

Unknown Object Grasping for Assistive Robotics

Elle Miller[†], Maximilian Durner, Matthias Humt, Gabriel Quere, Wout Boerdijk, Ashok M. Sundaram, Freek Stulp, and Jörn Vogel

Abstract—We propose a novel pipeline for unknown object grasping in shared robotic autonomy scenarios. State-of-the-art methods for fully autonomous scenarios are typically learning-based approaches optimised for a specific end-effector, that generate grasp poses directly from sensor input. In the domain of assistive robotics, we seek instead to utilise the user’s cognitive abilities for enhanced satisfaction, grasping performance, and alignment with their high level task-specific goals. Given a pair of stereo images, we perform unknown object instance segmentation and generate a 3D reconstruction of the object of interest. In shared control, the user then guides the robot end-effector across a virtual hemisphere centered around the object to their desired approach direction. A physics-based grasp planner finds the most stable local grasp on the reconstruction, and finally the user is guided by shared control to this grasp. In experiments on the DLR EDAN platform, we report a grasp success rate of 87% for 10 unknown objects, and demonstrate the method’s capability to grasp objects in structured clutter and from shelves.

I. INTRODUCTION

The ability to grasp a wide variety of objects is crucial for independent living. Daily activities such as eating, drinking, dressing, grooming, and others all rely on the underlying capability to retrieve and interact with objects. This ability is often compromised in individuals with motor impairment, which can arise from spinal cord injury, traumatic brain injury, multiple sclerosis, muscular atrophy, and various other neuromuscular diseases. A review of surveys, reflecting the views of over 200 people with motor impairments, found that tasks relating to picking up, fetching, and carrying objects were of highest priority for an assistive technology [1]. There is a large body of research concerning the grasping of *known* objects, for which full object information and models are available to assist with grasp planning. For assistive technologies however, the first application environment is the home, which may have thousands of items. In a given day, an individual may need to interact with food, clothes, hygiene products, electronics, personal belongings, appliances, kitchenware, all of which are likely to vary and update through time. The ability to grasp only known objects greatly reduces the potential applications, and thus the focus of this work is the grasping of completely unknown rigid objects.

This problem has seen recent success with deep neural networks that can generate grasp poses directly from RGB or point cloud input, e.g. [2]. While fully autonomous

All authors were with the DLR Institute of Robotics and Mechatronics during the project. [†]Elle Miller is now affiliated with the University of Edinburgh. This work is partly supported by the Bavarian Ministry of Economic Affairs, Regional Development and Energy (StMWi) by means of the project SMILE2gether (LABAY102). Contact: elle.miller@ed.ac.uk



Fig. 1: Grasping an unknown object in structured clutter with the DLR EDAN assistive robotic system. After selecting an object, the user guides the end-effector across a virtual hemisphere to their desired approach direction. Depicted is an author, not a target user.

methods could be applied in this setting, we argue that user satisfaction would be lower. A study with 10 traumatic spinal cord injured subjects [3] found that autonomy did not provide more satisfaction for a robotic pick-and-place task. The authors conclude that there is a need to appropriately channel the autonomy provided, to enhance user satisfaction:

“While able-bodied users may prefer to cede autonomy to robots, we believe that disabled users tend to see the robot not merely as an agent to retrieve objects but also as a quintessential tool to reassert their domain of interaction with their environment as well as engage and exercise their cognitive faculties to the fullest.”

Moreover, the user already has intuition of where the best grasp might be, and they may wish to align the grasp with high level task-specific goals, e.g. transporting an empty mug compared to drinking from a full mug. In light of these findings, we aim to design a system that (i) requires the user to actively interact with the environment and (ii) utilises their cognitive abilities for decision-making. Yet, using manual control to perform goal-directed multi-fingered object manipulation is complex and difficult due to the high number of Degrees of Freedom (DoF) in robotic manipulators. Hence, our system uses shared control so the user can more intuitively control the end-effector [4][5].

Our pipeline (Fig. 2) performs Unknown Object Instance Segmentation (UOIS) on a pair of stereo images, and generates a 3D mesh reconstruction of the object using a deep

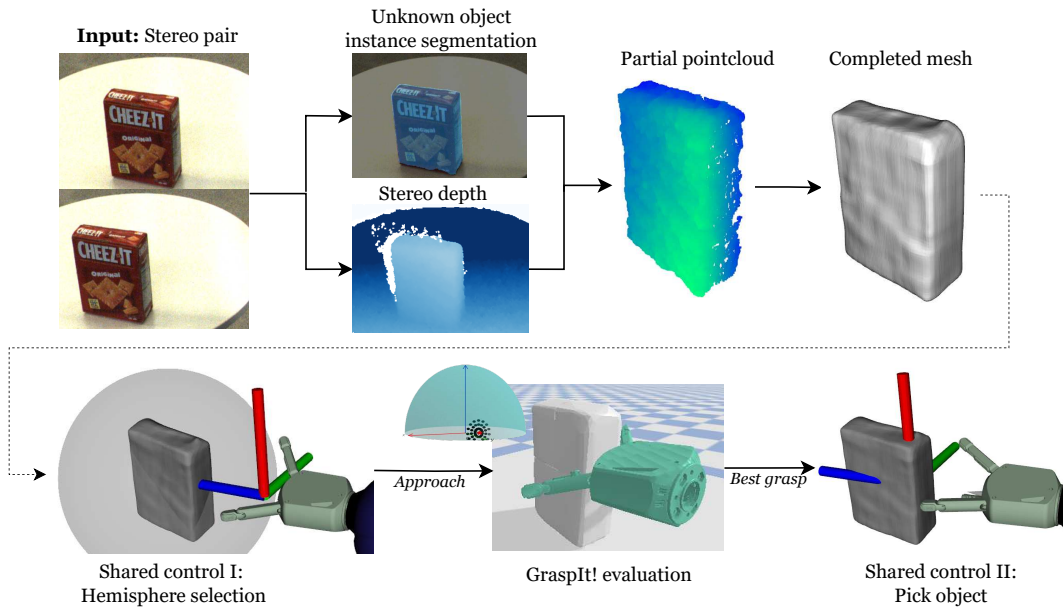


Fig. 2: Our unknown object grasping pipeline. The perception module transforms a pair of rectified stereo images to the completed mesh of an unknown object. Unknown object instance segmentation is performed using INSTR [6], and for shape completion we use the method presented in [7] which applies a Vector Quantized Deep Implicit Function (VQDIF) [8] trained on simulated Kinect data. In shared control [4], the user guides the end-effector across a virtual hemisphere centered around the unknown object, selects their desired approach position, and a physics-based grasp planning simulation finds the best local grasp at which the actual grasp is executed automatically.

implicit function network. The user then guides the robot end-effector across a virtual hemisphere centered around the unknown object, and selects their desired approach position. Next, a physics-based grasp planning simulation is run on the reconstructed mesh to find the most stable grasp near the user’s selected approach. Finally, the user is guided by shared control to the planner’s best grasp pose.

The main advantages of the pipeline are outlined below.

- i) *Human-centred*: The user is established as an active cognitive agent for enhanced satisfaction, grasping performance, and alignment with task-specific goals.
- ii) *End-effector agnostic*: There is no additional time or effort cost for changing or altering the end-effector.
- iii) *Customisable*: The number and region of local grasps sampled can be adjusted according to user ability.
- iv) *Scalable*: The system can be extended to more complex scenes, relying on human intuition to overcome difficulties (e.g. collision-avoidance, Fig. 1&5).

The novel contribution of this work is a pipeline for grasping unknown objects in shared autonomy or assistive settings, that utilises the valuable cognitive abilities of the user. We provide experimental verification of the method on the DLR EDAN assistive robotic system with the CLASH hand [9][10], showing high grasping success rates for singular rigid objects and extensions to more complex scenes.

II. RELATED WORK

A. Robotic grasping of unknown objects

Robotic grasping is the task of planning and executing grasp poses and associated finger configurations for a given object and end-effector. Analytic methods directly analyze

the object geometry, and data-driven methods typically employ machine learning. Many learning-based methods predict planar grasps (2D position, 1D rotation), suited for scenarios like bin picking [11] [12]. In complex environments however, such as the home, additional task and collision constraints make 6D grasp prediction crucial. [13] review the application of deep learning to 6D grasp pose synthesis, identifying the major methods to be sampling, regression, and reinforcement learning (RL). Sampling methods learn a function to evaluate the quality of a sampled grasp [14], regression (or end-to-end) methods learn a function to predict high-quality grasps from visual information [2][15], and RL aims to learn a policy that maximises total reward [16][17].

We did not explicitly use one of these approaches, in order to meet two criteria: the incorporation of *user intent* and scalability to *complex end-effectors*. As end-to-end methods simultaneously reason on perception, grasp planning, and control, it was unclear how to integrate a user’s input for further reasoning, which we deemed to be essential for this shared-autonomy system as discussed in Section I. Since these systems can not provide feedback on individual modules, this may also potentially result in lower user satisfaction and acceptance. We considered extending approaches such as Contact-GraspNet [2] to generate grasp candidates that the user could choose between, but found that many works were heavily optimised for simple rigid end-effectors. For example, [2] exploit the symmetry of a parallel-jaw gripper to formulate a simple loss function and reduce the problem dimensionality from 6D to 4D. These methods can not be trivially replicated for complex end-effectors such as our CLASH [10], an intrinsically compliant 7 DOF hand inspired by the tendon routing in humans.

Moreover, end-effectors may be updated through time for different requirements (e.g. EDAN previously used the DLR-HIT hand [18]). Thus, learning-based methods typically require problem reformulation, revised data collection, and retraining for different end-effectors. While it may appear that our proposed pipeline is complex with several failure points, we note that many other approaches also rely on RGB and pointcloud input, e.g. [2][15]. Finally, the recent realisation of semantic manipulation e.g. “*Grab the left ear of the elephant*” [19] provides interesting opportunities for human-centred grasping. However, we believe our method offers finer-grained grasp selection and increased satisfaction through active interaction for users in assistive settings.

B. Shape completion enabled grasping

Reconstructing a 3D model of an object using only a partial point cloud is a promising method for enabling robotic grasping. [14] use a 3D Convolutional Neural Network (CNN) for shape completion, trained on simulated depth images from 484 objects and achieve a 93% grasp success rate (GSR) on 15 YCB objects, most of which were already used during training. [20] propose a transformer-based encoder for shape completion, trained on real depth images from the YCB-Video dataset and achieve an 83% GSR across 6 YCB objects, all of which were seen during training. [21] propose a transformer-based encoder-decoder network for shape completion trained on the data from [14], and achieve a 76% GSR across 10 YCB objects, which also have been used at least partially during training. [7] recognise the challenges of reconstruction from partial and noisy depth data, and train an implicit function network [8] on simulated Kinect camera depth images to further bridge the simulation gap to real sensing conditions. We implement this method of shape completion in our work, and refer the reader to [7] for further information and related work. In order to recover an unknown object point cloud from a scene, we initially perform UOIS. In this context, [6] present a novel stereo-based transformer approach, addressing corrupted depth maps in real-world scenarios. To our knowledge, this is the first work combining RGB-based UOIS with shape completion for robotic grasping.

C. Assistive robotic grasping

Shared control grasping in assistive robotic systems has been studied previously. [22] focus on known objects with predefined grasp poses and approach vectors. [23] use a dense grasp database to provide the user with more freedom in selection, but also for known objects. Most similar to our work, [24] study grasping for unknown objects, but their approach is limited to objects that fall into shape primitive categories. Their system autonomously selects a single grasp for execution from semantic candidates (e.g. top, side, pinch) based on the user’s teleoperation of the end-effector. While this method also incorporates user intent, we believe our approach allows greater unknown object generalisation and user autonomy through finer-grained grasp direction selection beyond pre-defined semantic categories.

III. METHOD

Our proposed pipeline covers all relevant modules – from detecting object instances up to the prediction of robust grasp poses – for unknown object grasping. Given a pair of stereo images, the goal of perception is to create a complete mesh for each object in the scene (see upper part of Fig. 2). One should note that the use of mesh reconstruction limits this method to rigid objects. To this end, we perform UOIS and employ Instance Stereo Transformer (INSTR) [6] that creates pixel-wise instance masks for all arbitrary objects visible in the current view. Next, for higher robustness, each of the binary segmentation masks is filtered by contour area for noise. The user is presented with the filtered instances, and they select a target object by cycling through the options. We then construct an instance-specific point cloud by masking the Semi-Global Matching (SGM) [25] created depth map with the binary mask of the object of interest. After down sampling and outlier removal, the point cloud is reconstructed using the method presented in [7] which applies a VQDIF [8] trained on simulated Kinect data. Finally, we decimate the completed mesh to ~ 1000 triangular faces using the Quadric Error Metrics [26] for improved collision computation in the physics simulator.

The completed mesh is then forwarded to the grasping module. As shown in Fig. 2, the grasping module consists of three parts: hemisphere selection, grasp evaluation, and picking up the object. For both hemisphere selection and object picking, we employ Shared Control Templates (SCTs), previously developed for the EDAN system [4]. An SCT maps low-dimensional user input to end-effector motions and provides task space constraints, acting as as regional constraints as defined by [27], to guide the robots end-effector within the task. In this case a 3 DOF joystick is used to control the end-effector. The end-effector is constrained on a virtual hemisphere centered around the object of interest, as well as constrained to point towards the hemisphere origin. Those constraints allow the user to select their desired approach pose: they can toggle between moving the end-effector across the hemisphere surface (translational control), and rotating the end-effector around its axis pointing towards the hemisphere origin (rotational control). When satisfied with the approach pose, the user triggers a button.

Given this approach pose, we use the grasp planning simulation GraspIt! [28] on the reconstructed mesh to find the best local grasp. For a more realistic physics simulation, we used the PyBullet implementation [29] of GraspIt! to exploit the improved Bullet dynamics, kinematics, and collisions. We place the object on a surface and apply gravity instead of being fixed in space, to simulate object-surface dynamics. The unknown object is defined by a Unified Robot Description Format (URDF) file. By default, PyBullet uses a convex hull for mesh collision detection, which means that objects with concavities are (incorrectly) approximated as convex. To deal with this, we used the in-built Hierarchical Approximate Convex Decomposition (v-HACD) [30] to decompose the reconstructed mesh into multiple convex hull objects. The




Fig. 3: Grasp approach hemispheres, decided by object height. The power grasp allows for side-on grasps, and the precision grasp allows for greater sampling around the center point.

default mass is set to 0.3 kg. The lateral, rolling, and spinning friction coefficients are set to 0.3, 0.01, and 0.01 respectively. The inertia is recomputed by Bullet based on mass and volume of the collision shape. For grasp planning, we define a hemisphere from which approach positions are sampled from. We split the grasping strategy into *power* and *precision* grasps, depending on the object height (Fig. 3). The power grasp hemisphere is raised by 12 cm to allow for “side-on” grasps (without our end-effector colliding with the table). The original GraspIt! generates grasp approach angles evenly spaced across a sphere but we only sample local positions, which drastically reduces the amount of sampling positions and thus required sampling time. We generate a set of circumferences at different angular offsets from the user’s pose, and sample points from them. An example sampling is illustrated in Fig. 2 on the blue hemisphere in grasp evaluation, with 3 circumferences evenly spaced between a 0-10° angular offset from the user’s approach pose, which is shown in red. Note that the approach poses which exceed the hemisphere surface are filtered out. The final parameter we sample is how bent the fingers are on approach, which we call *finger flexion*: a value of 0 is where each finger is fully outstretched, and a value 1 corresponds to completely closed fingers. For power grasps, we sample finger flexions $\{0, 0.1\}$ per approach pose. For precision grasps, finger flexion is much more consequential to grasp success, and as such we sample $\{0.1, 0.2, 0.25, 0.3, 0.35\}$.

For each approach pose and finger flexion, the hand moves towards the center of the hemisphere until it has sufficient contact points (2 with the object for power grasps, 3 with object and/or world for precision grasps). Note that the object is fixed here. The position of the hand with a small “back-off” margin applied is then saved. The hand is then fixed, and the object is unfixed. Gravity is applied to the object, and the hand is closed with velocity control. Finally, the surface is moved vertically downwards. If there are hand-object contact points after 3 seconds, the grasp pose is deemed successful. Finally, the grasp with the highest ϵ value is selected (Fig. 4). If no grasp is found, the user is asked to select another approach position. The ϵ quality measure is the radius of the largest sphere that fits in the convex hull of the Grasp Wrench Space, and represents how robust the grasp is to external disturbance [28]. We calculate ϵ of the grasp pose at 0.5 seconds after the surface is moved rather than at 3 seconds, as a more accurate representation of grasp quality.

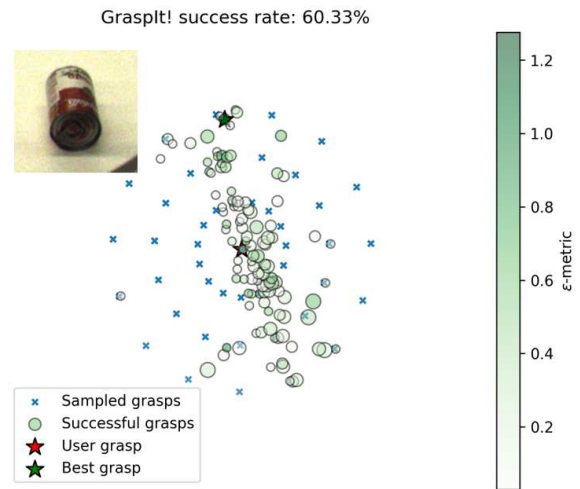


Fig. 4: Visualisation of GraspIt! simulation output for a soup can. From the user’s selected approach (top-down), local approach poses are sampled from 3 circumferences spaced between 0-10° from the input (blue crosses) projected onto the hemisphere. The hand approaches the soup from these poses but will keep moving until 3 contacts are reached, causing it to envelope around the object. Successful grasps are shown in circles, where the circle size is proportional to finger flexion (i.e. the smallest circles correspond to finger flexion 0.35, largest 0.1), and colour shows grasp quality.

IV. RESULTS

We implement our method on EDAN, the EMG-controlled Daily AssistaNt, a robotic wheelchair mounted with a torque-controlled, 8 DOF version of the DLR lightweight robot (LWR-III) [9]. The robot can be controlled by a joystick or via electromyographic (EMG) signals, and is designed to perform activities of daily living. We evaluate our method using 3 DOF joystick control, operated by one of the authors. Prior research shows that people with severe motor impairment can perform three-dimensional tasks from neural signals [31] or EMG [32]. Thus, we hypothesise that our method will also work with target users requiring alternative interfaces, but with longer control times. We believe that accuracy can be maintained if the sampling scheme is customised according to the precision abilities of the user, e.g. sample 0-30° from the selected approach for users with less control.

For stereo vision we use 2x Point Grey Chameleon CMLN-13S2C-CS cameras with 1/1.8” 4.5mm C-Mount wide angle lenses. For grasping we use the CLASH hand, a 7 DOF Compliant Low-cost Antagonistic Servo Hand also developed at DLR [10]. Our GraspIt! implementation is parallelised and run on a computer with 55 cores. For grasping, we select 10 objects from the YCB dataset [33]. The primary experiment is grasping singulated objects on a surface. Additionally, we compare autonomy modes and demonstrate the capability to grasp objects in structured clutter and from a shelf. We do not compare our results against baselines, e.g. Contact-GraspNet [2], because most related works have been trained on or optimised for parallel-jaw grippers, which would be non-trivial to re-implement for the CLASH hand (see Section II-A).

A. Singulated objects

For each of the 10 YCB objects we run 10 grasping trials, where the object’s position and orientation are randomised within an area the size of an A4 sheet of paper. A trial is successful if the object remains in the air 5 seconds after being grasped, lifted, and held stationary. The experimental results are presented in Table II. To evaluate INSTR, the mean Intersection over Union (IoU) and standard deviation for each object is provided. The shape completion metrics we use are the volumetric IoU and Chamfer-L1 distance (CD) mean and standard deviation. The mean GraspIt! success rate refers to the percentage of approaches that maintained object contact 3 seconds post gravity in simulation. Finally, the Grasp Success Rate (GSR) is the fraction of successful physical trials. To test the effectiveness of the complete pipeline, we only proceeded with attempts where a reconstructed mesh was recovered. As such, we discarded 5 attempts (2 *Banana*, 3 *Strawberry*) where INSTR failed to detect or severely under-segmented the object.

B. Autonomy modes

We show the benefits of our pipeline in an assistive setting and compare our approach (SC+GraspIt!) against varying degrees of autonomy (Table I). In manual mode, the user can toggle between controlling the end-effector translation, rotation, and closing the gripper. In Shared Control (SC) only mode, we remove GraspIt! from the pipeline and the user instead guides the end-effector along a line segment between their approach pose pose and the hemisphere origin, and manually triggers the gripper closing.

TABLE I: Comparison of results with different levels of autonomy.

| Autonomy | Cracker Box | | Strawberry | | Bowl | | Mean | |
|-------------|-------------|------|------------|-----|------|------|------|-----|
| | Time | GSR | Time | GSR | Time | GSR | Time | GSR |
| Manual | 62 | 90% | 50 | 90% | 57 | 100% | 56 | 93% |
| SC | 36 | 100% | 33 | 60% | 43 | 30% | 37 | 63% |
| SC+GraspIt! | 41 | 90% | 36 | 80% | 55 | 60% | 44 | 77% |

C. Complex scenes

We also demonstrate the scalability of our pipeline to more complex scenes. In structured clutter (Fig. 1), we grasped 5/5 objects (*Spam tin*, *Lemon*, *Pringles*, *Pear* and *Cracker box*) from the scene on the first trial. From the shelf (Fig. 5), we demonstrate grasping of the *Cracker box* despite degraded visual conditions. Both these scenes relied on the user’s intuition to select collision-free grasp directions for success.

V. DISCUSSION

A. Singulated objects

From Table II, we see that the UOIS method performed well with all objects having a mean IoU above 90. The applied shape completion approach is able to faithfully reconstruct most objects, but struggles with small ones (*Strawberry*, *Banana*), as they are relatively more affected by imprecise segmentation and noise. The volumetric IoU metric is deceptively small for objects with small volume




Fig. 5: Shelf grasping experiment. 1. INSTR prediction 2. Completion prediction 3. User’s selected hemisphere approach 4. GraspIt! predicted grasp 5. Successful grasp. More complex grasping scenarios utilise the user’s intuition for collision avoidance.

(*Mug*, *Bowl*), as even small pose discrepancies can lead to negligible overlap. The CD metric is not affected by this phenomenon and provides a more faithful representation of shape completion performance in these cases.

Comparing *Spam tin* with an IoU of 77.2 and CD of 0.325 to *Bowl* with an IoU of 2.8 and CD of 0.379, we observe a significant difference in IoU but only a modest increase in CD. To see how segmentation performance impacts the shape completion, we look to Fig. 6. There appears to be no clear correlation between segmentation IoU and shape completion IoU, except for a dense cluster of high IoUs in the upper right. This suggests that the shape completion model is robust to minor under- or over-segmentation. The GraspIt! success rate is moderately stable across objects, except the *Strawberry*, *Banana*, and *Bowl*. This indicates that for these objects the finger flexion parameter is very influential to success, where only small finger spreads during approach led to successful grasps. Considering run time, we can see that the major time is taken by the user to select a satisfactory approach pose (24s). Note that the timings do not include perception and object selection, which took <10s.

In the physical trials, we see reliable grasping across many objects. The worst performing object was the *Bowl* with a 6/10 GSR, followed by the *Banana*, *Strawberry*, and *Master Chef can* with 8/10 GSR. Therefore, does poor shape completion lead to grasp failure? Not necessarily, as in Fig. 7 we see that grasp failures and successes occurred throughout entire completion IoU range. Most successful grasps had a high shape completion IoU (>80), but there were also many successful grasps at low completion IoUs. For example, the *Mug* had a low mean completion IoU of 23.4 (likely due to its shiny surface and concavity), but since the overall object perimeter was preserved had 10/10 GSR for top-down grasping. For the bowl however, an accurate completion was likely very important for grasp planning as a specific

TABLE II: Experimental results on robot for 100 grasping trials of singulated objects. For each object, the mean and variance across 10 trials are presented for the segmentation IoU, completion volumetric IoU and Chamfer-L1 distance. The mean GraspIt! simulation success rate, and grasp pipeline times are also reported. The GSR is the physical trial result.

| YCB Object | Segmentation IoU | | Completion IoU | | | | GraspIt! Success rate (%) | Timings | | | GSR (10 trials) |
|-----------------|------------------|----------|----------------|----------|-----------|----------|---------------------------|------------|-----------|------------|-----------------|
| | \bar{x} | σ | \bar{x} | σ | \bar{x} | σ | | Select | GraspIt! | Pick | |
| Pringles | 96.4 | 0.6 | 87.0 | 2.3 | 0.091 | 0.012 | 59 | 15s | 4s | 10s | 100% |
| Master Chef can | 93.5 | 6.2 | 83.9 | 3.0 | 0.232 | 0.045 | 72 | 28s | 5s | 12s | 80% |
| Cracker box | 93.9 | 6.0 | 80.1 | 10.6 | 0.195 | 0.076 | 68 | 29s | 4s | 8s | 90% |
| Tomato soup can | 94.4 | 1.2 | 82.3 | 2.6 | 0.226 | 0.036 | 70 | 17s | 9s | 11s | 90% |
| Spam tin | 93.8 | 1.3 | 77.2 | 5.3 | 0.325 | 0.072 | 69 | 23s | 9s | 12s | 100% |
| Strawberry | 92.1 | 2.2 | 56.5 | 15.7 | 0.672 | 0.188 | 10 | 29s | 6s | 11s | 80% |
| Banana | 90.3 | 3.5 | 48.2 | 12.1 | 0.300 | 0.103 | 10 | 27s | 6s | 13s | 80% |
| Mug | 94.6 | 1.8 | 23.4 | 7.3 | 0.527 | 0.219 | 57 | 22s | 9s | 10s | 100% |
| Bowl | 92.0 | 3.7 | 2.8 | 1.7 | 0.379 | 0.144 | 33 | 28s | 7s | 20s | 60% |
| Softball | 94.5 | 1.5 | 82.8 | 3.9 | 0.314 | 0.067 | 57 | 22s | 7s | 10s | 90% |
| Averages | | | | | | | | 24s | 7s | 12s | 87% |




Fig. 6: Segmentation and completion IoU for the 100 trials of singulated objects.




Fig. 7: Histogram of shape completion IoU showing number of successful and failed trials for each region.

approach position and contacts are required for success. Thus, object complexity is likely an important confounding factor to both completion IoU and grasping success.

B. Autonomy modes

The autonomy experiments in Table I highlight interesting comparisons between time taken and reliability. Manual mode was the most reliable grasping method (28/30 GSR), but also took the most time (56s). The SC only mode was the fastest (37s), but had the lowest reliability (19/30 GSR). This is because when approaching the object with outstretched fingers along the hemisphere radius, it was difficult to visualise how the fingers would establish contact with the object at the resulting grasp pose. For example, there was a small window in which the strawberry could be grasped: too low and the outstretched fingers would collide with the surface, too high the strawberry would slip. Finally, SC+GraspIt! balanced the time and reliability of the other methods. Note that we do not include the timing of the perception pipeline (<10s) to contrast timings where the user is under assumed cognitive load. Considering the full pipeline, SC+GraspIt! becomes 44+10=54s, which is very close to the manual average of 56s. This result, in combination with the higher manual grasping reliability, may appear to undermine our pipeline. However, these measurements do not consider the duration spent under cognitive load. The user is assumed to be under

high cognitive load throughout the full length of manual mode (56s), but only during the hemisphere select phase in our method (24s), $\sim 44\%$ of the full pipeline time. Future work should explore methods of improving the grasping performance in order to reach the high reliability of manual mode, and quantify the differences in cognitive load between the approaches, e.g. using the NASA Task Load Index.

VI. CONCLUSION

In this work we present a novel pipeline for grasping unknown objects in an assistive setting, combining UOIS with shape completion. Importantly, our grasping methodology utilises the cognitive abilities of the user for enhanced satisfaction, grasping performance, and alignment with high level task-specific goals. We report a grasping success rate of 87% across 10 singulated objects, and demonstrate capability to grasp objects in structured clutter and from shelves. We find that the shape completion model is robust to minor segmentation failures, but can struggle with poor sensor data. In experiments comparing autonomy modes, we find that our pipeline reduces the time spent in assumed cognitive load by more than half, compared to manual mode. Future work could explore entire scene reconstruction and integrating collision-avoidance for dense clutter scenarios, and quantifying cognitive load across autonomy modes.

REFERENCES

- [1] C.A. Stanger et al. “Devices for assisting manipulation: a summary of user task priorities”. In: *IEEE Transactions on Rehabilitation Engineering* 2.4 (1994), pp. 256–265.
- [2] Martin Sundermeyer et al. “Contact-GraspNet: Efficient 6-DoF Grasp Generation in Cluttered Scenes”. en. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. Xi’an, China: IEEE, May 2021, pp. 13438–13444.
- [3] Dae-Jin Kim et al. “How Autonomy Impacts Performance and Satisfaction: Results From a Study With Spinal Cord Injured Subjects Using an Assistive Robot”. In: *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans* 42.1 (2012), pp. 2–14.
- [4] Gabriel Quere et al. “Shared Control Templates for Assistive Robotics”. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. 2020, pp. 1956–1962.
- [5] Bruno Siciliano, Oussama Khatib, and Torsten Kröger. Vol. 200. Springer, 2008. Chap. 43.3.2, p. 1092.
- [6] Maximilian Durner et al. “Unknown Object Segmentation from Stereo Images”. In: *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2021*. 2021.
- [7] Matthias Humt, Dominik Winkelbauer, and Ulrich Hillenbrand. “Shape Completion with Prediction of Uncertain Regions”. In: *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2023*. 2023.
- [8] Xingguang Yan et al. “ShapeFormer: Transformer-based Shape Completion via Sparse Representation”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2022.
- [9] Jörn Vogel et al. “EDAN: An EMG-controlled Daily Assistant to Help People With Physical Disabilities”. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2020, pp. 4183–4190.
- [10] Werner Friedl and Máximo A. Roa. “CLASH—A Compliant Sensorized Hand for Handling Delicate Objects”. In: *Frontiers in Robotics and AI* 6 (2020).
- [11] Lerrel Pinto and Abhinav Gupta. “Supersizing self-supervision: Learning to grasp from 50K tries and 700 robot hours”. In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. Stockholm, Sweden: IEEE Press, 2016, pp. 3406–3413.
- [12] Jeffrey Mahler and Ken Goldberg. “Learning Deep Policies for Robot Bin Picking by Simulating Robust Grasping Sequences”. In: *Conference on Robot Learning*. 2017.
- [13] Rhys Newbury et al. “Deep Learning Approaches to Grasp Synthesis: A Review”. en. In: *IEEE Transactions on Robotics* 39.5 (Oct. 2023), pp. 3994–4015.
- [14] Jacob Varley et al. “Shape completion enabled robotic grasping”. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2017, pp. 2442–2447.
- [15] Hao-Shu Fang et al. *AnyGrasp: Robust and Efficient Grasp Perception in Spatial and Temporal Domains*. arXiv:2212.08333 [cs]. June 2023.
- [16] Sergey Levine et al. “Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection”. In: *The International Journal of Robotics Research* 37.4-5 (2018), pp. 421–436.
- [17] Dmitry Kalashnikov et al. “Scalable Deep Reinforcement Learning for Vision-Based Robotic Manipulation”. In: *Proceedings of The 2nd Conference on Robot Learning*. Ed. by Aude Billard et al. Vol. 87. Proceedings of Machine Learning Research. PMLR, Oct. 2018, pp. 651–673.
- [18] H. Liu et al. “Multisensory five-finger dexterous hand: The DLR/HIT Hand II”. In: *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*. ISSN: 2153-0866. Sept. 2008, pp. 3692–3697.
- [19] Priya Sundareshan et al. *KITE: Keypoint-Conditioned Policies for Semantic Manipulation*. arXiv:2306.16605 [cs]. Oct. 2023.
- [20] Wenkai Chen et al. “TransSC: Transformer-based Shape Completion for Grasp Evaluation”. In: *CoRR* abs/2107.00511 (2021).
- [21] Seyed S. Mohammadi et al. “3DSGrasp: 3D Shape-Completion for Robotic Grasp”. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. 2023, pp. 3815–3822.
- [22] Katharina Muelling et al. “Autonomy infused teleoperation with application to brain computer interface controlled manipulation”. en. In: *Autonomous Robots* 41.6 (Aug. 2017), pp. 1401–1422.
- [23] Jörn Vogel et al. “Flexible, semi-autonomous grasping for assistive robotics”. In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2016, pp. 4872–4879.
- [24] Siddharth Jain and Brenna Argall. “Grasp Detection for Assistive Robotic Manipulation”. In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2016, pp. 2015–2021.
- [25] Heiko Hirschmuller. “Accurate and efficient stereo processing by semi-global matching and mutual information”. In: *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*. Vol. 2. IEEE. 2005, pp. 807–814.
- [26] Michael Garland and Paul S. Heckbert. “Surface Simplification Using Quadric Error Metrics”. In: *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH ’97*. USA: ACM Press/Addison-Wesley Publishing Co., 1997, pp. 209–216.
- [27] Stuart A Bowyer, Brian L Davies, and Ferdinando Rodriguez y Baena. “Active constraints/virtual fixtures: A survey”. In: *IEEE Transactions on Robotics (T-RO)* 30.1 (2013), pp. 138–157.
- [28] Andrew T Miller and Peter K Allen. “Grasplit! a versatile simulator for robotic grasping”. In: *IEEE Robotics & Automation Magazine* 11.4 (2004), pp. 110–122.
- [29] Carlyn Dougherty. *PyBullet Grasplit!*
- [30] Khaled Mamou and Faouzi Ghorbel. “A simple and efficient approach for 3D mesh approximate convex decomposition”. In: *Proceedings of the International Conference on Image Processing, ICIP 2009, 7-10 November 2009, Cairo, Egypt*. IEEE, 2009, pp. 3501–3504.
- [31] Leigh R. Hochberg et al. “Reach and grasp by people with tetraplegia using a neurally controlled robotic arm”. en. In: *Nature* 485.7398 (May 2012). Number: 7398 Publisher: Nature Publishing Group, pp. 372–375.
- [32] Annette Hagengruber and Jörn Vogel. “Functional Tasks Performed by People with Severe Muscular Atrophy Using an sEMG Controlled Robotic Manipulator”. In: *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. ISSN: 1558-4615. July 2018, pp. 1713–1718.
- [33] Berk Calli et al. “Yale-CMU-Berkeley dataset for robotic manipulation research”. In: *The International Journal of Robotics Research* 36.3 (2017), pp. 261–268.