

This work has been submitted to the IEEE for possible publication.
Copyright may be transferred without notice, after which this version may no longer be accessible.

3D Freehand Ultrasound using Visual Inertial and Deep Inertial Odometry for Measuring Patellar Tracking

Russell Buchanan
Institute of Perception, Action and Behaviour
University of Edinburgh
Edinburgh, UK
russell.buchanan@ed.ac.uk

S Jack Tu
NDORMS
University of Oxford
Oxford, UK
jack.tu@ndorms.ox.ac.uk

Marco Camurri
Faculty of Engineering
Free University of Bozen-Bolzano
Bolzano, Italy
marco.camurri@unibz.it

Stephen J Mellon
NDORMS
University of Oxford
Oxford, UK
stephen.mellon@ndorms.ox.ac.uk

Maurice Fallon
Oxford Robotics Institute
University of Oxford
Oxford, UK
mfallon@robots.ox.ac.uk

Abstract—Patellofemoral joint (PFJ) issues affect one in four people, with 20% experiencing chronic knee pain despite treatment. Poor outcomes and pain after knee replacement surgery are often linked to patellar mal-tracking. Traditional imaging methods like CT and MRI face challenges, including cost and metal artefacts, and there’s currently no ideal way to observe joint motion without issues such as soft tissue artefacts or radiation exposure. A new system to monitor joint motion could significantly improve understanding of PFJ dynamics, aiding in better patient care and outcomes.

Combining 2D ultrasound with motion tracking for 3D reconstruction of the joint using semantic segmentation and position registration can be a solution. However, the need for expensive external infrastructure to estimate the trajectories of the scanner remains the main limitation to implementing 3D bone reconstruction from handheld ultrasound scanning clinically.

We proposed the Visual-Inertial Odometry (VIO) and the deep learning-based inertial-only odometry methods as alternatives to motion capture for tracking a handheld ultrasound scanner. The 3D reconstruction generated by these methods has demonstrated potential for assessing the PFJ and for further measurements from free-hand ultrasound scans. The results show that the VIO method performs as well as the motion capture method, with average reconstruction errors of 1.25 mm and 1.21 mm, respectively. The VIO method is the first infrastructure-free method for 3D reconstruction of bone from wireless handheld ultrasound scanning with an accuracy comparable to methods that require external infrastructure.

Index Terms—3D freehand ultrasound, Deep Learning, Motion estimation, Inertial measurement unit

This work was supported by the EPSRC ORCA Robotics Hub (EP/R026173/1), the EU H2020 Project THING (Grant ID 780883), a Royal Society University Research Fellowship (Fallon), Versus Arthritis MedTech PoC Grant (22702) (Mellon & Tu) and Orthopaedic Research UK Early-career Research Fellowship (Tu)

I. INTRODUCTION

In the UK, approximately 100,000 people undergo total knee arthroplasty (TKA) each year. However, around 20% of TKA patients experience dissatisfaction afterwards [1], [2], typically due to ongoing pain. Abnormal patellofemoral joint (PFJ) kinematics can cause pain after TKA, specifically patellar maltracking [3], [4]. There are several surgical interventions to address patellar maltracking, but their outcomes vary, and they are only applied to patients with obvious patellar maltracking, such as subluxation and dislocation. A method for tracking the patella during knee flexion in patients with painful TKA could provide clinical insight and help identify less obvious forms of patellar maltracking. This, in turn, could be addressed by an appropriate intervention.

Ultrasound has been shown to have enormous potential for bone visualisation and is a low-cost and radiation-free alternative to MRI and X-ray imaging [5]. Consequently, there has been significant interest in using ultrasound for 3D reconstruction of bone. The most accurate 3D freehand ultrasound approaches require external infrastructure to facilitate device tracking, such as a motion capture system [6], [7] or electromagnetic markers [8] for US probe localisation; or the use of a robot arm for scanning [9]. Unfortunately, these methods require significant infrastructure which increases costs and reduces the mobility of a potential device and could hinder clinical implementation.

Freehand 3D ultrasound can be achieved without external infrastructure by measuring the elevational speckle decorrelation between frames [10]. However, it cannot fully capture the complexity of ultrasound image formation. Reference [11] presented a method for creating 3D reconstructions using a convolutions neural network (CNN) for motion estima-




Fig. 1. Flow diagram of complete reconstruction pipeline. The estimation module computes the scanning trajectory using either IMU-only, VIO or motion capture external tracking. Separately, 2D ultrasound points are segmented for the presence of bone. These points are projected in 3D using the estimated trajectory and accumulated into separate point clouds for each sweep over the bone. After manually segmenting out the surface, the point clouds are aligned using ICP and integrated into the surface estimate. Finally, once all the sweeps have been integrated, the final surface is sampled and converted to a mesh using the Poisson method.

tion between ultrasound frames. However, these image-based methods tend to have a significant reconstruction error in at least one plane that is difficult to remove without external localisation.

“State estimation” is a well-established research area for tracking the position and orientation of robots without external infrastructures in robotics [12]. Visual Odometry (VO) is one of the research areas within this field, which involves tracking visual features in the environment using a camera to estimate motion [13]. This approach is often combined with Inertial Measurement Units (IMUs), which are small, low-cost sensors capable of measuring acceleration and rotation rates. VIO has superior robustness to visual challenges, such as darkness or motion blur.

Recently, there have been several works applying deep learning to inertial data to enable low-drift, IMU-only odometry [14]–[16]. These methods train neural networks to learn motion models of, for example, pedestrians or robots. The network can then predict displacement or velocity from inertial data.

In this study, we used robotic state estimation techniques to track a handheld ultrasound scanner. Our objective was to explore the feasibility of infrastructure-free methods for generating a 3D reconstruction of the bone surface from ultrasound imaging. We also developed a motion model of the scanner by training a neural network with only IMU sensing. Finally, we compared the effectiveness of these two methods against a baseline established by external motion capture tracking of the scanner.

II. METHODS

We evaluated three different methods for tracking the ultrasound scanner: motion capture, visual-inertial odometry (VIO), and deep learning IMU-only odometry. An OptiTrack system was used to capture the motion of the ultrasound scanner with reflective markers. For VIO, we used a factor

graph-based approach [16] to combine data from a stereo camera and an IMU mounted on the scanner. To implement the IMU-only method, we trained a neural network to predict position increments based on IMU data and combined it with classical IMU integration, which was similar to previous work [17]. The full system pipeline is shown in Fig. 1

A. Experimental Setup

The experimental setup used is shown in Fig. 2. An Intel Realsense T265 stereo camera, which includes a Bosch BMI085 IMU, was mounted to a handheld ultrasound scanner (L7-38, Clarius Mobile Health Corporation, Burnaby, BC, Canada) with a 3D-printed mount. We used two phantom models of the same distal femur, with the patella in two different positions. The first model had the patella in the position of knee flexion, while the second had the patella in the position of knee extension. Our objective is to reconstruct these two models to visualise the different positions of the patella. To ensure that the bone models remained stationary during scanning, we bolted them to a metal plate. The phantoms were encased in a gelatine medium to enable ultrasound scanning. Retro-reflective markers were placed on the bone and the scanner to enable motion capturing. All the devices were connected to a single PC running a Robotic Operation System (ROS) to record all the sensor data and ensure that all the timestamps were recorded from a common clock.

B. Dataset Collection

A total of 150 scanning motions were taken, out of which 130 involved the user slowly panning across the phantom while repeating different starting positions and scanning directions. The scanning path was short and contained either 2 or 6 sweeps. The remaining 20 scans were longer and involved the user moving the scanner laterally while panning to cover the entire phantom. In all scans, the user picked up the scanner




Fig. 2. Experimental setup for collecting ultrasound dataset. A Realsense stereo camera is mounted on the Clarius ultrasound scanner and a mount holds the scanner in place. The human operator picks up the scanner and passes it over the bone model several times before returning it to the mount.

from a fixed cradle and returned it to the same position at the end.

An OptiTrack motion capture setup was used for ground truth, which has an expected error of less than 1 mm. However, marker occlusions can cause errors in the position estimation. To mitigate this issue, we applied a median filter with a window size of 9 to the camera data, which was recorded at a rate of 200 Hz. The filtered trajectory was used as a reference point for comparing other estimation methods and as training data for the network of the IMU-only method.

C. Tracking methods

The VIO and IMU-only methods are shown in Fig. 3 top and bottom, respectively. Both methods were applied to estimate the Maximum-A-Posteriori (MAP) state trajectory of a handheld ultrasound scanner in batch. The state estimator was implemented using the factor-graph library GTSAM [18] and built on the work from [17] for VIO, while the neural




Fig. 3. Two example factor graph frameworks for five states $j = 0, 1, 2, 3, 4$. On top is the VIO method which uses prior, visual and IMU factors. On the bottom is the learned inertial odometry method which only uses IMU and learned factors. The four types of factors are: prior factors constraining the start and end pose to be equivalent, IMU integration factor from [19], visual landmark factors from [17] and our learned displacement factor from [16].

network for the IMU-only was based on [16] and developed using the PyTorch library.

In this work, we assumed each trajectory has a “loop closure,” meaning the scanner trajectory ends with the same position and orientation as at the start. We also assume zero velocity at the start and end (as shown in Fig. 3). Two prior factors are used, one for the first pose and one for the last pose. We use standard IMU pre-integration theory to connect consecutive nodes in the graph, as in [19]. This relates to the pose, velocity, and biases between two consecutive states.

1) *VIO Method*: IMU data was recorded at 200 Hz while the stereo camera was sampled at 30 Hz with one in three frames selected as a keyframe. Therefore, nodes were added to the graph at 10 Hz, with IMU preintegration factors linking each consecutive pair.

2) *IMU-only Method*: the IMU data samples were buffered in windows of 1 s ($N = 200$). The inference was done at a rate of 30 Hz, which we empirically found to give the best performance. Therefore, each learned displacement links two states 1 s apart and the neural network starts producing outputs after 1 s of data is collected.

The IMU and OptiTrack data collected during the scanning were divided into two sets of 80% for training and 20% for validation, respectively. A single model was trained using all data. The Adam optimizer was used with a learning rate of 10^{-6} over 2000 epochs and the model minimized validation error. The training process took about 6 hours on an NVIDIA GeForce RTX 3080 Ti GPU. Two scan types were used for both phantoms, labelled as Scan Type A for cross-sections (x direction) and Scan Type B for longitudinal (y direction) sections.

D. Reconstruction

The reconstruction was adopted from a previous work [7] with weighted projective fusion for point cloud integration and surface reconstruction. The final output mesh can be directly compared with ground truth meshes of bone phantom models.

1) *Surface segmentation*: The bone surface was segmented from ultrasound images using the method in [20]. The seg-

mented pixels are then projected into 3D as points (as shown in Fig. 4). With the trajectory estimate from the tracking module, the projected points are accumulated into point clouds according to the “sweep” number over the bone. Finally, the reconstructed point clouds are reviewed and cleaned by removing mislabelled points. Several point clouds of the bone surface corresponding to a separate sweep were carried on for further evaluation.

2) *Alignment of Point Cloud Sweeps*: For VIO and IMU-only methods there is incremental drift while estimating the scanner’s pose. Knowing that we are scanning the same bone over several sweeps, we can correct this drift and align the sweep points using Iterative Closest Point (ICP) [21] with an open-source library *libpointmatcher* [22]. ICP was constrained to position and yaw (rotation about the gravity axis) only, considering the acceleration, pitch and roll estimate from the IMU were highly accurate over a short time period of a single scan. Although errors caused by tracking drift while moving across the bone cannot be corrected as the sweeps are treated as rigid point clouds, we assume that the drift incurred over this small distance is not enough to severely degrade the reconstruction.

While the Optitrack-estimated sensor pose is not subject to drift, aligning point clouds also slightly improves reconstruction for this configuration.

3) *Surface Reconstruction and Meshing*: The open source library Open3D [23] were used to perform Truncated Signed Distance Fields (TSDF) representation and integration algorithm [24] to fuse the aligned point clouds into a smooth mesh surface. TSDFs are a volumetric representation where voxel values correspond to the signed distance from the surface boundary, being positive outside the surface and negative when penetrating inside the surface. TSDF integration was computed with 2 mm resolution and 1 cm truncation with 1 mm voxel filtering and Point-to-Plane matching to improve robustness to outlier points.

Alignment and surface reconstruction were done in a loop whereby point clouds were first aligned and then integrated into the surface. A point cloud sampling of the TSDF surface is then extracted and used to align the next point cloud. Once all the sweeps have been aligned and integrated into the final surface, then the point cloud is converted to a mesh using the Poisson method [25].

III. RESULTS AND DISCUSSION

A. Tracking Results

Table I shows the level of estimation error of each tracking method compared to the Optitrack trajectory, which was considered the ground truth. Visual-inertial and deep learning inertial values are presented with and without the loop closure constraint from the start and end of prior factors.

Both the Absolute Pose Error (APE) and Relative Pose Error (RPE) translation components [26] were computed. APE gives the average error of each trajectory compared to the ground truth trajectory at any point in time. RPE, on the other hand, shows the rate of drift over a given distance. As shown in




Fig. 4. Ultrasound images collected through the Clarius API are in a size of 640×480 . The origin of the ultrasound image frame M was defined as pixel (0,0). A fixed, known transformation from M into the ultrasound scanner’s frame U was applied to all segmented pixels to project them into 3D points. Finally, the 3D points in U were transformed into the global coordinate system using the selected odometry source.

TABLE I
POSITION ESTIMATION ERROR [MM]

Method	APE				RPE (20 mm)			
	Bone 1		Bone 2		Bone 1		Bone 2	
Scan Type	A	B	A	B	A	B	A	B
VIO	8.7	3.8	3.5	4.2	2.3	1.9	1.6	1.2
VIO - Loop	8.3	3.8	4.9	4.2	2.3	1.9	1.6	1.2
IMU-Only	53.2	124.4	112.2	18.9	10	14.7	14.6	9.7
IMU-Only - Loop	49.0	41.9	132.7	34.7	9.9	11.8	15.3	9.7

Table I, RPE was computed over 20 mm segments, indicating that the value represents the average drift for every 20 mm travelled.

Adding the loop closure slightly decreases the average position error (APE) in most cases except for Bone 2 Scan A, but it has no effect on the RPE. This is because the VIO method is already highly accurate and drifts an average of 1.8 mm per 20 mm. In contrast, the IMU-only methods have larger errors than VIO. Even though adding the loop closure does reduce both RPE and APE for most trajectories, it remains higher than VIO.

The trajectories for all six scans for Bone 1 (Fig. 5) and Bone 2 (Fig. 6) were plotted and showed a clear difference in the magnitude of error between VIO and IMU. The VIO methods were almost perfectly aligned with Optitrack, whereas the IMU-Only method had more significant errors. However, for five of the six trajectories, the method largely followed the Optitrack trajectory with a slight offset. It is evident that the network has learned the scanning motion and is providing useful information to the estimator. In the reconstruction step, we corrected these errors using ICP.

B. Reconstruction Results

Table II shows the mean and worst distance errors of each method when compared to the ground truth mesh model of the phantom bone. The motion capture method and the visual-inertial method had average errors of 1.21 mm and




Fig. 5. Trajectories for the different algorithms while scanning Bone 1. Visual-inertial and learned inertial methods both use the loop closure constraint. In these plots, the trajectory is cropped to start when the scanner is in contact with the bone and cropped to end when scanning is finished. The first pose is aligned with ground truth in position and yaw, not roll or pitch.




Fig. 6. Trajectories for the different algorithms while scanning Bone 2. Visual-inertial and learned inertial methods both use the loop closure constraint. Alignment is the same as for Fig. 5.

1.25 mm respectively, with only a 3% difference between them. However, the IMU-only method had an average error of 1.85 mm, which was 35% higher than the motion capture method. This is because the IMU-only method has higher drift rates

Fig. 7 and Fig. 8 present the reconstructed models of Bone 1 and Bone 2 for each scanning type. The coloured mesh represents the distance error from the ground truth model, and the surface reconstruction is overlaid on top of the ground truth model for each bone. It is important to note that the mesh reconstructions may contain some errors, such as a flaring out of the mesh at the edges, which is observable in Bone 1 Scan A with OptiTrack (Fig. 7).

In this study, we demonstrated a slightly higher error rate compared to other studies, without requiring expensive external infrastructure. Reference [8] used EM trackers to track the scanner’s movement on the skin and reported an average error of 1.05 mm while scanning the distal femur and proximal tibia of 15 cadaveric specimens. Reference [9] used a robotic arm to perform the scanning, and they reported an average error of 0.85 mm across two distal femurs (one human model and one bovine) using their best results, with an average worst

TABLE II
MEAN/WORST RECONSTRUCTION ERROR [MM]

Method	Bone 1		Bone 2	
Sweep Type	A	B	A	B
Optitrack	0.99 / 7.71	1.00 / 7.50	1.53 / 9.95	1.76 / 5.11
VIO	0.85 / 5.2	1.10 / 3.93	1.70 / 6.43	1.80 / 6.32
IMU-Only	1.8 / 7.4	1.51 / 7.67	2.56 / 10.03	1.98 / 7.93

error of 9.06 mm.

Other reconstruction methods can cover a wider area of the bone, but they may introduce additional problems. For example, using a statistical shape model like in [8] can lead to biased reconstructions. Similarly, submerging the sensor in water as in [9] is not a practical solution. In contrast, our method involves using a general-purpose scanner held by a human and does not require an SSM. This approach makes our method more practical and versatile.

C. Future Work

Our work shows that using outward-facing VIO to track the pose of an ultrasound probe is accurate enough to perform markerless, handheld ultrasound reconstruction with ICP-based reconstruction. However, the IMU-only method is not currently practical due to its high drift rate. Although incorporating more prior information and accumulating point clouds over shorter distances may help, it could make ICP matching more challenging. Additionally, introducing an SSM could result in bias in the reconstruction, but it may be useful for estimating the joint position, provided that there are no bone abnormalities. To determine the impact of available data on the results, we tested models with less data and found that RPE only increased by 2.3% when half the data was used. We have also discovered that having more data does not significantly enhance the IMU-only tracking. Instead, we recommend exploring new neural network architectures.

IV. CONCLUSION

We introduced a new method for creating a 3D reconstruction of bones that doesn’t require the use of markers or external infrastructure. The technique involves using a handheld ultrasound scanner and combining it with recent robotic estimation techniques such as visual-inertial odometry and learning inertial motion tracking to determine the position of the scanner. Our experiment indicates that using outward-facing VIO to generate bone surfaces results in errors that are comparable to those obtained using external tracking infrastructure. Precise and streamlined 3D freehand ultrasound possesses significant potential across various clinical applications, for example, the method described in the current study can be used to visualise and measure the position of the patella with varying degrees of knee flexion in patients with painful TKA.

REFERENCES

- [1] K. K. Petersen, O. Simonsen, M. B. Laursen, T. A. Nielsen, S. Rasmussen, and L. Arendt-Nielsen, “Chronic postoperative pain after primary and revision total knee arthroplasty,” *The Clinical journal of pain*, vol. 31, no. 1, pp. 1–6, 2015.




Fig. 7. Reconstructions of Bone 1 which is in 40° flexion. We show results for each estimation method and for both scanning types A and B.




Fig. 8. Reconstructions of Bone 2 which is in full extension. The patella is now directly over the distal femur. We show results for each estimation methods and for both scanning types A and B.

- [2] P. J. van der Wees, J. J. Wammes, R. P. Akkermans, J. Koetsenruijter, G. P. Westert, A. van Kampen, G. Hannink, M. de Waal-Malefijt, and B. W. Schreurs, "Patient-reported health outcomes after total hip and knee surgery in a dutch university hospital setting: results of twenty years clinical registry," *BMC musculoskeletal disorders*, vol. 18, no. 1, pp. 1–10, 2017.
- [3] C. Belvedere, A. Ensini, A. Leardini, V. Dedda, A. Feliciangeli, F. Cenni, A. Timoncini, P. Barbadoro, and S. Giannini, "Tibio-femoral and patello-femoral joint kinematics during navigated total knee arthroplasty with patellar resurfacing," *Knee Surgery, Sports Traumatology, Arthroscopy*, vol. 22, pp. 1719–1727, 2014.
- [4] R. S. Pulavarti, V. V. Raut, and G. J. McLauchlan, "Patella denervation in primary total knee arthroplasty—a randomized controlled trial with 2 years of follow-up," *The Journal of Arthroplasty*, vol. 29, no. 5, pp. 977–981, 2014.
- [5] L. N. Nazarian, "The top 10 reasons musculoskeletal sonography is an important complementary or alternative technique to mri," *American Journal of Roentgenology*, vol. 190, no. 6, pp. 1621–1626, 2008.
- [6] E. Świątek-Najwer, K. Otto, P. Krowicki, K. Krysztoforowski, P. Keppler,

- and J. Kozak, "3d bone shape modelling basing on dataset recorded by ultrasound free-hand navigated probe," in *Information Technologies in Biomedicine, Volume 4*, E. Piętko, J. Kawa, and W. Wicławek, Eds. Cham: Springer International Publishing, 2014, pp. 45–56.
- [7] R. Jia, P. Monk, D. Murray, J. A. Noble, and S. Mellon, "CAT & MAUS: A novel system for true dynamic motion measurement of underlying bony structures with compensation for soft tissue movement," *J Biomech*, vol. 62, pp. 156–164, 2017.
 - [8] M. R. Mahfouz, E. E. Abdel Fatah, J. M. Johnson, and R. D. Komistek, "A novel approach to 3d bone creation in minutes," *The Bone & Joint Journal*, vol. 103-B, no. 6 Suppl A, pp. 81–86, 2021.
 - [9] W. Kerr, P. Rowe, and S. G. Pierce, "Accurate 3d reconstruction of bony surfaces using ultrasonic synthetic aperture techniques for robotic knee arthroplasty," *Computerized Medical Imaging and Graphics*, vol. 58, pp. 23–32, 2017.
 - [10] A. H. Gee, R. James Housden, P. Hassenpflug, G. M. Treece, and R. W. Prager, "Sensorless freehand 3d ultrasound in real tissue: Speckle decorrelation without fully developed speckle," *Medical Image Analysis*, vol. 10, no. 2, pp. 137–149, 2006. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1361841505000708>
 - [11] R. Prevost, M. Salehi, S. Jagoda, N. Kumar, J. Sprung, A. Ladikos, R. Bauer, O. Zettinig, and W. Wein, "3D freehand ultrasound without external tracking using deep learning," *Medical Image Analysis*, vol. 48, pp. 187–202, Aug. 2018.
 - [12] T. D. Barfoot, *State Estimation for Robotics*. Cambridge University Press, 2017.
 - [13] R. Mur-Artal and J. D. Tardós, "ORB-SLAM2: an open-source SLAM system for monocular, stereo and RGB-D cameras," *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
 - [14] C. Chen, X. Lu, A. Markham, and N. Trigoni, "IONet: Learning to Cure the Curse of Drift in Inertial Odometry," *32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, pp. 6468–6476, 2018.
 - [15] W. Liu, D. Caruso, E. Ilg, J. Dong, A. I. Mourikis, K. Daniilidis, V. Kumar, and J. Engel, "TLIO: Tight learned inertial odometry," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5653–5660, 2020.
 - [16] R. Buchanan, M. Camurri, F. Dellaert, and M. Fallon, "Learning inertial odometry for dynamic legged robot state estimation," in *Proceedings of the 5th Conference on Robot Learning*, vol. 164, 2021, pp. 1575–1584.
 - [17] D. Wisth, M. Camurri, and M. Fallon, "Vilens: Visual, inertial, lidar, and leg odometry for all-terrain legged robots," *IEEE Transactions on Robotics*, vol. 39, no. 1, pp. 309–326, 2022.
 - [18] F. Dellaert, "Factor Graphs and GTSAM: A Hands-on Introduction," Georgia Institute of Technology, Tech. Rep. September, 2012. [Online]. Available: <https://smartechn.gatech.edu/handle/1853/45226>
 - [19] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza, "On-Manifold Preintegration for Real-Time Visual-Inertial Odometry," *IEEE Transactions on Robotics*, vol. 33, no. 1, pp. 1–21, 2017.
 - [20] R. Jia, S. J. Mellon, S. Hansjee, A. P. Monk, D. W. Murray, and J. A. Noble, "Automatic bone segmentation in ultrasound images using local phase features and dynamic programming," in *IEEE 13th International Symposium on Biomedical Imaging (ISBI)*, 2016, pp. 1005–1008.
 - [21] F. Pomerleau, F. Colas, and R. Siegwart, "A review of point cloud registration algorithms for mobile robotics," *Foundations and Trends in Robotics*, vol. 4, pp. 1–104, 05 2015.
 - [22] F. Pomerleau, F. Colas, R. Siegwart, and S. Magnenat, "Comparing ICP Variants on Real-World Data Sets," *Autonomous Robots*, vol. 34, no. 3, pp. 133–148, 2013.
 - [23] Q.-Y. Zhou, J. Park, and V. Koltun, "Open3d: A modern library for 3d data processing," 2018. [Online]. Available: <https://arxiv.org/abs/1801.09847>
 - [24] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon, "KinectFusion: Real-time dense surface mapping and tracking," in *10th IEEE International Symposium on Mixed and Augmented Reality*, 2011, pp. 127–136.
 - [25] M. Kazhdan, M. Bolitho, and H. Hoppe, "Poisson Surface Reconstruction," in *Symposium on Geometry Processing*, A. Sheffer and K. Polthier, Eds. The Eurographics Association, 2006.
 - [26] R. Kümmerle, B. Steder, C. Dornhege, M. Ruhnke, G. Grisetti, C. Stachniss, and A. Klein, "On measuring the accuracy of slam algorithms," *Autonomous Robots*, vol. 27, no. 4, pp. 387–407, 2009.