.

A challenge when moving from simulation to the real world is that there are often discrepancies from simulation,

and sometimes even small differences can result in errors and task failure. We made use of several sensing and control schemes to allow the robot to both recognize and adjust to differences between the simulation and reality.

5.1.1 Grasping

When manipulating the hospital gown, the PR2 grasped a tool in its right gripper that held the top of the sleeve of the hospital gown. The tool was instrumented with a 6-axis force-torque sensor and a capacitive sensor that estimates the distance to the human body. The experimenters attached the tool to the gown before the start of each subtask.

5.1.2 Sensing

We used various sensors on the PR2 and the tool to estimate the pose of the person and track arm motion in real time. Using the PR2's head-mounted Kinect, the system perceives an ARTag mounted in a known position with respect to the person's wheelchair. The system estimates the pose of the person by assuming that the person is initially in the instructed pose produced by the optimized collaborative plan.

As we have noted in previous work (Kapusta et al. 2016b), visual perception of the person's body during dressing can be challenging as the body becomes occluded by the garment, so the system use force sensing and capacitive sensing to make inferences about the state of the dressing task and to more finely estimate the pose of the person's arm while executing the plan. The PR2 used the capacitive sensor on the tool to measure the distance between the tool and the person's arm during dressing and to track along the arm (maintaining a desired distance), as presented by Erickson et al. (2018b). This sensor's range is roughly from 0 to 10 cm, which works well for this purpose. With the capacitive sensor, the PR2 was able to move the garment along the arm with little contact despite movement in the arm pose or arm poses that did not match the expected pose. Our system also used the force-torque sensor to improve safety and detect anomalies, stopping the PR2's movement if the magnitude of forces exceeded 15 N.

5.1.3 Control

We implemented a proportional-derivative (PD) Cartesian controller running at 20 Hz on the PR2 to move the tool along the planned end effector trajectories, while maintaining a distance of 5 cm to the body using the capacitive sensor. The system waited until the person's arm was near the tool (the capacitive sensor read less than 5 cm) before starting to move. We used low stiffness for the PR2's arm to provide compliance and greater safety. We used TracIK (Beeson and Ames 2015) to feed joint-space input to the PR2's low-level

³ <https://github.com/gt-ros-pkg/hrl-assistive/>.



Fig. 6 The axes used by the controller are shown in red. The coordinate system is attached to the tool, so the axes translate and rotate with the tool. The trajectory for moving along the forearm, upper arm, and shoulder are illustrated in light blue (Color figure online)

PID controllers. The X axis points in the direction of the tool (the direction of the sleeve's opening) and the Z axis is orthogonal to the capacitive sensor, pointing away from the person's arm. These axes translated and rotated with the tool.

The robot obtained waypoints for the movement of the end effector from the end points of the linear end trajectories planned by TOORAD for each subtask. For example, for the subtask of dressing the entire left arm, the robot would obtain four waypoints: at the hand, at the elbow, at the proximal end of the upper arm, and at the top of the shoulder. Figure 6 shows the axes and the trajectory for dressing the entire left arm for one of the participants. The controller calculated error in the X and Y directions as the difference between the current pose of the end effector found via the robot's forward kinematics and the linear end effector trajectory defined with respect to the ARTag pose estimate. The controller calculated error in the Z direction as the difference between the distance to the body estimated by the capacitive sensor and a goal distance of 5 cm. The desired end effector orientation was the X axis pointing along the planned end effector trajectory and the Y axis parallel to the ground; orientation error was calculated as the angle in the axis-angle transformation between the current and desired orientations.

5.2 Study with participants with disabilities

We conducted the study with approval from the Georgia Institute of Technology Institutional Review Board (IRB), and obtained informed written consent from all participants. The study took place from March to May of 2018 at the Georgia Institute of Technology. The inclusion/exclusion criteria

were: all participants must have difficulty putting both arms through the sleeves of a jacket without assistance, be able to raise at least one arm against gravity, have no cognitive impairments, and be ≥ 18 years old. Six participants participated in the study. Three of the participants were female and three were male. The age range was 23–69. In the first session of the study, we administered questionnaires on the participant's current habits, needs, and physical capabilities as related to dressing and their views on the potential for robots to assist with dressing.

We additionally took measurements of their body geometry, the comfortable range of motion of their arms, and the wheelchair (we provided them a wheelchair if they did not use one). The body geometry measurements we used were shoulder height, the distance between shoulders, upper arm length and girth, forearm length and girth, wrist length, fist radius, head height and radius, and position of shoulders with respect to the wheelchair. We measured the range of motion of the shoulder pitch, shoulder yaw, and elbow bending angles, and we noted any dependence between the range limitations of these joints. We took sufficient wheelchair measurements to construct the geometric models shown in Fig. 5.

We invited all participants whose physical capabilities were matched to the capabilities of our system back for a second session of the study in which the participant and the robot collaborated to pull the two sleeves of the hospital gown onto the participant's body. Two of the six participants did not match these criteria because they could not comfortably hold their arms up against gravity. Four of the six participants took part in the second session. For these tests, TOORAD was run on a model personalized for the participant in advance of their arrival.

5.2.1 Participant details

Because the responses, needs, capabilities, and views of participants varied greatly depending on their disabilities, we now describe each participant's disabilities relevant to dressing. These descriptions come from the administered questionnaire in the first session of the study.

Participant 1 Participant 1 has cerebral palsy that causes involuntary movements and makes precise movements challenging. He/she uses a wheelchair, which we modeled as shown in Fig. 5a. He/she has a full range of motion of both arms and can attain and hold most poses with both arms, but movements while reaching poses can vary and are not fully voluntary. The participant said it can be hard to get his/her hands into the right spot, holes, etc., and said that the hardest part of dressing is jackets, buttons, zippers, shoes, and tasks requiring fine motor control. This participant received assistance from the robot with dressing both arms. Figure 5a

shows the configurations planned by TOORAD for this participant.

Participant 2 Participant 2 is an older adult who has post-polio syndrome. He/she has weakness in his/her arms and legs and has no feeling in his/her hands. He/she has limited range of motion with pain near the edges of the range and cannot raise his/her arms to be parallel with the ground and cannot move his/her right arm in front of his/her body. He/she quickly becomes tired and pained by raising his/her arms or holding his/her arms in a raised posture. This participant did not take part in the experiment with the robot because he/she could not comfortably raise his/her arms.

Participant 3 Participant 3 has partial paralysis of the right arm due to a C3-C8 spinal cord injury. Paralysis is of the arm distal from the elbow and of the bicep. He/she has no other impairments or limitations. This participant used his/her unimpaired arm to dress his/her impaired arm. The participant then received assistance from the robot in dressing his/her unimpaired arm. The configurations planned by TOORAD for this participant are shown in Fig. 5b.

Participant 4 Participant 4 has ALS that has caused muscle weakness in his/her shoulders, arms, hands, and legs. He/she can briefly raise arms against gravity but is unable to continue to hold them up. He/she has weakness in his/her wrists and hands that make grasping and pulling on garments, especially pants, difficult. This participant did not take part in the experiment with the robot because he/she she could not comfortably hold up his/her arms.

Participant 5 Participant 5 has an amputation of the dominant right hand, located in the right forearm, near the wrist, and has a torn right rotator cuff. He/she experiences tensing of the left shoulder from overwork that is uncomfortable and that can sometimes involuntarily pull his/her head, neck, and left shoulder together. He/she experiences discomfort reaching the left arm across the body. He/she has a prosthetic but does not consider it a helpful tool for dressing because it cannot grip sufficiently dexterously or strongly to be helpful. This participant received assistance from the robot with dressing both arms. The configurations planned by TOORAD for this participant are shown in Fig. 5c.

Participant 6 Participant 6 has cerebral palsy whose symptoms include motor impairments, fatigue and balance issues. Impairments are primarily of the right side of the body. He/she lacks some control in the right arm, specifically of the right elbow and wrist. His/her right shoulder has limited range of motion and decreased control. His/her flexion of the right elbow is not under independent voluntary control and is dependent on the pose of the right shoulder. His/her right

wrist is always in a flexed posture. He/she sometimes experiences involuntary spasms, most typically in the right arm. His/her left arm has reduced range of motion and he/she experiences balance issues when seated when reaching across his/her body from the left side to the right side. This participant used his/her less impaired arm to dress his/her impaired arm and received assistance from the robot with dressing his/her unimpaired arm. We found that our method of modeling range of motion was not well matched to the motion constraints of the participant's right arm. Figure 5d shows the configurations planned by TOORAD for this participant.

5.2.2 Questionnaire on habits, needs, capabilities, and thoughts on robot-assisted dressing

The purpose of this survey was to gain insight for future robot-assisted dressing research through a better understanding of target populations. In this section, we describe the types of questions asked in the questionnaire. A discussion of take-aways from the questionnaire responses is in Sect. 6.

We asked a series of questions regarding typical dressing habits, such as typical dressing scenarios, number and types of dressing tasks, and time-to-complete, to learn how a robot might fit into a daily routine. We further inquired about which dressing tasks and articles of clothing are most challenging and how participants might prefer a robot to help meet their dressing needs. We included questions to survey participants' physical capabilities, as related to dressing, including range-of-motion, involuntary movements, and fatigue. These responses might help with the design of methods or systems for specific target populations in the future. Section 6 provides a discussion of the responses to the questionnaire.

Table 3 summarizes the participants' answers to two 7-point-scale items asked in the questionnaire. Participants were generally positive towards receiving assistance from a robot with dressing.

5.2.3 Experimental protocol

Before the start of the dressing task, each participant was given verbal instructions by the experimenters and shown a visualization of the DART simulator with the human poses expected by the robot and how the robot would move. The experimenters provided each participant multiple view points of the poses in the simulator. Figure 5 shows snapshots from that visualization for each of the participants.

Each participant started the experiment seated in a wheelchair. Participants who did not bring their own, were provided a wheelchair by the experimenters. TOORAD used a model of the wheelchair in its simulation.

For this study, participants were asked to pull on the two sleeves of the hospital gown with assistance from the robot. Depending on the participant's level of impairment

Table 3 Responses to 7-point-scale questionnaire items from participants with disabilities

	Score: mean (std)
Pre-experiment questions ($N = 6$)	6.3 (0.74)
When I select clothing, the ease of don and doff is important	6.3 (0.74)
If a robot could assist me with dressing to the same effectiveness as my current methods, I would allow the robot to assist me with dressing	7.0 (0.0)
Post-experiment questions ($N = 4$)	6.3 (1.15)
The system successfully accomplished tasks	4.7 (1.53)
I felt comfortable using the robot-assisted dressing system	5.3 (1.53)
I would prefer receiving dressing assistance from a robot than another human	
I would like to receive this form of dressing assistance from a robot in the future	

Pre-experiment questions were administered in the first session of the study. Post-experiment questions were administered after interaction with the robot in the second session of the study. The second column provides the average and standard deviation of scores with 1 = strongly disagree, 2 = disagree, 3 = slightly agree, 4 = neither, 5 = slightly agree, 6 = agree, and 7 = strongly agree

and the assistance they desired, participants either received assistance from the robot with one or both sleeves. Two of the participants had one impaired arm and one unimpaired arm. These participants pulled the gown unaided onto their impaired arm and then received assistance from the robot with dressing their unimpaired arm. Two of the participants received assistance with pulling on both sleeves. TOORAD was used to plan the actions of the participant and the robot. For each sleeve with which the participant received assistance, the participant held his/her arm in a pose as the robot moved the sleeve along the arm. Figure 7 shows a sequence of images from dressing the right arm of a participant as well as the expected pose of the person and robot in simulation. For all four participants, TOORAD's optimized sequences of actions consisted of a single subtask for dressing each arm. Figure 5 shows the human and robot configurations planned by TOORAD for each of the four participants. We note that during informal testing during development, TOORAD split dressing a single arm into two subtasks, particularly for people with longer arms.

The PR2 began the experiment by approaching the participant. The experimenters then attached the PR2's tool to the

gown. For participants receiving assistance with both sleeves, the PR2 started with the gown held in its tool. For participants receiving assistance with only a single sleeve, the experiment started with the gown being handed to the participant to independently dress his/her impaired arm. For all participants, once the robot was in position, the participant started the task by moving his/her hand within 5 cm of the capacitive sensor on the tool.

Once the participant started the task, the PR2 moved the sleeve slowly up the participant's arm until the robot completed its trajectory or the 15 N threshold was reached. At that point the experimenters judged if the subtask was successful by examining if the arm was in the sleeve, the sleeve was pulled up the participant's shoulder and the sleeve was not caught on anything. If successful, the robot would move on to the next dressing subtask. If the subtask was unsuccessful, the entire dressing trial was deemed a failure and the experimenters would set up the next trial. Between subtasks, the experimenters would detach and reattach the PR2's tool from the gown so the robot could drive its base around freely without affecting the gown or the participant.

All movement of the robot's base was controlled via teleoperation by the experimenters. The experimenters moved the robot to the base pose planned by TOORAD.

We gave participants the opportunity to practice receiving assistance from the robot with putting on the sleeves of the hospital gown up to five times prior to the official test. During practice, the experimenters provided feedback to the participant on how to improve coordination with the robot. For example, if the participant was not attaining the correct arm pose, the experimenters demonstrated the arm pose planned by TOORAD. The experimenters also responded to questions from the participant. After the practice trials, the participants attempted to dress themselves four times with the robot's assistance. During the official test, the experimenters responded to questions, but otherwise did not provide feedback. At the end of the experiment, the experimenters administered a questionnaire with 7-point-scale questionnaire items as well as three open-ended questions.

5.2.4 Results: robot-assisted dressing system

All four participants successfully dressed both of their arms using assistance from the robot four times consecutively. Success criteria was: both arms were in the sleeves, the sleeves were pulled up to the shoulders, and the gown did not get caught on the person or wheelchair. Figure 8 shows a sequence of images from a participant successfully collaborating with the robot to dress his/her right arm.

One participant experienced a failed trial in addition to the four successful trials due to the experimenters moving the robot to the wrong base position. Another participant experienced a failed trial in addition to the four successful trials

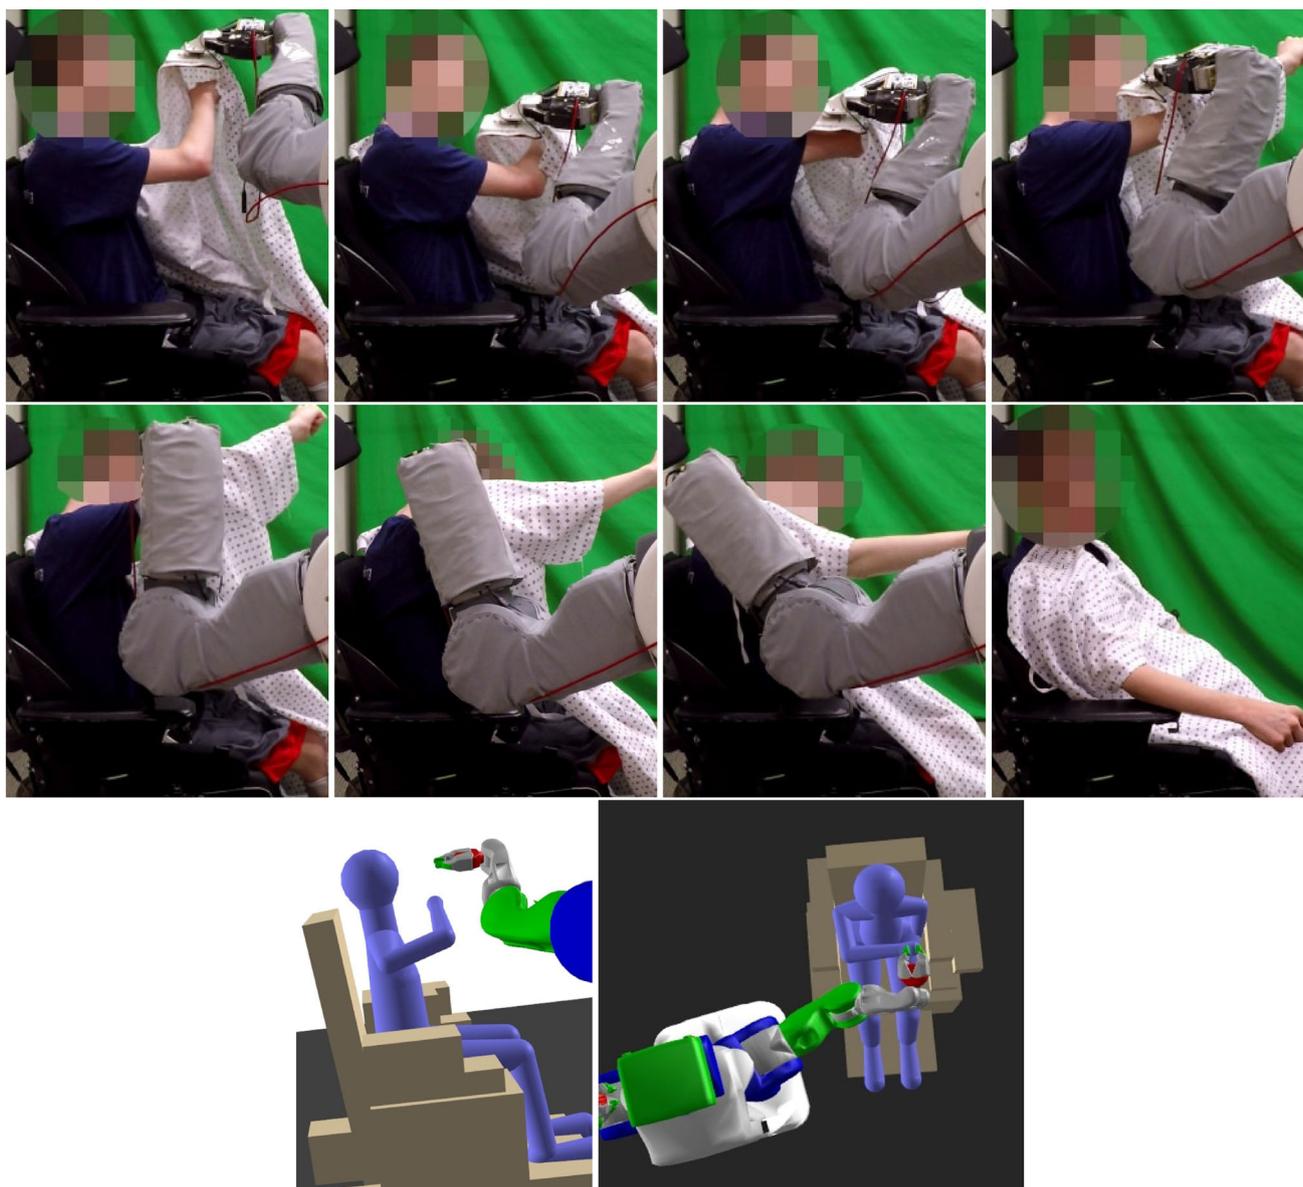


Fig. 7 (Top) a sequence of images from a participant receiving dressing assistance with the right arm of the hospital gown. The robot successfully pulled the sleeve up the arm as the participant's arm moved

involuntarily. (Bottom) views from the DART simulator showing the initial configurations of the person and robot for dressing the right arm

due to an erroneous IK solution that asked the robot to spin one of the roll joints of its arm 360° during the task. The two failed trials were retried and subsequently were successful. The average, maximum, and minimum times to complete dressing for each arm (excluding setup time) were 29.0 s, 37.9 s, and 24.0 s, respectively.

Table 3 summarizes a few of the participants' answers to 7-point-scale questionnaire items asked after interaction with the robot-assisted dressing system. Participants strongly agreed that the system successfully accomplished tasks but were more mixed on comfort and future assistance from

robots. In open-ended questions, participants were positive about the robot's ability to move along their arm, tracking their arm as it moved up or down. This tracking was particularly important for one participant who experiences involuntary movement of his/her arms. Figures 7 and 8 show sequences of images from two participants receiving dressing assistance from the robot, in which the participants moved their arms during the execution of the task.

Some of the feedback from participants was positive, including the following quotes: "This is pretty kick-ass", "You are not following the robot, the robot is following you",

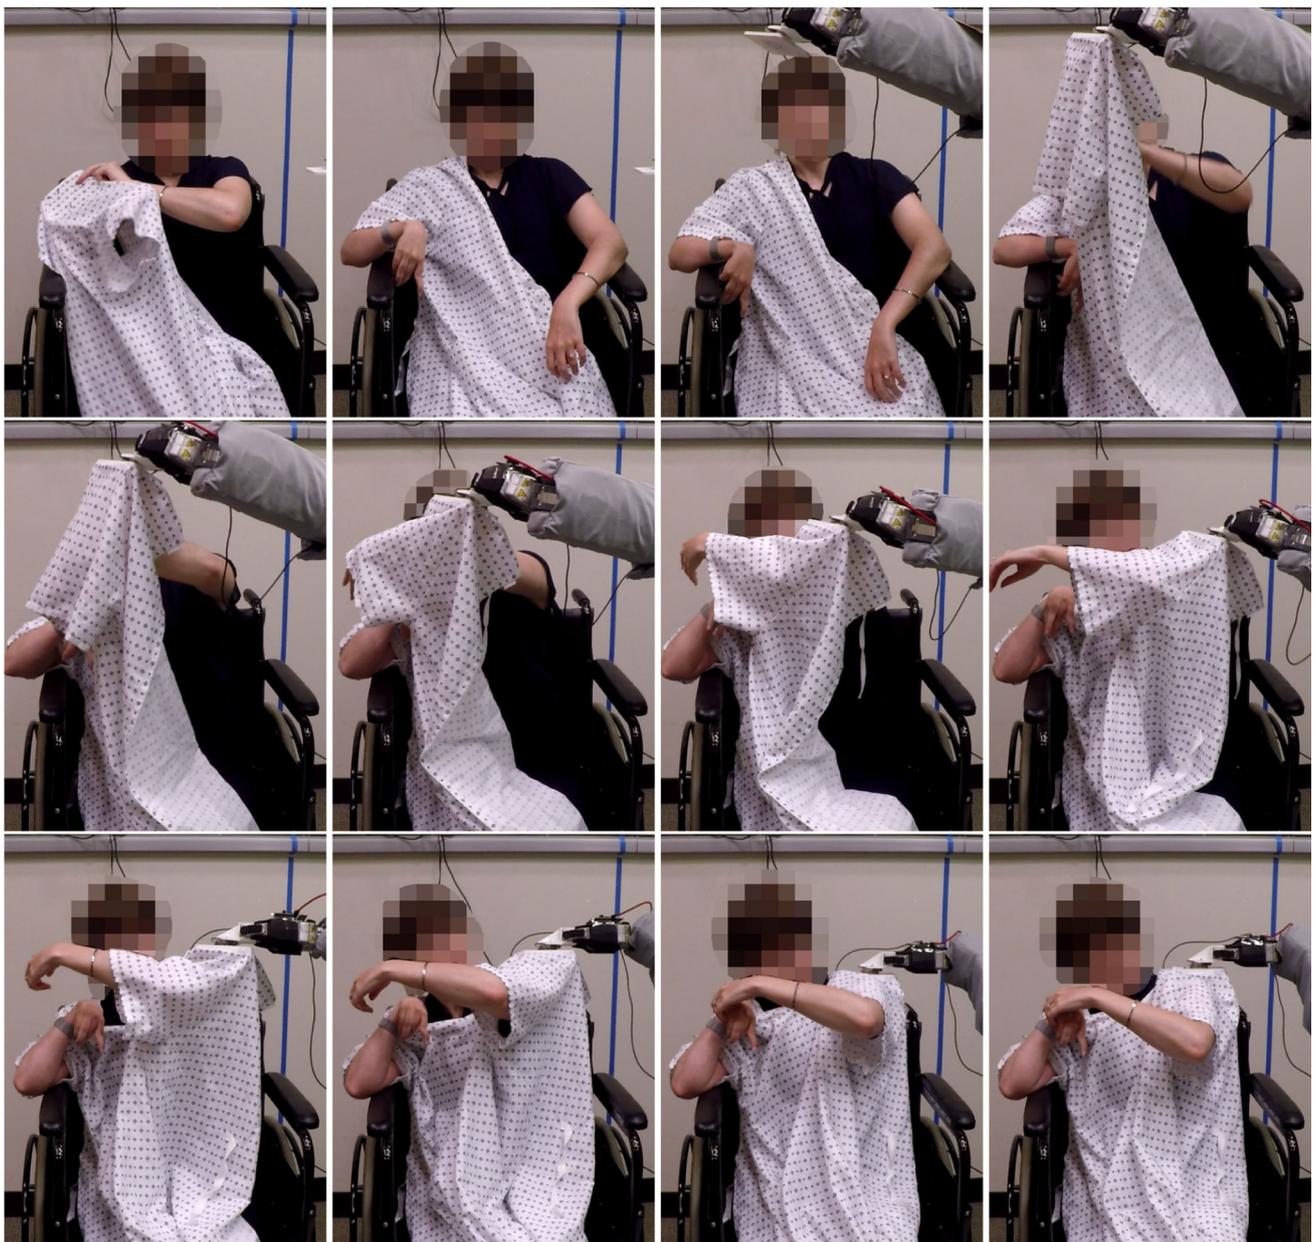


Fig. 8 A sequence of images (going left-to-right, top-to-bottom) showing a participant successfully putting on the sleeves of the hospital gown with assistance from the robot. This participant independently put on the right sleeve of the gown onto his/her more impaired arm and received

assistance with putting on the left sleeve onto his/her less impaired arm. The expected configurations of the participant's body is shown in Fig. 5d. The participant's arm pose differed from the expected configuration

and “It was pretty intuitive”. One participant expressed discomfort with the start position of the robot's arm: “I was uncomfortable with how close it was to my face”. One participant suggested that “Other garments would be great”, in addition to hospital gowns.

5.2.5 Results: confirmation in simulation

We also confirmed that the results of the optimization could be used for successful dressing assistance in simulation. This could potentially be used to check solutions prior to testing with people or as part of a process to further improve the solutions. We executed the trajectories of the robot's end effector holding a simulated gown in the PhysX simulator using the personalized models for each participant, pulling

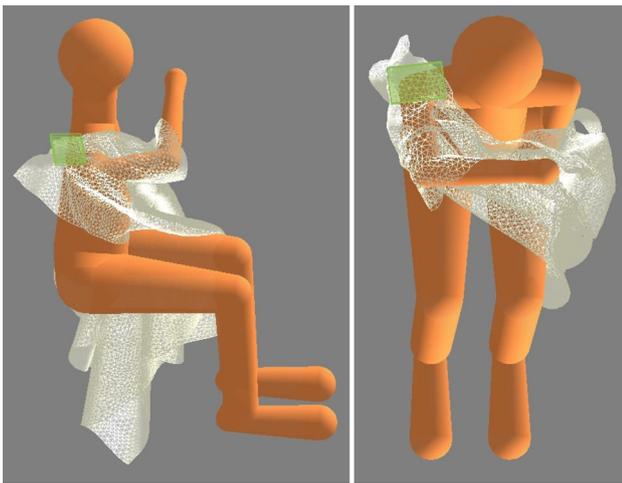


Fig. 9 The gown successfully pulled onto the right arm of a model of one of the participants using the robot's end effector trajectory and human configuration planned by TOORAD. This test was run in the PhysX simulator. The dimensions of the torso, legs, and feet were not personalized to the participant. (Left) view from the side. (Right) view from the top

each of the two sleeves onto the arms independently. In all cases, the sleeve of the gown was successfully pulled onto the human model's arm. Success was judged through manual visual inspection in the PhysX simulator. Figure 9 shows an example of the simulated robot having successfully pulled the right sleeve of the gown onto the right arm of a model of one of the participants.

6 Discussion

Although we only tested TOORAD on with a PR2 providing dressing assistance for the sleeves of a hospital gown, we expect aspects of the approach to generalize to other problems. We expect that the overall method, the equations, the modeling tools, the instructional visualization, and parts of the system would generalize to other garments, tasks, robot end effectors, and robots. Other robot models can be swapped in for the PR2 in DART and ikfast, requiring few other changes. Other robots' end effectors would require new tool models in PhysX and DART, and they might require modifications of the trajectory policies for the end effector. Other garments and dressing tasks would require new garment models in PhysX and the selection of new trajectory policies for subtasks. Many garments share properties with the hospital gown, such as requiring that a tube of cloth be pulled over an appendage and being suitable for subtasks. These garments would likely use some subtasks and trajectory policies similar to the ones used for dressing the sleeves of the hospital gown.

TOORAD's distinctive features with respect to the robot-assisted dressing literature include its optimization of the robot configuration and its use of geometric and physics simulation to perform its optimization offline with many iterations. Its objective functions allow TOORAD to plan configurations from which the robot can adapt to deviations from the planned trajectories. TOORAD could potentially be used in conjunction with other methods from the literature. For example, other methods use learning from real-world trials to perform trajectory optimization and improve dressing performance. These methods might be used to refine the trajectories planned by TOORAD after interaction with users.

6.1 Limitations

There are some limitations to our work with TOORAD and our evaluations. Although the problem formulation starts with exploring the whole space of possible actions by the robot and person, we constrain that space to achieve tractability as described in Sect. 3.1. However, actions that we removed from our actions space may be valuable for dressing. Feedback from participants suggests that our constrained action space may be preferred by some people, but that different actions may be preferred by others.

We implemented a robot-assisted dressing system with the purpose of testing TOORAD in the real world with participants with disabilities. Many avenues exist to improve that system. Use of the system required experimenter intervention to attach and reattach the robot's tool to the gown, move the robot's base, and configure the robot's arm prior to starting tasks. Previous work has addressed many of these actions (Kapusta et al. 2016a; Kapusta and Kemp 2019), and could be added to the system.

Although the method we used to model range-of-motion constraints worked sufficiently well for many of our participants, it does not capture some types of issues. For example, one of our participants experiences dependencies between joints where the range of motion of one joint depends on the configuration of another joint. We applied range-of-motion limitations to the optimization as a binary function: if the person's configuration is out of range, the configuration is discarded. Other binary functions for checking range-of-motion, such as that by Akhter and Black (2015) and Jiang and Liu (2017), could replace the range of motion modeling currently used in TOORAD.

6.2 Participant survey highlights and lessons learned

When conducting research, there is significant value to working with people from the population that the researchers intend to benefit. The six participants' responses to the survey of needs, habits, physical capabilities, and thoughts on robot-

assisted dressing could be helpful for future research. Here, we present some key takeaways and lessons learned from conducting the study with participants with disabilities.

Dressing assistance needs varied

The need for assistance for each type of garment varied by participant, with one exception. All participants indicated that they would like assistance with buttons (and many also with zippers), because they require fine motor control of both hands. Otherwise, participants wanted assistance with different garments and wanted different forms of assistance. All participants suggested that the robot should be able to help with many garments and dressing tasks.

Participants indicated consideration of the difficulty of dressing is important when selecting the garments they choose to wear. Several also suggested that if specialized garments would allow them to receive dressing assistance from a robot instead of from a human caregiver, they might be open to wearing such garments.

A robust task-execution system is helpful

We observed that having a system that could adapt the robot's movements to a dynamic environment using some sort of feedback mechanism was important for successful execution of the dressing tasks. For example, as noted in Sect. 5.2.4, one participant experienced large involuntary arm movements. Using a capacitive sensor as input in a feedback controller, the robot was able to follow the participant's arm as it moved, and successfully completed the task. Other forms of feedback and control to improve the robustness of task-execution may be valuable. We note that aspects of TOORAD increase the robustness of the solution configurations and trajectories of the person and robot. In particular, the representations of robot dexterity used in the robot configuration optimization helps plan robot configuration where the robot is likely to be able to reach poses near the planned trajectory. Feedback control without a good configurations of the person and robot might not work as well.

Keep users physically involved in the task

Feedback from all participants suggests that it is important for users to make use of their capabilities and to feel more involved. Some participants suggested that a high level of involvement would make them feel empowered and more like they are independently dressing themselves with aid from a robot. They presented this feeling in contrast to feeling dependent on dressing assistance from a human caregiver. Participants with fine control of the larger movements of their arms and body preferred that the robot hold a sleeve out for them to put their arm through, over pulling the sleeve onto their arm. Participants with decreased control of the gross movements of their arms preferred that the robot pull a sleeve onto their arm as they hold up their arm, over having the robot manipulate their body.

Personalization is important

In our study with six participants, we observed large variation in the needs and capabilities of participants. We found that our robot-assisted dressing system was only well matched to the capabilities of four of the six participants. The other two participants were unable to comfortably hold their arms raised, and the system could not manipulate their bodies to position their arms. Participants also differed on their preferred level of involvement in the task and suggested that they might want a different level of involvement depending on the day, their fatigue, and the task. It may be valuable to have different modes of assistance that depend on the needs of the particular user, and also on the current preference of the user. It may also be valuable to select a target population and design robot-assisted dressing systems and methods for that population.

Participants were interested in general-purpose mobile robots

A frequent question from participants was if the robot used to assist with dressing might be able to assist with other tasks around the house. Participants motivated this question both with a practical perspective (they desired assistance with many other tasks), as well as an economic perspective (if they were to spend money on a robot, they would want it to do many things). This view supports our approach of using a general-purpose mobile manipulator for robot-assisted dressing, which is more likely to be able to do other tasks than a fixed-base robot or a robot that has been customized solely for dressing.

Robot capabilities may limit possible assistance

Our system successfully provided dressing assistance to four of the participants by pulling one or two sleeves of a hospital gown onto the participants' bodies. However, our dressing system and our initial action policy constraints were not well suited to some participants. We suggest a few changes that may improve the ability for a dressing system to assist more participants. Our system used a PR2 robot that has weak arms with a max payload of ~ 1 kg. With this limit, it is challenging for the robot to manipulate a user's body. A stronger robot might allow other forms of assistance. The PR2 also has a relatively limited workspace. We found that the workspace sometimes causes the robot to require multiple base locations to perform a task or is insufficient to perform a task. A larger workspace may be beneficial.

Consideration of the entire dressing task is hard but may be valuable

TOORAD breaks up dressing tasks into mostly independent subtasks for computational efficiency. It additionally limits the action space of the person to holding still as the robot pulls on the garment. We observed that these two choices can limit the feasibility of some tasks. For example, stretching constraints for dressing the second arm after one is dressed may not be satisfiable for some garments or some

people with limited range of motion. Dressing both arms simultaneously or alternating dressing parts of each arm (e.g., dress left forearm, then right forearm, then left upper arm, etc.), might make dressing possible for more garments and people. We observed that simultaneously dressing both arms may be challenging for a robot with a limited workspace, like the PR2. Alternating dressing parts of each arm may require joint optimization of the two arms, which may present computational challenges.

Additional subtasks are sometimes necessary

For all participants with which we tested our system, TOORAD found that the optimal sequence of subtasks was a single subtask for pulling the sleeve onto each arm. However, we found that for some human models, such as those with longer arms, more subtasks were necessary. In part, this is because it can be difficult to fit the trajectories for dressing an entire arm into the robot's workspace at a single base location.

7 Conclusion

In this work, we have presented task optimization of robot-assisted dressing (TOORAD), a method to use optimization and simulation to plan actions for a robot to collaborate with a person with motor impairments in a dressing task. TOORAD makes use of geometric, kinematic, and physics simulation of the person, robot, and garment in its optimization. It uses personalized models of the person to match their geometries and physical capabilities and plans configurations of the robot to allow the task to be successful despite errors between simulation and the real world. TOORAD is able to explore a wide range of actions for dressing in simulation, some of which might be challenging to test in the real world, or might differ from what human caregivers do. We used TOORAD to optimize the actions of a person and a PR2 robot to collaborate in pulling the two sleeves of a hospital gown onto the person's body. We performed a study with six participants with disabilities who have difficulty dressing themselves to survey their needs, capabilities, habits, and views on robot-assisted dressing to guide current and future research in robot-assisted dressing. We implemented a robot-assisted dressing system using a PR2 robot and tested the actions planned by TOORAD by dressing four of the six participants in a hospital gown. The system used capacitive sensing to allow it to modify the trajectory planned by TOORAD to match unexpected error in the person's pose or movement of the person's arm, and TOORAD's optimization planned a robot configuration that allowed the robot to reach these end effector poses in the modified trajectory. The tasks were successful despite some participants experiencing large involuntary movements deviating from the expected arm poses. Participants were positive about how well the system was able to

track them through their movements. Our results suggest the need for personalized and robust assistance as people have varied needs and capabilities. Participants also stated that they prefer to use their capabilities and to be kept as involved in the task as possible, which further supports the need for personalized solutions. Our work provides suggestions for future directions in robot-assisted dressing research.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Akhter, I., & Black, M. J. (2015). Pose-conditioned joint angle limits for 3D human pose reconstruction. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 1446–1455).
- Anandalingam, G., & Friesz, T. L. (1992). Hierarchical optimization: An introduction. *Annals of Operations Research*, 34(1), 1–11.
- Beeson, P., & Ames, B. (2015). TRAC-IK: An open-source library for improved solving of generic inverse kinematics. In *15th International conference on humanoid robots (humanoids), 2015 IEEE-RAS, IEEE* (pp. 928–935).
- Berenson, D. (2013). Manipulation of deformable objects without modeling and simulating deformation. In *International conference on intelligent robots and systems (IROS), 2013 IEEE/RSJ, IEEE* (pp. 4525–4532).
- Bersch, C., Pitzer, B., & Kammel, S. (2011). Bimanual robotic cloth manipulation for laundry folding. In *International conference on intelligent robots and systems (IROS), 2011 IEEE/RSJ, IEEE* (pp. 1413–1419).
- Canadian Partnership for Stroke Recovery (2015). Stroke engine. <http://www.strokeengine.ca/>. Accessed 01 Jan 2018.
- Canal, G., Pignat, E., Alenyà, G., Calinon, S., & Torras, C. (2018). Joining high-level symbolic planning with low-level motion primitives in adaptive hri: Application to dressing assistance. In *2018 IEEE international conference on robotics and automation (ICRA), IEEE* (pp. 1–9).
- Clegg, A., Yu, W., Tan, J., Liu, C. K., & Turk, G. (2018). Learning to dress: Synthesizing human dressing motion via deep reinforcement learning. In *SIGGRAPH Asia 2018 technical papers, ACM* (p. 179).
- CMA (2018). The CMA evolution strategy. <https://pypi.python.org/pypi/cma>. Accessed 01 Jan 2018.
- Coumans, E., & Bai, Y. (2016–2018). PyBullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>. Accessed 1 May 2018.
- DART (2018). Dynamic animation and robotics toolkit (DART). <http://dartsim.github.io>. Accessed 01 Jan 2018.
- Diankov, R. (2010). Automated construction of robotic manipulation programs. PhD thesis, Carnegie Mellon University.
- Diankov, R., Ratliff, N., Ferguson, D., Srinivasa, S., & Kuffner, J. (2008). Bispase planning: Concurrent multi-space exploration. In *Proceedings of robotics: Science and systems IV* (p. 63).
- Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1), 269–271.

- Dogar, M. R., & Srinivasa, S. S. (2012). A planning framework for non-prehensile manipulation under clutter and uncertainty. *Autonomous Robots*, 33(3), 217–236.
- Elbanhawi, M., & Simic, M. (2014). Sampling-based robot motion planning: A review. *IEEE Access*, 2, 56–77.
- Erickson, Z., Clever, H. M., Turk, G., Liu, C. K., & Kemp, C. C. (2018a). Deep haptic model predictive control for robot-assisted dressing. In *2018 IEEE international conference on robotics and automation (ICRA)*, IEEE (pp. 1–8).
- Erickson, Z., Collier, M., Kapusta, A., & Kemp, C. C. (2018b). Tracking human pose during robot-assisted dressing using single-axis capacitive proximity sensing. *IEEE Robotics and Automation Letters*, 3(3), 2245–2252.
- Fahantidis, N., Paraschidis, K., Petridis, V., Dougeri, Z., Petrou, L., & Hasapis, G. (1997). Robot handling of flat textile materials. *IEEE Robotics & Automation Magazine*, 4(1), 34–41.
- Gao, Y., Chang, H. J., & Demiris, Y. (2015). User modelling for personalised dressing assistance by humanoid robots. In *International conference on intelligent robots and systems (IROS), 2015 IEEE/RSJ, IEEE* (pp. 1840–1845).
- Gao, Y., Chang, H. J., & Demiris, Y. (2016). Iterative path optimisation for personalised dressing assistance using vision and force information. In *International conference on intelligent robots and systems (IROS), 2016 IEEE/RSJ, IEEE* (pp. 4398–4403) IEEE. <http://hdl.handle.net/10044/1/39009>. Accessed 1 May 2018.
- Garrett, C. R., Lozano-Pérez, T., & Kaelbling, L. P. (2015). Backward-forward search for manipulation planning. In *International conference on intelligent robots and systems (IROS), 2015 IEEE/RSJ, IEEE* (pp. 6366–6373).
- Gienger, M., Janssen, H., & Goerick, C. (2005). Task-oriented whole body motion for humanoid robots. In *International conference on humanoid robots, 2005 5th IEEE-RAS, IEEE* (pp. 238–244).
- Goodin, H. J. (2003). The nursing shortage in the United States of America: An integrative review of the literature. *Journal of Advanced Nursing*, 43, 335–343.
- Grey, M. X., Garrett, C. R., Liu, C. K., Ames, A. D., & Thomaz, A. L. (2016). Humanoid manipulation planning using backward-forward search. In *International conference on intelligent robots and systems (IROS), 2016 IEEE/RSJ, IEEE* (pp. 5467–5473).
- Hansen, N. (2006). The CMA evolution strategy: A comparing review. In: J.A. Lozano, P. Larrañaga, I. Inza, E. Bengoetxea (Eds.), *Towards a new evolutionary computation: Advances in the estimation of distribution algorithms* (pp. 75–102). Berlin: Springer.
- Institute of Medicine. (2008). *Retooling for an aging America: Building the health care workforce*. Washington, DC: The National Academies Press.
- Jevtić, A., Valle, A. F., Alenyà, G., Chance, G., Caleb-Solly, P., Dogramadzi, S., & Torras, C. (2018). Personalized robot assistant for support in dressing. In *IEEE transactions on cognitive and developmental systems*.
- Jiang, Y., & Liu, C. K. (2017). Data-driven approach to simulating realistic human joint constraints. [arXiv:1709.08685](https://arxiv.org/abs/1709.08685).
- Kaelbling, L. P., & Lozano-Pérez, T. (2013). Integrated task and motion planning in belief space. *The International Journal of Robotics Research*, 32(9–10), 1194–1227.
- Kapusta, A., & Kemp, C. C. (2019). Task-centric optimization of configurations for assistive robots. *Autonomous Robots*. <https://doi.org/10.1007/s10514-019-09847-2>.
- Kapusta, A., Chitalia, Y., Park, D., & Kemp, C. C. (2016a). Collaboration between a robotic bed and a mobile manipulator may improve physical assistance for people with disabilities. In *IEEE RO-MAN workshop on behavior, adaptation and learning for assistive robotics (BALLAR)*, 2016.
- Kapusta, A., Yu, W., Bhattacharjee, T., Liu, C. K., Turk, G., & Kemp, C. C. (2016b). Data-driven haptic perception for robot-assisted dressing. In *25th IEEE international symposium on robot and human interactive communication (RO-MAN), IEEE* (pp. 451–458).
- Kim, J. O., & Khosla, K. (1991). Dexterity measures for design and control of manipulators. In *International workshop on intelligent robots and systems' 91. 'Intelligence for mechanical systems, proceedings IROS'91. IEEE/RSJ, IEEE* (pp. 758–763).
- Klee, S. D., Ferreira, B. Q., Silva, R., Costeira, J. P., Melo, F. S., & Veloso, M. (2015). Personalized assistance for dressing users. In A. Tapus, E. André, J. C. Martin, F. Ferland, & M. Ammi (Eds.), *Social robotics* (pp. 359–369). Berlin: Springer International Publishing.
- Koganti, N., Tamei, T., Ikeda, K., & Shibata, T. (2016). Bayesian nonparametric motor-skill representations for efficient learning of robotic clothing assistance. In *Workshop on practical Bayesian nonparametrics, neural information processing systems* (pp. 1–5).
- Kumar, S., Sukavanam, N., & Balasubramanian, R. (2010). An optimization approach to solve the inverse kinematics of redundant manipulator. *International Journal of Information and System Sciences (Institute for Scientific Computing and Information)*, 6(4), 414–423.
- Lindemann, S. R., & LaValle, S. M. (2005). Current issues in sampling-based motion planning. In *Robotics research* (pp. 36–54).
- Maitin-Shepard, J., Cusumano-Towner, M., Lei, J., & Abbeel, P. (2010). Cloth grasp point detection based on multiple-view geometric cues with application to robotic towel folding. In *IEEE international conference on robotics and automation (ICRA), IEEE* (pp. 2308–2315).
- Migdalas, A., Pardalos, P. M., & Värbrand, P. (2013). *Multilevel optimization: Algorithms and applications* (Vol. 20). Berlin: Springer.
- Miller, S., Van Den Berg, J., Fritz, M., Darrell, T., Goldberg, K., & Abbeel, P. (2012). A geometric approach to robotic laundry folding. *The International Journal of Robotics Research*, 31(2), 249–267.
- Mitzner, T. L., Chen, T. L., Kemp, C. C., & Rogers, W. A. (2013). Identifying the potential for robotics to assist older adults in different living environments. *International Journal of Social Robotics*, <https://doi.org/10.1007/s12369-013-0218-7>.
- OpenRAVE (2018). Open robotics automation virtual environment (OpenRAVE). <http://openrave.org>. Accessed 01 Jan 2018.
- Osawa, F., Seki, H., & Kamiya, Y. (2007). Unfolding of massive laundry and classification types by dual manipulator. *JACIII*, 11(5), 457–463.
- PhysX (2011). PhysX physics engine. <http://www.geforce.com/hardware/technology/physx>. Accessed 01 Jan 2018.
- Pignat, E., & Calinon, S. (2017). Learning adaptive dressing assistance from human demonstration. *Robotics and Autonomous Systems*, 93, 61–75.
- Sergeyev, Y. D., & Grishagin, V. A. (2001). Parallel asynchronous global search and the nested optimization scheme. *Journal of Computational Analysis and Applications*, 3(2), 123–145.
- Smits, R. (2006). KDL: Kinematics and dynamics library. <http://www.ros.org/kdl>. Accessed 01 Jan 2018.
- Stilman, M., & Kuffner, J. J. (2005). Navigation among movable obstacles: Real-time reasoning in complex environments. *International Journal of Humanoid Robotics*, 2(04), 479–503.
- Tamei, T., Matsubara, T., Rai, A., & Shibata, T. (2011). Reinforcement learning of clothing assistance with a dual-arm robot. In *International conference on humanoid robots (humanoids), 2011 11th IEEE-RAS, IEEE* (pp. 733–738).
- The Wright Stuff Inc (2017). The Wright Stuff healthcare products that make life easier. <https://www.thewrightstuff.com/>. Accessed 01 Jan 2018.
- Tilley, A. R. (2002). *The measure of man and woman: Human factors in design*. New York: Wiley.
- Topping, M. (1999). The development of handy 1, a robotic system to assist the severely disabled. In *International conference on rehabilitation robotics* (pp. 244–249).

- Topping, M., & Smith, J. (1998). The development of Handy 1, a rehabilitation robotic system to assist the severely disabled. *Industrial Robot*, 25(5), 316–20.
- Twardon, L., & Ritter, H. J. (2018). Learning to put on a knit cap in a head-centric policy space. *IEEE Robotics and Automation Letters*, 3, 764–771.
- Umali, A., & Berenson, D. (2017). A framework for robot-assisted doffing of personal protective equipment. In *2017 IEEE international conference on robotics and automation (ICRA), IEEE* (pp. 2874–2881).
- Wiener, J. M., Hanley, R. J., Clark, R., & Nostrand, J. F. V. (1990). Measuring the activities of daily living: Comparisons across national surveys. *Journal of Gerontology: Social Sciences*, 45(6), S229–237.
- Yamazaki, K., Oya, R., Nagahama, K., Okada, K., & Inaba, M. (2016). Bottom dressing by a dual-arm robot using a clothing state estimation based on dynamic shape changes. *International Journal of Advanced Robotic Systems*, 13(1), 5.
- Yoshikawa, T. (1984). Analysis and control of robot manipulators with redundancy. In *Robotics research: The first international symposium* (pp. 735–747). Cambridge: MIT Press.
- Yu, W., Kapusta, A., Tan, J., Kemp, C. C., Turk, G., & Liu, C. K. (2017). Haptic simulation for robot-assisted dressing. In *IEEE international conference on robotics and automation (ICRA), IEEE* (pp. 6044–6051).
- Zhang, F., Cully, A., & Demiris, Y. (2017). Personalized robot-assisted dressing using user modeling in latent spaces. In *International conference on intelligent robots and systems (IROS), 2017 IEEE/RSJ, IEEE* (pp. 3603–3610).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Ariel Kapusta received his Ph.D. in Robotics in 2018 as a student in the Healthcare Robotics Lab at the Georgia Institute of Technology. He earned a B.S. degree in Mechanical Engineering from Rice University in 2009. His research interests include assistive robotics, mobile manipulation, control systems for robotic manipulation, human-robot interaction, multi-modal sensing, and machine learning.



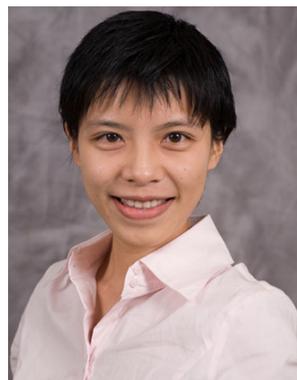
Zackory Erickson is a Ph.D. student at Georgia Institute of Technology, where he works in the Healthcare Robotics Lab and is advised by Professor Charlie Kemp. His research focuses on haptic perception and data-driven learning techniques to enable robots to more intelligently interact with and assist people. He received a BS at the University of Wisconsin-La Crosse with majors in both computer science and mathematics.



Henry M. Clever is a Ph.D. candidate at Georgia Institute of Technology, advised by Professor Charlie Kemp in the Healthcare Robotics Lab. His research involves machine learning for haptic and vision perception and human pose estimation to give robots better information to assist people with motor impairments. He received B.S. and M.S. degrees in mechanical engineering from the University of Kansas and New York University, respectively.



Wenhao Yu is a Ph.D. student in Computer Science at School of Interactive Computing, Georgia Institute of Technology advised by Karen Liu and Greg Turk. He received his master's degree from Georgia Institute of Technology in 2015 and his bachelor's degree from Shanghai Jiaotong University in 2013. His research interests are in character animation, reinforcement learning, transfer learning and robotics.



C. Karen Liu is an associate professor in School of Interactive Computing at Georgia Tech. She received her Ph.D. degree in Computer Science from the University of Washington. Liu's research interests are in computer graphics and robotics, including physics-based animation, character animation, optimal control, reinforcement learning, and computational biomechanics. She developed computational approaches to modeling realistic and natural human movements, learning complex control policies for humanoids and assistive robots, and advancing fundamental numerical simulation and optimal control algorithms. The algorithms and software developed in her lab have fostered interdisciplinary collaboration with researchers in robotics, computer graphics, mechanical engineering, biomechanics, neuroscience, and biology. Liu received a National Science Foundation CAREER Award, an Alfred P. Sloan Fellowship, and was named Young Innovators Under 35 by Technology Review. In 2012, Liu received the ACM SIGGRAPH Significant New Researcher Award for her contribution in the field of computer graphics.



Greg Turk received a Ph.D. in computer science in 1992 from the University of North Carolina at Chapel Hill. He was a postdoctoral researcher at Stanford University for two years, followed by two years as a research scientist at UNC Chapel Hill. He is currently holds the rank of Professor at the Georgia Institute of Technology, where he is a member of the School of Interactive Computing and the Graphics, Visualization and Usability Center. His research interests include computer

graphics, robotics, biological simulation and scientific visualization. He was the Technical Papers Chair for ACM SIGGRAPH 2008. In 2012 he received the Computer Graphics Achievement Award from ACM SIGGRAPH for his contributions to the field of computer graphics.



Charles C. Kemp is an Associate Professor at the Georgia Institute of Technology in the Department of Biomedical Engineering. He also has adjunct appointments in the School of Interactive Computing and the School of Electrical and Computer Engineering. He earned a doctorate in Electrical Engineering and Computer Science (2005), an M.Eng., and B.S. from MIT. In 2007, he founded the Healthcare Robotics Lab at Georgia Tech (<http://healthcarerobotics.com>), which focuses

on mobile manipulation and human–robot interaction. He is an active member of the Center for Robotics and Intelligent Machines (RIM@GT) and Georgia Tech’s multi-disciplinary Robotics Ph.D. program. He has received the 3M Non-tenured Faculty Award, the Georgia Tech Research Corporation Robotics Award, and the NSF CAREER award.