Robotic assisted 3D reconstruction with high resolution textures

Geng Lyu[1][0000−0002−5864−3952], Wenbin Hu[2][0000−0003−3073−5182], and Yang Luo[3][0000−0002−7478−1159]

[1] University of the West of England, Bristol
[2] University of Edinburgh
[3] Xi'an Jiaotong-Liverpool University

Abstract. 3D reconstruction remains a challenging scientific problem, encompassing various fields such as computer-aided geometric design, computer vision, computer graphics, and computational science. Traditional methods, whether active using mechanical or radiometric rangefinders or passive based on computer vision, predominantly focus on recovering 3D geometry rather than textures. However, highly detailed textures are crucial in fields requiring precise representation, such as digital museums and medical imaging. Recent advancements in 3D reconstruction based on macro photography and focal stacking have yielded impressive results. Building on these advancements, we aim to refine the process further. Inspired by the renowned Digital Michelangelo Project, which features well-designed hardware and software, we propose using robotic-assisted image capturing for microscopic 3D reconstruction. We have designed and implemented a robotic-assisted 3D reconstruction pipeline that captures microscopic image sets from subject surfaces, generating 3D reconstructed surfaces with microscopic-level details. This approach promises to enhance the precision and detail of 3D reconstructions, making them more suitable for applications requiring high fidelity and accuracy.

Keywords: First keyword · Second keyword · Another keyword.

1 Introduction

The 3D reconstruction has always been a difficult scientific problem, which involves a wide variety of fields, such as computer aided geometric design, computer vision, computer graphics, computational science. No matter either active methods which use mechanical or radiometrical rangefinders or passive methods based on computer vision, most of them focus on recovering the 3d geometry structure more than textures. However, highly detailed texture is essential in the areas that require precise representation, such as digital museums or medical imaging.

3D reconstruction methods base on macro photography and focal stack have achieve really good results, and we decided to try pushing it to a more delicate level. Inspired by the famous digital Michelangelo project [26] which has well

designed hardware setup and software, we decide that robotic assisted image capturing might be a good start for our microscopic 3D reconstruction.

We design and implement a robotic assisted 3d reconstruction pipeline which captures the microscopy image set from the subject surface, and generates a 3D reconstructed surface with microscopy level details.

2 Related works

2.1 Robotic control

Even for objects with simple geometry structure, it is almost impossible for human user to hold a microscope camera stable while maintaining contactless with the surface of the object. So we introduce a robotic arm to help with the image capturing.

Manipulator Motion Planning

Motion planning is a fundamental problem in robotics, where the robot has to find one collision-free motion from the starting state to the target state while satisfying intrinsic constraints such as joint and torque limits. Extended from path planning, which is the purely geometric problem, motion planning in a real robot concerns the dynamics, motion constraints, control inputs and even the uncertainties of sensory feedback. The modern motion planning algorithms and applications are covered by many textbooks[25, 39, 30] and surveys[38, 34, 28]. As one of the most thriving research topic in robotics, many modern algorithms have achieved remarkable results in addressing hard instances of the motion planning, such as real-time planning in complex environments [22][42].

Nowadays there are many open source motion planning frameworks which integrate mature implementations of motion planning algorithms. The frameworks are usually robot agnostic which means users can plan the motion of any robot once a proper model is provided. Some popular motion planning frameworks are: MoveIt[12], ExoTica[20], CoppelliaSim[36] and OpenRAVE[15].

Coverage Task Path Planning

Coverage path planning(CPP) is the problem of determining a path of the end-effector that covers the whole target surface or volume while avoiding obstacles. This task is the fundamental of many robotic applications in manufacturing industry such as polishing[40], painting[4][27] or surface defect inspection[31].

Nowadays most CPP algorithms solve the problem in two steps: First they decompose the target space into many simpler, non-overlapping sub-spaces(or mostly called cells) and then generate the planning in each of the cells[18]. Conventionally there are two off-line cell decomposition approaches which build the foundation of many CPP algorithms: Trapezoidal[24][10] and Boustrophedon decomposition[11]. Afterwards Morse-based cell decomposition approach[2] has been proposed for complete on-line coverage planning[1].

2.2 3D reconstruction

3D reconstruction is a always difficult problem in the past decades and there are plenty of researches about it [16]. Although there are special approaches like

shape-from-shading [43], photometric stereo [5], shape from texture [3], we focus on image based methods.

Structure from motion

Different from SLAM, structure from motion (SFM) takes unordered images collections for 3D reconstruction. And it has the flexibility of using images capture through different ways to enable tasks like reconstructing from internet photo collections [37], which is based on self-calibrating metric reconstruction systems [6].

Inspired by similar works, large scale reconstruction systems are also built and some of them can even take a hundred million internet photos [21].

Macro photogrammetry

Currently, macro photography [14] is widely used by artists for capturing small subjects like insects. Usually, a macro lens with a large reproduction ratio of at least 1:1 is used for this extreme close-up photography technology. Focal stack technique is used in macro photography applications like [17] to deal with the focal blur introduced by shallow depth of field. Naturally, some researchers make use of macro photography to collect data for 3D reconstruction of small subjects [19] and they managed to achieve really high accuracy within the bounding box.

RGB-D fusion

The availability of affordable depth scanners enables many applications of computer vision in the areas like robotics, human motion capture and scene modeling. However, a common drawback of these sensors is their lack of reliability and accuracy. On the other hand, they have the advantage of mobility and high frame rate compared with professional 3D scanners like laser radar, which means it is much easier to get a lot of low reliability depth maps of the scanned object from different camera poses. Therefore, researchers like [32, 33, 9] fuse the captured data to enhance the accuracy of the acquired depth map. Additionally, fusing the captured data can be use to combine different parts of the scanned object to create a 360 degree reconstructed model.

3 Methods

3.1 Curved surface scanning with robotic arm

To scan the curved surface reliably for capturing the microscopy image set, we turned to robot assisted capturing. The microscope is require to operate at a certain distance maintain being perpendicular to the subject surface to capture sharp images. Ideally, the robotic arm moves the digital microscope along the curved surface and capture a microscopy image sequence covers the entire region of interests during the motion as shown in Figure 1a, the accurate poses of the camera are also recorded at the moments when the associated microscopy image is captured. Another preferable requirement is the adjacent frames in the image sequence should have overlapping area large enough to ensure a true positive matching.

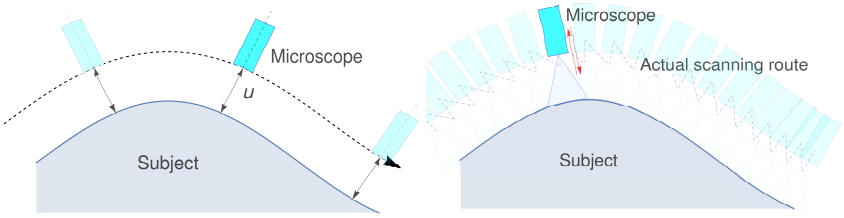

(a) Desired scanning route. (b) Actual scanning route.

Fig. 1: The desired scanning route and the actual scanning route which also applies the auto focusing. Noticing that the sampling points might not be a spatial continuous sequence in space since the scanning route is generated through depth-first traversal

In the robotic control area, there is a specific problem called coverage path planning for passing over all points in the region of interests. However, it requires a full awareness of the environment before planning.

We choose to use the robotic arm and its kinematics sensor for a rough surface geometry acquirement. First we choose some sampling points from the subject surface, and measure their 3D positions in the robot coordinating system via the robotic arm sensors by manually guiding the robotic arm to touch the subject surface.

Considering the working space and the reach range of the robotic arm, it is acceptable to approximate the subject's surface geometry by curve fitting and generate a surface point cloud G and corresponding normal map of the subject. The desired sampling points S_i on the surface can be generated by simply downsampling G via a proper voxel size. We sort the sampling points through depth-first traversal to get a relatively short scanning route.

According to the previous estimation of the depth of field (The depth of field of microscope we use is $0.26mm$ wide), the auto focusing movement is set to be $\pm 5mm$ from the proposed sampling point within in $3seconds$ to make sure that at least one image is supposed to be captured when the subject surface falls in the depth of field.

For each sampling point, the microscopy image and the camera pose when it is taken would be recorded for further processing.

3.2 Surface reconstruction

We treat this microscopy surface reconstruction problem as an image stitching problem on an unknown curved surface. The goals are:

- Seamless texture stitching at the tangent direction.
- Maintain the 3D geometry of the surface.

In our case, the captured data set contains two types of information. The image set $\{M_i\}$ captured through a microscope and their poses $\{U_i\}$ where they

are captured. Similar with the image stitching applied on 2D plane, this problem can be divided as the following stages:

Local tile registration. Estimating the transformations $\{T_{ij}\}$ between tile pairs M_i and M_j .

Global optimization. Find an alignment that minimize the error for the entire group.

Post processing. Generating the stitching result.

Local tile registration We define a microscopy image tile as a quad in the 3D space, which is a planar approximation of the certain part of the subject surface. Considering the depth of field of microscope is quite shallow, and the field of view is very narrow, the captured microscopy image shows a relatively flat area. We use microscopy image as the texture of a microscopy image tile M_i (which we refer as tile later, unless specifically explained).

The goal of this local registration is to estimate the transformation T_{ij} from the source tile M_i to the target tile M_j that aligns them in 3D space. Considering the microscopy images with overlapping area are close to each other on a smooth surface, T_{ij} is an affine transformation contains only translation and rotation, without scaling or shearing, and it is hard to calculate the accurate 3D affine transformation from perspective projection cause the difference between the normal directions of the two pairwise images can hardly be estimated.

After analyzing the captured microscopy image set and their associated camera poses recorded by the robotic arm sensor, we notice that the orientations extracted from the camera poses are considerably reliable.

Therefore, we decided to estimate the tangent direction component of T_{ij} from texture features, and get the normal direction component from robotic sensor data, as shown in figure 2.

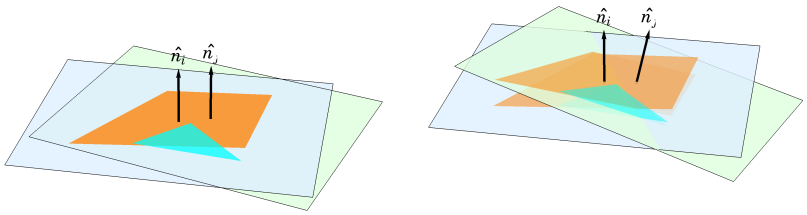

(a) planar local transformation. (b) adding normal direction difference.

Fig. 2: The local transformation composing. The normal direction difference is injected into the planar planar local transformation generated from 2D homography estimation via feature matching.

In most cases, the local transformation T_{ij} has almost no scaling or shearing and its matching confidence on tangent plane depends on the vision feature matching confidence. Therefore we can clip most of the false positive pairwise matches by applying a simple sanity check. All the confirmed local transformations are defined as $\{T_{ij}\}$. If a tile M_k does not have any confirmed adjacent images, it will be removed from $\{M_i\}$, as well as its initial pose U_k.

Given a microscopy data set consists of n microscopy images $\{M_i\}$ and the corresponding image poses $\{U_i\}$, a brutal pairwise matching would require to evaluate the match possibility between $N = \frac{n(n-1)}{2}$ potential image pairs, which wastes a great deal of calculation resource cause most of these microscopy images share no overlapping areas.

Besides, the features of microscopy images are rather different from these of landscape images, they might have similar patterns which might lead to false positive feature matching. Therefore, the pairwise matching ground truth of the N image pairs would be N_p positive matches and N_n negative matches, and the negative matches would outnumber the positive ones significantly, therefore we have

$$N_n >> N_p \tag{1}$$

The estimation results would be the mixture of N_{tp} true positive, N_{tn} true negative, N_{fp} false positive and N_{fn} false negative matching results, so we have

$$N = N_n + Np = (N_{fp} + N_{tn}) + (N_{tp} + N_{fn}) \tag{2}$$

and the registration accuracy:

$$a = \frac{N_{tp} + N_{tn}}{N} \tag{3}$$

Considering the true positive pairwise matches are outnumbered by the negative ones, it is more reasonable to evaluate the registration results with recall and precision.

Recall:

$$r = \frac{N_{tp}}{N_{tp} + N_{fn}} \tag{4}$$

Precision:

$$p = \frac{N_{tp}}{N_{tp} + N_{fp}} \tag{5}$$

A proper global registration method should be able to push precision to as high as possible while maintain acceptable high recall, an effective solution is reducing the false positive matching result amount N_{fp} or purging the registration results in the following optimization progress[9]. However, the lack of reliable odometry transformation and the discrete sampling points make it difficult to do so.

But on the other hand, the initial image poses $\{U_i\}$ are measured by the robotic arm sensors so they are much more reliable than these calculated from odometry chain in term of global position, though the local accuracy is not as

high as obometry transformation estimated by computer vision or other methods. So for each image in $\{M_i\}$, it is possible to find images with higher chance to overlap with it by their initial poses to reduce the false positive matches.

First we make a KD-tree from the initial image poses $\{U_i\}$. For certain microscope image M_i, we find the images $\{M_j\}$ that locating within a certain distance from it by searching in the KD-tree. Considering the adjacent images would have similar normal directions so the adjacent images should satisfy

$$(U_i\hat{n})^T U_j\hat{n} \geq 1 - \xi \tag{6}$$

in which the ξ is the tolerance of the normal direction difference. Now for each microscopy image M_i, there is a group of potential adjacent images $\{M_j\}$, the pairwise matching need only to be applied to these potential matching pairs.

Global optimization Global optimization is usually implemented by applying bundle adjustment for 2D image stitching. Considering we have the robotic sensor data as additional information and the alignment is made in 3D space, we choose to use pose graph for the global optimization.

The initial poses of tiles can be considered satisfying Gaussian distribution in the terms of the 3D positions, so it is reasonable to make use of it in the optimization progress too. We use g2o for the pose graph optimization[23].

In our modified graph-based optimisation, we define a pose graph G which has fixed sensor data nodes and free tile pose nodes, as shown in figure 3. For

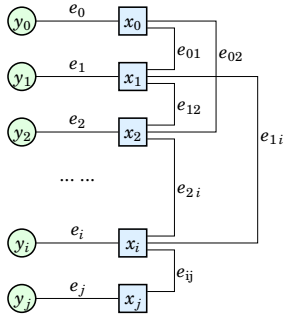

Fig. 3: The pose graph G in the global optimization progress. Nodes $\{x_i\}$ are the free tile nodes that would be optimized in the iteration and $\{y_i\}$ are the fixed nodes that record the initial poses of $\{x_i\}$.

each tile M_i, we add a free node x_i and a fixed node y_i in the pose graph, both of them have U_i as their initial pose. Also we add a edge e_i connecting x_i and y_i, which gives a constrain to keep the final optimized pose x_i close to its initial pose y_i. The transformation of edge e_i is set to I and it is defined as sensor edge. For each confirmed pairwise match, we add a matching edge e_{ij}, which connects nodes x_i and x_j and the edge transformation is set to be T_{ij}.

The goal of the optimization is to find a group of optimized poses $\{P_i\}$ for all tiles with confirmed adjacent image tiles, which can be described as minimizing the global objective

$$\begin{aligned} E(P) &= \eta \sum_i g(P_i, x_i) + (1-\eta) \sum_{i,j} h(P_i, P_j, x_{ij}) \\ &= \eta \sum_i f(P_i, U_i, I) + (1-\eta) \sum_{i,j} f(P_i, P_j, T_{ij}) \end{aligned} \tag{7}$$

in which g is the error function of sensor edges, h is the error function of matching edges and η is the confidence of the sensor data. They share similar structure so we define a function $f(P_i, P_j, T, \mu)$ to measure the inconsistency between poses P_i and P_j and the local transformation T with a parameter μ adjusting the weights of the alignment terms.

$$f(P_i, P_j, T, \mu) = f(T^{-1} P_j P_i^{-1}, \mu) \tag{8}$$

The inconsistency $\Delta = T^{-1} P_j P_i^{-1}$ is described by a loosen constrain $\delta = (a, b, c, \alpha, \beta, \gamma)$. (a, b, c) is the translation and (α, β, γ) is the Euler angles extracted via (x, y, z) order, which is widely used in 3D pose graph optimization. μ is used as a 6×6 information matrix in the optimization progress.

Considering the nodes are microscopy image tiles with flat geometry rather than dense point clouds and $\Delta \approx I$, we can use a simplified information matrix

$$\mu = \begin{bmatrix} w_a & 0 & 0 & 0 & 0 & 0 \\ 0 & w_b & 0 & 0 & 0 & 0 \\ 0 & 0 & w_c & 0 & 0 & 0 \\ 0 & 0 & 0 & w_\alpha & 0 & 0 \\ 0 & 0 & 0 & 0 & w_\beta & 0 \\ 0 & 0 & 0 & 0 & 0 & w_\gamma \end{bmatrix} \tag{9}$$

In matching edges, the vision feature matching provides a high alignment confidence on the tangent plane and the assumption that all adjacent image tiles are approximately coplanar, so the confidence at normal direction is also relatively high. On the other hand, sensor data provides high reliability in terms of the normal direction difference between tiles, but the rotation around normal direction is affected by the camera offset. Therefore we have

$$w_b = w_c = w_\alpha > w_a = w_\beta = w_\gamma \tag{10}$$

for matching edges $\{e_{ij}\}$,and

$$w_\beta = w_\gamma > w_a = w_b = w_c > w_\alpha \tag{11}$$

for sensor edges $\{e_i\}$. Considering the initial pose is relatively reliable, robust kernel is only applied to matching edges $\{e_{ij}\}$ to deal with the false positive pairwise matching results.

Post-processing Microscopes have narrow field of view and the illumination might be blocked by the microscope, which leads to intensity difference between image tiles especially when the illumination is not constant or the contrast of the subject surface is high. We apply a color correction following the method of Brown and Lowe [8] to address the color inconsistency.

To generate a triangle mesh we run a ball pivoting algorithm [7] for the center positions of all the aligned tiles. For each triangle of the generated mesh, the three microscopy images which are closest to the three triangle vertices are used to generate the texture by stitching these three images and cropping the trianglar area for texturing. Although the adjacent triangles have different textures, they appear to be seamless in the visualization.

4 Experiments

In the experiment, we use a Franka Emika robotic arm and a digital microscope for data capturing, and a stress ball with a smile face is used as the testing subject because it is rich of features and have a curved surface, as shown in Figure 4a. The raw microscopy image set (283 microscopy images) is visualized in Figure 4b.

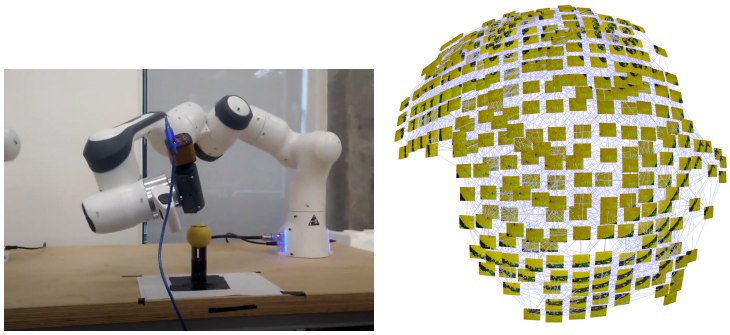

(a) Experiment setting. (b) Raw data.

Fig. 4: Experiment setup and the raw microscopy image set visualized in 3D. Noticing that the initial poses from the robotic arm sensor are not so accurate. The potential matching pairs are visualized as lines in 4b.

The final reconstruction result is shown in Figure 5a. We also reconstructed the same stress ball with Open3D from the 600 RGBD images captured through Realsense, the baseline result is shown in Figure 5b. We compare with Realsense cause its sensors have similar resolution with our microscope and it is an out-of-shelf solution from data capturing to reconstruction. For the area we captured, our result is considerably better than the reconstruction result of realsense.

In order to make a comparison, we also feed the image set we captured to open source tools like Bundler [13] and visualSFM [41], however they fail to give recognizable results.

(a) Our method. (b) Realsense & open3D.

Fig. 5: Part of the smile face stress ball and the comparison with the reconstruction result of Open3D from the RGBD sequence captured through Realsense.

Since there is no available microscopy data set captured on curved surface, we cannot evaluate the accuracy by comparing with the ground truth. So we evaluate our reconstruction result by rendering the virtual microscopy images from the camera poses associated with the real microscopy images, then comparing the real microscopy images and the synthesized microscopy images. (Figure 6)

For most camera poses, our reconstruction managed to preserve the details.

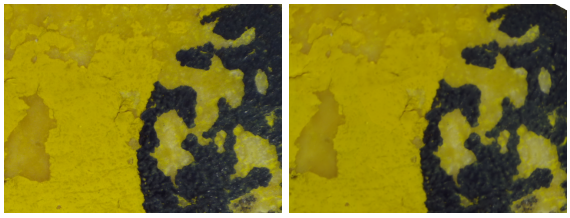

(a) original image. (b) recaptured image.

Fig. 6: An image with same resolution is captured from a similar camera pose to make the evaluation. Noticing that the specular light in the black area is eliminated.

5 Discussion

We managed to reconstruct the curved surface with high detail textures, but there are a lot of limitations in this pipeline, summarized as follow:

Unfocused microscopy images.

We apply auto focusing by moving the microscope forward and backward through the robotic arm and then pick the sharpest image, which is supposed to be the focused one, from the sequence. Sometimes, our auto-focus method output a clearly unfocused image. To do so, we use the method presented by [35], which computes a single floating point value to represent how blurry a given image is by calculating the variation of the Laplacian over the entire image, which we refer as

$$LAP\_VAR(I) = \sum_{m}^{M}\sum_{n}^{N}(|L(m,n) - \overline{L}|)^2 \tag{12}$$

where

$$\overline{L} = \frac{1}{NM}\sum_{m}^{M}\sum_{n}^{N}|L(m,n)| \tag{13}$$

Technically, this method works fine for synthesized focal stack sequence and real image sequence captured in regular circumstance, but an inevitable drawback is the $LAP\_VAR(I)$ value is not independent from the content. For instance, the $LAP\_VAR(I)$ of figure 7a is lower than that of figure 7b, but the former is obviously sharper than the latter. The reason why an unfocused image has higher $LAP\_VAR(I)$ is that it has a strong feature in its left part, the brown line which has huge effect on the Laplacian in that area.

(a) focused (b) unfocused (c) aligned sequence

Fig. 7: A microscopy image sequence captured in one auto-focusing action. Only the microscopy image with highest $LAP\_VAR(I)$ value is stored in a buffer and the saved as after the auto-focusing action is over in the practical running. 7a and 7b are microscopy images labelled with their $LAP\_VAR(I)$ value. 7c is the microscopy image sequence that aligned manually.

For a regular camera, the image sequence captured during the auto-focus action are almost perfectly aligned and the change of field can be ignored in most situation. However, our microscope does not have a focus motor so the image

sequence is captured while the robotic arm is moving the camera forward and backward. The shaking introduced by the robotic arm motion is considerably tiny for most regular use, but even a tiny disturbance is very obvious in the microscope camera because of its narrow field of view, which changes the content of the images significantly.

In theory, it is possible to address this problem by only evaluating the sharpness within the overlapping area of all the microscopy images in the sequence, as shown in Figure 7c. However, aligning the image sequence automatically requires reliable feature matching between images, but there is rarely recognizable feature in completely unfocused images.

Limited geometry complexity.

For our method, the amount of vertices of the final reconstructed textured mesh is equal with the amount of microscopy images, which is a drawback caused by the design of the processing pipeline. Additionally, the microscope can only be used on a relatively smooth surface, so it can hardly capture microscopy image set from a subject with complex geometry structure.

Missing sampling points.

In most cases, the robotic arm fails to move the microscope to all desired sampling points, as shown in Figure 8. By design the sampling points are supposed to be evenly distributed along the interpolated surface, but the robotic arm fails to capture microscopy images from some sampling points.

To figure out the reason why they are determined as unreachable by the robotic arm, we mark the missing sampling points and make a new sampling points list that consists of them only, run the automatic sampling again. Surprisingly, some of them become reachable and the robotic arm managed to capture microscopy images from these associated camera poses, and the reachable points vary each time we run the automatic sampling. We believe this is because the motion planer output is affected by the initial pose of the robotic arm. And some of them might be unreachable mechanically for the robotic arm due to its freedom limitation, since they remain unreachable in all our tests.

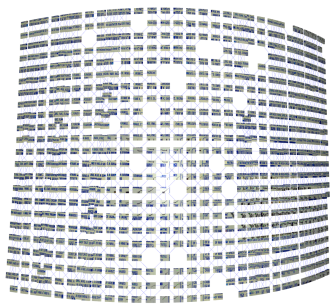

Fig. 8: A microscopy image set visualized in 3D. The microscopy image tiles are scaled for better view. The positive pairwise matches are visualized as lines. Noticing that there are holes in the image set

Although we regard the robotic arm as an universal solution for moving things around, they might have more limitations than we realized. Besides, it seems to be hard for the robotic arm to move its end effector for a short distance like $2mm$ and then stop.

6 Conclusion

In this research, we try to challenge a very specific problem in 3D reconstruction area, performing 3D reconstruction on images captured through microscopes. However, the short focal distance of microscopes leads to a quite narrow field of view, alongside with the shallow depth of field, image capturing for this task is quite different from that of regular photogrammetry, and the captured image set is also very unique. As a result, we have to design special robotic assisted image capturing method and a novel 3D reconstruction approach to address this problem all because of this specific camera choice.

Inspired by the level of detail (LOD) [29] loading strategy in computer graphics, we figured a possible compromise is applying typical macro photogrammetry to reconstruct a rough textured mesh, then generate sampling points from it to guide the robotic arm to capture a microscopic image set. At last, refine the texture of the existing mesh with the captured microscopic images. In theory, the geometry can be refined too. Mathematically, it is possible to repeat this refining progress with multiple zooming levels.

References

1. Acar, E.U., Choset, H.: Robust sensor-based coverage of unstructured environments. In: Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No.01CH37180). vol. 1, pp. 61–68 vol.1 (2001)
2. Acar, E.U., Choset, H., Rizzi, A.A., Atkar, P.N., Hull, D.: Morse decompositions for coverage tasks. The International Journal of Robotics Research **21**(4), 331–344 (2002)
3. Aloimonos, J.: Shape from texture. Biological cybernetics **58**(5), 345–360 (1988)
4. Atkar, P.N., Greenfield, A., Conner, D.C., Choset, H., Rizzi, A.A.: Uniform coverage of automotive surface patches. The International Journal of Robotics Research **24**(11), 883–898 (2005)
5. Basri, R., Jacobs, D., Kemelmacher, I.: Photometric stereo with general, unknown lighting. International Journal of computer vision **72**(3), 239–257 (2007)
6. Beardsley, P., Torr, P., Zisserman, A.: 3d model acquisition from extended image sequences. In: European conference on computer vision. pp. 683–695. Springer (1996)
7. Bernardini, F., Mittleman, J., Rushmeier, H., Silva, C., Taubin, G.: The ball-pivoting algorithm for surface reconstruction. IEEE transactions on visualization and computer graphics **5**(4), 349–359 (1999)
8. Brown, M., Lowe, D.G.: Automatic panoramic image stitching using invariant features. International journal of computer vision **74**(1), 59–73 (2007)

9. Choi, S., Zhou, Q.Y., Koltun, V.: Robust reconstruction of indoor scenes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5556–5565 (2015)
10. Choset, H., Lynch, K., Hutchinson, S., Kantor, G., Burgard, W., Kavraki, L., Thrun, S.: Principles of Robot Motion: Theory, Algorithms, and Implementations. MIT Press (2005)
11. Choset, H., Pignon, P.: Coverage path planning: The boustrophedon cellular decomposition. In: Field and Service Robotics. pp. 203–209. Springer London (1998)
12. Coleman, D., Sucan, I.A., Chitta, S., Correll, N.: Reducing the barrier to entry of complex robotic software: a moveit! case study. ArXiv **abs/1404.3785** (2014)
13. Dai, A., Nießner, M., Zollöfer, M., Izadi, S., Theobalt, C.: Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface re-integration. ACM Transactions on Graphics 2017 (TOG) (2017)
14. Davies, A.: Close-up and macro photography. CRC Press (2012)
15. Diankov, R.: Automated Construction of Robotic Manipulation Programs. Ph.D. thesis, Carnegie Mellon University, Robotics Institute (August 2010)
16. Evgenikou, V., Georgopoulos, A.: Investigating 3d reconstruction methods for small artifacts. International Archives of the Photogrammetry, Remote Sensing & Spatial Information Sciences (2015)
17. Fuchs, F., Koenig, A., Poppitz, D., Hahnel, S.: Application of macro photography in dental materials science. Journal of dentistry **102**, 103495 (2020)
18. Galceran, E., Carreras, M.: A survey on coverage path planning for robotics. Robotics and Autonomous Systems **61**(12), 1258 – 1276 (2013)
19. Gallo, A., Muzzupappa, M., Bruno, F.: 3d reconstruction of small sized objects from a sequence of multi-focused images. Journal of Cultural Heritage **15**(2), 173–182 (2014)
20. Ivan, V., Yang, Y., Merkt, W., Camilleri, M.P., Vijayakumar, S.: EXOTica: An Extensible Optimization Toolset for Prototyping and Benchmarking Motion Planning and Control, pp. 211–240. Springer International Publishing (2019)
21. Jared, H., Schonberger, J.L., Dunn, E., Frahm, J.M.: Reconstructing the world in six days. In: CVPR (2015)
22. Kavraki, L.E., Svestka, P., Latombe, J.., Overmars, M.H.: Probabilistic roadmaps for path planning in high-dimensional configuration spaces. IEEE Transactions on Robotics and Automation **12**(4), 566–580 (1996)
23. Kümmerle, R., Grisetti, G., Strasdat, H., Konolige, K., Burgard, W.: g 2 o: A general framework for graph optimization. In: 2011 IEEE International Conference on Robotics and Automation. pp. 3607–3613. IEEE (2011)
24. Latombe, J.C.: Robot Motion Planning. Kluwer Academic Publishers, USA (1991)
25. LaValle, S.M.: Planning Algorithms. Cambridge University Press, Cambridge, U.K. (2006)
26. Levoy, M., Pulli, K., Curless, B., Rusinkiewicz, S., Koller, D., Pereira, L., Ginzton, M., Anderson, S., Davis, J., Ginsberg, J., et al.: The digital michelangelo project: 3d scanning of large statues. In: Proceedings of the 27th annual conference on Computer graphics and interactive techniques. pp. 131–144 (2000)
27. Li, M.Z., Lu, Z.P., Sha, C.F., Huang, L.Q.: Trajectory generation of spray painting robot using point cloud slicing. In: Frontiers of Manufacturing and Design Science. Applied Mechanics and Materials, vol. 44, pp. 1290–1294. Trans Tech Publications Ltd (2 2011)
28. Lindemann, S.R., LaValle, S.M.: Current issues in sampling-based motion planning. In: Robotics Research. The Eleventh International Symposium. pp. 36–54 (2005)

29. Luebke, D., Reddy, M., Cohen, J.D., Varshney, A., Watson, B., Huebner, R.: Level of detail for 3D graphics. Morgan Kaufmann (2003)
30. Lynch, K.M., Park, F.C.: Modern Robotics: Mechanics, Planning, and Control. 1st edn. (2017)
31. Molina, J., Solanes, J.E., Arnal, L., Tornero, J.: On the detection of defects on specular car body surfaces. Robotics and Computer-Integrated Manufacturing **48**, 263 – 278 (2017)
32. Newcombe, R.A., Izadi, S., Hilliges, O., Molyneaux, D., Kim, D., Davison, A.J., Kohi, P., Shotton, J., Hodges, S., Fitzgibbon, A.: Kinectfusion: Real-time dense surface mapping and tracking. In: 2011 10th IEEE International Symposium on Mixed and Augmented Reality. pp. 127–136. IEEE (2011)
33. Or-El, R., Rosman, G., Wetzler, A., Kimmel, R., Bruckstein, A.M.: Rgbd-fusion: Real-time high precision depth recovery. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5407–5416 (2015)
34. Overmars, M.H.: Recent developments in motion planning. In: Computational Science — ICCS 2002. pp. 3–13. Berlin, Heidelberg (2002)
35. Pech-Pacheco, J.L., Cristóbal, G., Chamorro-Martinez, J., Fernández-Valdivia, J.: Diatom autofocusing in brightfield microscopy: a comparative study. In: Proceedings 15th International Conference on Pattern Recognition. ICPR-2000. vol. 3, pp. 314–317. IEEE (2000)
36. Rohmer, E., Singh, S.P.N., Freese, M.: Coppeliasim (formerly v-rep): a versatile and scalable robot simulation framework. In: Proc. of The International Conference on Intelligent Robots and Systems (IROS) (2013)
37. Schaffalitzky, F., Zisserman, A.: Multi-view matching for unordered image sets, or "how do i organize my holiday snaps?". In: European conference on computer vision. pp. 414–431. Springer (2002)
38. Schwartz, J., Sharir, M.: A survey of motion planning and related geometric algorithms. Artificial Intelligence **37**(1), 157 – 169 (1988)
39. Siciliano, B., Khatib, O.: Springer Handbook of Robotics. 2nd edn. (2016)
40. Tian, F., Li, Z., Lv, C., Liu, G.: Polishing pressure investigations of robot automatic polishing on curved surfaces. The International Journal of Advanced Manufacturing Technology **87**(1-4), 639–646 (2016)
41. Wu, C., et al.: Visualsfm: A visual structure from motion system (2011)
42. Yang, Y., Merkt, W., Ivan, V., Li, Z., Vijayakumar, S.: Hdrm: A resolution complete dynamic roadmap for real-time motion planning in complex scenes. IEEE Robotics and Automation Letters **3**(1), 551–558 (2018). https://doi.org/10.1109/LRA.2017.2773669
43. Zhang, R., Tsai, P.S., Cryer, J.E., Shah, M.: Shape-from-shading: a survey. IEEE transactions on pattern analysis and machine intelligence **21**(8), 690–706 (1999)