Grasp Intention Interpretation in Object Handover for Human-Robot Teaming

Hui Li, Akhlak Uz Zaman, and Hongsheng He(✉)

The University of Alabama, Tuscaloosa, AL 35487, USA
{hli98,azaman2}@crimson.ua.edu, hongsheng.he@ua.edu

Abstract. Effective human-robot collaboration requires social robots to adapt to individual human grasping habits to ensure smooth and safe object handovers. However, current robotic systems struggle to interpret diverse grasping behaviors, as individual habits can introduce variations even within the same grasp topology. This limitation affects the effectiveness of robotic systems in social contexts. This paper presents a grasp adaptation algorithm that enables robots to recognize and adjust to human grasping habits. The system identifies human grasping poses from RGB images and maps them to abstract representations consisting of 21 3D points each. These representations are then classified into one of six standard grasp topologies. Based on the identified topology, key points are selected from the abstract grasp to estimate the object's pose. A reinforcement learning model is subsequently employed to optimize the object handover process. Experimental results demonstrate that this approach significantly enhances both the fluidity and safety of human-robot object handovers.

Keywords: grasping habits adaptation · dexterous grasping · grasp topology · deep learning · reinforcement learning

1 Introduction

Human-robot collaboration, particularly in the context of object handovers, is essential for ensuring smooth and effective interactions in both structured and unstructured environments, such as healthcare and industrial settings [1,2,18]. For example, a robot delivering tools to a nurse or handing parts to a factory worker must adapt to the human's grasping preferences and workspace constraints to enhance both safety and efficiency. Extensive research has been conducted to address the handover task in human-robot collaboration. Studies have analyzed the trajectory and velocity of approach movements to ensure smooth transitions [9,13]. Additionally, object orientation and affordances have been optimized to make it easier for the receiver to grasp the object [4,15]. Some approaches also involve learning from human behavior to improve the naturalness and effectiveness of handovers [4]. However, less work has been done to

This research was funded by NSF grant #2420355 and #2402466.

H. Li et al. (Eds.): ICSR + InnoBiz 2024, LNAI 15170, pp. 346–354, 2025.
https://doi.org/10.1007/978-981-96-1151-5_35

adapt to the habits of the receiver, particularly the variability in human grasping behaviors, which are influenced by individual preferences and situational factors. Developing systems that can accurately recognize and adapt to these diverse human behaviors is crucial for making robots more intuitive and practical in real-world applications.

Grasp topology and taxonomy are key to understanding human grasping behaviors and robotic adaptation. Grasp topology describes the geometric configuration of the fingers or contact points with an object, such as pinching or cupping. Grasp taxonomy categorizes these topologies into structured classes based on factors like contact points and object shape. This classification aids in designing robots that can effectively recognize and adapt to human grasping behaviors in various tasks and environments. Previous research has shown that human grasp choices tend to cluster over a large set of objects, leading to the development of grasp taxonomies to simplify grasping choices. For example, Cutkosky's taxonomy identified 16 grasp types used by machinists [6], and Feix's taxonomy expanded this to 33 different grasp types [3,7].

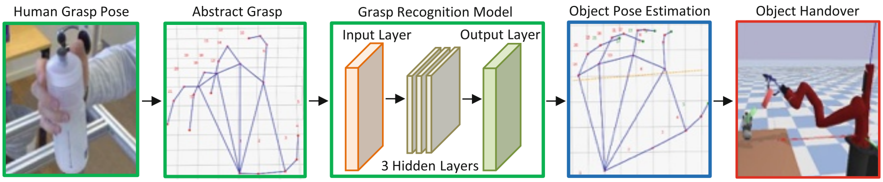

Fig. 1. Structure of the grasp adaptation system.

The FreiHAND dataset provides a large collection of annotated 3D hand poses, high-resolution RGB images, and key point annotations, which make it an invaluable resource for advancing research in hand tracking and grasp recognition [20]. In this paper, we extend the FreiHAND dataset to map individual grasping habits to a standard set of grasp topologies.

Reinforcement learning (RL) is a powerful approach for robot control, where robots learn to perform tasks through trial and error [8,16,19]. By receiving feedback in the form of rewards or penalties based on their actions, robots improve their behavior over time [17]. This method enables robots to develop adaptive and optimized control strategies for complex and dynamic environments. It enhances their ability to perform a wide range of tasks autonomously. In this paper, an RL model is designed to conduct the object handover task in a simulation environment, identical to the MagicHand system [10,11].

We propose a grasp adaptation algorithm, as illustrated in Fig. 1, that processes an RGB image of a human grasping pose, converts it into an abstract grasp representation, and classifies it into one of six standard grasp topologies. Key points are then selected from the abstract grasp, based on the identified grasp topology, to estimate the appropriate object pose. A reinforcement learning model is subsequently employed to optimize the object handover process. The contributions of this research include:

- We developed a grasp adaptation system that adjusts the robot's grasping strategy based on individual human grasping habits, enhancing the fluidity and safety of object handovers.
- We proposed a method for accurately estimating the appropriate object pose based on the identified grasp topology.
- We designed a reinforcement learning model to optimize the object handover process by learning adaptive strategies that align the object's position and orientation with the receiver's grasp pose while minimizing interaction errors

2 Human Grasping Habit Adaptation

To adapt to human grasping habits, three key challenges must be addressed: recognizing the human grasp, determining the object pose based on the grasp, and moving the object to the desired pose. The proposed system tackles these challenges through three integrated models: a grasp recognition model, an object pose estimation model, and a reinforcement learning model.

2.1 Recognition of Human Grasp Topology

Standard Grasp Topology. We classify grasp poses into six distinct grasp topologies [14], as illustrated in Fig. 2. The figure categorizes grasps into two main types: power grasps and precision grasps, based on object shapes and the involvement of virtual fingers (VF). Power grasps (e.g., circular or prismatic objects) prioritize security and stability, utilizing more virtual fingers and 3D wrapping. In contrast, precision grasps focus on dexterity and sensitivity, typically involving fewer virtual fingers and 2D wrapping.

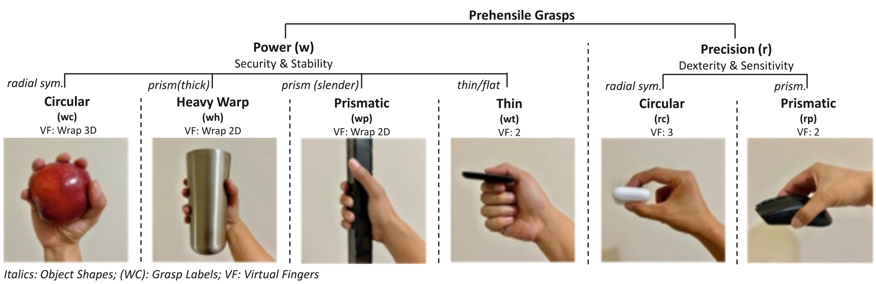

Fig. 2. The predefined grasp topology.

After establishing the grasp taxonomy, a standard grasp pose was selected for each topology by evaluating poses from the FreiHAND dataset. The evaluation was based on how well each pose matched the defined topologies. The pose that best fit each topology was chosen as the standard grasp pose. Figure 3 displays both the RGB image of the standard grasp pose and the corresponding grasp information, which includes abstract grasp key points (highlighted in green). The orange dotted axis represents the estimated object pose derived from the key points.

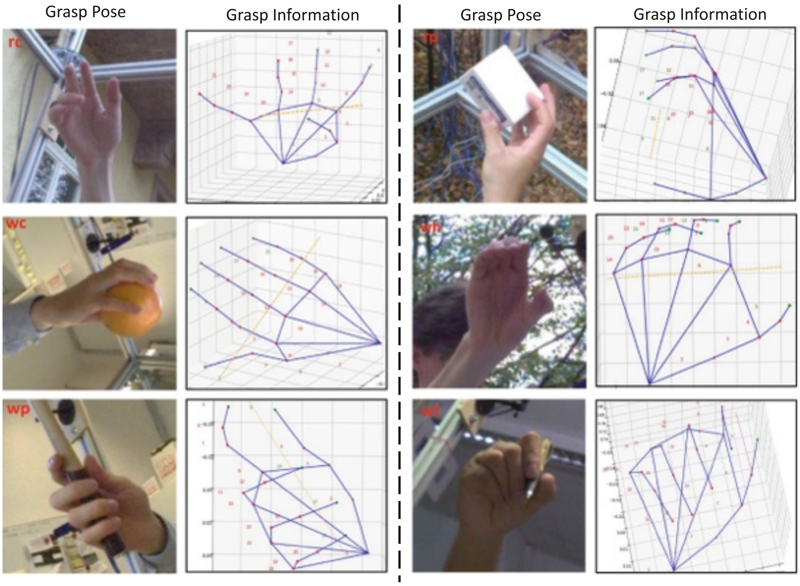

Fig. 3. Standard grasp topology: This figure displays images of the grasp topology, including the skeletons of each standard grasp topology and their corresponding key points, highlighted in green.

Grasp Topology Recognition. A multi-layer perceptron (MLP) deep neural network was developed to map abstract grasps to standard grasp topologies. Rectified Linear Units (ReLU) were used as activation functions in the input and hidden layers. The input layer comprises 63 neurons, corresponding to the 21 3D points in each abstract grasp. The network features three hidden layers with 1,024, 256, and 32 nodes, respectively. The output layer consists of six neurons with Softmax activation functions to classify the grasps. The model was trained using a refined FreiHAND dataset.

2.2 Object Pose Estimation

The pose of the object, including both position and orientation, is determined based on the key points associated with the standard grasp topology. For grasp topologies such as “wc”, “wh”, “wp”, “rc”, and “rp”, the object should be positioned within the grasp’s aperture, which is the space between the fingertips of the thumb and the fingers. The object position is defined as the midpoint between p_t, the point representing the tip of the thumb, and p_c, the closest fingertip to p_t. The object position is expressed as $p_m = 0.5[x_t + x_c, y_t + y_c, z_t + z_c]^T$. Since the orientation is typically aligned with the palm or fingers, two key points, p_s and p_e, are pre-selected from the abstract grasp to define the object’s orientation. These points are usually located on the palm. The orientation of the object is expressed as

$$\mathbf{l}(t) = [x_m, y_m, z_m]^T + t \cdot [x_e - x_s, y_e - y_s, z_e - z_s]^T \tag{1}$$

where x, y, and z are coordinates of the point and t is a scalar parameter. The points p_s, p_e, p_c, and p_t, are specific to each grasp topology and serve as key points provided for analysis.

For the grasp topology “wt” the object should be positioned between the fingertip of the thumb and the side of the index finger, specifically at the key point p_{pip}, which corresponds to the PIP joint (Proximal Interphalangeal Joint) of the index finger. The object’s position for this grasp topology is expressed as $p_m = 0.5[x_t + x_{\text{pip}}, y_t + y_{\text{pip}}, z_t + z_{\text{pip}}]^T$. The orientation of the object should be roughly parallel to the index finger. In this case, the key points p_s and p_e in (1) represent the PIP and MCP (Metacarpophalangeal) joints of the index finger, respectively.

2.3 Object Handover Using Reinforcement Learning

Once the object’s pose is determined, the robot must adjust the object to the specified position and orientation. We designed and developed a reinforcement learning model to achieve this goal. A simulation environment mirroring the MagicHand platform, which supports a variety of manipulation tasks, has been established [10,12].

The task is to handover the object to a human hand, simulated by a Schunk anthropomorphic robotic hand in a simulation environment, and positioning it at the target pose, represented by a green area. The action space of the proposed model is a six-dimensional vector, including movements and rotations of the robotic hand along the x, y, and z axes. The observation space consists of seven dimensions: the relative position and orientation between the object and the target, as well as the distance between them.

In this task, our goal is to position the object as close as possible to the target, defined as the center point of the green area, rewarding smaller distances between the object and the target. Additionally, we aim to align the object’s orientation with the target’s orientation, rewarding smaller differences in orientation along the x, y, and z axes. To ensure safe interaction, we impose penalties for collisions with the human hand. The reward function is expressed as

$$r = e^{-|d|} + e^{-|h|} + e^{-|l|} + e^{-|k|} - \alpha n_{\text{c}} \tag{2}$$

where d is the distance between the object and the target location, and h, l, and k represent the differences in orientation between the object and the target along the x, y, and z axes, respectively. The coefficient α is a constant, and n_c represents the number of contact points with the human hand. We chose an exponential function because its value changes more rapidly when the exponent is large and more slowly when the exponent is small. This approach encourages the robot to make larger adjustments when the current pose is far from the target, while allowing for more precise, gradual adjustments as the pose approaches the target.

The proximal policy optimization (PPO) algorithm is employed to train the model. PPO is a reinforcement learning algorithm designed to improve policy stability by limiting the size of policy updates. It uses a clipped objective function to ensure that the new policy does not deviate excessively from the old policy. The objective function is given by

$$J(\theta) = \mathbb{E}\left[\min\left(\frac{\pi_\theta(a|s)}{\pi_{\theta_{\text{old}}}(a|s)}\hat{A}(s,a), \text{clip}\left(\frac{\pi_\theta(a|s)}{\pi_{\theta_{\text{old}}}(a|s)}, 1-\epsilon, 1+\epsilon\right)\hat{A}(s,a)\right)\right]$$

where $\pi_\theta(a|s)$ and $\pi_{\theta_{\text{old}}}(a|s)$ are the probabilities of taking action a in state s under the new and old policies, respectively, and $\hat{A}(s,a)$ is the advantage function. The clipping function clip restricts the ratio of the new to old policy probabilities, with ϵ controlling the extent of the allowed change, thereby balancing exploration with stability.

3 Experiments

The proposed system was evaluated under different thresholds. The simulation environment was set up using PyBullet and Gym, and task simulations were conducted to test the final performance of the system.

3.1 Data Preparation

The FreiHAND dataset, comprising RGB images with 3D annotations, was refined by labeling each grasp pose with one of six predefined topologies. This updated dataset contains 600 annotated poses (100 per topology). Figure 4 illustrates sample images of these labeled poses alongside their abstract grasp representations.

3.2 Grasp Recognition

The revised dataset is divided into two parts: 540 grasp poses for training and validation, and 60 new grasp poses for testing. The proposed algorithm was evaluated using 4-fold cross-validation with hyperparameters of a batch size of 64, 500 epochs, and the Adam optimizer with a learning rate of 0.001. The training accuracy achieved 93.3% while the accuracy on testing set achieved 87.2%. The testing accuracy is relatively low because some of the grasp topologies are difficult to distinguish. For example, the "wh" and "rc" grasps have similar configurations, making them harder to recognize accurately.

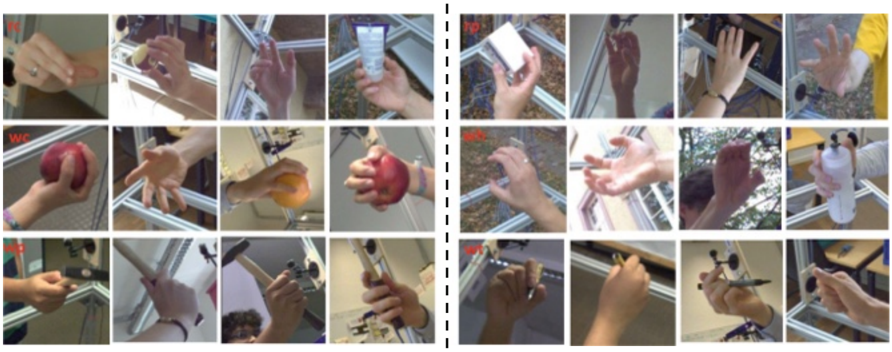

Fig. 4. The revised dataset consists of 600 grasp poses. Each pose is paired with a corresponding abstract grasp representation for each grasp topology.

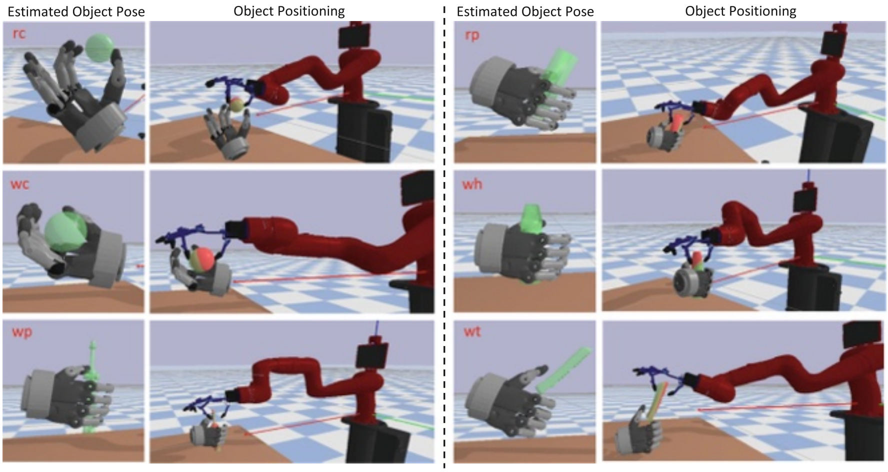

Fig. 5. Handover task for each grasp topology: In each task, the system first estimates the object pose based on the grasp pose of the simulated human hand. The model then attempts to place the object at the target pose.

3.3 Object Pose Estimation and Handover

The effectiveness and accuracy of an estimated object pose are evaluated through object-handover tasks in a simulation environment. An object pose is considered effective if the robot can successfully pass the object to the human hand, which should then be able to securely grasp it. For this evaluation, we used PyBullet [5] to simulate an AR10 robotic hand mounted on a Sawyer robot holding the object. The human hand, simulated by a Schunk robotic hand, positioned in front of the robot, performs a variant of one of the six grasp topologies. The system estimates the object pose based on the grasp pose of the simulated human hand and highlights the estimated pose as a green area. The handover task for each grasp topology is illustrated in Fig. 5, where the robot's goal is to position the object to align with the green area while avoiding collisions with the human hand.

The model was trained for 20,000 episodes with a learning rate of 1.6×10^{-6} and a batch size of 32, with varying target positions and orientations in each episode. Each grasp topology was tested 100 times achieving an overall success rate of 83%.

4 Conclusions

In conclusion, this paper presents a robust approach to enhancing human-robot cooperation through the development of a grasp adaptation system. By accurately recognizing diverse human grasping habits and classifying them into standard grasp topologies, the system can determine optimal object handover strategies for smooth handovers. The use of deep learning for grasp pose recognition and reinforcement learning for strategy optimization demonstrated strong performance in experimental settings. Although challenges remain in distinguishing similar grasp configurations, the results underscore the potential of the proposed system to improve human-robot interaction by making robots more adaptable.

References

1. Ajoudani, A., Zanchettin, A.M., Ivaldi, S., Albu-Schäffer, A., Kosuge, K., Khatib, O.: Progress and prospects of the human-robot collaboration. Auton. Robot. **42**, 957–975 (2018)
2. Billard, A., Kragic, D.: Trends and challenges in robot manipulation. Science **364**(6446), eaat8414 (2019)
3. Bullock, I.M., Feix, T., Dollar, A.M.: The Yale human grasping dataset: grasp, object, and task data in household and machine shop environments. Int. J. Robot. Res. **34**(3), 251–255 (2015)
4. Cakmak, M., Srinivasa, S.S., Lee, M.K., Forlizzi, J., Kiesler, S.: Human preferences for robot-human hand-over configurations. In: 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 1986–1993. IEEE (2011)
5. Coumans, E., Bai, Y.: Pybullet, a python module for physics simulation for games, robotics and machine learning (2016–2019). http://pybullet.org
6. Cutkosky, M.R., et al.: On grasp choice, grasp models, and the design of hands for manufacturing tasks. IEEE Trans. Robot. Autom. **5**(3), 269–279 (1989)
7. Feix, T., Romero, J., Schmiedmayer, H.B., Dollar, A.M., Kragic, D.: The grasp taxonomy of human grasp types. IEEE Trans. Hum.-Mach. Syst. **46**(1), 66–77 (2015)
8. He, W., Gao, H., Zhou, C., Yang, C., Li, Z.: Reinforcement learning control of a flexible two-link manipulator: an experimental investigation. IEEE Trans. Syst., Man Cybern.: Syst. **51**(12), 7326–7336 (2020)
9. Huber, M., Rickert, M., Knoll, A., Brandt, T., Glasauer, S.: Human-robot interaction in handing-over tasks. In: RO-MAN 2008-the 17th IEEE International Symposium on Robot and Human Interactive Communication, pp. 107–112. IEEE (2008)

10. Li, H., Tan, J., He, H.: MagicHand: context-aware dexterous grasping using an anthropomorphic robotic hand. In: 2020 IEEE International Conference on Robotics and Automation (ICRA), pp. 9895–9901. IEEE (2020)
11. Li, H., Yihun, Y., He, H.: MagicHand: in-hand perception of object characteristics for dexterous manipulation. In: International Conference on Social Robotics, pp. 523–532. Springer (2018)
12. Li, H., Zhang, Y., Li, Y., He, H.: Learning task-oriented dexterous grasping from human knowledge. In: 2021 IEEE International Conference on Robotics and Automation (ICRA), pp. 6192–6198. IEEE (2021)
13. Mason, A.H., MacKenzie, C.L.: Grip forces when passing an object to a partner. Exp. Brain Res. **163**, 173–187 (2005)
14. Rao, B.: Learning robotic grasping strategy based on natural-language object descriptions (2018)
15. Strabala, K., et al.: Toward seamless human-robot handovers. J. Hum.-Robot Inter. **2**(1), 112–132 (2013)
16. Stulp, F., Theodorou, E.A., Schaal, S.: Reinforcement learning with sequences of motion primitives for robust manipulation. IEEE Trans. Rob. **28**(6), 1360–1370 (2012)
17. Sutton, R.S.: Reinforcement learning: an introduction (2018)
18. Tavakoli, M., Carriere, J., Torabi, A.: Robotics, smart wearable technologies, and autonomous intelligent systems for healthcare during the COVID-19 pandemic: an analysis of the state of the art and future vision. Adv. Intell. Syst. **2**(7), 2000071 (2020)
19. Thuruthel, T.G., Falotico, E., Renda, F., Laschi, C.: Model-based reinforcement learning for closed-loop dynamic control of soft robotic manipulators. IEEE Trans. Rob. **35**(1), 124–134 (2018)
20. Zimmermann, C., Ceylan, D., Yang, J., Russell, B., Argus, M., Brox, T.: FreiHand: a dataset for markerless capture of hand pose and shape from single RGB images. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 813–822 (2019)