**Autonomous Scene Understanding, Motion Planning, and Task Execution
for Geometrically Adaptive Robotized Construction Work**

by

Kurt Michael Lundeen

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Civil Engineering)
in the University of Michigan
2019

Doctoral Committee:

Professor Vineet R. Kamat, Chair
Assistant Professor Clive R. D'Souza
Assistant Professor Branko Kerkez
Associate Professor SangHyun Lee
Assistant Professor Wes McGee
Associate Professor Carol C. Menassa

Kurt Michael Lundeen

klundeen@umich.edu

ORCID iD: 0000-0002-5190-4224

Table of Contents

List of Tables

List of Figures

Abstract

The construction industry suffers from such problems as high cost, poor quality, prolonged duration, and substandard safety. Robots have the potential to help alleviate such problems by becoming construction co-workers, yet they are seldom found operating on today's construction sites. This is primarily due to the industry's unstructured nature, substantial scale, and loose tolerances, which present additional challenges for robot operation. To help construction robots overcome such challenges and begin functioning as useful partners in human-robot construction teams, this research focuses on advancing two fundamental capabilities: enabling a robot to determine where it is located as it moves about a construction site, and enabling it to determine the actual pose and geometry of its workpieces so it can adapt its work plan and perform work.

Specifically, this research first explores the use of a camera-marker sensor system for construction robot localization. To provide a mobile construction robot with the ability to estimate its own pose, a camera-marker sensor system was developed that is affordable, reconfigurable, and functional in GNSS-denied locations, such as urban areas and indoors. Excavation was used as a case study construction activity, where bucket tooth pose served as the key point of interest. The sensor system underwent several iterations of design and testing, and was found capable of estimating bucket tooth position with centimeter-level accuracy.

This research also explores a framework to enable a construction robot to leverage its sensors and Building Information Model (BIM) to perceive and autonomously model the actual pose and

geometry of its workpieces. Autonomous motion planning and execution methods were also developed and incorporated into the adaptive framework to enable a robot to adapt its work plan to the circumstances it encounters and perform work. The adaptive framework was implemented on a real robot and evaluated using joint filling as a case study construction task. The robot was found capable of identifying the true pose and geometry of a construction joint with an accuracy of 0.11 millimeters and 1.1 degrees. The robot also demonstrated the ability to autonomously adapt its work plan and successfully fill the joint.

In all, this research is expected to serve as a basis for enabling robots to function more effectively in challenging construction environments. In particular, this work focuses on enabling robots to operate with greater functionality and versatility using methods that are generalizable to a range of construction activities. This research establishes the foundational blocks needed for humans and robots to leverage their respective strengths and function together as effective construction partners, which will lead to ubiquitous collaborative human-robot teams operating on actual construction sites, and ultimately bring the industry closer to realizing the extensive benefits of robotics.

Chapter 1
Introduction

Modern civilization is dependent upon built environments, such as buildings, roads, and utilities, but built environments are time consuming and dangerous to construct and maintain. Robots have the potential to help construct our environments at a lower cost, in a shorter timeframe, with higher quality, and with improved safety, but technological challenges have thus far thwarted such possibilities. This research seeks to address the fundamental challenges impeding robots from functioning as useful co-workers on today's construction sites. Specifically, this research seeks to enable construction robots to localize within their environment, perceive and model construction objects, adapt their work plans to the circumstances they encounter, and manipulate construction objects to perform useful construction work. An overview of the research is depicted in Figure 1.1.
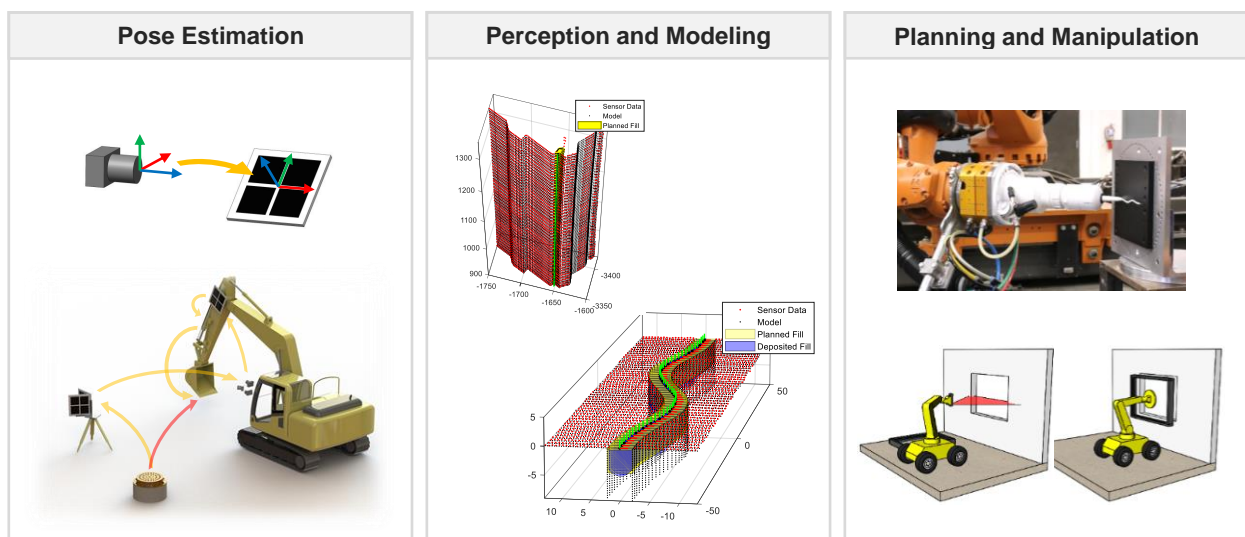
Figure 1.1 Research overview

1.1 Importance of Research

The construction industry suffers from a range of problems, such as high cost, poor quality, prolonged duration, and substandard safety. For example, the high cost of construction is indicated by the fact that the cost per unit output (i.e., unit of constructed facility) has continued to rise over recent decades while plateauing or dropping for nearly all other U.S. industries [1]. Similarly, the industry's substandard quality is reflected in the fact that rework accounts for 2 to 20 percent of a typical project's contract amount [2]. Construction projects are also lengthy and burdensome, often imposing considerable societal and opportunity costs during the months or years of their duration. Lastly, the reported incident rate for construction injuries is 20% higher than the average private U.S. industry [3], and construction consistently ranks highest among all U.S. industries for fatalities [4].

Robots possess great potential to help address many of these problems. Robotics has been extensively proposed in national reports and other literature as a means of improving construction productivity and safety [5,6]. Construction work tends to be dirty, dull, and dangerous, which are the characteristics often heralded as those best suited for robotic assistance. Construction robots can help to reduce project costs and shorten project timeframes by expediting construction processes, such as by reaching high work to eliminate the need for scaffolding. They can help to improve construction quality by measuring, cutting, or placing construction components with high precision. Robots can also help improve working conditions and job satisfaction by assisting with detested work. Lastly, construction robots can help reduce injuries by performing repetitive or strenuous tasks, such as drywall installation, or by assisting with work in dangerous places such as confined spaces or at heights. In all, construction robots have the potential to help reduce

construction costs, improve construction quality, shorten project duration, improve job satisfaction, and improve worker safety [7].

Together, construction robots and their human co-workers have the potential to form a complementary team that is greater than the sum of its parts. In human-robot collaborative work, the distribution of responsibility between the robot and the human can be represented as shown in Figure 1.2. The robot's capabilities can be described as a vector along three axes of competence: physical, cognitive, and planning. These respectively refer to a robot's ability to perform physical work, a robot's ability to perceive and reason so it can adjust its plan, and a robot's ability to do high-level operational planning (e.g., project scheduling). A null vector corresponds to a fully human operation, a unity vector represents a fully autonomous operation, and everything between represents varying degrees of semi-autonomy and human-robot collaboration. In construction, humans excel at high-level planning and troubleshooting, but they struggle with precise motion control and physical exertion—demands which robots are well suited to handle. Therefore, it is envisioned that humans will assume nearly all responsibility for high-level planning and troubleshooting, such as scheduling, resource allocation, and supply chain management, and they will benefit from the assistance of robots in such work as measuring, precise motion control, and heavy lifting.
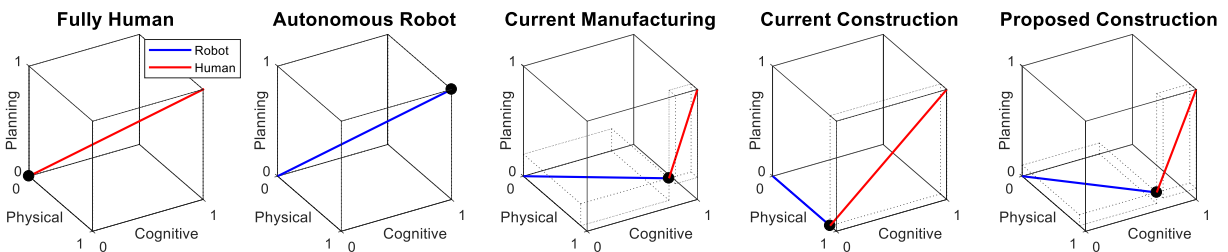
Figure 1.2 Human-robot responsibility distribution in construction

In industries like manufacturing, the structured nature of the environment and the production process have enabled robots to advance along the planning and cognitive axes of the responsibility distribution, but in industries like construction, robots have remained largely confined to the physical axis. Today's construction robots exist primarily in the form of mechanized equipment, i.e., human-operated machinery capable of assuming nearly all physical responsibility but very little planning or cognition.

Despite robots' potential to help, it has been difficult for the construction industry to move from mechanization to robotization because the nature of the industry poses additional challenges for robots. First, construction projects are unique, so although similarities exist between one project and the next, as well as between steps within a project, no two projects or steps are identical. Even in the case of a highly repetitive activity such as bricklaying, no two bricks are placed in the exact same location, which is a notable difference from manufacturing. Second, because the location changes for each project and the environment evolves throughout the construction process, it is not practical to instrument the environment or apply engineering controls to support robot operation to the degree commonly found in industries like manufacturing. Third, construction environments are rugged, cluttered, dynamic (i.e., perpetually changing), and unstructured, which poses many challenges for perception, scene understanding, and planning. Fourth, the construction robot must go to its workpieces, as opposed to having its workpieces brought to it, thereby producing a reversed spatial conveyance between the product and the robot [8]. This, along with the sheer size of construction projects, requires the construction robot to be mobile, which contributes to the robot's own spatial uncertainties, as well as those between itself and its workpieces. Fifth, construction materials tend to exhibit considerable geometric variation. Due to their substantial size and properties, construction materials are often susceptible to large

deflections and geometric irregularities. Lastly, construction tolerances are relatively loose, which tends to result in large tolerance stack-ups. This is in part due to the fact that many tolerances are left unspecified in project designs [9], and the distributed nature (i.e., location) of projects makes it difficult to perform strict quality control.

Thus, construction robots must be able to adapt to the large uncertainties and lack of structure in their surroundings. Toward that end, perceptual advances are needed to move the state of the art in the responsibility cube along the cognitive axis, making it a human-robot collaboration where the human assumes most of the high-level planning and troubleshooting responsibility, and the robot assumes most of the physical and cognitive responsibility. Since it is largely the construction industry's unstructured, chaotic nature that has kept robots from shifting further along the planning or cognition axes, there are two solutions to help robots become more impactful in construction: the industry itself can become more structured, or robots can be made more adaptive to compensate for the lack of structure. Some of both may need to occur, but in this research, I focus on making construction robots more perceptive and adaptive.

In order for a construction robot to provide utility as a construction partner, it must possess some autonomous capabilities. Here, the word *autonomous* is used to refer to a degree of autonomy, not an absolute form. It implies a grade of self-sufficient functionality, but is not meant to suggest that the robot is completely independent or works in solitude. In fact, the contrary is true; autonomy, to some degree, is needed for a robot to function as a useful team member. Analogously, a human who possesses no autonomous capabilities—someone who is incompetent or completely dependent upon others—offers little ability to contribute to a team. Like their human counterparts, construction robots need some level of autonomy to provide utility as construction partners.

Arguably one of the most fundamental perceptive capabilities needed of a construction robot is the ability to perceive where it is located. Thus, as a first step toward overcoming the aforementioned construction challenges and becoming a useful partner, a robot first needs the ability to estimate its own pose relative to the construction site. With such a pose estimate, the robot can project a model of itself into the construction project's BIM model so it can begin performing such actions as locomotion, granular navigation, and even some coarse construction work. Thus, the first thrust in this body of research is to develop and investigate the use of a new sensor system for estimating the pose of a mobile robot on a construction site, where an excavator serves a specific case study. A camera-marker sensor system is of particular interest since many construction environments, especially indoors and dense urban environments, cannot be reached by the Global Navigation Satellite System (GNSS).

However, even armed with a BIM model and a pose estimate, construction robots are still unable perform precision operations. Specifically, a robot cannot perform operations requiring more precision than the uncertainties propagated from the BIM model and the robot's pose estimate. For example, if a robot needs to position its tool on an object with millimeter precision, but the robot has centimeter-level uncertainty in its own jobsite pose estimate and centimeter-level uncertainty in the BIM model's representation of the actual object, then it is easy to see that the likelihood of successfully performing the operation would be prohibitively low. The robot may get lucky and place its tool in the correct position, but it is much more likely that it will fail. And if placing its tool in the wrong position has the potential to cause damage, then it may not be worth the risk to try. Thus, in order for the robot to perform precision work reliably, it needs precise models of its workpieces. One means of attaining such models is to perceive the workpieces in place and form models that are precise relative to the robot's local coordinate frame. Thus, the second thrust in

this body of research is to develop and explore a framework to allow a construction robot to perceive and model the actual pose and geometry of its workpieces. Of particular interest are techniques that leverage the availability of BIM models and can be generalized to a variety of construction objects and activities.

Lastly, the ultimate goal of almost any construction project is to produce a physical product, which almost invariably requires performing physical work. Thus, in addition to perceiving and modeling its workpieces, a construction robot needs the ability to adapt to the circumstances it encounters and manipulate its surroundings. As such, the third portion of this research explores how to enable a robot to use its newly perceived workpiece models to adapt its work plan and physically execute work. Of particular interest is a means by which work plans can be embedded into BIM models and then easily extracted and adapted to a variety of construction activities.

This research as a whole enables construction robots to operate with greater perceptive and cognitive capabilities, which greatly expands their utility as co-workers. The introduction of these advances in perceptive and cognitive capabilities permits the robot operation paradigm to shift from the task level to the activity level of Everett's taxonomy of construction work [10]. Rather than being confined to performing individual tasks, such as making a single measurement, robots can begin operating on the composite activity level, such as helping with the installation of a prefabricated window by combining task-level operations such as measuring, planning, gripping, moving, and placing.

To summarize, there are many technical challenges that must be overcome before robots can begin functioning as collaborative partners on actual construction sites. This research attempts to address such challenges by endowing construction robots with several fundamental capabilities, including

estimating their own pose, perceiving their surroundings, and performing work in an adaptive manner. This research helps address many of the construction industry's problems by synthesizing innovations from robotics, computer vision, and artificial intelligence with current construction practices. It revolutionizes how construction work is performed by augmenting human capabilities and advancing human-technology partnerships in the work environment. This research helps address the scalability of robots by enabling them to perceive, plan, and act in uncertain, real-world environments so they can serve as useful collaborators with human workers. This research also helps address the customizability of robots by enabling them to adapt to a variety of tasks and environments with minimal software modification. Ultimately, this research is a critical and significant step toward realizing the extensive benefits of robotics for the construction industry.

1.2 Background

If one considers advanced semi-autonomous construction equipment to be a primitive form of construction robot, then GNSS sensors are arguably the most common sensors for robot localization on today's construction sites. For example, equipment like dozers, graders, and scrapers are often outfitted with GNSS sensors to enable automatic blade adjustment based on the machine's global position and a desired reference grade [11,12]. Similarly, some surface mining equipment, such as haulers, are now beginning to operate autonomously (with remote supervision) using GNSS as their primary localization sensors [13-16].

GNSS-based sensor systems are particularly well-suited for earthwork and surface mining applications because the environments are open to the skies, and the sensor costs are reasonably proportionate to the equipment costs. However, GNSS-based sensor systems are not always economically justifiable for smaller, more fragmented sectors of the construction industry.

Furthermore, GNSS-based sensor systems are ineffective in some construction environments. For example, in dense urban environments, tall buildings and trees tend to block or bounce GNSS signals (i.e., multipath effect), causing degraded accuracy, intermittency, or complete loss of signal. And for underground or indoor construction environments, GNSS is altogether ineffective.

Thus, a ubiquitous, affordable sensor system is needed for construction robot localization that can reach even the places GNSS cannot. The first portion of this research explores whether cameras and fiducial markers can be used as an alternative sensor system for construction robot localization, especially due to their low cost and ability to function in a variety of environments. If common digital cameras and vinyl markers can be used for construction robot localization, then their reduced cost might make them accessible to a greater portion of the industry.

Once ubiquitous robot localization has been achieved, construction robots can begin performing some coarse work. However, in order to perform precision work, a construction robot needs a means of precisely perceiving and modeling its workpieces. Commercial solutions to this problem are effectively non-existent in the construction sector, but researchers have been working toward solutions. Still, current perception and modeling techniques fall short of enabling construction robots to perform detailed work with a useful degree of autonomy due to the following limitations. First, many perception and modeling techniques [17-20] require some degree of human intervention, often for data collection, data parsing, or data association, which limits the robot's utility as a co-worker. Second, many techniques [21-23] do not include object recognition; they often employ occupancy grid or convex hull models to represent occupied space with little or no additional meaning applied to the model. Third, techniques [24-26] often lack the modeling detail and degrees of freedom necessary to satisfy construction process tolerances. Fourth, techniques [27,28] that build models directly from sensor data, such as syntactic analysis, are difficult to

9

generalize for the wide variety of object geometries found among construction objects. Fifth, many techniques [17,23,27-29] are susceptible to modeling failure under partial object occlusion and incomplete sensor data. And lastly, many techniques are developed for specific construction objects, and work well for those particular objects, but are ineffective for the wide range of other objects commonly encountered in construction projects.

Thus, the goal of the second portion of this research is to explore workpiece perception and modeling techniques that permit a construction robot to autonomously perceive its workpieces through the registration of semantically rich BIM models with sufficient geometric fidelity to eventually perform precision construction work.

Once a construction robot has perceived its workpieces and formed precise geometric models, it needs a means of acting upon its newly formed models to produce useful work. Again, commercial solutions are effectively non-existent in the construction domain, but previous research has been conducted. However, such studies have limitations which preclude their methods from being widely adopted. For example, some studies [24,30] have demonstrated geometrically adaptive capabilities with accuracies on the order of several centimeters, but the methods used may not be sufficient for construction tasks with tighter tolerances. Similarly, other studies [26,31] have used simple workpiece models or workpiece identification methods, which might not be very generalizable to a wide range of tasks or those requiring high-fidelity models. Furthermore, some studies [32,33] have employed model registration techniques that do not lend well to objects of complex curvature or high geometric complexity, so the methods used in these studies are unlikely to be suitable for a wide range of object geometries. Other studies have used different sensing strategies, such as leveraging fiducial markers [34] or using force feedback [35], but visualization of natural features is anticipated to serve as the principal sensing modality for most construction

robots. Lastly, some studies [36,37] have demonstrated noteworthy adaptive construction capabilities, but have not disseminated the methods used, thereby preventing the results from being replicated, or the methods adapted or modified.

Thus, the goal of the third portion of this research is to explore a means by which a construction robot can adapt its work plan to its perceived workpiece models and perform precision work. Of particular interest is a means by which a work plan can be integrated directly into a BIM model, such that the BIM model can store information about the construction process and later be used by the robot to autonomously adapt its work plan to it as-perceived workpiece models.

1.3 Objectives

In general, the objectives of this research include developing a ubiquitous sensor system for construction robot localization and an adaptive manipulation framework to enable construction robots to adapt to the geometry of their workpieces and perform work. The specific objectives of this research are as follows.

- Test and characterize the performance capabilities of cameras and markers to assess their utility for construction robot localization.
- Develop and evaluate a real camera-marker sensor system for the localization of construction robots operating on actual construction sites via an iterative process of design, fabrication, and testing.
- Devise a means to enable construction robots to autonomously convert raw sensor data into meaningful models representing the actual pose and geometry of their workpieces.
- Develop and evaluate techniques for registering point cloud sensor data to construction workpiece BIM models, as a critical step in forming as-perceived workpiece models.
- Devise a means of enabling a construction robot to autonomously adapt its work plan to its as-perceived workpiece models.

- Devise a BIM-driven adaptive manipulation framework that synthesizes sensing, modeling, planning, and execution to enable construction robots to perform geometrically adaptive work.
- Evaluate a construction robot's ability to perform geometrically adaptive work despite discrepancies between its actual workpieces and their as-designed counterparts.

Together, such objectives will contribute to the overarching goal of enabling construction robots to function as useful partners in collaborative human-robot teams on actual construction sites.

1.4   Methodology

The methodology used in this research combines approaches from mechanical engineering, robotics, and computer vision with construction domain knowledge to address application-specific problems in construction engineering. Wherever existing approaches are found insufficient for adoption and integration, new approaches are developed to bridge gaps and form complete solutions.

Automation and robotics has the potential to benefit the construction industry at many levels, but the author particularly focuses on solutions at the lowest level, where the proverbial "rubber meets the road." This is due to the belief that construction robots have the greatest potential to contribute to the industry by assisting with dirty, detested, and dangerous work, such as heavy lifting, which tends to occur at the lowest levels of construction. Thus, this research focuses on robotic solutions to assist with the performance of actual construction processes, with a particular emphasis on real-world practicality. Accordingly, much of the research is experimental in nature, often involving design, prototyping, testing, and evaluation.

1.5 Dissertation Outline

This dissertation is a compilation of peer-reviewed scientific manuscripts which document an investigation into construction robot localization, scene understanding, motion planning, and task execution. The remainder of the dissertation is organized as follows.

Chapter 2 presents the development of a ubiquitous camera-marker sensor system for construction robot localization. In particular, digital cameras and vinyl markers are explored for their use as pose estimation sensors. Excavation serves as a case study construction activity where pose estimation enables utility avoidance and automatic grade control.

Chapter 3 describes the development of a framework to enable robots to leverage their sensors and BIM models to perceive and model of the actual pose and geometry of their workpieces. Joints between abutted objects are used as a case study for evaluating the robot's ability to autonomously model construction objects and features.

Chapter 4 describes the development of autonomous planning and execution methods to enable construction robots to physically perform work in an adaptive manner. Joint filling, and in particular caulking, is used as a case study for demonstrating and evaluating a robot's ability to adapt to the geometry of its workpieces and perform precision work.

Lastly, Chapter 5 provides a summarizing conclusion of the dissertation as a whole. Specifically, Chapter 5 describes the significance and contribution of the research, as well as discusses future research directions.

Chapter 2

Camera-Marker Sensor System for Construction Robot Localization

2.1 Introduction and Motivation

2.1.1 Motivation for Articulated Excavator Pose Estimation Systems

Excavation is a common construction activity wherein earth is removed to produce a desired ground contour or gain access to a buried object. Much of today's construction excavation is performed by excavators, or human-operated powered digging machines. In general, the operation of an excavator requires skill and experience, but even skilled operators may be susceptible to human error.

In US brownfield (i.e., urban) environments, the excavation process typically begins with a request of local utility providers to mark the location of underground utilities in a planned excavation area. Approximate utility locations are then marked on the ground's surface with flags or paint, and with little or no information indicating the depth of the utility. The uncertainty inherent in locating underground utilities is indicated by the fact that mechanized excavation is typically prohibited within a 0.9 meter (36 inch) wide band of the utility markings [38]. Additionally, the incidental disturbance of ground markings during construction processes can further exacerbate the uncertainty in actual utility location.

Not only does there exist uncertainty in the location of underground utilities, but the sheer number of utilities continues to grow with population, property development, and the advancement of

technologies like telecommunications. According to the Common Ground Alliance, an organization which collects and reports data on underground utility strikes, the US contains more than 30 million kilometers (19 million miles) [39] [40] of underground utilities including gas, electric, water, sewer, cable television, fiber optic, phone, drainage, traffic signals, and street lighting circuits. Shown in Figure 2.1 are two examples of the congestion found underground in some urban areas. Many of these utilities are now approaching the ends of their service lives and will need to be repaired or replaced in the near future. For example, the American Water Works Association reports that the majority of the US underground water supply infrastructure is at least fifty years old and will need to be replaced or expanded over the next twenty-five years [41].

Figure 2.1 Examples of urban underground congestion

The increasing volume of underground utilities and the uncertainty in their location are factors contributing to the rate of excavation incidents. The 2013 Common Ground Alliance Annual Damage Information Reporting Tool Report estimates there were over 4 million [42] excavation incidents resulting in damaged underground utilities in the US during the ten year period between 2004 and 2013. Similarly, the US Department of Transportation Pipeline and Hazardous Materials Safety Administration reports there were 786 [43] incidents of excavation-related damage to

natural gas and liquefied petroleum gas pipelines in the US during the same ten years. Those 786 gas pipeline incidents resulted in 39 fatalities, 127 injuries, and over 200 billion dollars in property damage [43]. As stated in the US Congressional Transportation Equity Act for the 21st Century, Title VII, Subtitle C, Section 7301, "…unintentional damage to underground facilities during excavation is a significant cause of disruptions in telecommunications, water supply, electric power, and other vital public services, such as hospital and air traffic control operations, and is a leading cause of natural gas and hazardous liquid pipeline accidents" [44].

Thus, changes are needed to help reduce incidental utility strikes. One potential solution lies in the establishment of accurate as-built records of underground utilities [45] which could be used to provide better estimates of utility locations onsite. It is anticipated, however, that the existence of accurate ground markings alone may have a limited effect on the reduction of utility strikes, since operators would continue digging without direct sight of the bucket or buried utilities. A second solution involves combining accurate utility records with sensory feedback to provide operators with real-time estimates of their excavator's pose (i.e., position and orientation) relative to the jobsite [46]. In such a system, sensors could be used to estimate the current pose of an excavator's components, combine the information with a model of underground utilities, and relay the information back to the operator to permit a more informed judgment. Potentially, the inclusion of warning systems or more sophisticated measures could also be implemented to further reduce the risk of utility strikes.

In addition to warning operators where they should not dig, excavator pose estimation systems can help operators identify where they should dig, as shown in Figure 2.2. In some applications, like drainage pipe installation for example, both the height and slope of the excavated surface are often critical. In such situations, accurate grading is typically accomplished through an iterative process

of digging and checking. Checking is usually performed by a human grade checker who manually measures the height of the ground's surface and informs the operator whether more soil should be removed. In addition to being slow, the process can be dangerous, as the grade checker is often subjected to the hazards of working near both trenches and heavy machinery. Pose estimation systems may serve to relieve a grade checker of such duty, thus potentially improving both safety and efficiency.

Grade Control (Where to Dig)

Utility Avoidance (Where Not to Dig)
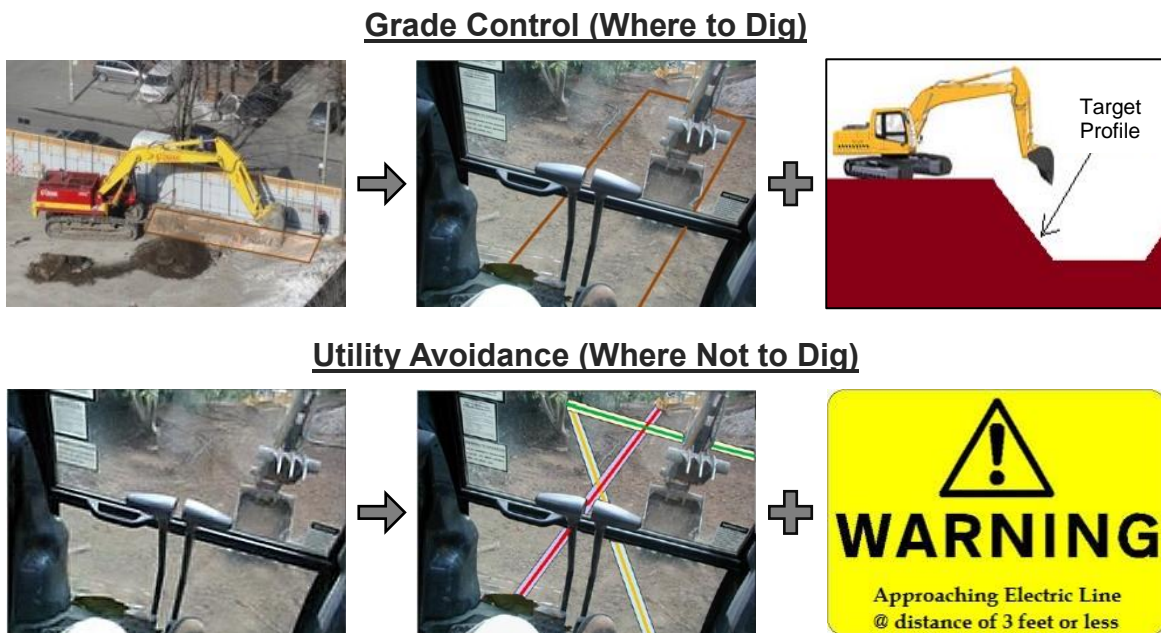
Figure 2.2 Benefits of excavator pose estimation: grade control and utility avoidance

2.1.2 Motivation for a Low-Cost, Ubiquitous Pose Estimation System

Excavator pose estimation systems are commercially available today [47] [48], but have not yet been widely adopted by the construction industry. Various types of pose estimation systems are available, but all possess shortcomings. Companies like Trimble and Leica offer packages that make use of Global Navigation Satellite Systems (GNSS) to permit a full six-dimensional (three position and three orientation) estimation of the machine's pose relative to the jobsite. These

17

systems likely offer the most potential to improve construction safety and productivity, but tend to be cost prohibitive due to the expensive GNSS components. Additionally, the use of these systems is limited to environments where GNSS communication is reliable, as GNSS-based systems are susceptible to failure in areas with tall trees, in urban environments with tall buildings, and indoors [49].

Alternatively, pose estimation systems can make use of robotic total stations, rather than GNSS, to provide full six-dimensional pose estimation relative to the jobsite. Since these systems do not use GNSS, they are capable of functioning reliably in urban environments and indoors. However, as each excavator requires a dedicated robotic total station during operation, total-station-based systems also tend to suffer from high cost.

Lastly, laser level pose estimation systems offer an alternative to GNSS and total station systems. Using a rotary laser level and a receiver mounted on the stick (dipper) of the excavator, these systems are capable of accurately monitoring digging depth. However, these systems are incapable of determining an excavator's position within the jobsite. Additionally, their functional range is often limited by the length of the laser receiver. Although laser level systems offer an affordable option capable of functioning in GNSS-denied areas, their potential to improve construction safety and productivity is hindered by limitations in their position estimation and functional range. Thus, the development of a low-cost, ubiquitous, six-dimensional excavator pose estimation system may help to fill a market need.

2.1.3 Motivation for a Marker-Based Pose Estimation Prototype

One potential approach to developing a low-cost, ubiquitous, six-dimensional excavator pose estimation system is to use optical cameras and fiducial markers as sensors. A fiducial marker, as

18

described here, refers to a two-dimensional visual object having a prescribed geometry and identification features to distinguish it from naturally occurring objects and other markers. Upon capturing video of such a marker, a camera's image can be analyzed to estimate the marker's pose relative to the camera, generally in the form of a homogeneous transformation, as illustrated in Figure 2.3. Given that the focus of this research is prototype development, only an overview of camera-marker pose estimation is provided here. Interested readers are referred to Feng and Kamat, 2013 [50] for more details.
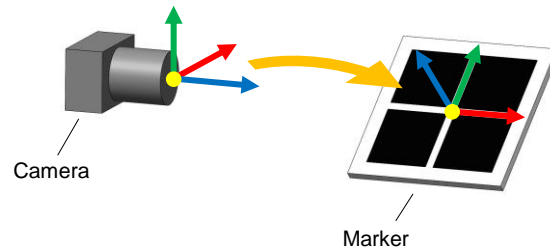
Figure 2.3 Camera-marker pose estimation

Such methods can be extended to the pose estimation of an object simply by attaching a marker to the object and viewing it with a camera. Thus, it follows that cameras and markers could hypothetically be used to estimate the pose of an excavator simply by attaching markers to the machine's components and viewing the markers with a camera, as described by Feng et al., 2015 [51]. A marker-based sensor system could potentially offer such benefits over traditional systems as requiring fewer and less expensive sensors, as well as providing functionality in urban environments and indoors. To evaluate the feasibility and capabilities of a marker-based excavator pose estimation system, a series of prototypes was developed and tested. The remainder of the chapter is dedicated to describing the development of such a physical marker-based excavator pose estimation system.

2.2 Technical Approach Overview

Prior to the work described in this chapter, preliminary studies were conducted in a simulated environment to evaluate whether cameras and markers could be used for excavator pose estimation. The studies offered encouraging results, but were limited in their ability to capture real-world effects. Interest in real-world performance and commercial application drove the need to develop and test physical prototypes. As such, the building and testing of field prototypes emerged as the primary objective of this work.

Rather than immediately delving into prototype development for excavation equipment, two experiments were first conducted to gain a better understanding of the pose estimation capabilities of cameras and markers. Section 2.3 is dedicated to the discussion of such experiments. The first experiment was conducted to determine whether the pose estimation capabilities of real cameras and markers were accurate enough to justify further prototype development, as well as to study the effects of marker yaw on pose estimation accuracy. Unexpected findings led to a second camera-marker experiment, which further explored the effects of marker yaw angle on pose estimation accuracy. The two experiments not only verified the system's accuracy and provided the confidence to move forward with prototype development, but also offered insights into camera-marker behavior and informed strategies for prototype design.

After conducting the two camera-marker experiments, work continued with an iterative process of excavator-centric prototype development. Section 2.4 is dedicated to the discussion of such prototypes. Four prototypes in total were iteratively developed to evaluate the feasibility and performance of a marker-based pose estimation system on an actual excavator. The first iteration advanced through conceptualization stages into design stages, but was prematurely terminated due

to the discovery of singularities in the mechanical system. The second design iteration provided much the same function as the first design, but utilized a different transmission mechanism free of singularities. The full prototype was found to perform satisfactorily, but was deemed impractical for use on actual excavators as it impeded normal operation. The third design iteration attempted to minimize interference with excavator operation by using different marker placement locations on the machine. To facilitate such a design, a new mapping approach was adopted. When experimentally evaluated, however, the third iteration prototype was found to provide insufficient accuracy. Lastly, the fourth design iteration employed a hybrid sensor suite by combining an electromechanical sensor with camera-marker sensors and the mapping approach developed in the third design. Experimental evaluation of the fourth iteration prototype yielded satisfactory results.

2.3   Identification of Camera-Marker Performance Characteristics

After having identified cameras and markers as an alternative sensor system for excavator pose estimation, it was important to gain a better understanding of the system's performance capabilities. Interviews with more than a dozen professional excavator operators indicated +/- 2.5 centimeters (1 inch) as a typical tolerance range for most excavation activities. As such, an experiment was conducted to identify the sensor system's pose estimation accuracy for comparison with customer requirements.

Since the relative pose between a camera and marker is defined by three components of translation (e.g. differential x, y, and z positions) and three components of orientation (e.g. differential roll, pitch, and yaw angles), a full characterization of pose error would require measurement of all six components. However, it was not feasible in the experiment to measure all six components quickly or accurately with the equipment available. Rather, a single component, vertical displacement, was

chosen as an approximate measure of the system's pose error. Vertical displacement was selected because ground truth measurements could be quickly and accurately ascertained using a rotary laser level and tape measure, and because vertical displacement, or depth, is generally the component of greatest interest in most excavation activities. Thus, the experiment involved estimating the vertical displacement between a camera and marker using the sensor system, and comparing the estimates with ground truth measurements.

While measuring the system's pose estimation accuracy, the effects of marker yaw were also studied by varying the marker's yaw angle. Marker yaw was defined as shown in Figure 2.4, with a zero yaw angle corresponding to the marker facing directly toward the camera. The experiment was composed of 756 test combinations, which were equally divided between 0, -15, and +30 degree marker yaw angles.
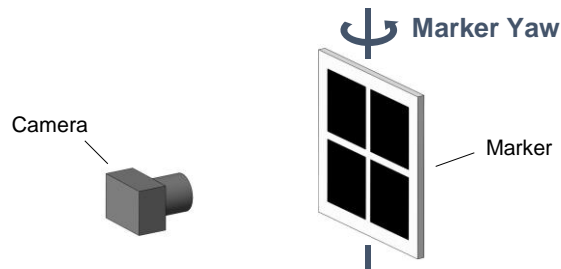
Figure 2.4 Marker yaw illustration

The markers used in the experiment were made from AprilTags developed by Olson [52], and measured 57 centimeters (22.5 inches) between outermost black edges. The markers were leveled and affixed to a plumb surface, which was free to rotate about a plumb axis. The camera used was a 1.3 megapixel Firefly® MV optical camera by Point Grey with a Fujinon DF6HA-1B 6mm F/1.2 fixed focal lens. The camera was mounted at the same height as the center of the marker. Half of the experiment's test combinations were conducted at a distance of 4.0 meters (13 feet), and the

other half at a distance of 7.6 meters (25 feet). A picture of the experimental setup is shown in Figure 2.5.
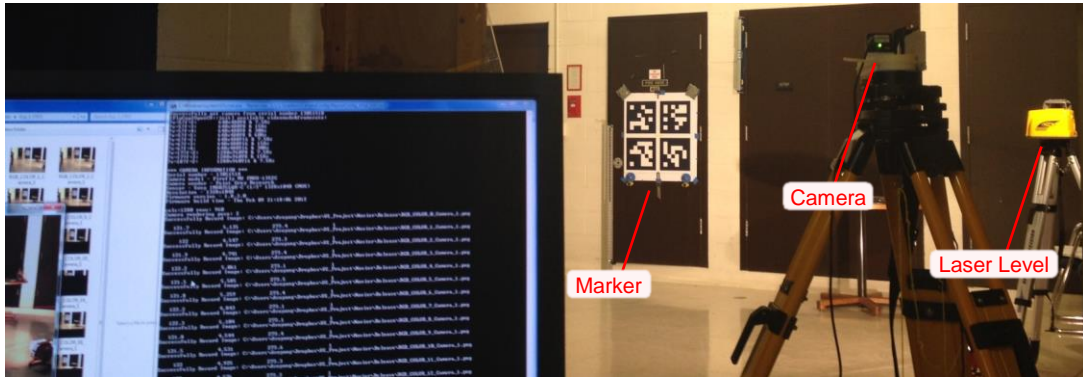
Figure 2.5 First camera-marker experimental setup

The experimental results are shown in Figure 2.6. The experimental results led to two interesting observations. First, the sensor system was found to exhibit considerably less error when the marker's yaw angle was nonzero (i.e., -15 or +30 degrees). Such results contradicted expectations, initiating a follow-up experiment to further explore the effects of marker yaw. Second, it was observed that the absolute error remained below 2.5 centimeters (1 inch) for 97% of nonzero marker yaw cases. For such nonzero marker yaw cases, the sensor system's accuracy appeared sufficiently close to the requirements of excavation to warrant further development.
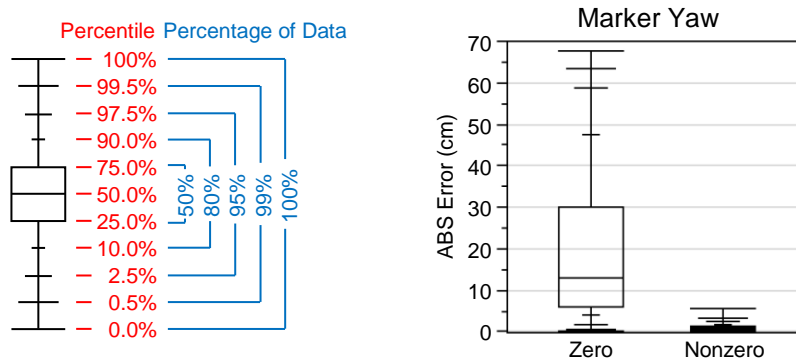
Figure 2.6 First camera-marker experimental results

A second experiment was then conducted to further investigate the seeming correlation between marker yaw and error suggested by the first marker experiment. The experimental setup was similar to that of the first experiment, as shown in Figure 2.7, but with a denser and wider range of marker yaw angles studied, namely -20, -10, 0, +5, +10, +15, +20, +30, +40, +50, +55, +60, +65, +70, and +80 degrees. Vertical displacement was again selected as the measure of error because ground truth measurements could be quickly and accurately ascertained using a rotary laser level and tape measure, and because vertical displacement, or depth, is generally the component of greatest interest in most excavation activities. The markers used in the experiment were made from AprilTags developed by Olson [52], and measured 57 centimeters (22.5 inches) between outermost black edges. The markers were leveled and affixed to a plumb surface, which was free to rotate about a plumb axis. The camera used was a 1.3 megapixel Firefly® MV optical camera by Point Grey with a Fujinon DF6HA-1B 6mm F/1.2 fixed focal lens. The camera was mounted at the same height as the center of the marker. The experiment was composed of 120 test combinations, half at a distance of 4.6 meters (15 feet) and half at 9.1 meters (30 feet).

Figure 2.7 Second camera-marker experimental setup

The marker yaw behavior observed in the first experiment was again produced in the second experiment, as can be seen in Figure 2.8. The system's ability to measure vertical displacement

24

appeared to suffer considerable error in the vicinity of zero marker yaw, and appeared to improve as marker yaw approached 70 degrees. However, at 65 degrees and beyond, the camera began to lose detection of the marker, as might be expected for any two-dimensional image turned nearly perpendicular to a viewer's line of sight. Thus, it appeared advisable to avoid marker yaws near the vicinities of 0 and 90 degrees in pose estimation applications.
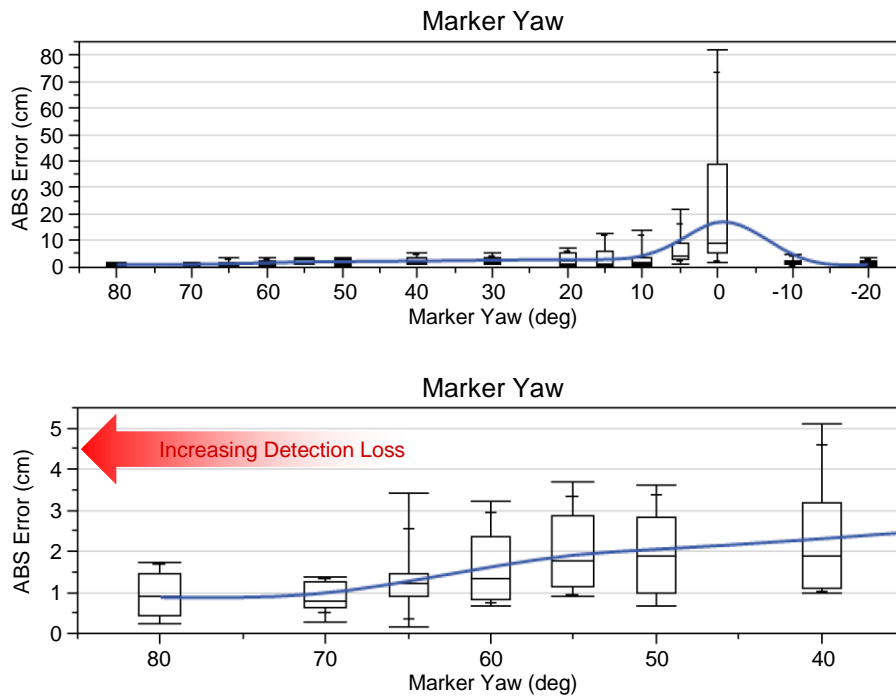
Figure 2.8 Second camera-marker experimental results

The results of the two experiments suggested that camera-marker sensor systems might be sufficiently accurate for excavator pose estimation applications provided marker yaw angles near 0 and 90 degrees are avoided. After gaining insight into the performance characteristics of marker-based sensors, it was possible to consider their physical implementation on excavation equipment.

2.4　Electromechanical Prototype Development

2.4.1　System Implementation Overview

To begin the development of a functional prototype for excavators, it was first noted that the teeth of the excavator bucket represent the cutting edge of the machine, making bucket tooth location and orientation the primary points of interest in most excavation activities. As previously stated, it is conceptually possible to determine the pose of a bucket's teeth by attaching a marker to the bucket and observing the marker with a camera. In the simplest case, tooth pose can be determined by placing a marker directly on the side of the bucket in alignment with the teeth, as shown in Figure 2.9. Additionally, a camera can be aimed toward the setup with a known relationship to the jobsite, for example, facing north on a level tripod one meter above a benchmark. Here, the term *benchmark* refers to a known jobsite reference point, sometimes denoted by a point marked on a stable and pervasive jobsite feature like a sewer manhole rim. In such a manner, a kinematic chain can be established to relate tooth pose to the rest of the jobsite. If a model of the jobsite includes underground utility information, the bucket's pose can also be determined relative to the underground utilities.
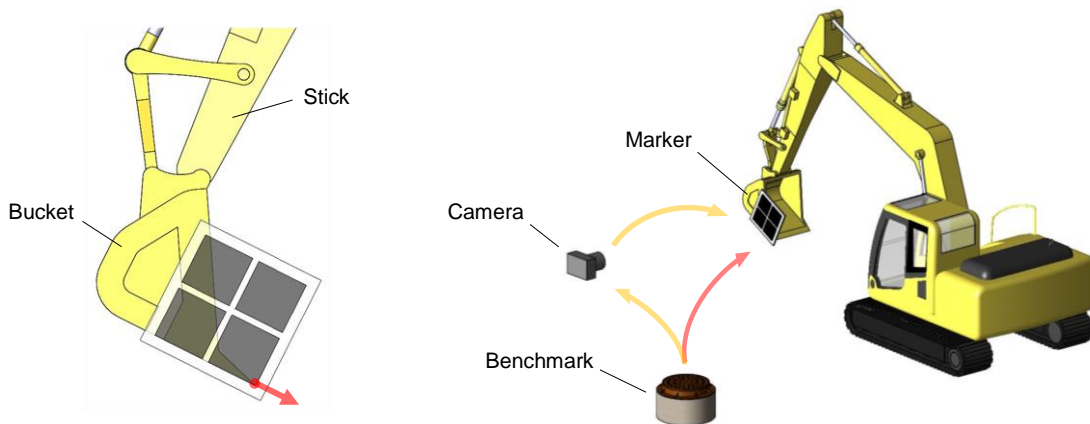
Figure 2.9 Conceptual marker-based sensor setup

26

However, there may exist benefits to mounting the camera on the excavator cabin since it is a likely location for the occurrence of data processing and interfacing with the operator. In such close proximity to the marker, a second camera, with known relation to the first camera, can be added to help expand the sensor's vertical range. As shown in Figure 2.10, this cabin-mounted camera setup may be accomplished through the use of one or more site-fixed markers called sentinel markers. Sentinel markers, as defined here, are markers stationed nearby with a known or measured relationship to the jobsite. Lastly, a third camera, having known relationships to the other two cameras, may be added to view the sentinel marker. With such a setup, the camera cluster is free to move with the excavator, provided at least one sentinel marker is within view of at least one camera.
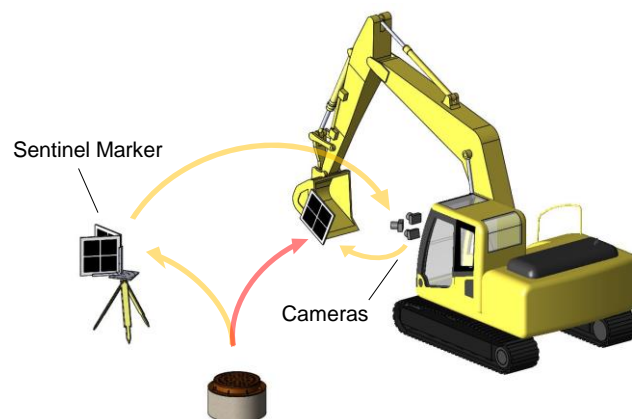
Figure 2.10 Cabin-mounted camera setup with sentinel marker

However, since the bucket penetrates the ground's surface while digging, it may be impractical to place a marker near the bucket teeth, as the ground would likely occlude the line of sight between the camera and marker. In fact, placement of a marker anywhere on the bucket may be impractical due to the likelihood of occlusion or damage during digging. Thus, an alternative method is needed to estimate tooth pose without placement of a marker directly on the bucket.

2.4.2 Four-Bar Linkage Design

One solution for implementing marker-based tooth pose estimation without placing a marker directly on the bucket is to form a short kinematic chain involving the last two links of the excavator: the stick and the bucket. One marker can be used to provide an estimate of the stick's pose, and another to estimate the bucket's rotation relative to the stick, as shown in Figure 2.11. The first marker, termed stick marker, is rigidly attached to the stick with a known relationship to the bucket's axis of rotation. The second marker, termed rotary marker, is attached at a location removed from the vicinity of the bucket. The rotary marker is constrained with one degree of rotational freedom and a known angular relationship to the bucket. If the bucket's geometry is also known, or measured onsite, then all necessary information is available to deduce tooth pose.
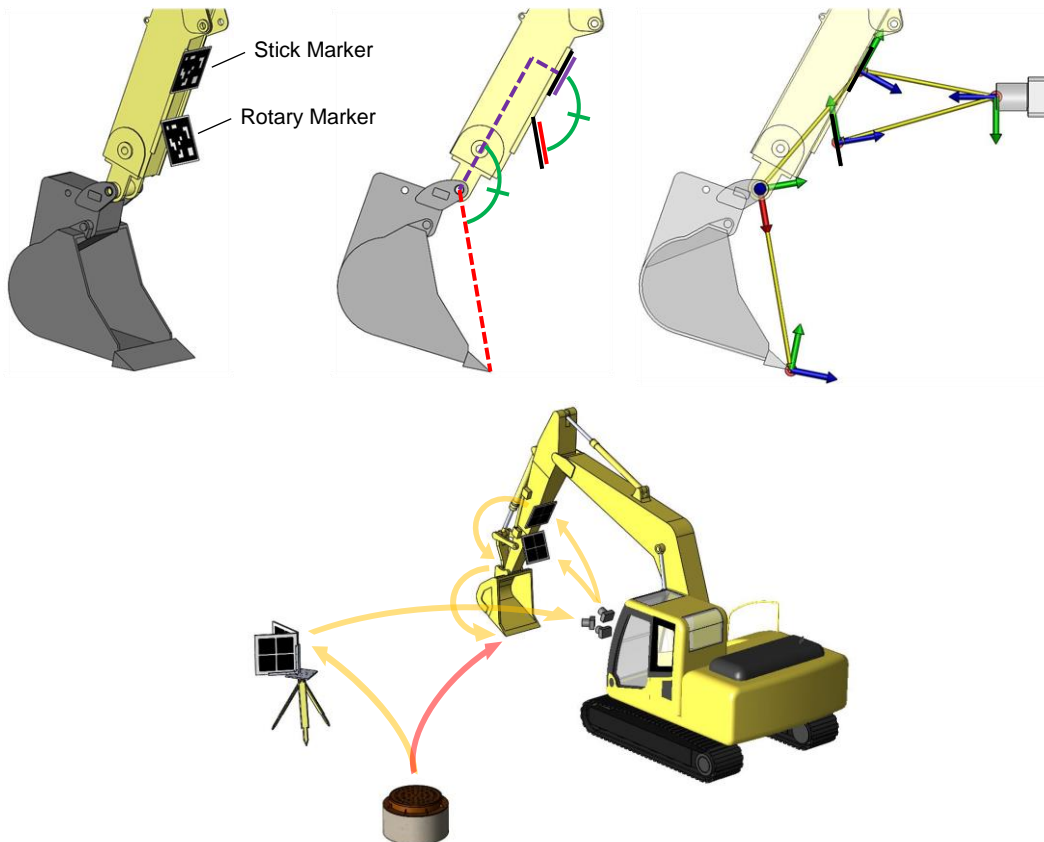
Figure 2.11 Marker-based kinematic chain setup

28

The four-bar linkage was one mechanism briefly explored for implementing such a system. An initial design for an experimental prototype was developed, as shown in Figure 2.12. The primary design objective involved transmitting the bucket's angle to a removed location and outputting the angle through the pose of a marker, all without permanently modifying the excavator. The design features a parallelogram four-bar linkage with strong magnets for attachment to the stick and bucket.
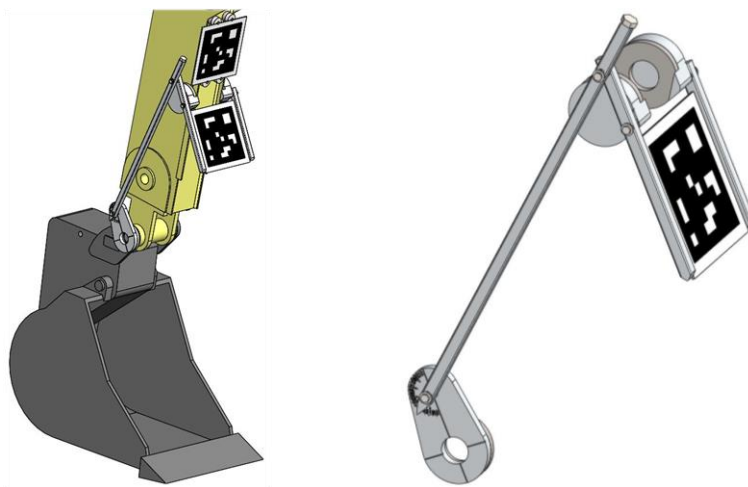
Figure 2.12 Four-bar linkage design

During the design process, it was realized that the operational range of an excavator bucket commonly exceeds 180 degrees, while the singularities inherent in a parallelogram four-bar linkage limit its functional range to less than 180 degrees. Although it may have been possible to avoid singularities through the use of a non-parallelogram four-bar linkage, such a mechanism would have introduced a nonlinear relationship between the bucket and marker. To avoid the complexities associated with nonlinear mapping, the non-parallelogram four-bar linkage was avoided. Thus, the four-bar linkage design was discarded for an alternative free of singularities and nonlinearities.

2.4.3   Synchronous Belt Design

The synchronous belt and pulley, sometimes referred to as toothed belt and pulley, was identified as an alternative mechanism for transmitting bucket angle to a removed location without singularities or nonlinearities. In addition to providing continuous rotation, the belt and pulley system offers the benefits of low maintenance and gearing capabilities through the employment of different sized pulleys. A belt-driven prototype was designed and fabricated for the purpose of marker-based sensor experimentation, as shown in Figure 2.13. As with the four-bar linkage, the primary design objective involved accurately transmitting bucket angle to a marker at a removed location without permanently modifying the excavator. In order to accurately transmit bucket angle through the device, alignment between components was critical.
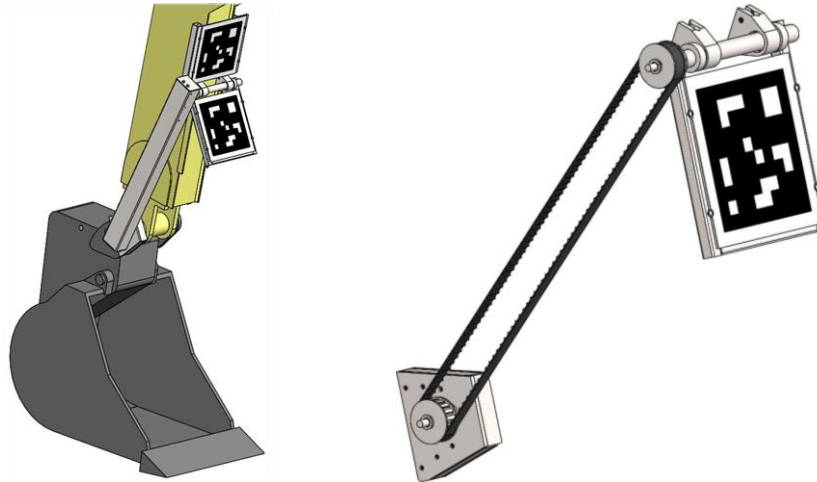
Figure 2.13 Synchronous belt design

At the base of the device is a bracket which clamps to the bucket through the use of set screws, as shown in Figure 2.14, thus providing sufficient anchoring for experimentation without permanently modifying or damaging the excavator. Before clamping, the bracket is aligned with the bucket's teeth and axis of rotation. A cylindrical pocket in the rear of the bracket is fitted over

the bucket's pin for alignment with the bucket's axis of rotation. The bracket is aligned with the

bucket's teeth through the use of an alignment tool.
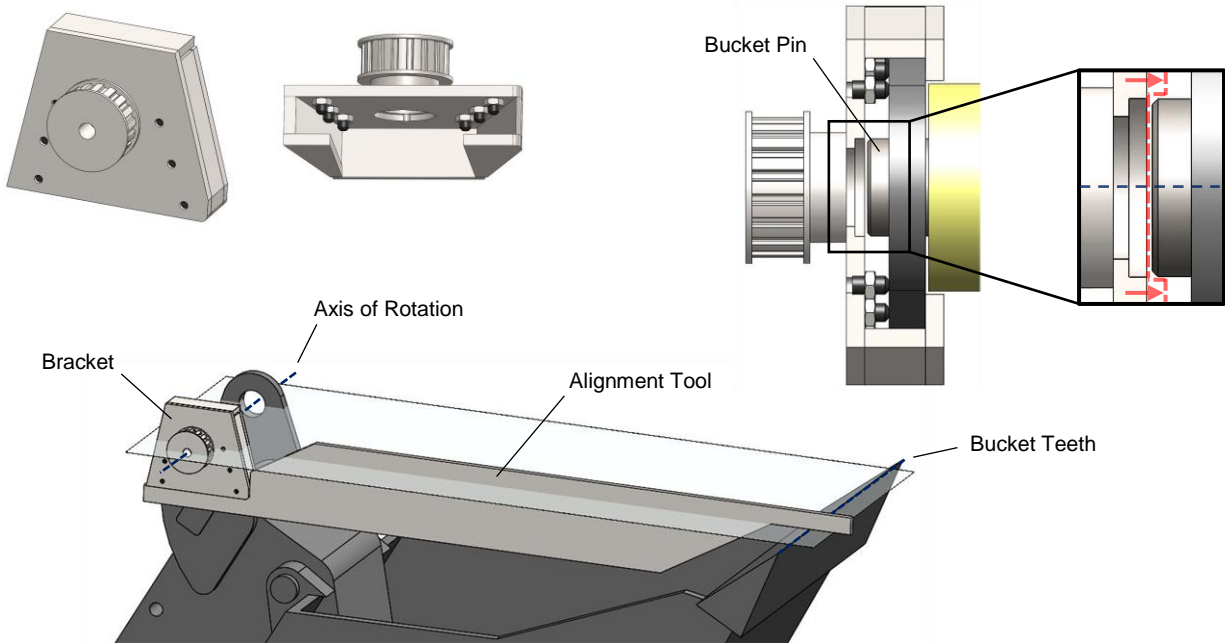
Figure 2.14 Alignment and clamping of synchronous belt device to excavator bucket

Attached to the bracket is a synchronous pulley. During fabrication, the pulley is aligned with the

bracket using dowel pins and accurately located holes, as shown in Figure 2.15. The pulley drives

a synchronous belt, which transmits bucket angle to an upper pulley. The upper pulley is attached

to a belt tensioning system which helps maintain accurate alignment between pulleys.
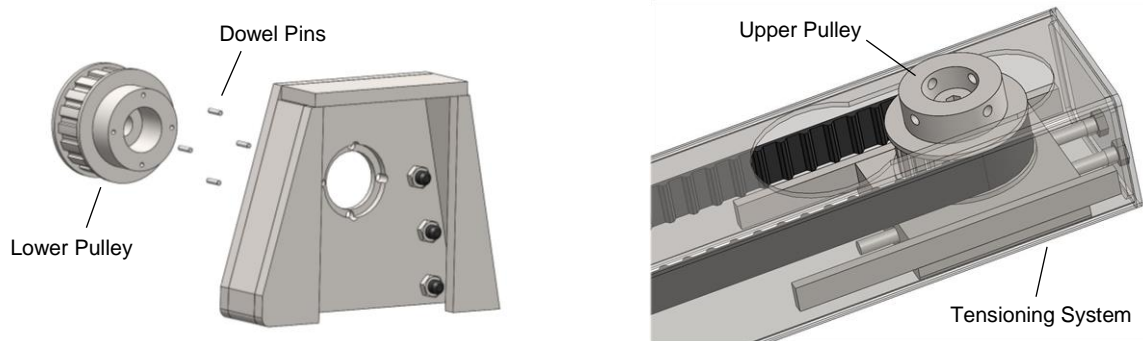
Figure 2.15 Pulley alignment methods

The upper pulley then attaches to a shaft, which drives the rotation of the marker, as shown in Figure 2.16. Cone point set screws are used to connect the upper pulley to the marker shaft. During installation, the cones press into the conical recesses in the shaft, thereby aligning and securing the two parts. This also provides a detachment point, allowing the prototype to be disassembled into two parts for easy installation and removal. Lastly, the rotary marker shaft is clamped to the excavator stick with set screws to provide sufficient anchoring for experimentation without permanently modifying the machine.
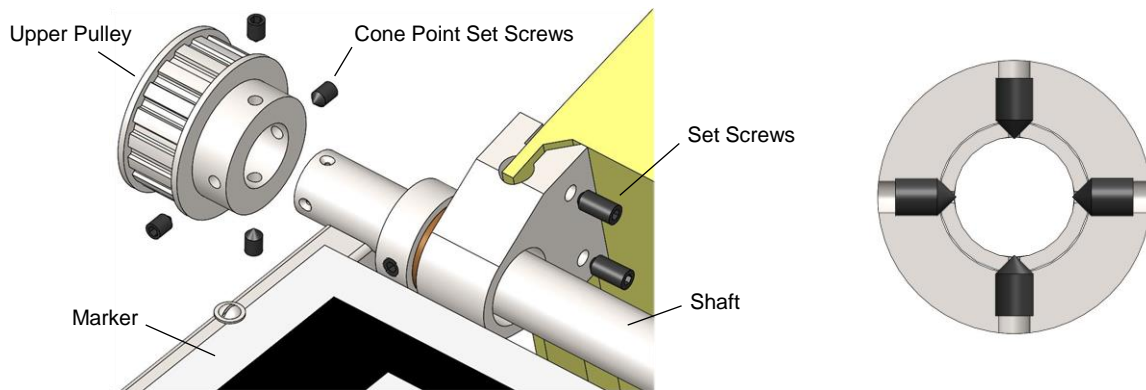
Figure 2.16 Alignment and attachment of upper pulley and rotary marker shaft

The synchronous belt prototype was fabricated, installed, and tested on a Caterpillar 430E IT Backhoe Loader, as shown in Figure 2.17. The experiment served three primary purposes: to determine whether the prototype functioned mechanically as intended, to measure the accuracy of the entire pose estimation system, and to obtain open-ended feedback from a professional excavator operator about the prototype. The primary function of the sensor system is to provide an estimate of bucket tooth pose, which is defined by three components of position (e.g. x, y, and z) and three components of orientation (e.g. roll, pitch, and yaw) relative to a jobsite coordinate frame. In order to obtain an approximation of the system's pose estimation accuracy, a single component of bucket tooth pose was measured: tooth height. Tooth height was selected because

ground truth measurements could be quickly and accurately ascertained using a rotary laser level and tape measure, and because tooth height, or digging depth, is generally the component of greatest interest in most excavation activities.

The experiment involved placing the backhoe in random poses, using the sensor system to estimate bucket tooth height, and comparing sensor estimates with ground truth measurements to obtain a measure of error. The experiment's pass/fail criterion was set at 2.5 centimeters (1 inch) of absolute error. The markers used in the experiment were made from AprilTags, as developed by Olson [52], and measured 43.2 centimeters (17.0 inches) between outermost black edges. The cameras used were 1.3 megapixel Firefly® MV optical cameras by Point Grey with Fujinon DF6HA-1B 6mm F/1.2 fixed focal lenses. The cameras were mounted on a tripod near the operator's cabin, at approximately 1.5 meters (5 feet) above the ground. The experiment involved a total of 30 test combinations which were divided equally into three sets of horizontal distances between the camera and bucket teeth: 3.7 meters (12 feet), 4.9 meters (16 feet), and 6.1 meters (20 feet). These distances were chosen because they appeared fairly representative of the excavator's typical operational range. The experimental setup is shown in Figure 2.17.
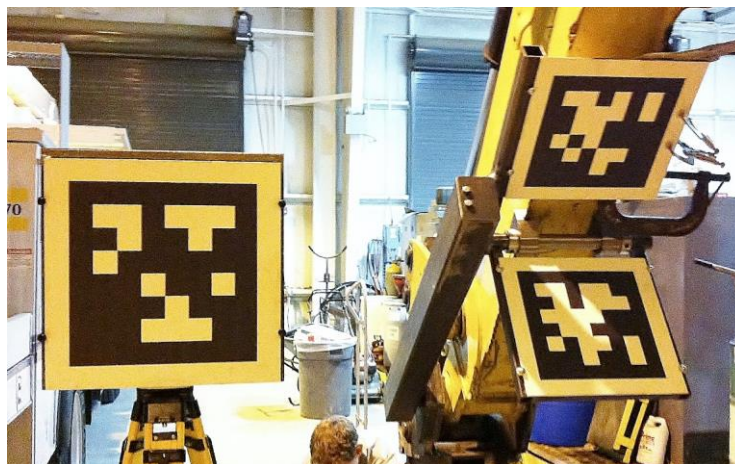
Figure 2.17 Synchronous belt prototype installed on backhoe

The experiment was initiated, but inferior camera calibration methods, software bugs, and lack of tag redundancy in the markers led to poor sensor performance, and the experiment was quickly terminated. After months of further camera and software development, the experiment was repeated. The results of the repeated experiment are shown in Figure 2.18, where it can be seen that all thirty estimates of tooth height fell within the maximum allowable 2.5 centimeters (1 inch) of absolute error.
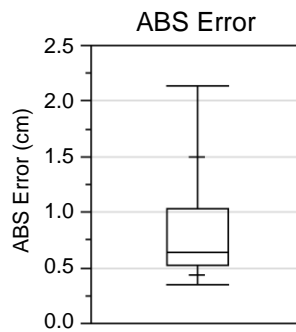
Figure 2.18 Synchronous belt experimental results

Despite the system's demonstrated ability to estimate tooth height, it was ultimately determined that the synchronous belt design was not practical for use on excavators. The professional excavator operator assisting with the experiment expressed several concerns about the design. The first concern pertained to the hindrance of the bucket's clamping functionality by the presence of the device. Specifically, operators sometimes clamp objects between the bucket and stick for gripping purposes, but the locations of the markers in the synchronous belt design interfered with such action. Additional concerns included the device's lack of compatibility across different excavators, its need to be removed to change buckets, and its susceptibility to damage as a consequence of its high profile. Thus, an alternative design was sought to circumvent such issues.

2.4.4   Bucket Linkage Generalized Mapping Design

One design alternative having the potential for a low mechanical profile and high compatibility across platforms was termed *generalized mapping*. Generalized mapping, as described here, involves the use of a calibration process to map relationships between various components and bucket tooth pose. After establishing such a relationship, tooth pose can be inferred through measurement of the mapped components.

One implementation of generalized mapping is bucket linkage generalized mapping. Bucket linkage generalized mapping entails placing a stick marker on the stick and a key marker on a bucket transmission linkage, as shown in Figure 2.19. A third marker is also placed on the side of the bucket and aligned with the bucket teeth for the duration of a calibration process. The system is then calibrated by moving the bucket through its range of motion and recording intermediate marker poses obtained by the camera system. The intermediate pose measurements provide a mapping which relates the transformation between stick marker and key marker to the transformation between stick marker and tooth pose. After the calibration process, the bucket marker is removed. The sensor system then functions by observing the stick and key markers, and mapping their relative pose to that of the bucket teeth.
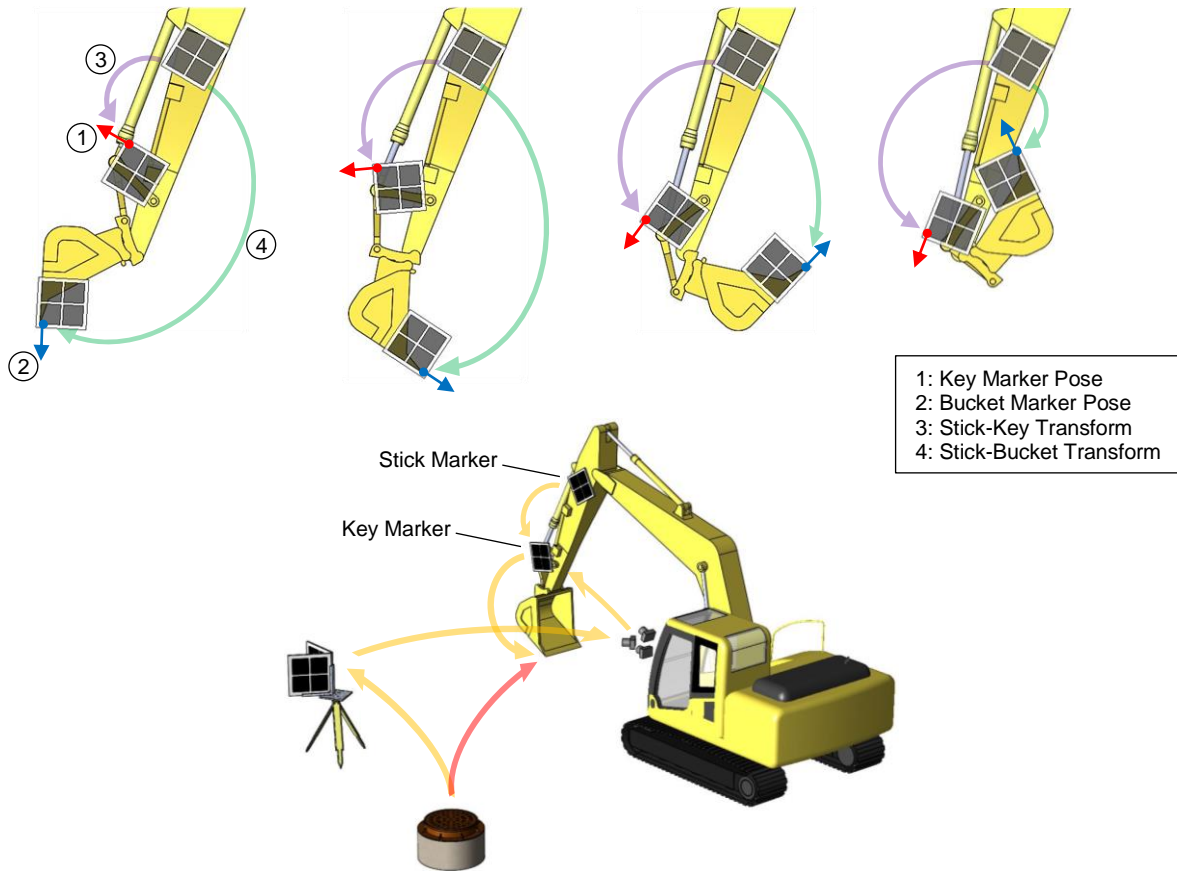
Figure 2.19 Calibration process and implementation of bucket linkage generalized mapping

One advantage of bucket linkage generalized mapping is that the placement of the stick and key markers is made nearly arbitrary by the calibration process. The markers can be placed on their respective components with little regard for alignment and location, as long as they are attached rigidly and with a clear line of sight from the cameras. Another advantage of the system is that its mechanical components are relatively flat, providing the potential for a low profile. In general, the less the prototype protrudes from the excavator's components, the less likely it will be to interfere with operation or get damaged. Although the markers used for experimental testing were of considerable surface area, it is not unreasonable to presume the markers could be made smaller through further sensor development. Lastly, the use of a calibration process allows for high

compatibility between excavators. Effectively, the only criterion for compatibility is that the markers are able to be mounted on the machine.

A prototype of the bucket linkage generalized mapping system was installed and tested on a Caterpillar 430E IT Backhoe Loader, as shown in Figure 2.20. The purpose of the experiment was to measure the system's pose estimation accuracy. As with the synchronous belt prototype experiment, the system's accuracy was approximated by measuring one of the six components of bucket tooth pose: tooth height. The markers used in the experiment were made from AprilTags developed by Olson [52], and measured 58.4 centimeters (23.0 inches) between outermost black edges. The cameras used were 1.3 megapixel Firefly® MV optical cameras by Point Grey with Fujinon DF6HA-1B 6mm F/1.2 fixed focal lenses. The cameras were mounted on a tripod near the operator's cabin, approximately 1.5 meters (5 feet) above the ground. The experiment involved calibrating the sensor system, placing the backhoe in random poses, using the sensors to estimate bucket tooth height, and comparing estimates with ground truth measurements obtained using a total station. The experiment's pass/fail criterion was set at 2.5 centimeters (1 inch) of absolute error.

Figure 2.20 Bucket linkage generalized mapping prototype installed on backhoe

The experiment was conducted, but the prototype exhibited unreasonably large errors. The experiment was abandoned and the data went unrecorded. It was concluded that the resolution of the relationship between key marker pose and bucket tooth pose was insufficient to accurately estimate tooth pose with the camera sensor system being used. It can be seen in Figure 2.21 that the key marker's radius of motion is considerably less than that of the bucket teeth. Thus, a small change in translation of the key marker corresponds to a large change in translation of the bucket teeth. Additionally, the resolution is further degraded by the nonlinear transmission of angle through the non-parallelogram four-bar linkage, which is evident in the total angles swept out by each marker. As a result, the error inherent in the sensor system was magnified by the low-resolution relationship between key marker pose input and bucket tooth pose output. Mounting the key marker further from the linkage's axis of rotation may have helped alleviate the resolution issue, but would likely have caused the marker to extend an impractical distance from the excavator. Thus, an alternative sensor was sought to provide a measure of bucket angle with higher resolution.
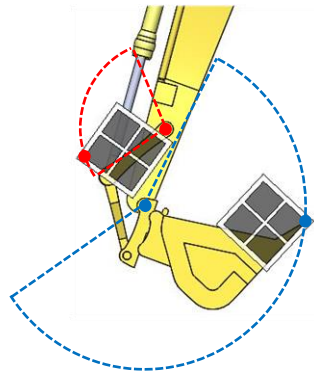
Figure 2.21 Smaller radius and swept angle of the key marker

2.4.5 Cable Potentiometer Generalized Mapping Design

The cable-driven rotary potentiometer, sometimes referred to as *string potentiometer*, was identified as an alternative sensor capable of measuring bucket angle with sufficient resolution. In addition to its resolution, the cable potentiometer was identified as an economical, robust sensor that could easily be installed on an excavator without permanently modifying the machine. Specifically, the cable potentiometer could be mounted on the bucket's hydraulic cylinder to provide a measure of stroke length, which could then be used to determine tooth pose through the generalized mapping approach.

Cable potentiometer generalized mapping is accomplished much the same as bucket linkage generalized mapping. A calibration process is performed in which the bucket is moved through its range of motion and measurements of stroke length are mapped to the relative transformation between the stick marker and bucket teeth, as shown in Figure 2.22. After the calibration process, the marker is removed from the bucket. The system then functions by measuring stick marker pose and cylinder stroke length, and mapping such measurements to bucket tooth pose.
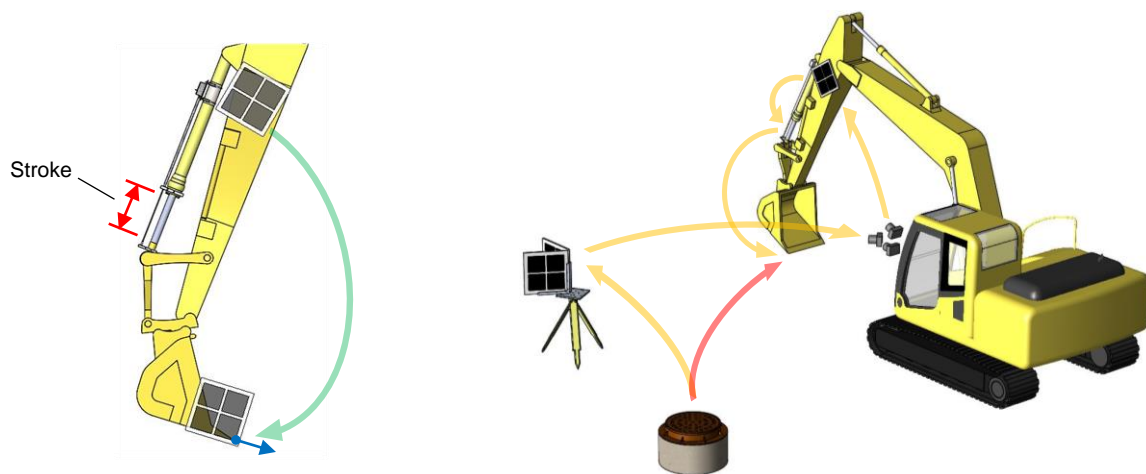
Figure 2.22 Calibration and implementation of cable potentiometer generalized mapping

A prototype was designed for the purpose of testing the cable potentiometer generalized mapping concept, as shown in Figure 2.23. The prototype attaches to the bucket's hydraulic cylinder through the use of clamps to avoid permanent modification of the machine. The device includes a cable potentiometer for measuring linear displacement, a microcontroller for signal conversion, a radio for wireless transmission, and a battery for power, all mounted inside an enclosure for protection.
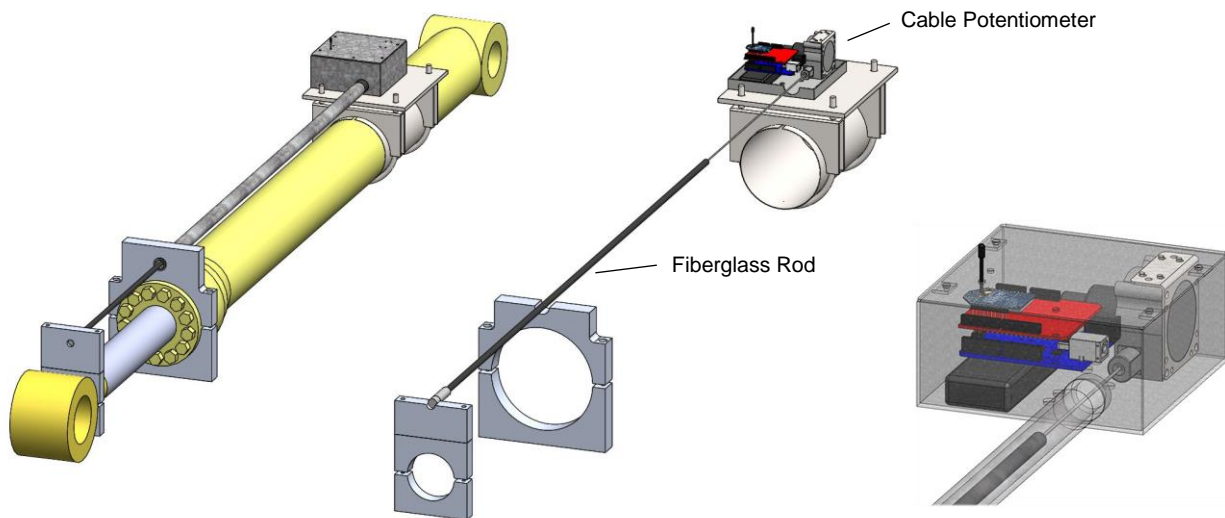
Figure 2.23 Cable potentiometer generalized mapping design

The sensor's cable is connected to a fiberglass rod, which telescopes through a rigid pipe lined with plastic to reduce abrasion. If brought into contact with a foreign object, the rod's compliant properties allow it to temporarily deflect and return to its natural shape. In the case of failure, however, the rod is expected to shear at the pipe and prevent the cable from being extracted beyond its physical limit. In this way, the fiberglass rod protects the cable potentiometer by acting as a sacrificial component. Additionally, the rod's low strength also prevents damage to the hydraulic cylinder by limiting the forces transferred to the clamps.

A cable potentiometer generalized mapping prototype was built, installed, and tested on a Caterpillar 430E IT Backhoe Loader, as shown in Figure 2.24. The experiment served three primary purposes: to determine whether the prototype functioned mechanically as intended, to measure the accuracy of the entire pose estimation system, and to obtain open-ended feedback from a professional excavator operator about the prototype. The system's accuracy was approximated by measuring three of the six components of bucket tooth pose: x-, y-, and z-position. The markers used in the experiment were made from AprilTags developed by Olson [52], and measured 58.4 centimeters (23.0 inches) between outermost black edges. The cameras used were 1.3 megapixel Firefly® MV optical cameras by Point Grey with Fujinon DF6HA-1B 6mm F/1.2 fixed focal lenses. The cameras were mounted on a tripod near the operator's cabin, approximately 1.5 meters (5 feet) above the ground. The experiment involved calibrating the sensor system, placing the backhoe in random poses, using the system to estimate bucket tooth position, and comparing estimates with ground truth measurements obtained using a total station. The experiment's pass/fail criterion was set at 2.5 centimeters (1 inch) of absolute error.

Figure 2.24 Cable potentiometer generalized mapping prototype installed on backhoe

A total of eight trials were conducted. For each trial, three components of tooth position (x, y, and z) were measured and compared with ground truth measurements. The three components of position were defined by a local orthogonal coordinate system, with y corresponding to the zenith direction, and x and z arbitrarily oriented in the level plane. Nearly all the data collected satisfied the pass/fail criterion of 2.5 centimeters (1 inch) of absolute error, as shown in Figure 2.25. Although x- and z-position estimates only marginally met the requirements, performance as a whole was deemed satisfactory since the system exhibited exceptional accuracy in the zenith direction (y-position), which is typically the most critical in excavation applications.
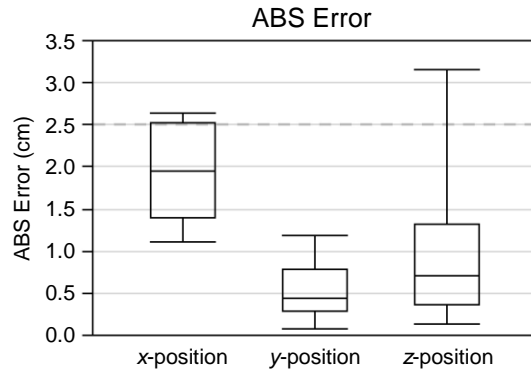
Figure 2.25 Cable potentiometer experimental results

The cable potentiometer generalized mapping prototype was also tested on a PC138USLC Komatsu Excavator operating on an active construction site, as shown in Figure 2.26. A computer screen was mounted in the operator's cabin, providing a display of bucket height relative to a desired trench grade. The system was used to assist the operator in trenching operations.

Figure 2.26 Cable potentiometer prototype installed on Komatsu excavator

Shown in Figure 2.27 is a short section of trench in which the operator conducted a side-by-side comparison of traditional human-guided grading versus grading guided by the sensor system. After performing the excavation, the operator visually compared the depths of the two surfaces and judged the sensor-guided grading satisfactory in their expert opinion. After several days of function, the prototype incurred an impact to the fiberglass rod. The rod sheared off at the pipe opening as intended, and was thus considered a controlled failure. The failure, however, marked the end of the experiment.
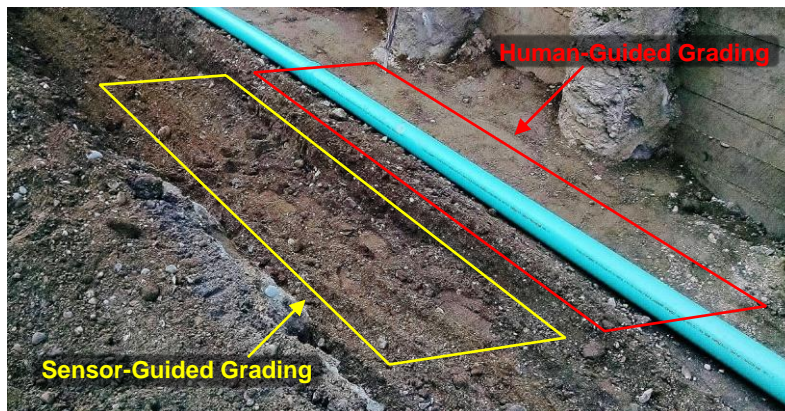
Figure 2.27 Comparison of sensor guided grading versus traditional grading

2.5 Discussion

The experiments described here led to several major findings. First, it was suggested plausible by the camera-marker experiments that cameras and markers could be capable of providing sufficient accuracy for excavator pose estimation applications. Second, it was indicated by the camera-marker experiments that marker yaw angles near 0 and 90 degrees should be avoided in pose estimation applications. Third, it was found that although purely marker-based sensor systems may be sufficiently accurate for excavator pose estimation, it may be more practical to employ a combination of marker-based sensors and traditional sensors so as to exploit the strengths of each type. Lastly, it was demonstrated through physical experiments that cameras and markers have the capability to provide actual utility in the monitoring and guidance of construction equipment, especially in terms of reducing cost and providing functionality in GPS-denied areas.

This work is significant because it serves as a guide for the development of future marker-based pose estimation systems for excavators and other construction equipment. For example, this work describes several fundamental configurations for mounting cameras and markers on a machine and a jobsite, it describes lessons learned with regard to implementation, and it identifies future work needed. Most importantly, this work is significant because it demonstrates the feasibility of using marker-based sensors for the pose estimation of excavators and other equipment in real-world construction applications.

2.6 Conclusions and Future Work

Excavator pose estimation systems offer the potential to improve excavation safety and performance, but their widespread adoption is impeded by the high cost and limitations of traditional excavator pose estimation systems. A marker-based sensor system, composed of

fiducial markers and optical cameras, offers a potential low-cost, GNSS-free alternative. However, because markers cannot be placed directly on an excavator's bucket for occlusion and durability reasons, a device is needed to transmit bucket motion and provide information necessary to deduce tooth pose.

Several design iterations were developed to address this problem, including a four-bar linkage design, a synchronous belt design, a bucket linkage generalized mapping design, and a cable potentiometer generalized mapping design. Prototypes were fabricated, installed, and tested. Promising experimental results suggest the feasibility of marker-based sensor technology for excavator pose estimation.

Marker-based sensors appear to offer a potential solution for the economical localization and pose estimation of articulated construction equipment, especially in environments with limited GNSS capabilities, but their specific application and implementation remain unclear. Although several variations of marker-based pose estimation systems were demonstrated with reasonable success, a robust system has yet to be developed for use on actual construction sites.

Future work is needed to identify the proper implementation of marker-based sensors in excavator pose estimation systems. For example, it is still unclear whether it is best to place cameras on an excavator in view of a jobsite, or on a jobsite in view of an excavator. One potential advantage of the former is the centralization of power, sensors, processors, and interface, while a potential advantage of the latter is the possibility of viewing multiple machines with the same camera. Also regarding implementation, experiments have indicated that it may be advantageous to employ a combination of marker-based sensors and traditional sensors, but the best combination and specific role of each sensor type remains unclear.

Another challenge needing to be addressed is camera-marker occlusion. Considering the busy nature of a construction site, it is not unreasonable to expect that a person, machine, or object could occasionally block the line of sight between a camera and marker. Thus, it may be necessary to evaluate the severity of the concern and address it, possibly through the incorporation of redundancies. Similarly, as an excavator moves, it may prove challenging to keep markers in the cameras' field of view. Again, one solution might exist in the use of multiple site-based cameras or multiple site-based markers for redundancy.

Lighting conditions may also pose a challenge, as low light levels can lead to marker detection loss. Thus, some form of active lighting may prove necessary for nighttime operation, and possibly even during dawn and dusk or in poorly lit indoor environments. Lastly, the harsh environment of the typical jobsite may present challenges for the sensor system's components. Thus, a fully functional system will likely require considerable design efforts to ensure electrical and mechanical robustness.

Chapter 3
BIM-Based Construction Workpiece Perception and Modeling

3.1 Introduction

The construction industry is often considered an industry of slow change, hazardous conditions, old technology, and stagnant productivity levels. Robotics offers the potential to change that by reducing construction project cost, shortening project lead time, improving construction quality, and improving worker safety [7]. However, the construction industry's adoption of robotics has proven slower than other industries, such as manufacturing. This is largely attributable to technological challenges arising from the unique characteristics of the construction industry [34]. One such challenge is the construction robot's need to perceive its workpieces and adapt its plan in order to reliably perform quality work. The objective of this research is to develop a means by which a construction robot can perceive and model the workpieces in its immediate environment so it can ultimately adapt its plan and autonomously perform detailed construction work.

The remainder of the chapter is outlined as follows. Section 3.2 further expands upon the challenges confronting construction robots and the importance of scene understanding in robotized construction work. Section 3.3 describes the relevant prior work and the gaps addressed by this chapter. Section 3.4 outlines the central contributions of this chapter. Section 3.5 describes the technical approaches employed in this work. Section 3.6 describes the experiments and results used to evaluate the effectiveness of the technical approaches. Section 3.7 provides a discussion

of the experimental results. Section 3.8 provides a summarizing conclusion and outline of future work.

3.2   Need for Scene Understanding and Adaptive Manipulation

To illustrate the construction robot's need to perceive its workpieces and adapt its plan to perform quality work, it is instructive to contrast the construction industry with the manufacturing industry. Although today's manufacturing robots possess a variety of sensing capabilities, many are able to perform work with little or no sensing of their workpieces. This is typically made possible by reducing stochastic variation through tight process and environmental controls. As such, a manufacturing robot can neglect stochastic variation and simply estimate the relative pose (i.e., position and orientation) between its tool and a point of interest on its workpiece (hereafter referred to as *tool-to-point-of-interest pose*) from a kinematic chain, as illustrated in Figure 3.1 and Equation 3.1, where T_b^a is a homogenous transformation matrix describing frame b in frame a.
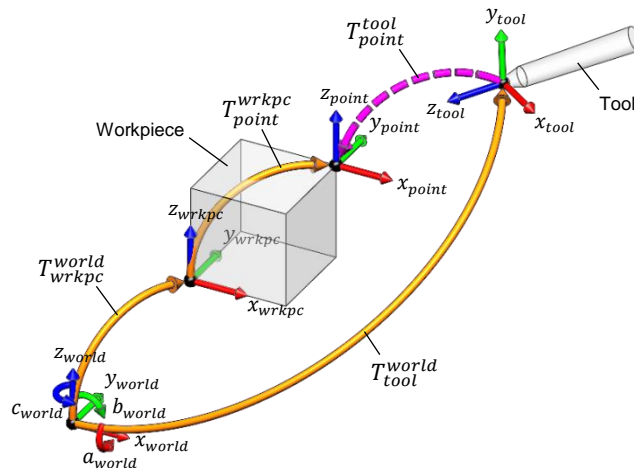
Figure 3.1 Kinematic chain used to determine the relative pose (dashed magenta arrow) between a robot's tool and the point of interest on its workpiece

$$T_{point}^{tool} = \left(T_{tool}^{world}\right)^{-1} T_{wrkpc}^{world} T_{point}^{wrkpc} \tag{3.1}$$

48

In order for a robot to reliably perform value-added work, the uncertainty in tool-to-point-of-interest pose (in addition to other sources of uncertainty) must be less than the allowable process uncertainty [53]. This requirement is described by the law of propagation of uncertainty, as shown in Equation 3.2, where the standard deviation is taken as the measure of uncertainty, Σ is the covariance matrix, and J is the Jacobian matrix which projects uncertainty into a common frame.

$$\Sigma_{tool}^{world} + \Sigma_{wrkpc}^{world} + J_{wrkpc}^{world} \Sigma_{point}^{wrkpc} J_{wrkpc}^{world \top} < \Sigma_{process,allow}^{world} \qquad (3.2)$$

The manufacturing robot, its workpieces, and its environment are all tightly controlled, which translates to low variation and low uncertainty. Low uncertainty in tool pose, workpiece pose, and workpiece geometry results in low uncertainty in tool-to-point-of-interest pose. Since the uncertainty in tool-to-point-of-interest pose is less than that allowed by the process, the manufacturing robot is able to estimate tool-to-point-of-interest pose directly from the kinematic chain and still perform work of sufficient quality and reliability.

A closer inspection reveals how such sources of uncertainty are managed in manufacturing. Uncertainty in world-to-tool pose is often managed through such means as rigid robot anchoring, precise robot installation or calibration, rigid robot links, and high precision joint sensors. Uncertainty in world-to-workpiece pose is managed through the controlled delivery of the workpiece to the robot. Fine control can often be achieved through the positioning, orienting, and holding of the workpiece via mechanical constraints such as stops, clamps, and jigs. Uncertainty in workpiece geometry (i.e., workpiece-to-point-of-interest pose) is managed through tight design tolerances and tight process controls. This allows the workpiece's geometric errors to be neglected when it arrives at the robot, despite the accumulation of errors from previous manufacturing steps. Additionally, the manufacturing workpiece is typically rigid enough that material deflections

remain small and the workpiece's designed geometry remains sufficiently representative of its true geometry. Lastly, uncertainty in workpiece geometry is also managed through definition. Albeit a different form of uncertainty, the standard practice of fully defining the manufacturing product helps to reduce ambiguity about workpiece geometry.

A look at the construction industry, on the other hand, reveals that large uncertainties in the robot, its workpieces, and its environment generally preclude the use of the kinematic chain in Equation 3.1 for estimating tool-to-point-of-interest pose. Estimating tool-to-point-of-interest pose from a kinematic chain is likely to prove ineffective for a construction robot due to large stochastic variation in the construction process and environment. If the construction robot were to estimate its tool-to-point-of-interest pose from the kinematic chain, then the combined uncertainty of the chain's components would generally exceed the allowable process uncertainty, and the construction process would be in violation of Equation 3.2. Violation of Equation 3.2 implies insufficient reliability in the construction process, or the inability to perform quality work reliably [53].

A closer inspection of the construction robot sheds light on the uncertainty present in its kinematic chain. Uncertainty in the world-to-tool pose arises from several sources. First, the sheer scale and nature of the typical construction project suggests that the construction robot will most likely need to be mobile. Although the pose uncertainty of a mobile robot's base varies depending on the sensors and estimation method used, such uncertainty is generally much larger than that of an anchored robot. A second consequence of the robot's mobility is that its base is likely to be less rigid than that of an anchored robot. Some uncertainty might be reduced through such means as outriggers, but the rigidity of a mobile robot's base is generally less than that of an anchored robot base. Third, since the construction robot must go to the workpiece to perform work, as opposed to

having the workpiece brought to it, it is expected that the construction robot will generally require a longer reach than the manufacturing robot. Thus, the need to reach long distances and exert large forces, combined with size restrictions imposed by the need to access tight spaces, suggests that the construction robot will likely have greater deflection and uncertainty in its links.

Uncertainty is also present in the world-to-workpiece pose estimate. Although construction tolerances and manufacturing tolerances both vary widely, construction tolerances are generally looser [9]. As a result, there tends to be greater variation in the placement of construction components. In addition to uncertainty caused by loose tolerances, many construction details are not explicitly defined during the conventional design process, which has the potential to contribute considerably to uncertainty in the world-to-workpiece pose estimate. For example, consider the location of a stud in a wall. Although it is common practice to define the spacing between studs in construction plans, it is not common practice to define where the stud spacing begins, so a typical stud may actually be found up to 20 centimeters (8 inches) in either direction from its expected location.

Lastly, uncertainty is also present in the workpiece's geometry (i.e., workpiece-to-point-of-interest pose). It is common for some construction components to deviate from their designed geometry as a result of deflection, often due to their material properties and considerable length. Similarly, uncertainty arises from geometric variation in construction materials. For example, lumber often has such natural defects as knots and wanes, such warpages as bows and cups, and such manufacturing defects as skip marks and raised grain [54]. Workpiece uncertainty also arises from the accumulation of errors produced during previous construction processes. That is, the construction industry's loose tolerances often result in large process errors, and the product of one process often becomes the workpiece of a later process. So although a construction feature (e.g.,

window opening) may have been constructed in the correct location, its geometric errors (e.g., width, height, perpendicularity) can still be critical to subsequent construction processes (e.g., window installation).

Thus, given the large uncertainties faced by robots in the construction industry, it is expected that construction robots will need to employ different tool-to-point-of-interest pose estimation approaches than those typically employed by manufacturing robots. To that end, this work proposes new methods for the sensing and modeling of construction components to enable a construction robot to perceive and adapt to unexpected circumstances.

3.3 Prior Work

Past research has been conducted in the sensing and modeling of construction objects and environments. Such researchers employed various sensing and modeling techniques for the purposes of project progress monitoring, as-built documentation, material quantity estimation, material tracking, obstacle avoidance, and object manipulation.

Cheok et al. [17] demonstrated that a scanning laser range finder can be used for progress monitoring and planning in earthmoving applications. As an example, the authors used a scanning laser range finder to generate a 3D point cloud of a sand pile's surface so as to model the pile's geometry and estimate its volume. However, the method, as it was employed, was not well suited for autonomy because it required manual sensor placement at various locations around the pile, manual cropping of the data to remove background objects, and manual registration of the point clouds to form an aggregated dataset.

To circumvent the challenges and computational burden associated with fitting models and extracting meaning from dense point clouds, Cho et al. [18] proposed leveraging human cognition to assist in the modeling of a construction environment. In this approach, a human selects a geometric primitive (e.g., cuboid or cylinder) or a complex model (e.g., excavator or bridge) from a database and manually aims a laser rangefinder at strategic points on the physical object to obtain individual point measurements. The selected model is then fit to the sparse point cloud. Cho et al. [19] [20] used such an approach to measure street poles, a trench, and two connected pieces of pipe. McLaughlin et al. [21] used the same approach, but developed three additional construction workspace model types: partitions, convex hulls, and tight-fitting bounding boxes. Kim et al. [22] [23] applied the aforementioned partition and convex hull methods for such purposes as obstacle avoidance, accident prevention, and material tracking. However, all such variations of this sensing approach rely on human operation, which is not suitable for autonomous applications.

Kim et al. [29] later demonstrated that convex hull modeling could be used in combination with a more automated form of sensing, flash laser distance and ranging (flash LADAR), in which a 2D array of measurements can instantly be obtained by pointing a flash LADAR sensor in the general direction of interest. The approach was only demonstrated with large objects and coarse models for such purposes as obstacle avoidance, and did not include semantic recognition of the objects it detected. Thus, the approach is not suitable for the autonomous recognition and manipulation of detailed construction components.

Stentz et al. [24] developed an autonomous excavator for excavating soil and loading it into a parked truck. The excavator used two laser rangefinders for sensing the soil terrain and the nearby truck. A 2D grid of height values was fit to the terrain data, and a parametric truck model was fit to the truck data. The terrain model was used to determine where to remove soil, and the truck

53

model was used to determine where to dump the soil. The approach reportedly worked well, but was only demonstrated with large objects and coarse models. The approach may not be suitable for the recognition and manipulation of detailed construction components.

Work has also been done in the sensing and modeling of small, targeted construction components. Kim et al. [25] sensed and modeled individual stones as a means of performing quality control for the production of construction materials such as asphalt. The authors scanned the stones with a combination of projected laser beam and charged-coupled device (CCD). Although individual objects were targeted and modeled in this work, only coarse models were applied. The highly irregular stones were simply modeled as cuboids of corresponding principle dimension (i.e., tight-fitting bounding boxes). Kahane and Rosenfeld [26] used a similar projected laser and CCD to measure the gap width between adjacent wall tiles to aid in the placement of tiles by an autonomous tiling robot. Although the gaps were targeted and measured accurately, they were only modeled by a single parameter: gap width. Such simple models may not be sufficient for detailed construction manipulation tasks.

Authors like Sicard and Levine [27] and Kim et al. [28] employed a method called polygonal approximation and syntactic analysis to extract 2D models from weld joint data obtained from a projected laser beam and CCD. In this approach, data points were evaluated one by one and combined into linear segments via polygonal approximation. The segments were then merged according to prescribed syntactic rules used for describing weld joints. However, the syntaxes developed in such studies were only capable of handling a handful of simple joint types because the complexity of segmentation and syntactic rules tends to increase with increasing component complexity. Furthermore, such methods are susceptible to failure under conditions of partial object occlusion and incomplete sensor data.

All of these studies suffer from limitations which inhibit their employment in the robotic perception and manipulation of detailed construction components. Such limitations include the need for human intervention, the lack of semantic object recognition, the lack of modeling detail necessary to perform detailed work, the inability to handle a variety of complex components, and the susceptibility to modeling failure under partial object occlusion and incomplete sensor data. The research described herein seeks to address such limitations by enabling a construction robot to autonomously perceive its workpieces through the fitting of complex models to sensor data with sufficient fidelity to ultimately adapt its plan and perform detailed construction work.

3.4 Research Contribution

This chapter offers five central contributions. First, this chapter provides insight into the underlying sources of uncertainty confronting the construction robot and driving the need for adaptive perception and modeling. Second, this chapter introduces a technique referred to as the Clustering and Iterative Closest Point (CICP) construction component model fitting technique, whereby a clustering segmentation algorithm and an Iterative Closest Point (ICP) search algorithm are applied to fit a model of a construction component to a robot's sensor data. Third, this chapter introduces a modified search algorithm for such general purposes as fitting a geometric model to a point cloud, fitting a point cloud to another point cloud, or fitting a point cloud to a map. Specifically, the modified algorithm is called the Generalized Resolution Correlative Scan Matching (GRCSM) search algorithm and is a generalization of the Multi-Resolution Correlative Scan Matching (MRCSM) search algorithm [55]. Fourth, this chapter introduces a technique referred to as the GRCSM construction component model fitting technique, in which the GRCSM search algorithm is applied in the fitting of a construction component model to a robot's sensor data. Lastly, this chapter provides an initial investigation into the capabilities and limitations of

the GRCSM and CICP construction component model fitting techniques as model fitting tools for construction robot perception via three experiments.

3.5 Technical Approach

This work focuses on the detailed modeling of targeted workpieces for robotic manipulation. Specifically, this work employs model fitting techniques in which complete models of construction components are fit to point clouds so that contextual meaning can be directly applied to the data. Although the techniques can be applied to either 2D or 3D models and sensor data, 2D models and sensor data are used here.

Despite the existence of numerous construction tasks to which such modeling techniques could be applied, construction joint filling is used as an example application in this work due to its commonness in construction activities. Activities such as welding, caulking, drywall finishing, tile grouting, spray insulating, and pipe soldering may be classified as joint filling. For the purpose of this chapter, a construction *joint* is defined as the region where two or more construction components meet. Generally such a joint is composed of two or more construction components, the gap between the components, backing material, and the filler or sealant material. Examples include a joint required by design, such as an expansion joint, or an unintentional joint-like feature, such as a crack. Here, a *workpiece* refers to one of the construction components involved in the formation of the joint, upon which joint work is to be performed. Lastly, joint *work* refers solely to the placement of filler or sealant material into the joint, although the treatment of an actual construction joint might require additional tasks such as scraping, cleaning, masking, priming, backing, tooling, and cleaning.

3.5.1 General Setup

Many interdependent technological advancements must come together to support the successful operation of a construction robot on a real-world construction project. Rather than attempting to address all aspects of the problem at once, this research focuses on a portion of the problem, leveraging the assumption that the remaining pieces of the interdependent system are also place. The general setup of the considered environment is described as follows.
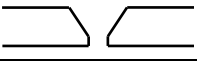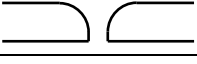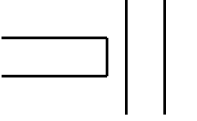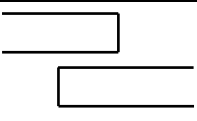
First, it is assumed that the construction robot has access to a Building Information Model (BIM) containing the designed geometry of the construction project, which is reasonable considering the rapidly expanding use of BIMs in architecture, engineering, and construction [56]. Second, it is assumed that the robot has been placed, has been teleoperated, or has autonomously driven in such a manner that it has reached the near vicinity of the object of interest. Furthermore, it is assumed that the robot's pose estimation error is small enough that it can aim its sensor in the direction it expects to find the object, and despite pose errors, still detect the actual object within the sensor's operational window. This is a reasonable assumption considering that it is not uncommon for today's indoor mobile robots to localize themselves within several centimeters and a single degree [57]. Furthermore, advancements continue to be made in sensor technology, Simultaneous Localization and Mapping (SLAM) methods, and computational power, which are likely to result in improved pose estimation in the future. Lastly, it is assumed that the sensor is capable of sensing the object of interest. In reality, highly reflective or transparent materials tend to degrade the data quality of many sensors, and specialized sensing modalities may be necessary in such cases. In the experiments described here, materials were intentionally chosen on the basis of sensor compatibility.

3.5.2 Modeling

In order for a robot to identify or manipulate a construction feature, a feature model is first needed. A construction joint is used as an example here to demonstrate feature modeling. For construction joints, the majority of the feature's descriptive information lies transverse to the joint. Thus, one modeling strategy would be to sense and model the feature in thin 2D slices, and then assemble the slices to create a 3D model. Such an approach is employed in this work.

An inspection reveals that most construction joints can be modeled by an abutment of two simple shapes in one of three configurations: butt, T (tee), or lap. Furthermore, the edges of the abutting construction components can most often be approximated by one of three edge preparations: square, bevel (V for double bevel), or flare. Shown in Table 3.1 are simple representations of some of the most common construction joint types. Given that the square butt joint is arguably the simplest and most common joint type, it is used as the illustrative example in this work.
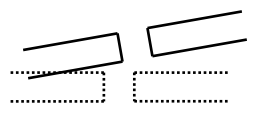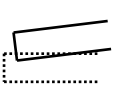
Table 3.1 Common construction joint types characterized by abutment configuration and edge preparation

Butt	Square	
	V (Vee)	
	Flare	
T (Tee)	Square	
Lap	Square	

It is worth noting that no model can perfectly describe a physical entity. Furthermore, the greater the modeling accuracy needed, the greater the model complexity required. In general, the more

degrees of freedom a model possesses, the more accurately it is able to represent its real-world counterpart. For example, if the two components of a joint model were fixed relative to one another, but were free to translate and rotate together, then a 2D model would possess three degrees of freedom in the sensor frame: $x_{combined}$, $y_{combined}$, and $\theta_{combined}$. This is illustrated for the case of a square butt joint in Row 1 of Table 3.2. If instead, full freedom were permitted in the pose of each component separately, then the model would possess six degrees of freedom: x_{left}, y_{left}, θ_{left}, x_{right}, y_{right}, and θ_{right}. An example is shown in Row 2. To achieve even greater modeling accuracy, additional degrees of freedom could be added to each component, such as preparation angle and component thickness, as shown in Row 3. Yet even greater accuracy could be obtained by including variables to describe such characteristics as edge curvature, corner curvature, and edge roughness, as shown in Row 4.

Table 3.2 Examples of varying model complexity for a square butt joint

1	Combined Pose	
2	Separated Poses	
3	Prep Angle, Thickness	
4	Curvature, Roughness	

The techniques presented in this chapter tend to be best suited for situations in which modeling accuracy requirements permit the fitting of fixed models to the data, such as situations resembling Row 1 or Row 2 of Table 3.2. For situations requiring more accuracy in the model of each individual component, such as situations resembling Row 3, techniques such as polygonal

approximation and syntactic analysis [27] may be better suited. Situations requiring high modeling accuracy, such as those resembling Row 4, are likely to require different techniques altogether.

Given that a joint's primary feature of interest is the gap formed by two construction components, a model in which each of the two component models is free to translate and rotate offers considerable descriptive capability for the gap. Since each of the two individual components has three degrees of modeling freedom, and the gap is defined by a combination of those two components, the gap model then has six degrees of freedom. The square butt joint model is shown in Figure 3.2, where it has been decomposed into an individual component model, a combined component model, a separated component model, and a gap model. Since the sensor is only expected to detect the top and inner component edges at most, the bottom edge will be omitted in subsequent models for simplicity.
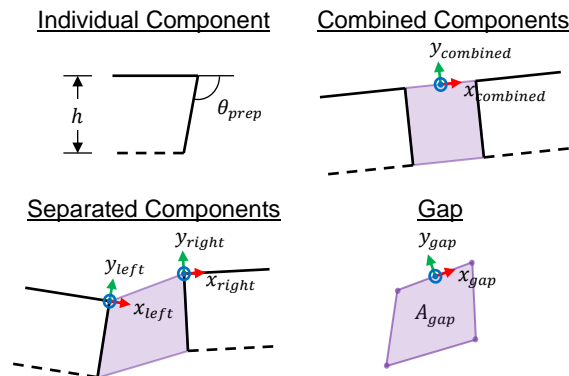
Figure 3.2 Square butt joint model decomposed into an individual component model, a combined component model, a separated component model, and a gap model

3.5.3 Sensing

A wide range of sensors could be used for the sensing of such construction components. In addition to considering the sensor modality type, such as optical, ultrasonic, or laser, one might also

consider the dimension of the output data. Since the physical construction joint exists in 3D, the joint could be sensed with a 3D sensor. However, since the majority of the joint's descriptive information lies transverse to the direction of the joint, an alternative strategy would be to sense the joint in 2D scans approximately transverse to the direction of the joint. Such a strategy is employed in this work.

The methods described in this chapter were developed for use with low precision sensors for several reasons. First, high precision sensors generally have higher associated cost, so economic feasibility may preclude the use of high precision sensors for many construction applications. As such, methods that are compatible with low precision sensors may be more impactful. Second, since there is generally a tradeoff between sensor precision and vibration or impact resistance, and since construction environments tend to be relatively rugged, robust (i.e., low precision) sensors may be favorable for many construction applications. Lastly, it is presumed that methods that are compatible with low precision sensors will also be compatible with high precision sensors, but the converse may be less likely to hold true.

3.5.4   Model Fitting

In this chapter, a distinction is drawn between two general classes of model fitting approaches for fitting models to sensor data. In the first class, geometric primitives (e.g., line segments, planes, cuboids, etc.) are fit to data points and contextual meaning is subsequently inferred from the primitives. In the second class, complete models are fit to the data points so that any contextual meaning associated with the model can be directly related to the data. Although each approach has its own pros and cons, and different approaches are likely to be more suitable for different applications, the second class modeling approaches are employed in this work.

One of the benefits of second class modeling approaches is that they do not require sophisticated rules for segmenting data, regardless of the complexity of the model. For example, the weld joint modeling approach employed by Sicard and Levine [27], which is a first class modeling approach, uses polygonal approximation to combine 2D points into linear segments. For each individual segment, points are combined until a point is reached that deviates from linearity more than a specified threshold, at which point a break is added between segments. However, if the desired model contains an arc, for example, polygonal approximation will erroneously segment the data points into several lines. Sicard and Levine [27] acknowledge that their segmentation method is not suitable for modeling curves.

Similarly, another benefit of second class modeling approaches is their avoidance of complex rules in the construction of their models. After the aforementioned segmentation step, Sicard and Levine [27] use syntactic analysis, or a pre-defined set of rules, to combine the linear segments into a complete joint model. Again, the syntactic rules are likely to erroneously merge segments along a curve during the model construction step. In addition to such segmentation issues, the complexity of syntactic rules may grow unmanageable for highly complex models, and each complex model may require a unique set of rules. Such issues are avoided altogether in the second class approaches.

The general idea behind second class modeling approaches is that effectively any arbitrarily shaped model can be fit to the data, regardless of complexity. One might appreciate such a point if they consider the extensive variety and complexity of trim board profiles found in building construction. Thousands of trim board profiles are used in building construction today [58], of which a few examples are shown in Figure 3.3. Not only may it be difficult to develop first class

modeling syntax for one such complex profile, but the difficulty may be compounded by the sheer number of available trim board profiles.

Figure 3.3 Several examples of trim board profiles found in building construction

For example, if a robot were charged with the task of caulking the joint between a piece of door casing and the wall behind it, it might be beneficial for the robot to identify the casing profile in its entirety. If the robot were simply programmed to search for a near-perpendicular inside corner representative of the joint between the casing and the wall, the robot may encounter multiple such topologies. Near-perpendicular inside corners may be present in the casing profile itself, as well as in the abutment of nearby doorframe components like the jamb and stop. In such a case, the robot would still be confronted with the need to distinguish the appropriate corner. Conversely, if a complete casing profile were fit to the data, identification of the casing-to-wall corner would become straightforward.

Another benefit of second class modeling approaches is their robustness to incomplete sensor data. In the case of first class approaches, models are only built from the data available, whereas for second class approaches, complete models are fit to the data. Because first class approaches can

only build models from the data available to them, the handling of incomplete sensor data requires additional consideration. However, it is possible to fit a model to incomplete sensor data using a second class modeling approach, as long as sufficient data is obtained.

Shown in Figure 3.4 is an example joint configuration in which it might not be possible to obtain complete sensor data. If the attribute of interest is the cross-sectional area of the gap between abutting components, it might not be possible to obtain such information directly from the sensor data. In such a case, the sensor data might be considered incomplete because the vertical depth of the gap cannot be measured directly. That is, the thickness of the abutting workpiece cannot be determined from the sensor data. However, it may still be possible to estimate the gap area if complete workpiece models are fit to the sensor data so that the designed thickness can be used to estimate the depth of the gap.
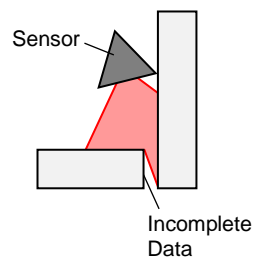
Figure 3.4 Incomplete sensor data example

Despite their benefits, second class modeling approaches do have shortcomings. Such approaches, which often function via search algorithms, are generally slower than first class approaches, which typically involve the execution of elementary operations on a finite number of data points over a finite number of passes. Second class approaches are also generally less suitable for measuring an individual component's geometric errors since the approaches involve fitting complete models to

the data. Nonetheless, the work detailed in this chapter focuses on the second class modeling approaches, namely fitting complete models to data so inherent meaning can be directly applied.

Two construction component model fitting techniques are described here for fitting geometric models of construction components to sensor data. The first construction component model fitting technique is referred to as Clustering and Iterative Closest Point (CICP) and the second construction component model fitting technique is referred to as Generalized Resolution Correlative Scan Matching (GRCSM). These construction component model fitting techniques can be generalized to 3D data and 3D models, but are presented here in the context of 2D data and 2D models. As previously noted, it is assumed that the construction robot has access to a BIM containing geometric information describing the construction components it is attempting to sense and model.

3.5.4.1 *Model Fitting Using Clustering and Iterative Closest Point*

The Clustering and Iterative Closest Point (CICP) construction component model fitting technique is composed of three primary stages: building a point cloud model, segmenting the data into clusters, and fitting the clustered data to the point cloud model. First, a point cloud model must be created to represent the construction component. This can be accomplished by converting a parametric model of the component into a bitmap representation, and then into a list of 2D or 3D points. Alternatively, the list of 2D or 3D points can be obtained by interpolating intermediate points between the model's vertices. In either case, it is important to ensure that the spacing between adjacent points satisfies the model fitting tolerance requirements. For example, if a modeling tolerance of $t = \pm 0.5\ mm$ is required, then as a first approximation, the spacing between

adjacent points in the point cloud model should not exceed $d = \|t\| = \| \pm 0.5\ mm\| = 1\ mm$ for ICP methods. Further restriction may also be necessary to satisfy orientation tolerances.

The second stage of the CICP technique involves segmenting the raw data into clusters for the purpose of isolating individual components in the data. This allows extraneous data points to be ignored and models to be fit to individual components alone. Various methods could be used to cluster the data, such as range differential thresholding. Some of the sensors applicable to the sensing of construction components provide data points in the form of bearing angles and range distances. Differential thresholding exploits the scan order of such data and divides the data wherever the absolute range difference between two consecutive points exceeds a threshold. An example threshold used by the author is three standard deviations of the sensor's uncertainty. Alternatively, methods like k-means clustering [59] and agglomerative hierarchical cluster tree clustering [60] can be used, and are recommended by the author.

After the data has been partitioned, clusters containing fewer data points than a threshold quantity are removed. It is reasoned that clusters containing far fewer points than necessary to describe the targeted component are either something other than the component itself or would be expected to describe the component poorly. In general, the elimination of clusters helps to reduce processing time by reducing the number of model fitting iterations. Furthermore, the elimination of clusters helps to remove outliers and garbage data since they tend to be clustered in small groups.

Shown in Figure 3.5 is a square butt joint clustering example in which three clusters, denoted by color, have been correctly identified in the data. The three clusters correspond to the joint's left workpiece, the joint's right workpiece, and the background object. The removed outliers and removed clusters are denoted by black 'x' markers.
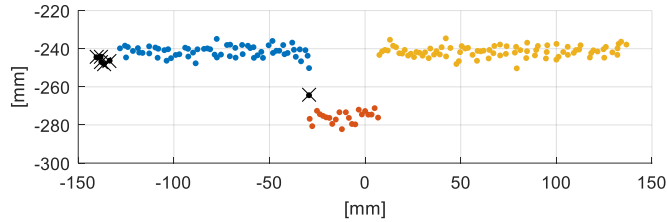
Figure 3.5 Square butt joint clustering example

In the third stage of the CICP construction component model fitting technique, point cloud models are fit to combinations of data clusters using an Iterative Closest Point (ICP) [61] point cloud registration method. Despite ICP's utility as a fast and general search algorithm for registering point clouds with unknown pairing, it suffers from the well-known problem of local minima entrapment [62]. Such issues can be alleviated, however, by initializing the search with a more accurate pose estimate. For example, in the case of construction component modeling, an initial pose estimate can be obtained from the robot's pose estimate or the centroid of the data cluster. If an accurate pose estimate is not available, then the search can be repeated over a grid of initial pose estimates and the fit of smallest error selected as the component's estimated pose.

For the case of a single construction component, the fitting process is fairly simple. Each cluster is fit one by one to the component model, and the pose corresponding to the fit of smallest error is selected as the component's estimated pose. If multiple initial poses are used for each cluster, then the fit of smallest error is selected as each cluster's estimated pose.

For the case of the construction joint, which is composed of two components, the process occurs differently. First, a combined model is built from the two component models such that their relative pose matches the BIM. Various combinations of cluster pairs are then fit to the combined model. The number of combinations is related to the number of clusters according to the common k-

combination relationship shown in Equation 3.3, where k is the component quantity and n is the cluster quantity.

$$\binom{n}{k} = \frac{n!}{k!(n-k)!} \tag{3.3}$$

The pose corresponding to the fit of smallest error is selected as the combined pose estimate. Similarly, the two clusters corresponding to the best fit are identified as the component clusters, and all other clusters are discarded. Each component cluster is then fit, one at a time, to its associated component model using the combined pose estimate as an initial guess. The resulting pose estimates are then adopted as the component pose estimates, and the joint gap is modeled directly from the transformed component models.

3.5.4.2 *Model Fitting Using Generalized Resolution Correlative Scan Matching*

A second construction component model fitting technique was developed for the fitting of models to sensor data that avoids the problem of local minima entrapment. This model fitting technique makes use of a search method called the Generalized Resolution Correlative Scan Matching (GRCSM) search algorithm, which is a generalization of the Multi-Resolution Correlative Scan Matching (MRCSM) search algorithm originally developed for robot localization [55]. Like MRCSM, GRCSM is a brute force search algorithm that uses multiple resolution levels to quickly narrow its search. However, the two methods differ in that MRCSM performs a search over a fixed number of resolution levels (i.e., two levels), while GRCSM is generalized to search over any number of resolution levels. GRCSM automatically determines and implements the appropriate number of resolution levels based on the search window size and desired fit tolerance.

The GRCSM construction component model fitting method is composed of three primary stages: determining the appropriate number of resolution levels, building a component model lookup table for each resolution level, and fitting the sensor data to the component model by searching for a best fit. The first step in determining the appropriate number of resolution levels is to determine the search-range-to-tolerance ratio r for each degree of freedom in the search, as shown in Equation 3.4, where w is the search window size and t is the loosest allowable fit tolerance.

$$r = \frac{w}{t} \tag{3.4}$$

The number of resolution levels L is then determined as shown in Equation 3.5, where s is the search test quantity per level for each degree of freedom and *ceil* is a function that rounds a number up to the nearest integer.

$$L = ceil\left(\frac{ln(r)}{ln(s)}\right) \tag{3.5}$$

As an example, imagine a 2D search that has $f = 3$ degrees of freedom (3DOF): x, y, and θ. If the search has windows of $w_x = \pm 200\ mm$, $w_y = \pm 200\ mm$, and $w_\theta = \pm 45°$; tolerances of $t_x = \pm 1\ mm$, $t_y = \pm 1\ mm$, and $t_\theta = \pm 1°$; and $s = 3$ tests per level, then it would have corresponding resolution levels of $L_x = 5$, $L_y = 5$, and $L_\theta = 4$.

For ease of implementation, the author sets the number of resolution levels equal for all degrees of freedom, namely to the largest value (5 in the example above). Additionally, the author used $s = 3$ tests per level for all experiments. Fewer tests per level generally correspond to shorter search times, as one might expect, but $s = 3$ tests per level was found to yield faster search times than $s = 2$ tests per level in practice.

The second step in the GRCSM method involves building a component model raster lookup table for each resolution level, starting with the highest resolution table. The table and its entries can be thought of as a uniform grid of cells. The dimensionality of the table corresponds to the dimensionality of the search (e.g., 2D or 3D). The size of the high resolution table is determined along each dimension as shown in Equation 3.6, where c is the number of high resolution grid cells. In practice, margins can be added to the high resolution table in integer multiples of s^{L-1}, but c should still refer to the expression shown in Equation 3.6.

$$c = s^L \tag{3.6}$$

The scale q (e.g., $1\ mm/grid$) of a cell along any dimension is given by Equation 3.7.

$$q = \frac{w}{c} \tag{3.7}$$

The component model is then projected onto the high resolution lookup table. If the model already exists as a bitmap representation, or can easily be converted to such form, then it can be directly inserted into the table, provided appropriate scaling. Alternatively, the model's vertices can be projected onto the lookup table, and a method such as Bresenham's line algorithm [63] [64] can be applied to interpolate intermediate points between the model's vertices. After a bitmap representation of the model has been established, a probabilistic distribution is applied to the model to reflect the stochastic nature of the sensor, as suggested by Olson [55]. For simplicity, the author employs a radial Gaussian distribution with standard deviation equivalent to that of the sensor's error. An example bitmap representation and its stochastic counterpart are shown in Figure 3.6.
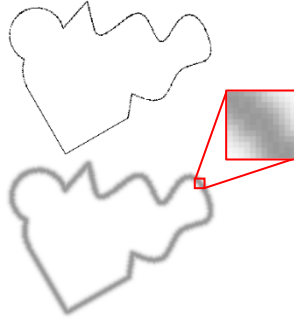
Figure 3.6 Example profile bitmap representation (top) and stochastic representation (bottom)

After the high resolution lookup table is complete, lower resolution lookup tables are produced by iteratively employing Olson's [55] technique for generating a low resolution lookup table. The generation of lower resolution tables is explained using a continuation of the aforementioned example. In such a case, the first cell of the lower resolution table is set to the maximum value found in the first $s \times s$ (3×3) block of cells in the high resolution counterpart. Similarly, the second cell of the lower resolution table is set to the maximum value found in the second $s \times s$ (3×3) block of the high resolution table, and so on. As such, each new table becomes $1/s$ ($1/3$) the size of its higher resolution counterpart. This process continues until L tables, including the high resolution table, are obtained. The lowest resolution table then becomes of size $s \times s$ (3×3). An example set of lookup tables is shown in Figure 3.7, where dashed lines have been superimposed on the lowest two resolution levels to help illustrate the process.
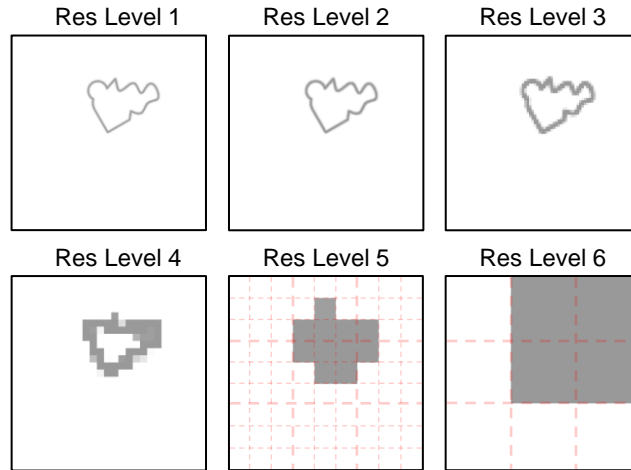
Figure 3.7 Component model lookup tables of various resolution levels

It is expected that such lookup tables may only need to be built once for construction features with uniformly designed cross sections. This may be especially beneficial for such components as construction joints, which are likely to be designed with a uniform cross section, but which actually need to be measured and modeled at many locations along their length.

The third step in the GRCSM method involves fitting the sensor data to the component model by searching for a best fit between the data and the lookup tables. Using the expected sensor-to-joint pose as a reference, the data points are projected onto the lookup table and the corresponding values are summed to provide a fit score. The process can be explained with the aid of the aforementioned example. To start, the data points are transformed and projected onto the lowest resolution table using various pose combinations. For each degree of freedom, the immediate search window i is given by Equation 3.8, where v is the current resolution level.

$$i = \frac{qc}{s^{L-v}} \qquad (3.8)$$

The immediate search window is divided into s uniform intervals, such that test points lie at the center of each interval. For the x-position degree of freedom in the example above, the test points would be located at $16.7\ mm$, $50\ mm$, and $83.3\ mm$. Similarly, s test points are also created for the other degrees of freedom. The data is then projected onto the table for each of the f^s ($3^3 = 27$) pose combinations. The table values are summed, and the total fit score is recorded for each combination. The scores are then added to a list, and the pose of highest score is removed from the list and explored at a new resolution level. If, for example, the pose $x = 83.3\ mm$, $y = 50\ mm$, and $\theta = 0°$ were found to be the best fit at the lowest resolution level, then the test points along the x-position dimension at the new resolution level would become $72.2\ mm$, $83.3\ mm$, and $94.4\ mm$. This process is repeated until the highest scoring pose found on the list coincides with the highest resolution level (i.e., Level 1). Such pose is then selected to represent the best fit between the data and the model. An example search process is shown in Figure 3.8. Also displayed are the best fit result for each resolution level leading to the ultimate best fit at the highest resolution level.
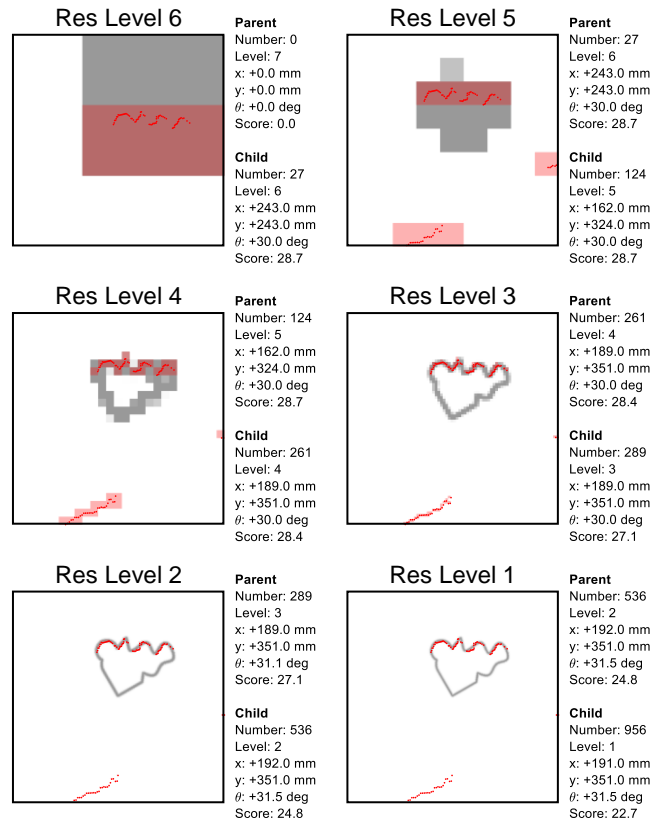
Figure 3.8 Example GRCSM search process

It should be noted that during the process, it is possible to move bi-directionally along the resolution spectrum and jump more than a single level, given that search transitions are dictated by the highest score found in the score list. Furthermore, it should be noted that the scores in the score list have both an associated pose and resolution level. That is, the same pose can be explored at more than one resolution level.

For the case of a single construction component, the GRCSM model fitting process is straightforward. However, for the case of a construction joint, which is composed of two components, the process occurs differently. First, a combined model is built from the two component models such that their relative pose matches the BIM. The GRCSM method is then applied to the combined model. The pose corresponding to the fit of smallest error is selected as

74

the combined pose estimate. The process is then repeated for each component individually, using the combined pose estimate to initialize the search. Furthermore, the search window for the individual components is limited to a fraction of the designed gap width in order to avoid erroneously fitting a model of one component to the sensor data of other components. The resulting pose estimates are then adopted as the component pose estimates, and the joint gap is modeled directly from the transformed component models.

3.6 Experiments and Results

Experiments were conducted to evaluate the effectiveness of the CICP and GRCSM construction component model fitting approaches in the fitting of construction component models to robot sensor data. Two types of construction features were sensed and modeled. The first construction feature was a single generic component of complex, arbitrary geometry. The second construction feature was a square butt joint composed of two simple components. The computational processing for all experiments was performed in MATLAB using an Intel® Core™ i7-4900MQ 2.80 GHz central processing unit.

3.6.1 Arbitrary Object: A Single Complex Component

A physical experiment was conducted to evaluate the construction component model fitting approaches' ability to model an arbitrary construction object of known, complex geometry. In this experiment, the object's designed geometry was assumed known and sufficiently representative of the actual object. As such, the construction component model fitting techniques were not used to estimate the object's geometry, but rather, the object's pose.

3.6.1.1 Arbitrary Object Experimental Setup

First, the object was designed using parametric modeling software to serve as a construction component in a BIM. The model's cross-sectional profile was then exported in a bitmap format and scaled such that each pixel corresponded to $1\,mm$ of the scaled model. The bitmap representation of the object's profile is shown in Figure 3.9, along with a plot illustrating its geometric dimensions.
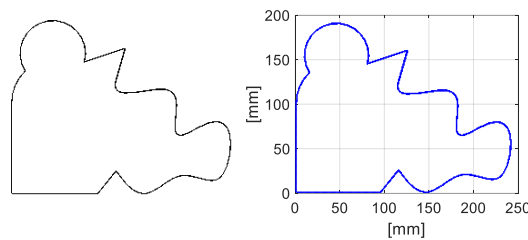
Figure 3.9 Bitmap representation (left) and dimensioned plot (right) of the complex, arbitrary object's cross-sectional profile

A physical object was then fabricated to match the model. The physical object was cut from $25\,mm$ ($1\,in$) thick foam board. The physical object was placed on a planar surface and sensed using a Hokuyo URG-04LX-UG01 laser rangefinder (2D), as shown in Figure 3.10. The relative pose between object and sensor was measured with a tape measure to obtain ground truth information. A piece of wood was placed behind the target object to simulate a background object, such as a wall. The sensor's scan window was limited to $\pm45°$ in order to target the object and reduce extraneous data. The object was moved to four different poses, and a single scan was collected at a rate of $10\,Hz$ from the laser rangefinder at each pose. Before collecting data, the sensor was allowed to warm up for more than $90\,min$, as recommended by Okubo et al. [65] for a Hokuyo URG-04LX laser rangefinder.
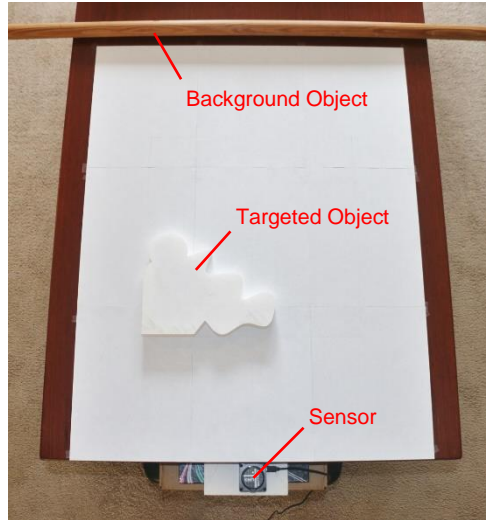
Figure 3.10 Arbitrary object experimental setup (top view)

Models were then fit to the sensor data using the CICP and GRCSM construction component model fitting techniques. An intentionally erroneous pose was fed to the model fitting algorithms to simulate such errors as robot pose error and workpiece pose error. The algorithms were then expected to determine the pose of the actual object despite the initial estimation errors. The agglomerative hierarchical cluster tree approach was used to cluster the data in the CICP construction component model fitting technique. The number of output clusters was set to three, and any clusters comprising less than 5% of the data were removed. For both the ICP search and the GRCSM search, the loosest allowable fit tolerances for x, y, and θ were set to $\pm 0.5\ mm$, $\pm 0.5\ mm$, and $\pm 0.1°$, respectively. For the GRCSM search, the test quantity per level was set to four, and the search windows for x, y, and θ were set to $\pm 364.5\ mm$, $\pm 364.5\ mm$, and $\pm 45°$, respectively. A value of $3\ mm$ was used as an estimate of the sensor's standard deviation.

The dataset was processed a third time to provide a comparison between the MRCSM and GRCSM search algorithms. To remind the reader, MRCSM was originally developed for two resolution

levels, so MRCSM was simulated in the experiment by restricting the number of GRCSM resolution levels to two, with an equal number of tests for each resolution level.

3.6.1.2 Arbitrary Object Experimental Results

The CICP and GRCSM construction component model fitting results for the four test cases are shown in Figure 3.11. As can be seen, the CICP construction component model fitting technique exhibited gross model fitting errors in two of the four test cases.
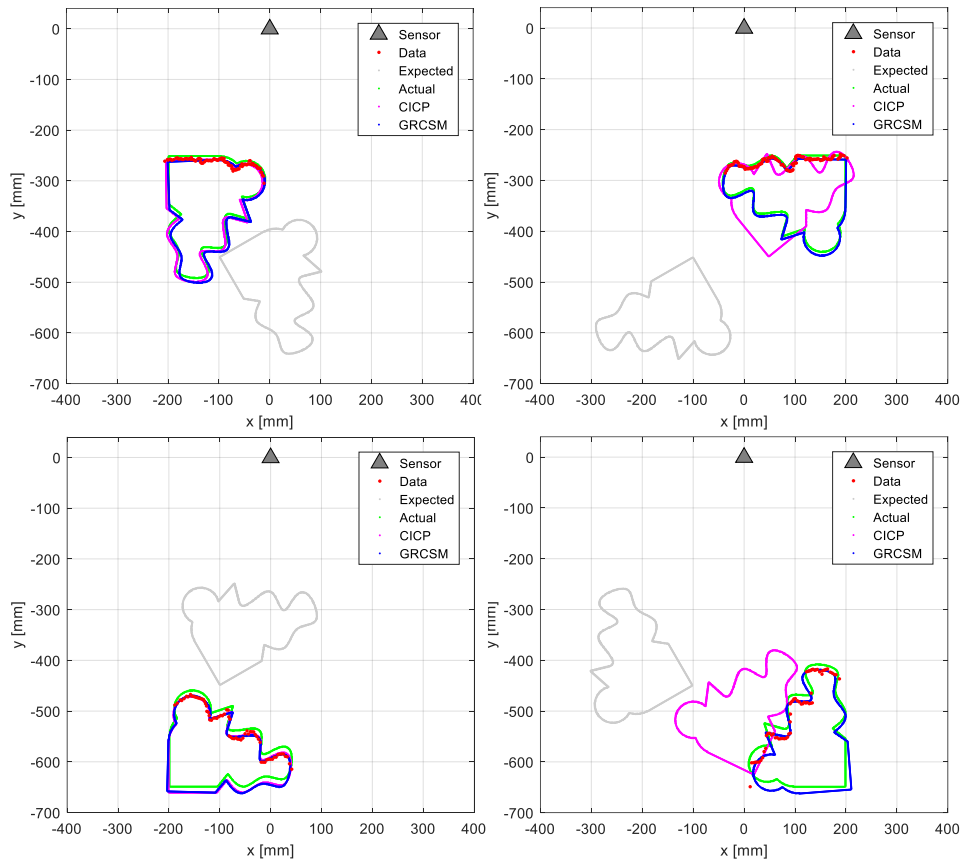
Figure 3.11 Model fitting results for an arbitrarily shaped object using the CICP and GRCSM construction component model fitting techniques

The mean processing times were found to be $1.8\ s/scan$ for the CICP construction component model fitting technique and $0.46\ s/scan$ for the GRCSM construction component model fitting

78

technique. Additionally, the mean processing time was found to be $50 \, s/scan$ when using

MRCSM to perform the best fit search, suggesting that the GRCSM search algorithm executed the

search an average of 110 times faster than the MRCSM implementation. The model fitting

performance statistics for the four test cases are shown in Table 3.3, where the statistics include

gross modeling errors.

Table 3.3. Absolute model fitting pose estimation error statistics for an arbitrarily shaped object using the CICP and GRCSM construction component model fitting techniques.

	CICP				GRCSM			
	Min	Max	Mean	Std Dev	Min	Max	Mean	Std Dev
$\|err_x\| \, (mm)$	1.7	180	85	95	1.9	15	6.4	5.9
$\|err_y\| \, (mm)$	8.3	200	61	94	5.7	11	8.4	2.4
$\|err_\theta\| \, (deg)$	0.0	140	42	67	0.9	5.2	2.5	1.9

3.6.2 Square Butt Joint: Two Simple Components

Experiments were conducted to evaluate the ability of the construction component model fitting

approaches to estimate the geometry of a construction joint. The first experiment was conducted

in a virtual environment, and the second experiment was conducted in a physical environment.

3.6.2.1 *Virtual Joint Experimental Setup*

The first joint experiment was conducted in a virtual environment to evaluate the construction

component model fitting approaches' ability to estimate the geometry of a square butt joint that

deviates from its expected, or designed, geometry. A randomized square butt joint was generated

by specifying information about the gap and the two components that comprise the joint. The

transverse parameters of the expected joint are listed below, where a, b, and c refer to

counterclockwise rotations about x, y, and z, respectively.

- Gap width (top corner to top corner): $25 \, mm$
- Workpiece thickness: $25 \, mm$
- Workpiece width: $200 \, mm$
- Expected pose in world frame: $[\bar{x} \quad \bar{z} \quad \bar{a} \quad \bar{b} \quad \bar{c}]^\mathsf{T} = [0 \, mm \quad 50 \, mm \quad 0° \quad 0° \quad 0°]^\mathsf{T}$

The joint was extended along a third dimension by uniformly spacing two-dimensional cross sections, or slices, along the world y-axis. The joint was made continuous by interpolating between particular slices designated as transition slices. The joint's longitudinal parameters were as follows:

- Slice spacing: $5 \, mm$
- Transition slice spacing: $8 \, slices$
- Transition slice quantity: $12 \, transition \, slices$

Randomized geometric deviations were applied to the transition slices as follows, where $\mathcal{N}(\mu, \sigma^2)$ represents random values drawn from a Gaussian (normal) distribution with mean μ and standard deviation σ:

- Transition slice spacing deviation: $\mathcal{N}(0, (1 \, slice)^2)$
- x-position deviation: $\mathcal{N}(0, (10 \, mm)^2)$
- y-position deviation: $\mathcal{N}(0, (3 \, mm)^2)$
- b-orientation deviation: $\mathcal{N}(0, (1°)^2)$
- Gap width deviation: $\mathcal{N}(0, (25/3 \, mm)^2)$
- Upper corner angle deviation: $\mathcal{N}(0, (3°)^2)$
- Workpiece angle deviation: $\mathcal{N}(0, (1°)^2)$
- Workpiece thickness deviation: $\mathcal{N}(0, (0.5 \, mm)^2)$
- Workpiece preparation angle deviation: $\mathcal{N}(0, (1°)^2)$

A random joint was generated from such parameters. The joint was composed of 83 cross-sectional slices and took the form shown in Figure 3.12. Dashed lines denote the joint's 83 cross-sectional slices, and solid lines denote the joint's 12 transition slices.
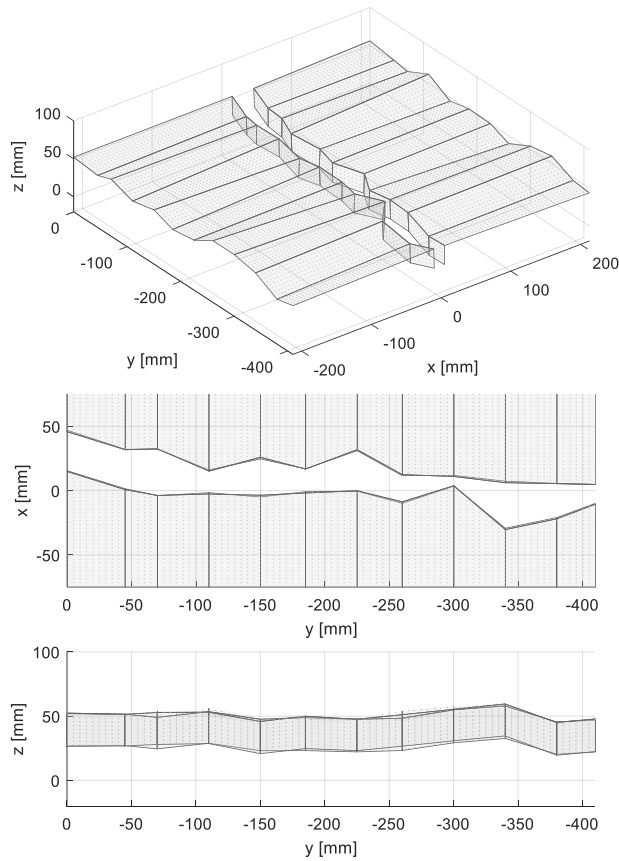
Figure 3.12 Randomized virtual square butt joint

A virtual sensor was used to generate virtual sensor data of the joint. The virtual sensor was modeled as a 2D scanning laser rangefinder that outputs range and bearing measurements. Loosely modeled after the Hokuyo URG-04LX-UG01, the sensor was assigned bearing increments of $0.352°$ ($10\ bits$) and a range standard deviation of $3\ mm$. Uncertainty in the sensor's bearing measurement was neglected. The sensor's scan window was set to $\pm30°$. The virtual sensor scanned the joint while translating along the world y-axis, $250\ mm$ directly above the expected

center of the joint. The virtual sensor data was generated by incrementing the sensor's bearing, projecting a beam onto the incident joint face, finding the point of intersection, and adding a normally distributed random range error corresponding to the sensor's standard deviation. The resulting virtual sensor data is shown in Figure 3.13.
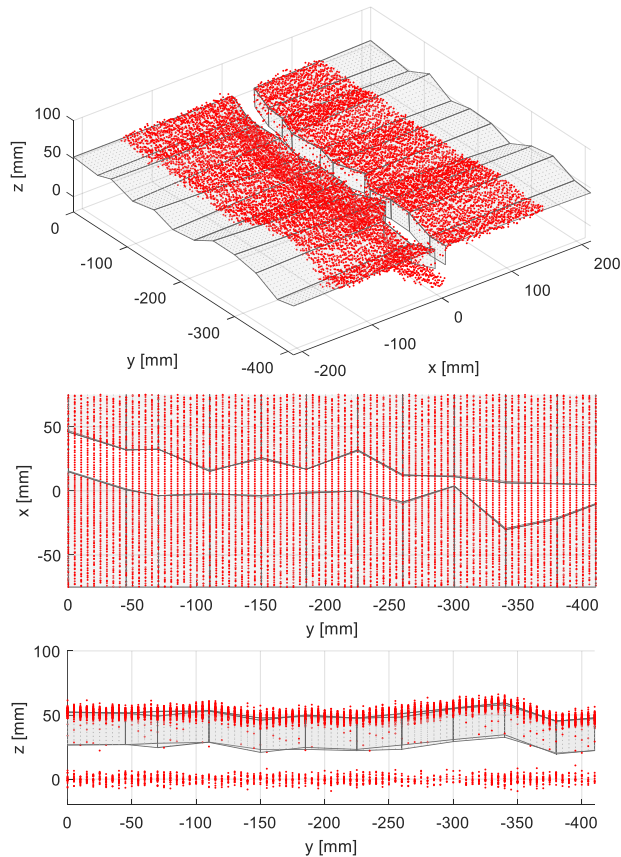
Figure 3.13 Virtual square butt joint sensor data

The construction component model fitting techniques were then applied to model the poses of the two workpieces at each slice along the length of the joint. The agglomerative hierarchical cluster tree approach was used to cluster the data in the CICP construction component model fitting technique. The number of output clusters was set to five, and any clusters comprising less than 5% of the data were removed. For both the ICP search and the GRCSM search, the loosest allowable

82

fit tolerances for x, y, and θ were set to $\pm 0.5\ mm$, $\pm 0.5\ mm$, and $\pm 0.3°$, respectively. For the GRCSM search, the test quantity per level was set to three, and the search windows for x, y, and θ were set to $\pm 121.5\ mm$, $\pm 121.5\ mm$, and $\pm 45°$, respectively.

The dataset was processed a third time to provide a comparison between the MRCSM and GRCSM search algorithms. MRCSM was simulated in the experiment by restricting the number of GRCSM resolution levels to two, with an equal number of tests for each resolution level.

Because the gap is defined by the two workpieces, the gap's pose and geometry was extracted directly from the workpiece model results. The 2D gap geometry (e.g., cross-sectional area) was converted to 3D (e.g., differential volume) by projecting the information half the distance between slices in both directions along the joint.

3.6.2.2   *Virtual Joint Experimental Results*

The gap modeling results for both the CICP and GRCSM construction component model fitting techniques are shown in Figure 3.14, where the results are represented as a collection of gap pose estimates and differential gap volumes estimates, and the gap origin is defined as the midpoint between upper corners of the workpiece models.
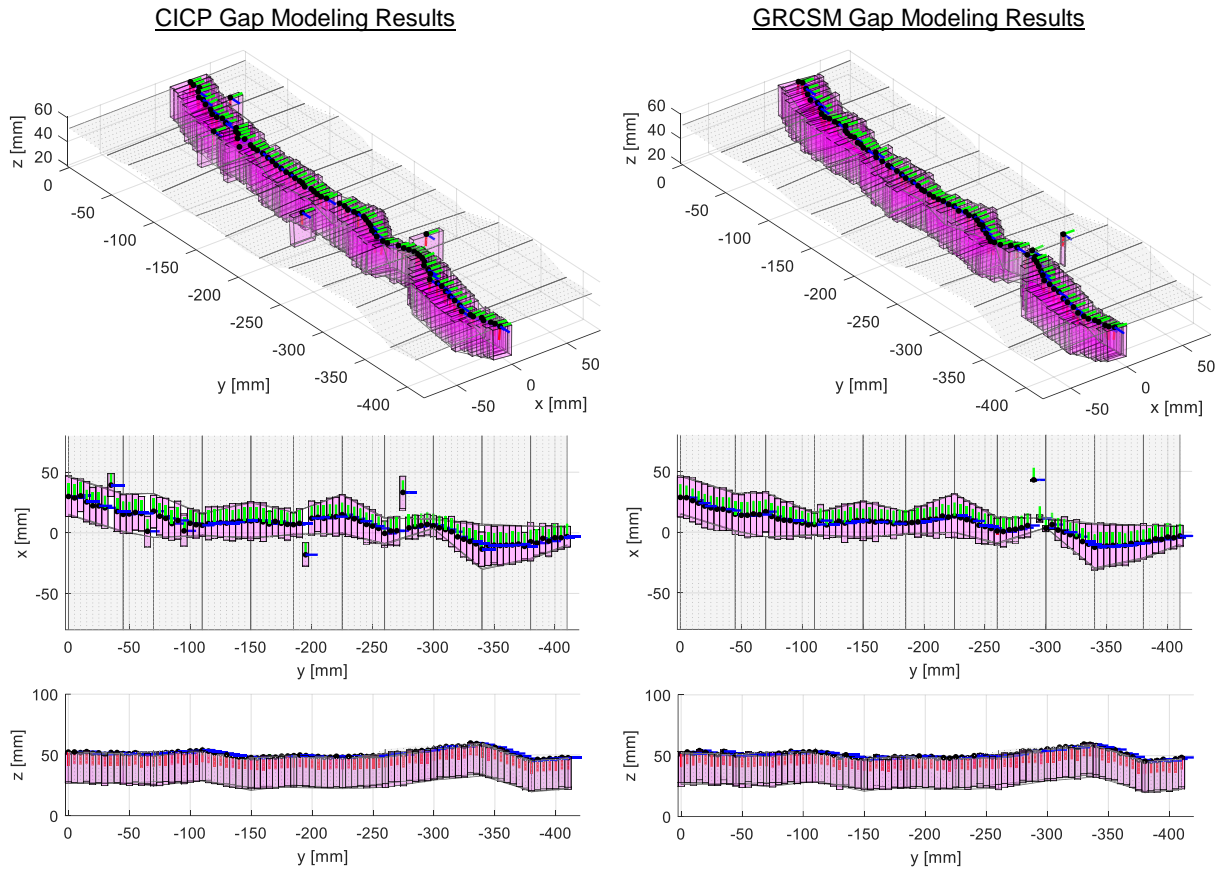
Figure 3.14 Model fitting results for the randomized virtual joint using the CICP and GRCSM construction component model fitting techniques

Shown in Figure 3.15 is a plot of the estimated poses and differential volumes of the gap at the various slices along the joint, where the pose corresponds to the midpoint between the upper corners of the workpiece models.
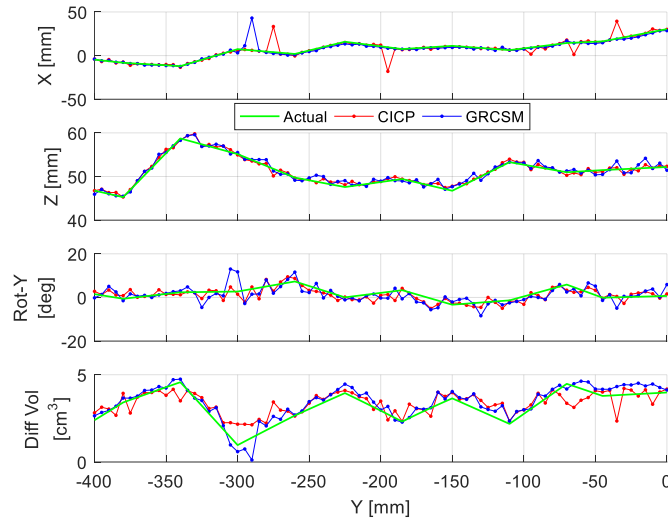
Figure 3.15 Gap pose estimates for the randomized virtual joint using the CICP and GRCSM construction component model fitting techniques

Model fitting pose estimation statistics for the virtual joint are shown in Table 3.4, where the statistics do include gross modeling errors.

Table 3.4 Absolute model fitting pose estimation error statistics for the virtual joint using the CICP and GRCSM construction component model fitting techniques

	CICP				GRCSM			
	Min	Max	Mean	Std Dev	Min	Max	Mean	Std Dev
$|err_x|\ (mm)$	0.0	29	2.0	5.0	0.0	37	1.6	4.0
$|err_z|\ (mm)$	0.0	2.0	0.5	0.4	0.0	2.2	0.6	0.5
$|err_\theta|\ (deg)$	0.0	5.5	1.7	1.3	0.1	10	2.3	2.0
$|err_V|\ (cm^3)$	0.0	1.5	0.4	0.3	0.0	1.3	0.3	0.2

The mean processing times were found to be $0.20\ s/scan$ for the CICP construction component model fitting technique and $0.082\ s/scan$ for the GRCSM construction component model fitting technique. Additionally, the mean processing time was found to be $630\ s/scan$ when using MRCSM to perform the search, suggesting that the GRCSM search algorithm executed the search an average of 3100 times faster than the MRCSM implementation.

3.6.2.3 *Physical Joint Experimental Setup*

A physical experiment was conducted to evaluate the ability of the construction component model fitting techniques to estimate the geometry of a square butt joint that deviates from its expected, or designed, geometry. A square butt joint of known geometry was cut from $25\ mm$ ($1\ in$) thick foam board, as shown in Figure 3.16. The joint was positioned such that a distance of $51\ mm$ ($2\ in$) was maintained from the nearest background object.
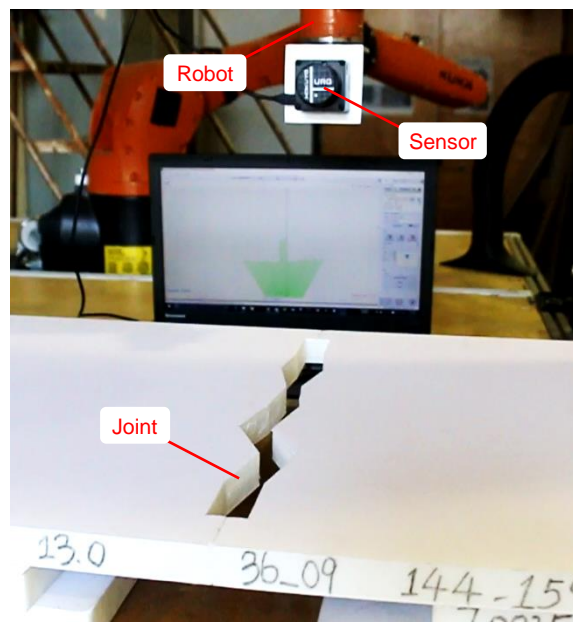
Figure 3.16 Physical square butt joint experimental setup

The joint was replicated in the virtual environment for reference during analysis. The virtual replica of the physical joint was composed of 89 cross-sectional slices, and is shown in Figure 3.17. Dashed lines denote the joint's 89 cross-sectional slices, and solid lines denote the joint's 9 transition slices.
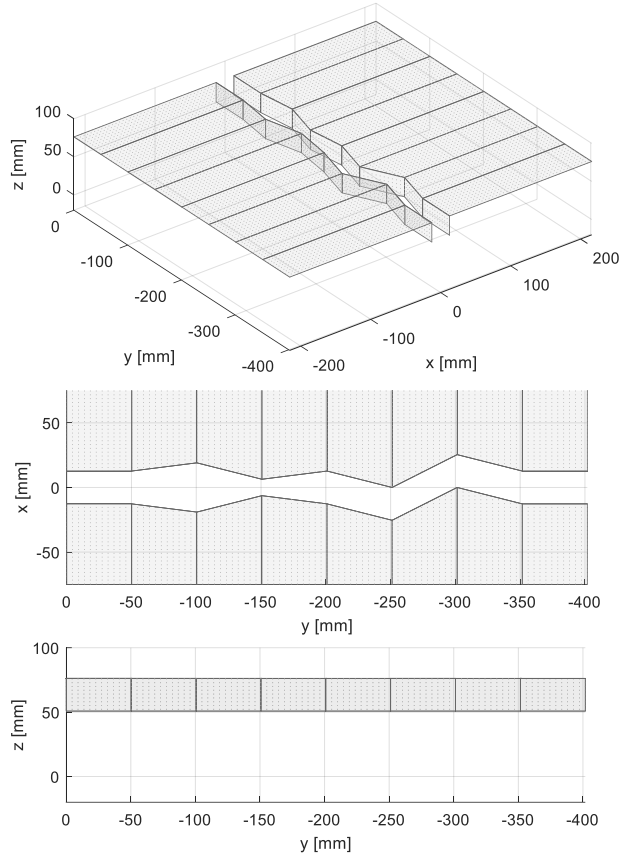
Figure 3.17 Replica of the physical joint in the virtual environment

To sense the physical joint, a KUKA KR6 R900 SIXX robot was used to move a Hokuyo URG-04LX-UG01 (2D) laser rangefinder along the length of the joint. The sensor was translated at a constant velocity of $46 \, mm/s$, at a distance of $250 \, mm$ directly above the expected center of the joint. The relative pose between object and sensor was measured with a steel rule to obtain ground truth information. Scans were obtained at a rate of $10 \, Hz$. The sensor's scan window was limited to $\pm 30°$ to target the object and reduce extraneous data. Before collecting data, the sensor was allowed to warm up for more than $90 \, min$, as recommended by Okubo et al. [65] for a Hokuyo URG-04LX laser rangefinder. The collected sensor data is shown in Figure 3.18, where each scan was assumed to lie in a plane perpendicular to the joint.

87
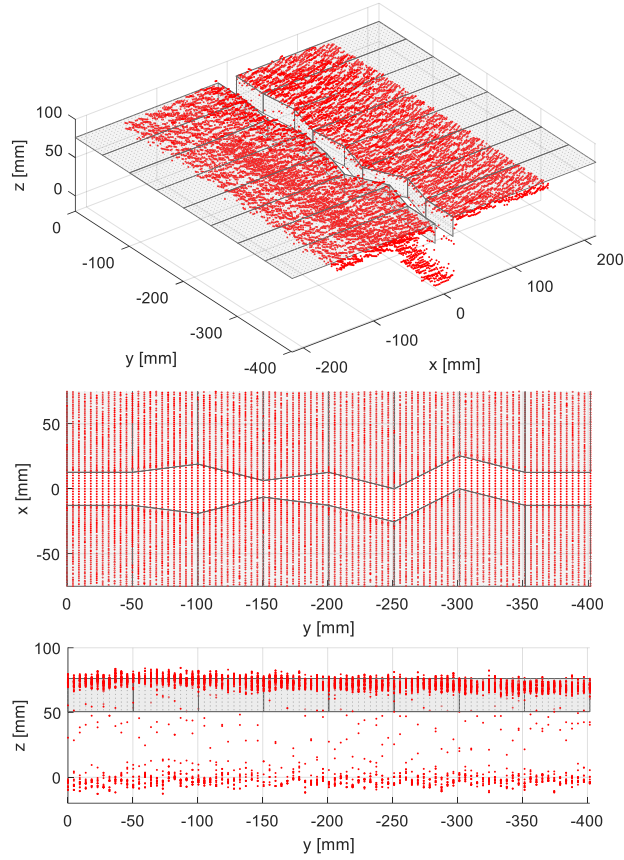
Figure 3.18 Real sensor data shown in the virtual environment for visualization

The construction component modeling fitting techniques were then applied to model the poses of the two workpieces at each slice along the length of the joint. The agglomerative hierarchical cluster tree approach was used to cluster the data in the CICP construction component model fitting technique. The number of output clusters was set to five, and any clusters comprising less than 5% of the data were removed. For both the ICP search and the GRCSM search, the loosest allowable fit tolerances for x, y, and θ were set to $\pm 0.5\ mm$, $\pm 0.5\ mm$, and $\pm 0.3°$, respectively. For the GRCSM search, the test quantity per level was set to three, and the search windows for x, y, and θ were set to $\pm 120\ mm$, $\pm 120\ mm$, and $\pm 45°$, respectively. A value of $3\ mm$ was used as an estimate of the sensor's standard deviation.

The dataset was processed a third time to provide a comparison between the MRCSM and GRCSM search algorithms. MRCSM was simulated in the experiment by restricting the number of GRCSM resolution levels to two, with an equal number of tests for each resolution level.

Because the gap is defined by the two workpieces, the gap's pose and geometry was extracted directly from the workpiece model results. The 2D gap geometry (e.g., cross-sectional area) was converted to 3D (e.g., differential volume) by projecting the information half the distance between slices in both directions along the joint.

3.6.2.4 *Physical Joint Experiment Results*

The gap modeling results for both the CICP and GRCSM construction component model fitting techniques are shown in Figure 3.19, where the results are represented as a collection of gap pose estimates and differential gap volumes estimates, and the gap origin is defined as the midpoint between upper corners of the workpiece models.
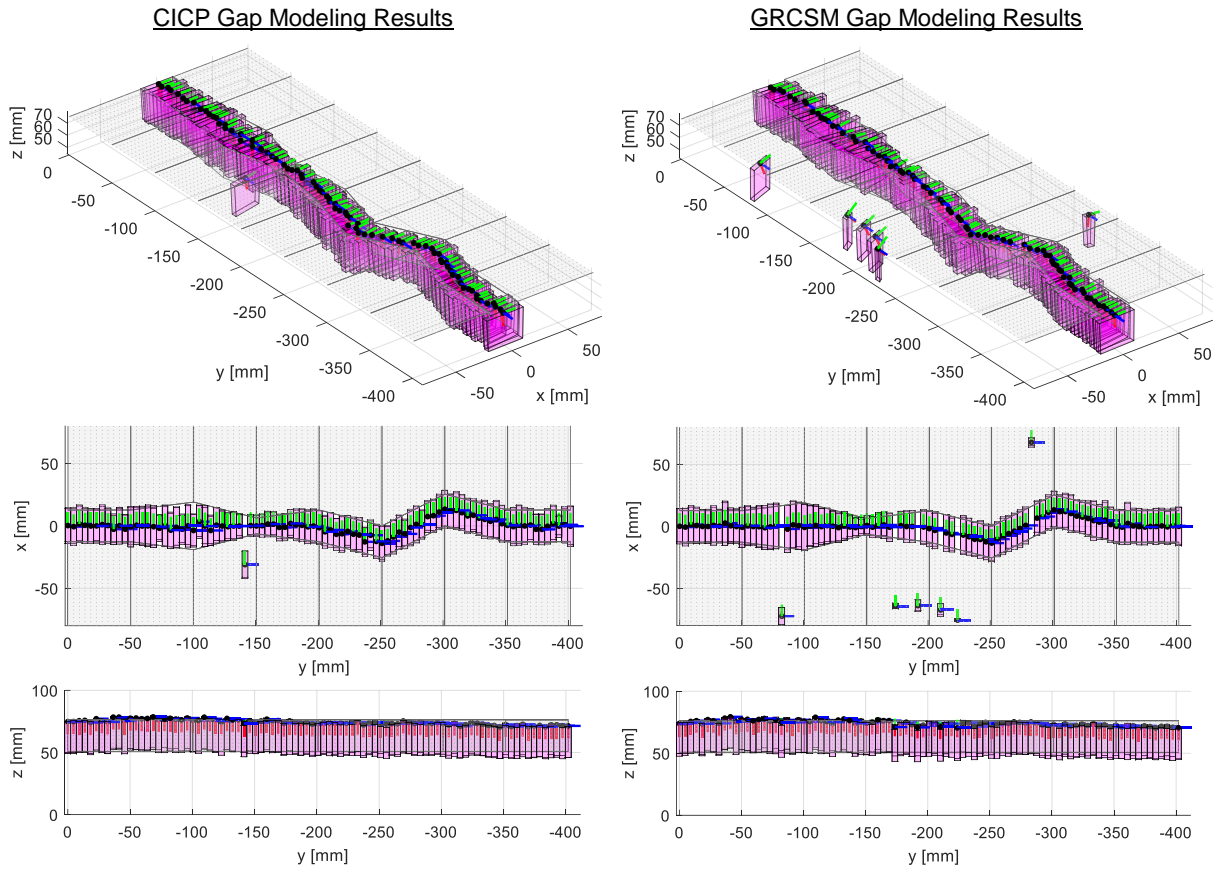
Figure 3.19 Model fitting results for the physical joint using the CICP and GRCSM construction component model fitting techniques

Shown in Figure 3.15 is a plot of the estimated poses and differential volumes of the gap at the various slices along the joint, where the pose corresponds to the midpoint between the upper corners of the workpiece models.
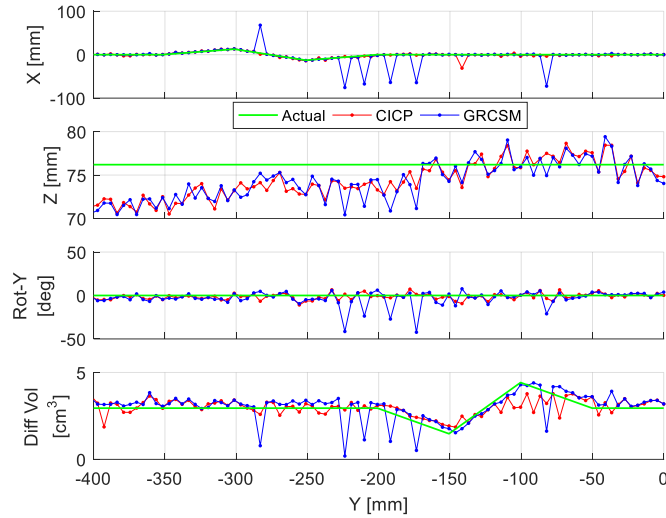
Figure 3.20 Gap pose estimates for the physical joint using the CICP and GRCSM construction component model fitting techniques

Model fitting pose estimation statistics for the physical joint are shown in Table 3.5, where the statistics do include gross modeling errors.

Table 3.5 Absolute model fitting pose estimation error statistics for the physical joint using the CICP and GRCSM construction component model fitting techniques

	CICP				GRCSM			
	Min	Max	Mean	Std Dev	Min	Max	Mean	Std Dev
$\|err_x\|\ (mm)$	0.0	31	1.5	3.4	0.0	72	5.1	17
$\|err_z\|\ (mm)$	0.0	5.7	2.3	1.5	0.0	5.8	2.4	1.7
$\|err_\theta\|\ (deg)$	0.0	11	2.9	2.3	0.0	42	5.2	7.2
$\|err_V\|\ (cm^3)$	0.0	1.5	0.3	0.3	0.0	2.8	0.4	0.5

The mean processing times were found to be $0.235\ s/scan$ for the CICP construction component model fitting technique and $0.081\ s/scan$ for the GRCSM construction component model fitting technique. Additionally, the mean processing time was found to be $490\ s/scan$ when using MRCSM to perform the search, suggesting that the GRCSM search algorithm executed the search an average of 2400 times faster than the MRCSM implementation.

3.7 Discussion

As suggested by Figure 3.11, the CICP and GRCSM construction component model fitting techniques appear capable of estimating the pose of an arbitrary construction object of known, complex geometry. GRCSM appears to outperform CICP both in terms of success rate and processing speed. Excluding cases of gross error, the model fitting accuracy of the two techniques appears comparable. A close inspection of Figure 3.11 suggests that a portion of the error is likely due to ground truth measurement, as opposed to modeling, since the results appear to fit the data better than the ground truth. This suggests that the true modeling errors may have actually been less than that reported in Table 3.3. Furthermore, the GRCSM search algorithm was found to execute the model fitting search an average of 110 times faster than the MRCSM implementation.

As indicated by Figure 3.14, Figure 3.15, and Table 3.4, the CICP and GRCSM construction component model fitting techniques also appear capable of estimating the geometry of a virtual construction joint. It is difficult to determine which model fitting technique performed best, but judging from Figure 3.14, the author is of the opinion that the GRCSM modeling results appear slightly more accurate than the CICP modeling results. GRCSM outperformed CICP in terms of processing speed, and the GRCSM search algorithm appears to have executed the model fitting search an average of 3100 times faster than the MRCSM implementation.

As indicated by Figure 3.19, Figure 3.20, and Table 3.5, the CICP and GRCSM construction component model fitting techniques also appear capable of estimating the geometry of a physical construction joint with real sensor data. As opposed to the arbitrary object and virtual joint experiments, the author is of the opinion that the CICP modeling results were slightly more accurate than the GRCSM modeling results for the physical joint experiment. However, GRCSM

outperformed CICP in terms of processing speed. A close inspection of Figure 3.18 and Figure 3.20 reveals that a portion of error along the z-dimension can be attributed to ground truth measurement error, as opposed to modeling error, which indicates that the actual z-dimension modeling errors may have been less than that reported in Table 3.5. Furthermore, the GRCSM search algorithm was found to execute the model fitting search an average of 2400 times faster than the MRCSM implementation.

This research provides two construction component model fitting techniques which enable a construction robot to perceive and model the workpieces in its immediate environment so it can ultimately adapt its plan and autonomously perform construction work. Despite the large uncertainties in robot pose, workpiece pose, and workpiece geometry, this framework enables the robot to perceive the workpiece directly and perform quality work. As opposed to past construction component modeling techniques, the techniques described herein eliminate the need for human intervention, provide a level of semantic recognition, offer sufficient modeling detail for detailed construction work, handle a variety of complex component geometries, and offer improved robustness to partial object occlusion and incomplete sensor data. Regardless of differences in performance, both the CICP and GRCSM construction component modeling techniques appear to be promising tools for the geometric estimation of construction components, especially for situations involving full automation, detailed construction work, incomplete sensor data, and complex object geometry.

3.8   Conclusions and Future Work

The objective of this research was to explore the extent to which a construction robot can perceive and model construction work components, which is a critical step in making adaptive manipulation

decisions to accomplish work. This chapter began by providing insight into the underlying sources of uncertainty confronting the construction robot and driving the need for adaptive perception and modeling. Two construction component model fitting techniques were introduced to meet such a need, namely the Clustering and Iterative Closest Point (CICP) construction component model fitting technique and the Generalized Resolution Correlative Scan Matching (GRCSM) construction component model fitting technique. Three experiments were presented to evaluate the ability of these techniques to model construction components. The first experiment involved the pose estimation of a physical arbitrary component, the second experiment involved the geometric modeling of a virtual construction joint, and the third experiment involved the geometric modeling of a physical construction joint.

It was found that the GRCSM search algorithm is significantly faster than MRCSM in executing the model fitting searches described in this chapter. It was also confirmed that the CICP and GRCSM construction component model fitting techniques are capable of estimating the pose and geometry of an arbitrarily shaped object and a construction joint. However, it was found that the CICP and GRCSM construction component model fitting techniques are both susceptible to modeling error. The occurrence of gross error appears to increase with initial pose error for the CICP construction component model fitting technique. The GRCSM construction component model fitting technique tends to exhibit increased modeling error when other objects, such as a background object, are sensed too close to the feature of interest. This is likely caused by the fact that the GRCSM construction component model fitting technique used in this work does not employ any segmentation methods to distinguish and isolate different objects in the sensor data.

Future work is needed to compare the performance of the CICP and GRCSM construction component model fitting techniques with other model fitting techniques, such as polygonal

approximation and syntactic analysis, for construction feature modeling. It may also prove beneficial to explore the use of clustering in the GRCSM construction component model fitting technique for such purposes as improving robustness to nearby objects. Additional work is also needed to evaluate how the construction component model fitting techniques handle cases with multiple similar objects or unexpected objects in the scene. Lastly, additional work is needed to convert a component model into a robot plan and physically execute the plan to evaluate the robot's ability to autonomously perform work on a construction feature. The author is currently exploring methods to perform such robot planning, as well as to explore techniques for improving plan execution through such means as incorporating feedback control.

Chapter 4
BIM-Based Motion Planning and Task Execution

4.1   Introduction

Despite the potential for robots to help alleviate many of the construction industry's problems [1-4], the very nature of the industry poses significant challenges for construction robots as compared to manufacturing environments. First, the robot must go to its workpieces instead of having its workpieces brought to it, which produces a reversed spatial conveyance between the robot and its product [8]. This and the sheer size of construction projects requires the robot to be mobile, which contributes to the spatial uncertainty between the robot and its workpieces. Second, construction materials tend to exhibit considerable geometric variation. Due to their substantial size and properties, materials are often susceptible to large deflections and geometric irregularities. Lastly, construction tolerances are relatively loose, which tends to result in considerable tolerance stack-ups. This is largely due to the fact that design tolerances, even when tightly specified, are challenging to enforce during construction, and the distributed nature of projects often makes it impractical to perform strict quality control [66].

These factors produce additional challenges for robots by contributing to pose (i.e., position and orientation) uncertainty between the robot and a point of interest on its workpiece. Hypothetically, a robot could combine its jobsite pose estimate with the as-designed pose and geometry of its workpiece to estimate the relative pose between itself and the point of interest, but such an estimate would naturally inherit the propagated uncertainties from its constituent sources, namely robot

96

pose, workpiece pose, and workpiece geometry [53]. The resulting relative pose uncertainty between the mobile robot and its workpiece would likely exceed that permitted by the construction process, thereby preventing the robot from performing work reliably. For example, in the case of a mobile robot installing a prefabricated window into a rough opening, the uncertainties in the robot's pose, the opening's pose, and the opening's actual geometry generally yield a relative pose uncertainty exceeding that allowed for window installation. If the robot were to attempt to insert the window, there would be an unacceptably high likelihood of collision between the window and the sides of the opening. More generally, if a robot were to attempt to perform work without sensing its workpieces directly, then the uncertainty in its kinematic chain would, in general, exceed the allowable process uncertainty, and the robot would be unable to perform the construction process with sufficient reliability [53]. Readers may refer to [67] for more details about relative pose uncertainties between construction robots and their workpieces.

To overcome these spatial uncertainties, construction robots must be equipped with the ability to perceive their workpieces in situ and adapt their work plans to the workpiece pose and geometry encountered. This is a critical first step toward enabling robots to become useful construction partners. The research described in this chapter builds upon previous work [67] in which the author explored how to enable a construction robot to sense and model the actual pose and geometry of its workpieces, but that body of research stopped short of addressing the planning and execution steps need for the robot to act upon its newly perceived workpiece models. Thus, the objective of this research is to develop a means by which a robot can adapt its work plan to its as-perceived workpiece models so it can overcome spatial uncertainties and perform detailed construction work.

This research and its methods demonstrate that it is possible for a construction robot to perceive the pose and geometry of its workpieces so it can adapt its plan and perform useful work. This is

expected to expand the scope of tasks that can be performed by robots, advance their capacity to assist humans in the execution of construction work, and reduce the burden of modifying construction processes and environments for robot operation. In general, this research helps bring the construction industry closer to realizing the benefits of robotics. These geometrically adaptive capabilities may be applicable to such tasks as demolition, rebar installation, formwork installation, trench backfill, material handling, prefabricated unit installation, block laying, structural member installation, welding, fireproofing, rough carpentry, pipe installation, conduit installation, ductwork installation, flooring and wall tile installation, window and door installation, caulking, painting, cabinet installation, finish carpentry, furniture installation, inspection, and others. Equipped with such geometrically adaptive capabilities, construction robots may begin extending their contributions from a purely physical capacity to a rudimentary form of cognition, such as identifying and reconciling geometric discrepancies.

Despite the applicability of such geometrically adaptive capabilities to a range of construction activities, the author uses joint filling, and in particular caulking, as the specific case study in this research to demonstrate and evaluate the robot's capabilities. However, the basic framework (i.e., collecting sensor data, registering BIM models, adapting a work plan, and executing the adapted plan) is applicable to a wide range of construction tasks and objects. With only minor application-specific modifications, such as tooling and plan details, the approaches described here could be applied to other joint filling tasks, such as welding, drywall finishing, tile grouting, and pipe soldering. With minor modifications to modeling and moderate modifications to planning, the approaches could also be applied to pick-and-place operations, such as material handling, block laying, tile installation, and window installation, although the fundamental adaptive framework could remain unchanged. Lastly, it is worth noting that this research does not attempt to address

such challenges as robot localization or navigation, but rather focuses on enabling a robot to adapt its work plan to the encountered geometry of its workpieces once it has reached their vicinity.

The specific objectives of this research are as follows.

- Establish a framework to enable a construction robot to perform geometrically adaptive, model-driven work.
- Establish a formulation to enable a construction robot to adapt its joint filling work plan to its as-perceived workpiece models.
- Evaluate a robot's ability to perform geometrically adaptive, model-driven work using custom hardware developed for robotic joint filling activities.

The remainder of the chapter is outlined as follows. Section 2 describes the relevant related work and gaps addressed by this research. Section 3 describes the technical approaches employed in this work. Section 4 describes the experiments used to evaluate the technical approaches. Section 5 provides the experimental results. Section 6 offers a discussion of the experimental results. Section 7 provides summarizing conclusions.

4.2 Related Work

Past research has been conducted into providing construction robots with the ability to adapt to the actual pose and geometry of their workpieces so they can perform work. However, such studies have limitations which preclude their methods from being adopted widely. For example, the methods used in some studies may not be applicable to tasks requiring tight tolerances or dexterous maneuvering. As an example, Stentz et al. [24] demonstrated that a robotic excavator could use laser rangefinders to adapt its plan to the topology of nearby soil and the pose of a nearby truck for a digging and dumping operation. However, the system's accuracy was on the order of several centimeters, and the demonstration was designed to avoid the need for dexterous maneuvers (e.g.,

no structures or utilities to reach under or around). Similarly, Keating et al. [30] demonstrated that a robot could print foam formwork for a cast-in-place concrete structure, where previous courses were sensed and depths of subsequent courses were adjusted accordingly. However, the spatial tolerances of the deposited foam were on the order of several centimeters, and the experiment was designed such that the robot could maintain a vertical nozzle orientation throughout the entire spray. Thus, such methods may not be suitable for construction tasks requiring tight tolerances or dexterous maneuvering.

Similarly, some studies have employed simple workpiece models or workpiece identification methods which might not be very generalizable to a range of tasks or high-fidelity models. For example, Jung et al. [68] demonstrated the rudimentary ability of a robot to identify and pick up tiles using stereo cameras and a suction gripper in a robotic tile installation operation, but the tile detection algorithms relied on identifying a square planar surface of a specific color (e.g., red, green, or blue) representing the tile, which may not be highly generalizable to a variety of tiles or construction objects. Similarly, Kahane and Rosenfeld [26] used a 2D projected laser and CCD to measure the width between adjacent tiles and enable a wall tiling robot to accurately control the placement of new tiles relative to previous tiles, but the gap was modeled as a simple scalar length of laser line discontinuity. Lastly, Wang et al. [31] showed that a camera and four 1D laser sensors could be used by a construction and mining robot to identify boreholes and adapt its tool pose to fill the holes with emulsion explosives. However, each borehole was identified by the dark circle at the mouth of the opening, and was modeled using a single point and orientation. Thus, such methods may not be very generalizable to a wide range of tasks or those requiring high-fidelity models.

Furthermore, some studies have employed model registration techniques that may not lend well to objects of complex curvature or high geometric complexity. For example, Giftthaler et al. [69] and Lussi et al. [32] demonstrated that a robot could be used to construct a wire mesh wall in an adaptive manner by sensing and modeling the as-encountered positions of wire nodes in a given course, and then modifying the planned poses of subsequent courses in an attempt to return the wall to its as-designed geometry. Although the edge detection approach worked well for objects like wires that have two distinct edges in an image, it might not work well for a wide range of object geometries. Similarly, Kermorgant [33] developed a magnetic tracked mobile robot capable of climbing a ship's hull and using a 2D laser profiler to adaptively track and weld the joints between the hull's cladding, but the joint's profile edges were modeled using an iterative line fitting algorithm, which would likely not work well for profiles containing complex curvature or high geometric complexity. Thus, the methods used in such studies are unlikely to be suitable for a wide range of object geometries, such as those characterized by complex curvature or high geometric complexity.

Other studies have used different sensing strategies or modalities to help robots perform work adaptively. For example, Feng et al. [34] demonstrated that a robot could use a camera and fiducial markers to adaptively identify, pick up, and install dry blocks on a block wall. However, the approach required the blocks and environment to be fitted with fiducial markers, which may not be ideal for real-world construction applications. Similarly, Lublasser et al. [35] showed that a robot could deconstruct a multi-layered façade wall using joint torque feedback to adapt to the actual geometry and composition of the wall's layered components. The study demonstrated that force feedback can be used to perform work requiring delicate finesse or when visual data are unattainable. However, it was necessary to manually position the robot at the onset of the

experiments because the robot was unable to sense its workpieces before making physical contact. Thus, construction robots are expected to benefit from multiple sensing modalities, but visualization of natural features is anticipated to be the principal modality for most construction applications.

Some studies have demonstrated noteworthy adaptive construction capabilities, but have not disseminated the methods used. For example, Helm et al. [36] demonstrated that a robot could construct a dry block wall in an adaptive manner by sensing the wall's top course with a 2D laser rangefinder and modifying the installation poses of subsequent blocks accordingly. Similarly, Willmann et al. [37] demonstrated that timber structures could be robotically constructed by sensing previously installed components, adapting the planned pose of subsequent pieces, cutting appropriate angles on the pieces, and installing them on the structure accordingly. However, the technical details of such studies are lacking, thereby preventing the results from being replicated, or the methods adapted or modified.

Geometrically adaptive research has also been conducted in other fields, such as manufacturing [70], domestic service [71], and space exploration [72]. In some studies, especially in manufacturing where robots are often specialized to perform only one class of tasks (e.g., welding) on only one class of objects or features (e.g., joints), the methods employed may not be very generalizable to a wide variety of tasks and object geometries [73,74]. Unlike manufacturing, in which standardization, centralization, and repetition permit a high degree of specialization among robots, the relatively distributed and individualized nature of construction projects requires construction robots to have greater versatility in the tasks they can perform and workpiece geometries they can perform work upon. Lastly, some approaches from manufacturing register complete 3D CAD models to determine the relative pose of the robot's workpieces [75,76]. This

works well in manufacturing applications where tight tolerances ensure that a workpiece's 3D geometry effectively matches its as-designed CAD model, but such approaches are not expected to work well in construction applications where the geometry of an individual workpiece can deviate substantially from its as-designed BIM model [66]. As such, construction approaches require additional model decomposition steps to capture the geometric deviations present within a single workpiece.

The research described in this chapter attempts to address the broad set of limitations and gaps identified above by proposing a generalizable framework to enable construction robots to accurately sense their workpieces, model their workpieces, adapt their work plans, and perform detailed work using techniques that are highly accurate and applicable to a wide range of objects.

4.3   Technical Approach

4.3.1   General Setup

The adaptive capabilities presented here may be applied to a variety of construction activities, but in order to demonstrate and evaluate these capabilities, joint filling, and in particular caulking, is used as the case study due to its commonality across a range of construction activities (e.g., welding, soldering, drywall joint filling, etc.). Although the treatment of an actual construction joint might require additional steps such as cleaning, priming, backing, and tooling, this research focuses on the dispensing of material to fill the joint. These other activities would require the robot to have similar capabilities to those developed here but different custom hardware to perform the tasks. Treatment of the joint's termination conditions (e.g., beginning and end) is also not considered in this work.

The premise for this research is as follows. A construction robot is operating on a construction site. The robot has been instructed to perform a specific task on specific workpieces. The robot has access to the most up-to-date Building Information Model (BIM) of the project, which contains the as-designed pose and geometry of the workpieces, as well as basic information about the work to be performed, but is not necessarily updated with all changes incurred during the construction phase. The robot reaches the vicinity of its workpieces via manual placement, teleoperation, or autonomous navigation. The robot has an estimate of its pose relative to the jobsite coordinate frame through such means as human measurement, direct sensor readings (e.g., GPS), or one of various autonomous robot localization approaches. Naturally, the robot's pose estimate contains error, but the estimate is accurate enough that the robot can aim its sensor in the direction it expects to find its workpieces and still detect them somewhere in the sensor's field of view. This is reasonable considering that it is not uncommon for today's indoor mobile robots to localize with only a couple centimeters and a single degree of pose error [57]. Lastly, the robot is equipped with sensors it can use to perceive its workpieces directly and form as-encountered workpiece models that are accurate relative to the robot's local coordinate frame. In this research, I focus on 2D laser sensors due to the inherently spatial nature of their output and their relative accuracy and precision, although the presented methods are applicable to any sensor capable of producing spatial data (e.g., point cloud data). It is also worth noting that I use a relatively accurate and precise (repeatability of $2\ \mu m$) sensor in our experiments to provide a baseline measure of performance under relatively ideal conditions, but performance should be expected to degrade with lower quality sensors.

4.3.2 Adaptive Framework Overview

To enable the robot to adapt to the geometry it encounters and perform work, I have developed an adaptive framework composed of four primary stages: workpiece sensing, workpiece modeling,

work planning, and work execution. The primary stages of the adaptive framework are shown in Figure 4.1 and are detailed in subsequent sections.
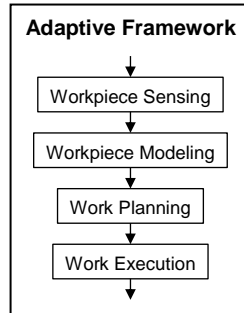
Figure 4.1 Adaptive framework overview

4.3.3 Workpiece Sensing

The first stage of the adaptive framework is sensing, as shown in Figure 4.2. Sensing involves executing a sensor data collection plan to collect sensor data, but it might also require planning if a sensor data collection path is not provided. The sensing stage of the adaptive framework is not treated in detail here, but readers can refer to Lundeen et al. [67] for more details. For simplicity, in this research it is assumed the robot is provided with an appropriate sensor data collection plan, which it executes to collect sensor data.

Figure 4.2 Workpiece sensing process

As an example, in the experiments described in this chapter, the data collection plan consists of a series of sensor poses at increments of 2 mm along the length of the as-designed joint. The sensor poses are defined such that t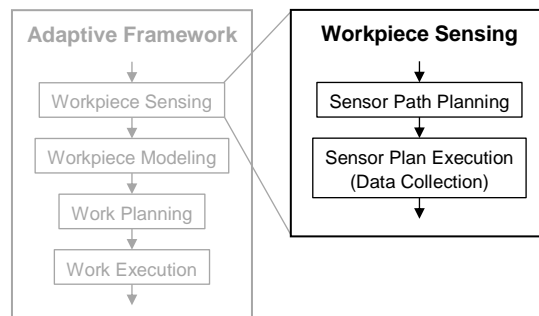he sensor is aimed normal to the broad surfaces of the workpieces, and the center of the sensing window corresponds with the center of the as-designed joint. Additionally, the data collection plan includes a trigger which instructs the sensor to collect one sweep (i.e., profile) of sensor data at each pose.

4.3.4 Workpiece Modeling

After collecting sensor data, the robot uses the data to model the actual pose and geometry of its workpieces, as outlined in Figure 4.3. To do so, the robot leverages its jobsite pose estimate and its sensor's field of view to identify and extract the object geometry it expects to encounter in its sensor data. The robot then uses ray casting to isolate the surfaces of the objects it expects to see in its data. A BIM model of the entire expected scene is then registered to the sensor data using one of various model registration techniques. In this research, the author employs the Generalized Resolution Correlative Scan Matching (GRCSM) search algorithm, which is a breadth-first search that uses multiple resolution levels to narrow its search. The GRCSM search algorithm functions by building rasterized lookup tables of the object models, projecting the sensor data onto the lookup tables at various test poses, and summing the corresponding table values to produce a correlation score between the sensor data and the model. The correlation scores are then used to determine which pose to explore further at the next resolution level, and this process is repeated until a relative pose is found at the level of highest resolution, which is taken to be the estimated relative pose between the as-designed object and the actual object. After registering the entire scene, the sensor data are projected onto each object model using the registration results, and the correlation scores returned by the GRCSM search algorithm are used to associate each data point

with the particular object that best describes it. The scene is then decomposed and individual objects are re-registered to their associated sensor data independently. Finally, using the individual registration results, the robot is able to form as-perceived models of its workpieces that are accurate relative to the robot's local coordinate frame. Readers are encouraged to refer to Lundeen et al. [67] for more details about workpiece modeling.
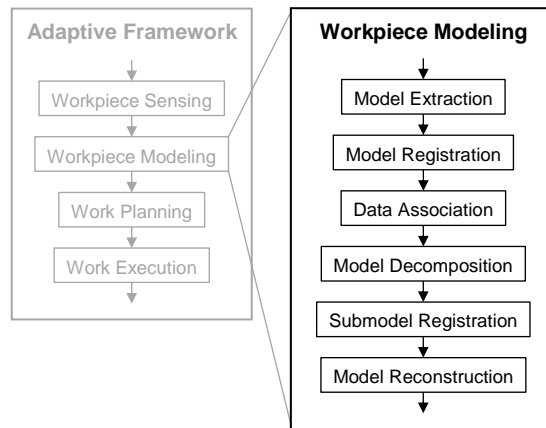
Figure 4.3 Workpiece modeling process

4.3.5 Work Planning

Once the robot has formed as-perceived models of its workpieces, it is then able to adapt its work plan to the models. The work plan can be stored in the BIM model, can exist as an independent algorithm, or can be formed from some combination thereof. The work plan can also be of varying degrees of complexity. The process used in this research to form a joint filling work plan is summarized in Figure 4.4 and detailed thereafter.
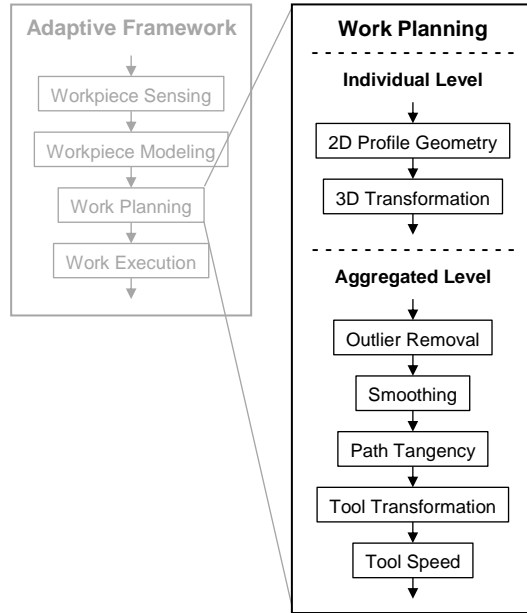
Figure 4.4 Work planning process

4.3.5.1   Individual Profile Level

In the case of filling a square butt joint between two workpieces, the work plan for an individual joint profile can be formulated as follows. The formulation corresponds to the *2D Profile Geometry* step in Figure 4.4 and leverages Figure 4.5 as a visual reference. It is described generically so as to remain applicable to analysis performed in 2D or 3D. Points and vectors are described using the column vector convention (e.g., $\boldsymbol{p} = \begin{bmatrix} x \\ y \end{bmatrix}$).

Figure 4.5 Individual profile model for a square butt joint fill plan

To start, points p_1, p_2, p_3, and p_4 are assigned to the workpiece corners with p_1 and p_3 on the side to receive fill material. The gap opening's tangent vector, $\vec{t}_{gap,ind}$, can then be defined from p_1 to p_3 as follows.

$$\vec{t}_{gap,ind} = p_3 - p_1 \tag{4.1}$$

Any arbitrary vector can be converted into a unit vector using the following formula.
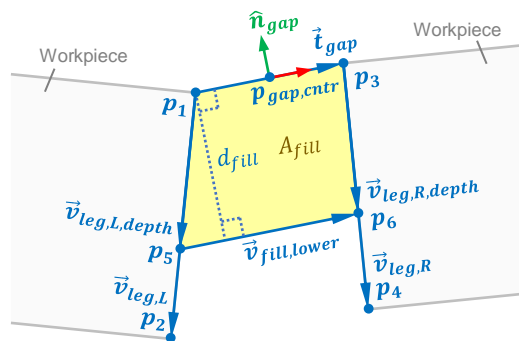
$$\hat{v} = \frac{\vec{v}}{\|\vec{v}\|} \tag{4.2}$$

Thus, the gap opening's tangent unit vector, $\hat{t}_{gap,ind}$, can be determined as follows.

$$\hat{t}_{gap,ind} = \frac{\vec{t}_{gap,ind}}{\|\vec{t}_{gap,ind}\|} \tag{4.3}$$

If the workpiece points (i.e., p_1, p_2, p_3, and p_4) are 2D points defined in the plane of the sensor data, then the gap opening's normal unit vector, $\hat{n}_{gap,ind}$, can be determined as follows.

$$\hat{n}_{gap,ind} = Rot_{90°}\hat{t}_{gap,ind} = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \hat{t}_{gap,ind} \tag{4.4}$$

Alternatively, if the points are defined relative to any other coordinate frame (e.g., 3D world coordinates), then the gap opening's normal unit vector can be determined by cross multiplying the sensing plane's normal unit vector (i.e., vector projecting out of the plane of Figure 4.5) with the gap opening's tangent unit vector, as follows.

$$\hat{n}_{gap,ind} = \hat{n}_{sense,plane} \times \hat{t}_{gap,ind} \tag{4.5}$$

The position of the center of the gap's opening can be determined as follows.

$$p_{gap,cntr,ind} = p_1 + \vec{t}_{gap,ind}/2 \qquad (4.6)$$

The width of the gap opening, $w_{gap,opening}$, can be determined from the norm of the gap opening's tangent vector, as follows.

$$w_{gap,opening} = \left\| \vec{t}_{gap,ind} \right\| \qquad (4.7)$$

In joint filling applications, such as caulking and welding, it is sometimes desirable to have a contoured upper fill profile, such as concave, convex, or recessed [77,78]. However, a flush profile is arguably the simplest and most common among joint filling applications, so it is the profile used here. As such, the width of the fill material across its upper edge, $w_{fill,upper}$, is simply equal to the width of the gap opening.

$$w_{fill,upper} = w_{gap,opening} \qquad (4.8)$$

The depth of fill material varies depending on the specific application. In welding, for example, it is often desirable to obtain complete workpiece depth penetration [78], while in caulking a common guideline is one-half the joint's width [77]. Thus, the fill depth, d_{fill}, can be specified using a depth-to-width ratio, $R_{depth,width}$, as follows.

$$d_{fill} = w_{fill,upper} R_{depth,width} \qquad (4.9)$$

Next, vectors can be defined along the workpiece's inner legs, as shown below for the left leg.

$$\vec{v}_{leg,L} = p_2 - p_1 \qquad (4.10)$$

Depth-adjusted vectors forming each side of the fill profile can then be determined, as shown below for the left leg.

$$\vec{v}_{leg,L,depth} = -\hat{v}_{leg,L}d_{fill}/(\hat{n}_{gap,ind}^{\top}\hat{v}_{leg,L}) \tag{4.11}$$

The fill material's lower edge vector can then be determined as follows.

$$\vec{v}_{fill,lower} = -\vec{v}_{leg,L,depth} + \vec{t}_{gap,ind} + \vec{v}_{leg,R,depth} \tag{4.12}$$

The width of the fill material across its bottom edge can be determined from the norm of that vector.

$$w_{fill,lower} = \left\| \vec{v}_{fill,lower} \right\| \tag{4.13}$$

Since the fill profile's upper and lower edges are parallel, the profile's cross-sectional area can be determined from the trapezoidal area formula as follows.

$$A_{fill} = d_{fill}(w_{fill,upper} + w_{fill,lower})/2 \tag{4.14}$$

At this point, in order for the robot to perform work on its 3D workpieces, the analysis must be transformed to 3D. This marks a transition from the *2D Profile Geometry* step to the *3D Transformation* step in the work planning process outlined in Figure 4.4. To begin, it should be noted that any 3D rigid-body rotation can be represented by a 3×3 orthonormal rotation matrix belonging to the SO(3) special orthogonal group [79], as shown below.

$$\boldsymbol{R} = \begin{bmatrix} r_{11} & r_{12} & r_{13} \\ r_{21} & r_{22} & r_{23} \\ r_{31} & r_{32} & r_{33} \end{bmatrix} \tag{4.15}$$

Next, there are three vectors that define the orientation of the gap's opening at a particular cross section: the gap opening's tangent unit vector, the gap opening's normal unit vector, and the sensing plane's normal unit vector (i.e., vector projecting out of the plane of Figure 4.5). Because

these three vectors are orthonormal, they can be used to construct the gap opening's orthonormal rotation matrix. If the geometric analysis has thus far been performed in 2D, then the gap opening's rotation matrix can be expressed in 3D as follows.

$$R_{gap,ind} = \begin{bmatrix} \hat{t}_{gap,ind} & \hat{n}_{gap,ind} & 0 \\ & & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{4.16}$$

Alternatively, if the geometric analysis has thus far been performed in 3D, then the gap opening's rotation matrix can be expressed as follows.

$$R_{gap,ind} = \begin{bmatrix} \hat{t}_{gap,ind} & \hat{n}_{gap,ind} & \hat{n}_{sense,plane} \end{bmatrix} \tag{4.17}$$

The 3D rigid-body homogeneous transformation that defines the gap's opening at a particular cross section can then be expressed as follows.

$$T_{gap,ind} = \begin{bmatrix} R_{gap,ind} & p_{gap,cntr,ind} \\ 0 \quad 0 \quad 0 & 1 \end{bmatrix} \tag{4.18}$$

The 3D orientation of the gap's opening can be expressed as a combination of three successive rotations about a fixed coordinate frame in the order of Rot_x, Rot_y, and then Rot_z. These three principle rotations can be determined as shown below, where $Atan2(Y,X)$ is a two-argument arctangent function having a domain spanning all four quadrants [80] and r_{ij} is an element of a rotation matrix as shown in Equation 4.15.

$$Rot_x = Atan2(r_{32}, r_{33}) \tag{4.19}$$

$$Rot_y = Atan2(-r_{31}, \sqrt{r_{11}^2 + r_{21}^2}) \tag{4.20}$$

$$Rot_z = Atan2(r_{21}, r_{11}) \tag{4.21}$$

In the case that the solution approaches the singularity at $Rot_y \approx +90°$, the following solution can be used.

$$Rot_x = Atan2(r_{12}, r_{22}) \qquad (4.22)$$

$$Rot_y = +90° \qquad (4.23)$$

$$Rot_z = 0 \qquad (4.24)$$

In the case that the solution approaches the singularity at $Rot_y \approx -90°$, the following solution can be used.

$$Rot_x = Atan2(r_{12}, r_{22}) \qquad (4.25)$$

$$Rot_y = -90° \qquad (4.26)$$

$$Rot_z = 0 \qquad (4.27)$$

The gap opening's 3D (6 DOF) rigid-body pose at a given cross section can be described as follows. It should be noted that an equally acceptable convention is to list the rotations in the opposite order, but the author uses the following notation.

$$\boldsymbol{X} = \begin{bmatrix} x \\ y \\ z \\ Rot_x \\ Rot_y \\ Rot_z \end{bmatrix} \qquad (4.28)$$

4.3.5.2 *Aggregated Profile Level*

In addition to planning at the individual profile level, planning also occurs at the aggregated profile level, which corresponds to the *Aggregated Profile Level* portion of the work planning process outlined in Figure 4.4. Shown in Figure 4.6 is an example of an aggregated set of individual joint profiles.
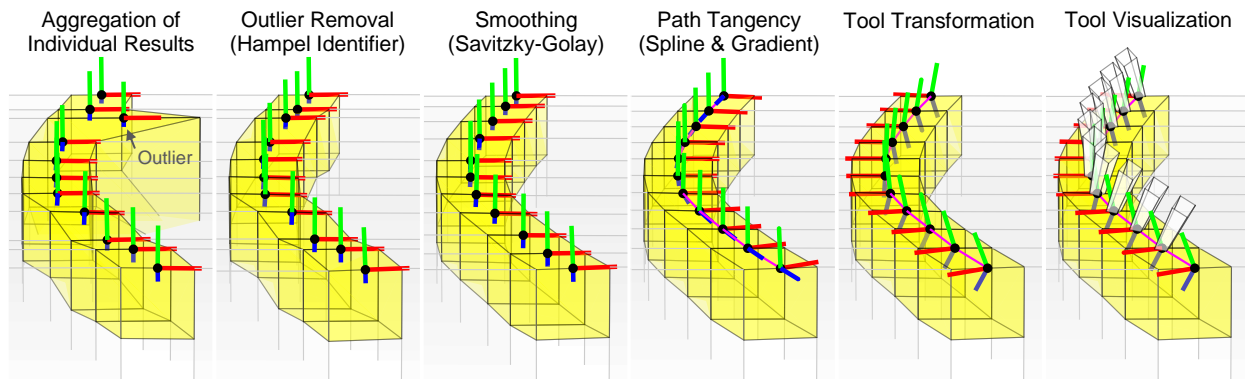

Figure 4.6 Fill planning at the aggregated model level

Outlier Removal. The first step in planning at the aggregated level is processing the raw results from the model registration process. Namely, each sensed profile along the length of the joint has an estimated pose for each workpiece at that particular cross section. The profile pose values for each workpiece are filtered using a Hampel identifier [81] to remove statistical outliers. Shown in Figure 4.6 is an example illustrating the removal of one outlier. The author uses a central moving window of 7 samples (i.e., 3 samples on either side of center) and a threshold of 3 standard deviations, although such parameters could be tuned for specific applications.

Smoothing. Next, each workpiece's profile pose values are smoothed using one of various data smoothing approaches. If the application requires real-time processing, then a simple moving average filter may be suitable due to its high processing efficiency. If the application permits post

processing, but still requires high processing speed, then a central moving average filter may provide higher fidelity due to its lack of phase shift along the length of the joint. In applications where post processing is permitted, but the joint is anticipated to be highly irregular or the sensor data sparse, a filter like the Savitzky-Golay filter [82] may be well suited. Like a moving average filter, the Savitzky-Golay filter uses a moving window but functions by fitting a polynomial to the data in each window. The Savitzky-Golay filter is more computationally intensive than the moving average, but when the ratio of sensor spacing to feature size becomes large, the Savitzky-Golay filter tends to retain more of the joint's features than the moving average filter, as illustrated in the example in Figure 4.7. The author uses the Savitzky-Golay filter in this research with a second degree (i.e., quadratic) polynomial. As a result of this smoothing step, the gap poses and fill profiles are smoothed, as shown in Figure 4.6.

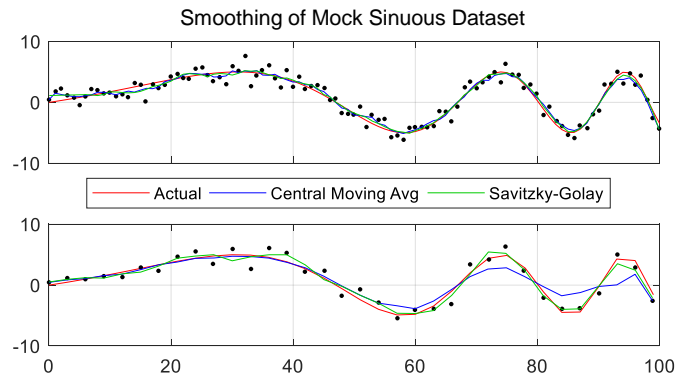Figure 4.7 Differences in feature retention between filtering methods

Path Tangency. After filtering and smoothing the profiles at the aggregated model level, the following equation is used to find the gap opening's tangent vector for each profile, where p_1 and p_3 refer to the corners of the butted workpiece edges on the side of the workpieces from which the joint is to be filled (see Figure 4.5).

$$\vec{t}_{gap,agg} = \boldsymbol{p_3} - \boldsymbol{p_1} \tag{4.29}$$

Additional joint profile characteristics, such as the cross-sectional fill area, are also recalculated for each profile using the filtered profiles and the relevant equations from the individual profile analysis, which corresponds to the *2D Profile Geometry* step in Figure 4.4.

Next, spline interpolation is used to fit a spline to the gap's center along the length of the joint. In this research, the author uses a spline with cubic interpolation, uniform spacing, and not-a-knot end conditions [83]. The spline used exhibits C^2 smoothness (i.e., continuous second derivative) to ensure the tool's orientation, which is a function of the spline's first derivative, exhibits C^1 smoothness. This yields finite angular accelerations for the tool. If a finite angular jerk is desired, it may be achieved via a spline with C^3 smoothness.

With an interpolated spline fitted to the gap's center positions, it then becomes possible to reorient the gap's coordinate frames to suit the 3D aggregated model. Whereas each coordinate frame previously had only one degree of rotational freedom (i.e., about the sensing plane's normal vector), the coordinate frames can now be reoriented such that the tool both contacts the workpieces and remains tangent to the path. To do so, it is first necessary to find the spline's numerical gradient vectors for each point along the spline. For the spline's interior points, a gradient vector, $\vec{g}$, can be approximated as follows.

$$\vec{g}_k = (\boldsymbol{p_{k+1}} - \boldsymbol{p_{k-1}})/2 \tag{4.30}$$

The first and last vectors in the array can be approximated as follows, where N is the total number of points in the array.

$$\vec{g}_1 = \boldsymbol{p_2} - \boldsymbol{p_1} \tag{4.31}$$

$$\vec{g}_N = p_N - p_{N-1} \tag{4.32}$$

Each local gradient vector is then converted into a unit vector using Equation 4.2. The gap opening's normal unit vector in the aggregated model can be determined by cross multiplying the local gradient unit vector by the gap opening's tangent unit vector from the aggregated model, as follows.

$$\hat{n}_{gap,agg} = \hat{g} \times \hat{t}_{gap,agg} \tag{4.33}$$

A third orthonormal vector can be constructed by cross multiplying the gap opening's normal unit vector by the local gradient unit vector, as follows.

$$\hat{v}_{gap,third,agg} = \hat{n}_{gap,agg} \times \hat{g} \tag{4.34}$$

With three orthonormal vectors, it is possible to construct the gap opening's 3D rotation matrix, as follows.

$$R_{gap,agg} = [\hat{v}_{gap,third,agg} \quad \hat{n}_{gap,agg} \quad \hat{g}] \tag{4.35}$$

Using the gap opening's rotation matrix and center spline points, it is possible to construct the gap opening's 3D homogeneous transformation matrix, as follows.

$$T_{gap,agg} = \begin{bmatrix} R_{gap,agg} & p_{gap,cntr,agg} \\ 0 \quad 0 \quad 0 & 1 \end{bmatrix} \tag{4.36}$$

The gap opening's 6 DOF pose vector can also be determined by applying Equations 4.19 through 4.28, which correspond to the *3D Transformation* step in Figure 4.4. An example of a spline and its tangentially oriented coordinate frames can be seen in Figure 4.6.

Tool Transformation. Depending on the definition of the tool's coordinate frame, it may be necessary to apply a homogeneous transformation to relate the gap opening frame to the tool frame. If so, one way of constructing a rotation matrix to describe the orientation of the tool's frame with respect to the gap's frame is shown below. In this approach, three successive rotations are applied about the gap's fixed frame in the order of Rot_x, Rot_y, and Rot_z, where θ represents an angle of rotation about the corresponding axis of rotation, and $c\theta$ and $s\theta$ are shorthand notations commonly used in robotics for $\cos(\theta)$ and $\sin(\theta)$, respectively. Readers can refer to [79,80] for more details about rotation matrices and approaches for constructing them.

$$R_{tool}^{gap} = R_z R_y R_x = \begin{bmatrix} c\theta & -s\theta & 0 \\ s\theta & c\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} c\theta & 0 & s\theta \\ 0 & 1 & 0 \\ -s\theta & 0 & c\theta \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & c\theta & -s\theta \\ 0 & s\theta & c\theta \end{bmatrix} \tag{4.37}$$

The homogenous transformation that describes the tool's frame with respect to the gap's frame can then be constructed as follows, where $\boldsymbol{p}_{tool}^{gap}$ is the position (i.e., translation) of the tool frame with respect to the gap frame. Readers can refer to [79,80] for more details about homogeneous transformations and their construction.

$$T_{tool}^{gap} = \begin{bmatrix} R_{tool}^{gap} & \boldsymbol{p}_{tool}^{gap} \\ 0 \quad 0 \quad 0 & 1 \end{bmatrix} \tag{4.38}$$

Each tool frame can then be determined as follows by applying the transformation relating the gap frame to the tool frame. Here, the relative transformation is successively applied to the gap transformation to yield the tool transformation.

$$T_{tool} = T_{gap,agg} T_{tool}^{gap} \tag{4.39}$$

118

Examples of tool frames along the length of the joint can be seen in Figure 4.6. Miniature caulk nozzles are also shown in Figure 4.6 to help readers visualize the tool frame definitions used in these experiments, as well as visualize the relationships between gap frames, tool frames, and the physical tool itself.

Tool Speed. At this point, the tool's pose has been determined along the joint, thus defining the robot's path, but the element of time is needed to convert the path into a trajectory. Similarly, the fill's cross-sectional area has been determined at each waypoint along the path, but the robot's plan still lacks a means of producing the desired fill output. In the case of viscous material deposition, the volume of material deposited per unit of time, dV/dt, can be expressed as a function of the volume of material deposited per unit of distance along the path, dV/ds, and the unit of distance traversed along the path per unit of time, ds/dt.

$$\frac{dV}{dt} = \frac{dV}{ds}\frac{ds}{dt} \quad \Longrightarrow \quad \frac{ds}{dt} = \frac{dV}{dt}\frac{ds}{dV} \tag{4.40}$$

The differential term ds/dt can be re-expressed as the tool's speed along the path, $\|\vec{v}_{tool}\|$; the differential term dV/dt can be re-expressed as the volumetric flowrate of material dispensed from the tool, $\dot{V}_{fill}$; and the differential term dV/ds can be re-expressed as the fill material's cross-sectional area, $A_{fill}(s)$, which varies as a function of position along the path. In order to output fill material of varying cross-sectional area, either the caulk flowrate or the tool's speed must be varied. For simplicity, the author chooses to fix the volumetric flowrate and vary the tool's speed along the path. As such, the following equation can be used to determine the necessary tool speed along the path to achieve the desired cross-sectional area of deposited fill material.

$$\|\vec{v}_{tool}(s)\| = \dot{V}_{fill}/A_{fill}(s) \tag{4.41}$$

The robot then has a series of positions, orientations, and tangential speeds, which together serve as its adapted joint fill plan.

4.3.6 Hardware and Calibration

In addition to being able to adapt its work plan, the robot also needs the ability to physically manipulate its environment in order to perform real construction work. Because caulking was selected as the case study for the experiments in this research, a caulk dispensing tool was developed for the robot, as can be seen in Figure 4.8.

Figure 4.8 Caulk dispensing tool

The caulk dispensing tool is designed to accommodate a $300\ mL$ ($10.1\ FL. OZ.$) tube of caulk. As shown in Figure 4.9, the caulk dispensing tool is a pneumatic-driven pressure vessel constructed from PVC components available from a common hardware store. The PVC components are assembled using PVC cleaner, primer, and adhesive. Airtight seals are created via Teflon tape on the threaded joints and rubber grommets between the vessel and the end of the caulk tube. The internal gauge pressure is regulated using a 0-$25\ psi$ mechanical on/off pressure regulator. An adjustable backstop is included in the design to compensate for variation in the lengths of caulk

tubes and fabrication variation in the construction of the tool. The caulk tube's external nozzle is cut at a 45° angle with an outer width measuring 8 mm, as shown in Figure 4.9, to suit the size of the gaps used in the experiments. The cut yields an elliptical orifice with a minor diameter of 5.7 mm, major diameter of 8.0 mm, and cross-sectional area of 36 mm^2.
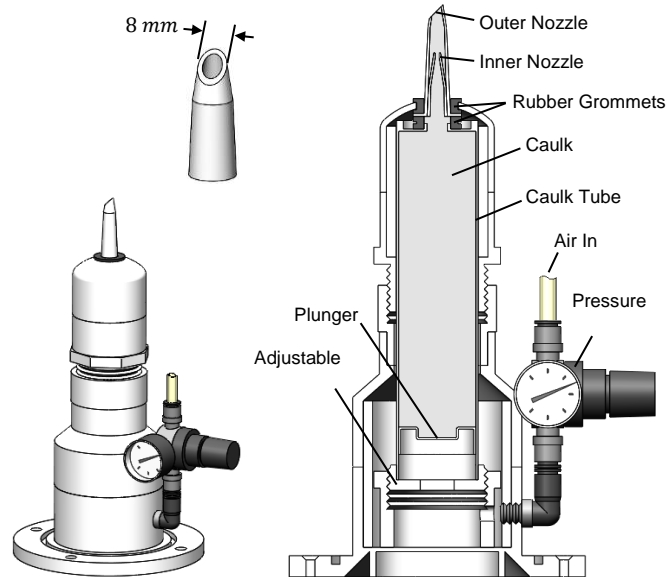

Figure 4.9 Caulk dispensing tool design

4.3.6.1 *Caulk Density Calibration*

In joint filling applications, such as caulking, a common characteristic of interest is the volume of fill material deposited. There are various ways to measure the volume of an irregularly shaped substance (e.g., water displacement), but one way to do it quickly, easily, and with a digital readout is to dispense the material onto a digital scale and divide its weight by density to estimate volume. To enable the utilization of such an approach, it was decided to calibrate the density of the caulk. The specific caulk selected for use in the experiments was DAP® Alex® Painter's Acrylic Latex Caulk. To calibrate the caulk's density, the tube's nozzle was cut, the tube's plunger was driven to expel air from the tube, the caulk was struck clean across the nozzle, the inner diameter of the tube

was measured, the distance from the back of the plunger to the end of the tube was measured, and the tube was weighed. Then, the majority of caulk was dispensed from the tube. The caulk was again struck off across the nozzle, the tube was weighed, and the distance from the back of the plunger to the end of the tube was measured. The tube's inner diameter and plunger displacement were used to estimate the volume of caulk dispensed, and the weight difference was used to estimate the weight of the caulk dispensed. This process was repeated for a total of eight tubes of caulk. The caulk's mean specific volume was found to be $631\ mm^3/g$ with a standard deviation of $8\ mm^3/g$.

4.3.6.2 Caulk Flowrate Calibration

In order for the robot to deposit the desired amount of caulk into the joint, the relationship between input air pressure and output caulk volumetric flowrate needed to be determined. To determine this relationship, a series of tests were conducted wherein the pressure regulator was set to one of various pressures, and the weight of caulk dispensed was measured over a period of time. The process was recorded on video, which was later reviewed to determine the starting and ending times, weights, and pressures. The starting time was determined from the gauge's sudden jump from $0\ psi$ to the operating pressure. A spring-loaded trigger mechanism was used to momentarily interrupt the flow of caulk from the nozzle to create a well-defined ending time, weight, and pressure. The starting and ending pressures were averaged to yield a mean pressure. The weight of caulk dispensed over the course of the process was divided by the duration of the process to yield a dispensed weight per unit of time that was converted to a dispensed volume per unit of time via the calibrated caulk density. A plot of output caulk volumetric flowrate versus mean input air pressure can be seen in Figure 4.10.
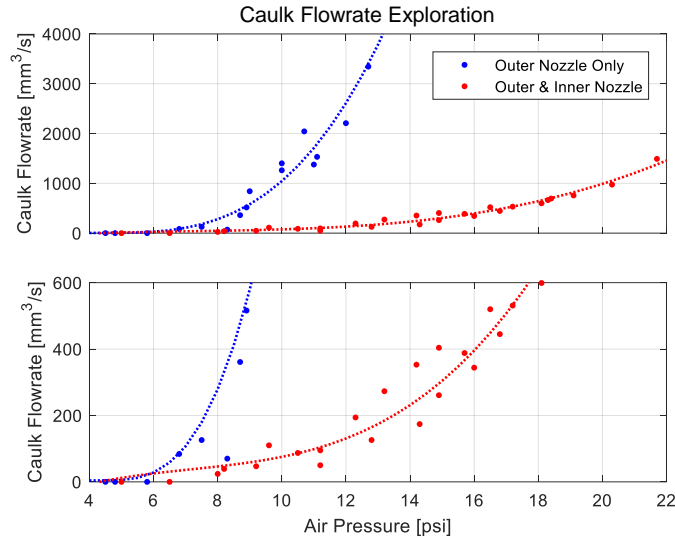
Figure 4.10 Exploration of output caulk flowrate as a function of input air pressure

It had been observed in preliminary experiments that a caulk flowrate in the vicinity of $200\ mm^3/s$ was needed to produce a reasonable robot speed for the size of gaps used in the experiments. As can be seen in Figure 4.10, an input air pressure of roughly $7.5\ psi$ could hypothetically produce the desired caulk flowrate. However, due to the nature of the on/off pressure regulator, the mean pressure was commonly observed to vary approximately $\pm 0.1\ psi$ from the mean value within each individual experiment, and approximately $\pm 0.1\ psi$ from one experiment to the next. Given the sensitivity of the caulk flowrate (as indicated by the slope of the trend line), the variation in the regulator's output, and the variation in the flowrate (as indicated by the vertical deviation of the data from the trend line), it was anticipated that the output of such a system would exhibit prohibitively large variation. Thus, an alternative option was explored in which a smaller nozzle was placed inside the larger nozzle, as shown in Figure 4.11. The smaller nozzle had an orifice diameter of $2.4\ mm$ and a corresponding cross-sectional area of $4.5\ mm^2$. Thus, the cross-sectional area of the new orifice was reduced to only 13% of the original orifice's area.

Figure 4.11 Internal nozzle used to achieve a shallower caulk flowrate profile

Through the use of the internal nozzle, a shallower caulk flowrate profile was achieved, as shown in Figure 4.10. It can be seen that the new profile has a shallower slope at the $200\ mm^3/s$ crossing than the previous profile, thereby indicating a reduced sensitivity to errors in input air pressure. For the new profile, a caulk flowrate in the vicinity of $200\ mm^3/s$ could be achieved using an input pressure of approximately $13\ psi$, but the author chose to use an operating pressure of $12.5\ psi$ because they felt the pressure gauge could be more accurately read with the needle centered between two tick marks instead of occluding a mark.

Given that $12.5\ psi$ had been selected as the operating pressure for the system, it then became necessary to calibrate the caulk flowrate more thoroughly around that particular pressure. A series of tests were conducted using the same procedures as in the previous tests but at pressures more closely grouped around $12.5\ psi$. The results are shown in Figure 4.12. Using a least squares fitted first order polynomial, the calibrated caulk flowrate at $12.5\ psi$ was found to be $239\ mm^3/s$. In the case that the mean pressure during a single experiment lies at an exact value of $12.5\ psi$, one can expect a corresponding flowrate standard deviation of $35\ mm^3/s$. In the case that the mean pressure during a single experiment lies within the range of $12.5 \pm 0.1\ psi$, one can expect a corresponding flowrate standard deviation of $40\ mm^3/s$. It is worth noting that those are fairly large standard deviations, approaching 50% error at 3 standard deviations. It is also worth noting that the system exhibits a local sensitivity of $278\ mm^3/s \cdot psi$, or a mean output change of nearly

$28 \ mm^3/s$ (12% error) for an input fluctuation of only $0.1 \ psi$. Generally speaking, the pneumatic caulk dispensing tool exhibits considerable variance.
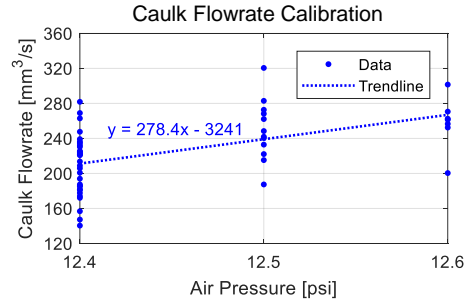

Figure 4.12 Caulk flowrate calibration at 12.5 psi

Lastly, the robots used in the experiments were designed to receive speed commands in the form of a percentage of maximum speed capacity. Thus, it was necessary to calibrate the relationship between speed percentage (%) inputs and physical speed (mm/s) outputs. A series of trials were conducted in which the robot's physical speed output was measured for various percentage inputs. Specifically, the robot was instructed to move a predefined distance, and the duration was measured. Three repetitions were conducted for each input speed percentage, as shown in Figure 4.13. Thus the calibrated robot's speed was determined to be $205 \ mm/s$ at 100% speed, or any corresponding percentage thereof.
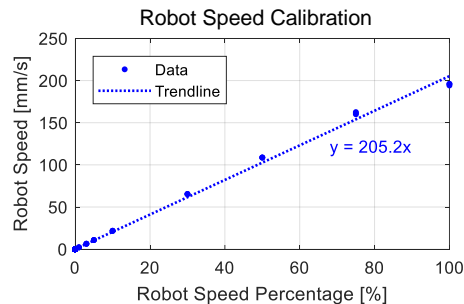

Figure 4.13 Robot speed percentage calibration

4.4   Experiments

Experiments were conducted to assess a construction robot's ability to adaptively fill an unexpected joint using the methods developed in this research. The experiments are first described in detail here, and the results are presented in the subsequent section.

In the experiments, the robot was provided a BIM model containing the as-designed poses and geometries of two abutted workpieces, and the robot was instructed to sense, plan, and adaptively caulk the joint between them. Experiments were conducted using three types of joints: a random joint, a joint formed between objects of complex geometry, and a set of accurately fabricated joints specifically designed for assessing the robot's performance capabilities. In all experiments, the robot expected to find a straight, uniform joint with a $4\ mm$ gap width, but the author placed a joint of unexpected geometry in front of the robot to see if it could adapt to the actual geometry it encountered. The robot aimed its sensor in the direction it expected to find the joint and collected sensor data. The robot then registered workpiece models to the sensor data, adapted its work plan, and executed its adapted plan to physically fill the joint.

Two Kuka KR120 robot arms were used in the experiments. The caulk dispensing tool was mounted on one arm and was loaded with DAP® ALEX® Painter's Acrylic Latex caulk. An LMI Technologies Gocator® 2350 2D laser profiler was mounted on the other arm and was used for sensing the workpieces. The sensor had a trapezoidal sensing window with a near field-of-view of $158\ mm$, a far field-of-view of $365\ mm$, and a range of $400\ mm$. Each GRCSM model registration search was performed over search windows of $204.8\ mm$ ($\pm\ 102.4\ mm$), $204.8\ mm$ ($\pm\ 102.4\ mm$), and $90°$ ($\pm\ 45°$). The robot's poses and sensor data were collected and synchronized using the robot's Programmable Logic Controller (PLC), and then asynchronously

126

read from the PLC using a Rhino® plugin called Super Matter Tools [84]. All computational processing was performed in MATLAB using an Intel® Core™ i7-4900MQ 2.80 GHz central processing unit running on a single thread. An overview of the experimental apparatus can be seen in Figure 4.14.
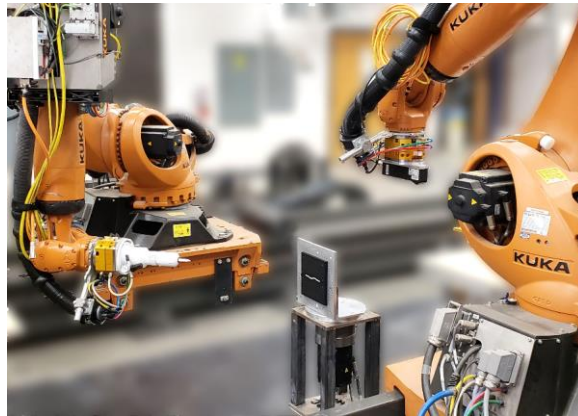

Figure 4.14 Experimental apparatus overview

4.4.1 Random Joint

To investigate whether a robot could model and adaptively fill a joint of unexpected geometry, a random joint was created by breaking a foam board in half. The two halves were laid side by side and spaced approximately 4 mm apart to form a gap, as shown in Figure 4.15.
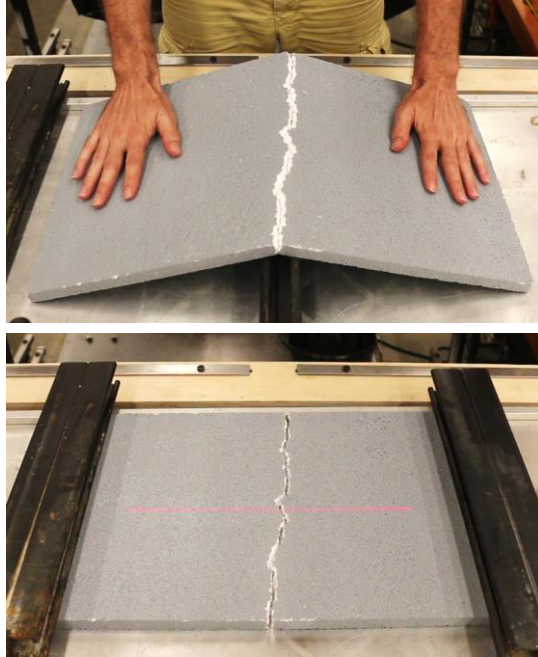
Figure 4.15 Random joint experimental setup

The robot then aimed its sensor toward the joint's as-designed location and collected sensor data. Sensor data were collected at intervals of approximately $5\ mm$, yielding a total of 66 sensor profiles along the length of the $330\ mm$ joint. After collecting sensor data, the robot built an as-encountered model of the joint, adapted its fill plan, and executed its adapted plan to fill the joint.

4.4.2 Complex Workpiece Geometry

To investigate the applicability of the adaptive methods to objects of greater geometric complexity, a mock door frame was constructed with a $4\ mm$ gap between the door jamb and the door casing, as shown in Figure 4.16.
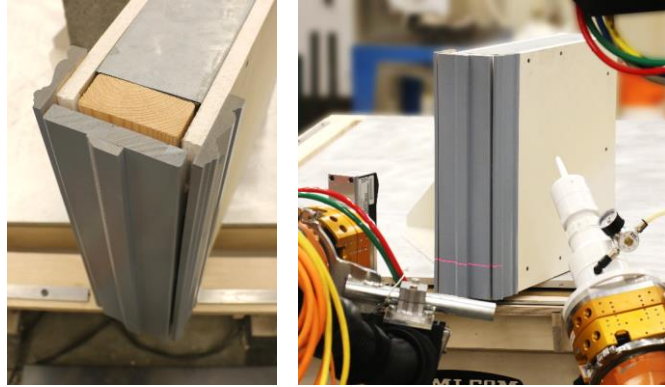
Figure 4.16 Complex workpiece geometry experimental setup

Here, *geometric complexity* refers to complexity from the perspective of robot modeling. To help illustrate the geometric complexity of the workpieces, Figure 4.17 contrasts the cross-sectional profiles of the door jamb and casing with the square-edged boards used in the other experiments of this chapter.
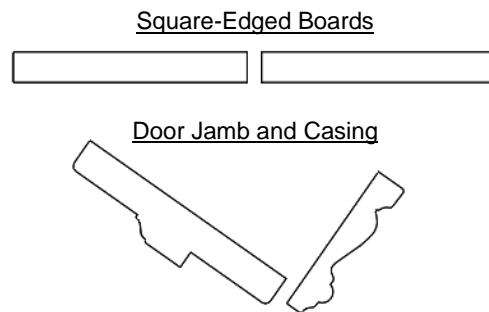

Figure 4.17 Contrast of workpiece profile complexity

As can be seen, the square-edged boards can be described by four line segments at 90° angles, which could be modeled using a variety of approaches such as line fitting and syntactic analysis [27,28]. In contrast, the door casing profile is composed of a larger number of line segments, relative angles, and curves, which are not conducive to line-fitting methods, and for which syntactic rules become difficult to formulate. Furthermore, line fitting and syntactic methods are not able to handle partial sensor data well and require a specific set of rules to be manually written

for each type of construction component. Thus, an approach that uses geometry directly from the BIM model may be more suitable for construction robots, given the wide variety of component geometries found in construction. The key purpose of this experiment is to demonstrate and highlight the generalizability of the proposed methods to a variety of object geometries.

In the experiment, the robot aimed its sensor toward the door frame's as-designed location and collected sensor data. Sensor data was collected at intervals of approximately $4.5\ mm$, yielding a total of 103 sensor profiles along the length of the $450\ mm$ joint. After collecting sensor data, the robot built an as-encountered model of the workpieces, adapted its fill plan, and executed its adapted plan to fill the joint.

4.4.3   Performance Evaluation

Five joint specimens were designed and fabricated to assess the performance capabilities of the construction robot in performing adaptive work, as shown in Figure 4.18. Each specimen pair had an outer width of $200\ mm$ and a joint measuring $100\ mm$ in length. Each joint was designed such that its geometry could be described by a simple mathematical function to provide a convenient and accurate means of calculating ground truth values at any point along the joint. The specimens were milled using a CNC router to minimize discrepancies between their as-designed geometry and their physical embodiment. Black acetal thermoplastic was chosen for the specimen material due to its machinability and dimensional stability.
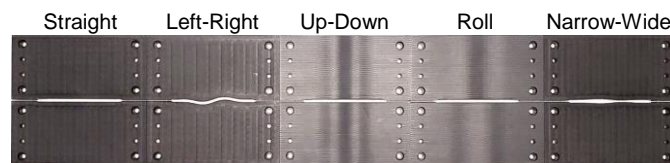

Figure 4.18 Performance assessment joint specimens

The joint specimens were designed to comprehensively test the operation's various degrees of freedom. The joints are referenced by their characteristic shapes, namely *straight*, *left-right*, *up-down*, *roll*, and *narrow-wide*. The straight joint matches the geometry of the as-designed joint, and thus serves as the experimental control. The straight joint and the as-designed joint both have a rectangular cross section with a width of $4\,mm$ and depth of $15\,mm$. A coordinate frame is assigned to the specimen as shown in Figure 4.19. As such, the following mathematical functions describe the characteristic geometry of the workpieces' upper butted corners in millimeters.

$$y_{strght,wrkpc,1} = 2 \tag{4.42}$$

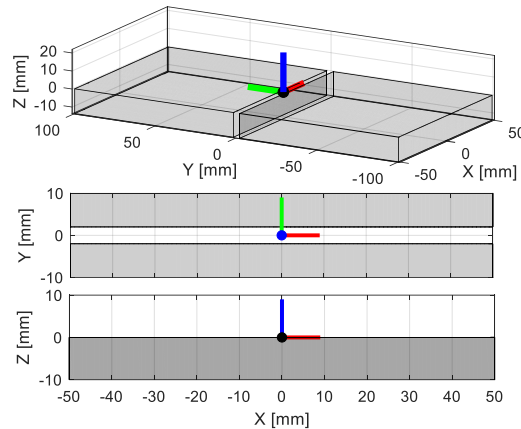$$y_{strght,wrkpc,2} = -y_{strght,wrkpc,1} \tag{4.43}$$

Figure 4.19 Straight joint specimen's geometry

The left-right joint was designed to test the robot's ability to compensate for discrepancies in the joint's y-position and yaw angle. The geometry of the left-right joint can be seen in Figure 4.20. The following mathematical functions describe the geometry of the workpieces' upper butted corners in millimeters. The constants were chosen so as to produce a sinusoid with a maximum absolute yaw angle of 20°, thereby deviating as much as 20° from the design. The underlying

formulation is complex and not central to the research, so only the numerical values are presented for the purposes of experimental reproducibility.

$$y_{lr,wrkpc,1} = 1.4481[2\,sin(\pi x/50) + sin(2\pi x/50)] + 2 \qquad (4.44)$$
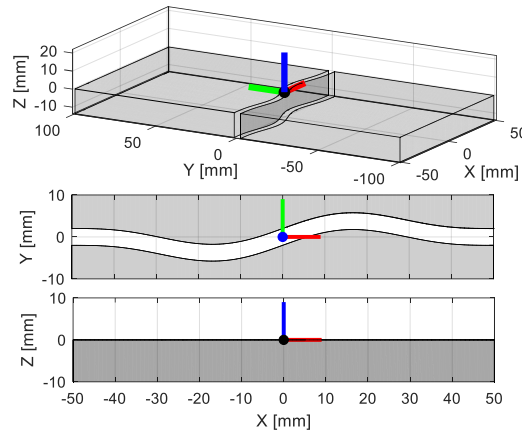
$$y_{lr,wrkpc,2} = y_{lr,wrkpc,1} - 4 \qquad (4.45)$$

Figure 4.20 Left-right joint specimen's geometry

The up-down joint was designed to test the robot's ability to compensate for discrepancies in the joint's z-position and pitch angle. The geometry of the up-down joint can be seen in Figure 4.21. The following mathematical functions describe the geometry of the workpieces' upper butted corners in millimeters. The constants were chosen so as to produce a sinusoid with a maximum absolute pitch angle of 10°, thereby deviating 10° from the design. The underlying formulation is complex and not central to the research, so only the numerical values are presented for the purposes of experimental reproducibility.

$$z_{ud,wrkpc,1} = 0.7014[2\,sin(\pi x/50) + sin(2\pi x/50)] \qquad (4.46)$$

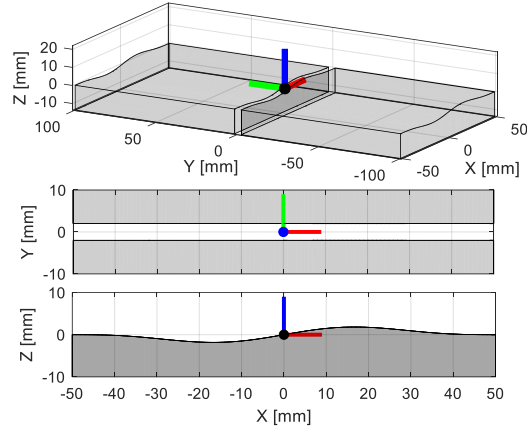$$z_{ud,wrkpc,2} = z_{ud,wrkpc,1} \qquad (4.47)$$

132

Figure 4.21 Up-down joint specimen's geometry

The roll angle joint was designed to test the robot's ability to compensate for discrepancies in the joint's roll angle. The geometry of the roll angle joint can be seen in Figure 4.22. The following mathematical functions describe the geometry of the workpieces' upper butted corners in millimeters. The constants were chosen so as to produce a sinusoid with a maximum absolute roll angle of 20°, thereby deviating 20° from the design. The underlying formulation is complex and not central to the research, so only the numerical values are presented for the purposes of experimental reproducibility.

$$z_{roll,wrkpc,1} = -0.2802[2\,sin(\pi x/50) + sin(2\pi x/50)] \tag{4.48}$$

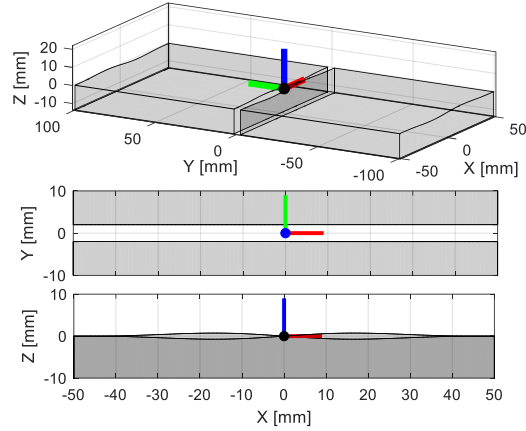$$z_{roll,wrkpc,2} = -z_{roll,wrkpc,1} \tag{4.49}$$

Figure 4.22 Roll angle joint specimen's geometry

The narrow-wide joint was designed to test the robot's ability to compensate for discrepancies in the joint's width (via tool speed). The geometry of the narrow-wide joint can be seen in Figure 4.23. The following mathematical functions describe the geometry of the workpieces' upper butted corners in millimeters. The constants were chosen so as to produce a joint with a minimum width of $2\ mm$ and a maximum width of $6\ mm$, thereby deviating $\pm 2\ mm$ from the designed width of $4\ mm$. The underlying formulation is complex and not central to the research, so only the numerical values are presented for the purposes of experimental reproducibility.

$$y_{nw,wrkpc,1} = -0.3849[2\ sin(\pi x/50) + sin(2\pi x/50)] + 2 \qquad (4.50)$$

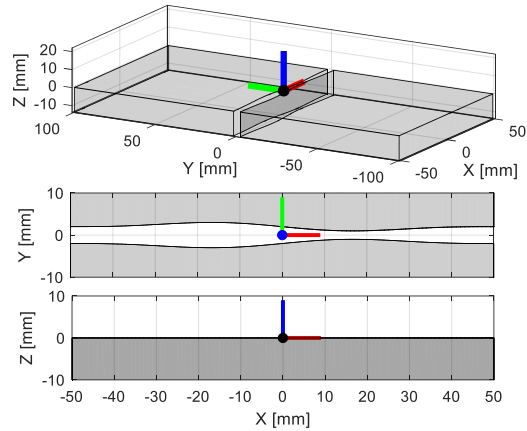$$y_{nw,wrkpc,2} = -y_{nw,wrkpc,1} \qquad (4.51)$$

Figure 4.23 Narrow-wide joint specimen's geometry

One at a time, the specimens were mounted on a frame in the approximate vicinity of the as-designed location of the joint, as shown in Figure 4.14. Before each experiment, the specimen was sensed on three orthogonal sides to establish the specimen's ground truth pose. The specimen's ground truth pose and as-designed geometry were then combined to determine the ground truth position of any point on the specimen's surface.

To begin the experiment, the robot aimed its sensor toward the joint's as-designed location and collected sensor data, as shown in Figure 4.24. Sensor data were collected at intervals of approximately 2 mm, yielding a total of 51 sensor profiles along the length of each joint.
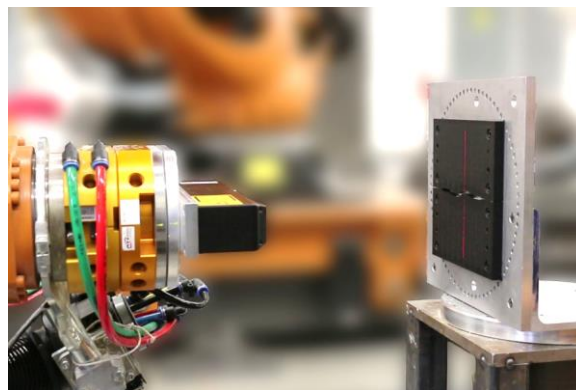

Figure 4.24 Sensor data collection example

After collecting sensor data, the robot built an as-encountered model of the joint and adapted its fill plan accordingly. The adapted fill plan was then executed using the caulk dispensing tool, an example of which is shown in Figure 4.25. After each experiment, the specimens were rescanned to measure the actual volume of caulk deposited into the joint.
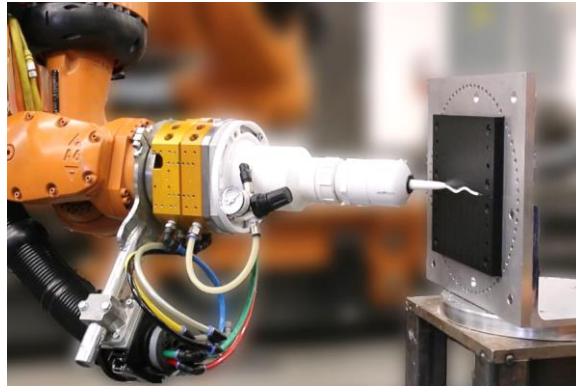
Figure 4.25 Example joint filling plan execution

To validate the system's ability to measure the volume of caulk deposited into the joint, a validation study was conducted in which a square rod measuring 4.05 mm on each side was placed into the joint, and was scanned and measured. The system was used to estimate the rod's cross-sectional area, and was found capable of doing so with a mean absolute error of 0.26 mm^2 (1.6% error) and standard deviation of 0.11 mm^2 (0.7% error), as shown in Figure 4.26.

Figure 4.26 Caulk volume measurement validation

4.5 Results

As previously noted, experiments were conducted for three types of joints: a random joint, a joint formed between objects of complex geometry, and a set of accurately fabricated joints specifically designed to assess the robot's performance capabilities. This section pr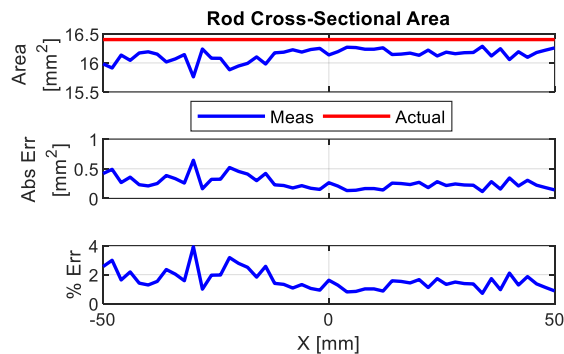ovides a concise presentation of the results, and the subsequent section discusses and compares the results to put them into perspective.

4.5.1 Random Joint

A random joint was formed by breaking a foam board in half and laying the two halves side by side with a small space between. Then the joint was scanned, the joint was modeled, and the work was planned. Shown in Figure 4.27 are the raw data collected during the sensor data collection step, the as-perceived joint model from the workpiece modeling step, and the adapted fill plan from the work planning step.

Figure 4.27 Random joint: perceived model and fill plan

The robot then executed its plan to fill the joint, as shown in Figure 4.28. A video of the experiment can be found here [85]. As can be seen, the robot demonstrated the general ability to identify the joint's true location and geometry, adapt its plan for filling it, execute its adapted plan, and fill the

joint with caulk. The mean processing time for the modeling and planning stages was found to be 1.1 $s/profile$, or a total of $70\ s$ for the entire joint (66 profiles).

Figure 4.28 Random joint: visual results

4.5.2   Complex Workpiece Geometry

A joint was formed between two objects of complex cross-sectional geometry, namely a door jamb and door casing. The joint was scanned, the joint was modeled, and the work was planned. Shown in Figure 4.29 are the raw data collected during the sensor data collection step, the as-perceived joint model from the workpiece modeling step, and the adapted fill plan from the work planning step.
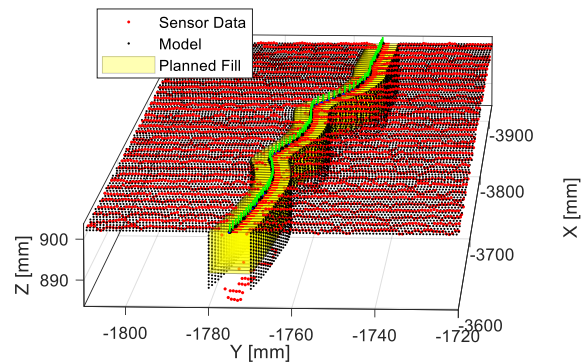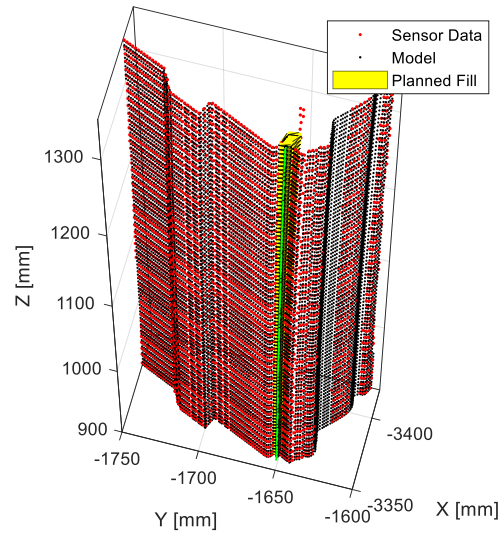
Figure 4.29 Complex geometry joint: perceived model and fill plan

The robot then executed its plan to fill the joint, as shown in Figure 4.30. A video of the experiment can be found here [86]. As can be seen, the robot demonstrated the general ability to identify the joint's true location and geometry, adapt its plan for filling it, execute its adapted plan, and fill the joint with caulk. The mean processing time for the modeling and planning stages was found to be 0.18 $s/profile$, or a total of 17 s for the entire joint (103 profiles).
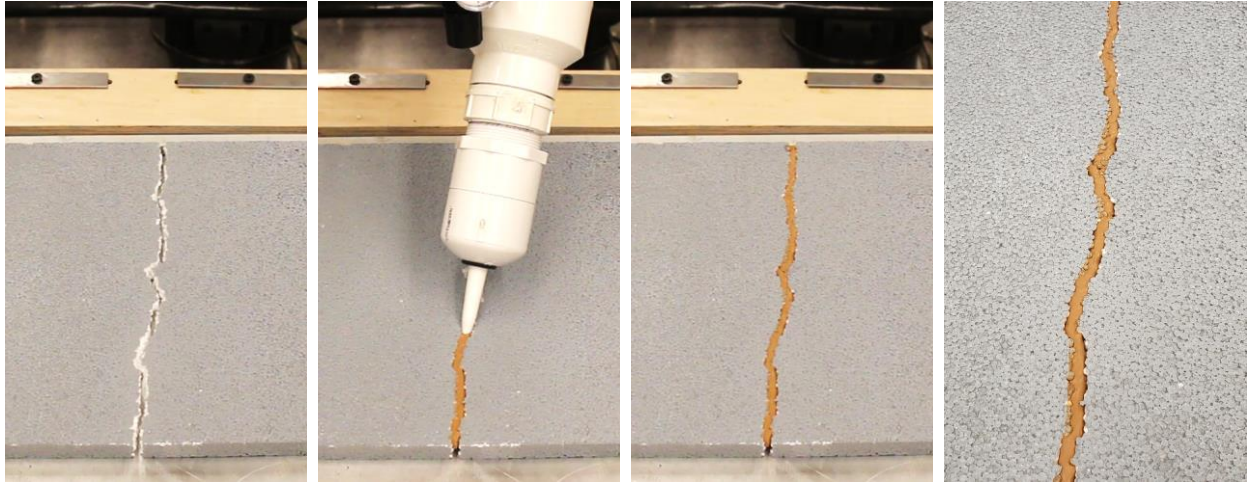

Figure 4.30 Complex geometry joint: visual results

4.5.3 Performance Evaluation

Experiments were conducted to assess the how accurately the robot could adapt to the pose and geometry of its workpieces. The experiments were conducted using five joint specimens (i.e., straight, left-right, up-down, roll, and narrow-wide) designed to comprehensively test each of the operation's degrees of freedom. The following subsections concisely present the same set of results for each of the five joint specimens, and the subsequent section discusses and compares the results to put them into perspective.

4.5.3.1 Straight Joint

The straight joint acts as a control specimen since the joint's actual geometry does not deviate from its as-designed geometry. Shown in Figure 4.31 are the raw data collected during the sensor data collection step, the as-perceived joint model from the workpiece modeling step, the adapted fill plan from the work planning step, and the actual volume of caulk deposited from the plan execution step. These results offer a visually intuitive sense of what the robot perceived, planned, and performed.
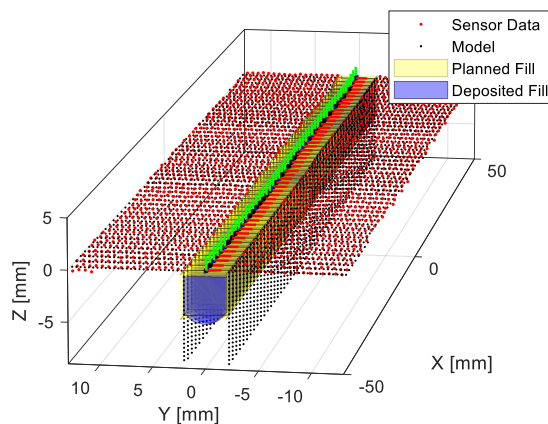
Figure 4.31 Straight joint: perceived model, fill plan, and deposited fill

Shown in Figure 4.32 is a comparison between the as-perceived joint model and the actual model along the length of the joint for each of the process's degrees of freedom, thereby providing a comparison between the true joint and the robot's perception of it. Also shown are the cross-sectional areas (i.e., differential volumes) of the actual caulk deposited into the joint.
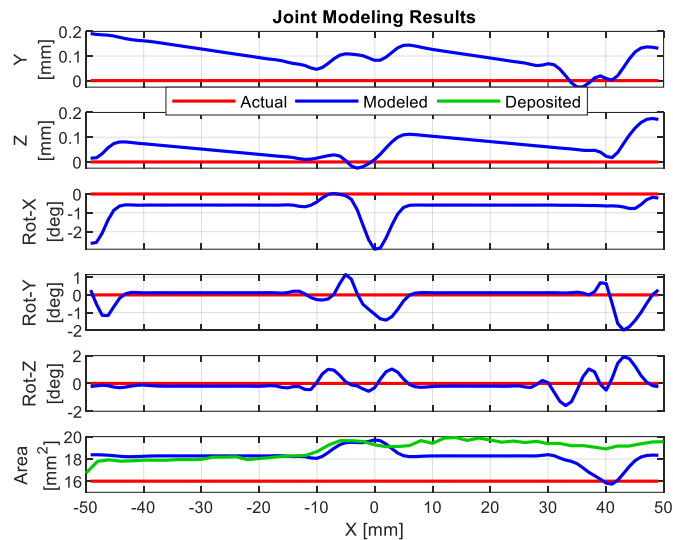
Figure 4.32 Straight joint: perceived model profiles

Shown in Figure 4.33 is the absolute modeling error along the length of the joint for each of the process's degrees of freedom, thereby providing a measure of how much error exists between the true joint and the robot's perception of it. Such information may be useful, for example, in applications with direction-specific tolerances. Also shown is the absolute error between the volume of the caulk deposited and the theoretically ideal amount for the specimen.
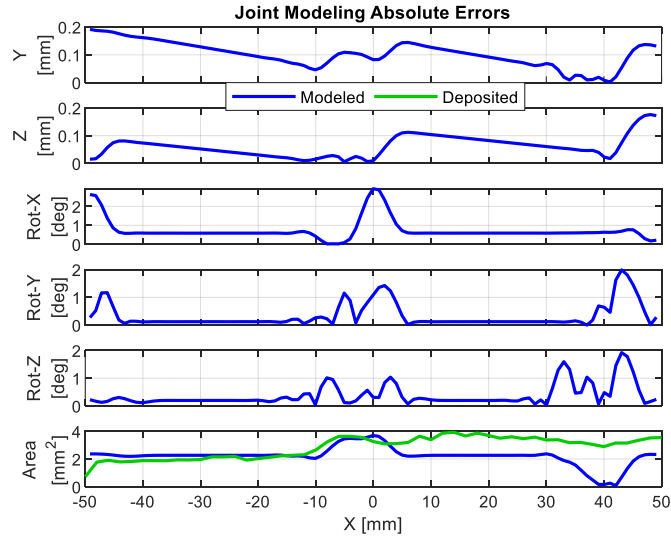
Figure 4.33 Straight joint: absolute modeling errors

Shown in Figure 4.34 are consolidated measures of modeling error along the joint, where the process's various degrees of freedom have been combined into single norm measures of position and orientation error via the 2-norm (e.g., Euclidean distance). Such information may be useful, for example, in applications with non-directional tolerances, where the magnitude of a deviation cannot exceed some limit regardless of direction.
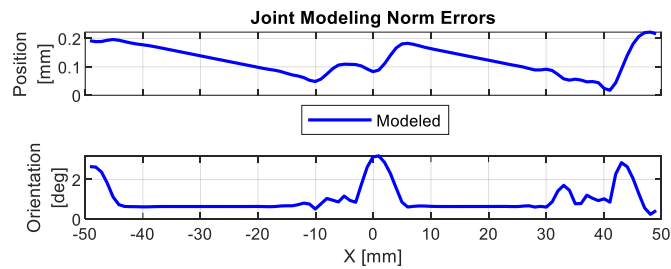
Figure 4.34 Straight joint: norm modeling errors

Shown in Table 4.1 are the summary statistics for the various modeling errors along the entire joint. As can be seen, the robot was found capable of identifying the true position and orientation of the joint's center with a mean norm positioning error of 0.12 mm and orientation error of 0.99°.

142

Table 4.1 Straight joint: modeling error statistics

	Abs Error Mean	Abs Error Std Dev	Norm Error Mean	Norm Error Std Dev
Position			0.12	0.05
Y (mm)	0.10	0.05		
Z (mm)	0.06	0.04		
Orientation			0.99	0.68
Rot-X (deg)	0.74	0.56		
Rot-Y (deg)	0.34	0.43		
Rot-Z (deg)	0.38	0.39		
Area (mm²)	2.22	0.66		

Shown in Figure 4.35 are images of the actual caulk deposited into the joint as a result of the robot's adaptive joint filling operation. Although an evaluation of the caulk quality is tangential to the scope of this research, the robot appeared to have successfully filled the joint with an appropriate amount of caulk and with minimal surface residue, and was visually acceptable compared to any hand-deposited caulk in similar joints.
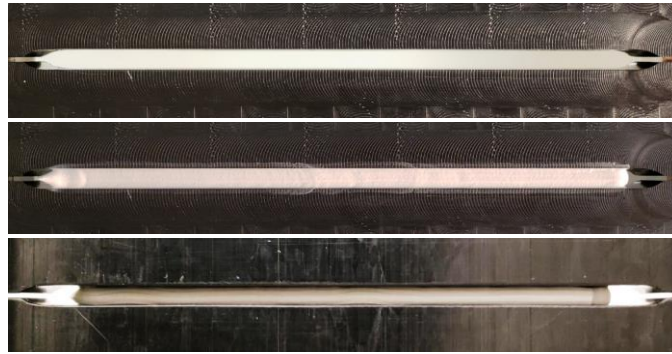
Figure 4.35 Straight joint: visual results

The mean processing time for the modeling and planning stages was found to be $0.42\ s/profile$, or a total of $22\ s$ for the entire joint (51 profiles).

4.5.3.2 *Left-Right Joint*

The left-right joint tests the robot's ability to compensate for discrepancies in the joint's y-position and yaw angle. Shown in Figure 4.36 are the raw data collected during the sensor data collection step, the as-perceived joint model from the workpiece modeling step, the adapted fill plan from the work planning step, and the actual volume of caulk deposited from the plan execution step. These results offer a visually intuitive sense of what the robot perceived, planned, and performed.
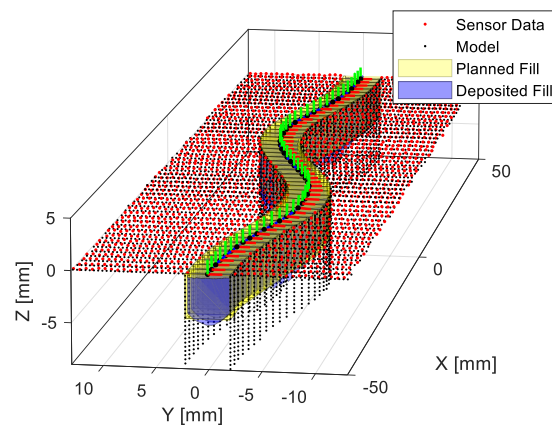
Figure 4.36 Left-right joint: perceived model, fill plan, and deposited fill

Shown in Figure 4.37 is a comparison between the as-perceived joint model and the actual model along the length of the joint for each of the process's degrees of freedom, thereby providing a comparison between the true joint and the robot's perception of it. Also shown are the cross-sectional areas (i.e., differential volumes) of the actual caulk deposited into the joint.
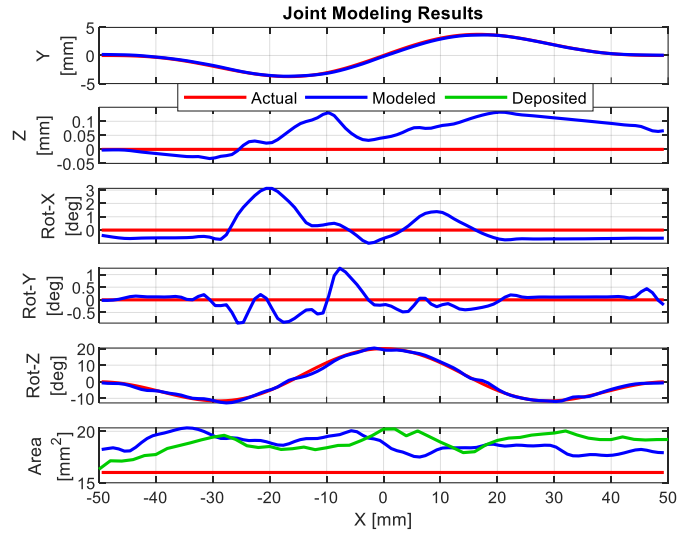
Figure 4.37 Left-right joint: perceived model profiles

Shown in Figure 4.38 is the absolute modeling error along the length of the joint for each of the process's degrees of freedom, thereby providing a measure of how much error exists between the true joint and the robot's perception of it. Also shown is the absolute error between the volume of the caulk deposited and the theoretically ideal amount for the specimen.
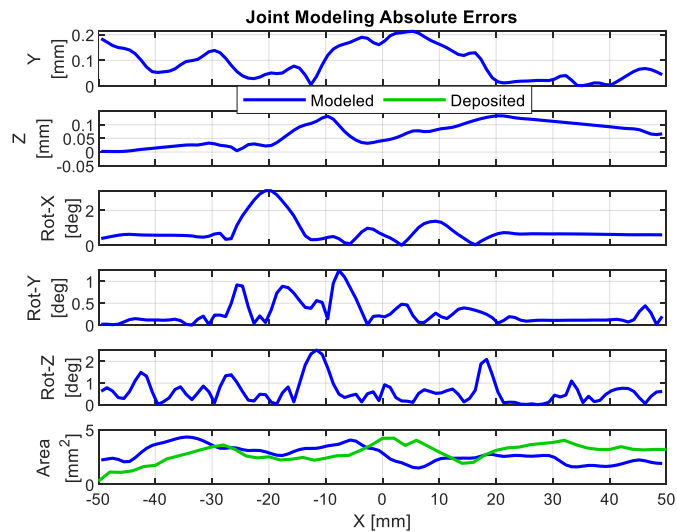
Figure 4.38 Left-right joint: absolute modeling errors

145

Shown in Figure 4.39 are consolidated measures of modeling error along the joint, where the process's various degrees of freedom have been combined into single norm measures of position and orientation error via the 2-norm (e.g., Euclidean distance).
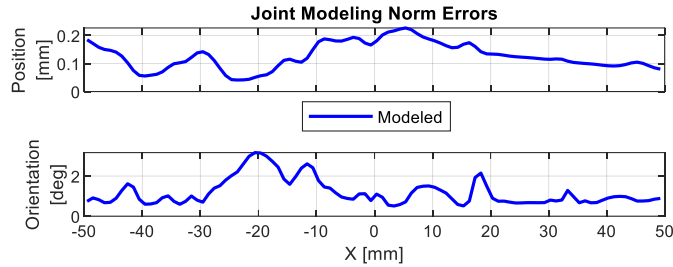
Figure 4.39 Left-right joint: norm modeling errors

Shown in Table 4.2 are the summary statistics for the various modeling errors along the entire joint. As can be seen, the robot was found capable of identifying the true position and orientation of the joint's center with a mean norm positioning error of $0.13\ mm$ and orientation error of $1.19°$, despite the discrepancies in y-position and yaw angle between the actual joint and as-designed joint.

Table 4.2 Left-right joint: modeling error statistics

	Abs Error Mean	Abs Error Std Dev	Norm Error Mean	Norm Error Std Dev
Position			0.13	0.05
Y (mm)	0.09	0.07		
Z (mm)	0.07	0.04		
Orientation			1.19	0.66
Rot-X (deg)	0.80	0.64		
Rot-Y (deg)	0.26	0.26		
Rot-Z (deg)	0.62	0.54		
Area (mm²)	2.73	0.75		

Shown in Figure 4.40 is a series of images showing the motion of the caulk dispenser as the robot executed its adapted fill plan. A video of the experiment can be found here [87].

146
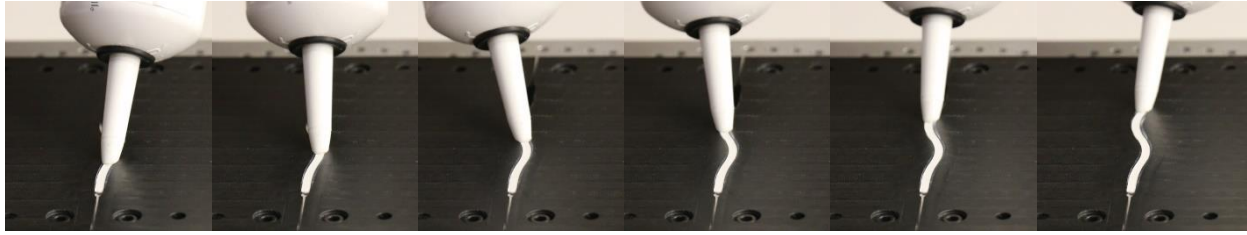
Figure 4.40 Left-right joint: fill action sequence

Shown in Figure 4.41 are images of the actual caulk deposited into the joint as a result of the robot's adaptive joint filling operation. The robot appears to have successfully filled the joint with an appropriate amount of caulk and with minimal surface residue.

Figure 4.41 Left-right joint: visual results

An inclined view of the deposited caulk can be seen in Figure 4.42.

Figure 4.42 Left-right joint: inclined view of visual results

The mean processing time for the modeling and planning stages was found to be $0.39\ s/profile$, or a total of $21\ s$ for the entire joint (51 profiles).

4.5.3.3 Up-Down Joint

The up-down joint tests the robot's ability to compensate for discrepancies in the joint's z-position and pitch angle. Shown in Figure 4.43 are the raw data collected during the sensor data collection step, the as-perceived joint model from the workpiece modeling step, the adapted fill plan from the work planning step, and the actual volume of caulk deposited from the plan execution step. These results offer a visually intuitive sense of what the robot perceived, planned, and performed.
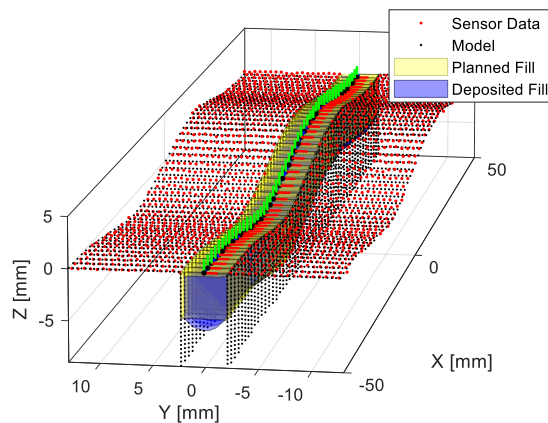
Figure 4.43 Up-down joint: perceived model, fill plan, and deposited fill

Shown in Figure 4.44 is a comparison between the as-perceived joint model and the actual model along the length of the joint for each of the process's degrees of freedom, thereby providing a comparison between the true joint and the robot's perception of it. Also shown are the cross-sectional areas (i.e., differential volumes) of the actual caulk deposited into the joint.

148

Figure 4.44 Up-down joint: perceived model profiles

Shown in Figure 4.45 is the absolute modeling error along the length of the joint for each of the

process's degrees of freedom, thereby providing a measure of how much error exists between the

true joint and the robot's perception of it. Also shown is the absolute error between the volume of

the caulk deposited and the theoretically ideal amount for the specimen.
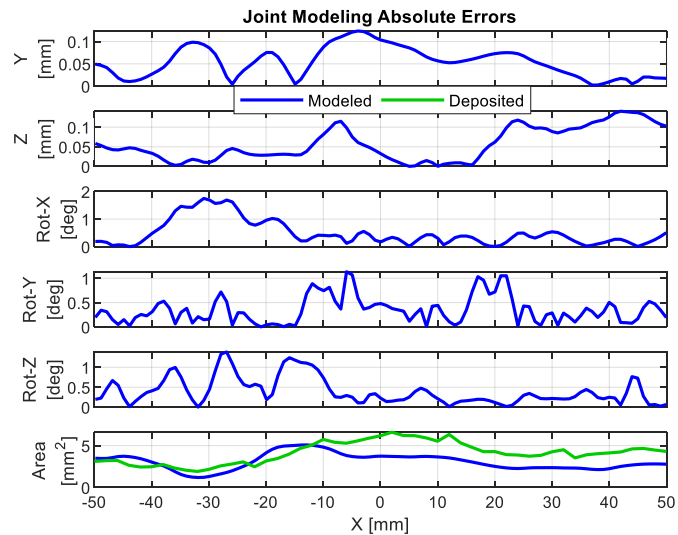
Figure 4.45 Up-down joint: absolute modeling errors

Shown in Figure 4.46 are consolidated measures of modeling error along the joint, where the process's various degrees of freedom have been combined into single norm measures of position and orientation error via the 2-norm (e.g., Euclidean distance).

Figure 4.46 Up-down joint: norm modeling errors

Shown in Table 4.3 are the summary statistics for the various modeling errors across the entire joint. As can be seen, the robot was found capable of identifying the true position and orientation of the joint's center with a mean norm positioning error of 0.09 mm and orientation error of 0.83°, despite the discrepancies in z-position and pitch angle between the actual joint and as-designed joint.

Table 4.3 Up-down joint: modeling error statistics

	Abs Error Mean	Abs Error Std Dev	Norm Error Mean	Norm Error Std Dev
Position			0.09	0.03
Y (mm)	0.05	0.03		
Z (mm)	0.06	0.04		
Orientation			0.83	0.47
Rot-X (deg)	0.48	0.46		
Rot-Y (deg)	0.35	0.28		
Rot-Z (deg)	0.40	0.33		
Area (mm²)	3.06	0.97		

Shown in Figure 4.47 are images of the actual caulk deposited into the joint as a result of the robot's adaptive joint filling operation. The robot appears to have successfully filled the joint with

an appropriate amount of caulk, but it did leave a thick layer of excess caulk in the joint's concave regions.
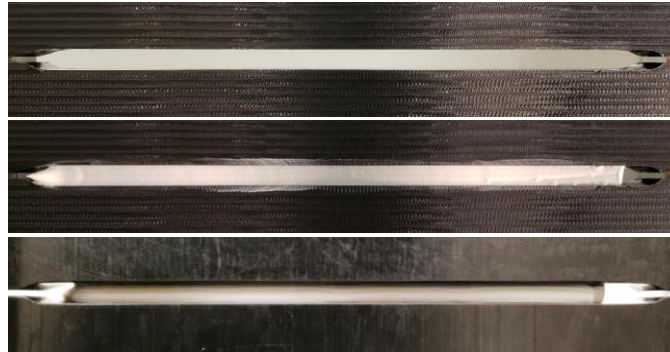
Figure 4.47 Up-down joint: visual results

The mean processing time for the modeling and planning stages was found to be $0.44\ s/profile$, or a total of $24\ s$ for the entire joint (51 profiles).

4.5.3.4   Roll Angle Joint

The roll angle joint tests the robot's ability to compensate for discrepancies in the joint's roll angle. Shown in Figure 4.48 are the raw data collected during the sensor data collection step, the as-perceived joint model from the workpiece modeling step, the adapted fill plan from the work planning step, and the actual volume of caulk deposited from the plan execution step. These results offer a visually intuitive sense of what the robot perceived, planned, and performed. It is worth noting the excess caulk deposited on the surface of the specimen toward the outer side of each roll.
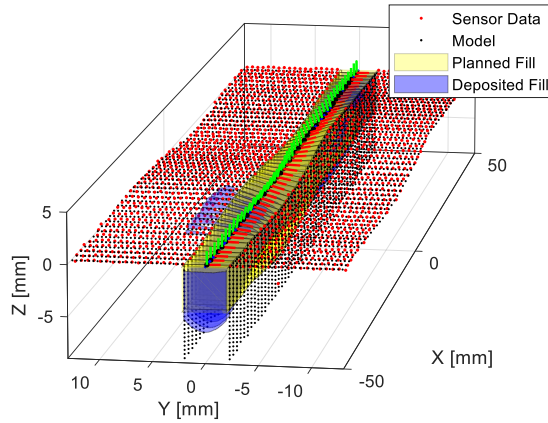
Figure 4.48 Roll angle joint: perceived model, fill plan, and deposited fill

Shown in Figure 4.49 is a comparison between the as-perceived joint model and the actual model along the length of the joint for each of the process's degrees of freedom, thereby providing a comparison between the true joint and the robot's perception of it. Also shown are the cross-sectional areas (i.e., differential volumes) of the actual caulk deposited into the joint.
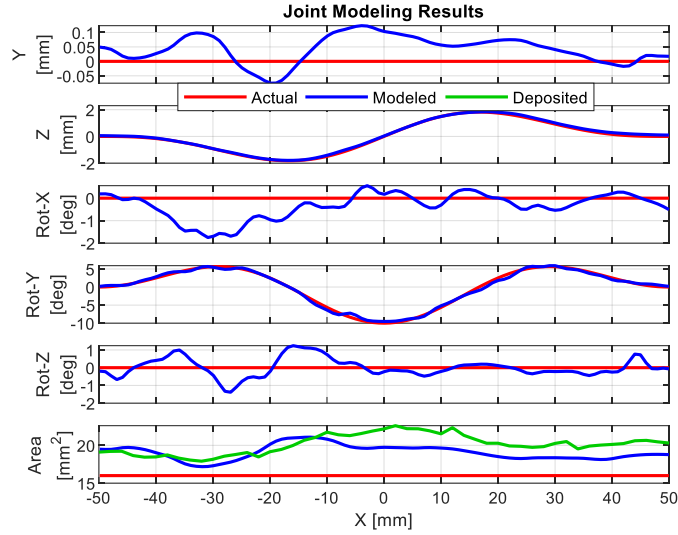

Figure 4.49 Roll angle joint: perceived model profiles

Shown in Figure 4.50 is the absolute modeling error along the length of the joint for each of the process's degrees of freedom, thereby providing a measure of how much error exists between the

true joint and the robot's perception of it. Also shown is the absolute error between the volume of

the caulk deposited and the theoretically ideal amount for the specimen.

Figure 4.50 Roll angle joint: absolute modeling errors

Shown in Figure 4.51 are consolidated measures of modeling error along the joint, where the

process's various degrees of freedom have been combined into single norm measures of position

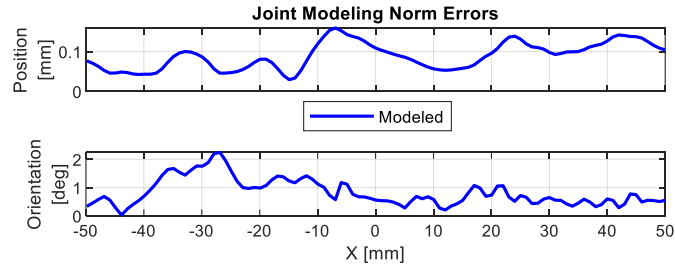and orientation error via the 2-norm (e.g., Euclidean distance).

Figure 4.51 Roll angle joint: norm modeling errors

Shown in Table 4.4 are the summary statistics for the various modeling errors across the entire

joint. As can be seen, the robot was found capable of identifying the true position and orientation

of the joint's center with a mean norm positioning error of $0.11\ mm$ and orientation error of $1.46°$, despite the discrepancies in roll angle between the actual joint and as-designed joint.

Table 4.4 Roll angle joint: modeling error statistics

	Abs Error Mean	Abs Error Std Dev	Norm Error Mean	Norm Error Std Dev
Position			0.11	0.04
Y (mm)	0.05	0.03		
Z (mm)	0.09	0.05		
Orientation			1.46	0.76
Rot-X (deg)	1.07	0.87		
Rot-Y (deg)	0.62	0.47		
Rot-Z (deg)	0.34	0.29		
Area (mm²)	2.98	0.94		

Shown in Figure 4.52 are images of the actual caulk deposited into the joint as a result of the robot's adaptive joint filling operation. The robot appears to have successfully filled the joint with an appropriate amount of caulk, but it did leave a thick ridge of excess caulk on the specimen toward the outer side of each roll maneuver.

Figure 4.52 Roll angle joint: visual results

The mean processing time for the modeling and planning stages was found to be $0.55\ s/profile$, or a total of $29\ s$ for the entire joint (51 profiles).

4.5.3.5 Narrow-Wide Joint

The narrow-wide joint tests the robot's ability to compensate for discrepancies in the joint's width (via tool speed). Shown in Figure 4.53 are the raw data collected during the sensor data collection step, the as-perceived joint model from the workpiece modeling step, the adapted fill plan from the work planning step, and the actual volume of caulk deposited from the plan execution step. These results offer a visually intuitive sense of what the robot perceived, planned, and performed.
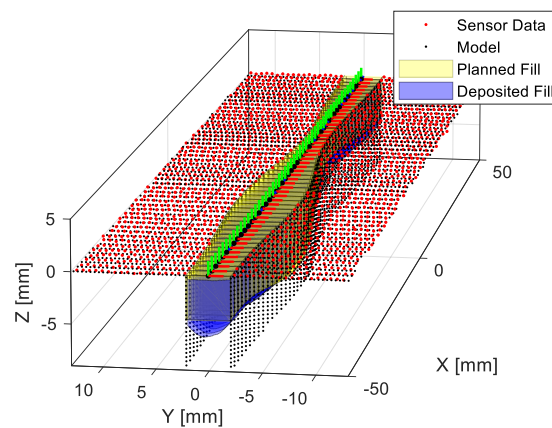
Figure 4.53 Narrow-wide joint: perceived model, fill plan, and deposited fill

Shown in Figure 4.54 is a comparison between the as-perceived joint model and the actual model along the length of the joint for each of the process's degrees of freedom, thereby providing a comparison between the true joint and the robot's perception of it. Also shown are the cross-sectional areas (i.e., differential volumes) of the actual caulk deposited into the joint.
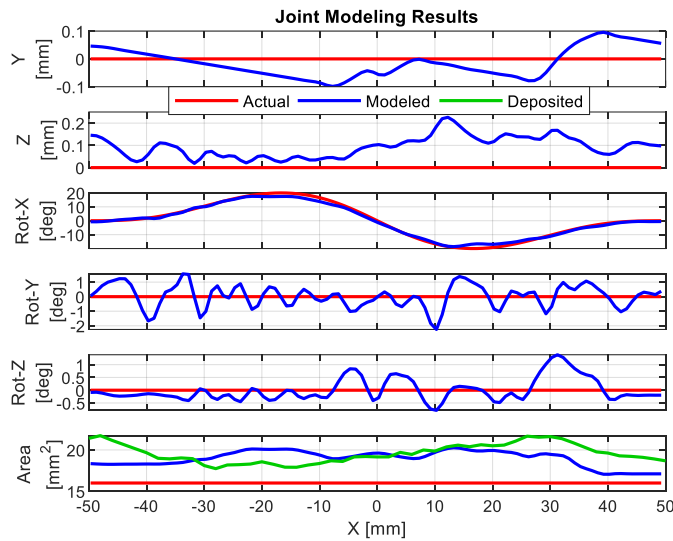
155
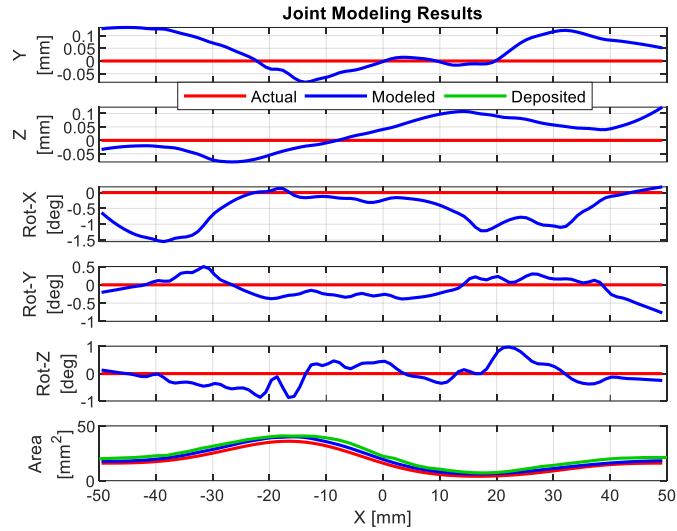
Figure 4.54 Narrow-wide joint: perceived model profiles

Shown in Figure 4.55 is the absolute modeling error along the length of the joint for each of the process's degrees of freedom, thereby providing a measure of how much error exists between the true joint and the robot's perception of it. Also shown is the absolute error between the volume of the caulk deposited and the theoretically ideal amount for the specimen.
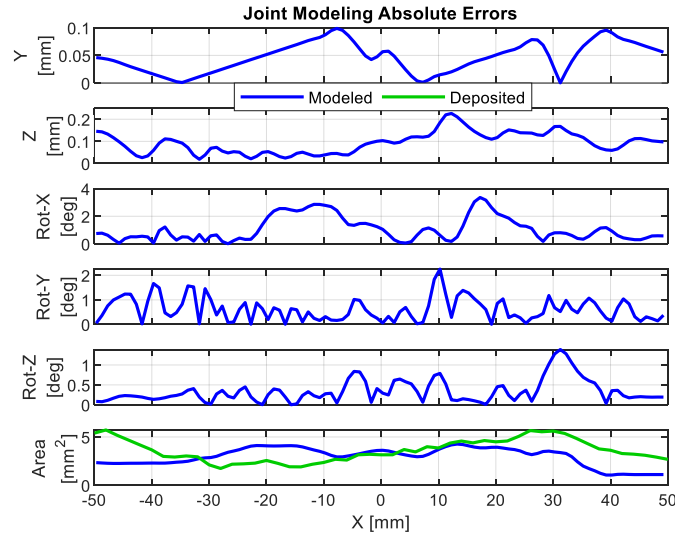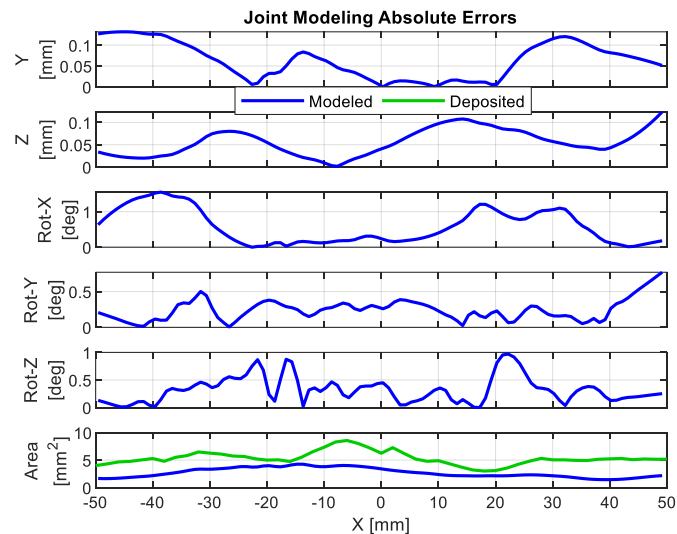

Figure 4.55 Narrow-wide joint: absolute modeling errors

Shown in Figure 4.56 are consolidated measures of modeling error along the joint, where the process's various degrees of freedom have been combined into single norm measures of position and orientation error via the 2-norm (e.g., Euclidean distance).
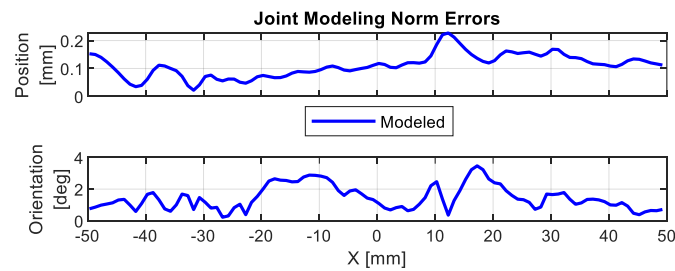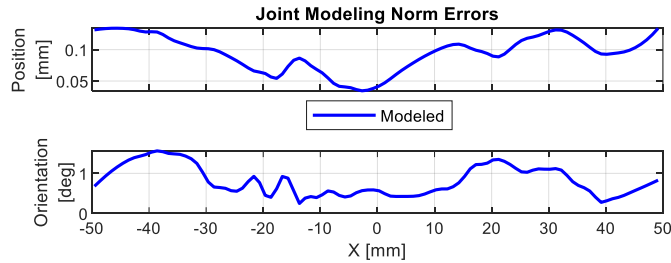

Figure 4.56 Narrow-wide joint: norm modeling errors

Shown in Table 4.5 are the summary statistics for the various modeling errors across the entire joint. As can be seen, the robot was found capable of identifying the true position and orientation of the joint's center with a mean norm positioning error of 0.09 mm and orientation error of 0.82°, despite the discrepancies in width between the actual joint and as-designed joint.

Table 4.5 Narrow-wide joint: modeling error statistics

	Abs Error Mean	Abs Error Std Dev	Norm Error Mean	Norm Error Std Dev
Position			0.09	0.03
Y (mm)	0.06	0.04		
Z (mm)	0.06	0.03		
Orientation			0.82	0.37
Rot-X (deg)	0.58	0.48		
Rot-Y (deg)	0.24	0.15		
Rot-Z (deg)	0.32	0.23		
Area (mm²)	2.70	0.87		

Shown in Figure 4.57 are images of the actual caulk deposited into the joint as a result of the robot's adaptive joint filling operation. The robot appears to have successfully filled the joint with an appropriate amount of caulk and with minimal surface residue.
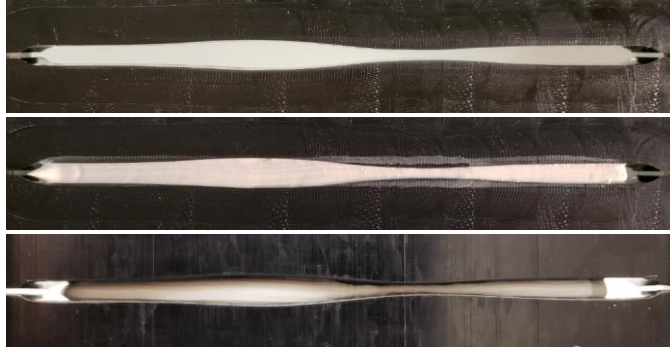
Figure 4.57 Narrow-wide joint: visual results.

The mean processing time for the modeling and planning stages was found to be $0.47\ s/profile$, or a total of $25\ s$ for the entire joint (51 profiles).

A video showing an overview of additional experiments can be found here [88,89].

4.6 Discussion

The robot demonstrated the general ability to identify the joint's true location and geometry, adapt its plan for filling it, execute its adapted plan, and fill the joint with caulk despite discrepancies between the as-designed joint and the actual joint. More generally, the robot demonstrated the basic ability to sense its workpieces, build as-encountered models of its workpieces, adapt its work plan to the as-encountered models, and physically execute its adapted plan to perform work. The caulk deposited by the robot was not without defects, but visual inspection indicates it to be of a level of quality approaching that of a layperson's, and likely within an acceptable range for many real-world applications.

It is worth noting that humans typically rely on visual feedback while performing joint filling operations. To illustrate this point, imagine glancing once at a joint and attempting to weld or caulk it with your eyes closed. In the experiments described here, the robot operated without visual

feedback. That is, the robot sensed the joint once, planned its operation, and executed its plan in an open loop, as opposed to a closed feedback loop. Furthermore, it is common for humans to supplement their visual sensory input with tactile input, which this robot did not use. Such factors might be kept in mind when considering the robot's performance.

Regarding caulk defects, a thin layer of residue was left on either side of the gap in all five experiments. In the up-down joint experiment, a thick layer of excess caulk was left on the face of the specimen in the joint's concave regions, as can be seen in Figure 4.47. It is believed this was a consequence of caulk seeping from the sides of the nozzle through a gap formed between the nozzle's planar surface and the specimen's concave surface.

In the roll angle experiment, a thick ridge of excess caulk was left on the specimen toward the outer side of each roll, as can be seen in Figure 4.52. This is believed to have been caused by an over-simplified fill plan model. As shown in Figure 4.58, the roll orientation of the caulk nozzle is dictated by the gap opening's tangent vector, $\vec{t}_{gap,ind}$, which extends between the upper corners of the two workpieces in the fill plan model. When the gap opening is narrower than the outer width of the caulk nozzle, there is potential for the nozzle to be physically impeded from reaching the corners of the workpieces. In particular, this occurs when an acute angle is formed between the gap opening's tangent vector and the top surface of a workpiece, as demonstrated in Figure 4.58. To improve the robot's performance in such situations, the author recommends including additional fidelity in the fill plan model, as suggested by Figure 4.58, to account for the fact that the contact points lie on the corners of the workpieces for obtuse angles but the corners of the nozzle for acute angles.
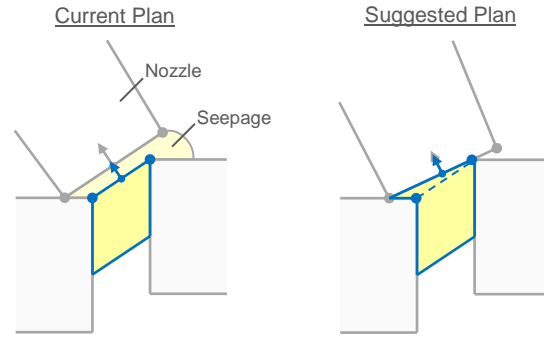
Figure 4.58 Fill plan seepage and suggested improvement

As can be seen in the experimental results, discrepancies exist between the caulk volume planned by the robot and that actually deposited. A portion of such error can be attributed to the sensitivity and variation in the hardware's output flowrate, as suggested by the calibration data in Figure 4.12. Specifically, the pneumatic-drive system was susceptible to pressure fluctuations from the pressure regulator, especially given that a mechanical on/off regulator was used. As a first step toward improving caulk dispensing performance, an electronically controlled pressure regulator could potentially be used instead. However, any pneumatic-drive system would still be susceptible to the complex friction interactions between the plunger and the inside of the caulk tube. A pneumatic-drive system would also be susceptible to the viscous effects related to such factors as caulk temperature, age, and cure. Pneumatic-drive systems are also susceptible to air seeping past the plunger, infiltrating the caulk, and forming air bubbles, which can affect the quality of the caulk deposited into the joint. The author suspects that the aforementioned effects could largely be avoided through the use of a positive displacement pump system, which would likely control the flowrate with greater precision, reduce variation in the output caulk, and ultimately improve joint filling performance.

Furthermore, proper treatment of a joint often includes additional steps and materials. For example, it is sometimes necessary to include cleaning, priming, backing, or tooling [77]. Inclusion of such capabilities could improve the performance of such a system for caulking applications.

However, this research does not propose or claim to provide a complete solution for robotic joint filling. In fact, in many regards, this research is not as much about joint filling as it is about enabling a construction robot to leverage its sensors and BIM model to perceive and model the actual geometry of its workpieces, adapt its work plan, and execute work. A secondary objective of this research is to evaluate the accuracy with which a robot can model its workpieces and adapt its plan, thereby providing an initial benchmark of the performance capabilities that might be achievable for similar tasks using such a framework. Joint filling is merely used as a case study for the purpose of providing a tangible example and demonstration. Thus, the key takeaway from the experiments is that, despite discrepancies between the as-designed joint and actual joint, the robot was able to model the joint and identify the true position and orientation of its center with a mean norm positioning error of $0.11\ mm$ and orientation error of $1.1°$ when averaged across all experiments. Such levels of positioning and orientation accuracy would likely be sufficient for a wide range of construction tasks. More generally, the model-driven adaptive framework presented here was shown to be capable of enabling a robot to adapt to the geometry of its workpieces and perform work by decomposing the BIM model, registering the model to sensor data, adapting the work plan, and executing the adapted plan.

4.7 Conclusions and Future Work

A critical first step in extending the capabilities of construction robots is providing them with the ability to perceive the pose and geometry of their workpieces and perform work in an adaptive

manner. Such adaptive capabilities expand the scope of construction tasks that can be performed by robots and advance their capacity to assist humans in the execution of construction work. This research helps bring the construction industry closer to realizing the benefits of robotics.

The specific contributions of this research are as follows.

- A basic framework is presented for enabling a construction robot to perform geometrically adaptive, model-driven work.
- A formulation is presented for enabling a construction robot to adapt its joint filling work plan to its as-perceived workpiece models.
- The fundamental ability of a construction robot to perform geometrically adaptive, model-driven work on a physical system is demonstrated and evaluated using custom hardware for robotic caulking applications.

A basic framework was proposed to enable a construction robot to adapt to the as-encountered pose and geometry of its workpieces. The framework's latter stages, namely planning and execution, were proposed as a complement to the initial stages of sensing and modeling. The geometrically adaptive model-driven framework was experimentally evaluated to assess the performance capabilities of a robot in executing adaptive work, and to provide a baseline measure of the level of performance that might be achievable with such a system. Joint filling, and in particular caulking, was used as the case study construction activity for testing and demonstrating such adaptive capabilities. A joint of unexpected geometry was placed in front of a robot to see if the robot could successfully adapt its work plan to the circumstances it encountered and perform work. Numerical and pictorial results of the experiments were presented to provide readers with the ability to assess the suitability of such methods for their own application.

It was found that the robot was capable of identifying the true position and orientation of an unexpected joint with a mean norm positioning error of $0.11\ mm$ and a mean norm orientation error of $1.1°$, averaged across five performance assessment experiments. Aside from cosmetic defects, the robot generally proved capable of adaptively filling the joints.

Future work is needed to explore whether such joint filling methods could be extended to include backing (e.g., backer rod) as an additional joint component, since many joint filling applications require such components. Similarly, work is needed to explore the system's sensitivity to unanticipated objects in the sensor data. Work is also needed to address joint termination conditions, such as the beginnings and ends of joints. Research should also be conducted to explore whether a positive displacement drive, as opposed to a pneumatic drive, can be used to improve the robot's caulking performance. Similarly, the work plan can be expanded to accommodate the joint conditions shown in Figure 4.58, and visual or tactile feedback can be incorporated into the robot's operation. Future work is also needed to extend the adaptive methods to 3D, since many objects have geometries more conducive to 3D analysis than 2D. Lastly, future work is needed to expand the adaptive framework to other construction tasks by further generalizing the methods so they can be applied to a wider range of construction activities.

Chapter 5

Conclusion

5.1 Significance of Research

With the ability to localize and perform work in a geometrically adaptive manner, construction robots can begin functioning as useful partners on actual construction sites. This research establishes a key fundamental framework that allows robots to begin performing actual construction work, as summarized in Figure 5.1.
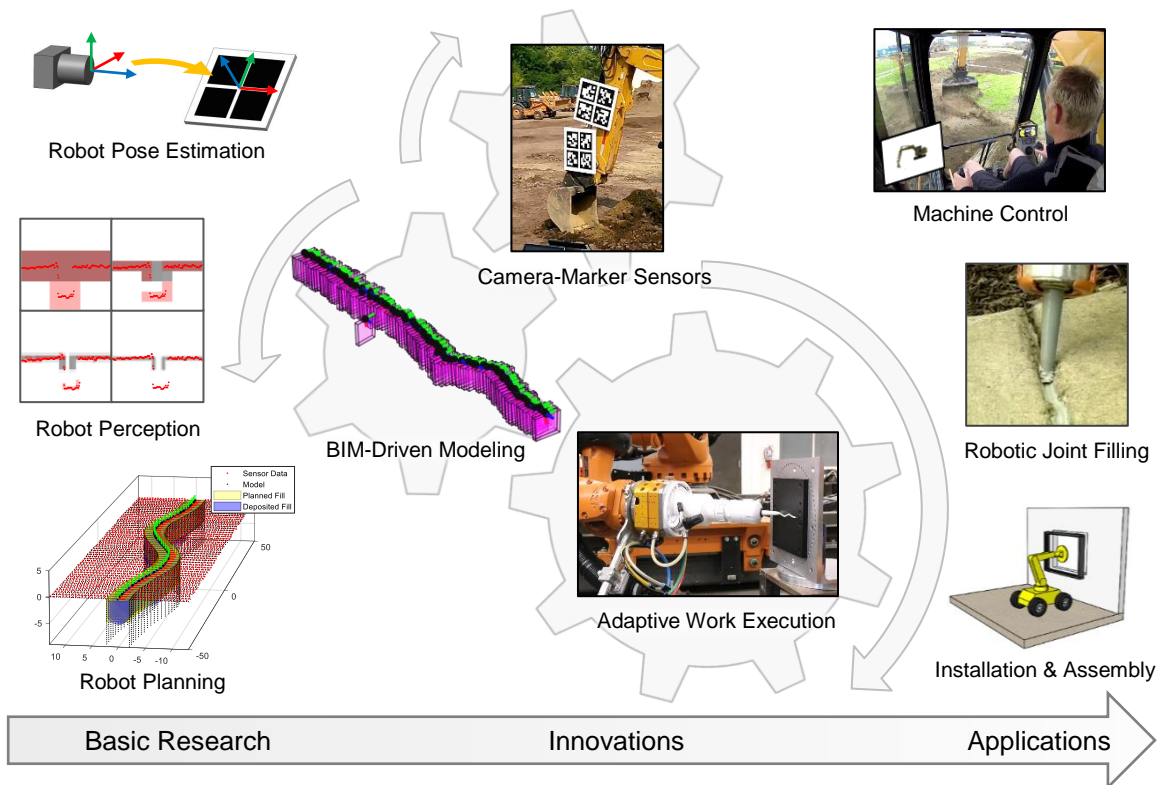
Figure 5.1 Research summary

Additionally, the notion of centralizing such capabilities around the BIM model is a simple yet powerful idea. It is envisioned that this BIM-based geometrically adaptive framework will permit a rapid proliferation of robotic capabilities by allowing the same general methods to be applied to a wide variety of construction objects and activities. Lastly, this research is expected to serve as a seminal work for future BIM-based geometrically adaptive manipulation research, and a cornerstone for many subsequent technological advances in construction robotics.

5.2   Research Contributions

In general, this research provides a ubiquitous sensor system for construction robot localization and an adaptive manipulation framework to enable construction robots to adapt to the geometry of their workpieces and perform work. The specific contributions of this research are as follows.

- Four iterations of camera-marker sensor prototypes were designed, fabricated, and tested for robot localization in construction applications.
- Two construction component model fitting techniques were developed to enable a construction robot to perceive and model the actual pose and geometry of its workpieces.
- The formulation of an adapted work plan was demonstrated using joint filling as an example construction task.
- The fundamental ability of a construction robot to perform geometrically adaptive, model-driven work on a physical system was demonstrated and evaluated using custom hardware.
- A basic framework was established and detailed for enabling a construction robot to perform geometrically adaptive, model-driven work.

5.3 Future Directions

5.3.1 Localization

Further research should be conducted into camera-marker localization for construction robots to determine whether it is better to use robot-mounted cameras and site-mounted markers, or vice versa. To achieve complete sensor coverage, it will likely be necessary to use a large number of one or the other. As cameras are more expensive than markers, mounting cameras on robots may permit fewer total cameras, thereby lowering equipment costs. However, site-mounted cameras could offer the potential for other functions, such as asset tracking and security.

Additionally, other localization approaches should be explored, such as registering sensor data to BIM models and filtering to estimate robot pose. To obtains such pose estimates, sensors such 2D LIDAR sensors, 3D LIDAR sensors, structured light sensors, monocular cameras with structure from motion, and stereo cameras with feature extraction can be used to generate point cloud data, which can be registered with BIM models and filtered with any of the Kalman filter or particle filter variants to estimate the robot's pose. Such methods offer the benefit that neither the environment nor the robots need be outfitted with markers. However, research is needed to determine how suitable such methods are for construction robot applications.

5.3.2 Workpiece Perception

In order to collect sensor data of its workpieces, a construction robot must decide where to aim its sensor and how to move between desired sensor poses. In this research, such a sensor data collection plan is provided manually. There are opportunities to explore how a construction robot can autonomously determine where it needs to aim its sensor to collect the information it needs,

and how to plan a path that reaches all such desired poses without encountering collisions or violating any kinematic constraints along the way.

Additionally, once a robot has collected sensor data and registered its workpiece models, it might be helpful if the robot could evaluate the uncertainty in its modeling results. If the modeling uncertainty were to exceed an allowable tolerance for a particular object and operation, then the robot could continually collect more sensor data and re-register until it has achieved sufficient confidence in its modeling results. Such an extension could be explored to make geometric adaptive manipulation more robust for real-world applications.

5.3.3 Workpiece Modeling

During the model registration processes described in this research, data association was performed using nearest-neighbor point correspondences. There may be opportunities to explore whether sensor data association could be improved through the use of machine learning to identify point-object correspondences. It may be possible to use an RGBD sensor, apply machine learning to identify objects in the RGB data, and then apply the semantic object labels to the depth data (i.e., the point cloud).

5.3.4 Adaptive Planning

In this research, work plans are devised without taking into consideration obstacles or kinematic constraints (e.g., joint limits, joint singularities, etc.). In reality, construction work requires considerable maneuvering and obstacle avoidance. There are opportunities to explore how such obstacles and kinematic constraints can be autonomously incorporated into work plans to produce a more robust solution for onsite operation.

5.3.5 Work Execution

This research primarily focuses on enabling robots to become more useful construction partners by adapting their work plans to the geometry of their workpieces so they can help perform work with some degree of autonomy. While perceiving workpieces and adapting to their geometry substantially advances a robot's utility as a construction partner, it may not be the only criterion by which to judge the effectiveness of a partner. For example, it may also be important to be cooperative, coordinated, and considerate (e.g., for safety reasons). Thus, there are research opportunities to explore other means of enabling robots to become better partners, likely centered more on human-robot interaction aspects.

Bibliography

[1]   P.M. Goodrum, C.T. Haas, P. Teicholz, U.S. Construction Labor Productivity Trends, 1970-1998, Journal of Construction Engineering and Management, 127 (5) (2001) 427 - 429, http://dx.doi.org/10.1061/(ASCE)0733-9364(2001)127:5(427).

[2]   CII, IR252-2b - A Guide to Construction Rework Reduction, Construction Industry Institute, Austin, Texas, 2011, p. 41, Accessed 11/05/2018, https://www.construction-institute.org/resources/knowledgebase/knowledge-areas/construction-execution/topics/rt-252/pubs/ir252-2b.

[3]   U.S. Bureau of Labor Statistics, Industry Injury and Illness Data 2007-2016, Accessed 05/23/2018, https://www.bls.gov/iif/oshsum.htm.

[4]   U.S. Bureau of Labor Statistics, Census of Fatal Occupational Injuries 2007-2016, Accessed 05/23/2018, https://www.bls.gov/iif/oshcfoi1.htm.

[5]   C. Balaguer, Nowadays trends in robotics and automation in construction industry: Transition from hard to soft robotics, 21st International Symposium of Automation and Robotics in Construction, Citeseer, 2004, https://doi.org/10.22260/ISARC2004/0009.

[6]   W. Perry, A. Broers, F. El-Baz, W. Harris, B. Healy, and W.D. Hillis, Grand challenges for engineering, National Academy of Engineering, Washington, D.C., 2008.

[7]   J. O'Brien, K. Saidi, A. Lytle, Robotics in Construction, in: B. Siciliano, O. Khatib (Eds.), Springer Handbook of Robotics, Springer-Verlag, Berlin Heidelberg, ISBN 9783540239574, 2008, pp. 1079-1099.

[8]   C. Feng, Y. Xiao, A. Willette, W. McGee, V.R. Kamat, Towards autonomous robotic in-situ assembly on unstructured construction sites using monocular vision, International Symposium on Automation and Robotics in Construction, Vol. 31, 2014, pp. 163-170, http://dx.doi.org/10.22260/ISARC2014/0022.

[9]   C.T. Milberg, Application of Tolerance Management to Civil Systems, Civil and Environmental Engineering, University of California, Berkely, 2006, ISBN 9781109922721.

[10]   J.G. Everett, Construction automation--basic task selection and development of the CRANIUM, Massachusetts Institute of Technology, 1991.

[11]   Caterpillar, Cat | Grade--Grading System, Autograde & Machine Control Technology | Caterpillar, 2019, Accessed 02/16/2019, https://www.cat.com/en_US/by-industry/construction/cat-connect/technologies/grade-technology.html.

[12]   Trimble, Grade Control for Dozers | Trimble Civil Engineering and Construction, 2019, Accessed 02/16/2019, https://construction.trimble.com/products-and-solutions/grade-control-dozers.

[13]   Caterpillar, Cat | Realizing the Benefits of Autonomous Haulage | Caterpillar, 2019, Accessed 02/16/2019, https://www.cat.com/en_US/articles/customer-stories/mining/viewpoint/autonomous-haulage-gains.html.

[14]   Caterpillar, Cat | Boosting Efficiency Through Autonomous Dozing | Caterpillar, 2019, Accessed 02/16/2019, https://www.cat.com/en_US/articles/customer-stories/mining/viewpoint/autonomous-dozing-boosts-efficiency.html.

[15]   Komatsu, Innovation | Komatsu America Corp, 2019, Accessed 02/16/2019, https://www.komatsuamerica.com/equipment/innovation#Advanced%20Technology.

[16]   Komatsu, Komatsu first to operate Autonomous Haulage System over dedicated private LTE network, 2019, Accessed 02/16/2019, https://www.komatsuamerica.com/our-company/press-releases/2019-01-21-ahs-lte.

[17]   G. Cheok, W.C. Stone, R. Lipman, Automated Earthmoving Status Determination, Robotics 2000, ISBN 9780784404768, 2000, pp. 111-119, http://dx.doi.org/10.1061/40476(299)14.

[18]   Y.-K. Cho, C.T. Haas, K.A. Liapi, S.V. Sreenivasan, Rapid visualization of geometric information in a construction environment, Fifth International Conference on Information Visualisation, IEEE, ISBN 9780769511955, 2001, pp. 31-36, http://dx.doi.org/10.1109/IV.2001.942036.

[19]   Y.-K. Cho, C.T. Haas, K. Liapi, S.V. Sreenivasan, A framework for rapid local area modeling for construction automation, Automation in Construction, 11 (6) (2002) 629-641, http://dx.doi.org/10.1016/S0926-5805(02)00004-3.

[20]   Y.-K. Cho, C.T. Haas, Rapid Geometric Modeling for Unstructured Construction Workspaces, Computer-Aided Civil and Infrastructure Engineering, 18 (4) (2003) 242-253, http://dx.doi.org/10.1111/1467-8667.00314.

[21]   J. McLaughlin, S.V. Sreenivasan, C. Haas, K. Liapi, Rapid Human-Assisted Creation of Bounding Models for Obstacle Avoidance in Construction, Computer-Aided Civil and Infrastructure Engineering, 19 (1) (2004) 3-15, http://dx.doi.org/10.1111/j.1467-8667.2004.00333.x.

[22]   C. Kim, C.T. Haas, K.A. Liapi, Rapid, on-site spatial information acquisition and its use for infrastructure operation and maintenance, Automation in Construction, 14 (5) (2005) 666-684, http://dx.doi.org/10.1016/j.autcon.2005.02.002.

[23]   C. Kim, C.T. Haas, K.A. Liapi, C.H. Caldas, Human-Assisted Obstacle Avoidance
         System Using 3D Workspace Modeling for Construction Equipment Operation, Journal
         of Computing in Civil Engineering, 20 (3) (2006) 177-186,
         http://dx.doi.org/10.1061/(ASCE)0887-3801(2006)20:3(177).

[24]   A. Stentz, J. Bares, S. Singh, P. Rowe, A robotic excavator for autonomous truck loading,
         Autonomous Robots, 7 (2) (1999) 175-186, http://dx.doi.org/10.1023/A:1008914201877.

[25]   H. Kim, A.F. Rauch, C.T. Haas, C. Browne, Dimensional Ratios for Stone Aggregates
         from Three-Dimensional Laser Scans, Journal of Computing in Civil Engineering, 16 (3)
         (2002) 175-183, http://dx.doi.org/10.1061/(ASCE)0887-3801(2002)16:3(175).

[26]   B. Kahane, Y. Rosenfeld, Real-time "Sense-and-Act" operation for construction robots,
         Automation in Construction, 13 (6) (2004) 751-764,
         http://dx.doi.org/10.1016/j.autcon.2004.05.004.

[27]   P. Sicard, M.D. Levine, Joint recognition and tracking for robotic arc welding, IEEE
         Transactions on Systems, Man, and Cybernetics, 19 (4) (1989) 714-728,
         http://dx.doi.org/10.1109/21.35336.

[28]   J.S. Kim, Y.T. Son, H.S. Cho, K.I. Koh, A robust method for vision-based seam tracking
         in robotic arc welding, Tenth International Symposium on Intelligent Control, 1995, pp.
         363-368, http://dx.doi.org/10.1109/ISIC.1995.525084.

[29]   C. Kim, H. Son, H. Kim, S.H. Han, Applicability of flash laser distance and ranging to
         three-dimensional spatial information acquisition and modeling on a construction site,
         Canadian Journal of Civil Engineering, 35 (11) (2008) 1131-1341,
         http://dx.doi.org/10.1139/L08-082.

[30]   S.J. Keating, J.C. Leland, L. Cai, N. Oxman, Toward site-specific and self-sufficient
         robotic fabrication on architectural scales, Science Robotics, 2 (5) (2017),
         http://dx.doi.org/10.1126/scirobotics.aam8986.

[31]   W. Wang, H. Chi, S. Zhao, Z. Du, A control method for hydraulic manipulators in
         automatic emulsion filling, Automation in Construction, 91 (2018) 92-99,
         http://dx.doi.org/10.1016/j.autcon.2018.03.001.

[32]   M. Lussi, T. Sandy, K. Doerfler, N. Hack, F. Gramazio, M. Kohler, J. Buchli, Accurate
         and adaptive in situ fabrication of an undulated wall using an on-board visual sensing
         system, IEEE International Conference on Robotics and Automation (ICRA), 2018,
         http://dx.doi.org/10.3929/ethz-b-000227968.

[33]   O. Kermorgant, A magnetic climbing robot to perform autonomous welding in the
         shipbuilding industry, Robotics and Computer-Integrated Manufacturing, 53 (2018) 178-
         186, http://dx.doi.org/10.1016/j.rcim.2018.04.008.

[34] C. Feng, Y. Xiao, A. Willette, W. McGee, V.R. Kamat, Vision guided autonomous robotic assembly and as-built scanning on unstructured construction sites, Automation in Construction, 59 (2015) 128-138, http://dx.doi.org/10.1016/j.autcon.2015.06.002.

[35] E. Lublasser, L. Hildebrand, A. Vollpracht, S. Brell-Cokcan, Robot assisted deconstruction of multi-layered façade constructions on the example of external thermal insulation composite systems, Construction Robotics, 1 (1) (2017) 39-47, http://dx.doi.org/10.1007/s41693-017-0001-7.

[36] V. Helm, S. Ercan, F. Gramazio, M. Kohler, Mobile robotic fabrication on construction sites: DimRob, IEEE/RSJ 2012 International Conference on Intelligent Robots and Systems (IROS), 2012, pp. 4335-4341, http://dx.doi.org/10.1109/IROS.2012.6385617.

[37] J. Willmann, M. Knauss, T. Bonwetsch, A.A. Apolinarska, F. Gramazio, M. Kohler, Robotic timber construction—Expanding additive fabrication to new dimensions, Automation in Construction, 61 (2016) 16-23, http://dx.doi.org/10.1016/j.autcon.2015.09.011.

[38] Infrastructure Resources LLC, One Call and State Law Directory (USA), Excavation Safety Guide and Directory, Vol. 1, 2014, pp. 49-55, http://www.excavationsafetyguide.com/display_issue.php.

[39] CGA, Damage Prevention, 2015, Accessed Jan 26, http://www.commongroundalliance.com/damage-prevention.

[40] A. Patel, A. Chasey, Integrating GPS and Laser Technology to Map Underground Utilities Installed Using Open Trench Method, 2010 Construction Research Congress, 2010, p. 628, http://dx.doi.org/10.1061/41109(373)63.

[41] AWWA, Buried No Longer: Confronting America's Water Infrastructure Challenge, American Water Works Association, 2012, pp. 3-4, http://www.awwa.org/Portals/0/files/legreg/documents/BuriedNoLonger.pdf.

[42] J. Bowman, P. Giovannoni, 2013 CGA DIRT Analysis and Recommendations, Vol. 10, Common Ground Alliance, 2013, pp. 2-4, http://www.cga-dirt.com/annual/2013/DIRT_Report_for_2013_Final_20141008_REDUCED.pdf.

[43] PHMSA, Pipeline Incidents by Cause, Pipeline and Hazardous Materials Safety Administration, US Department of Transportation, 2015, https://hip.phmsa.dot.gov/analyticsSOAP/saw.dll?Portalpages.

[44] Transportation Equity Act for the 21st Century, Vol. 112, US, 1998, http://www.gpo.gov/fdsys/pkg/PLAW-105publ178/pdf/PLAW-105publ178.pdf.

[45] J.H. Anspach, The Case for a National As-Builting Standard, American Society of Civil Engineers, 2011, http://content.asce.org/files/pdf/10Asbuiltstandard.pdf.

[46]   S.A. Talmaki, Real-Time Visualization for Prevention of Excavation Related Utility Strikes, Civil Engineering, University of Michigan, 2012, http://hdl.handle.net/2027.42/96133.

[47]   Trimble, Grade Control for Excavators, 2015, Accessed Jan 26, http://construction.trimble.com/products/machine-control/grade-control-for-excavators.

[48]   Leica, Machine Control, Digging Solutions, 2015, Accessed Jan 26, http://www.leica-geosystems.com/en/Digging-Solutions_63797.htm.

[49]   J.D. Bossler, J.B. Campbell, R.B. McMaster, C. Rizos, Manual of Geospatial Science and Technology, 2nd ed., CRC Press, Boca Raton, FL, ISBN 9781420087338, 2010.

[50]   C. Feng, V.R. Kamat, Plane Registration Leveraged by Global Constraints for Context-Aware AEC Applications, Computer-Aided Civil and Infrastructure Engineering, 28 (5) (2013) 325-343, http://dx.doi.org/10.1111/j.1467-8667.2012.00795.x.

[51]   C. Feng, S. Dong, K.M. Lundeen, Y. Xiao, V.R. Kamat, Vision-Based Articulated Machine Pose Estimation for Excavation Monitoring and Guidance, International Symposium on Automation and Robotics in Construction (ISARC), ISARC, Oulu, Finland, 2015, https://doi.org/10.22260/ISARC2015/0029.

[52]   E. Olson, AprilTag: A Robust and Flexible Visual Fiducial System, 2011 IEEE International Conference on Robotics and Automation (ICRA), IEEE, 2011, pp. 3400-3407, http://dx.doi.org/10.1109/ICRA.2011.5979561.

[53]   D.T. Koenig, Process Control Engineering and Quality Control in Job Shops, Manufacturing Engineering: Principles for Optimization, ASME Press, New York, 2007, pp. 279-284, http://dx.doi.org/10.1115/1.802493.ch10.

[54]   K. Leonard, Carpentry, Third ed., American Technical Publishers, Inc., USA, ISBN 9780826907356, 1997.

[55]   E.B. Olson, Real-time correlative scan matching, IEEE International Conference on Robotics and Automation, IEEE, Kobe, Japan, ISBN 9781424427888, 9781424427895, 2009, pp. 4387-4393, http://dx.doi.org/10.1109/ROBOT.2009.5152375.

[56]   D. Holzer, The BIM Manager's Handbook: Guidance for Professionals in Architecture, Engineering and Construction, John Wiley & Sons Inc, Wiley-Blackwell, 2016, http://dx.doi.org/10.1002/9781118982310.

[57]   R. Kümmerle, B. Steder, C. Dornhege, M. Ruhnke, G. Grisetti, C. Stachniss, A. Kleiner, On measuring the accuracy of SLAM algorithms, Autonomous Robots, 27 (4) (2009) 387-407, http://dx.doi.org/10.1007/s10514-009-9155-6.

[58]   E. Allen, J. Iano, Fundamentals of Building Construction: Materials and Methods, John Wiley & Sons, Hoboken, New Jersey, ISBN 9780470074688, 2011.

[59]   S.P. Lloyd, Least squares quantization in PCM, IEEE Transactions on Information Theory, 28 (2) (1982) 129-137, http://dx.doi.org/10.1109/TIT.1982.1056489.

[60]   O. Maimon, L. Rokach, Clustering Methods, Data Mining and Knowledge Discovery Handbook, Springer, New York, ISBN 9780387244358, 2005, pp. 321-352.

[61]   Y. Chen, Medioni, Gérard, Object modelling by registration of multiple range images, Image and Vision Computing, 10 (3) (1992) 145-155, http://dx.doi.org/10.1016/0262-8856(92)90066-C.

[62]   P.J. Besl, H.D. McKay, A method for registration of 3-D shapes, IEEE Transactions on Pattern Analysis and Machine Intelligence, 14 (2) (1992) 239 - 256, http://dx.doi.org/10.1109/34.121791.

[63]   J.E. Bresenham, Algorithm for computer control of a digital plotter, IBM Systems Journal, 4 (1) (1965) 25-30, http://dx.doi.org/10.1147/sj.41.0025.

[64]   A.A. Gohara, M.A. Pazoti, F.A. Silva, D.R. Pereira, A.O. Artero, M.A. Piteri, A New Approach for Urban Roads Detection Using Laser Data and Aerial Digital Images, Journal of Urban and Environmental Engineering, 10 (1) (2016) 83-97, http://dx.doi.org/10.4090/juee.2016.v10n1.083097.

[65]   Y. Okubo, C. Ye, J. Borenstein, Characterization of the Hokuyo URG-04LX laser rangefinder for mobile robot obstacle negotiation, SPIE, Vol. 7332, 2009, http://dx.doi.org/10.1117/12.818332.

[66]   C. Milberg, I. Tommelein, Role of tolerances and process capability data in product and process design integration, Construction Research Congress, 2003, pp. 1-8, http://dx.doi.org/10.1061/40671(2003)93.

[67]   K.M. Lundeen, V.R. Kamat, C.C. Menassa, W. McGee, Scene understanding for adaptive manipulation in robotized construction work, Automation in Construction, 2017 (82C) (2017) 16-30, http://dx.doi.org/10.1016/j.autcon.2017.06.022.

[68]   D.H. Jung, J. Park, M. Schwartz, Towards on-site autonomous robotic floor tiling of mosaics, IEEE 14th International Conference on Control, Automation, and Systems (ICCAS), 2014, pp. 59-63, http://dx.doi.org/10.1109/ICCAS.2014.6987959.

[69]   M. Giftthaler, T. Sandy, K. Dörfler, I. Brooks, M. Buckingham, G. Rey, M. Kohler, F. Gramazio, J. Buchli, Mobile robotic fabrication at 1:1 scale: the In situ Fabricator, Construction Robotics, 1 (1-4) (2017) 3-14, http://dx.doi.org/10.1007/s41693-017-0003-5.

[70]   M. Dinham, G. Fang, Detection of fillet weld joints using an adaptive line growing algorithm for robotic arc welding, Robotics and Computer Integrated Manufacturing, 30 (3) (2014) 229-243, http://dx.doi.org/10.1016/j.rcim.2013.10.008.

[71]   L. Sukhan, J. Daesik, K. Eunyoung, H. Suyeon, H. JungHyun, A real-time 3D workspace modeling with stereo camera, IEEE/RSJ 2005 International Conference on Intelligent Robots and Systems, 2005, pp. 2140-2147, http://dx.doi.org/10.1109/IROS.2005.1545105.

[72]   R.P. Hoyt, J.I. Cushing, J.T. Slostad, G. Jimmerson, T. Moser, G. Kirkos, M.L. Jaster, N.R. Voronka, SpiderFab: An architecture for self-fabricating space systems, AIAA SPACE 2013 Conference and Exposition, ISBN 9781624102394, 2013.

[73]   P. Kah, M. Shrestha, E. Hiltunen, J. Martikainen, Robotic arc welding sensors and programming in industrial applications, International Journal of Mechanical and Materials Engineering, 10 (1) (2015) 1-16, http://dx.doi.org/10.1186/s40712-015-0042-y.

[74]   R.P. Manorathna, P. Phairatt, P. Ogun, T. Widjanarko, M. Chamberlain, L. Justham, S. Marimuthu, M.R. Jackson, Feature extraction and tracking of a weld joint for adaptive robotic welding, IEEE 2014 13th International Conference on Control Automation Robotics & Vision (ICARCV), 2014, pp. 1368-1372, http://dx.doi.org/10.1109/ICARCV.2014.7064515.

[75]   P. Maiolino, R. Woolley, D. Branson, P. Benardos, A. Popov, S. Ratchev, Flexible robot sealant dispensing cell using RGB-D sensor and off-line programming, Robotics and Computer Integrated Manufacturing, 48 (2017) 188-195, http://dx.doi.org/10.1016/j.rcim.2017.04.004.

[76]   M. Rodrigues, M. Kormann, C. Schuhler, P. Tomek, An intelligent real time 3D vision system for robotic welding tasks, IEEE 2013 9th International Symposium on Mechatronics and its Applications (ISMA), 2013, pp. 1-6, http://dx.doi.org/10.1109/ISMA.2013.6547393.

[77]   ASTM International, Standard Guide for Use of Joint Sealants, Vol. C1193 - 16, 2016, http://dx.doi.org/10.1520/C1193-16.

[78]   ANSI/AWS, D1.1/D1.1M:2015, Structural Welding Code--Steel, ISBN 9780871718648, 2015.

[79]   M.W. Spong, S. Hutchinson, M. Vidyasagar, Robot Modeling and Control, Wiley New York, ISBN 9780471649908, 2006.

[80]   J.J. Craig, Introduction to Robotics: Mechanics and Control, 3rd ed., Pearson/Prentice Hall, Upper Saddle River, N.J., ISBN 9780201543612, 2005.

[81]   H. Liu, S. Shah, W. Jiang, On-line outlier detection and data cleaning, Computers & Chemical Engineering, 28 (9) (2004) 1635-1647, https://dx.doi.org/10.1016/j.compchemeng.2004.01.009.

[82]   A. Savitzky, M.J.E. Golay, Smoothing and differentiation of data by simplified least squares procedures, Analytical Chemistry, 36 (8) (1964) 1627-1639, http://dx.doi.org/10.1021/ac60214a047.

[83] C. De Boor, A Practical Guide to Splines, Springer-Verlag New York, ISBN 9780387903569, 1978.

[84] D. Pigram, W. McGee, Formation Embedded Design: A methodology for the integration of fabrication constraints into architectural design, 31st Annual Conference of the Association for Computer Aided Design in Architecture, ACADIA, Banff, Alberta, Canada, ISBN 9781613645956, 2011, pp. 122-131, http://hdl.handle.net/10453/19363.

[85] K.M. Lundeen, V.R. Kamat, C.C. Menassa, W. McGee, Random Joint Filling Experiment, Multi-View, YouTube, 2018, Accessed 07/13/2018, https://youtu.be/yjCbchW__-0.

[86] K.M. Lundeen, V.R. Kamat, C.C. Menassa, W. McGee, Door Frame Joint Filling Experiment, Multi-View, YouTube, 2018, Accessed 07/13/2018, https://youtu.be/SJnAzx-uzrs.

[87] K.M. Lundeen, V.R. Kamat, C.C. Menassa, W. McGee, Left-Right Joint Filling Experiment, Multi-View, YouTube, 2018, Accessed 06/08/2018, https://youtu.be/Zqr6Mr_NHes.

[88] R. Coelius, Autonomous robot construction is here, The Michigan Engineer News Center, 2017, Accessed 12/15/2017, https://news.engin.umich.edu/2017/08/autonomous-robot-construction-is-here.

[89] R. Coelius, On-Site Construction Robots | Vineet Kamat, YouTube, 2017, Accessed 06/08/2018, https://youtu.be/b61ekbcfUfE.