Undergraduate Thesis

Simulating Assistive Robotics Tasks

Healthcare Robotics Lab

Vamsee Gangaram

Mentor: Zackory Erickson

Research Advisor: Charles Kemp

Second Reader: Karen Liu

Table Of Contents

Abstract

In this thesis, I summarize two published research papers [1][2] to which I contributed as an undergraduate researcher. My contributions to this research primarily consisted of implementing realistic human joint limitations and better cloth visualization in Assistive Gym [1], as well as testing out various capacitive sensor designs for the multidimensional capacitive sensing work [2]. Physics-based simulation offers an opportunity for robots to learn to better provide safe and efficient assistance to people. By training robotic controllers in accurate simulations, we can drastically improve data collection and training times as compared to data collection with real robots and real people. Simulation also provides robots with a safe environment to learn, practice, and make mistakes, without having to put real people at risk. In a previous work Erickson et al. introduced Assistive Gym, a simulation framework based on the PyBullet physics engine to simulate various assistive tasks with robot and human interaction [1]. The six assistive tasks modeled are drinking, eating, itch scratching, dressing, bed bathing, and arm manipulation. We also model various human limitations as well as active human cooperation which results in better learned assistance policies. We include four common assistive robots as options for training in the six environments and show how they can be benchmarked for each assistive task. Another work from Erickson et al. on using multidimensional capacitive sensing for dressing and bathing tasks [2] is summarized, and we describe how this sensor can be modelled in simulation to incorporate into Assistive Gym in the future. Overall, Assistive Gym is shown to be an encouraging framework for training assistive robots in simulation and is open source for the research community to build upon.

Introduction

Recently, there has been an increase in the study of using simulated environments for training robotic controllers, as physics engines have significantly improved and robots are becoming increasingly more expensive to test. By training robotic controllers in accurate simulations, we can drastically improve training times by eliminating the need to physically reset an environment, such as refilling a cup with water after each trial, and can train at much higher rates than in real life. Training and testing in simulation can also save on expensive testing and degradation of the physical robot as they are usually costly machines. In addition, it allows for quicker experimentation to develop the optimal robotic controller for a task by virtually altering the robot rather than changing the physical build [3]. However, even with training the controllers in simulation, there are differences in sensing, actuation, and in dynamic interactions between the robot and environment compared to the physical world. This difference can be expressed as the reality gap [4]. Prior studies have focused on minimizing the reality gap in various ways by either improving the simulator itself or changing the network learning architecture to be more robust.

Many research works focused on simulation to reality tasks for either object recognition or locomotion tasks, however there are no works currently that focus on the problem of simulation to reality for assistive tasks with collaboration with humans such as feeding or bathing disabled patients. In order to address this gap in research, we created Assistive Gym [1], a simulated environment framework with PyBullet to be able to simulate various assistive tasks with robot and human interaction. In developing this framework, we incorporated many features that can be used the create baseline robot policies for various assistive tasks. For the dressing and bathing assistive tasks we describe how a previous work used capacitive sensing to complete these tasks in reality [2] and how recent research has been able to model capacitive sensing in simulation [5], which could be implemented in Assistive Gym for better results. We incorporate simulation environments for six key activities of daily living and four different robots, with realistic human models that can be used for co-optimization of the assistive tasks.

Literature Review

Simulation to reality training for robots has been focused on by researchers for many years due to savings on training time and expensive testing and experimentation. However there is yet to be a standard way to train robots in simulation that fare well in a real environment due to the reality gap. This is especially the case with human robot interaction for the assistive tasks we will be investigating. The main improvements that researchers have made in this field of work can be split into two categories: improving the simulation environment and creating better learning network architectures to generalize well to reality.

There are many physics engines used for simulations, however a study by Erez et al. demonstrated that the engines MuJoCo and Bullet outperform gaming oriented physics engines such as PhysX that trade physical accuracy for stability [6]. Based on this study we will use Bullet for simulation tasks as it is relatively accurate compared to the others and can also be run from a python wrapper. The physics engine itself however isn't precise enough to accurately model the inconsistencies with robotic movement in the physical environment. Tan et al. explored bridging this reality gap by demonstrating their method to learn robotic locomotion from scratch in simulation and transfer the controllers to their robot in the real world successfully [7].

There have been recent research publications that demonstrate simulation environments for various robotic tasks including manipulation, navigation, or visual tasks as seen in Fan et al.

and Savva et al.'s works [8] [9]. Zamora et al. extended OpenAI Gym to other robot simulators, which is the same learning control framework we utilize [10]. Fan et al. introduced a different simulation framework named SURREAL, which incorporates robot manipulation tasks such as block stacking. In contrast to such robot simulation frameworks, Assistive Gym aims to provide simulation environments for assistive tasks that are shown to be helpful for daily living.

Assistive Gym is developed based on many of the above works, and aims to provide a unified simulation framework that others can use to train robot and human models in for many assistive tasks. The provided environments and tasks form a basis to train policies, compare between baselines, and add new assistive tasks.

Specifically for the robot assisted dressing task, many previous works have explored using vision based systems to complete this task in reality [11] [12] [13] [14]. However a limitation with this approach is that visual occlusions from clothing could hinder the accuracy and performance of the dressing task using vision based systems. Thus capacitive sensing is also explored as a successful option for completing the dressing task. For single-axis proximity sensing a single electrode can be used to approximate distance to an object based on the parallel plate capacitor equation as explored by Erickson et al. [15], however for multidimensional sensing the sensor design needs to be adapted. In a prior work we expanded upon this capacitive sensing based approach to work for multidimensional movement in the dressing task as well as completing the bed bathing task successfully [4]. While capacitive sensing is not implemented in Assistive Gym currently, it is proven to be a viable method of completing the dressing and bed bathing task in reality with potential to be simulated as described in detail later.

Methodology

Assistive Gym

Assistive Gym is a simulation framework that we proposed for simulating robot and human assistive daily living tasks. The simulation environments in Assistive Gym are built using Pybullet, an open source physics engine, allowing for cloth and softbody simulation, variable robot and human abilities such as joint limitations, and real time visualizations. The framework is integrated with OpenAI Gym which allows for using their learning algorithms on the robots and humans created in Assistive Gym to complete tasks.

There are four robots modelled in Assistive Gym: the PR2, Sawyer, Baxter, and Jaco robots. These were picked as they are the most commonly used robots for physical human-robot interaction in the real world. We created a male and female human model based on 50th percentile values with the option to change the properties such as size and joint limitations in Assistive Gym [16]. Human limitation options are also included in Assistive Gym as the people most likely in need of assistive robotics are disabled patients. The limitations modelled are head and arm tremors, joint limitations, and joint weakness. Head and arm tremors are implemented by oscillating the joints within a 20 degree limit, joint limitations by scaling the pose independent limits for each joint by a factor between 0.25 and 1, and joint weakness by scaling the maximum torque a joint can apply by a factor between 0.5 and 1. These human models and extent of limitations are randomly selected for each of the assistive task environments.

The six assistive task environments that are modeled in Assistive Gym are based on activities of daily living that are most common for disabled people:

- Feeding: A robot's goal is to hold a spoon of food (small spheres) and navigate it into the human's mouth minimizing spillage.

- Drinking: Similar to the feeding environment however the robot's goal is to hold a cup of water (smaller spherical particles) and pour the water into the human's mouth minimizing spillage.
- Itch Scratching: A robot's goal is to rub it's end effector on a randomly identified point on the human's arm without applying more than 10N of force.
- Arm Manipulation: A human is lying in a bed with his/her arm hanging off the bed. A robot's goal is to move the human's arm onto the bed.
- Bed Bathing: A human is lying in a bed with markers on their arm to be wiped off. A robot's goal is to rub the markers off the human's arm with a washcloth tool.
- Dressing: A robot's goal is to place a hospital gown on the human by pulling the sleeve through the human's arm. (Fig. 1)

Fig 1: Baxter robot dressing a human in Assistive Gym.

With simulating human-robot interaction, it is important that the robots learn safe policies for interacting with the human as to not create discomfort. Thus we implemented realistic human joint limits which are pose dependent, based on a previous work [17]. They created a function which estimates if the current joint configuration of an arm is a valid pose or not. Jiang et al. trained a fully connected neural net to classify poses based on Akhter and Black's work which allows for much faster classification of valid poses with around 95% accuracy [18]. We implemented this neural net in Assistive Gym to model realistic joint limits by running the human model joint pose through the neural net at each time step of the simulation and if the position is invalid we set the position back to the last valid position. We believe this will allow for a more accurate pose estimation for humans and will thus encourage the robot to not learn to put the human in uncomfortable positions to achieve its goal. (Fig. 2)

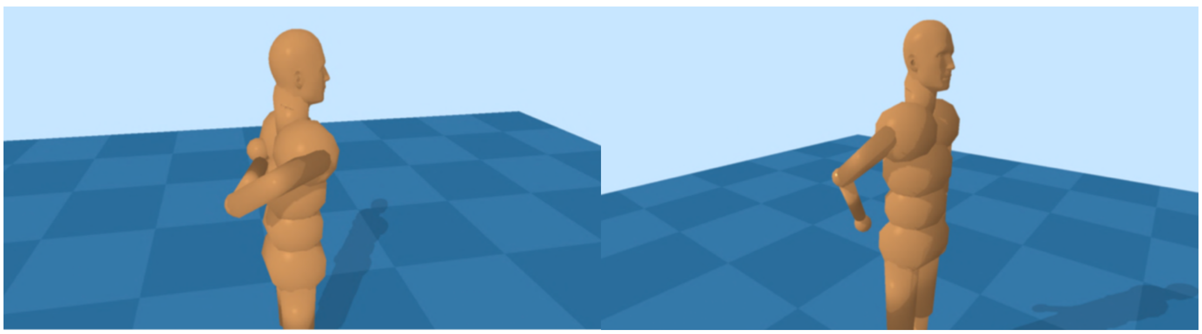

Fig 2: Left shows human pose with joint limits set independently and right shows human with realistic joint limits set which would be a much more comfortable position.

Depending on where the base position of a robot is set there can be a difference in the success rate of the robot reaching its end goal state. Thus we used task centric optimization of robot configurations (TOC) and joint limited weighted kinematic isotropy (JLWKI) as described in Kapusta et al.'s work to determine the optimal robot base position to initialize the simulation to [19]. A 100 base positions of the robots are randomly sampled in Assistive Gym for each task and the position with a collision free inverse kinematic solution to the most number of goal end effector poses is selected. In the result of a tie for possible base positions we use the position with the highest JLWKI across goal end effector poses. The base position for all robots in Assistive Gym are optimized at the start of each simulation run for all assistive tasks.

Human preferences are also accounted for in Assistive Gym as we ideally aim for the robots to learn policies that provide high quality of care to patients. For example patients would not want food spilled on them nor would they want robots to apply too much force on their body. To implement human preferences in Assistive Gym we created a reward function that represents how well the robot is performing its task in relation to the modeled human preferences described below. The robot reinforcement learning algorithm we use will attempt to maximize both the human preference reward as well as the task completion reward, ideally learning a policy that is according to the human preferences and completes the assistive task. The human preferences that we rate during simulation for the reward function are the following:

- High end effector velocities by the robot
- The robot applying force away from the target assistance location (e.g. Human mouth for feeding assistance)
- The robot applying high forces near the target assistance location
- The robot spilling food or water on the human
- Food or water entering the mouth at high velocity
- Cloth garment applying high force to the body during dressing
- The robot applying high pressure on the human using the end effector tools

To train the robots to learn policies to complete tasks in Assistive Gym we use a deep reinforcement learning algorithm called proximal policy optimization (PPO). Our implementation of the policy gradient algorithm uses a fully connected neural network of two hidden layers with 64 nodes using a tanh activation function. The observations of the environment state at each time step of the simulation is recorded by the robot to determine an action and receive a reward for the next state. For collaboration between both the human and robot to complete an assistive task in simulation we train a policy for each concurrently using the reward functions for human preference and task completion.

Multidimensional Capacitive Sensing

In a separate work from developing Assistive Gym, we trained a PR2 robot to complete dressing and bathing tasks using multidimensional capacitive sensing [4]. Traditionally, capacitive sensing is where electrodes create an electric field and when this field is interrupted by an external conductive material then energy is transferred to this external conductor resulting in an increase in capacitance of the electrodes [20]. Capacitive sensing can also sense through dielectric material to the human body which is conductive [21]. Thus the capacitance measure from a capacitive sensor is inversely proportional to the distance of the sensor from the human body.

In the multidimensional capacitive sensing work we designed the sensor to have six electrodes arranged in a flat grid structure with roughly 15cm sensing distance. (Fig. 3) This enables the sensing of not only vertical and horizontal distance but also pitch and yaw angle of the sensor with respect to the human limb. To convert the capacitance measurements into distance and orientation measurements, we attached the sensor onto a PR2 robot and collected training data over participants' arms and legs to train a connected neural network with four hidden layers each with 400 nodes and ReLU activation. This model takes a time window of capacitance data as input and outputs the estimated position and orientation of the sensor with respect to a person's limb. We were then able to run dressing and bathing trials where the goal was to pull a gown onto a participant's arm and wipe a participant's arm/leg respectively. With the participants gradually moving their arm vertically or horizontally during each dressing trial, the PR2 was successfully able to dress all participants with minimal contact. For the bathing task the robot was successfully able to wipe the full length of a participant's arm 7 out of 8 times and 6 out of 8 times for their leg [4].

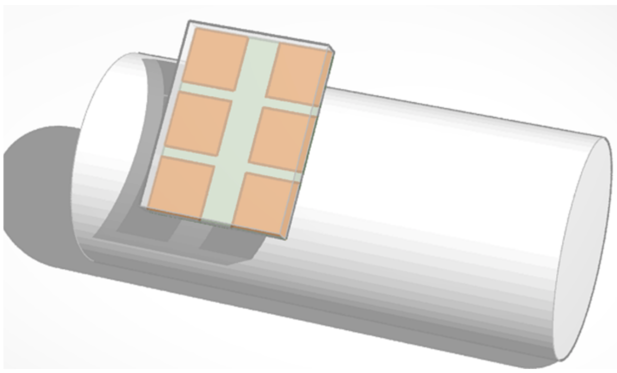

Fig 3: Diagram of multidimensional capacitive sensor over a cylinder representing an arm/leg

In a recent work Clegg et al. explored training robot and human control policies in simulation to perform a dressing task, and in doing so also simulated capacitive sensing for this task [5]. To simulate a capacitive sensor, Clegg et al. modeled six electrodes attached to the robot end effector with a proximity detection sphere of 15cm based on our physical sensor. The simulated capacitive sensor was shown to have a beneficial impact on the trained simulated dressing task with a 99% success rate compared to 79% without the capacitive sensor. While capacitive sensing isn't included in Assistive Gym yet, the demonstrated success of simulating capacitive sensing for robot assisted dressing by Clegg et al. shows that implementing this feature could prove useful in improving success of the robot assisted tasks.

Evaluation

Here we discuss the policies trained for the six assistive tasks using four different robot models. We trained each policy using PPO, with a total of 50,000 simulation trials to iteratively improve the policy. The policies were trained on Amazon Web Services (AWS), where the training times varied from two hours for the itch scratching task and close to six days due to the computation complexity of simulating dynamic cloth movement.

In the first scenario where the human model assumed a static pose, we trained policies for each of assistive tasks with four different robots resulting in 24 total policies. With PPO all the robots were able to learn reasonably successful policies for the itch scratching, feeding, drinking, and bed bathing tasks. (Fig. 3) However for the dressing and arm manipulation tasks we see that it is much more challenging for any of the robots to successfully complete the tasks. The dressing task can be successfully completed when the arm hole of the cloth is positioned close to the human model's fist, and the arm manipulation task had little success since it tried to use the edge of its end effector resulting in high pressure applied to the arm of the human model and thus incurring high penalties. With these policies we can compare the success rates as well as the average rewards for each robot to see which one would fare well for different types of assistive tasks. For example we can infer that the PR2 would perform worse than any of the other robots in reality for the itch scratching task as its simulation reward score is much lower than the others.

Task	PR2	Jaco	Baxter	Sawyer	Success
Itch Scratching	55.1	280.8	225.4	136.8	54%
Bed Bathing	86.7	104.4	88.4	109.0	24%
Feeding	100.5	83.8	108.5	95.6	88%
Drinking	182.5	85.7	263.3	436.0	72%
Dressing	11.5	-17.0	5.6	-27.6	27%
Arm Manipulation	-162.4	-177.5	-228.1	-210.6	8%

Fig 3: Table of average reward for each of the robot polices for the six different tasks with a static human model.

For the second scenario we trained the robot control policies with the assumption that patients generally have some motor function and would want to help the robot to complete their tasks since it directly benefits them. We model this by using co-optimization where the human and robot models are trained simultaneously. In training both control policies we use the same reward function across both during PPO however each agent/model has their own observation and action sets (e.g. the robot can see its own joint angles and same with the human model). We retrain these policies across the different robots and assistive tasks resulting in Fig. 4. In most cases, excluding the arm manipulation task, the success rate of the trained policies is significantly higher than when trained without human modelled collaboration. The arm manipulation task wasn't trained for in this scenario as we assume patients that have motor movement will be able to lift their arm back onto their bed.

Task	PR2	Jaco	Baxter	Sawyer	Success
Itch Scratching	80.9	443.2	83.3	131.2	68%
Bed Bathing	90.2	193.6	175.5	166.2	81%
Feeding	122.8	106.1	108.3	112.5	99%
Drinking	493.4	402.6	466.8	464.0	79%
Dressing	-1.3	13.0	30.0	56.9	89%

Figure 4: Table of average reward for each of the robot polices for the six different tasks with a human model collaboration.

Conclusion

In developing Assistive Gym, we show that the simulation framework can be used to train robots for multiple assistive tasks and is focused on physical human robot interaction. There are four robots and six assistive tasks implemented, which allow for benchmarking between robots. In addition, Assistive Gym allows for new environments and assistive tasks to be programmed as well as the testing of robot learning algorithms. With human preferences and accurate human joint modelling programmed in, Assistive Gym shows promise for simulating human collaboration with robot assistive tasks. In the future we can possibly implement simulated capacitive sensing and even transfer the controllers trained in Assistive Gym to robots in the real world as described in Clegg et al.'s work [5]. We plan to release Assistive Gym as an open source framework with the hopes of the research community to further the development and improvement of this work.

Works Cited

[1] Erickson, Z., Gangaram, V., Kapusta, A., Liu, C. K., & Kemp, C. C. (2019). Assistive Gym: A Physics Simulation Framework for Assistive Robotics. *arXiv preprint arXiv:1910.04700*.

[2] Erickson, Z., Clever, H. M., Gangaram, V., Turk, G., Liu, C. K., & Kemp, C. C. (2019, June). Multidimensional capacitive sensing for robot-assisted dressing and bathing. In *2019 IEEE 16th International Conference on Rehabilitation Robotics (ICORR)* (pp. 224-231). IEEE.

[3] Timperley, C. S., Afzal, A., Katz, D. S., Hernandez, J. M., & Le Goues, C. (2018, April). Crashing simulated planes is cheap: Can simulation detect robotics bugs early?. In *2018 IEEE 11th International Conference on Software Testing, Verification and Validation (ICST)* (pp. 331-342). IEEE.

[4] Nolfi, S., Bongard, J., Husbands, P., & Floreano, D. (2016). Evolutionary robotics. In *Springer Handbook of Robotics* (pp. 2035-2068). Springer, Cham.

[5] Clegg, A., Erickson, Z., Grady, P., Turk, G., Kemp, C. C., & Liu, C. K. (2020). Learning to Collaborate From Simulation for Robot-Assisted Dressing. *IEEE Robotics and Automation Letters*, *5*(2), 2746-2753.

[6] Erez, T., Tassa, Y., & Todorov, E. (2015, May). Simulation tools for model-based robotics: Comparison of bullet, havok, mujoco, ode and physx. In *2015 IEEE international conference on robotics and automation (ICRA)* (pp. 4397-4404). IEEE.

[7] Tan, J., Zhang, T., Coumans, E., Iscen, A., Bai, Y., Hafner, D., ... & Vanhoucke, V. (2018). Sim-to-real: Learning agile locomotion for quadruped robots. *arXiv preprint arXiv:1804.10332*.

[8] Fan, L., Zhu, Y., Zhu, J., Liu, Z., Zeng, O., Gupta, A., ... & Fei-Fei, L. (2018, October). Surreal: Open-source reinforcement learning framework and robot manipulation benchmark. In *Conference on Robot Learning* (pp. 767-782).

[9] Savva, M., Kadian, A., Maksymets, O., Zhao, Y., Wijmans, E., Jain, B., ... & Parikh, D. (2019). Habitat: A platform for embodied ai research. In *Proceedings of the IEEE International Conference on Computer Vision* (pp. 9339-9347).

[10] Zamora, I., Lopez, N. G., Vilches, V. M., & Cordero, A. H. (2016). Extending the openai gym for robotics: a toolkit for reinforcement learning using ros and gazebo. *arXiv preprint arXiv:1608.05742*.

[11] Pignat, E., & Calinon, S. (2017). Learning adaptive dressing assistance from human demonstration. *Robotics and Autonomous Systems*, *93*, 61-75.

[12] Klee, S. D., Ferreira, B. Q., Silva, R., Costeira, J. P., Melo, F. S., & Veloso, M. (2015, October). Personalized assistance for dressing users. In *International Conference on Social Robotics*(pp. 359-369). Springer, Cham.

[13] Koganti, N., Tamei, T., Ikeda, K., & Shibata, T. (2017). Bayesian nonparametric learning of cloth models for real-time state estimation. *IEEE Transactions on Robotics*, *33*(4), 916-931.

[14] Tamei, T., Matsubara, T., Rai, A., & Shibata, T. (2011, October). Reinforcement learning of clothing assistance with a dual-arm robot. In *2011 11th IEEE-RAS International Conference on Humanoid Robots* (pp. 733-738). IEEE.

[15] Erickson, Z., Collier, M., Kapusta, A., & Kemp, C. C. (2018). Tracking human pose during robot-assisted dressing using single-axis capacitive proximity sensing. *IEEE Robotics and Automation Letters*, *3*(3), 2245-2252.

[16] Tilley, A. R. (2001). *The measure of man and woman: human factors in design*. John Wiley & Sons.

[17] 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).

[18] Jiang, Y., & Liu, C. K. (2018, May). Data-driven approach to simulating realistic human joint constraints. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*(pp. 1098-1103). IEEE.

[19] Kapusta, A., & Kemp, C. C. (2019). Task-centric optimization of configurations for assistive robots. *Autonomous Robots*, *43*(8), 2033-2054.

[20] Barrett, G., & Omote, R. (2010). Projected-capacitive touch technology. *Information Display*, *26*(3), 16-21.

[21] Cheng, J., Bannach, D., & Lukowicz, P. (2008, June). On body capacitive sensing for a simple touchless user interface. In *2008 5th International Summer School and Symposium on Medical Devices and Biosensors* (pp. 113-116). IEEE.