Robot-to-Human Construction Tool Handover Grasp Prediction for 6-DOF Robotic Arm with Parallel Gripper

Yu-Lun Kuo[1], Guan-Yong Xiong[2], Jacob J. Lin[1] *, Ci-Jyun Liang[2]

[1]Department of Civil Engineering, National Taiwan University, Taiwan
[2]Department of Civil Engineering, Stony Brook University, USA

R12521717@ntu.edu.tw, R12521614@ntu.edu.tw, jacoblin@ntu.edu.tw, Ci-Jyun.Liang@stonybrook.edu

Abstract -

Workers often spend significant time locating tools on construction sites, resulting in productivity losses. With the continuous advancement of robotic technologies, robot-to-human handover offers a potential solution to address this issue. However, existing robot handover methods face challenges in handling the irregular shapes and dangerous nature of tools, as well as the dynamic and complex environments of construction sites. To tackle these challenges, this research proposes a method that enables robots to effectively predict and grasp construction tools while ensuring safety. A grasp dataset consisting of four commonly used construction tools—pliers, screwdrivers, wrenches, and safety glasses—was generated in a simulation environment. During the annotation process, grasping areas for robots and humans are distinguished to reduce potential risks in handover processes. The dataset includes multi-angle RGB-D images and detailed annotated grasp poses specifically designed for robotic grasps. This proposed dataset provides a foundation for training and validating grasp prediction models, enabling robots to handle tools with irregular shapes in various scenarios. The proposed approach integrates data annotation, scene generation, grasp pose prediction, and motion planning to achieve safe and efficient handovers. Validation experiments conducted in a simulation environment using a UR5 robotic arm equipped with a parallel gripper demonstrate a 92.5% success rate. This research provides a foundation for the application of robot-to-human handovers on construction sites and highlights the potential of using virtual datasets to address real-world challenges.

Keywords -

Robot-to-human Handover; 6-DOF Grasp Dataset; Robotic arm; Grasp pose prediction

1 Introduction

The construction industry is facing labor shortage [1], and workers on construction sites often spend significant time searching for tools, further reducing productivity [2]. With the continuous advancement of robotic technologies, more construction sites are adopting automation and robotic systems to improve efficiency [3]. For example, an increasing number of construction robots have been deployed on construction sites [4]. In this context, robot-to-human handover has become a critical approach to addressing this issue.

Robot-to-human handover consists of three main steps: object searching, localization, grasping, and delivering tools. The entire process must ensure both efficiency and safety [5]. However, construction sites are highly dynamic and complex environments. Most current construction robots perform pre-programmed tasks, making them less adaptable to the variability of construction sites [6]. Compared to the operational environments of existing robot-to-human handover methods, construction sites present greater challenges in terms of environmental complexity, requiring robots to adapt to diverse scenarios and respond flexibly [7]. In addition, construction tools vary widely in size and shape, often exhibiting irregular geometries and potential hazards. Ensuring safety during the handover process is therefore critical.

To address the issues mentioned above, this paper proposes a grasp prediction method for hazardous construction tools. We develop a 6-DOF grasp pose dataset and introduce a method to distinguish between robotic and human grasping regions. The proposed method integrates a grasp dataset from GraspNet [8], grasp pose prediction, and motion planning, enabling robots to effectively grasp handheld tools while ensuring safety. The rest of this paper is organized as follows. Section 2 reviews related work on grasp pose prediction and robot-to-human object handover, identifying research gaps. Section 3 describes the proposed model architecture. Section 4 outlines the experimental setup. Section 5 presents the results. Section 6 concludes the paper and discusses future work.

2 Related Work

This study addresses the challenges of safe and efficient grasping and handover in human-robot interaction, particularly on construction tools, which involve higher safety risks. Consequently, two main research areas are reviewed: (1) Grasp Pose Prediction and (2) Robot-to-Human Handover. We focus on these domains because the quality of a robot's initial grasp directly influences the risk of subsequent handovers—an issue of critical importance in high-risk environments such as construction sites. By examining previous methods, we establish a foundation for our proposed approach to ensure efficient tool handovers through accurate and context-aware grasp prediction.

2.1 Grasp pose prediction

Grasp pose prediction is a critical component of robotic manipulation, especially in tasks that require safe and stable object handovers. Several methods have been proposed to tackle this problem in recent years. Traditional methods often rely on geometric heuristics or predefined rules [9]. However, these methods are limited in their ability to generalize to objects with complex geometries. To address these limitations, some researchers have leveraged CNNs and deep learning techniques, enabling the model to adapt to various shapes [10][11][12]. While deep learning has shown promising results, these methods typically lack context-awareness. Recent research such as Co-Grasp [13] incorporates human preferences in the grasping process, enhancing collaboration between robots and humans.

2.2 Robot-to-Human Handover

Robot-to-human handover refers to a collaborative interaction where one party, the giver, transfers an object to another, the receiver. This process requires precise coordination of both cognitive and physical actions to ensure a smooth exchange [5]. Therefore, researchers have developed a range of approaches to address the difficulties involved in robot-to-human handovers.

Some studies rely on heuristic rules, such as positioning the robotic arm's grasping point on the opposite side of the object from the human hand pose [14]. However, such approaches are not suitable for all types of objects. Other methods remove potential collision-prone grasp points from the human affordance map [15]. While these methods are effective for objects with regular shapes, the inherent risks associated with construction tools place a greater emphasis on safety considerations during the handover process. Additionally, some researchers have proposed using tactile and auditory cues to ensure safety during handovers [16]. However, their work does not address the prediction and selection of optimal grasping points.

In this paper, we propose a method for robotic arms to predict grasping positions specifically for construction tools. Our approach prioritizes safety and provides a foundation for secure and efficient handovers after the grasping phase.

3 Methodology

While traditional grasping methods often rely on simple geometric heuristics or rectangular grasp annotations, they tend to fall short in handling objects with irregular geometries and in providing the necessary context-awareness for dynamic environments. In contrast, GraspNet leverages 6-DOF grasp pose annotations, which allow for more precise and adaptable grasp predictions. This capability is particularly crucial for our application on construction tools, which frequently exhibit complex shapes and potential safety hazards. Therefore, despite the existence of simpler strategies, we chose GraspNet as our baseline due to its demonstrated robustness and superior performance in generating accurate grasp poses for irregular objects.

Based on the limitations in existing studies, we developed a comprehensive workflow to achieve safe construction tool handovers. The proposed methodology is composed of four main components: data annotation, assembly scene generation, grasp prediction, and robot motion planning. These components are designed to work in synergy, enabling safe robot-to-human tool handovers. The proposed framework is shown in Figure 1.

For tool selection, we chose tools commonly found on construction sites. Additionally, the shape of the tools was an important consideration; the selected tools exhibit significant diversity and irregularity in their shapes. In the annotation process, we differentiated between human and robotic grasping regions. This distinction ensures that the predicted grasp poses are focused on designated locations. After annotating the data, we generated the scenes using a Realsense D435 camera to capture images in a simulation environment. Each scene includes RGB-D images, object 6D poses, and masks. For grasp prediction, we employed the GraspNet-baseline [8] model as our prediction framework. We fine-tuned it using the generated dataset and conducted experiments in a simulation environment to verify the reliability of the predicted grasp poses. This approach ensures that the workflow effectively addresses safety in robot-to-human handover.

3.1 Data Annotation

We selected four tools from the Occupational Safety and Health Administration (OSHA) Hand and Power Tools Manual classified as hazardous during operation: adjustable wrench, screwdriver, pliers, and safety glasses. This highlights the importance of ensuring safety during the handover process. The selected tools are illustrated in Figure 2. The screwdriver and adjustable wrench are sourced from the YCB dataset [17], while the pliers and safety glasses are newly introduced objects in this dataset.

Among the selected tools, wrenches, pliers, and screwdrivers are widely used handheld tools on construction

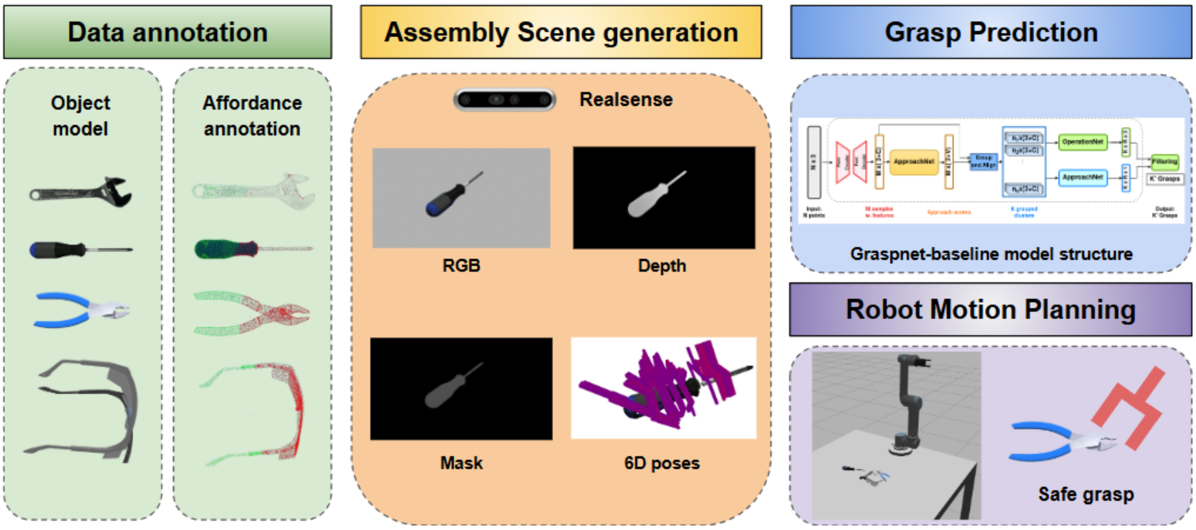

Figure 1. The overview architecture of the proposed handover strategy

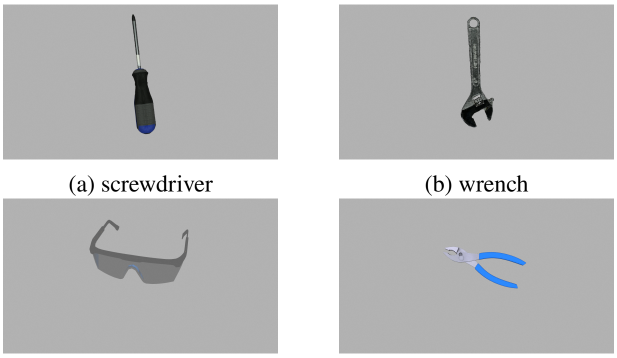

(a) screwdriver (b) wrench

(c) safety glasses (d) plier

Figure 2. Commonly used construction tools on construction sites

sites. These tools were specifically chosen for their diverse shapes and sizes, providing a comprehensive evaluation of the versatility of the proposed handover method. Additionally, safety glasses, as an essential protective tool on construction sites, are included to assess their grasping and handover performance in the construction site scenario.

The annotation process consists of two steps. First, based on the tool's usage scenario, we manually annotate the grasping areas for both human and robot interactions. The annotated data is then saved for grasp label generation. For example, in the case of a wrench, humans typically hold the handle for operation, while the other end may pose safety risks due to its sharpness. Therefore, the robot grasping area is annotated on the sharp end of the wrench, while the human grasping area is annotated on the handle.

The second step involves annotating the grasp labels. Various grasp datasets employ different annotation methods based on their intended usage. Examples of well-known datasets, such as Jacquard [18] and Cornell grasping dataset [19], are shown in Table 1. These datasets use rectangle annotations for grasp labels, which are less flexible compared to GraspNet. GraspNet employs annotations in the form of 6-DOF poses, which are particularly advantageous in complex construction environments. The additional precision provided by 6-DOF annotations enhances the robotic arm's ability to execute accurate and reliable movements in such challenging settings. As a result, we chose GraspNet as the annotation format for our grasp labels.

The annotated data are stored as point clouds, with green areas representing human grasping regions and red areas representing robot grasping regions. This separation not only ensures that the robot and human grasping positions are distinguished during the grasp prediction process but also facilitates safe handling by enabling humans, as receivers, to take the object securely from the handle. This dual focus enhances safety in robot-to-human handover scenarios. The annotation for each tool is shown in Figure 3.

3.2 Assembly Scene Generation

In order to construct a dataset for model training, we placed the tools flat on a table and selected the Realsense D435 camera for data collection. Camera calibration was performed prior to data collection to ensure accurate camera poses. The camera captured RGB-D images from 33 different poses, allowing us to gather more information about the tools from various angles. Corresponding masks were manually annotated, and both the camera and object

Table 1. Properties of public grasp datasets

Dataset	Grasp label	6D pose	Total images	Modality	Data source
Cornell [19]	Rect.	No	1035	RGB-D	1 Camera
Jacquard [18]	Rect.	No	54K	RGB-D	Sim
GraspNet [8]	6-DOF	Yes	97K	RGB-D	1 Camera
Ours	6-DOF	Yes	132	RGB-D	Sim

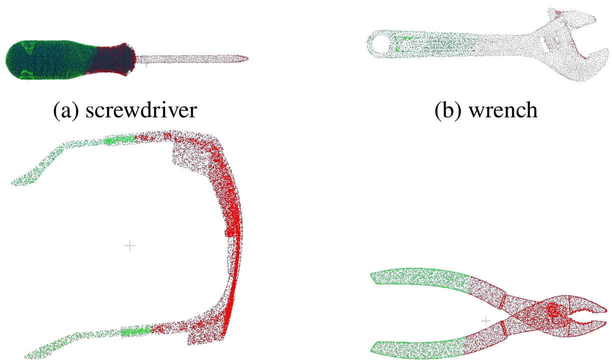

(a) screwdriver (b) wrench (c) safety glasses (d) plier

Figure 3. Visualization of annotated 3D models of every tool

6D poses were recorded. The dataset included a total of 132 images, along with their corresponding camera and object 6D poses.

3.3 Grasp Prediction

After preparing the training data, we used GraspNet-baseline as the model to predict grasp poses. The structure of the model is shown in Figure 4. During training, the initial learning rate was set to 0.01, and the batch size was 4. The learning rate was reduced to half its value every 5 epochs.

Once the training was completed, we first visualized the predicted grasp poses to verify their validity. Next, the RGB-D images from the simulation environment were used as input to the model, which generated the predicted grasp poses. In the following step, motion planning was employed to test the accuracy of these predicted grasp poses.

3.4 Robot Motion Planning

Once the dataset and scene setup were complete, we employed a fine-tuned model that takes RGB-D images as input and predicts multiple candidate grasp poses in 6-DOF, each accompanied by a confidence score. Among these predicted grasp poses, we select the one with the highest score as the optimal grasping position. The UR5 robotic arm then executes the grasp at this selected pose. This ensures stable and precise grasping by leveraging the model's scoring mechanism to identify the most reliable pose for each tool.

4 Experiment

To evaluate our approach, we compare its performance against the original GraspNet model, which serves as a benchmark. Specifically, the benchmark uses the official GraspNet implementation with the original pretrained weights, without being fine-tuned on our newly collected dataset. In contrast, our model starts from the same pretrained weights but is then fine-tuned on a custom construction-tool dataset containing domain-specific annotations. By juxtaposing the benchmark and our fine-tuned model, we can measure the impact on grasp prediction performance.

We first conduct robotic experiments to demonstrate that our predicted grasp poses can align well with simulation-world grasping. The success rate is calculated as the ratio of successful grasps to total attempts. A grasp is considered successful if the robot securely holds the tool without slippage, avoids collisions during both grasping and returning to the home position, and remains in stable contact with the tool throughout the entire motion. Failures are categorized into three primary types: (1) slippage during the grasping process, (2) collisions with the tool or environment, and (3) the tool dropping while the robot returns to the start pose.

4.1 Experimental Setup

The experiments were conducted in the Gazebo simulation environment using a UR5 robotic arm equipped with a parallel gripper. A RealSense D435 camera was employed for visual input, providing RGB-D data. The detailed experimental setup is depicted in Figure 5. All tools were placed flat on a table to simulate realistic conditions.

4.2 Evaluation Procedure

We tested both the benchmark and our fine-tuned model by predicting the top ten grasp poses (based on each model's grasp scores) for each tool, then executing each

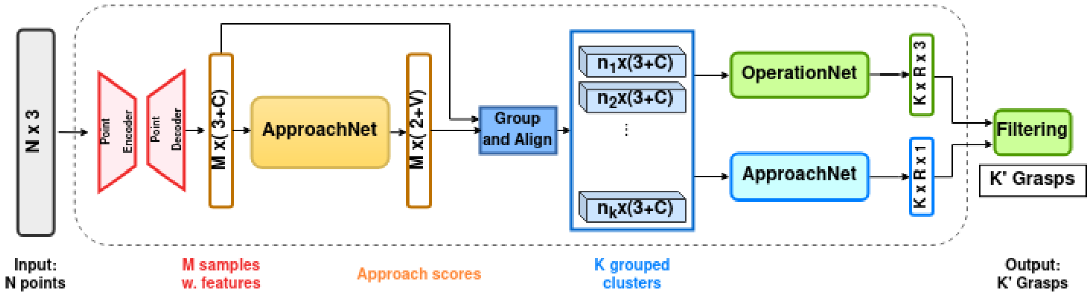

Figure 4. The overview of Graspnet-baseline [8] network

pose in simulation. To determine whether a grasp was successful, we checked:

1. Whether the predicted grasp pose lay within the designated "robot grasp region" of the annotated tool model.
2. Whether the grasp was maintained without slippage throughout the approach and retraction phases.
3. Whether no collision occurred between the robot and the tool (or the environment) at any point.

Any violation of these conditions such as grasping outside the annotated region, colliding with the tool, or dropping the tool during motion was recorded as a failure. In addition to evaluating the top 10 highest-scoring grasp poses, we also tested the highest-scoring pose specifically to check if the model's top-ranked grasp was valid. This procedure provides insight into each model's robustness and flexibility for a range of grasp candidates.

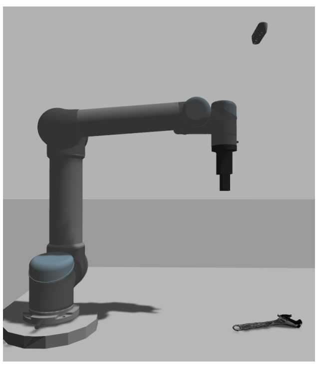

Figure 5. The experiment scene in Gazebo

5 Results

5.1 Benchmarking with GraspNet-baseline

Tables 2 summarizes the results of the benchmark model. The benchmark model achieve a success rate of 30% . For certain objects such as the screwdriver and plier, it exhibits a high collision failure rate. Specifically, with the screwdriver, half of the top 10 grasps failed due to collisions. In particular, when using the original GraspNet model, many predicted grasps lie outside the designated robot-grasp region and often result in collisions with the object, which reduces the overall success rate.

Table 2. Benchmark results: Success rates for the original GraspNet-baseline (top 10 grasp scores)

Object	Success rate
Screwdriver	20%
Adjustable wrench	40%
Plier	20%
Safety glasses	40%

5.2 Performance of Our Fine-Tuned Model

The predicted grasp poses were visualized and are shown in Figure 6, demonstrating the fine-tuned model's ability to generate feasible and collision-free grasp predictions. The model consistently produced grasp poses within the annotated regions for all tools, ensuring that the predictions were both accurate and aligned with safety requirements. The next section presents the results of the experiments, evaluating the predicted grasp poses in successfully grasping the objects.

The predicted grasp scores for each tool are summarized in Table 3. The success rates of the two experiments are detailed in Tables 4 and 5. For the highest-scoring grasp poses, all tools achieved a success rate of 100%, indicating that the proposed dataset and the trained model can generate highly reliable grasp poses. For the top 10 scoring grasp poses, the success rates were 100%, 80%, 90%, and 100% for the screwdriver, adjustable wrench, pliers, and safety glasses, respectively, with an overall average success rate of 92.5%. These results demonstrate the effectiveness of the proposed dataset in training the model to predict accurate grasp poses.

Despite the overall high success rates, some failure cases

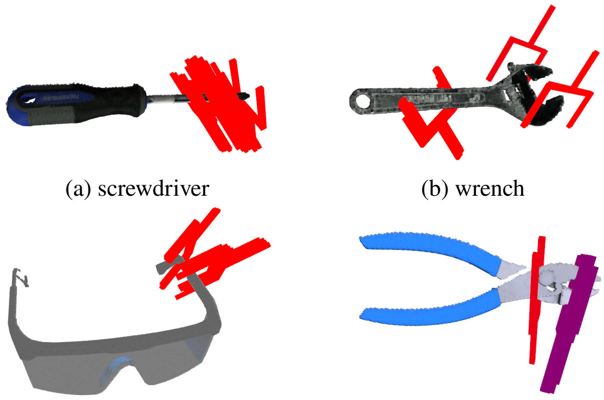

(a) screwdriver (b) wrench (c) safety glasses (d) plier

Figure 6. Qualitative results of the predicted grasp poses (top 10 highest-scoring grasps)

Table 3. Predicted grasp scores from fine-tuned model

Object	Highest	Top 10 average
Screwdriver	0.446	0.241
Adjustable wrench	0.680	0.467
Plier	0.385	0.376
Safety glasses	0.665	0.550

Table 4. Success rate of grasps with the highest grasp scores (fine-tuned model)

Object	Success rate
Screwdriver	100%
Adjustable wrench	100%
Plier	100%
Safety glasses	100%

Table 5. Success rate of grasps with top 10 grasp scores (fine-tuned model)

Object	Success rate
Screwdriver	100%
Adjustable wrench	80%
Plier	90%
Safety glasses	100%

were observed, particularly with the adjustable wrench, which achieved a success rate of 80% in the top 10 grasp experiments. Two failure cases occurred: one involved a collision between the robotic arm and the wrench during the grasping process, while the other involved slippage from the gripper. Figure 7 illustrates these two failure cases, with the left image showing a collision and the right image showing slippage.

The primary cause of these failures is the wrench's low height (1.4 cm) and flat geometry, which makes accurate grasp planning challenging. When the robotic arm executes the motion planning for grasping, the reduced thickness of the wrench can lead to misalignment, resulting in slippage or collisions. Similarly, the pliers slipped when lower-scoring grasp poses were tested. To further mitigate collisions and slippage, we plan to adopt a more rigorous collision detection strategy that discards grasp poses likely to cause interference. During the handover process, we will integrate reinforcement learning to optimize grasp stability and minimize the human operator's joint torque and displacement. By enhancing the robot's ability to adaptively maintain a secure grip while simultaneously reducing strain on the human arm, we aim to improve both safety and convenience in the tool handover process.

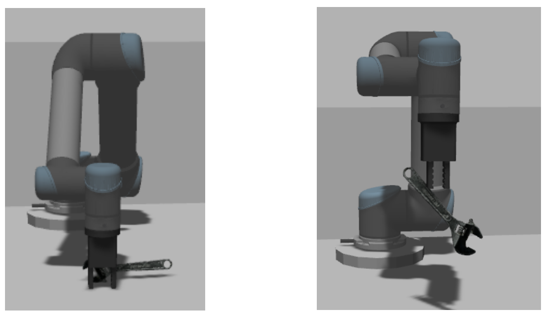

(a) collision (b) slippage

Figure 7. Failure cases of robotic arm grasping wrench

The computational efficiency of the approach is a critical factor for practical applications. The inference process requires an average of 1.6 seconds, while the motion planning process, including moving to the object, executing a stable grasp, and returning the arm to its home position takes 10.4 seconds. Although the current motion planning process takes longer to test the grasp stability. Integrating a reinforcement learning strategy during the handover phase will help reduce the overall planning time and enhance the system's for real-world deployments.

Overall, the experimental results confirm that the proposed dataset enables the GraspNet-Baseline model to predict reliable grasp poses. Despite occasional failures with specific tools like the adjustable wrench, the dataset demonstrated its ability to adapt to diverse tool shapes and sizes. These results validate the applicability of the dataset and model for safe and accurate robotic grasping in construction tool handover scenarios.

6 Conclusion and Future Work

In this paper, we proposed a method for robot-to-human hand tool handover grasp prediction and developed a dataset that includes four types of handheld tools commonly used in the construction industry. The dataset consists of multi-angle RGB-D images and detailed annotations. During the annotation process, we introduced a method to distinguish between human and robotic grasping regions, ensuring that the robotic arm predicts grasp poses concentrated on the potentially dangerous end of the tool. This approach minimizes safety risks during human handling by keeping the safer end of the tool accessible. Furthermore, we trained the dataset using the GraspNet-baseline model and validated the predicted grasp poses in a simulation environment. The experimental results demonstrated that our dataset aligns well with grasping requirements.

Despite the promising results, there are several limitations. First, the study is based on simulation experiments with a relatively limited dataset. The dataset doesn't capture the diversity of tools in real-world construction sites (e.g., variations in covers, handles, and materials), which may affect the model's generalizability. We plan to expand the dataset by incorporating a wider range of tool types and shapes, and to conduct real-world experiments to validate the transferability.

Finally, while the study focuses solely on grasp prediction, we acknowledge that human-robot interaction factors such as grip force control and handover timing are critical for ensuring safe tool handover in dynamic construction environments. In our future work, we plan to integrate sensor-based feedback (e.g., tactile sensors for grip force and vision-based human motion tracking) and adaptive control algorithms to dynamically adjust the handover process. This approach will not only enhance the grasp prediction module but will also optimize the overall handover process, providing a more robust and practical solution for real-world applications.

Overall, by addressing these limitations, our work aims to advance the state-of-the-art in robotic tool handover and bring the approach closer to practical deployment in real-world human-robot collaborative tasks.

References

[1] Mani Golparvar-Fard, Feniosky Peña-Mora, and Silvio Savarese. Automated progress monitoring using unordered daily construction photographs and ifc-based building information models. *Journal of Computing in Civil Engineering*, 29(1):04014025, 2015. doi:10.1061/(ASCE)CP.1943-5487.0000205.

[2] Karianne Berntsen. Did you know construction workers spend 38 hours annually looking for their tools? On-line: https://www.abax.com, June 2020. Accessed: 01/12/2025.

[3] Nathan Melenbrink, Justin Werfel, and Achim Menges. On-site autonomous construction robots: Towards unsupervised building. *Automation in Construction*, 119:103312, 2020. doi:https://doi.org/10.1016/j.autcon.2020.103312.

[4] Bo Xiao, Chen Chen, and Xianfei Yin. Recent advancements of robotics in construction. *Automation in Construction*, 144:104591, 2022. ISSN 0926-5805. doi:https://doi.org/10.1016/j.autcon.2022.104591.

[5] Valerio Ortenzi, Akansel Cosgun, Tommaso Pardi, Wesley P. Chan, Elizabeth Croft, and Dana Kulić. Object handovers: A review for robotics. *IEEE Transactions on Robotics*, 37(6):1855–1873, 2021. doi:10.1109/TRO.2021.3075365.

[6] Ming Zhang, Rui Xu, Haitao Wu, Jia Pan, and Xiaowei Luo. Human–robot collaboration for on-site construction. *Automation in Construction*, 150:104812, 2023. ISSN 0926-5805. doi:https://doi.org/10.1016/j.autcon.2023.104812.

[7] Ci-Jyun Liang, Xi Wang, Vineet R. Kamat, and Carol C. Menassa. Human–robot collaboration in construction: Classification and research trends. *Journal of Construction Engineering and Management*, 147(10):03121006, 2021. doi:10.1061/(ASCE)CO.1943-7862.0002154.

[8] Hao-Shu Fang, Chenxi Wang, Minghao Gou, and Cewu Lu. Graspnet-1billion: A large-scale benchmark for general object grasping. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11441–11450, 2020. doi:10.1109/CVPR42600.2020.01146.

[9] A.T. Miller, S. Knoop, H.I. Christensen, and P.K. Allen. Automatic grasp planning using shape primitives. In *2003 IEEE International Conference on Robotics and Automation (Cat. No.03CH37422)*, volume 2, pages 1824–1829 vol.2, 2003. doi:10.1109/ROBOT.2003.1241860.

[10] Joseph Redmon and Anelia Angelova. Real-time grasp detection using convolutional neural networks.

In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1316–1322, 2015. doi:10.1109/ICRA.2015.7139361.

[11] Sulabh Kumra and Christopher Kanan. Robotic grasp detection using deep convolutional neural networks. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 769–776, 2017. doi:10.1109/IROS.2017.8202237.

[12] Edward Johns, Stefan Leutenegger, and Andrew J. Davison. Deep learning a grasp function for grasping under gripper pose uncertainty. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4461–4468, 2016. doi:10.1109/IROS.2016.7759657.

[13] Abhinav K. Keshari, Hanwen Ren, and Ahmed H. Qureshi. Cograsp: 6-dof grasp generation for human-robot collaboration. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9829–9836, 2023. doi:10.1109/ICRA48891.2023.10160623.

[14] Chongxi Meng, Tianwei Zhang, and Tin lun Lam. Fast and comfortable interactive robot-to-human object handover. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3701–3706, 2022. doi:10.1109/IROS47612.2022.9981484.

[15] Zixi Wang, Zeyi Liu, Nicolas Ouporov, and Shuran Song. Contacthandover: Contact-guided robot-to-human object handover. On-line: `https://arxiv.org/abs/2404.01402`, Accessed: 04/2024.

[16] Hongrui Yu, Vineet R. Kamat, Carol C. Menassa, Wes McGee, Yijie Guo, and Honglak Lee. Mutual physical state-aware object handover in full-contact collaborative human-robot construction work. *Automation in Construction*, 150:104829, 2023. ISSN 0926-5805. doi:https://doi.org/10.1016/j.autcon.2023.104829.

[17] Berk Calli, Arjun Singh, James Bruce, Aaron Walsman, Kurt Konolige, Siddhartha Srinivasa, Pieter Abbeel, and Aaron M Dollar. Yale-cmu-berkeley dataset for robotic manipulation research. *The International Journal of Robotics Research*, 36(3):261–268, 2017. doi:10.1177/0278364917700714.

[18] Amaury Depierre, Emmanuel Dellandréa, and Liming Chen. Jacquard: A large scale dataset for robotic grasp detection. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3511–3516, 2018. doi:10.1109/IROS.2018.8593950.

[19] Yun Jiang, Stephen Moseson, and Ashutosh Saxena. Efficient grasping from rgbd images: Learning using a new rectangle representation. In *2011 IEEE International Conference on Robotics and Automation*, pages 3304–3311, 2011. doi:10.1109/ICRA.2011.5980145.