

Introspective Perception for Long-term Aerial Telemanipulation with Virtual Reality

Jongseok Lee^{1,2}, Ribin Balachandran¹, Konstantin Kondak¹, Andre Coelho³, Marco De Stefano¹, Matthias Humt¹, Jianxiang Feng⁴, Tamim Asfour² and Rudolph Triebel^{1,2}

¹Institute of Robotics and Mechatronics, German Aerospace Center (DLR)

²Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology (KIT)

³Apptronik Inc

⁴Agile Robots AG

Corresponding author: Jongseok Lee (email: jongseok.lee@dlr.de).

This work is also supported by the Helmholtz Association's Initiative and Networking Fund (INF) under the Helmholtz AI platform grant agreement (ID ZT-I-PF-5-1) and by the European Commission under the contract 644271 EU2020 AEROARMS and 824990 EU2020 RIMA.

ABSTRACT This article presents a novel telemanipulation system to advance aerial manipulation in dynamic and unstructured environments. The proposed system features not only a haptic device, but also a virtual reality (VR) interface that provides real-time 3D displays of the robot's workspace, as well as a haptic guidance to its remotely located operator. To realize this, multiple sensors, namely a LiDAR, cameras, and IMUs are utilized. For processing of the acquired sensory data, pose estimation pipelines are devised for industrial objects of known and unknown geometries. We further propose an active learning pipeline in order to increase the sample efficiency of a pipeline component that relies on Deep Neural Networks (DNNs) based object detection. All of these algorithms jointly address various challenges encountered during the execution of perception tasks in industrial scenarios. In the experiments, exhaustive ablation studies are provided to validate the proposed pipelines. Methodologically, these results commonly suggest how an awareness of the algorithms' own failures and uncertainty ('introspection') can be used to tackle the encountered problems. In addition, outdoor experiments are conducted to evaluate the effectiveness of the overall system in enhancing aerial manipulation capabilities. In particular, with flight campaigns over days and nights, from spring to winter, and with different users and locations, we demonstrate more than 70 robust executions of pick-and-place, force application, and peg-in-hole tasks with the DLR cable-Suspended Aerial Manipulator (SAM). As a result, we show the viability of the proposed system in future industrial applications. The project website can be accessed at <https://sites.google.com/view/vr-sam/>.

INDEX TERMS Virtual Reality and Interfaces, Telerobotics and Teleoperation, Deep Learning in Grasping and Manipulation, and Aerial Systems: Perception and Autonomy.

I. INTRODUCTION

EVERY critical infrastructure needs to be inspected and maintained. Often, these jobs are required within areas that are difficult to reach. Some concrete examples are inspecting oil and gas pipes in refineries, maintenance of antennas at cell towers, and repairing the blades of wind turbines. For such tasks, the current state-of-the-art is to deploy human technicians, who reach the high altitude areas through specialized equipment such as ropes, scaffolding, and cranes, and perform the required manipulation tasks. The

current solutions are dangerous and costly due to inherent risks in ensuring the safety of technicians. One of the research directions for these applications is aerial manipulation [2]. An aerial manipulation system is composed of robotic manipulators and a controlled flying platform [3]–[6]. The platform allows coarse positioning, while the manipulator allows dexterous grasping and manipulation for complex tasks. Hence, these aerial platforms extend the mobility of robotic manipulators, which can be deployed at high altitudes



Figure 1: Left: the cable-Suspended Aerial Manipulator, dubbed SAM [1] during field experiment. Right: a ground station where an operator remotely controls the robotic arm through a haptic interface. In real-world applications of bilateral teleoperation, the operator is often remotely located without visual contact to the robot.

above ground, increasing safety for human workers while reducing costs.

In this article, the real-world applications of aerial manipulators are envisioned for several industrial scenarios in dynamic and unstructured environments. For these industrial applications of aerial manipulators, our current interests are in bilateral teleoperation concepts, that is, a human operator remotely controls the robotic manipulator from a safe area on the ground and receives visual and haptic feedback from the robot. This increases the safety of the human operator while robots perform their tasks in dangerous environments [7], [8]. Such a concept is motivated by having a robotic system with a human-in-the-loop, where the system can leverage human intelligence to reliably accomplish its missions. To realize this, existing work has focused on relevant components of the system, namely force feedback teleoperation under time delays [9], [10], shared autonomy [11], human-machine interfaces [12]–[14], and robotic perception for aerial manipulators [15], [16].

Building upon the aforementioned developments, we propose a novel virtual reality (VR)-based telemanipulation system for an aerial manipulation system operating in industrial scenarios. Fig. 1 and Fig. 2 illustrate the main idea. The proposed system is intended for real-world scenarios, where the remotely located robot performs aerial manipulation tasks, while its human operator is inside a ground station without having direct visual contact with the robot (Fig. 1). To this end, we propose a system that involves not only a haptic device to enable the *sense of touch* for the operator, but also a VR to increase the *sense of vision* (Fig. 2). While the live video streams can also provide a certain level of situation awareness to the operator, several studies confirm that adding a virtual environment where one can change its sight-of-view, zoom in and out, and further provide haptic guidance,

supports the operator in accomplishing the tasks [17]–[21]. Our own field studies also confirm that the enhancement of live video streams with 3D visual feedback and haptic guidance can enhance the manipulation capabilities of aerial robots.

The main novelty of our VR based concept is its realization with a *fully on-board perception* system for a *floating-base robot*, which *does not rely* on any external sensors like Vicon, or any pre-generated maps in *outdoor environments*. Instead, *multiple sensors*, namely LiDAR, a monocular camera, a pair of stereo cameras, and inertial measurement units (IMU) are jointly utilized (Tab. 1). To achieve this, we propose object pose estimation and active learning pipelines. First, in order to virtually display industrial objects with known geometry, we provide a simple extension of a marker tracking algorithm [22] by combining it with the Simultaneous Location And Mapping (SLAM) on-board. Second, if the objects of interest are geometrically unknown, we devise a LiDAR based pose estimation pipeline that combines LiDAR Odometry And Mapping (LOAM [23]) with a pose graph, a point cloud registration algorithm [24], and a detector based on deep neural network (DNN) [25]. For both cases, the combinations are facilitated by an introspection module that identifies the reliability of pose estimation. Finally, we present a pool based active learning pipeline, which uses an explicit representation of DNN’s uncertainty to generate the most informative samples for a DNN to learn from. This enhances the sample efficiency of deploying DNN based algorithms in outdoor environments. We identify certain real-world challenges and describe in detail how these introspective approaches can mitigate these challenges.

With the DLR’s SAM platform [1], the feasibility and benefits of the proposed idea are examined. To this end, we first present ablation studies on the designed pipelines with

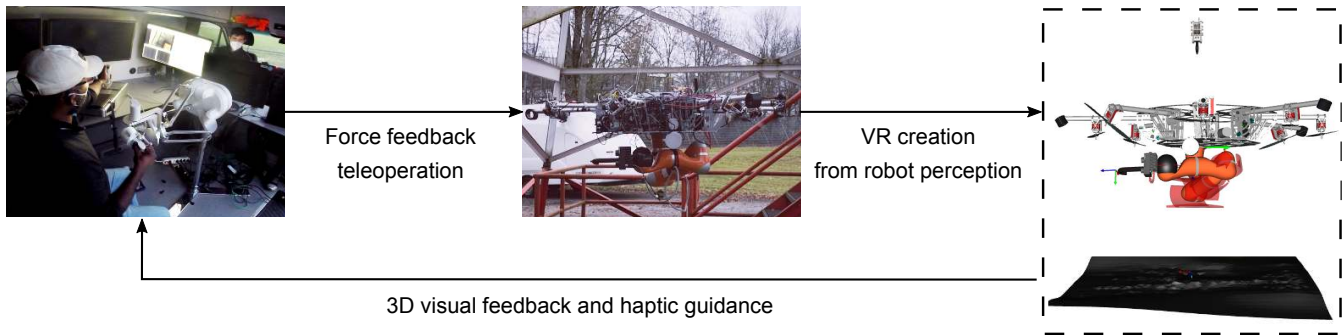


Figure 2: The proposed telemanipulation system with VR from robot perception and active learning. In the proposed system, the robot creates VR of its workspaces as a 3D visual feedback to the human operator, and further provides a haptic guidance. The main novelty of this work is the realization of such a system for real-world scenarios.

indoor and outdoor data-sets from the robot sensors. Here, the influence of each component is examined with regard to mitigating the identified challenges, and we show the feasibility of creating real-time VR, which can closely match the real workspaces of the robot. Moreover, the effectiveness of the proposed method is shown through outdoor experiments within the industrial scenario considered. This scenario, which was designed within the scope of the EU project AEROARMS [26], is relevant to inspection and maintenance applications for the gas and oil industry. It involves pick-and-place and force-exertion tasks during the mission, which is to deploy a robotic crawler to automate pipe inspection routines. In addition, the SAM platform that performs peg-in-hole tasks with a margin of error less than 2.5 mm is considered, which is one of the standard manipulation tasks in industrial settings. By executing over 70 executions of the aforementioned tasks over days and nights, from spring to winter, and with different users and locations, the benefits of our VR based telemanipulation concept are illustrated for enhancing aerial manipulation capabilities in real-world industrial applications.

In summary, the key contributions of this work are as follows. First, we propose an advanced VR based telemanipulation system for aerial manipulation, which provides a 3D visual feedback and haptic guidance. The system neither requires any external sensors nor pre-generated maps, has been evaluated outside laboratory settings, and can cope with the challenges of a floating-base system. Moreover, multiple sensors are fused to exploit their respective strengths for the given perception tasks. Second, we devise object pose estimation and active learning pipelines to realize the proposed system in dynamic and unstructured environments. Challenges to existing methods are reported, and several ablation studies are provided to validate the proposed approaches. Methodologically, this work suggests the relevance of robotic introspection in the realization of VR based telepresence robots with aerial manipulation capabilities. Finally, we performed exhaustive flight experiments over extended durations, including 40 task executions in outdoor environments, 27 task executions within a user validation

study, and the operation of the system at night. Thus, we establish the proposed concept as a viable future option for real-world industrial applications.

The paper starts with a survey of related work (Section II) and provide the system description of the hardware of the SAM robot, human-machine interfaces, sensor choices, and integration (Section A). We formulate the problem of VR creation and identify challenges in realizing the system (Section B). Then, the designed pipelines are presented, which are to address these challenges (Section IV). In Section A, we provide ablation studies to validate the designed framework, while Section B contains the results of our flight experiments. We report the lessons learned in Section D and conclude the work with some future extensions in Section VI.

Relation to Previous Publications This paper extends the author’s previous publications, namely [27] and [28]. In terms of methodology, we provide a LiDAR based pose estimation pipeline (Section B). This extension enables the creation of VR without relying on markers, which is required in industrial scenarios. The developed active learning pipeline for object detection (Section C) extends and brings the theoretical framework [28] to practical applications. Furthermore, with respect to experimental contributions, this article provides new ablation studies that are associated with the new methods. Most importantly, exhaustive outdoor experiments for manipulation tasks are performed in addition to examine the benefits of the proposed VR based concept over extended durations and characterize its technical readiness for industrial applications. We note that this article, previously appeared in [29], is being archived at IEEE Transactions on Field Robotics, following the suggestions of the editor-in-chief.

II. Related Work

The proposed VR based concept advances the area of VR interfaces for robotics. The comparison of this work to existing works is summarized in Tab. 1. The current literature from different domains of robotic research is discussed, which are, pose estimation (Section A and B), and active learning with DNNs (Section C). Importantly, we stress that

	Outside the laboratory settings?	No external sensors or pre generated map?	Floating-base manipulation system?	Multiple exteroceptive sensors?
AeroVR [13]	X	X	✓	X
ARMAR-6 [30]	X	✓	X	X
ModelSegmentation [31]	X	✓	X	X
AvatarDrone [12]	X	X	✓	X
PaintCopter [32]	X	X	✓	X
AR [33]	X	✓	X	X
AR [34]	X	✓	X	X
GraspLook [35]	X	✓	X	X
The proposed system	✓	✓	✓	✓

Table 1: Comparisons between the existing VR based robotic systems and the proposed system.

this work is not to advance the state-of-the-art methods in these two areas. Rather, the aim is to apply and extend them to realize a working system for the given industrial scenarios. For example, the provided extension of a marker tracking algorithm with visual-inertial SLAM is not the main contribution of this paper. Finally, we further locate our work within the literature of aerial robotic perception in field applications.

Virtual Reality Interfaces In the past, several VR interfaces have been widely used in robotics, including aerial systems [36]. So far, the presented approaches often create the VR either by using external sensors such as Vicon and a-priori generated maps. In particular, a priori generated maps were utilized for the applications of VR in aerial painting [32]. For aerial manipulation, the Vicon system [13] and a portable sensor kit [12] were used instead of sophisticated onboard perception. Recently, many VR techniques have gained interest in the robotic manipulation community. In this way, many works [37], [38] let a human perform a demonstration in VR, and transfer the demonstrated manipulation skills to real robots. These works greatly show the synergy between VR and robotics. As this paper demonstrates the feasibility of creating VR with on-board sensors only, the work can contribute to many of these works in showing how one can create a VR for robotics.

In contrast, many researchers aimed to provide VR of the remote scene by applying 3D reconstruction techniques [31], [39]. For example, [31] presents an approach using a RGB-D camera. As the main challenge of reconstruction based methods is the limited bandwidth in communication, [31] proposes an object recognition pipeline, that is, replace the detected object with sparse virtual meshes and discard the dense sensor data. [30] uses RGB-D sensor to construct VR for affordance-based manipulation with a humanoid, while [33] and [34] create augmented reality for a drone and a manipulator, respectively. [19] conducts a user study and argues that the point clouds of RGB-D sensors are noisy and inaccurate (with artifacts), which motivates for point cloud pre-processing methods for telepresence applications [19]. In

contrast, our approach is based on scene graphs (Section B) with pose estimation, which is an alternative to 3D reconstruction methods. Finally, the main novelties are illustrated in Tab. 1, which are the realizations of a VR based telemanipulation system for outdoor environments using multiple sensors together. No external sensors or pre-generated maps are used, while dealing with specific challenges of a floating base manipulation system, that is, the surface that holds a robotic arm is constantly changing over time, thereby inducing motions for the attached sensors.

Object Pose Estimation One of the crucial components in the proposed framework are algorithms for estimating object poses. This is because we utilize a scene graph representation, which requires 6D pose of the objects for creating a 3D display, as opposed to a 3D reconstruction of the remote site. As the literature is vast, we refer to the survey [40] for a comprehensive review. In this work, the main novelty is the working solutions for the considered application, which is tailored towards realizing the proposed VR system. For this, the two scenarios are discussed below. These are visual object pose estimation for objects of known geometry, and LiDAR based method for unknown geometry.

If the object is known and accessible a-priori, one of the robust solutions is to use fiducial marker systems. Fiducial markers, which create artificial features in the scene for pose estimation, are widely used in robotics. The use-cases are for creating the ground truths [41], where environments are known [42], and for simplifying the problem in lieu of sophisticated perception [43]. However, as the herein aim is on real-time VR creation, this use-case demands stringent requirements on their limitations in run-time, inherent time-delays and robustness. Therefore, an extension of ARToolKit-Plus [22] is provided with an on-board visual-inertial SLAM system.

For LiDAR, point cloud registration is often used for pose estimation. By finding the transformation between the current scans and a CAD model of an object, we can obtain a 6D pose of an object. In general, point cloud registration algorithms can be classified as local [24], [44], [45] or global [46],

and model based [47] or learning based [48], [49]. As CAD models of objects are often not available in the given industrial scenario, a DNN based detector and the idea of LOAM with pose graphs are combined, in order to obtain robust object pose estimates that cope with occlusions, moving parts and view point variations in the scene.

Active Learning for Neural Networks The motivations are the considerations of field robotic applications of DNN based object detectors. Here, the need for labeled data can cause overhead in development processes, especially when considering a long-term deployment of learning systems in outdoor environments. For example, weather conditions can change depending on the seasons, and we need to efficiently create labeled data. Active learning provides a principled way to reduce manual annotations by explicitly picking data that are worth being labeled. One way to autonomously generate the "worth" of an unlabeled sample is to use the uncertainty of DNNs. In the past, for robot perception, we find active learning frameworks using random forests, Gaussian processes, etc [50], [51] while for DNNs, [52] pioneered an active learning approach based on Bayesian Neural Networks, i.e., a probabilistic or stochastic DNN [53], which offers a principled method for uncertainty quantification. Recent works can also be found on active learning for DNN based object detectors [54], [55], where the focus is on adaptations of active learning to existing object detection frameworks. These include new acquisition functions (or selection criteria) and how uncertainty estimates are generated.

For uncertainty quantification in DNNs, the so-called Monte Carlo dropout (MC dropout [56]) has gained popularity recently. The main advantage of MC dropout is that it is relatively easy to use and scale to large data-set. However, MC dropout requires a specific stochastic regularization called dropout [57]. This limits its use on already well-trained architectures, because the current DNN based object detectors are often trained with other regularization techniques such as batch normalization [58]. The deep ensemble [59] is another scalable framework with a relaxed assumption on the model. Unfortunately, a deep ensemble requires the training of several large DNN models to form an ensemble. This technique is generally popular, but it is difficult to utilize in active learning due to inefficiency in training. In this article, a previous work [28] on uncertainty quantification of DNNs is utilized. The main motivations are scalability to large architectures and data-sets, training-free feature that requires no changes in network architectures and no retraining, and the ability to model every layer of DNNs as Bayesian. These aspects can make the given framework well suited for active learning in practice, and thus, this work attempts to provide an extension to active learning for its real-world applications in robotics.

III. System Description, Problem Statement and Identified Challenges

This paper investigates how a robot can create VR of a remote scene using on-board sensors and computations. This is to

improve situational awareness of the human operator in real-world applications. To set the scene for the work, we first describe the system integration that is needed to implement the proposed VR based telemanipulation concept. Then, the problem of VR creation, using on-board sensors with a scene graph approach, is formulated. The limitations of the off-the-shelf methods are then presented, which hinder the realization of the proposed system in outdoor environments.

A. System Description

This section describes the used robotic systems with a focus on robot hardware, haptic device and VR interfaces, and sensors. Main features of the system is also discussed. Fig. 3 depicts an overview of our physical hardware.

Robot Hardware DLRs' SAM [1] is a novel aerial manipulation system for inspection and maintenance applications. SAM is composed of three modules, namely a carrier, a suspended cable platform, and a seven-degree-of-freedom (DoF) industrial robotic arm - KUKA LWR [60]. The purpose of the carrier is to transport the manipulation system to a desired location. In this work, we use a crane that provides safety, versatility, robustness and applicability for the considered industrial scenario¹. Then, a platform attached to the carrier via a rope autonomously damps out disturbances induced by the carrier, the environment, and the manipulator. This oscillation damping control is performed using eight propellers and three winches. Another important component of our system is the seven Dof torque controlled KUKA LWR [60], which features significantly more powerful versatile manipulation capabilities than many smaller existing manipulators. The main feature of the cable-suspension concept is that the weight of the SAM is supported by the carrier. Thus, the energy required to carry an aerial robot arm can be reduced. This allowed us to reduce the overall size, from a helicopter based system [5] to a relatively smaller robot, which enables operation in confined spaces. The helicopter based system had two rotors with overall diameter of 3.7m, while SAM can fit within 1.5m diameter. Moreover, the cables from the carrier can also be used to power SAM, which gives theoretically unlimited operation time. In the appendix, more details are provided regarding the platform control and different architectures. [1] can also be referenced for conceiving design and control aspects of SAM in detail. Recently, similar concepts are being investigated in [61], [62].

Haptic Devices and Virtual Reality Interfaces In this work, two haptic devices, namely a space qualified haptic device called the Space Joystick RJo [10], and also a six Dof force feedback device, Lambda (Force Dimension), are integrated to teleoperate the LWR in SAM. The VR interface of this work is based on Instant Player [63], which is a

¹On the other hand, there is no free lunch. Cranes may not be able to reach all desired locations, as they require available access routes by ground. There are also several industrial tasks where smaller robotic arms with less DoF may be sufficient


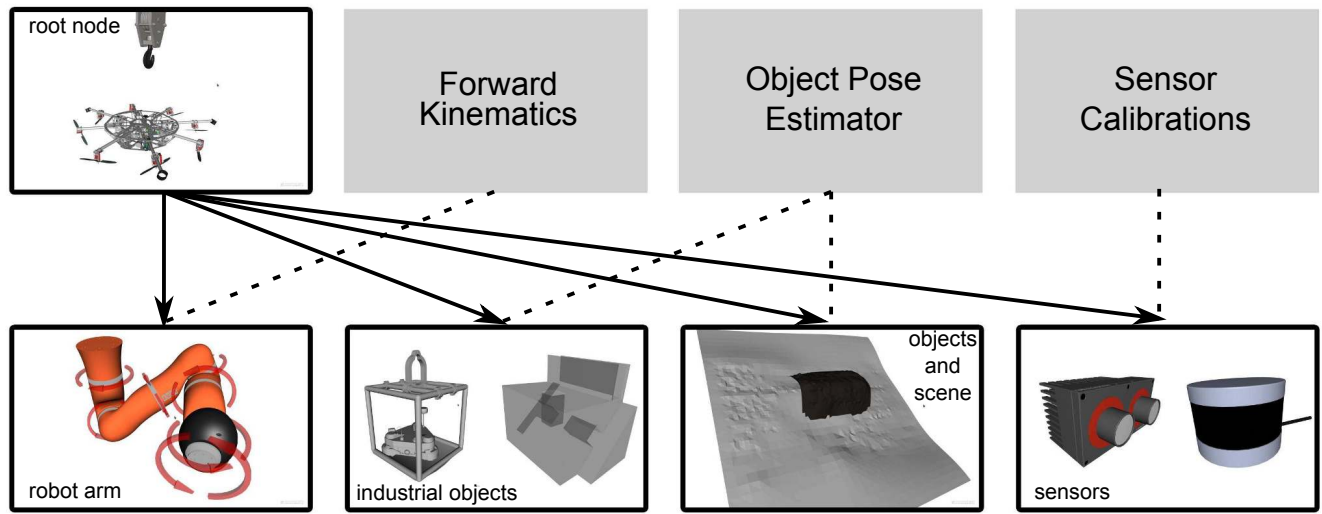


Figure 3: The concept of SAM with its integrated sensors and human-machine interfaces. Left: the concept involves the carriers such as a manned helicopter or a crane, which transports SAM to a desired location. Middle: SAM is equipped with a stereo camera at the end effector or the manipulator, a monocular camera as well as a LiDAR on the cable suspended platform. Right: haptic interfaces are integrated for teleoperating the robotic arm. Robot hardware constitutes of SAM, and the carriers (helicopter or crane), while the used haptic devices are Force Dimension Lambda and Space Joystick. Finally, the integrated sensors are a monocular camera (dubbed mako), a 3D LiDAR, and a stereo eye-in-hand camera (dubbed hc).


lightweight software that runs on standard laptops without GPUs (enhancing portability). Instant Player also supports various hierarchies of a scene graph to create the required display. Facebook’s head mounted display Oculus is also integrated as an option and uses Ubiquiti Bullet for the WiFi connection. The robot is equipped with advanced control strategies, namely whole-body teleoperation and adaptive shared control. The time-domain passivity approach [10] is employed to obtain stable teleoperation control under communication time delays, packet loss, and jitters. These control methods advance aerial manipulation capabilities [64] and [65]. The former presents a passivity based framework to enable time-delayed teleoperation of different hierarchically-sorted tasks through the use of multiple input devices. The latter presents a method to stabilize on-line adaptation of control authorities for the operator and the virtual assistance system in haptic shared control.

Sensor Choices and Integration We integrate several sensors to measure the robot’s own states, as well as to perceive the environment. More specifically, a KUKA LWR [60] is equipped with torque and position sensors, which measure its joint torques and angles. Furthermore, we integrate other sensors in the SAM for perception tasks. First, a camera (Allied Vision: mako) is integrated in the frame of the SAM to stream the overall operational space of the


robotic arm. This is because the operator prefers an eye-to-hand view, which is more natural to a human. The camera provides color images of 1292 by 964 px at 30Hz. Secondly, a stereo camera is integrated near the tool-center-point (tcp) of the robotic arm. This eye-in-hand set-up avoids occlusion of the camera view by the robotic arm and ensures proximity to the considered objects. These are crucial for the success of our image processing algorithms, i.e., the accuracy of visual marker tracking depends on the size of the markers and their distance to the sensors, while the depth sensing from the stereo depends on the baseline. We use a commercial 3D vision sensor the Roboception Rcvisard that provides built-in visual-inertial SLAM. The SLAM system originates from [66], [67] but we refer the reader to the company for more details. Rcvisard streams 1280 by 960 px images at 25Hz and SLAM estimates can be acquired at 200Hz by fusing it with an IMU. Lastly, as a step towards industrial application of SAM, we mount Velodyne PUK-LITE LiDAR on the frame of SAM, which provides 3D point clouds of the scene at 10Hz. We intend to use LiDAR for 3D object pose estimation, as well as SAM navigation in outdoor environments. Note that the minimum range is set to 0.9m while the maximum range of 100m is utilized. We designed and integrated the sensor stacks so that close-range perception is not affected. All perception algorithms are executed on the NVIDIA Jetson TX2.



(a) scene graph with flat hierarchy



(b) virtual reality



(d) side view

Figure 4: Illustration of the scene graph representation for the proposed VR framework. The root node is the base frame of the robot, while robot arm, industrial objects, scenes, and sensors are object nodes with transformation matrices as the edges. Forward kinematic provides state of the robot arm, and the fixed transformations to the robot sensors are obtained from extrinsic sensor calibrations. Then, the 6D estimates of object pose are obtained online using robot perception. Therefore, this work focuses on the object pose estimation for the realization of the proposed system.

In Appendix VI, more details on platform control, telemanipulation systems, and IT architectures are presented.

B. Problem Formulation and Identified Challenges

Assume that SAM is performing manipulation tasks far away from the human operator. So, the operator does not have direct visual contact to the scene, and the robot has to enhance the situational awareness of the operator. For this, SAM creates a VR of its environment and workspaces using on-board sensing and computations, and further provides haptic guidance via virtual fixtures [68]. Followed by the system level requirements, the problem formulation and the challenges of realizing such VR based telemanipulation concept are introduced next (see Fig. 4 and Fig.5).

The system level requirements are highlighted as follows. Firstly, the created VR has to accurately match the real remote site in real-time. This is because the operator needs visual feedback that reflects reality, and the performance of haptic guidance depends on the positioning accuracy and runtime. The latter is due to potential movements of the robot while hovering. Second, the robustness of the created VR is crucial to give a sense of trust to the human operator and further provide reliable haptic guidance. This means that the abnormalities in the object pose estimators are to be coped with, which often arises in outdoor environments. Last, the algorithms must run on-board the robot, and only send the transformation matrices through WiFi network (apart from an initialization phase, where surface reconstructed 3D models are sent). This is to avoid overloading of the communication channel for stable bilateral teleoperation. For example, both the sparse LiDAR point clouds and the dense stereo point clouds must be processed first, and only the pose of the objects must be sent through the WiFi network. The pose information contains only six float values, while continuous streaming of the point clouds require much more memory that grows with the number of points.

For VR creation that addresses aforementioned requirements, this work relies on a scene graph approach (shown in Fig. 4). A scene graph is general data structure with graph or tree like representations. It is used by the VR/AR softwares [63], in order to produce the real-time 3D visualizations. Mathematically, let S be a scene graph. It constitutes of sets of nodes and edges, denoted by (V, E) . The nodes V are any 3D models, while the edges E represent the spatial relationships. The root node V_{root} is chosen to be the robot's base frame, which is a fixed coordinate of the SAM platform. Then, a flat hierarchy of the scene graph [63] is assumed. This means the root node is a single "parent" to all other "child" nodes. In the given scenarios, the models to be displayed in VR are the sets of industrial objects, reconstructed external scenes, robotic arm, and the robot sensors. For a node of robotic arm V_{LBR} , the corresponding edges $E_{\text{LBR}}^{\text{root}}$ are readily provided by the forward kinematics. Similarly, the edge of three sensors, $E_{\text{hc}}^{\text{root}}$, $E_{\text{mako}}^{\text{root}}$ and $E_{\text{LiDAR}}^{\text{root}}$, are the outputs of the extrinsic camera calibration. These spatial relations or the

relevant transformation matrices are therefore fixed for the sensor nodes $V_{\text{hc}}^{\text{root}}$, $V_{\text{mako}}^{\text{root}}$ and $V_{\text{LiDAR}}^{\text{root}}$.

On the contrary, the spatial relations of industrial objects and external scenes are constantly changing, leading to the problem of pose estimation. Here, we divide the problem formulation into two sub-problems. The first sub-problem is when the object is known a-priori with available 3D models $(V_{\text{o1}}^{\text{root}}, E_{\text{o1}}^{\text{root}})$, while the second sub-problem is when the object is semantically known a-priori, but no primitives on the geometry exist $(V_{\text{o2}}^{\text{root}}, E_{\text{o2}}^{\text{root}})$. For the former, the corresponding edges are to be estimated. The latter involves the estimation of both the nodes and the edges. As articulated in Section A, the available raw sensor data are RGB camera images $I \in \mathbb{R}^{H \times W \times 3}$, where H and W are the image height and width, respectively. The images are obtained either from the eye-in-hand stereo camera (denoted by hc), or a monocular camera at the base (denoted by mako). A LiDAR, which is located also at the base, generates scans that are represented by the point clouds $P = (p_1, p_2, \dots, p_N) \in \mathbb{R}^{3 \times N}$. We also have a visual-inertial SLAM system at the end effector of the robotic arm, which outputs the rotation matrix R and translation vector t between the coordinate frames of the camera and a fixed world frame. In summary, the problem of VR creation can be formulated as estimating $E_{\text{o1}}^{\text{root}}$, $V_{\text{o2}}^{\text{root}}$ and $E_{\text{o2}}^{\text{root}}$ using the available sensory data from different cameras, an IMU and a LiDAR.

For this problem, several existing approaches can be applied. However, several practical challenges of directly applying these approaches have been identified from the field work (depicted in Fig. 5). For objects of known geometry, we can resort to marker based object pose estimation methods. For this, we cannot assume the holistic view of the markers. Violations of this assumption are caused by shadows, loss-of-sight or partial views of the markers. This results in failures when using off-the-shelf marker tracking methods. Moreover, in an industrial scenario, we cannot assume the availability of precise CAD models. Thus, point cloud registration methods cannot be directly employed. Tracking is also subject to occlusions and moving objects in front of the LiDAR, e.g. the robotic arm, and significant view point changes also result in less accurate 6D object poses while employing off-the-shelf methods such as iterative closest point algorithm. Finally, while deploying data-driven approaches for field robotics, the key to its success is the preparation of the data. The main challenges are the variations of scenes encountered during long-term deployment; darkness in the evenings or snow in winter are such examples. This means data has to be repeatedly collected for varying environment conditions, which is a laborious process. So, the question is how to make the data collection procedure more efficient so that DNNs can generalize. In the next section, these challenges are revisited after providing mathematical formulations of existing methods, which is then followed by our extensions to resolve these challenges.




Figure 5: Identified challenges for realizing our VR based concept. Top: (a,d,g,j) show the challenges associated with a marker tracking algorithm. Middle: (b,e,h,k) depict the challenges associated with directly applying point cloud registration methods for pose estimation. In particular for (e), precise geometry of objects are not available for pipe inspection scenario as an example, and therefore, its CAD models must be reconstructed online. Bottom: (c,f,i,l) visualize different scenes that a learning based method must cope with, when deployed for real-world applications. For example, a DNN trained in a laboratory, may not generalize to the scenes with (f) stair cases.




Figure 6: The proposed extension of ARToolKitPlus. Left: the position and orientation estimates of camera motion from a SLAM system infer the object when the markers are not detected. Right: linear and angular velocity estimates of the SLAM system are used with the time delay term t_d to predict the motion of the camera in $t + t_d$ seconds.

IV. The Proposed Methods

The aim is to create a 3D display of the robot and the objects so that the human operator can remote control the robotic arm from a distance. If done in real-time, the operator can see the VR and perform the tasks. Haptic guidance through virtual fixtures can further help the human operator during the execution of challenging manipulation tasks. So far, we have formulated the problem and also outlined the practical challenges. As previously discussed, the scene graph creation problem relies on the accurate, fast, and reliable 6D object pose estimation algorithms for the industrial objects of known and unknown geometries. This section describes the proposed pipeline for the object pose estimation.

A. The Proposed Pipeline for Objects of Known Geometry

Once the objects to be actively manipulated are known a-priori, i.e., the CAD models are available and the objects are physically accessible, the fiducial marker systems [22] can be exploited. These systems consist of a marker, which is a physical plane with black and white squared shapes (similar to QR codes), and a detection with a decoding algorithm. The key idea is to artificially create features in a plane that are physically attached to an object. Then, we can compute the pose of a camera in relation to a coordinate of the plane via a homography. Concretely, using the eye-in-hand camera (hc), the objective is to find the transformation matrix of the markers T_m^{hc} , expressed in the camera coordinate system, which consists of the rotation matrix R_m^{hc} and the translation vector t_m^{hc} . To do so, four corner points of the markers are

extracted, which are expressed in the marker coordinates $p_m = (x_m, y_m, 0)^T$ (hence $z_m = 0$ and given the size), and the image plane with pixels $p_{image} = (u_m, v_m, w_m)^T$. Then, the optimizer:

$$h(t) = \arg \min_h \sum_{i=1}^4 \rho(p_{i,image}(t), H_{image}^m(t) p_{i,m}(t)) \quad \text{where} \quad (1)$$

$$H_{image}^m(t) = \left(R_m^{hc}(t) + \frac{t_m^{hc}(t)}{d} n^T \right),$$

is the solution to the homography problem. Here, t denotes time, H_{image}^m is the homography matrix with h being its vector form, and ρ is a distance based cost function. Knowing the homography matrix, the desired rotation and translations can be obtained given the parameters of the intrinsic camera calibration: d and n . Typically, an algebraic formulation is used with the Direct Linear Transformation (DLT) algorithm [69]. We note that fiducial marker systems are widely adopted as ground truths in the robotics community due to their precision [41].

Challenges However, many existing fiducial markers systems [22], [41]–[43] do not address this work’s application scenarios, where aerial manipulation tasks in outdoor environments are considered. For example, shadows that are created by the robot can often destroy certain shapes of the markers and as a result, the methods would fail as the artificial visual features in the markers are occluded. Similarly, the eye-in-hand camera can lose view on the marker as the manipulator and the base can move rapidly. Lastly, the time

Algorithm 1: Robust marker localization algorithm with Visual-Inertial SLAM

```

     $I$     camera images from the eye-in-hand camera (hc).
     $m$     target marker identification.
input :   $i$     identification numbers of additional markers  $i = 1, 2, \dots, n$ .
            $t_d$    time delay parameter, computed online or prespecified.
            $\mathbf{T}_{hc}^w$  SLAM estimates of hc camera w.r.t a world coordinate.
output:  $\mathbf{T}_m^{hc}$  6D pose of the target marker  $m$  w.r.t the hc camera.

1 begin
2   /* Initialization */
3    $\mathbf{T}_m^{hc}(0), \mathbf{T}_i^{hc}(0) \leftarrow \text{multiART+}(I) \forall i$ ; // detect all the markers (Eq. 1)
4    $\mathbf{T}_m^i \leftarrow \text{marker\_init}(\mathbf{T}_m^{hc}(0), \mathbf{T}_i^{hc}(0)) \forall i$ ; // save all the relative poses
5   /* Main Loop */
6   while True do
7      $\mathbf{T}_m^{hc}(t), \mathbf{T}_i^{hc}(t) \leftarrow \text{multiART+}(I) \forall i$ ; // detect the markers (Eq. 1)
8     if all markers detected then
9        $\mathbf{T}_{m,i}^{hc}(t), \mathbf{T}_{m,i}^i(t) \leftarrow \text{trafo2m}(\mathbf{T}_m^{hc}(t), \mathbf{T}_i^{hc}(t), \mathbf{T}_m^i) \forall i$ ; // transform to target
10       $\mathbf{T}_m^{hc}(t) \leftarrow \text{ransac\_avg}(\mathbf{T}_m^{hc}(t), \mathbf{T}_{m,i}^{hc}(t)) \forall i$ ; // ransac and average
11       $\mathbf{T}_m^i \leftarrow \text{init\_update}(\mathbf{T}_m^{hc}(t), \mathbf{T}_i^{hc}(t)) \forall i$ ; // update all the relative poses
12    else if not all marker detected then
13       $\mathbf{T}_{m,i}^{hc}(t), \mathbf{T}_{m,i}^i(t) \leftarrow \text{trafo2m}(\mathbf{T}_m^{hc}(t), \mathbf{T}_i^{hc}(t), \mathbf{T}_m^i) \forall i$ ; // transform to target
14       $\mathbf{T}_m^{hc}(t) \leftarrow \text{ransac\_avg}(\mathbf{T}_m^{hc}(t), \mathbf{T}_{m,i}^{hc}(t)) \forall i$ ; // ransac and average
15    else if no marker detected then
16       $\mathbf{T}_m^{hc}(t) \leftarrow \text{slam\_integrate}(\mathbf{T}_w^{hc}(t), \mathbf{T}_{hc}^w(t - \Delta t), \mathbf{T}_m^{hc}(t - \Delta t))$ ; // Eq. 2
17       $\mathbf{T}_m^{hc}(t + t_d) \leftarrow \text{delay\_integrate}(\mathbf{T}_w^{hc}(t + t_d), \mathbf{T}_{hc}^w(t), \mathbf{T}_m^{hc}(t))$ ; // Eq. 3
18  end
19 end

```

delays inherent in these systems must be corrected in order to create a real-time virtual display of the scene. Next, the proposed solution to these challenges are described.

Our Solution To tackle these problems, we propose a robust marker localization pipeline (depicted in Algorithm 1) as an extension to ArtoolKitPlus [22]. As an overview, the proposed pipeline utilizes multiple markers as well as the robots' SLAM system. To explain, multiple markers are placed on an object, where there exist a predefined target marker ID m and n additional markers with unique identifications, i.e., $i = 1, 2, \dots, n$. This results in total $k = n + 1$ markers. At initialization, the algorithm detects all the markers, where *multiART+* is the function that executes a variant of the marker tracking method: ArtoolKitPlus [22]). Using the image of the eye-in-hand camera I (either the left or right camera of the stereo setup), we obtain the initial 6D pose of the target marker \mathbf{T}_m^{hc} as well as all n additional markers \mathbf{T}_i^{hc} at $t = 0$. Then, we save the relative poses of all n markers to the target marker m (denoted \mathbf{T}_m^i for $i = 1, 2, \dots, n$). This step is executed within the function *marker_init*. Then, the 6D pose of the target marker \mathbf{T}_m^{hc} can be obtained in the main loop of Algorithm 1. The first step is to execute *multiART+*. Then, if all the k markers are detected, we transform the 6D pose of n additional markers to the target marker: $\mathbf{T}_m^{hc} = \mathbf{T}_i^{hc} \mathbf{T}_m^i$ (executed with a function *trafo2m*). As

this results in n additional 6D poses of the target marker m , we note them as $\mathbf{T}_{m,i}^{hc}$ for $i = 1, 2, \dots, n$. Then, RANSAC [70] is applied to these estimates to remove outliers, and then we perform averaging to reduce the variance (*ransac_avg*). Then, the relative transformations \mathbf{T}_m^i are updated. If at least one marker is detected, the same step is applied to estimate the target marker without updating the relative transformations \mathbf{T}_m^i . We also note that RANSAC is skipped when less than three points are available. The described steps have two advantages. First, the accuracy and orientation ambiguity of ArtoolKitPlus can be improved with RANSAC, and second, the algorithm is robust to loss-of-sight of the target marker, i.e., detecting only one of the markers is enough to still estimate the 6D pose of the target. Similar steps have been presented in the past with several variants [42], [43], [71].

However, the algorithm must be robust to loss-of-sight on all the markers, and further compensate for the time delay. This is achieved by extending the algorithm with SLAM estimates. The overview is shown in Fig. 6. As a first step, we propose to address the problem of complete loss-of-sight on all the markers by integrating SLAM estimates of camera motion with respect to its inertial coordinate, i.e., utilize the estimated transformation between the hc camera to a world coordinate of our SLAM system w : $\mathbf{T}_{hc}^w(t)$. If no markers are detected in the main loop of the algorithm, one can still

estimate the target marker $\mathbf{T}_m^{hc}(t)$ by integration (executed within the function *slam_integrate*):

$$\mathbf{T}_m^{hc}(t) = \mathbf{T}_w^{hc}(t)\mathbf{T}_{hc}^w(t-\Delta t)\mathbf{T}_m^{hc}(t-\Delta t). \quad (2)$$

In Eq. 2, $\mathbf{T}_w^{hc}(t)\mathbf{T}_{hc}^w(t-\Delta t)$ is a relative transformation of camera motion from time t-1 to t and we assume a static object. In a similar fashion, the time delay of the system t_d can be computed (executed with a function *delay_computation*) and corrected with SLAM algorithm by kinematics:

$$\mathbf{T}_m^{hc}(t+t_d) = \mathbf{T}_w^{hc}(t+t_d)\mathbf{T}_{hc}^w(t)\mathbf{T}_m^{hc}(t), \quad (3)$$

which is executed within a function *delay_integrate*. Note that the time delay is present in any perception system (e.g. rectifying an image), fiducial marker systems as well as the communication delays. In Eq. 3, $\mathbf{T}_w^{hc}(t)$ and $\mathbf{T}_i^{hc}(t)$ are computed using SLAM and multi-marker tracking. On the other hand, $\mathbf{T}_w^{hc}(t+t_d)$ can be computed using linear and angular velocity estimates of SLAM, multiplied by the delay time t_d . These two steps have several advantages. The algorithm is robust to the failure modes found in fiducial marker systems as it copes with missing marker detection, and time delays are incorporated using velocity signals and computed delay time. Furthermore, maximum run-time of the algorithm can be pushed upto 200Hz, which is the rate of visual-inertial SLAM estimates. The algorithm also deals with drifts of SLAM estimates by using relative motion estimates only when the marker detection is lost. Note that the proposed method is simple but can be an effective way of exploiting the commodity vision sensors with SLAM modules in order to improve the robustness of the fiducial marker systems.

B. The Proposed Pipeline for Objects of Unknown Geometry

Whenever we cannot assume the availability of the markers, the 6D pose of the objects can be estimated using depth sensors such as a LiDAR with the point cloud registration methods. For example, within the intended industrial application, the markers cannot be used for estimating the pose of the pipe. This is because in oil and gas refineries, the pipes are often very long while their inspection points are generally unknown a-priori. Concretely, given the incoming streams of point clouds $\mathbf{P}(t)$ and the point clouds of the object \mathbf{Q} from a CAD model, the goal is to find the rotation and translation between $\mathbf{P}(t)$ and \mathbf{Q} . Here, the point clouds $\mathbf{P}(t)$ contain the points $\mathbf{p}_i(t)\forall i$ with their coordinates lying at the weighted centroid \mathbf{c}_p . Similarly, the point clouds \mathbf{Q} contain the points $\mathbf{q}_i(t)\forall i$ with its coordinate system defined at the weighted centroid \mathbf{c}_q . Defining \mathbf{c}_q to be aligned with the LiDAR coordinate system l , the 6D pose of an object o can be obtained by matching the two point clouds: $\mathbf{p}_i = \mathbf{R}_o^l \mathbf{q}_i + \mathbf{t}_o^l$. This goal of finding \mathbf{R}_o^l and \mathbf{t}_o^l is often formulated as an optimization problem:

$$\mathbf{R}_o^l, \mathbf{t}_o^l = \arg \min_{\mathbf{R}_o^l, \mathbf{t}_o^l} \sum_i \rho(\|\mathbf{p}_i - \mathbf{R}_o^l \mathbf{q}_i - \mathbf{t}_o^l\|), \quad (4)$$

where ρ is again a distance based cost function, e.g. typically a mean squared error.

Commonly, the solution is obtained by first computing the rotation and then the translation. Centering all the points:

$$\bar{\mathbf{p}}_i = \mathbf{p}_i - \mathbf{c}_p \quad \text{and} \quad \bar{\mathbf{q}}_i = \mathbf{q}_i - \mathbf{c}_q \quad \text{such that} \quad (5)$$

$$\mathbf{R}_o^l = \arg \min_{\mathbf{R}_o^l} \sum_i \rho(\|\bar{\mathbf{p}}_i - \mathbf{R}_o^l \bar{\mathbf{q}}_i\|), \quad (6)$$

the goal is to find the rotation matrix that aligns the centered point clouds. Defining the correlation matrix as $\mathbf{C} = \sum_i \bar{\mathbf{p}}_i \bar{\mathbf{q}}_i^T$ and its singular value decomposition as $\mathbf{C} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$, the rotation can be estimated by the orthogonal Procrustes algorithm, while the translation can be obtained from the weighted centroids \mathbf{c}_p and \mathbf{c}_q after rotation:

$$\mathbf{R}_o^l = \mathbf{U}\mathbf{V}^T \quad \text{and then} \quad \mathbf{t}_o^l = \mathbf{c}_q - \mathbf{R}_o^l \mathbf{c}_p. \quad (7)$$

This assumes the correspondences between each points to be known. In practice, however, the correspondences are often not known and the Iterative Closest Point (ICP) algorithm is often used [24], [44], [45]. Intuitively, the ICP algorithm iterates the following steps: (1) finding the closest point in the transformed point cloud for each point: $\min \rho(\mathbf{P}, \mathbf{Q})$, (2) estimating the transformation using Eq. 7, and (3) applying the found transformation to all points and iterating all steps until a convergence criterion is reached. As ICP algorithm is subject to local minima, ICP is often initialized by employing the global registration methods [46] or using higher level features at the first step of the ICP algorithm.

Challenges Unfortunately, such a strategy does not fully address the current use-case of the point cloud registration methods. This is due to aerial manipulation tasks for an inspection and maintenance scenario. We outline the resulting failure modes in Fig. 5. First, the strategy assumes the availability of a precisely known object geometry \mathbf{Q} , which can be obtained through CAD models. Unfortunately, this assumption is invalid in our industrial scenario, as the CAD models of the objects that belong to external environments are unknown, e.g. CAD model of an oil or gas pipe. Although refineries may have a 3D geometry of the site, there exist erosion and other changes to their initial model. Second, a holistic view of the object cannot be assumed. In the setup, robotic arm and other objects can occlude the object of interest, resulting in partial and overlapping view of the point cloud. Lastly, as we deal with a floating base system where the base of a robotic arm is not fixed, view point challenges can occur. This challenges the out-of-box point cloud registration methods for LiDAR systems, which contains sparse point clouds.

Our Solution To this end, we propose a 6D object pose estimator using a LiDAR. The pipeline is depicted in Fig. 7. For an overview, the proposed algorithm constitutes an initialization step and a multi-process main loop. At initialization, the CAD model of the scene is reconstructed online by exploiting an object detector. In the main loop, three

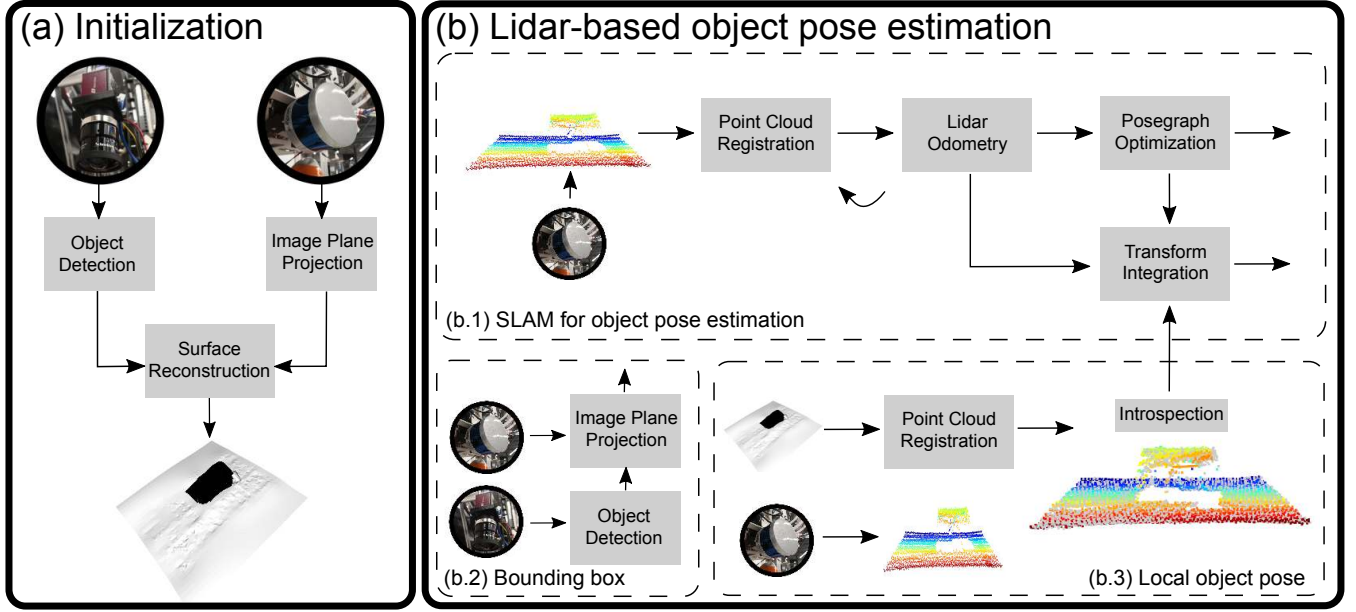


Figure 7: The proposed LiDAR pose estimation pipeline. (a) Initialization. Using an object detector, the point clouds that belong to the objects of interest can be obtained for a surface reconstruction technique. (b) Main loop. The object poses are estimated by combining a SLAM technique with local object poses from a point cloud registration method. An object detector computes the regions of occlusions and dynamic objects for masking out.

parallel processes are created: a bounding box estimator that computes the locations of the moving and occluding objects to be masked out, a SLAM pipeline that computes the poses of the objects in a global reference frame, and a local object pose estimator that estimates the object poses locally. The combination of local and global methods is to mitigate the challenges related to non-holistic view of the object. The SLAM estimates can deal with large perspective changes by matching the scans sequentially, but suffer from drift. The local method, whenever reliable, can be exploited to reduce the drift of the SLAM system. Lastly, what motivates the multi-process architecture is the efficiency, i.e., LiDAR odometry can be executed at a faster rate than the other processes that can be executed only at a slower rate.

First, the pipeline is initialized by creating the CAD model of the object online. This is because one cannot always assume a known geometry of the object in the targeted application scenario, i.e., the target point cloud \mathcal{Q} is not available, and consequently its CAD model \mathcal{O} for the VR. Yet, from the specification of the given task, e.g. pick and place an inspection robot on a pipe, what we know a-priori is the semantics of the objects of interest, e.g. a pipe. Therefore, one can still create the CAD model of the object \mathcal{O} online by finding the point clouds that belong to the objects of interest $\mathcal{P}_o(0) \in \mathcal{P}(0)$ and applying a surface reconstruction technique once. For this, we train a DNN based object detector [72] using the eye-to-hand camera (mako). Defining this DNN as a parametric function f_θ with its input as an image \mathcal{I} , the goal of a 2D object detector is to classify and locate objects in an image; for object semantics c , e.g.

$c \in \{\text{pipe, robotic arm, cage}\}$, the classification probability p_c (a score between 0 and 1), and location as a bounding box $\mathbf{b}_c \in \mathbb{R}^4$ in the given image, the 2D object detector returns the tuples:

$$\{c, p_c, \mathbf{b}_c\} = f_\theta(\mathcal{I}) \quad \text{with} \quad \mathbf{b}_c = [u_{c,1} \quad v_{c,1} \quad u_{c,2} \quad v_{c,2}]^T.$$

Here, the bounding box is described by two points in the image with heights ($h = u_{c,1}$ and $h = u_{c,2}$) and widths ($w = v_{c,1}$ and $w = v_{c,2}$) that are the upper left corner and the bottom right corner of the box that contains the object c . Defining the target object $c = o$ and using the extrinsic calibration parameter between the LiDAR and the eye-to-hand camera T_{mako}^l to transform all point clouds $\mathcal{P}(0)$ into the image plane, we can find the point clouds that belong to o :

$$\mathcal{P}_o(0) = \left\{ \mathbf{p}_j = (x_j, y_j, z_j) \right\} \quad \text{such that} \\ j = \left\{ i \mid u_{o,1} \leq x_{i,image} \leq u_{o,2}, v_{o,1} \leq y_{i,image} \leq v_{o,2} \forall i \right\}.$$

This means that all the LiDAR scans $\mathcal{P}_i(0)$ to the image plane are transformed, which results in $\mathbf{p}_{i,image} = (x_{i,image}, y_{i,image})$. Then, we obtain the indices j of the point clouds that lie inside the bounding box of our target object \mathbf{b}_o and crop the original point cloud $\mathcal{P}(0)$ to obtain the point clouds $\mathcal{P}_o(0)$ that only contain the information about our target object. Applying a surface reconstruction technique [73], the CAD model of the target object \mathcal{O} can be created. In this way, we can still create the VR of the scenes with objects of unknown geometry.

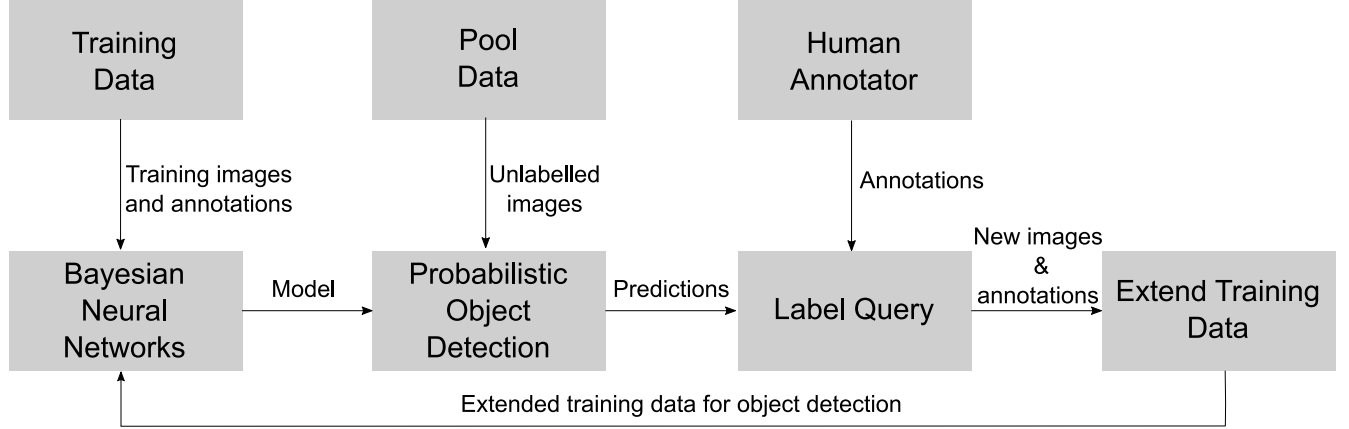


Figure 8: Active learning pipeline for generating labelled data more efficiently. Instead of randomly selecting the images to be labelled, we query the most informative samples from Bayesian Neural Networks - an uncertainty-aware Deep Neural Networks for the state-of-the-art object detection frameworks.

Next, the main loop of our algorithm is described, where the first process is the bounding box estimator. This process tackles the problem of occluding and moving objects $c = u$ by estimating the bounding box of other objects u in the LiDAR coordinate system, which is for actively removing the points that belong to the occluding and moving objects u . Concretely, similar to before, the object detector and the extrinsic calibration can be used to obtain:

$$P_u(t) = [p_u = (x_u, y_u, z_u)] \quad \text{such that}$$

$$u = \{i \mid u_{u,1} \leq x_{i,image} \leq u_{u,2}, v_{u,1} \leq y_{i,image} \leq v_{u,2} \forall i\},$$

where $P_u(t)$ is the point cloud that belongs to objects u at time t . Then, the bounding boxes of the objects can be computed in xy-plane of the LiDAR coordinate system. Defining this as b'_u , and examining all point clouds:

$$b'_u = [\min(x_u) \quad \min(y_u) \quad \max(x_u) \quad \max(y_u)]^T. \quad (8)$$

Note that the projection of all the point clouds to the image plane can be inefficient for embedded CPU computations. Therefore, a separate process is assigned. Also, moving averages and a margin are also implemented. Lastly, these bounding boxes are used to mask out occluding and moving objects in all other modules.

Then, in the main loop, a LiDAR based SLAM system [23] is employed to address the problem of view point changes. Again, a naive strategy is to perform point cloud registration between the reconstructed CAD model of the target object O and the incoming point cloud scans P . However, if the initially constructed object O is significantly different from the current point cloud P , the point cloud registration method may perform poorly due to the smaller overlaps between the two scans. Therefore, our key idea is that a LiDAR odometry pipeline that performs the registration between the consecutive point clouds and sets the coordinate of the constructed object

O as a global reference, do not suffer from significant view point changes. However, as this approach suffers from drift, i.e., accumulation of errors, two mechanisms are introduced. The first is a posegraph optimization:

$$\{T_i\} = \arg \min \lambda \sum_i \sum_{(P,Q) \in \mathcal{K}_i} \|T_i P - T_{i+1} Q\|^2$$

$$+ \sum_{i < j} \sum_{(P,Q) \in \mathcal{K}_{ij}} \rho(\|T_i P - T_j Q\|),$$

where λ determines the weight of a cost between two consecutive scans within the keyframes, and ρ is a robust function, e.g. set to L2 norm in our case. Here, a framework that performs robust pose graph optimization, is applied, which is less prone to the errors of pairwise registration [74]. Second, we propose to combine a local object pose that is obtained by registering the point cloud of incoming scans with the target object O . Whenever the confidence estimates of the local object poses are high (or above a specified threshold), we reset the SLAM system with initialization from the local object pose estimator. In this way, we account for the drift of the SLAM system.

C. The Proposed Active Learning Framework

So far, the proposed object pose estimators are described for realizing the VR based telemanipulation system. Here, our pipeline relied on a DNN based object detector. This has been used for the online creation of a CAD model, and to rule out any occlusions and moving objects. As our entire system relies on a DNN for the VR creation, we next propose an active learning framework to obtain the required performance in DNNs within the context of field robotics.

Challenges The problem is on the training and the deployment of DNN based systems for various environments including both the indoor and the outdoor conditions. The challenge lies in realizing a DNN based system for long-

term operations in outdoor environments. This is due to the necessity and manual preparations of large amounts of high-quality, annotated data, which can cover the variety of operation conditions. For example, as illustrated in Fig. 5, a DNN trained from annotated images of laboratory environments may not be generalized to outdoor environments. Similarly, seasonal variations of the scene from summer to winter can cause similar effects in the deterioration of generalization performance. Therefore, each change in the scene may require an iterative process of collecting and annotating the data. As this can be a long and tedious process, we attempt to find a principled solution that guides the process of gathering the required data for the field deployments of the DNN based systems.

Our Solution To this end, we now describe the proposed pool based active learning framework [76] (also known as "experiment design"). To explain, active learning is a class of machine learning paradigm, where labelled data is not available for a supervised learning problem. Instead of obtaining annotations for all available unlabelled data, active learning attempts to only label fewer but the most informative data. Intuitively, the aim of active learning is to create a learning system that chooses by itself what data it would like the user to label. As opposed to the heuristic choice of the user, active learning enables a DNN to select small amounts of data, guiding the user in the data creation process. In practical applications of robotics, this indicates that first, a pool of unlabelled data needs to be collected by the robot. As we use visual learning methods, a camera set-up could replace the deployment of the robot for data collection. We note that upto this stage, the procedure is similar to a standard supervised learning settings in robotics. Then, instead of annotating all the available data manually, fewer images are then autonomously selected by the active learning algorithm. Deep learning models are trained from these fewer images, and finally deployed to the robot.

Fig. 8 illustrates the overall idea behind active learning. In a pool based approach, a model is trained on an initial training set, which is often small. Then, the model selects a subset of data points from a pool of unlabeled data, and asks a human to label the selected data points. The selection involves a decision making process, which is performed through the choice of an acquisition function. Based on the updated training set, a new model is trained. Repeating this process, we can reduce the amount of labeling required to train a learning system. We present such a system for DNNs, which relies on an uncertainty quantification technique for DNNs. These are namely Bayesian Neural Networks (BNNs) and probabilistic object detection. To explain, our algorithm 2 depicts the working principle overall. Using an initial training set \mathcal{D}_{init} , we train a BNN which is denoted as $p(\theta|\mathcal{D}_{init})$. Then, for a user specified number of query steps Q , we first select the most informative, top K samples from the pool of unlabelled images: \mathcal{D}_{pool} . This is achieved by estimating the uncertainty from BNNs (denoted $p(y_i^* | x_i^*, \mathcal{D}_{train})$). Then,

we label the selected images, and the BNN is updated with the new training set. For more explanation in detail, we next present each of these components namely the BNNs, the uncertainty of BNN based object detector, and query step through the acquisition function.

1) Bayesian Neural Networks for Uncertainty Quantification

One of the crucial components of the proposed active learning framework is BNNs. Intuitively, BNNs are Bayesian reasoning applied to DNNs which allows for the uncertainty quantification in DNN predictions. We note that our previous works on BNNs [28], [75] are being extended to active learning framework for object detection. While we refer to [28], [75] for in-depth treatment, our description below aims to provide the basic formulation within the context of its application to active learning.

For this, consider a supervised learning set-up with input-output pairs $\mathcal{D} = \{\mathcal{X}, \mathcal{Y}\} = \{(x_i, y_i)\}_{i=1}^N$, where $x_i \in \mathbb{R}^D$, $y_i \in \mathbb{R}^K$. Similar to previous sections, we define a DNN as a parametrized function $f_\theta : \mathbb{R}^D \rightarrow \mathbb{R}^K$, where $\theta \in \mathbb{R}^P$ is a vectorized form of all DNN weights or parameters, e.g. all the weights of convolution kernel or the weights and biases of a multi-layer perceptron. In a standard DNN, we typically aim to minimize the loss function: $\min_{\theta} \frac{1}{|\mathcal{D}|} \sum_{(x,y) \in \mathcal{D}} \mathcal{L}(f_\theta(x), y) + \frac{\delta}{2} \theta^T \theta$ where δ is an L_2 -regularizer, and $\mathcal{B} \subset \mathcal{D}$ denote mini-batches. The resulting solution is a single hypothesis of a local maximum-a-posteriori (MAP) solution $\hat{\theta}$. To the contrary, BNNs explicitly express DNNs as probability distributions over DNN model parameters θ given the data $p(\theta|x, y)$, which is also known as the posterior distribution over the DNN model parameters:

$$p(\theta|x, y) = \frac{p(y|x, \theta)p(\theta)}{p(y|\theta)} = \frac{p(y|x, \theta)p(\theta)}{\int p(y|x, \theta)p(\theta)d\theta}. \quad (9)$$

As a direct application of Bayes theorem, where a prior distribution over the model parameters $p(\theta)$ is specified, along with the likelihood $p(y|x, \theta)$ and the model evidence $p(y|\theta)$. Once the posterior distribution over the weights is obtained, the prediction of an output for a new input datum x^* can be obtained by marginalizing the likelihood $p(y|x, \theta)$ with the posterior distribution. This step is called Bayesian Model Averaging, which can be used for active learning:

$$p(y^*|x^*, \mathcal{D}) = \int p(y^*|x^*, \theta)p(\theta|\mathcal{D})d\theta. \quad (10)$$

This indicates that the uncertainty estimates for a DNN prediction y^* can be obtained through combining different hypotheses of model parameters, resulting in the predictive distribution $p(y^*|x^*, \mathcal{D})$. Another implication of the formulation is the reliance on posterior probabilities $p(\theta|\mathcal{D})$ for uncertainty quantification.

Unfortunately, estimating the weight posterior is a challenging task, and has been one of the central topic in research of BNNs [53]. While the reasons are multitude, one of the primary reasons is the lack of a closed form solution due to the

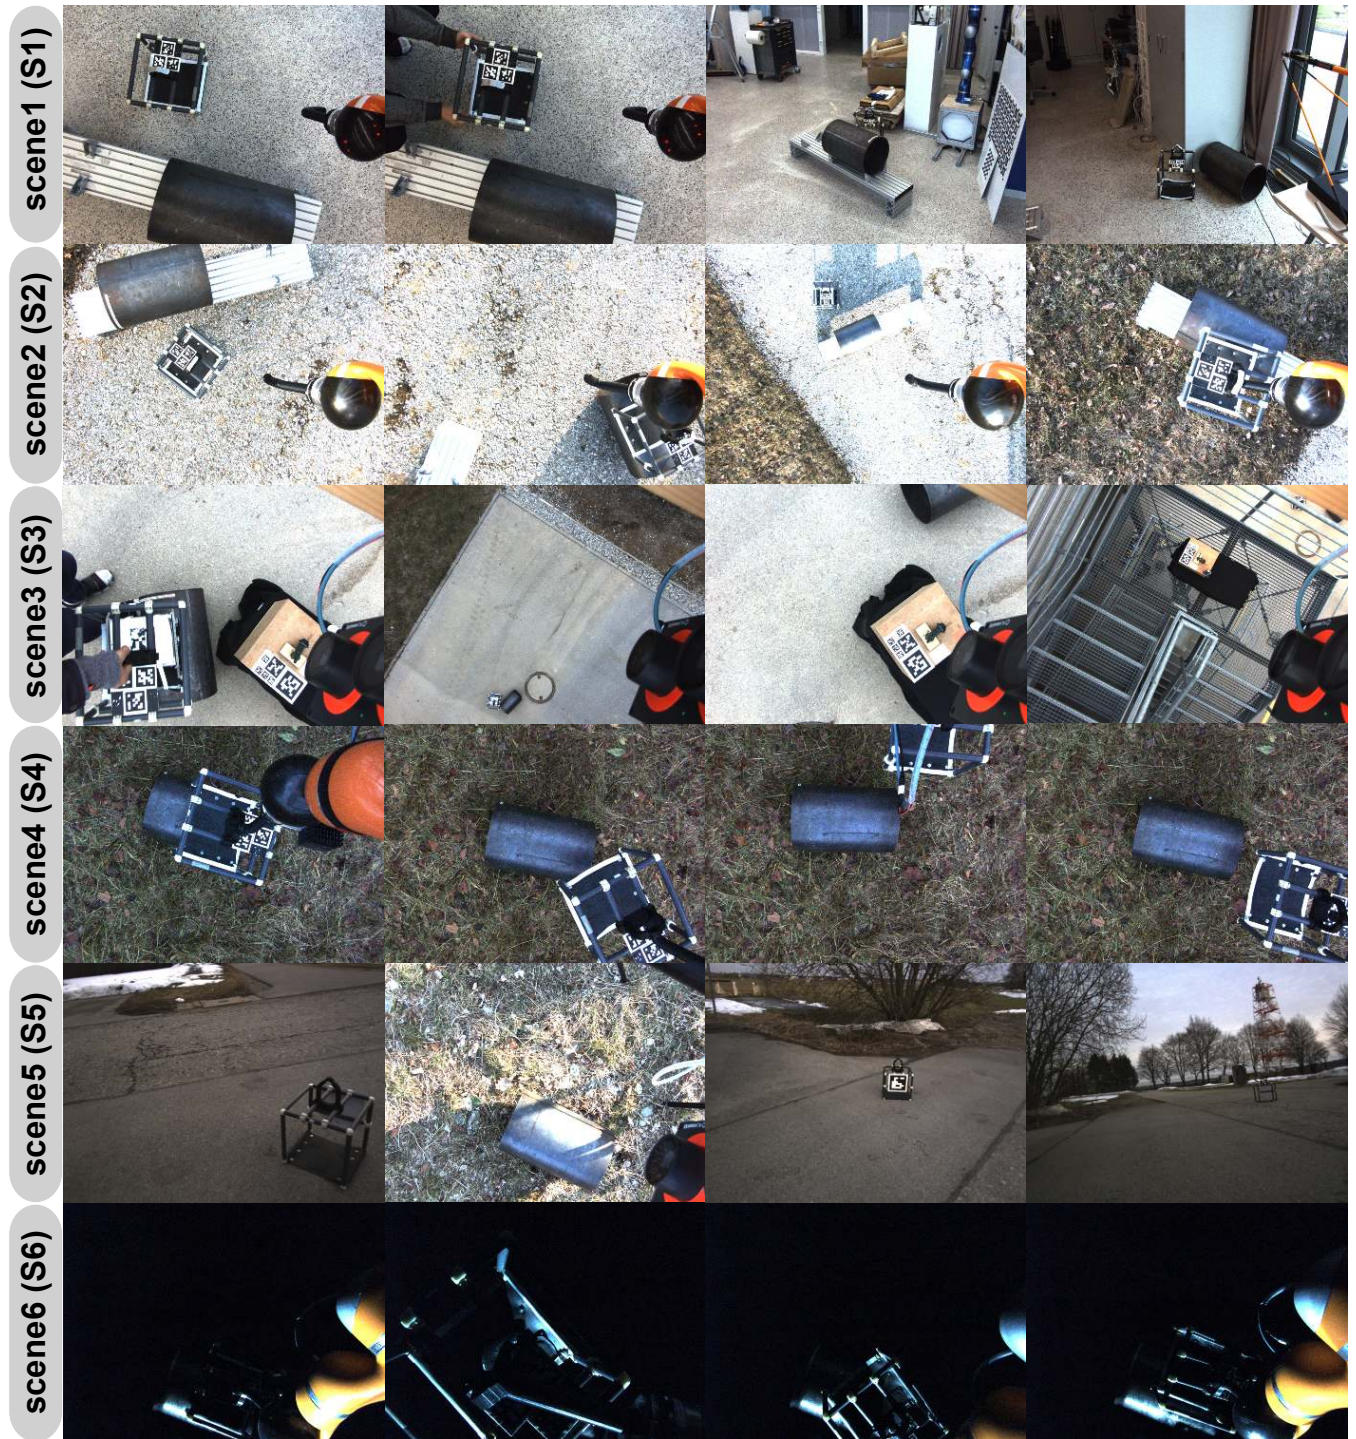


Figure 9: Collected sets of pool data from different scenes. Six different scenes cover indoor environments, varying backgrounds and height, and the scenes with snow and at night.

Algorithm 2: Deep Active Learning using Bayesian Neural Networks

```

     $\mathcal{D}_{init}$    The initial annotated training data.
input :  $\mathcal{D}_{pool}$  The unlabelled data.
         $Q$       The number of query steps.
         $K$       The size of the query per step.
output:  $f_{\theta}$    The trained DNN based object detector.
         $\mathcal{D}_{train}$  The annotated training data.

1 begin
2   /* Initialization */
3    $p(\theta|\mathcal{D}_{init}) \leftarrow \text{create\_BNN}(\mathcal{D}_{init})$ ; // Apply [28], [75]
4    $\mathcal{D}_{train} \leftarrow \text{update\_data}(\mathcal{D}_{init})$ ; // Initialize the training set from  $\mathcal{D}_{init}$ 
5   /* Main Loop */
6   for all the number of query steps  $Q$  do
7      $p(\mathbf{y}_i^* | \mathbf{x}_i^*, \mathcal{D}_{train}) \leftarrow \text{prob\_detector}(\mathcal{D}_{pool}) \forall i$ ; // Evaluate uncertainty on a pool set (Eq. 11 and 10)
8      $\mathcal{D}_{selected} \leftarrow \text{query}(\mathcal{D}_{pool}, K)$ ; // Query from the pool set (Eq. 12)
9      $\mathcal{D}_{new} \leftarrow \text{generate\_annotations}(\mathcal{D}_{selected})$ ; // The user or human supervisor annotates the images
10     $\mathcal{D}_{train} \leftarrow \text{update\_data}(\mathcal{D}_{new})$ ; // Update the training set by adding new annotated data
11     $p(\theta|\mathcal{D}_{train}) \leftarrow \text{create\_BNN}(\mathcal{D}_{train})$ ; // Apply [28], [75]
12  end
13 end

```

nonlinearities of DNNs that prohibit the validity of conjugate priors [77]. As a result, the use of approximation techniques of Bayesian inference such as variational inference or Monte-Carlo Monte-Carlo (MCMC) sampling have been researched with a focus on dealing with the high dimensionality of DNN weight space and the scalability with respect to large amounts of data that DNNs typically assume. For the computations of the weight posterior, the proposed pipeline relies on Laplace Approximation [28], [75]. These works are well suited for active learning in robotics, due to the demonstrated scalability to large architectures and dataset [28]. The automation of the hyperparameter tuning via Bayesian Optimization [75] can also be exploited in every query steps of active learning.

2) Uncertainty Estimation for Object Detectors

Having obtained the posterior probabilities of BNNs, the uncertainty estimates can now be computed for the underlying object detector. A key challenge is the adaption of BNNs for the object detector that may rely on several post-processing steps [78]. As we use an anchor based detectors such as Retinanet [25] (as these types of object detectors can provide real-time performance on the Jetson TX2 as oppose to regional-proposal approaches or end-to-end pipelines), one needs to deal with miss-correspondence between the anchor predictions and final outputs, and (ii) hard cut-off behavior in non-maximum suppression (NMS) step [25].

For these, the BayesOD framework [78] is employed, which infers the output distributions from the BNNs predictive distributions. In BayesOD, the samples of the BNNs predictive distributions are clustered in anchor level in order to derive the uncertainty estimates of the object detection. For this,

one can assume that the clusters contains M number of anchors. We further assume the highest classification score as the center of this cluster (indexed by 1) and other anchors of the cluster are considered as measurements to provide information for the center (denoted as $\hat{\mathbf{c}}_i$ and $\hat{\mathbf{b}}_i$). Then, the uncertainty estimates for classification $p_{[\hat{\mathbf{c}}_1, \dots, \hat{\mathbf{c}}_M]}(\mathbf{c}|\mathbf{x}^*, \mathcal{D}_{train})$ and regression $p_{[\hat{\mathbf{b}}_1, \dots, \hat{\mathbf{b}}_M]}(\mathbf{b}|\mathbf{x}^*, \mathcal{D}_{train})$ are:

$$\begin{aligned}
 p_{[\hat{\mathbf{c}}_1, \dots, \hat{\mathbf{c}}_M]}(\mathbf{c}|\mathbf{x}^*, \mathcal{D}_{train}) &\propto p_{\hat{\mathbf{c}}_1}(\mathbf{c}|\mathbf{x}^*, \mathcal{D}_{train}) \prod_{i=2}^m p(\hat{\mathbf{c}}_i|\mathbf{c}, \mathbf{x}^*, \mathcal{D}_{train}), \\
 p_{[\hat{\mathbf{b}}_1, \dots, \hat{\mathbf{b}}_M]}(\mathbf{b}|\mathbf{x}^*, \mathcal{D}_{train}) &\propto p_{\hat{\mathbf{b}}_1}(\mathbf{b}|\mathbf{x}^*, \mathcal{D}_{train}) \prod_{i=2}^m p(\hat{\mathbf{b}}_i|\mathbf{b}, \mathbf{x}^*, \mathcal{D}_{train}),
 \end{aligned} \tag{11}$$

where $p_{\hat{\mathbf{c}}_1}(\mathbf{c}|\mathbf{x}^*, \mathcal{D}_{train})$ indicates the per-anchor predictive distribution of the cluster center and $\prod_{i=2}^m p(\hat{\mathbf{b}}_i|\mathbf{b}, \mathbf{x}^*, \mathcal{D}_{train})$ is the likelihood term. We refer to more details on the BayesOD in [78].

3) Acquisition Functions for Query Generation

Another crucial component of active learning is the acquisition function, which relies on the uncertainty estimates from the BNN based object detector, in order to rank the images in the pool set. In other words, the defined acquisition function uses the uncertainty estimates to evaluate how informative each images in the pool set are. In case of an object detector, as there could be several object instances in an image, the information scores for each detected instances within an image are aggregated into one final score. Once such scores are obtained for all images in the pool set, the top K images can be queried for human annotation, which is then stacked




Figure 10: A summary statistics. Root Mean Squared Error (RMSE) and run-time are reported for the baseline methods and the proposed extension to the ARToolKitPlus. Lower the better for both the measures.

into the training set. The model is then retrained with the new and larger training set, and the process is repeated. As the acquisition function is a mechanism for the selection of active learning, its design can influence the performance of the learning framework.

Within a BNN based object detector, the uncertainty estimates can be obtained for both the classification and the bounding box regression [79]. Hence, one of the design choices are on how to effectively combine the two different types of uncertainty measures - one on semantic uncertainty, and the other on spatial uncertainty [79]. Defining the combination function $comb(\cdot)$ as a weighted sum or max operation [54], and also the aggregation function $agg(\cdot)$ as a sum or average operation [80], it is chosen;

$$\mathcal{A}(\mathbf{x}_k) = agg_{j \in N_k} (comb(\mathcal{U}_{j,cls}, \mathcal{U}_{j,reg})), \quad (12)$$

where $\mathcal{U}_{j,cls}$ and $\mathcal{U}_{j,reg}$ are the information score of the j -th detection instance on an image, for the classification and regression tasks, respectively. A mechanism of this acquisition function is to first combine both the semantic and spatial uncertainty by either a weighted sum or maximum operation, and then to sum or maximize the combined score per detection instance. What motivates the given choice is the handling of the problem itself. The combination operation relates to the need to combine the two different tasks per instance of an object detector, and the aggregation operation relates to the multiple instances in a single image [81]. What remains is then the design of the information scores for both classification and regression tasks: $\mathcal{U}_{j,cls}$ and $\mathcal{U}_{j,reg}$ respectively. Then,

$$\mathcal{U}_{j,cls} = \sum_{i=1}^{|\mathcal{C}|} \mathcal{H}(p(c_i|\mathbf{x}^*, \mathcal{D}_{train})) \quad \& \quad \mathcal{U}_{j,reg} = \mathcal{H}(p(\mathbf{b}|\mathbf{x}^*, \mathcal{D}_{train})), \quad (13)$$

which rely on the Shannon Entropy measure $\mathcal{H}(\cdot)$ - an indicator of how uncertain a distribution is. In case of classification, we assume categorical distributions over the classes c_i , while we assume multivariate Gaussian distributions for the bounding box regression \mathbf{b} . Importantly, what motivates the

optimization of the given entropy measure is its equivalence to maximize the information gain of a model [52].

V. Experiments and Evaluations

In this work, a VR based telemanipulation system is proposed, which is to provide real-time 3D displays of the robots' workspace and also haptic guidance to a human operator. The main contribution is the realization of such a system using robotic perception and active learning methods. This section therefore evaluates the proposed pipelines by examining how the created VR can match the real remote scenes, and if the identified challenges (in Fig. 4) are addressed by the proposed pipelines. Then, the results from the field experiments are presented, in order to characterize the effectiveness of the overall system in advancing aerial manipulation for real-world applications.

A. Ablation Studies and Evaluations

Several ablation studies are provided for the insights behind the presented algorithms. In particular, empirical evaluations of the algorithms devised, when faced with the challenges described in Fig. 4, are the objective.

1) Visual Object Pose Estimation for Known Objects

In summary, marker tracking algorithms can be exploited to create VR with known sets of objects. Here, the identified challenges are the shadows, the loss-of-sight or the partial views of the markers, which can cause the mismatch between the real remote scene and the VR. To address these challenges, the SLAM estimates of commodity visual-inertial sensors have been integrated, and here, validation of the devised algorithm is performed. To this end, the accuracy, the run-time and the robustness of the proposed algorithm are examined.

Experiment Setup For this, the ground truth of the relative poses between the objects and the camera are measured using a Vicon tracking system. Then, the algorithms are evaluated on the sequences that simulate the peg-in-hole insertion tasks. The Vicon measurements represent the ground truth of the object poses for indoor environments. To evaluate the effectiveness of the proposed algorithm against the identified challenges the observed failure modes of the

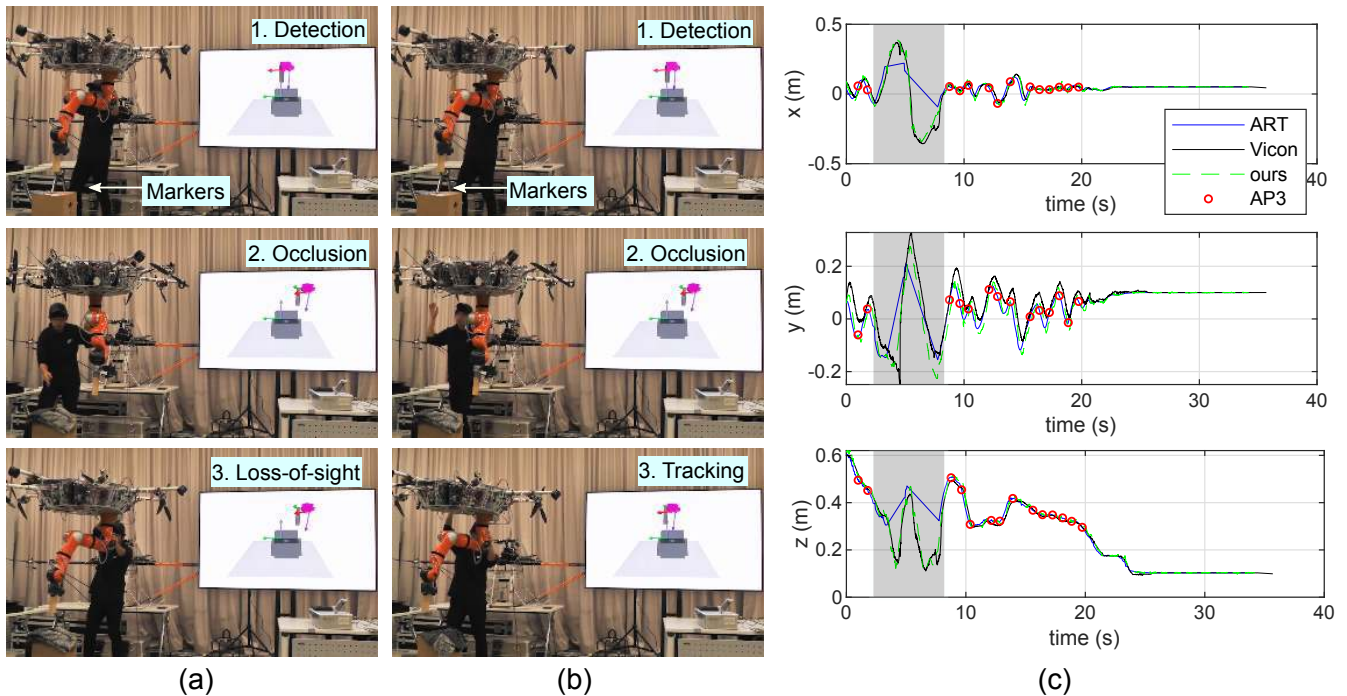


Figure 11: Qualitative results of the proposed marker tracking algorithm. (a) The existing marker tracking algorithms under the loss-of-sight of the markers are evaluated. (b) The proposed extension of the marker tracking algorithm with SLAM estimates is evaluated under the same scenario of the loss-of-sight. (c) The estimated positions from the baselines and the proposed algorithm are compared. These results indicate that the proposed algorithm can cope with the loss-of-sight of the markers and the time-delays, thereby justifying the design choices of the algorithm. Three markers of size 2.5cm, a marker of 6.25cm and a marker of 10cm are used in this evaluation scenario.

existing marker tracking systems are manually created. The baselines are Apriltag3 [41] (denoted AP3) and ARTToolKit-Plus [22] (denoted ART), which represent plug-in and play alternatives. Particularly, as the proposed algorithm extends the ARTToolKitPlus with SLAM estimates, this choice of the baseline enables a direct comparison. Five repetitions of these experiments are conducted in total.

Results The quantitative and qualitative results are reported in Fig. 10 and Fig. 11 respectively. In Fig. 11, the estimated trajectories of the relative poses are compared with the Vicon measurements. As depicted, the proposed algorithm is robust against the lost-of-sight problems of object localization with a hand-eye camera. On the other hand, the alternatives, namely AP3 and ART produce jumps as no markers are detected (between $t=2$ to $t=8$ as an example) when the camera loses the sight of the markers. This is due to the design of the algorithm where the SLAM estimates of the camera pose are integrated out, whenever the markers are not detected. Furthermore, ART suffers from a time delay, while AP3 has both the time delay, and the slow run-time. Moreover, the proposed algorithm can compensate the time delay, resulting in slightly more accurate estimates. The corresponding Root Mean Squared Errors (RMSE) is reported in Fig. 10 along with the run-time. We observe that AP3 is slow when using high-resolution images, and this results in more errors as the trajectories

are compared. In the approach, these trajectories are relevant for creating the VR with object localization methods. In this experiment, the proposed method yields the least errors and fastest run-time. We attribute the former to the robustness against the loss-of-sight of the markers, while the latter is due to the integration of the inertial sensors. These analysis of the accuracy, the robustness and the run-time validates the proposed algorithm. Moreover, the success of the algorithm is visually demonstrated in the video attachment, in addition to Fig. 11 (a-b).

2) LiDAR Object Pose Estimation for Objects of Unknown Geometry

For external scenes without the availability of precise geometry, a LiDAR based object pose estimation pipeline has been proposed. The features of the pipeline are to deal with the challenges outlined in Fig. 4. Thus, the aim now is to validate the components of the pipeline using the visual-inertial-LiDAR data-sets collected.

Experiment Setup For this, the point clouds and the visual data are collected in various situations. Within the controlled lab environment the following scenarios are created: a "nominal" scenario where the sensors ideally are pointed to an object statically, a "shaking" scenario in which imperfect




Figure 12: Statistical analysis of the baselines and the proposed method per scenario. Lower the better for the RMSE and run-time, and higher the better for the number of valid point cloud matches. These results indicate that the proposed algorithm can cope with the identified challenges, thereby justifying the design choices of the algorithm.

hovering of the robot creates sensor movements, a "rotation" scenario where the robot rotates around the object, and a "occlusion" scenario in which the robot arm and other objects moves to occlude the target object. These scenarios represent the identified challenges during a manipulation task (e.g. see Fig. 4). To further evaluate the proposed pipeline in a realistic use-case, additional scenarios are considered in outdoor environments. These are: a "night" scenario where the sensor data were collected during a manipulation task at night, a "tower 1" and "tower 2" scenario where the data are acquired similarly at two different locations. Importantly, extensive evaluations given under varying conditions are motivated by the industrial scenarios considered where this paper aims to build a working system that goes beyond the proof-of-concept prototypes.

For the baselines, the off-the-shelf methods such as point-to-point ICP [24], [82], point-to-plane ICP [44], [45] and the combination with the global registration methods [46] are compared. The pairwise registration is denoted pICP (with coarse-to-fine matching strategy) whereas pFPFH denotes

the global registration method. These are to examine the vanilla object pose estimators without specific measures to address the identified challenges. For brevity, only the best performing ones between the point-to-point and the point-to-plane methods are reported. Furthermore, we compare our method without different components to evaluate the contributions of each module to the final performance. These are the pure odometry (odom), odometry with posegraph backend (backend), combination of odometry and local object poses (comb) and the proposed object pose estimator (all). For better insights, we evaluate these methods with masked out dynamic parts of the scene, while existing works motivate the importance of masking out the dynamic parts of the scene in SLAM context. With these baselines, closely following [44], [82], the RMSE of the translation (RMSE t) and the rotation (RMSE R), the number of valid matches, and the run-time of each algorithms are measured.

Results The quantitative and qualitative results are reported in Fig. 12 and Fig. 13 respectively. In these experiments, the proposed object pose estimator yielded the least RMSE for

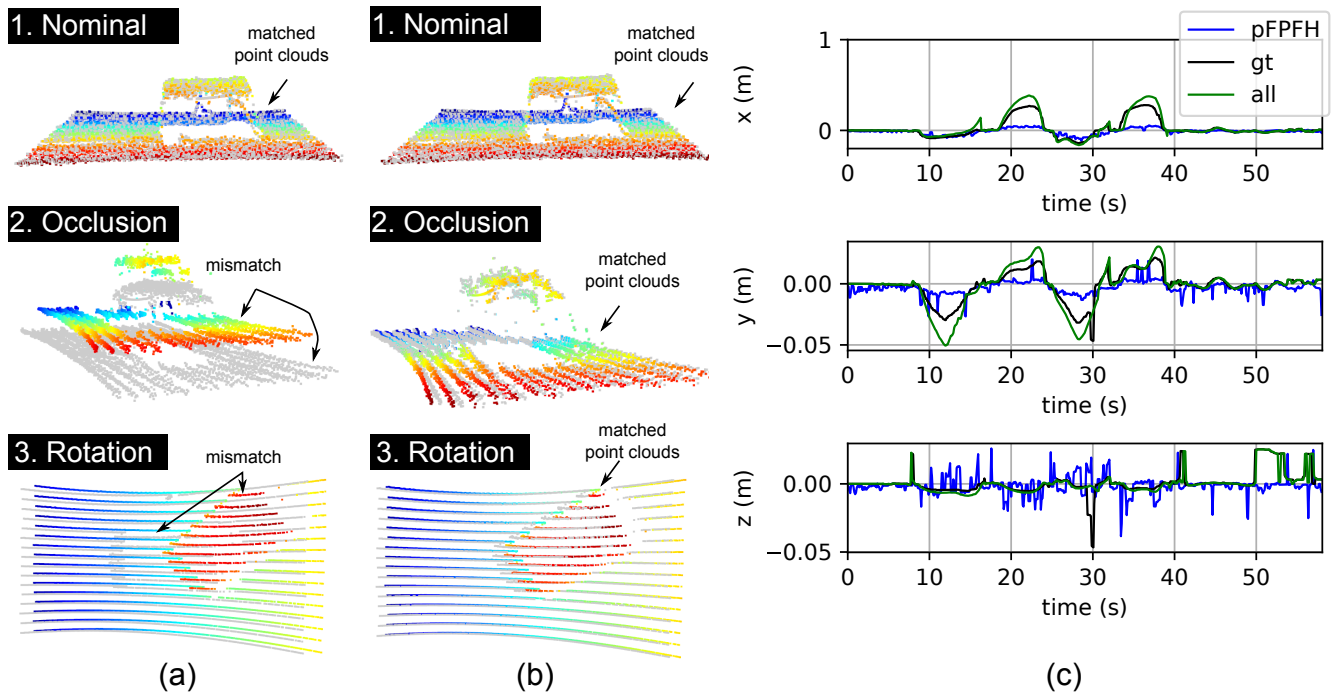


Figure 13: Qualitative results for the LiDAR based object pose estimators. (a) Evaluation of a vanilla object pose estimator. (b) Evaluation of the proposed object pose estimator. (c) Plots of the estimated positions from a baseline (pFOFH) and our approach (all) against the ground truth (gt) measurements. The qualitative results over occlusion and rotation scenarios justify the design choices of the proposed method. The monotonic gray indicates the source point cloud while the colored point cloud refers to the current scan. All point clouds are cropped for the visualization purpose.

both translation and rotation. Odometry with and without the back-end suffers from drift over time, while the vanilla methods such as pICP and pFPFH perform poorly when the view point changes are significant. The latter is qualitatively shown in Fig. 13, while the number of matches in Fig. 12 also indicates their relatively poor registration between the target and the source point clouds. supports the claims of this work on the identified challenges. Moreover, as shown in Fig. 12, it can be seen that all the components introduced namely pose-graph, local object poses and odometry, contributes to the accuracy of the estimates. With respect to the run-time, the odometry-based methods are capable of real-time, which we attribute to the significant overlap between two consecutive LiDAR scans that helps the ICP algorithm to converge faster. On the other hand, pFPFH is the slowest in terms of run-time because it relies on several components such as feature extraction, correspondence matching, and refinement through ICP. All of these results are consistently observed in seven scenarios with varying degrees of severity. these experiments justify and validate the proposed methods and its design choices. Importantly, the key takeaway is the effectiveness of combining object pose estimators with SLAM methods for a floating-base system, which can handle the failure cases of conventional object pose estimators via introspection.

3) Bayesian Active Learning for Field Robotics

Lastly, the proposed Bayesian active learning framework is evaluated within the context of field robotics. To recap, the main challenge is to deal with large variations in the environment, which can affect the performance of an object detector that has never seen such data in the training set. The natural question to evaluate here is the amount of labeling efforts that the active learning framework can save. As an application of the Bayesian active learning paradigm for field robotics, we focus on the impact of the system performance rather than the algorithmic advances.

Experiment Setup To this end, the visual data in various locations and conditions have been gathered. These are not only the (i) laboratory environments, but also (ii) the outdoor environments in different locations. These environments are indicated as scenes 1-6 or S1-6, which are visualized in Fig. 8. To evaluate the performance of the system, the manual annotations within these images have been created. The objects are a cage, a pipe, and a robotic arm. In total of about 20k images, we randomly labeled 5k images. The splits are performed at the ratio of 7.5:2:0.5, respectively, to a pool, test, and validation set. This is to simulate the real-world scenario where the training data is initially limited (e.g. the data collected in summer and having to test in the winter). We add uniform sampling from pool data (denoted random) and MC dropout [56] as the baseline. Although deep




Figure 14: Results of active learning (AL) compared to the training set-up that uses 100% of all data over different scenes. AL is used for the acquisition sequence of 5%, 10%, 25%, 50% and 75% of all data. Higher the better for mAP. The results show that one can save upto 75% of data, in order to reach more than 95% of the total performance.

ensembles [59] are another popular baseline, the suitability to active learning is limited due to the excessive training time. Here, the sampling strategy chosen from [81], and therefore, the only difference between the baseline methods are the uncertainty estimates.

Implementation details are as follows. the Pytorch implementations are used, namely the Retinanet implementation from Detectron2 [25] and the official implementation of BayesOD [78] with slight modifications for better performance. These modifications include the use of Bayesian inference only for the bounding box regression, instead of also applying to the classification head. The learning rate has also been tuned to obtain better convergence. The monte-carlo samples of 30 are used for computing the uncertainty estimates, and the rank of 100 and 50 BO iterations are used. The latter is applicable to only the Laplace Approximation, which was applied to all the layers in the Retinanet. On the other hand, only the existing dropout layers within Retinanet have been used for MC dropout. Such implementations are motivated by the promise of each methods. MC dropout assumes dropout layer, and have been popular in practice as one could make use of existing dropout layers, while Laplace Approximation can directly render every layer as Bayesian, given an already trained parameters.

Results Firstly, it is evaluated, how much data annotations one can save by comparing the training set-ups that uses 100% of the annotated training data against the acquisition sequences of 5%, 10%, 25% 50% and 75% of the total data. Repeated sequentially over each scenes, the performance of the resulting object detector with mAP as a metric, are measured. The test set contains samples from each scenes and therefore, this repetition shows how the performance gap due to scene variations are being closed. The results are depicted in Fig. 14. We observe that the gap between the active learner (AL) and Retinanet with 100% of annotated data (denoted 100%), reduces as we increase the acquisition size from 5% to 75% of the total data. In particular, in these scenarios,

AL with only 25% of the total annotated data can reach more than 95% of the 100%s mAP values, which results in saving upto 75% of the annotated data. This results are due to the redundancy in the data. We believe that this result can motivate AL for field robotic applications, where data preparation can be inherently more expensive than laboratory settings.

Next, the design choices of the proposed active learning pipeline are examined by comparing the method with the selected baselines. The results are depicted in Fig. 15. Here, the transfers between the scenes are assumed. As the robots may operate at different environments, we attempt to evaluate by starting the active learning with a neural network in an indoor scene and acquiring the data over different outdoor scenes. For all results, we acquire 5% of the data per iteration and repeat 10 iterations to reach 50% of all data. In total, three random seeds are used to compute the standard deviation (also in Fig. 14) for statistical significance. Examining the mean mAP over all the iterations, the data suggest that the performance increases over using the MC dropout. The results are consistently observed in several settings with different magnitude of the improvements. We attribute the reason to post-hoc nature of our Laplace Approximation based approach. In detail, the methods based on variational inference, such as MC dropout, may be at a disadvantage in active learning settings, where each loop involves training a DNN. Naturally, since variational inference rather learns uncertainty during training, finding hyperparameters that deliver good performance over many loops is difficult. On the other hand, post-hoc methods such as ours, the uncertainty estimates are obtained after training the DNN. This decoupling enables us to extensively search for hyperparameters, which is feasible within each active learning loop. Therefore, we interpret these results to show the validity of the design choices of the proposed active learning pipeline. In summary, the key take ways are the redundancy of the data when training a neural network




Figure 15: Comparisons of the proposed pipeline (INF) with other baselines such as random sampling (random) and MC dropout (MCD) over six different scenes. The mean mAP during the active learning process is displayed along with the labels of the curve. The standard deviation is shown in shade. Higher the mean mAP, the better.

in dynamic and unstructured environments, and an active learning framework with well-calibrated uncertainty estimates can produce a practical and positive impact by guiding the data preparation steps towards efficiency.

B. Field Testing and User Validation

While the previous focus was on the validation of the methods for VR creation, the flight experiments with SAM is now presented. The main purpose is to evaluate the benefits of the proposed system in advancing aerial manipulation capabilities. To this end, we examine two industrial scenarios in outdoor environments. Then, the robustness of the proposed system is examined by varying environments and users.

Experiment Setup The design of our experiment setup is to account for real-world applications of aerial manipulators. As a first step, the description of two industrial scenarios with SAM are provided, which involve the following aerial manipulation tasks in dynamic and unstructured environments:

- **Industrial task 1: Peg-in-hole Insertion** As one of the benchmarks for manipulation, this task involves inserting an object (attached to the end-effector) into a hole. An example is shown in Fig. 16. Industrial tasks such as valve opening and closing in high altitude areas, or in-air assembly of structures require the execution of this task. In this work, the peg-in-hole task with an error margin of less than 2.5mm is considered. This is a challenging task for aerial manipulation, since the robotic arm is on a floating base.
- **Industrial task 2: Pick-and-place and Force Exertion** Two other benchmarks for manipulation are combined, which are pick-and-place and force exertion, in the second task. In particular, our scenario, designed under

the scope of the EU project AEROARMS, involves deployment and retrieval of a robotic inspection crawler. An example is depicted in Fig. 17. It requires grasping of a cage (as a carrier of the crawler robot), placing the cage on a pipe, and pressing the cage onto the pipe while the crawler moves in and out of it for pipe inspection.

Note that, for the execution of these two tasks, the operator is located far away from the robot without direct visual contact to the workspace of the robot. More concretely, as shown in Fig. 16 and Fig. 17, the robot operates in an outdoor environment, while its operator remotely commands the robotic arm from a ground station. This simulates a real-world application scenario of a teleoperated aerial robot.

To evaluate the feasibility and benefits of the proposed telemanipulation system in advancing aerial manipulation capabilities of SAM, four sets of experiments are considered:

- **Set 1: Repetitions of Peg-in-hole Insertion** Several repetitions of the peg-in-hole insertion task are performed (as shown in Fig. 16). Here, we vary the conditions by executing the manipulation tasks with three modes, namely (a) VR and haptic guidance mode (denoted VR+HG), (b) VR mode (denoted by VR+Tele), and (c) only with live camera streams (denoted CAM). Eight repetitions are performed for each mode, and the total time for a successful execution is chosen as an evaluation metric.
- **Set 2: Repetitions of Pick-and-place** Several repetitions of the pick-and-place task are performed (similar to the environment in Fig. 16 without the crane and the markers on the pipe). Here, we also consider three modes, namely (a) VR and haptic guidance mode (denoted by VR+HG), (b) VR mode (denoted by VR+Tele), and (c) only with

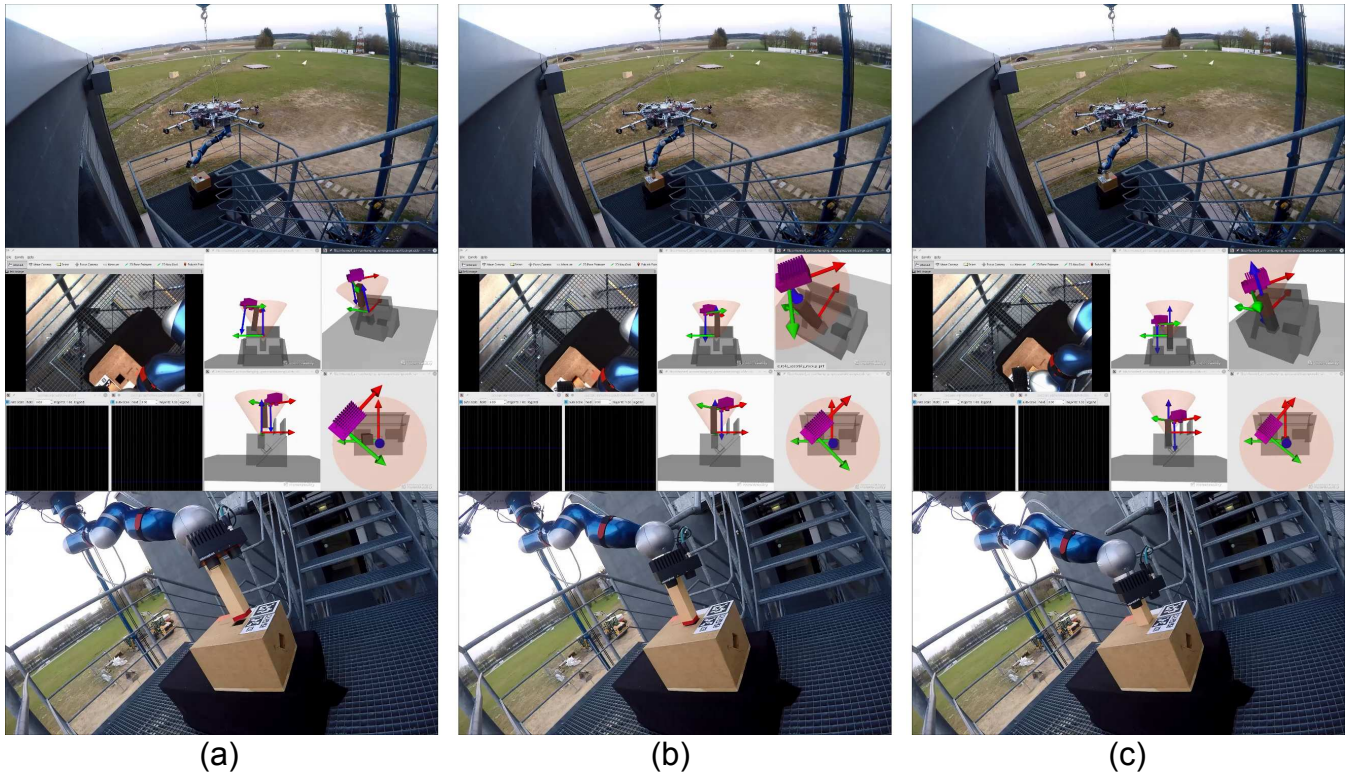


Figure 16: Qualitative results for peg-in-hole task with the aerial manipulator. (a) Approach phase, (b) precise positioning and (c) successful insertion. Top: overview of the robot’s remote workspace, where SAM is suspended by a crane. Middle: the operator view with live video streams, and the created VR, which displays four different view points simultaneously, and provides haptic guidance in both position and orientation. Bottom: close view on the robot’s workspace. With the proposed VR based system, the operator is able to achieve precise peg-in-hole task with the robot in outdoor environments. The markers of size 2.5cm (x3), 6.25cm (x1) and 10cm (x1) are used in this scenario.

live camera streams (denoted CAM). Six repetitions are performed for each mode, with the total time for a successful execution as an evaluation metric.

- Set 3: User Validation** In a laboratory setting, a user validation study is conducted with three subjects. The variations of the users are to demonstrate that the considered manipulation tasks can be performed by different users. The considered tasks are force exertion onto a pipe for three seconds (denoted by validation task 1), and also placing a cage on a pipe with and without moving base (denoted as validation tasks 2 and 3 respectively). With VR and haptic guidance mode, the total time for a successful execution is chosen as an evaluation metric.
- Set 4. Operations at Night with VR** For both the industrial tasks, we perform experiments at night without sunlight. With flash light from an external source, the feasibility and benefits of the proposed system are demonstrated. At night in outdoor environments, this functionality of being able to perform manipulation tasks is important to increase the range of operation hours

including emergency services for several industrial use-cases of aerial manipulators.

With these sets of experiments, the aim is to examine the following aspects. For Set 1, the VR+HG mode are examined for achieving manipulation tasks with a high precision. The proposed marker tracking algorithm is utilized here. With Set 2, the benefits of VR in providing depth information to the operator with haptic guidance only for the orientation, are examined. The LiDAR based object pose estimator is utilized here for pipe localization, while the pose of the cage is monitored with the marker tracking method. Set 3 aims for a user validation, while with Set 4, we attempt to push the limits of the proposed system. Again, the use-case of the developed system is to augment the live video streams by providing 3D visual feedback and haptic guidance. The use of VR interface only, is not the intended use-case of the system. Besides, in industrial scenarios of pipe inspection and maintenance, the pipes are often very long and their inspection points are unknown a-priori. Therefore, the proposed LiDAR based pose estimation method is used to localize the pipe. This use-case justifies the assumption that the object is semantically known, but no geometry is available.

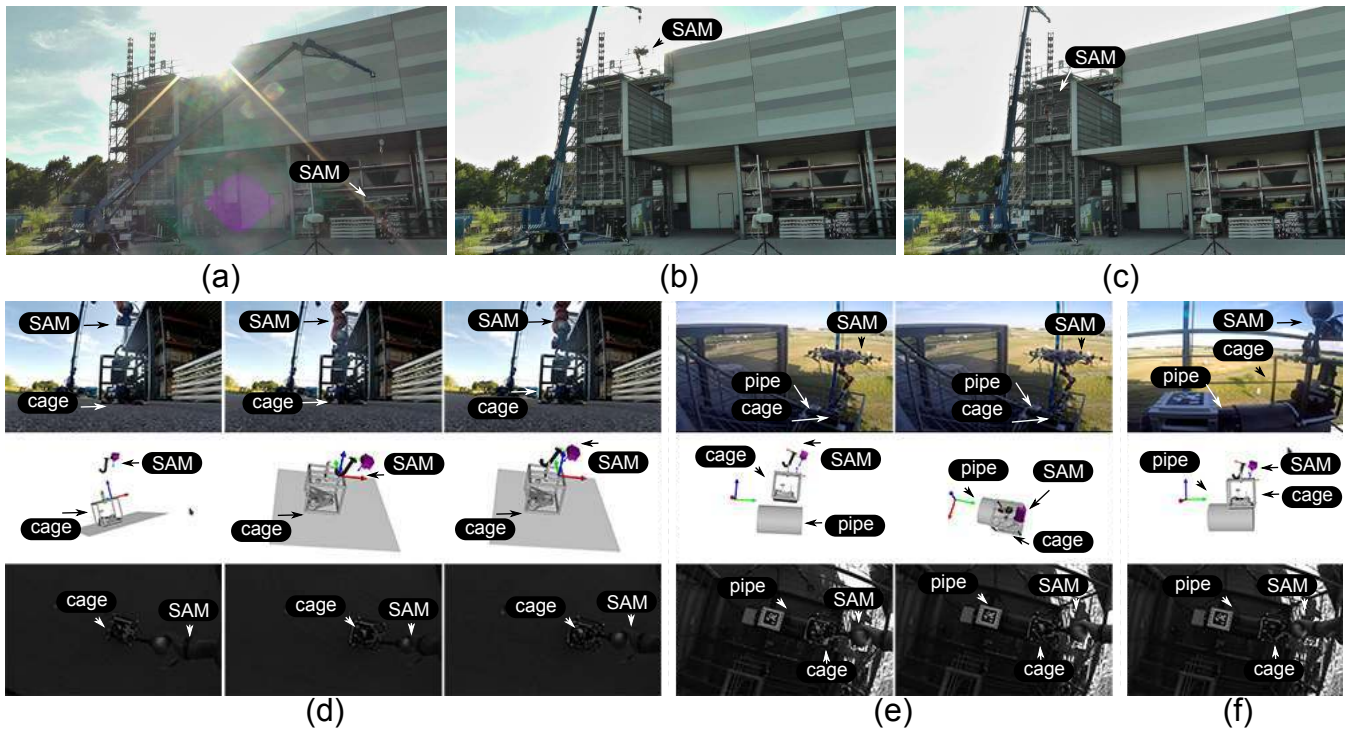


Figure 17: Qualitative results for deployment of robotic crawler within an industrial inspection and maintenance scenario. (a-c) The carrier brings the robot from the ground to the remote location. (d) The robot picks the cage that carries the inspection robotic crawler. (e) The robot places the cage on a mock up of an industrial pipe. (f) The robot exerts force on the cage, so that the crawler can roll out of the cage without falling. The operator can use the VR (bottom), which contains 3D information as opposed to 2D camera images (middle). Live video stream is also subject to over and under exposure when under a shadow on a bright day. With the proposed VR based system, the robot is able to execute advanced aerial manipulation tasks for the considered real-world application. Four markers of size 5cm are used for the cage hosting the mobile robot. The first experiment used a marker of 25cm on the pipe with its CAD model. Later these primitives were replaced by the proposed LiDAR based pose estimation algorithm.

Results The results of Sets 1, 2 and 3 are depicted in Fig. 18. First and foremost, the comparison study of the peg-in-hole insertion tasks with 24 successful executions shows that VR+HG requires the least execution time, while VR+Tele and CAM took similar mean execution times. Note that the executions with CAM used an automated initialization of the end-effector orientation, which was to make the task execution successful. The superiority of VR+HG is expected as the human operator is assisted by the haptic guidance system. Similarly, the comparison study of the pick-and-place task shows similar trend, where the statistics are computed using 18 successful executions in an outdoor environment. As a third point, the user validation demonstrates that all the tasks can be executed by different users with different performance characteristics. Results from the scenario with a moving base, namely validation task 3, required more time for execution, which indicates that the tasks are more challenging with a moving base. Overall, these studies indicate the performance benefits of the system including feasibility and robustness of the proposed system.

The qualitative results of Set 4 are depicted in Fig. 19. The figures show the live-stream view from the eye-in-hand camera, and also from the VR. Lights are provided from an external source and the camera exposures are tuned to achieve balance between noise, brightness and stability of streaming. Poses of the end-effector are plotted to illustrate task executions. These plots are also similar for the peg-in-hole and the pick-and-place tasks from previous sets of experiments. Notably in (b) of Fig. 19, peg-in-hole insertion is best characterized in z-axis between 50s and 60s. In (d) of Fig. 19, the placements are observed in z-axis between 28s and 35s, while the effects of haptic guidance is shown between 8s and 15s. These experiments show that the proposed concept can also work under the unfavorable lighting conditions, thereby extending the operation range of the aerial manipulators. Additional plots for the interaction wrenches during the manipulation task executions can be found in the appendix.




Figure 18: Performance of SAM in terms of execution time. Left: the results of the comparison study is depicted, where we compare pure VR based telemanipulation (VR+Tele), VR with haptic guidance (VR+HG) and a solution using only a camera (CAM). The statistics are computed over 24 and 16 successful executions in outdoor environments for the peg-in-hole and pick-and-place tasks, respectively. Right: the results of user validation is shown, where three users performed three pre-designed tasks, namely force exertion and placing the cage on the pipe. The statistics are computed over 27 successful executions in an indoor environment. Lower the execution time, better the performance.

C. Discussion

The results obtained with ablation studies and field experiments suggest successful development and deployment of the proposed VR based telemanipulation system for advancing aerial manipulation. For providing both the sense of touch and the sense of vision to the human operator, the proposed system featured not only a haptic device, but also a VR interface that provides a real-time 3D display of the robot’s workspace as well as a haptic guidance. In the experiments, it is shown that the system neither requires any external sensors nor pre-generated maps, copes with the challenges of a floating-base manipulation systems, i.e., induced motions of attached sensors due to coupling between the manipulator and the base, fuses multiple sensors whenever appropriate, and has been exhaustively evaluated outside laboratory settings. These features of the proposed VR system are requirements for several industrial applications of aerial manipulation technologies. To the best of our knowledge, using on-board robotic perception only, this work is the first of its kind to demonstrate the feasibility of such VR based concept in dynamic and unstructured environments, which includes several outdoor locations, days and nights, as well as different seasons.

To build such a system, several methodological insights are provided, from the identification of practical challenges to their working solutions, both of which are validated using the real data from robot’s sensors. In particular, the object pose estimators are subject to non-holistic view of the objects, which includes loss-of-sight, partial view and occlusions as examples. For this, we have combined the object pose estimators with ego-motion tracking of the environments using real-time SLAM estimates. In the custom data-sets that emulate these challenges, the results show that the identified problems can be coped with, which has resulted in the pipelines that meets the requirements of VR creation in accuracy, run-time and robustness. Moreover, when one aims

for a long-term deployment of a learning system in outdoor environments, we find that data collection and preparation become a practical problem. To this end, a pool based active learning pipeline has been evaluated, which used a previous work on uncertainty quantification [28]. In a field robotics settings, the results show that only 25% of total data is enough to reach 95% of a solution with all data points and other baseline approaches can be outperformed, overall improving the sampling-inefficiency of DNNs.

Overall, the experiments of this work illustrate several benefits of the proposed VR based telemanipulation system for advancing aerial manipulation capabilities in real-world applications. Intuitively, a virtual environment allows the human operator to change its sight-of-view, zoom in and out, and provides a haptic guidance. In the presented comparison study (with 40 task executions in outdoor environments; a single user), the results show significant reduction in the total execution time when using the proposed system with haptic guidance. The user validation study (three users with 28 total task executions) suggests that three different users can execute the tasks successfully with varying degrees of performance. Moreover, with the demonstration of the operations at night, the range of operation hours has been extended for the current aerial manipulation systems. All these results are obtained within two industrial scenarios that requires advanced aerial manipulation capabilities, namely pick-and-place, force application and peg-in-hole, which goes beyond a contact based inspection. Therefore, these results demonstrate the viability of the proposed VR based telemanipulation concept for industrial applications in the real-world.

D. Lessons Learned

During the flight campaigns with SAM, we learned a few lessons, which we would like to share with the community. These lessons learned are centered around the proposed VR

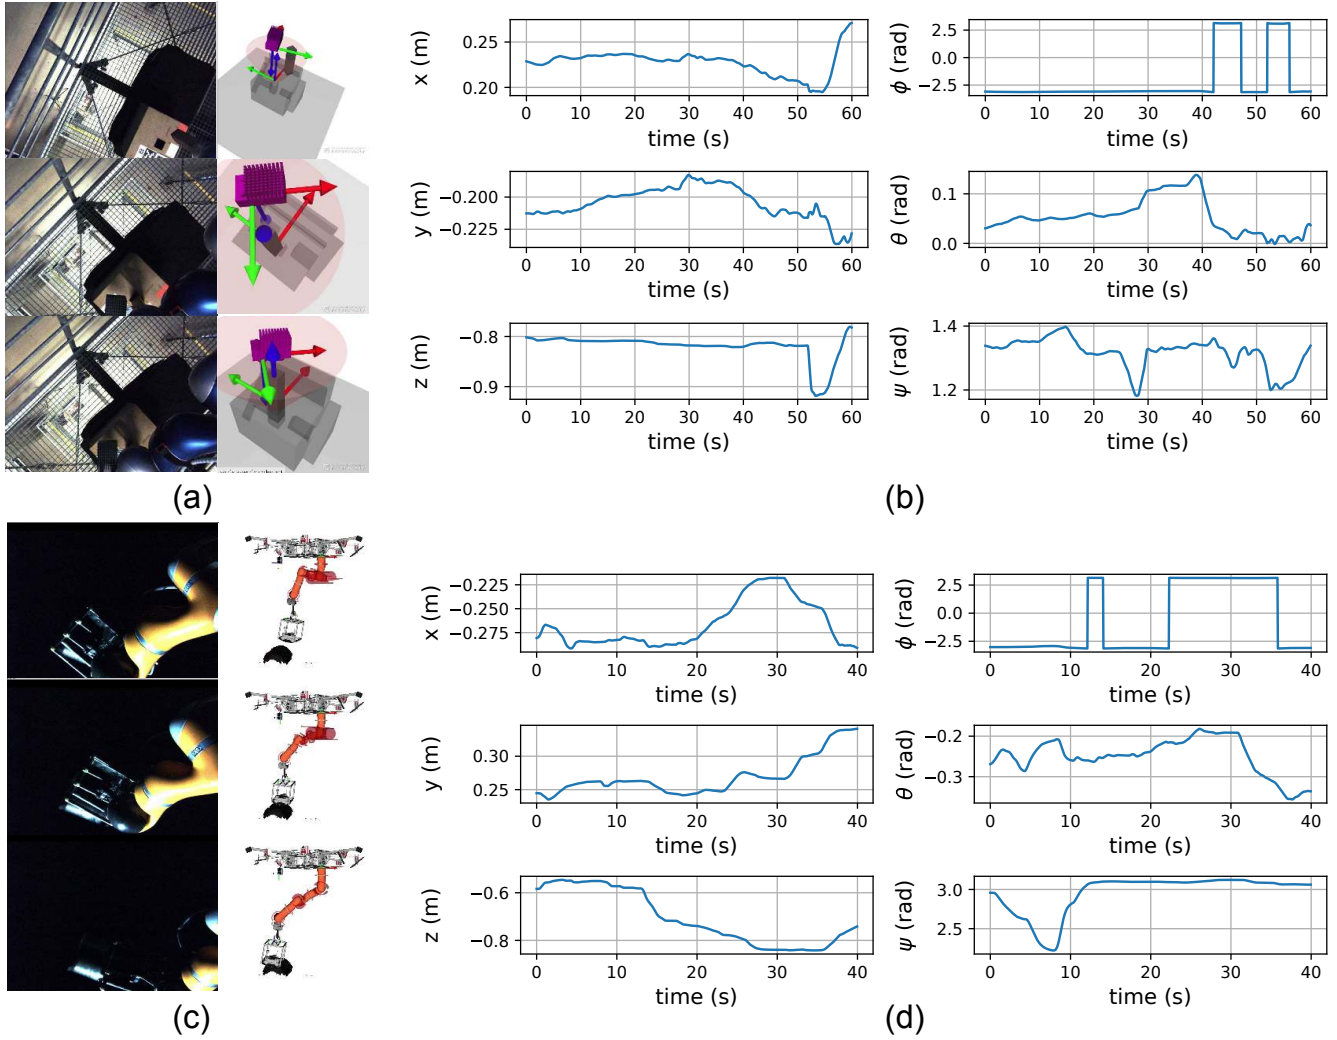


Figure 19: Qualitative visualization of aerial manipulation performed during the night. (a, c) The camera view and the VR. (b, d) The translation and the orientation of the robots’ end-effector. The corresponding pairs are (a,b) and (c,d). These results suggest that the proposed system can extend the operational range of SAM, which further establishes its viability for real-world industrial applications. External views are depicted in Fig. 20.

based telemanipulation system. Note that the focus herein is on the use-cases of the proposed system, the design choices, and the limitations.

On Use-cases of VR with Haptic Guidance for Aerial Manipulation The necessity of VR with haptic guidance for SAM (or robots with similar morphology) largely depends on the choice of the haptic device and difficulty of the manipulation tasks. In the initial flight experiments using only the 2-DoF Space Joystick Rjo, the operator could not easily complete the considered manipulation tasks by only relying on live camera streams [27]. On the other hand, at the later stages of development, it was much easier for the operator to complete the tasks, when we augmented the system with the 6-DOF haptic device Lambda. With the 6-DoF device Lambda and a whole-body controller of the suspended platform to handle occlusions and enhance the

camera’s field of view (e.g., in [64]), the operators could also complete the tasks using only live video streams, despite slower execution time.

However, while the necessity of VR and haptic guidance may depend on the system and the complexity of tasks, we find that the combination of VR, haptic guidance and live video stream resulted in the best performing system. Intuitively, the live video stream can provide situational awareness to the human operator, but suffers from over- and underexposure depending on the light conditions, camera jitters due to the movement of the platform under severe winds, lack of complete 3D information and inability to provide haptic guidance. The proposed VR system can complement the live video stream as it does not degrade with outdoor conditions, provides complete 3D information with an option to change the field of view, and supports



Figure 20: Aerial manipulation at night. The views of the scene from an external camera are depicted. Left: SAM performing peg-in-hole insertion task at night. Right: SAM placing a cage onto a metal pipe for the deployment of a robotic crawler. SAM and the objects are highlighted in white.

haptic guidance. Another benefit is that VR enables seeing the "full model" instead of the limited field of view of the camera at its current position, which includes configuration of the robotic arm.

On Scene Graph Verses 3D Reconstruction The VR creation from robot perception can either rely on scene graph or 3D reconstruction techniques, where the choice of the approach largely depends on the validity of either static-base or floating-base assumptions. For example, a ground based mobile manipulators can first stop, and then perform manipulation. If the scene and objects are static, the relative motion between the sensors and the objects can be easily estimated, and the real-time capability from the perception algorithm is not required. In such a scenario, relying on the outputs of 3D sensors such as RGB-D or stereo would be the simplest option to implement. The robot can map the environments and the objects first to ensure a good field of view, e.g., avoiding occlusions, and then use the map to create a VR. On the other hand, if the relative motion between the sensors and the objects are consistently changing, e.g., in a floating-base system like ours, we find that the scene graph approach can be better suited. The scene graph approach can rely on the object pose estimators that are fast and accurate, and the existing corner cases such as occlusions and loss-of-sight can be handled by using the proposed pipelines. Another consideration is bandwidth, i.e., the object poses require only 6D vectors while streaming point clouds is more expensive. The 6D pose representation can also be plugged in directly for the shared controllers with position based visual servoing (as in VR+HG).

On Inherent Uncertainty in VR Creation The proposed VR from robot perception cannot match the reality perfectly. In spite of this limitation, the considered task could be successfully completed even for several challenging outdoor environments. What attributed to the successful deployment of the proposed VR system was identifying when the VR was prone to failures (see Fig. 5). The provided object

pose estimators mitigates the identified failure cases by combining a standard object pose estimator with tracking of the environments. Here, the combination is facilitated by a module that identifies the failure cases, e.g., self-evaluation of point cloud processing methods, and missed detection of the markers while using visual-inertial systems. Moreover, in the proposed active learning pipeline, a more explicit representation of uncertainty is used to improve the data preparation steps for our DNN based component. Therefore, we find that reliability awareness of an algorithm is crucial for the robotic systems to achieve complex tasks in dynamic and unstructured environments.

The current use of DNN's uncertainty has been off-line, like pool based active learning, while its use on-board the robot could potentially offer several more benefits. In this regard, combining a real-time uncertainty estimation method [83] with a reliability-aware shared control architecture [84], could be an interesting direction of future research for reliable operations of complex systems in unstructured and dynamic environments. Lastly, a full-scale user study is envisioned, which is tailored on telepresence robots with aerial manipulation capabilities, in outdoor environments.

VI. Conclusion

In this article, the real-world applications of aerial manipulation in dynamic and unstructured environments are envisioned. A novel telemanipulation system has been proposed, which involves not only a haptic device for the sense of touch, but also a virtual reality (VR) for enhancing the sense of vision and further providing haptic guidance. To create such system, we identified challenges while using off-the-shelf methods, and devised several extensions to address them. These techniques include pose estimation pipelines for industrial objects of both known and unknown geometries, and also a deep active learning pipeline to efficiently collect and annotate training data. Empirically, we validated the proposed methods using data-sets collected from the robot's sensors. With these,

the influence of each component is examined with regard to mitigating the identified challenges, and we demonstrate the feasibility of creating the real-time and accurate VR. Methodologically, the key to success was an awareness of the algorithms' own failures and uncertainty – also known as robotic introspection. One example is the combination of object pose estimation and SLAM, which is facilitated by a module that identified the failure cases. Another example is the active learning pipeline, where information gain is computed from an explicit representation of uncertainty. Most importantly, with the DLR's SAM platform, we conducted exhaustive experiments over extended durations in which we executed over 70 complex aerial manipulation tasks to characterize the performance of the resulting system. The obtained experimental results show that the proposed system can reduce the execution time of both pick-and-place and peg-in-hole insertion tasks by approximately 1.8 times. The system is also demonstrated to operate at night without any direct sun light. Therefore, the effectiveness of the proposed telemanipulation system is demonstrated for future industrial applications of aerial manipulation technology.

APPENDIX: Platform Design, Control, Teleoperation and IT Architectures

In this section, we present the details about platform design and control, teleoperation system and IT architectures.

A. Platform Control

The control framework of SAM is depicted in Fig. 21. It includes three separate controllers for three sets of actuators. Each of these blocks are to fulfill three different control tasks. The first controller is a propeller based oscillation damping and yaw motion control using IMU as a single main sensor. The main task herein is to damp out oscillations and control yaw motion. Oscillations occur due to the forces and moments caused by the robotic arm, which interacts with the environments. In outdoor settings, severe wind, motion of the carrier, and other external disturbances cause such oscillations. Damping out these undesired motions are to perform precise manipulation tasks with the robotic arm. Similarly, yaw motion controller is to change the orientation of the platform, which can position the manipulator in a convenient pose. To do so, the robot actuation is performed by eight propellers attached to each BLDC motors. The Electronic Speed Controllers (ESCs) regulates BLDC motors to rotate at specific speed. The used sensors are again a single IMU attached to a fixed point of the platform. The control signals are generated by the propeller controller, which is essentially a PID control algorithm. These control signals are mapped to PWM signals per each motors using the known configurations of motors and propellers. The frequency rate is 200Hz in a real-time computer. Secondly, a feedforward position controller is used to control three winches. These winches are connected to cables that suspend the platforms. Maxon motors are used without any feedback signals. The

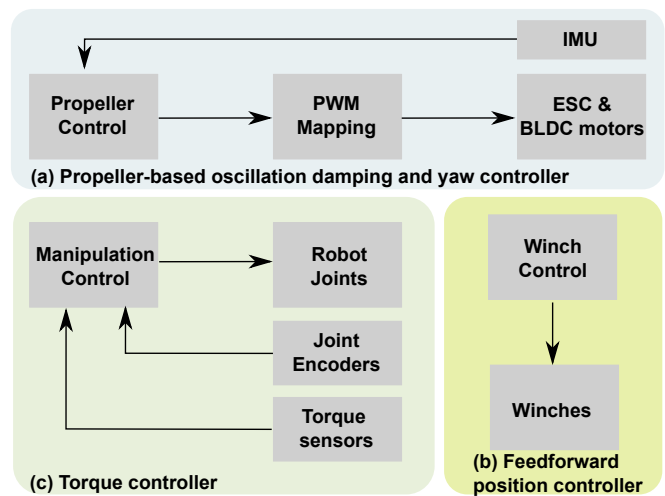


Figure 21: The control framework of SAM. Propeller based oscillation damping and yaw controller, torque controller for robotic manipulator, and feedforward position controller for the winches. Oscillation damping and yaw controller use propeller based actuation to stabilize the platform while SAM is performing manipulation task. IMU provides feedback signals. Impedance based torque control is performed for the robotic arm, where joint encoders and torque sensing are used as feedback. The three winches can adjust the length of the cable suspending the platform.

main feature is to control the length of the cables. With these, the pitch and the roll orientation of the platform can be adjusted with slower dynamics. Another advantage is to move the platform up and down, without moving the carrier. Here, a simple feedforward position controller is integrated where desired relative cable length are converted into the motor movements. The frequency rate of the controller is again 200Hz in a real-time computer. Lastly, the robotic arm attached to the platform, is controlled using impedance based torque control algorithm. Torque control is the current golden standards for such robotic arms. The main task here is to perform grasping and manipulation tasks, using the torque control capabilities of the robot. This means that, in teleoperation mode, the robotic arm must follow the command from the human operator, while autonomously taking care of local redundancy of the joints. Joint encoders and torque sensors provide such feedback signals. The manipulator's internal joint torque controller uses sampling rate of 3kHz.

B. Teleoperation System

The used bilateral teleoperation system is depicted in Fig. 22. The main challenge is ensuring a stable bilateral teleoperation with force feedback. Instabilities can be caused by time delays, packet loss and jitters, which are characteristics of imperfect communication.

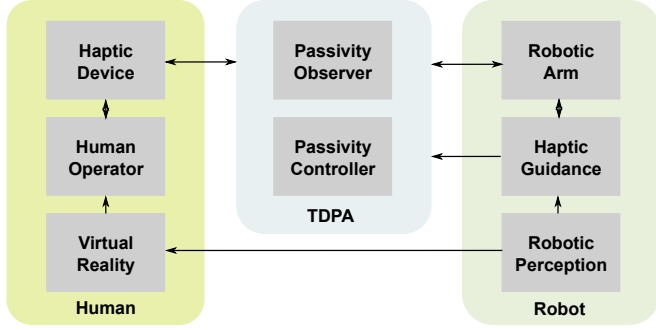


Figure 22: Teleoperation architecture based on Time-Domain Passivity Approach (TDPA). The TDPA approach ensures stability under imperfect communication between the haptic device and the robot, and consists of two components, namely passivity observer and passivity controller. The proposed telepresence system also includes haptic guidance, robotic perception and VR for 3D visual feedback.

In the paper, a two channel architecture with time-domain passivity approach is used, which works as follows. The human operator sends both position $v_m(t)$ (velocity analogously) from the haptic device to the robot at time t . Due to the time-delay T , the robot receives them as $v_{sd}(t) = v_m(t - T)$ where G_h is a scaling factor that can be tuned to match both system dynamics. A local impedance controller then generate reference force $f_s(t)$ based on command position. If K terms represent controller gains and $x_s, v_s \in \mathbb{R}^2$ are the respective feedback signals from the robot which enables position tracking, then the total commanded force can be written as:

$$f_s(t) = K_{ds}(v_{sd}(t) - v_s(t)) + K_{ps}(x_{sd}(t) - x_s(t)). \quad (14)$$

The computed force $f_{m0}(t) = G_s f_s(t - T)$ and measured forces at the end effector $f_{me}(t) = G_e f_e(t - T)$ are sent back to the haptic device resulting in the force feedback term $f_m(t)$ in Eq. 15 where we additionally add a feedforward term with $v_m(t)$. The feedforward terms add transparency to the system and is known to be advantageous over a 2-channel architecture:

$$f_m(t) = f_{m0}(t) + f_{me}(t) + K_{dm}v_m(t). \quad (15)$$

As the signals pass through communication channels, time-delays, jitter and packet losses are typically present and can cause instability of overall system. To cope with this issue, we use TDPA which constitutes two components namely Passivity Observer (PO) and Passivity Controller (PC). Briefly speaking, PO monitors the energy flow of a network whereas PC dissipates the energy introduced by the network. A key underlying idea is PC's control law ensures passivity of the system by damping out the energy that is more than the stored amount. Since passivity is a sufficient condition for stability, TDPA ensures stability in trade-off to performance.

Therefore, POPCs are placed for delayed signals at the robot side $v_m(t - T)$, and haptic device side $f_s(t - T)$. For brevity, let us denote the haptic device signals received and sent as $u_m(k - D)$ and $y_m(k)$, and the robot input and output signals as $u_s(k - D)$ and $y_s(k)$. Here, k is a discrete time and D is a discrete time delay. Then,

$$u_{m,c}(k) = \begin{cases} u_m(k - D) & \text{if } W_m(k) > 0 \\ u_m(k - D) - \frac{W_m(k)}{T_s y_m^2(k)} y_m(k) & \text{else} \end{cases} \quad (16)$$

is the governed control law at haptic device. The same rationale applies at the robot side. In Eq. 16, T_s is the sampling time and $W_m(k)$ is the energy flow at haptic device, which is observed by the PO. In this way, PC modifies the delayed signal so that passivity condition $W_m(k) \geq 0$ for all k is met:

$$\begin{aligned} W_m(k) &= E_{s,in}(k - D) - E_{m,out}(k) + E_{m,PC}(k) \\ E_{s,in}(k - D) &= E_{s,in}(k - D - 1) + T_s P_{s,in}(k - D) \\ 5E_{m,out}(k) &= E_{m,out}(k - 1) + T_s P_{m,out}(k). \end{aligned} \quad (17)$$

PO essentially estimates $W_m(k)$ for PC's control law. This is achieved by Eq. 17 which uses the delayed energy $E_{s,in}(k - D)$ input from the robot side, the energy exiting at the haptic device side $E_{m,out}(k)$ and the dissipated energy by PC $E_{m,PC}(k)$. As the signals being exchanged are velocities v and forces f , the energy can be computed by inner products and sampling time. The power contributions should take into account the direction of energy flow. For example, $P_{s,in} = 0$ if $P_{s,in} \leq 0$ and otherwise, $P_{s,in} = f_s(k)(-v_m(n - D))$. Taking into account time-delays, jitter and packet-loss, TDPA works on energy level and it ensures stability in teleoperation.

Besides, the presented perception algorithms are executed in the robot, providing the information about the object poses. This information is feed into the ground station for creating VR. Another usage of perception is haptic guidance via virtual fixtures. Virtual fixtures [85] are artificial walls that, by means of force feedback, helps the human operator for high performance task execution. Once the human operator is trying to move outside the artificial walls, certain computed forces are activated and sent to the haptic device through TDPA. This then limits the motion of the human operator by inserting certain forces and moments in the haptic device. Because these artificial walls are obtained from the perception system of the robot, the proposed telepresence system supports the haptic guidance. More details about virtual fixtures and other means of haptic guidances are in [63], [86], [87].

C. IT Architectures

In Figure 25, an overview of the used IT architecture is shown. Broadly, the set-up can be divided into the ground station components and the robot itself. To emulate real industrial scenarios of telepresence robots, the connection between the robot and the ground station is established through a WiFi router. The ground station constitutes of a laptop (Dell




Figure 23: Wrench forces and moments from the peg-in-hole experiments. The poses of the end-effector are alternatively depicted in Fig. 19. In F_z , between 50s and 60s, a drop in force is observed due to the successful peg-in-hole insertion.




Figure 24: Wrench forces and moments from the pick-and-place experiments. The corresponding poses of the end-effector are depicted in Fig. 19. In M_z , the haptic guidance activation can be seen, which leads to correction of yaw angle.




Figure 25: SAM’s IT architecture. The broad division is between the ground station and the robot. The ground station hosts haptic device, VR and the human operator. WiFi routers are used to communicate to the robot, which simulates the real teleoperation scenarios with imperfect communication and delays. The robot hosts computers, actuators, sensors and other auxiliary components such as power distribution, battery, etc. SAM has three computers. The robot control unit (RCU) deals with the control of robotic manipulator and hosts real-time linux as its operating system. The flight control computers (FCC) deals with platform control, including propellers, winches and additional servo motors as landing gear. FCC hosts QNX real-time linux system. Lastly, NVIDIA Jetson TX2 is used as a vision processing unit. The TX2 hosts non real-time linux (ubuntu tegra) but has a GPU for deploying the deep learning models.

Latitude 5591), a haptic device (Force dimension Lambda or DLR SpaceJoystick Ryo), VR headset (Meta occulus) and a computer monitor. VR headset is optionally used. From the robot side, the Flight Control Computer (FCC) is employed, which is a product from the DLR aerial robotic spin-off Elektra UAS. FCC is a QNX based real-time system and contains a field-programmable gate array (FPGA) based safety switch. FCC is connected to winches, servo motors (Futaba S3152) and ESCs with custom written drivers. Oscillation damping controller, yaw controller and on-and-off of servos are executed within FCC. This modules read data from IMU (Xsens MTi 100-Series). In addition, FCC is also connected to the manual command transmitter via a radio link. The robot control unit (RCU) is based on Kontron KTH81 Flex board and uses real-time linux patch of open suse operating system. Ethernet for Control Automation Technology (EtherCAT) protocol is used to communication with the robotic arm. We note that EtherCAT is a standardized real-time bus that enables synchronous actuation of all the joint motors. RCU executes TDPA, haptic guidance and impedance controller, while reading joint and torque information from the robotic arm. The last computing module is NVIDIA Jetson TX2 which contains all the sensor drivers, perception software stacks, and other GPU processing modules for running deep learning models. The sensors are all connected via Ethernet interface. For this, Cogswell carrier board is employed, which supports five Ethernet ports with Power of Ethernet (PoE) functionality. With this, Mako camera is easily powered. The carrier board also handles high data throughput from all these sensors. The robot is additionally equipped with safety switch for the robotic arm, power distribution system and battery. All the computers and WiFi routers are connected via an Ethernet switch (Netgear GS105). The communication between the ground station and the robot is through the point-to-point communication channel by opening an access point. Ubiquiti Bullet M5 is employed for the access point.

ACKNOWLEDGMENT

The authors would like to thank Michael Panzirsch and Nicolai Bechtel for the support regarding Lambda, Nidhish Raj for participating in the user validation study, and other members of the flying robots, like Min Jun Kim and Yura Sarkisov, and Elektra UAS team for the support. Special thanks to Martin Schuster and Lukas Meyer for providing detailed feedback on the manuscript. This work is also supported by the Helmholtz Association's Initiative and Networking Fund (INF) under the Helmholtz AI platform grant agreement (ID ZT-I-PF-5-1) and by the European Commission under the contract 644271 EU2020 AEROARMS and 824990 EU2020 RIMA.

References

- [1] Y. S. Sarkisov, M. J. Kim, D. Bicego, D. Tsetsrukou, C. Ott, A. Franchi, and K. Kondak, "Development of sam: Cable-suspended aerial manipulator," in *2019 IEEE International Conference on Robotics and Automation*, May 2019, pp. 5323–5329.
- [2] A. Ollero, M. Tognon, A. Suarez, D. Lee, and A. Franchi, "Past, present, and future of aerial robotic manipulators," *IEEE Transactions on Robotics*, vol. 38, no. 1, pp. 626–645, 2022.
- [3] J. Fishman, S. Ubellacker, N. Hughes, and L. Carlone, "Dynamic grasping with a" soft" drone: From theory to practice," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021, pp. 4214–4221.
- [4] K. Bodie, M. Brunner, M. Pantic, S. Walser, P. Pfändler, U. Angst, R. Siegwart, and J. Nieto, "Active interaction force control for contact-based inspection with a fully actuated aerial vehicle," *IEEE Transactions on Robotics*, vol. 37, no. 3, pp. 709–722, 2020.
- [5] K. Kondak, F. Huber, M. Schwarzbach, M. Laiacker, D. Sommer, M. Bejar, and A. Ollero, "Aerial manipulation robot composed of an autonomous helicopter and a 7 degrees of freedom industrial manipulator," in *2014 IEEE International Conference on Robotics and Automation*, May 2014, pp. 2107–2112.
- [6] S. Kim, S. Choi, and H. J. Kim, "Aerial manipulation using a quadrotor with a two dof robotic arm," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Nov 2013, pp. 4990–4995.
- [7] T. Hulín, M. Panzirsch, H. Singh, R. Balachandran, A. Coelho, A. Pereira, B. M. Weber, N. Bechtel, C. Riecke, B. Brunner *et al.*, "Model-augmented haptic telemanipulation: Concept, retrospective overview and current use-cases," *Frontiers in Robotics and AI*, vol. 8, p. 76, 2021.
- [8] G. Hirzinger, B. Brunner, K. Landzettel, N. Sporer, J. Butterfaß, and M. Schedl, "Space robotics—dlr's telerobotic concepts, lightweight arms and articulated hands," *Autonomous Robots*, vol. 14, no. 2, p. 127–145, 2003.
- [9] R. Balachandran, M. Panzirsch, M. De Stefano, H. Singh, C. Ott, and A. Albu-Schaeffer, "Stabilization of user-defined feedback controllers in teleoperation with passive coupling reference," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3513–3520, 2021.
- [10] J. Artigas, R. Balachandran, C. Riecke, M. Stelzer, B. Weber, J. Ryu, and A. Albu-Schaeffer, "Kontur-2: Force-feedback teleoperation from the international space station," in *2016 IEEE International Conference on Robotics and Automation*, May 2016, pp. 1166–1173.
- [11] C. Masone, M. Mohammadi, P. Robuffo Giordano, and A. Franchi, "Shared planning and control for mobile robots with integral haptic feedback," *The International Journal of Robotics Research*, vol. 37, no. 11, pp. 1395–1420, 2018.
- [12] D. Kim and P. Y. Oh, "Toward avatar-drone: A human-embodied drone for aerial manipulation," in *2021 International Conference on Unmanned Aircraft Systems*, 2021, pp. 567–574.
- [13] G. A. Yashin, D. Trinitatova, R. T. Agishev, R. Ibrahimov, and D. Tsetsrukou, "Aerovr: Virtual reality-based teleoperation with tactile feedback for aerial manipulation," in *2019 19th International Conference on Advanced Robotics*, 2019, pp. 767–772.
- [14] Y. Wu, J. Song, J. Sun, F. Zhu, and H. Chen, "Aerial grasping based on vr perception and haptic control," in *2018 IEEE International Conference on Real-time Computing and Robotics*, 2018, pp. 556–562.
- [15] M. Karrer, M. Kamel, R. Siegwart, and M. Chli, "Real-time dense surface reconstruction for aerial manipulation," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct 2016, pp. 1601–1608.
- [16] A. Pumarola, A. Vakhitov, A. Agudo, F. Moreno-Noguer, and A. Sanfeliu, "Relative localization for aerial manipulation with pl-slam," in *Aerial Robotic Manipulation*. Springer, 2019, pp. 239–248.
- [17] W. Pryor, L. J. Wang, A. Chatterjee, B. P. Vagvolgyi, A. Deguet, S. Leonard, L. L. Whitcomb, and P. Kazanzides, "A virtual reality planning environment for high-risk, high-latency teleoperation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 11 619–11 625.
- [18] J. Mellet, M. Allenspach, E. Cuniato, C. Pacchierotti, R. Siegwart, and M. Tognon, "Evaluation of human-robot interfaces based on 2d/3d visual and haptic feedback for aerial manipulation," *arXiv preprint arXiv:2410.15398*, 2024.
- [19] F. D. Pace, G. Gorjup, H. Bai, A. Sanna, M. Liarokapis, and M. Billingham, "Leveraging enhanced virtual reality methods and environments for efficient, intuitive, and immersive teleoperation of robots," in *2021 IEEE International Conference on Robotics and Automation*, 2021, pp. 12967–12973.
- [20] D. Whitney, E. Rosen, E. Phillips, G. Konidaris, and S. Tellex, "Comparing robot grasping teleoperation across desktop and virtual reality with ros reality," *Robotics Research*, pp. 335–350, 2020.

- [21] K. Huang, D. Chitrakar, F. Ryden, and H. J. Chizeck, "Evaluation of haptic guidance virtual fixtures and 3d visualization methods in telemanipulation - a user study," *Intell. Serv. Robotics*, vol. 12, pp. 289–301, 2019.
- [22] D. Wagner and D. Schmalstieg, "Artoolkitplus for pose tracking on mobile devices," in *Proceedings of 12th Computer Vision Winter Workshop*, 01 2007, p. 139–146.
- [23] J. Zhang and S. Singh, "Low-drift and real-time lidar odometry and mapping," *Autonomous Robots*, vol. 41, no. 2, p. 401–416, feb 2017.
- [24] P. J. Besl and N. D. McKay, "Method for registration of 3-d shapes," in *Sensor fusion IV: control paradigms and data structures*, vol. 1611, 1992, pp. 586–606.
- [25] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 2980–2988.
- [26] A. Ollero, G. Heredia, A. Franchi, G. Antonelli, K. Kondak, A. Sanfeliu, A. Viguria, J. R. Martinez-de Dios, F. Pierri, J. Cortes, A. Santamaria-Navarro, M. A. Trujillo Soto, R. Balachandran, J. Andrade-Cetto, and A. Rodriguez, "The aeroarms project: Aerial robots with advanced manipulation capabilities for inspection and maintenance," *IEEE Robotics Automation Magazine*, vol. 25, no. 4, pp. 12–23, Dec 2018.
- [27] J. Lee, R. Balachandran, Y. S. Sarkisov, M. De Stefano, A. Coelho, K. Shinde, M. J. Kim, R. Triebel, and K. Kondak, "Visual-inertial telepresence for aerial manipulation," in *2020 IEEE International Conference on Robotics and Automation*, 2020, pp. 1222–1229.
- [28] J. Lee, M. Humt, J. Feng, and R. Triebel, "Estimating model uncertainty of neural networks in sparse information form," in *International Conference on Machine Learning*, 2020, pp. 5702–5713.
- [29] J. Lee, R. Balachandran, K. Kondak, A. F. Coelho, M. De Stefano, M. Humt, J. Feng, T. Asfour, and R. Triebel, "Virtual reality via object pose estimation and active learning: Realizing telepresence robots with aerial manipulation capabilities," *Field Robotics*, vol. 3, pp. 323–367, 2023.
- [30] C. Pohl, K. Hitzler, R. Grimm, A. Zea, U. D. Hanebeck, and T. Asfour, "Affordance-based grasping and manipulation in real world applications," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020, pp. 9569–9576.
- [31] S. Kohn, A. Blank, D. Puljiz, L. Zenkel, O. Bieber, B. Hein, and J. Franke, "Towards a real-time environment reconstruction for vr-based teleoperation through model segmentation," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018, pp. 1–9.
- [32] A. S. Vempati, H. Khurana, V. Kabelka, S. Flueckiger, R. Siegwart, and P. Beardsley, "A virtual reality interface for an autonomous spray painting uav," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2870–2877, July 2019.
- [33] C. Liu and S. Shen, "An augmented reality interaction interface for autonomous drone," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020, pp. 11 419–11 424.
- [34] D. Puljiz, F. Krebs, F. Bosing, and B. Hein, "What the hololens maps is your workspace: Fast mapping and set-up of robot cells via head mounted displays and augmented reality," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020, pp. 11 445–11 451.
- [35] P. Ponomareva, D. Trinitatova, A. Fedoseev, I. Kalinov, and D. Tsetserukou, "Grasplook: a vr-based telemanipulation system with r-cnn-driven augmentation of virtual environment," in *2021 20th International Conference on Advanced Robotics*, 2021, pp. 166–171.
- [36] M. Wonsick and T. Padir, "A systematic review of virtual reality interfaces for controlling and interacting with robots," *Applied Sciences*, vol. 10, no. 24, p. 9051, 2020.
- [37] A. Haidu and M. Beetz, "Automated acquisition of structured, semantic models of manipulation activities from human vr demonstration," in *2021 IEEE International Conference on Robotics and Automation*, 2021, pp. 9460–9466.
- [38] Z. Zhang, Y. Zhu, and S.-C. Zhu, "Graph-based hierarchical knowledge representation for robot task transfer from virtual to physical world," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020, pp. 11 139–11 145.
- [39] D. Ni, A. Song, X. Xu, H. Li, C. Zhu, and H. Zeng, "3d-point-cloud registration and real-world dynamic modelling-based virtual environment building method for teleoperation," *Robotica*, vol. 35, no. 10, p. 1958–1974, 2017.
- [40] Z. He, W. Feng, X. Zhao, and Y. Lv, "6d pose estimation of objects: Recent technologies and challenges," *Applied Sciences*, vol. 11, no. 1, p. 228, 2021.
- [41] J. Wang and E. Olson, "AprilTag 2: Efficient and robust fiducial detection," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems*, oct 2016, pp. 4193–4198.
- [42] D. Malyuta, C. Brommer, D. Hentzen, T. Stastny, R. Y. Siegwart, and R. Brockers, "Long-duration fully autonomous operation of rotorcraft unmanned aerial systems for remote-sensing data acquisition," *Journal of Field Robotics*, vol. 37, pp. 137 – 157, 2020.
- [43] M. Laiacker, F. Huber, and K. Kondak, "High accuracy visual servoing for aerial manipulation using a 7 degrees of freedom industrial manipulator," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct 2016, pp. 1631–1636.
- [44] J. Park, Q.-Y. Zhou, and V. Koltun, "Colored point cloud registration revisited," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 143–152.
- [45] S. Rusinkiewicz and M. Levoy, "Efficient variants of the icp algorithm," in *Proceedings third international conference on 3-D digital imaging and modeling*, 2001, pp. 145–152.
- [46] Q.-Y. Zhou, J. Park, and V. Koltun, "Fast global registration," in *European Conference on Computer Vision*, 2016, pp. 766–782.
- [47] F. Pomerleau, F. Colas, and R. Siegwart, "A review of point cloud registration algorithms for mobile robotics," *Foundations and Trends in Robotics*, vol. 4, no. 1, pp. 1–104, 2015.
- [48] Y. Wang and J. M. Solomon, "Deep closest point: Learning representations for point cloud registration," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 3523–3532.
- [49] Z. Zhang, Y. Dai, and J. Sun, "Deep learning based point cloud registration: an overview," *Virtual Reality & Intelligent Hardware*, vol. 2, no. 3, pp. 222–246, 2020.
- [50] A. Narr, R. Triebel, and D. Cremers, "Stream-based active learning for efficient and adaptive classification of 3d objects," in *2016 IEEE International Conference on Robotics and Automation*, 2016, pp. 227–233.
- [51] D. Mund, R. Triebel, and D. Cremers, "Active online confidence boosting for efficient object classification," in *2015 IEEE International Conference on Robotics and Automation*, 2015, pp. 1367–1373.
- [52] D. J. MacKay, "Information-based objective functions for active data selection," *Neural computation*, vol. 4, no. 4, pp. 590–604, 1992.
- [53] J. Gawlikowski, C. R. N. Tassi, M. Ali, J. Lee, M. Humt, J. Feng, A. Kruspe, R. Triebel, P. Jung, R. Roscher *et al.*, "A survey of uncertainty in deep neural networks," *arXiv preprint arXiv:2107.03342*, 2021.
- [54] J. Choi, I. Elezi, H.-J. Lee, C. Farabet, and J. M. Alvarez, "Active learning for deep object detection via probabilistic modeling," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10 264–10 273.
- [55] H. H. Aghdam, A. Gonzalez-Garcia, J. v. d. Weijer, and A. M. López, "Active learning for deep detection neural networks," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 3672–3680.
- [56] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *International Conference on Machine Learning*, 2016, pp. 1050–1059.
- [57] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [58] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International Conference on Machine Learning*, 2015, pp. 448–456.
- [59] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [60] A. O. Albu-Schäffer, S. Haddadin, C. Ott, A. Stemmer, T. Wimböck, and G. Hirzinger, "The dlr lightweight robot: design and control concepts for robots in human environments," *Ind. Robot*, vol. 34, pp. 376–385, 2007.
- [61] F. Kong, G. Zambella, S. Monteleone, G. Grioli, M. G. Catalano, and A. Bicchi, "A suspended aerial manipulation avatar for physical interaction in unstructured environments," *IEEE Access*, 2024.
- [62] E. Niddam, J. Dumon, L. Cu villon, S. Durand, S. Querry, A. Hably, and J. Gangloff, "Design of a suspended manipulator with aerial elliptic

- winding,” *IEEE Robotics and Automation Letters*, vol. 9, no. 9, pp. 7939–7946, 2024.
- [63] H. Thomas, H. Katharina, P. Carsten, S. Chun-Yi, R. Subhash, and L. Honghai, “Interactive features for robot viewers,” in *Intelligent Robotics and Applications*, vol. 7508, 2012.
- [64] A. Coelho, Y. S. Sarkisov, J. Lee, R. Balachandran, A. Franchi, K. Kondak, and C. Ott, “Hierarchical control of redundant aerial manipulators with enhanced field of view,” in *2021 International Conference on Unmanned Aircraft Systems*, 2021, pp. 994–1002.
- [65] R. Balachandran, H. Mishra, M. Panzirsch, and C. Ott, “A finite-gain stable multi-agent robot control framework with adaptive authority allocation,” in *2021 IEEE International Conference on Robotics and Automation*, 2021, pp. 1579–1585.
- [66] K. Schmid, P. Lutz, T. Tomić, E. Mair, and H. Hirschmüller, “Autonomous vision-based micro air vehicle for indoor and outdoor navigation,” *Journal of Field Robotics*, vol. 31, no. 4, pp. 537–570, 2014.
- [67] P. Lutz, M. G. Müller, M. Maier, S. Stoneman, T. Tomić, I. von Bargaen, M. J. Schuster, F. Steidle, A. Wedler, W. Stürzl *et al.*, “Ardea—an mav with skills for future planetary missions,” *Journal of Field Robotics*, vol. 37, no. 4, pp. 515–551, 2020.
- [68] L. B. Rosenberg, “Virtual fixtures: Perceptual tools for telerobotic manipulation,” in *Proceedings of IEEE Virtual Reality Annual International Symposium*, 1993, pp. 76–82.
- [69] A. M. Andrew, “Multiple view geometry in computer vision,” *Kybernetes*, 2001.
- [70] M. A. Fischler and R. C. Bolles, “Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography,” *Commun. ACM*, vol. 24, no. 6, pp. 381–395, Jun. 1981.
- [71] C. Nissler, S. Büttner, Z. Marton, L. Beckmann, and U. Thomasy, “Evaluation and improvement of global pose estimation with multiple apriltags for industrial manipulators,” in *2016 IEEE 21st International Conference on Emerging Technologies and Factory Automation*, Sep. 2016, pp. 1–8.
- [72] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You only look once: Unified, real-time object detection,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 779–788.
- [73] M. Kazhdan, M. Bolitho, and H. Hoppe, “Poisson surface reconstruction,” in *Proceedings of the fourth Eurographics symposium on Geometry processing*, vol. 7, 2006.
- [74] S. Choi, Q.-Y. Zhou, and V. Koltun, “Robust reconstruction of indoor scenes,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 5556–5565.
- [75] M. Humt, J. Lee, and R. Triebel, “Bayesian optimization meets laplace approximation for robotic introspection,” in *IROS 2020 Long-Term Autonomy Workshop*, October 2020. [Online]. Available: <https://elib.dlr.de/137021/>
- [76] D. A. Cohn, Z. Ghahramani, and M. I. Jordan, “Active learning with statistical models,” *Journal of Artificial Intelligence Research*, vol. 4, pp. 129–145, 1996.
- [77] C. M. Bishop, *Pattern Recognition and Machine Learning*. Berlin, Heidelberg: Springer-Verlag, 2006.
- [78] A. Harakeh, M. Smart, and S. L. Waslander, “Bayesod: A bayesian approach for uncertainty estimation in deep object detectors,” in *2020 IEEE International Conference on Robotics and Automation*, 2020, pp. 87–93.
- [79] D. Feng, A. Harakeh, S. L. Waslander, and K. C. J. Dietmayer, “A review and comparative study on probabilistic object detection in autonomous driving,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, pp. 9961–9980, 2022.
- [80] S. Roy, A. Unmesh, and V. P. Namboodiri, “Deep active learning for object detection,” in *British Machine Vision Conference*, vol. 362, 2018, p. 91.
- [81] J. Feng, J. Lee, M. Durner, and R. Triebel, “Bayesian active learning for sim-to-real robotic perception,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 10 820–10 827.
- [82] P. Babin, P. Giguere, and F. Pomerleau, “Analysis of robust functions for registration algorithms,” in *2019 IEEE International Conference on Robotics and Automation*, 2019, pp. 1451–1457.
- [83] J. Lee, J. Feng, M. Humt, M. G. Müller, and R. Triebel, “Trust your robots! predictive uncertainty estimation of neural networks with sparse gaussian processes,” in *Conference on Robot Learning*, 2022, pp. 1168–1179.
- [84] R. Balachandran, H. Mishra, M. Cappelli, B. Weber, C. Secchi, C. Ott, and A. Albu-Schaeffer, “Adaptive authority allocation in shared control of robots using bayesian filters,” in *2020 IEEE International Conference on Robotics and Automation*, 2020, pp. 11 298–11 304.
- [85] A. Bettini, P. Marayong, S. Lang, A. M. Okamura, and G. D. Hager, “Vision-assisted control for manipulation using virtual fixtures,” *IEEE Transactions on Robotics*, vol. 20, no. 6, pp. 953–966, Dec 2004.
- [86] M. Sagardia and T. Hulin, “Multimodal evaluation of the differences between real and virtual assemblies,” *IEEE Transactions on Haptics*, vol. 11, no. 1, pp. 107–118, Jan 2018.
- [87] T. W. Martins, A. Pereira, T. Hulin, O. Ruf, S. Kugler, A. Giordano, R. Balachandran, F. Benedikt, J. Lewis, R. Anderl, K. Schilling, and A. Albu-Schäffer, “Space factory 4.0 - new processes for the robotic assembly of modular satellites on an in-orbit platform based on industrie 4.0 approach,” in *69th International Astronautical Congress*, October 2018.