Proceedings of the 2018 IEEE
International Conference on Robotics and Biomimetics
December 12-15, 2018, Kuala Lumpur, Malaysia

Visionless Tele-Exploration of 3D Moving Objects

Kevin Huang*, Patrick Lancaster*, Joshua R. Smith and Howard Jay Chizeck

Abstract— **This paper presents methods for improved tele-operation in dynamic environments in which the objects to be manipulated are moving, but vision may not meet size, biocompatibility, or maneuverability requirements. In such situations, the object could be tracked through non-geometric means, such as heat, radioactivity, or other markers. In order to safely explore a region, we use an optical time-of-flight pretouch sensor to detect (and range) target objects prior to contact. Information from these sensors is presented to the user via haptic virtual fixtures. This combination of techniques allows the teleoperator to "feel" the object without an actual contact event between the robot and the target object. Thus it provides the perceptual benefits of touch interaction to the operator, without incurring the negative consequences of the robot contacting unknown geometrical structures; premature contact can lead to damage or unwanted displacement of the target. The authors propose that as the geometry of the scene transitions from completely unknown to partially explored, haptic virtual fixtures can both prevent collisions and guide the user towards areas of interest, thus improving exploration speed. Experimental results show that for situations that are not amenable to vision, haptically-presented pretouch sensor information allows operators to more effectively explore moving objects.**

I. INTRODUCTION

This paper presents methods for improved tele-operation in dynamic environments in which the objects to be manipulated are moving, and vision may be degraded. Example motivating application scenarios include (1) robotic surgery, in which vision may be limited inside the body, and manipulation targets may move due to breathing, heart beats, and other biological motion, and (2) underwater robotic tele-manipulation, in which visual sensing is degraded due to murky water, and objects may move due to currents. In these applications, target/environment geometry must be well understood prior to interaction, but contact should not be used as an information gathering mechanism (because of the potential for damage or unwanted target displacement), and long range visual sensors cannot be relied upon.

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1256082, as well as IIS-1427419. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

*These two authors contributed equally to the study.

Kevin Huang is with Trinity College, Dept. of Engineering, 300 Summit St, Hartford, CT 06106, USA kevin.huang@trincoll.edu

Patrick Lancaster and Joshua R. Smith are with the University of Washington Departments of Computer Science and of Electrical Engineering, Seattle, WA. planc509, jrs@cs.washington.edu

Howard Jay Chizeck is with the University of Washington Department of Electrical Engineering, 185 Stevens Way, Paul Allen Center - Room AE100R, Campus Box 352500, Seattle, WA 98195, USA. chizeck @uw.edu

This project relaxes the assumption of static interaction targets, a common simplification in robotics that is not valid in less structured settings. Various forms of pretouch sensing are highly maneuverable and are robust to poor visibility, and haptic virtual fixtures can provide intelligently informed, non-visual cues in exploration. The use of multiple pretouch sensors and sensor modalities can reduce uncertainty and increase understanding of the environment and thus the ability to safely interact with it. We utilize a pretouch sensor based on optical time of flight measurement [1], see Figure 1. The sensor detects and ranges target objects prior to contact; information from these sensors is presented to the user via haptic virtual fixtures. This combination of techniques allows the tele-operator to "feel" the object without actually creating contact between the robot and the object/environment. Thus it provides the perceptual benefits of touch interaction to the operator, without relying on the negative consequences of the robot actually contacting unknown geometrical structures.

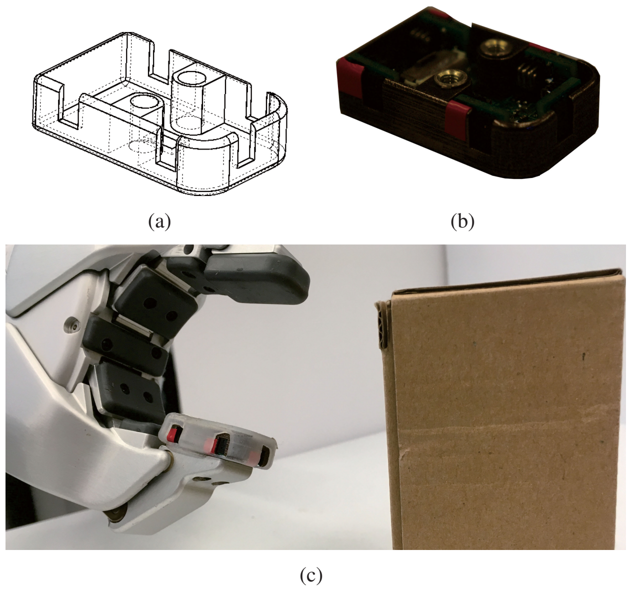

Fig. 1. The utilized optical pretouch sensor. (a) Computer generated prototype of housing. (b) Sensor casing and circuitry. It is packaged into a small form factor. (c) The robot uses the sensor mounted on its lower fingertip to measure an object.

In addition to the use of pretouch sensing, sensor-driven feedback pathways are used to intelligently inform the human operator. The haptic pathway is one of particular interest due to associated low reaction times [2] as well as intuitive nature when coupled with user motion input.

978-1-7281-0377-8/18/$31.00

A. Problem Description

The goal of this work is to integrate compact pretouch sensing and haptic feedback into teleoperative refinement of a moving object to simulate real-world scenarios. In some telerobotic applications, the robotic task can involve first identifying and recognizing the precise geometries of an object with which to interact. Sometimes it is sufficient to use modern range sensors, which can build a baseline geometry map of the remote scene in real-time. There are drawbacks, however. For example, the objects of interest may be poorly observed. This can be due to several factors, including occlusions, optical properties, glancing incidence angles and other factors [3]. In other scenarios, such sensors are not suitable for the task — the remote environment may completely interfere with the sensor's mode of operation (e.g. IR sensors outdoors or in highly radioactive settings), or the environment may have physical or biological constraints. One cost of high-bandwidth geometry information is form factor, making such sensors difficult to utilize in confined spaces. Biological systems are also oftentimes spatially compact and furthermore are sensitive to the material and chemical properties of foreign bodies. A more compact sensor solution is needed.

B. Contributions

To the best of the authors' knowledge, this work is the first to:

1) Use pretouch sensor arrays to explore moving targets whose geometry information is not initially known
2) Demonstrate the efficacy of haptic virtual fixtures for sequential point-cloud data of a moving object
3) Provide a method for real-time fusion of haptic guidance and forbidden region virtual fixtures

II. Background

A. Pretouch Sensing

Fingertip sensors are typically highly maneuverable, allowing them to explore areas that would be occluded in the relatively static view frames of long range depth sensors or cameras. Furthermore, by collecting data in closer proximity to the area or object of interest, fingertip sensors can often yield data that is more precise than sensors that measure from afar.

Tactile sensors can provide precise information about object pose. In [4], tactile measurements were used to estimate the 6 degree of freedom (DOF) pose of an object. Both works described and utilized the Scaling Series algorithm as a means to integrate tactile measurements with an object model for pose estimation. One problem with this approach is that it is unlikely for a robot to have a model of the object *a priori*. In general, tactile sensing necessarily requires contact, which could displace the object.

The "pretouch" modality of sensing occurs at a distance scale intermediate to that of long-range depth sensors and tactile sensors. Many types of pretouch sensing have been previously explored, such as electric field, optical, and acoustic. Electric field sensing only works well for objects that are highly capacitive or possess a dielectric constant significantly different from that of air. In [5], Mayton et.al demonstrated that servoing an arm and fingers equipped with electric field sensors can successfully gather geometric information about an object of interest. A number of works have explored the use of optical pretouch sensors for robot grasping [6] [7]. An optical sensor for the purpose of grasping, surface classification and slip detection was developed in [8]. In [9], optical proximity measurements were used to estimate object pose from a probabilistic model. Optical methods are affected by surface properties of the object and ambient lighting, and can completely fail when trying to sense highly specular or transparent objects. Jiang and Smith presented the first fingertip-mounted acoustic sensor in [10] and detailed a framework for using it and other pretouch sensors to perform exploration in [11] and a telerobotic implementation in [12]. Although acoustic sensing is well-suited for many of the failure cases of electric field and optical sensing, it does not detect open foams, rough fabrics, and fur well.

B. Telerobotics and Haptic Virtual Fixtures

Teleoperation has applications in science, medicine, education, and the military Several intriguing applications are described by [13] [14] [15]. Teleoperation is beneficial in situations that are too dangerous or otherwise inaccessible for human beings, but too unpredictable for autonomy. Using depth sensors, a real-time model of a real world scene can be constructed and interacted with via a haptic simulation [16] [17]. When the haptic interaction point (HIP) in the virtual representation is properly registered to a telerobot, this same methodology allows the user to transparently feel what the robot 'touches', or even what the robot would 'touch' in a potential pose, as described by Leeper et al. [18].

In addition to transparent reflection of forces through haptic display, force feedback can be extended to intelligently modify and regulate the motion commands of the human user. More precisely, haptic virtual fixtures can be used to limit tool motion to aid in task execution, and are enforced in software. During a task, the system monitors the user's motion commands, analyzes against software bounds and constraints, and finally modifies the user's motion based on violation of virtual fixtures. Much work in virtual fixtures deal with constraints and bounds defined *a priori*, which can be done with or without a human supervisor, as in [19] [20] [21].

C. Sensor Based Exploration

Environment perception is crucial for most dexterous tasks — a model of the real-world must first be constructed, and oftentimes this model is constructed from information gathered from sensors. Saxena et al. used a model-based approach, where fixed objects and obstacles were partially identified and then fitted to a known model [22]. This group, along with [23] and [24], use a library of 3D models for actual interaction.

Oftentimes, *a priori* models are not available, especially in the case of new environments and objects. As stated by

Aleotti et al. without actively exploring the environment, a robot can fail to build a complete 3D model because of occlusion [25]. An active exploration scheme, which utilizes maneuverable sensors, can be used to populate and refine previously unseen surfaces.

In [26], an RGB-D sensor was mounted to the gripper of the robot. This information was used to maintain and update a voxel-representation of the environment which was actively explored in occluded regions. Using this active exploration approach, grasping task speed was increased and success rates were more than tripled. Similarly, in [25], a laser-range finder was mounted to the robot end-effector. Multiple viewpoints were stitched together in order to build a 3D model of an object for grasp planning, and different methods of 3D registration were compared. In [12], a small form factor pretouch sensor robust to optical transparency was used to explore unobserved surfaces for the purposes of grasping transparent objects.

III. Haptic Virtual Fixtures

The haptic feedback for this task needs to help accomplish two primary goals:

1) Guide the operator to the unobserved target
2) Prevent collisions with obstacles and the target

In order to achieve the former, a guidance virtual fixture in the form of a distance proportional attractive radius can be implemented. This is in contrast to several path following approaches which restrict user motion along a predefined path or provide virtual force feedback in directions orthogonal to the desired path [20]. To achieve the latter, forbidden region virtual fixtures can be constructed with 3DOF proxy-method haptic rendering, as demonstrated by [27]. These forbidden regions will prevent the robot from coming within a given radius of any sensed geometry point.

While effective real-time implementations of these algorithms exist, the challenge here is twofold: firstly determining the proper parameters for the separate virtual fixtures, and secondly to resolve the two haptic cues. First consider setting up the virtual fixtures separately. To that end, first define the following variables

r_{fr}	forbidden region radius
r_g	no-guidance radius
r_o	anticipated target boundary radius
$\vec{f}_G$	guidance force vector
$\vec{f}_{fr}$	forbidden region force vector
$\vec{f}_{\text{net}}$	calculated force feedback

The determination of valid ranges for r_{fr} can be found in [12]. From this, $\vec{f}_{fr}$ can be calculated and proxy motion determined, as per [16].

Examine first the potential forces generated via the target geometries and location. This aids in determining the attractive guidance radius. On one hand, the operator should be encouraged to move toward the object if too far away. On the other, when the operator is within sensing range of the moving target, the attractive force should not produce a collision. The goal then is to determine a radius from the object at which guidance force starts and stops. With that said, bounds for both safe and efficient target refinement are sought. The following analyses are performed with a two-dimensional representation, but conclusions and implications can easily be extended to three dimensions.

First, consider the goal of maintaining the robot end effector within effective sensing range, and begin with the setup shown in Figure 2.

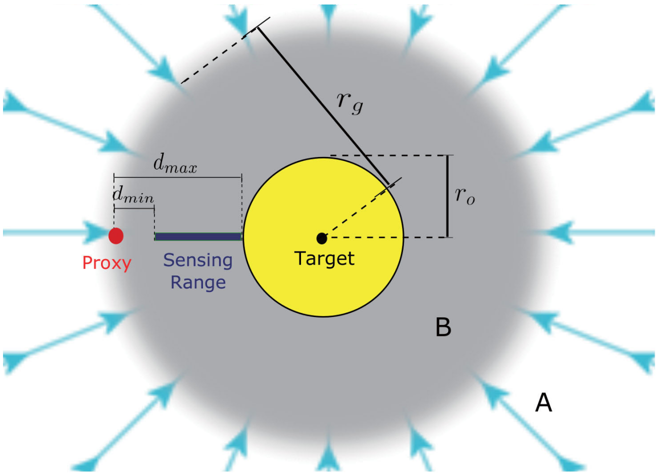

Fig. 2. Basic end effector sensor exploration as the robot end effector approaches a surface. The robot end effector follows the location of the proxy which is attracted to within r_g of the target center, at which point the target is within sensing range. A proxy in region A would experience an attractive force towards the target, while one in B would receive no guidance feedback since it would be deemed to be within sensing range of the target.

From this figure, it is immediately clear that the attractive force must bring the proxy/end effector to within d_{max} of r_o of the target point. This can be expressed simply as

$$r_g \leq r_o + d_{max} \quad (1)$$

This will ensure that the operator and robot will be guided to within sensing distance of object surface. Note that the end effector may be configured in a way such that none of the pretouch sensors within the sensor array are facing exactly normal to the anticipated target surface. This upper bound on r_g encompasses all other cases.

Similarly, a maximum lower bound on r_g to enable the end effector to detect the surface at r_o is easily found as

$$r_g \geq r_o + d_{min} \quad (2)$$

This bound does not ensure safe guidance force radius; rather it is the smallest radius at which the robot may be able to observe the surface at r_o. In practice, the geometry and arrangement of the sensor array with respect to the proxy and r_o are needed to determine a lower bound on r_g to ensure safe guidance.

Now suppose that r_{fr} is determined. Let it be depicted as a positive scalar $r_{fr} \in \mathbb{R}^+$ such that

$$d_{max} \geq r_{fr} \geq d_{min} \quad (3)$$

Observe that this radius then helps define an effective value of r_g. This can be shown by assuming that a point has been sensed on the object surface, and consider a radius of

$$r_i = r_o + r_{fr}$$

as is shown in Figure 3

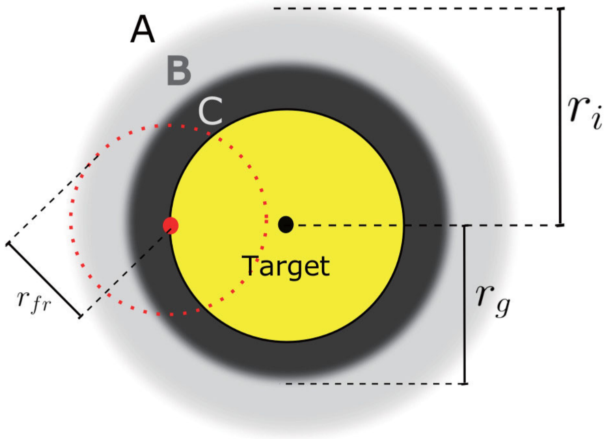

Fig. 3. Guidance and forbidden region virtual fixture overlap. This scenario shows when $r_g \leq r_i$ (when they are equal, region B does not exist). A proxy attempting to enter region C would feel only force repelling it from the target. A proxy in region A would result in only attractive guidance. Finally, a proxy somewhere in region B would result in no net force.

In Figure 3, region C corresponds to forbidden region feedback only, A represents only guidance feedback, and in region B, little to no net force is experienced due to virtual fixtures associated with the target. When inequalities 1, 2 and 3 are satisfied, force feedback from virtual fixtures associated with the target is as described above. Now consider forbidden region virtual fixtures due to obstacles. Forbidden region virtual fixtures from objects described as point cloud geometries are amenable to monolithic processing [27]. Consider resolving forbidden region force feedback from obstacles with haptic guidance towards the target. Suppose an obstacle obstructs a direct path to the target, as shown in Figure 4

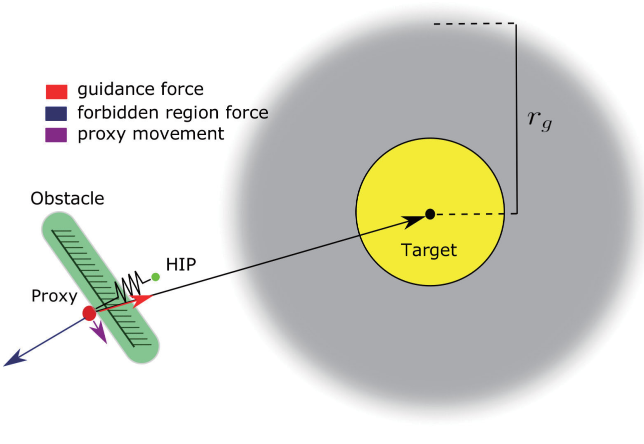

Fig. 4. Obstacle with forbidden region virtual fixture obstructing guidance virtual fixture. The reader is referred to [28] and [27] for a description of HIP and proxy methods for haptic rendering.

The danger occurs if the HIP moves far enough away from the proxy such that once the proxy has a clear, direct path to the HIP, it moves a great distance in a short period of time. This is a problem since the robot may move too fast for the pretouch sensor to detect an obstacle in time. While limiting proxy or robot velocity may mitigate this problem, it can reduce responsiveness to moving objects or anything else that requires fast motion. Instead, the proposed method can overcome the problem by accelerating the rate of navigation around obstacles as well as limiting the haptic guidance force to magnitude M. (M should be chosen such that the majority of force feedback is due to forbidden regions when the HIP is a certain distance from the proxy). This is outlined as pseudocode here.

Algorithm 1 Guidance Force Component

for each haptic update **do**
  Listen for latest published target frame
  Compute translation vector $\vec{t}_T^P$ proxy to target
  if $\left|\vec{t}_T^P\right| > r_g$ **then**
    Let $\vec{f}_G = k_G \vec{t}_T^P \left(1 - \frac{r_g}{|\vec{t}_T^P|}\right)$
    with fixed $k_G \in \mathbb{R}^+$
  else
    Let $\vec{f}_G = \vec{0}$
  end if
  if $|\vec{f}_G| \geq M$ **then**
    Set $\vec{f}_G = \vec{f}_G \frac{M}{|\vec{f}_G|}$
  end if
end for

The initial guidance force vector is thus acquired. Next, the initial force vector, $\vec{f}_{fr}$, generated from the forbidden region virtual fixtures can be obtained as shown in [27]. With these force vectors, the problem then becomes resolving the two signals to determine $\vec{f}_{\text{net}}$.

Algorithm 2 Net Force Determination

for each haptic message **do**
  Compute a plane Pl normal to $\vec{f}_G$
  Project $\vec{f}_{fr}$ onto Pl, call it $\vec{f}_{Pl}$
  $\vec{f}_{\text{net}} = \vec{f}_{fr} + \vec{f}_G + \vec{f}_{Pl}$
  Publish $\vec{f}_{\text{net}}$ to haptic device
end for

The effect of this method is twofold:

1) The user is guided towards the target as desired in free motion - when in conflict with forbidden region, the forbidden region effects are dominant.
2) Proxy movement is expedited along convex surfaces towards the target — the HIP is encouraged to move around obstacles obstructing a clear path to the target.

The proxy in this method is still susceptible to concave surfaces or crevices in which the proxy may get stuck. With a reduced guidance force, the user need only overcome a slight force if an alternative route is readily apparent.

IV. Target Tracking and Occupancy Grid

Methods to achieve object tracking can take advantage of many different types of information. Vision systems such as the Kinect can leverage geometric and color information to segment an object from the background, and then track it using the resultant reference model. However, there are many situations in which objects are tracked without real-time geometric or color information because it is either not available or not as robust as other methods. Small, highly maneuverable proximity sensors could quickly build a map instead. In this work, an augmented reality (AR) tag indicates an area of interest with little to no prior geometric information.

Voxel grid representations of space efficiently represent occupied regions and are user friendly. In this work, Octomap [29], an open source voxel grid that is based on the OcTree data structure, was used. A voxel grid presents the user with a clear indication of whether a region is occupied, unoccupied, or unknown. If the user were instead viewing a point cloud, he or she may spend an excessive amount of time exploring a region.

V. Experimental Implementation

A. System Setup

The setup includes a bilateral teleoperation architecture with a master console and telerobotic slave device interacting with a remote task environment. Software components were developed and executed in ROS.

1) Telerobotic Platform: The slave robotic platform is the PR2 from Willow Garage. The end effector of the PR2 was outfitted with a pretouch sensor array.

2) Pretouch Element: The robot is equipped with an optical pretouch sensor on one of its fingertips in order to facilitate exploration. It consists of six STMicroelectronics VL6180X [30] proximity sensing modules, each of which can sense distance within 1 to 10 cm at a rate of 33 Hz. There are two sensing modules on each side of the finger, one at the front, and one on the pad of the finger. The sensor is completely integrated into the robot's gripper.

3) Master Console: The master console both presents feedback to the human operator as well as receives user input motion commands, as shown in Figure 5(a). Visual feedback (a voxel representation of the target object) is rendered on an LED monitor while haptic feedback is presented through a 3DOF haptic device, the Geomagic® Touch™. The feedback reflects forces calculated to both prevent unwanted collisions (forbidden region virtual fixtures) as well as guide the user towards the exploration target (guidance virtual fixtures).

4) Telerobotic Task: A controllable circular conveyor belt, as shown in Figure 5(b), was designed to enable granularity in movement speed for experiments. An object could be translated and rotated with the conveyor belt, which sits flush to a tabletop. With this setup, a moving exploration target will be refined with the pretouch optical sensors on the slave device.

(a) (b)

Fig. 5. The experimental setup. (a) The user operates a 6-DOF input device to control the robot. A location to search and explored occupied voxels will be displayed in RVIZ. (b) The robot carries out the teleoperator's commands as a conveyer belt changes the pose of the object.

B. Experiment Workflow

This work evaluates the effect of haptic feedback on operator performance in exploring a completely unseen object. Two sets of 10 one-minute object refinement trials were conducted, one with haptic feedback and one without. Visual feedback was provided by RVIZ. At the start of each trial, the user is shown the kinematic model of the robot and a marked location to be explored. This location is tracked using an AR tag, which simulates a location tracked through non-geometric means.

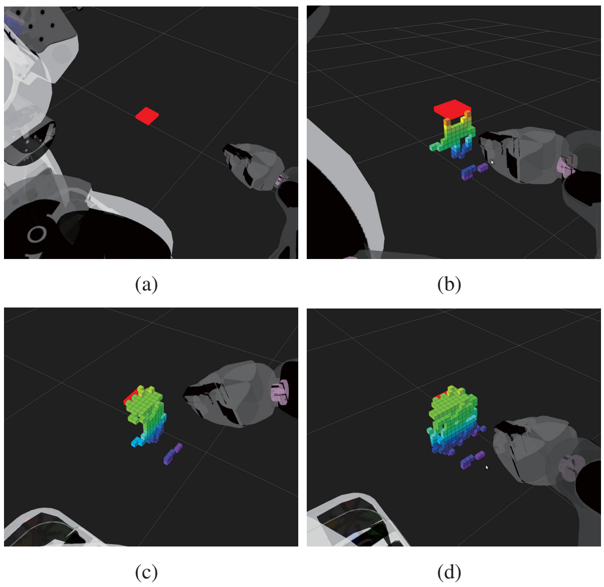

Fig. 6. The marker and explored occupied voxels are shown to the user during the experiment. (a) Initially, the user is only shown the marker of the location to be explored. (b) Additional voxels are discovered as the user explores one side of the object. (c) Voxels on top of the object are explored. (d) Voxels on the front and back of the object are found as the user continues to explore until the end of the experiment.

In each trial, the object moves back and forth at a speed of 0.5 $\frac{cm}{s}$ in front of the robot, with pauses in motion ranging from 0 to 4 seconds. During the trial the user attempts to explore an unseen box with a width, height and length of

9 cm, and any occupied voxels that he or she discovers are displayed and recorded. An example of the visual feedback that the user receives is illustrated in Figure 6. The distance between the end effector and the marked location was also recorded. Note that the red marker in Figure 6 does not represent any surface geometrical information about the unseen target. It reflects the general target location.

VI. Results

Throughout the 10 trials with haptic feedback, the user explored more than twice the number of occupied voxels when compared to visual feedback only. In order to compare the fit of the discovered voxels to the ground truth, the iterative closest points algorithm was performed between the collected occupied voxel centers and the true voxel centers of the explored cube.

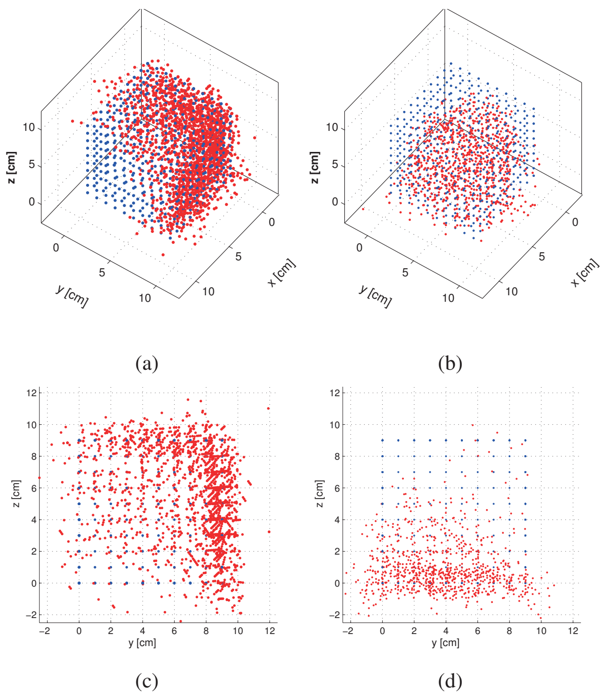

Fig. 7. Cumulative discovered occupied voxels gathered across all trials in red with the true voxel centers in blue: (a) and (b) show an isometric view of the results with and without haptic virtual fixtures respectively. (c) and (d) present side views of the data shown in (a) and (b) respectively.

In Figure 7(a) the discovered voxels across all 10 haptic feedback trials have been fitted to the ground truth model and plotted from an isometric perspective in red. Ground truth is plotted in blue. Figure 7(c) offers a side view of the same data. Figures 7(b) and 7(d) offer the corresponding information for the non-haptic feedback trials. The experiment also showed that haptic feedback allowed the user to bring the end effector closer to the object. In Figure 8, the distance between the fingertip and the center of the upper surface of the target box is plotted. Specifically, the average distance and standard deviation across all trials is plotted for each point in time.

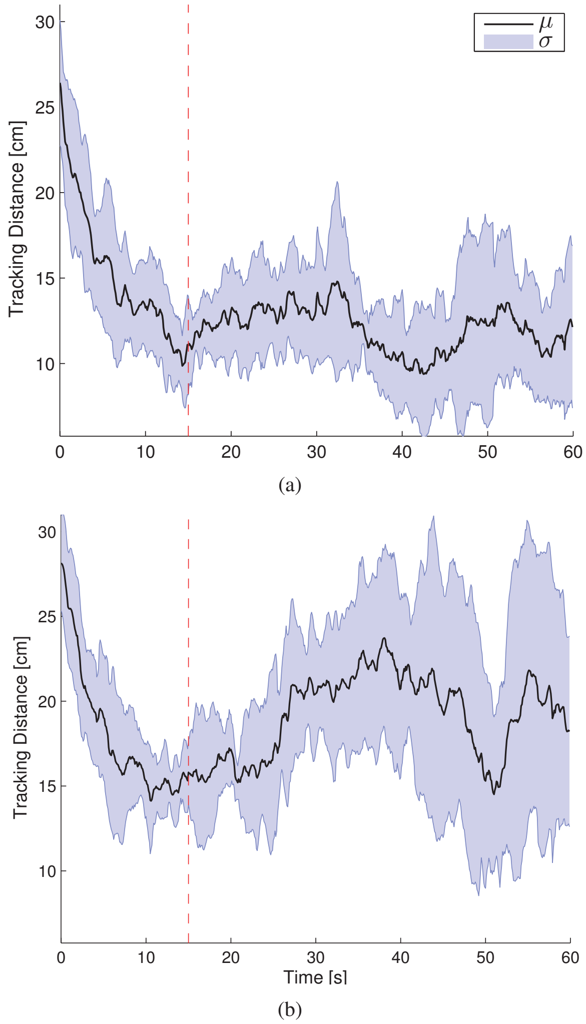

Fig. 8. Distance between the end effector and the tracked object during the 10 trials. The average distance across time is plotted in black, while the blue envelope illustrates the standard deviation of the distance. In each trial, the user spent approximately the first 15 seconds approaching the object, which is marked with a red dashed line. (a) With haptic feedback. (b) With visual feedback only.

Table I summarizes the findings of the experiment.

TABLE I
EXPERIMENTAL RESULTS OF THE EXPLORATION TASKS

	Marker Only	**Haptic Feedback**
Total Trials	10	10
Avg Sensed Voxels	88.90±21.46	192.5±24.38
Avg Distance [cm]	18.70±9.621	12.60± 6.121
Avg ICP RMS [cm]	0.9418 ±0.3396	0.7779 ± 8.801e-04
Total Collisions	7	3

VII. Conclusions

This paper presents a system for exploring dynamic objects and regions that cannot be well observed by cameras or long range-depth sensors due to space constraints, accessibility, or occlusions. Given a region to search, the teleoperator could use pretouch sensing to recover geometric information about an object as haptic feedback both encourages more aggressive exploration and prevents collisions. As the search continues, explored locations are added to the OctoMap and displayed to the user. The results of this work showed haptic feedback allowed the user to collect significantly more occupied voxels than when only utilizing visual feedback.

Indeed, the haptic feedback trials yielded less than half of the collisions that the non-haptic feedback trials did, and the user was able to consistently get closer to the object of interest, resulting in over twice the exploration coverage while inducing fewer collisions. While the haptic feedback encouraged the user to approach the object when far away, it also allowed the user to search more aggressively and focus on exploration, rather than having to search more conservatively in order to avoid collisions. The increased exploration of voxels was likely a direct result of heightened operator confidence while operating with haptic virtual fixtures; by spending more time close to the object, the object was more likely to be within the sensor's measurement range.

References

[1] B. Yang, P. Lancaster, and J. R. Smith, "Pre-touch sensing for sequential manipulation," in *Robotics and Automation, 2000. Proceedings. ICRA'17. IEEE International Conference on*. IEEE, 2017.

[2] A. Peon and D. Prattichizzo, "Reaction times to constraint violation in haptics: comparing vibration, visual and audio stimuli," in *World Haptics Conference (WHC), 2013*, April 2013, pp. 657–661.

[3] S. Izadi, D. Kim, O. Hilliges, D. Molyneaux, R. Newcombe, P. Kohli, J. Shotton, S. Hodges, D. Freeman, A. Davison, *et al.*, "Kinectfusion: real-time 3d reconstruction and interaction using a moving depth camera," in *Proceedings of the 24th annual ACM symposium on User interface software and technology*. ACM, 2011, pp. 559–568.

[4] A. Petrovskaya, O. Khatib, S. Thrun, and A. Y. Ng, "Bayesian estimation for autonomous object manipulation based on tactile sensors," in *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*. IEEE, 2006, pp. 707–714.

[5] B. Mayton, L. LeGrand, and J. R. Smith, "An electric field pretouch system for grasping and co-manipulation," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*. IEEE, 2010, pp. 831–838.

[6] D. Guo, P. Lancaster, L.-T. Jiang, F. Sun, and J. R. Smith, "Transmissive optical pretouch sensing for robotic grasping," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*. IEEE, 2015, pp. 5891–5897.

[7] K. Koyama, Y. Suzuki, A. Ming, and M. Shimojo, "Integrated control of a multi-fingered hand and arm using proximity sensors on the fingertips," in *Robotics and Automation (ICRA), 2016 IEEE International Conference on*. IEEE, 2016, pp. 4282–4288.

[8] A. Maldonado, H. Alvarez, and M. Beetz, "Improving robot manipulation through fingertip perception," in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*. IEEE, 2012, pp. 2947–2954.

[9] K. Hsiao, P. Nangeroni, M. Huber, A. Saxena, and A. Y. Ng, "Reactive grasping using optical proximity sensors," in *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*. IEEE, 2009, pp. 2098–2105.

[10] L.-T. Jiang and J. R. Smith, "Seashell effect pretouch sensing for robotic grasping," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 2851–2858.

[11] L.-T. Jiang and J. Smith, "A unified framework for grasping and shape acquisition via pretouch sensing," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, May 2013, pp. 999–1005.

[12] K. Huang, L.-T. Jiang, J. R. Smith, and H. J. Chizeck, "Sensor-aided teleoperated grasping of transparent objects," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 2015, pp. 4953–4959.

[13] P. F. Hokayem and M. W. Spong, "Bilateral teleoperation: An historical survey," *Automatica*, vol. 42, no. 12, pp. 2035–2057, 2006.

[14] D.-S. Kwon, J.-H. Ryu, P.-M. Lee, and S.-W. Hong, "Design of a teleoperation controller for an underwater manipulator," in *Robotics and Automation, 2000. Proceedings. ICRA'00. IEEE International Conference on*, vol. 4. IEEE, 2000, pp. 3114–3119.

[15] D. Hainsworth, "Teleoperation user interfaces for mining robotics," *Autonomous Robots*, vol. 11, no. 1, pp. 19–28, 2001.

[16] F. Ryden and H. Chizeck, "A proxy method for real-time 3-dof haptic rendering of streaming point cloud data," *Haptics, IEEE Transactions on*, vol. 6, no. 3, pp. 257–267, 2013.

[17] F. Ryden, S. Kosari, and H. Chizeck, "Proxy method for fast haptic rendering from time varying point clouds," in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, 2011, pp. 2614–2619.

[18] A. Leeper, S. Chan, K. Hsiao, M. Ciocarlie, and K. Salisbury, "Constraint-based haptic rendering of point data for teleoperated robot grasping," in *Haptics Symposium (HAPTICS), 2012 IEEE*, March 2012, pp. 377–383.

[19] R. Kumar, A. Kapoor, and R. H. Taylor, "Preliminary experiments in robot/human cooperative microinjection," in *Intelligent Robots and Systems, 2003.(IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on*, vol. 4. IEEE, 2003, pp. 3186–3191.

[20] A. Bettini, S. Lang, A. Okamura, and G. Hager, "Vision assisted control for manipulation using virtual fixtures: experiments at macro and micro scales," in *Robotics and Automation, 2002. Proceedings. ICRA '02. IEEE International Conference on*, vol. 4, 2002, pp. 3354–3361 vol.4.

[21] P. Marayong, M. Li, A. Okamura, and G. Hager, "Spatial motion constraints: theory and demonstrations for robot guidance using virtual fixtures," in *Robotics and Automation, 2003. Proceedings. ICRA '03. IEEE International Conference on*, vol. 2, Sept 2003, pp. 1954–1959 vol.2.

[22] A. Saxena, L. Wong, M. Quigley, and A. Y. Ng, "A vision-based system for grasping novel objects in cluttered environments," in *Robotics Research*. Springer, 2011, pp. 337–348.

[23] C. Goldfeder, M. Ciocarlie, J. Peretzman, H. Dang, and P. K. Allen, "Data-driven grasping with partial sensor data," in *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*. IEEE, 2009, pp. 1278–1283.

[24] D. Rao, Q. V. Le, T. Phoka, M. Quigley, A. Sudsang, and A. Y. Ng, "Grasping novel objects with depth segmentation," in *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*. IEEE, 2010, pp. 2578–2585.

[25] J. Aleotti, D. L. Rizzini, and S. Caselli, "Perception and grasping of object parts from active robot exploration," *Journal of Intelligent & Robotic Systems*, vol. 76, no. 3-4, pp. 401–425, 2014.

[26] G. Kahn, P. Sujan, S. Patil, S. Bopardikar, J. Ryde, K. Goldberg, and P. Abbeel, "Active exploration using trajectory optimization for robotic grasping in the presence of occlusions," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 2015, pp. 4783–4790.

[27] F. Ryden and H. Chizeck, "Forbidden-region virtual fixtures from streaming point clouds: Remotely touching and protecting a beating heart," in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, Oct 2012, pp. 3308–3313.

[28] D. C. Ruspini, K. Kolarov, and O. Khatib, "The haptic display of complex graphical environments," in *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., 1997, pp. 345–352.

[29] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "Octomap: An efficient probabilistic 3d mapping framework based on octrees," *Autonomous Robots*, vol. 34, no. 3, pp. 189–206, 2013.

[30] *Proximity and ambient light sensing (ALS) module*, STMicroelectronics, 8 2014, rev. 6.