Towards visual feedback loops for robot-controlled additive manufacturing

Abstract. Robotic additive manufacturing methods have enabled the design and fabrication of novel forms and material systems that represent an important step forward for architectural fabrication. However, a common problem in additive manufacturing is to predict and incorporate the dynamic behavior of the material that is the result of the complex confluence of forces and material properties that occur during fabrication. While there have been some approaches towards verification systems, to date most robotic additive manufacturing processes lack verification to ensure deposition accuracy. Inaccuracies, or in some instances critical errors, can occur due to robot dynamics, material self-deflection, material coiling, or timing shifts in the case of multi-material prints. This paper addresses that gap by presenting an approach that uses vision-based sensing systems to assist robotic additive manufacturing processes. Using online image analysis techniques, occupancy maps can be created and updated during the fabrication process to document the actual position of the previously deposited material. This development is an intermediary step towards closed-loop robotic control systems that combine workspace sensing capabilities with decision-making algorithms to adjust toolpaths to correct for errors or inaccuracies if necessary. The occupancy grid map provides a complete representation of the print that can be analyzed to determine various key aspects, such as, print quality, extrusion diameter, adhesion between printed parts, and intersections within the meshes. This valuable quantitative information regarding system robustness can be used to influence the system's future actions. This approach will help ensure consistent print quality and sound tectonics in robotic additive manufacturing processes, improving on current techniques and extending the possibilities of robotic fabrication in architecture.

Keywords: Robot control • 3D printing • Vision-based sensing

1 Introduction

With the advent of new, robotically enabled fabrication methods comes the ability to design new materials and geometries with new functionalities. Additive manufacturing methods have been rapidly developed for several new materials from the starting point of thermoplastics to silicones (Rodrigue et al. 2015), concrete (Lloret et al. 2015), and hydrogels (Barry et al. 2009), and many more. Approaches to robotic 3D printing have also started to go beyond in-plane, 2D layer-based methods towards freeform 3D material depositions (Hack and Lauer 2014; Laarman et al. 2014; Soler et al. 2017). These advancements bring about the ability to design not only novel forms, but also new performative material systems that produce variations according to local stresses or to respond to environmental conditions. These new material systems are possible in mono-material additive manufacturing and their potentials are greatly expanded with a move towards multi-material 3D printing. However, to fully leverage these functional capacities, a greater control of the material deposition is required.

Additive manufacturing processes comprise a complex ecology of interactions between a diverse set of parameters, including, but are not limited to the rheological characteristics of the material, the rate of material flow from the extrusion nozzle, the rate of cooling or curing of the material, and the structural capacities of the pre-cooled or -cured material as additional material is deposited. The complexity and non-linear interactions of these parameters make it difficult to predict and ensure the accurate deposition of material during the 3D printing process without extensive and computationally intensive simulations of the entire process prior to commencing the print.

In order to accommodate some of these parameters and overcome the challenges several researchers have taken an open-loop approach of explicit tool-pathing and control of the robotic processes. Hack et al. (2013) implemented a series of explicit robot motions including amplifying the z-direction movements, short stops for cooling, and air pressure changes to increase the level of control and predictability over the material behavior. McGee et al. (2017) implemented an explicit “pressing” motion path with additional tolerances embedded within it to ensure fully-fused joints between all 3D printed connections in a tensile mesh.

Approaches to robotic systems can be divided into two categories; open-loop systems and closed-loop systems (Vidal-Calleja et al. 2010). To a large extent, open-loop systems are the predominantly used approach in the field of architecture and design. As an alternative to the previously described open-loop approaches, there exists an opportunity to equip the robotic system itself with sensing capabilities and decision-making agency (Paul et al. 2009) to ascertain the accuracy of previously deposited material and adjust future toolpaths (Paul et al. 2015). Research suggests that closed-loop systems, are critical to the success of robotic fabrication processes in which the material does not behave predictably or in situations where conditions can change or lead to inaccuracies (Giftthaler et al. 2017). In robotic additive manufacturing processes, surveying the position of the previously deposited material, would give the system agency to take corrective action to these inaccuracies by recalibrating robot

trajectories and accurately executing tectonic motion paths like the pressing motion developed previously.

Vision-based feedback systems are generally classified depending on the number of cameras, the position of the camera with respect to the robot or the design of the error minimization function used to control the robot (Hutchinson et al. 1996). When considering the position of the camera, two main configurations exist; end-effector mounted (eye-in-hand) or fixed in the workspace (bird's eye). Feedback approaches are further classified into position based, image based, hybrid based and motion-based feedback systems. In a position-based feedback system, the online analysis of the scene is performed in what is referred to as the task space, whereby image information is used to reconstruct or map the scene via the a priori calibrated camera model (Zhang 1999). This provides geometric information of the scene in the industrial robot's base coordinates.

To leverage geometric information of the scene, metric maps are utilized to build accurate representations of the environment. An algorithm proposed by Elfes (1989) and Moravec (1988) known as Occupancy Grid (OG) mapping, utilizes grids to model the environment's free and occupied space. In 3D grid map representations, techniques such as voxel hashing and octrees discretize space into voxels. Each voxel possesses a position in space and a probability of occupancy; occupied, free-space or unknown (Paul et al. 2015). It is common to combine probabilistic approaches in conjunction with mapping methods as sensors are subject to measurement noise. Algorithms such as extended Kalman filters, Bayesian filters, and particle filters have been widely explored to provide improved state estimates (Thrun et al. 2005).

This paper proposes a framework for online visual-feedback in robotic additive manufacturing processes (shown in Fig. 2) via image analysis and probabilistic OG-maps (Elfes 1989). The framework sequentially takes an image as its input, segments the current area, utilizes the *a priori* calibrated camera model to build a map in task space and probabilistically fuses the information into an overall OG map of the printed areas.

2 Methodology

The research described here seeks to verify the accuracy of previously deposited material during additive manufacturing processes. The design process generates a set of task space waypoints for the Tool Center Point (TCP) to track. The robot is controlled so that it tracks the trajectory while images are acquired via a calibrated camera. As shown in Fig. 4, gathered images are processed such that the deposited material is segmented from the background. The classified images are transformed into task space, and sequentially fused into a probabilistic OG map.

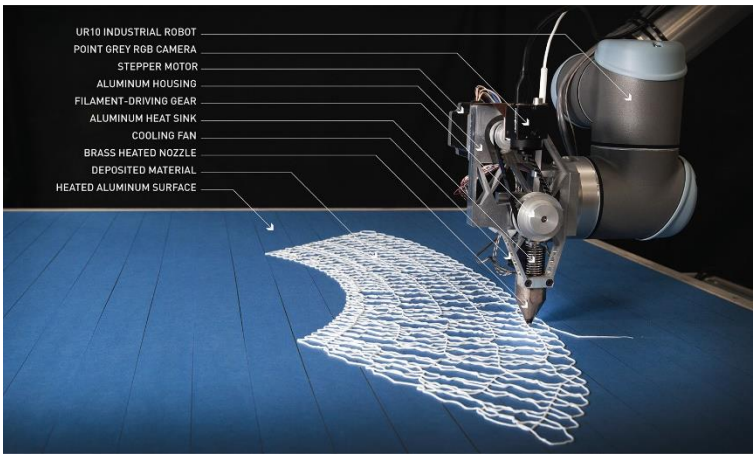

Fig. 1. UR10 robot with wrist-mounted filament extruder system, Point Grey Blackfly RGB camera and heat bed (in blue) while printing part B30.

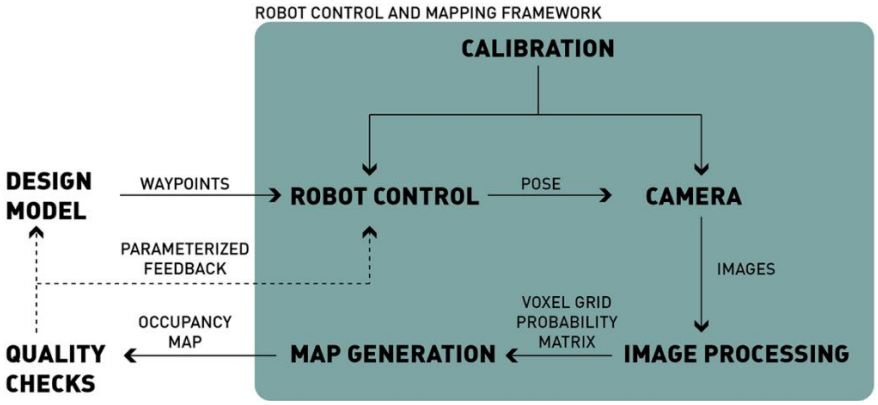

Fig. 2. Overview of the proposed robot control and mapping framework for closed-loop additive manufacturing processes.

2.1 Robot Eye-to-Extruder Calibration

The eye-in-hand configuration proposed in this work possesses a camera located on the filament extruder system as shown in Fig. 1. As the TCP moves through a given trajectory, the images gathered from the camera are unable to be related unless the transformation between the camera and robot base is known.

Accurate geometric information of the environment obtained from the camera requires two types of calibration, intrinsic and extrinsic. Intrinsic calibration determines the intrinsic matrix that represents the projective transformation from the 3D camera coordinate system into a 2D image coordinate system. This calibration utilizes the perspective projection camera model and procedure proposed by Zhang (1999) and

implemented by Bouguet (2000). Extrinsic calibration is used to determine the camera's pose (i.e. rotation and translation) in a defined coordinate system, or robot base, ${}^{0}T_{c}$.

Hand-eye calibration is the process that ascertains the camera coordinate system relative to the reference frame of the robot's TCP. Well-known approaches can be based on determining the rotation and translation consecutively (Daniilidis 1999), simultaneously (Strobl and Hirzinger 2006), or considering time-offsets (Furrer et al. 2018).

The hand-eye calibration implemented in this work requires a set of known points, $P \in \mathrm{R}^{3}$, defined in robot base coordinates and corresponding points defined in camera coordinates. Determining the rotation and translation can be formulated as an optimization problem. As shown in Fig. 3, a calibration target with known dimensions is needed to obtain the 3D points in camera coordinates from the points in the 2D image. To recover the points defined in camera coordinates requires P to be visible from known camera location.

Given the homogeneous transformation describing the base location, ${}^{0}T_{b}$ of the 6 degree-of-freedom robot with joint angles as a vector, $\mathbf{q} = [q_1, \dots q_6]^T$, the last link location can be computed using forward kinematics, ${}^{b}T_{f}(\mathbf{q})$. Consequently, the addition of an end-effector, i.e. the extruder system, requires a known transformation of the TCP relative to the last link, ${}^{f}T_{n}$. Thus, it is possible to determine the location of the TCP relative to the base location of the robot using, ${}^{0}T_{n}(\mathbf{q}) = {}^{0}T_{b}\,{}^{b}T_{f}(\mathbf{q})\,{}^{f}T_{n}$. Moving the TCP to P, yields a series of known transformations in robot base coordinates.

The translation is the difference between the two sets of points and incorporates the optimal rotation obtained via singular-value decomposition. Given ${}^{0}T_{n}(\mathbf{q})$ and the calculated transformation ${}^{0}T_{c}$, the camera relative to the TCP, ${}^{n}T_{c}$, can be obtained.

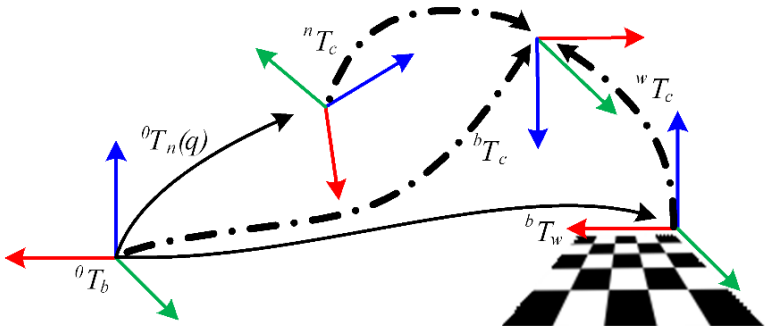

Fig. 3. Eye-in-hand calibration to determine the extrinsic calibration of a camera so the pixels in the camera coordinate frame can be transformed to world or robot base coordinate frame.

2.2 Deposited Material Segmentation via Image Processing

To segment the image into deposited material and background, the image undergoes a series of image processing steps. Initially, the image undergoes a color space conversion. Since the RGB color space suffers from variations in light intensity, the HSV color space is used as data represented in HSV is less susceptible to these changes. A threshold is applied to the image, then normalized to effectively occupy its range of pixel intensity values. Median filtering is applied to the image to remove noise from normalization. The image is then converted to a binary representation and undergoes median filtering again to remove noise that may have been a product of conversion.

The output of the process is a segmented image that categorizes the original values from the data into deposited material or background, ultimately represented as a binary matrix. This output is then fused into one probability map for further analysis asdiscussed in the next section.

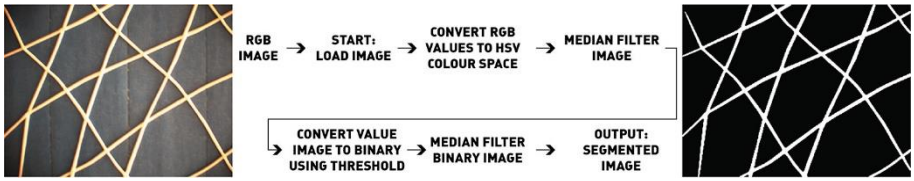

Fig. 4. Image processing procedure, takes in color images, processes them and then outputs the classified deposited material and background.

2.3 Deposited Material Segmentation via Image Processing

Mapping in robotics is the process of building a representation of the environment, commonly based on sensors. Given the intrinsic calibration and transformation, 0Tc (Sect. 2.1), images can be sequentially fused over time to infer whether a location in task space is occupied by deposited material. A Deposited Material Occupancy Map (DMOM) is used to represent the printed space.

The images obtained for fusion are a mapping of 3D points to a 2D surface, forming a 2D representation of the environment. Since this process involves the loss of depth information, depth to the printing surface can be obtained as the camera pose is known. Extending this approach to 3D map representations like octrees can be achieved using the method described by Paul et al. (2015).

The Occupancy Grid (OG) mapping method addresses the problem of generating a consistent map from noisy or incomplete sensor data and possesses potential use for data fusion (Stepan et al. 2005). As the eye-in-hand configuration limits the field of view of the camera, the probabilistic nature of OG mapping allows the images captured to be fused. Each grid cell in the OG map, DMOM in this case, is individually treated as having a mutually-exclusive probability of existence of deposited material. Since there is overlap in the data observed by the camera, the certainty of deposited material occupying a grid cell is dependent on whether deposited material is observed

at a location, and the frequency it appears at that location. The probability is updated using a Bayes update (Thrun et al. 2005), where higher trust is given to more recent measurements. The probability that a grid cell is occupied is independent of the location of the image pixel and distance from the camera. Due to this independence, it is assumed that at each point that is classified, regardless if it contains deposited material or not, is equally trustworthy.

3 Experiment setup

The approach is tested in a pilot study towards the production of a site specific, partially dynamic tensile surface installation covering over 25m^2 as shown in Fig. 5. The installation was designed and simulated using the Kangaroo Physics plug-in for Grasshopper, developed by Piker (2013). The overall form of the installation is divided into 59 panels to facilitate the printing of each panel within the reach of the robot.

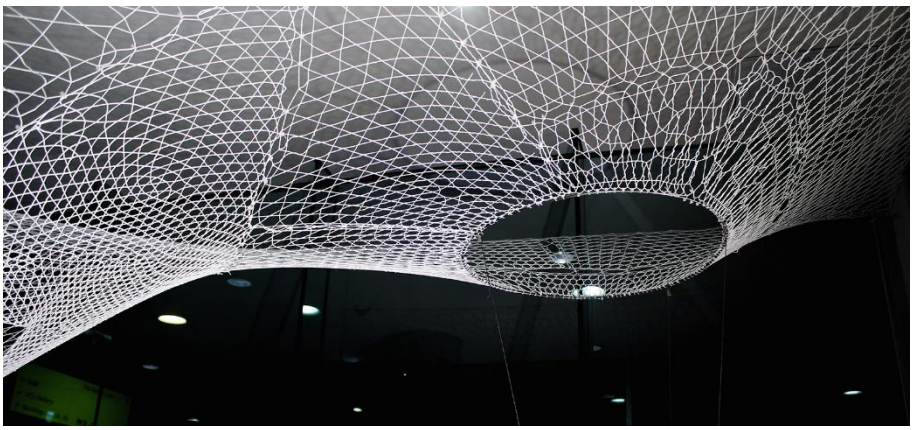

Fig. 5. Detail of robotic 3D printed functionally graded net installed on site

The topology of the tensile mesh is functionally graded to respond to the simulated differences in tension forces across the surface. By employing a bespoke computational method, the resulting graded mesh is flattened to embed the 3D geometry within the 2D pattern for printing McGee et al. (2017).

3.1 Physical Hardware Setup

To facilitate the manufacturing of the tensile meshes, a custom work cell was constructed. The setup consists of a Universal Robot (UR10) robot mounted to a frame housing a custom heat bed measuring 1:2m2. The extruder end-effector consists of a stepper motor that feeds filament to the nozzle, where it is melted by two heat cartridges embedded in the nozzle, with a thermocouple for temperature regulation.

A downward-facing global shutter Point Grey Blackfly RGB camera is mounted to the extruder. The resolution of the camera is 648 × 488 pixels, which results in an

area of 138 × 104 mm when the print nozzle is in contact with the heat bed. The offset position of the camera relative to the TCP, ${}^{n}T_{c}$ was estimated as [−43, −72, 187] mm.

3.2 Software

The toolpaths are developed from the form-found digital mesh output from the Kangaroo physics system. These meshes are translated into the final, continuous line toolpaths through a series of scripts which eliminate odd-valence vertices by replacing each edge in the mesh with a vertex, making connections between the mesh faces that share the original edge, a process which always results in an even-valence at the new vertices. The toolpaths are also constructed to ensure that the order of lines produces full cross over joints at each intersection for improved structural performance, rather than joints where the two polyline segments attempt to meet at a point which are failure prone.

The UR10 toolpath for each panel is initially solved through Robots Grasshopper plug-in. A list of target joint angles and their corresponding TCP locations are produced from the plug-in and passed to ROS to communicate with the robot. Additionally, the commands are distributed through ROS to overcome the limited number of points that can be sent through the Grasshopper plug-in. ROS enables the state of the hardware to be queried at 125 Hz providing a means to have positional feedback, whilst the images are acquired at 7 Hz and processed in MATLAB

3.3 Measurement Test using Generated Map

To test the validity of the map, quality experiments have been conducted including: intersection detection, filament thickness, and alignment assurance.

The detection of intersections between lines is necessary so that the additional "pressing" process can occur to create structurally sound joints at each crossing. To perform intersection detection, a point in the image is found and denoted as the start of the line. From this point, a series of smaller image segments are generated in order to locate the lines in 3D space. The result is a vector between a start and end point that can be iteratively checked to determine if there are intersections with any of the previous lines found before.

The thickness of the extruded filament affects the strength of the overall print, and should remain consistent throughout the print. However, unintended variances can occur due to issues with extruder speed or filament feed, heating element inconsistency, or imperfections in the original filament. The map can hence be analyzed to check that the extruded filament line width is within expected tolerances.

During the printing process, it is possible that the extruded filament fails to fully adhere to the heat bed surface, resulting in misaligned or otherwise incorrect geometries. To detect if there are any misalignment issues and to incorporate quality control, the map can be compared against the expected location of the print in the map. If

there are large variances, particularly global orientation differences, then it is likely there is an alignment issue and the print should be aborted and the user notified.

4 Results

Several experiments have been conducted to collect data during the printing process. The intersection points are taken from the geometry generation process and form the baseline ground truth from which the observed intersection points can be compared. Varying illumination sources on the printed filament and background presented challenging conditions for developing a robust method of detection. Two image segmentation results are shown in Fig. 6. The lighting in the room was a mixture of natural and artificial light that was not specifically controlled.

The proposed approach demonstrates the segmentation of RGB images into binary representations of deposited material and printing surface. However, it has been shown that the process is unable to distinguish between the current print and residual filament from previous prints due to similarities in color properties. Noise from varying illumination sources challenges the robustness of the image processing procedure. Thus, probabilistic fusion during map generation allows each pixel to possess a better estimate of the state (as areas of the print are observed multiple times).

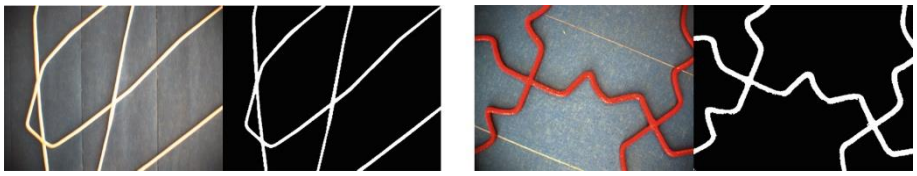

Fig. 6. Given two input images of deposited materials: cream-colored and red with varying lighting conditions on the background, the image processing results are shown.

A series of images shown in Fig. 7 demonstrate the progression of the print based upon the observed data fused by using the Bayes update. Each grid cell in the DMOM represents 0.2132 mm^2, with each possessing an intensity ranging from a maximum value of 1 (white) and a minimum value of 0 (black), and the initial state and final state of unseen grid cells possess a value of 0.5 (gray). The DMOM after 35 min of printing and 1575 processed images is shown in Fig. 7. As more images are fused, the certainty about the existence of printed material at that location is increased, converging upon known to be empty (i.e. a value of 0), or known to be occupied (i.e. a value of 1).

The printed piece shown in Fig. 8 contains a total of 466 intersections, with 15.88% of the intersections considered as unseen due to the trajectory, 5.15% of the intersections considered as unseen due to misalignment and 7.73% of the intersections considered as seen, although incorrectly mapped. Since the eye-in-hand configuration relies on the given trajectory to observe the print, Fig. 8 illustrates an incomplete DMOM shown with missing, incorrect and misaligned intersection points, due to incorrect mapping and unobserved locations.

The average thickness of the print is determined using the probabilistic OG map. A set of print thicknesses were obtained by randomly sampling 30 locations of the print, then measuring the thickness. The average thickness is 9.03 grid cells with a standard deviation of 1.20 grid cells, translating to 1.93 mm and 0.26 mm respectively.

Applying a binary classification test, the statistical measures of performance: accuracy, sensitivity, and specificity are obtained by comparing the classified pixels to the ground truth. A true (or false) positive occurs when a grid cell is correctly (or incorrectly) identified as containing; a true (or false) negative occurs when a grid cell is correctly (or incorrectly) identified as not containing deposited material. The accuracy of the print is 91.8%, indicating a high percentage of correctly classified grid cells. The sensitivity (40.2%) indicates a relatively weak ability to correctly identify grid cells that contain deposited material, and the specificity is 95.8% which indicates a strong ability to identify grid cells that are free of deposited material.

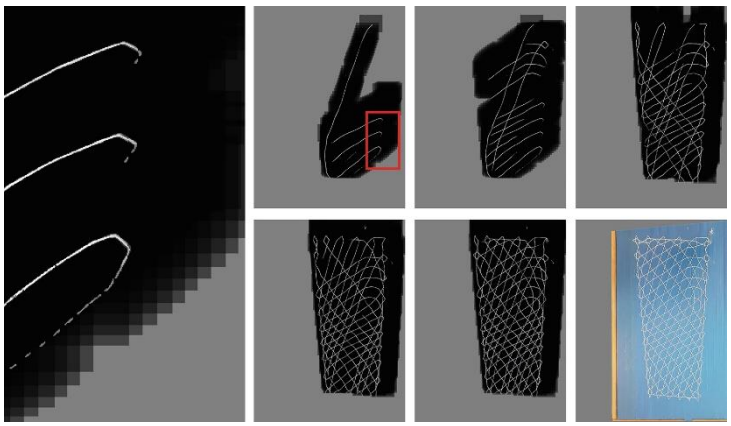

Fig. 7. Part A28. Results of a series of extrusion classifications as they are fused together into the task space DMOM. (Left) Detail of DMOM resolution. (Right, Top row from left) After 225 images and 5 min of printing; After 450 images and 10 min of printing; After 900 images and 20 min printing; (Right, Bottom row from left) After 1350 images and 30 min printing; Final DMOM; the actual complete print.

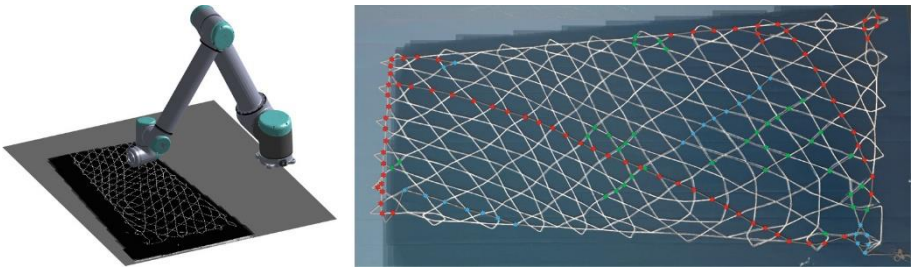

Fig. 8. (Left) Simulation of the robot with the final DMOM; (Right) Actual photo of final print overlaid with top-down view of the DMOM. Red markers indicate unseen intersections due to camera field of view and trajectory, blue markers indicate unseen intersections due to misalignment, green markers indicate detected intersections at incorrect locations due to misalignment.

5 Discussions and Conclusions

This paper has presented an approach to online feedback using a map-building technique that utilizes data from the robot and camera. The system probabilistically fuses the obtained data to construct an overall map of the print during the additive manufacturing process, enabling quality assurance processes to be applied. Throughout the research, the team has gained several insights towards this process. First, vibration of the extruder and camera system were noted during the extrusion process due to robot dynamics. This produces slight inaccuracies and misalignments with the camera's captured images but has no noticeable effect on the print quality. Future work is to close the control feedback loop by triggering pre-programmed localized actions, such as the pressing motion necessary for the tensile mesh extrusion when the system predicts an upcoming intersection.

References

Barry, R.A., Shepherd, R.F., Hanson, J.N., Nuzzo, R.G., Wiltzius, P., Lewis, J.A.: Direct-write assembly of 3D hydrogel scaffolds for guided cell growth. Adv. Mater. 21, 2407–2410 (2009)

Bouguet, J.Y.: Matlab camera calibration toolbox. Caltech Technical report (2000)

Daniilidis, K.: Hand-eye calibration using dual quaternions. Int. J. Robot. Res. 18(3), 286 298 (1999)

Elfes, A.: Using occupancy grids for mobile robot perception and navigation. Computer 22, 46–57 (1989)

Furrer, F., Fehr, M., Novkovic, T., Sommer, H., Gilitschenski, I., Siegwart, R.: Evaluation of combined time-offset estimation and hand-eye calibration on robotic datasets. In: Field and Service Robotics, pp. 145–159 (2018)

Giftthaler, M., Sandy, T., Dörfler, K., Brooks, I., Buckingham, M., Rey, G., Kohler, M.,Gramazio, F., Buchli, J.: Mobile robotic fabrication at 1:1 scale: the In situ Fabricator. Constr. Robot. 1(1–4), 3–14 (2017)

Hack, N., Lauer, W.V.: Mesh mould: robotically fabricated spatial meshes as reinforced concrete formwork. Archit. Des. 84(3), 44–53 (2014)

Hack, N., Lauer, W.V., Langenberg, S., Gramazio, F., Kohler, M.: Overcoming repetition: robotic fabrication processes at a large scale. IJAC 3(11), 285–299 (2013)
Hutchinson, S., Hager, G.D., Corke, P.I.: A tutorial on visual servo control. IEEE Trans. Robot. Autom. 12(1), 649–774 (1996)
Laarman, J., Jokic, S., Novikov, P., Fraguada, L.E., Markopoulou, A.: Anti-gravity additive manufacturing. In: Gramazio, F., Kohler, M., Langenberg, S. (eds.) Fabricate 2014: Negotiating Design & Making, pp. 192–197. UCL Press, London (2014)
Lloret, E., Shahab, A.R., Linus, M., Flatt, R., Gramazio, F., Kohler, M., Langenberg, S.: Complex concrete structures: merging existing casting techniques with digital fabrication. Comput. Aided Des. 60, 40–49 (2015)
McGee, W., Thun, G., Velikov, K., Tish, D.: Infundibuliforms: kinetic systems, additive manufacturing, and tensile surface control. In: Sheil, B., Menges, A., Glynn, R., Skavara, M. (eds.) Fabricate 2017, pp. 84–91. UCL Press, London (2017)
Moravec, H.: Sensor fusion in certainty grids for mobile robots. AI Mag. 9(2), 61–74 (1988)
Paul, G., Kirchner, N., Liu, D.K., Dissanayake, G.: An effective approach to simultaneous mapping and surface-type identification of complex 3D environments. J. Field Robot. 26(11–12 SI), 915–933 (2009)
Paul, G., Quin, P., To, A.W.K., Liu, D.K.: A sliding window approach to exploration for 3D map building using a biologically inspired bridge inspection robot. In: Proceedings of IEEE Cyber Technology in Automation, Control, and Intelligent Systems, pp. 1097–1102 (2015)
Piker, D.: Kangaroo: form finding with computational physics. Archit. Des. 83(2), 136–137 (2013)
Rodrigue, H., Bhandari, B., Wang, W., Ahn, S.H.: 3D soft lithography: a fabrication process for thermocurable polymers. J. Mater. Process. Technol. 217, 302–309 (2015)
Soler, V., Retsin, G., Jimenez Garcia, M.: A generalized approach to non-layered fused filament fabrication. In: Nagakura, T., Tibbits, S., Mueller, C., Ibanez, M. (eds.) ACADIA 2017 Disciplines Disruption, pp. 562–571 (2017)
Stepan, P., Kulich, M., Preucil, L.: Robust data fusion with occupancy grid. IEEE Trans. Syst. Man Cybern. Part C 35(1), 106–115 (2005)
Strobl, K.H., Hirzinger, G.: Optimal hand-eye calibration. In: 2006 IEEE/RSJ International-Conference on Intelligent Robots and Systems, pp. 4647–4653 (2006)
Thrun, S., Burgard, W., Fox, D.: Probabilistic Robotics (Intelligent Robotics and Autonomous Agents). The MIT Press, Cambridge (2005)
Vidal-Calleja, T., Andrade-Cetto, J., Sanfeliu, A.: Action selection for single camera SLAM. IEEE Trans. Syst. Man Cybern. Part B 40(6), 1567–1581 (2010)
Zhang, Z.: Flexible camera calibration by viewing a plane from unknown orientations. In: Proceedings of IEEE International Conference on Computer Vision, Kerkyra, vol. 1, pp. 666–673 (1999)