Machine Vision for In-Process Inspection on an Automated Peptide Manufacturing Platform

by

Abigail Campbell
B.S. Mechanical Engineering, University of Utah (2018)

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Master of Engineering in Advanced Manufacturing and Design
at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
September 2020

Author . . . . . .
Department of Mechanical Engineering
August 14th, 2020

Certified by . . . . . .
Brian W. Anthony
Principal Research Scientist, Mechanical Engineering

Accepted by. . . . . .
Nicolas Hadjiconstantinou
Chairman, Department Committee on Graduate Theses

This page is intentionally left blank.

Machine Vision for In-Process Inspection of an Automated Peptide Manufacturing Platform

by

Abigail Campbell

B.S. Mechanical Engineering, University of Utah (2018)

Submitted to the Department of Mechanical Engineering

in partial fulfillment of the requirements for the degree of

Master of Engineering in Advanced Manufacturing & Design

Abstract

Automated inspection of a manufacturing line utilizing machine vision is a powerful tool used to increase efficiency, reduce variation, and maintain a high-quality standard in modern manufacturing systems. The aim of this thesis is to develop and implement a machine vision system for Mytide Therapeutics' automated peptide manufacturing platform.

Throughout the peptide manufacturing process, vials are used for liquid and solid handling of the peptide compounds through the disparate process steps. To ensure the vials of peptides are transferred between these steps properly, several functions were developed to analyze the images at key points in the process. These images are analyzed to ensure a vial is present when required, is gripped properly, and is transferred successfully by the robot to and from each station. Moreover, a function to estimate the volume of resin inside of a vial before and after peptide synthesis is developed as a method of collecting data relevant to the process that would not be collected otherwise. The proposed algorithms are designed to provide key insights to the manufacturing process, ensure the process runs smoothly, reduce overall system downtime, and inform future system improvements.

This thesis presents an overview of the machine vision system, details about the algorithms developed to perform the image analysis, and the methods for implementation of the image analysis functions into the manufacturing platform.

Thesis Supervisor: Brian Anthony

Title: Senior Research Scientist, Mechanical Engineering

This page is intentionally left blank.

Acknowledgments

Throughout this project, I have received a great deal of support from some amazing people.

First, I want to thank Mytide for the opportunity to work with them on this project and for their continuing support. Thank you to Chase, Kevin, and Dale for being great mentors and for your insight and guidance throughout this project. Also, thank you to Liam, Justin, and Joel for always taking the time to answer my questions. It has been a privilege to be a part of the team.

I would also like to acknowledge my supporters at MIT. Thanks to my advisor, Dave Hardt, for being an encouraging mentor and being invested in my success during my time here. Thank you to Jose Pacheco and Brian Anthony for your support, career advice, and for giving me the opportunity to be a part of this program. I have loved my time at MIT and will forever cherish the knowledge I have gained and the relationships I have made here.

I also want to thank my thesis teammates, Liudi, and Ben. I learned so much from both of you and am so glad that we got to share this experience together this summer. Liudi, I couldn't have asked for a better lunch buddy and writing party partner. Thank you for the scones, bubble tea, and plant advice. Your friendship has made this crazy time one I can look back on with great memories.

Next, I want to thank my family. To my parents, thank you for your love, patience, encouragement, and faith in me. Thanks for always answering my late-night phone calls, sending me homemade banana bread, putting up with my stubbornness, and giving me the best foundation that I could ask for. You are my rocks, and I couldn't be prouder to be your daughter. To my sisters, Jane, Sam, and Isabel, you guys are my best friends and I wouldn't be the person I am today without you. I look forward to many more years of vaguely relevant movie quotes and beating you in Catan. To my brother-in-law, Alex, thank you for never failing to believe in me, and to my niece, Parker, you are the light of my life.

I want to thank my Grandpa Peterson for always encouraging me, being excited about each new step I take, and paving the way for me to follow in your footsteps. Thank you to my Grandma Peterson for always looking out for me, and my Nana and Grandpa Campbell, whose love and support has made all of this possible. I also want to thank my family here in Boston, Pat, Erin, Eliza, Annie, and Jake for being my home away from home and always making sure I was taken care of and well-fed.

I have also had the privilege of being mentored by many incredible people throughout my education at the University of Utah and career at Kennecott, Bard, LANL, and Wheeler. Your invaluable support has pushed me to where I am today and I owe much of my joy in engineering to you.

Finally, I want to thank my MEng cohort. Even though we didn't get the time together that we would have liked, you guys made this year an amazing experience and have become some of my best friends. Go forth and smash it.

This page is intentionally left blank.

Table of Contents

List of Figures

List of Tables

List of Equations

Chapter 1: Background and Project Motivation

1.1. Introduction and Statement of Purpose

This thesis outlines the implementation a machine vision system on an automated peptide manufacturing platform at Mytide Therapeutics in Boston, Massachusetts. Mytide Therapeutics is developing a fully automated peptide synthesis platform to manufacture and deliver custom peptides to customers at a higher speed and purity than the industry standard. Implementing a machine vision system will increase the versatility of the synthesis platform with automatic fault condition detection and visual data collection to inform future process improvements.

This thesis is part of a project conducted by Liudi Yang, Ben Russell, and the author to fulfill the requirements for a Masters of Engineering in Advanced Manufacturing Design at The Massachusetts Institute of Technology. The Mytide thesis group was given the task of increasing the process throughput with improved synthesis and purification methods and with decreased downtime of the automated peptide manufacturing system at Mytide. The project took place Mytide Therapeutics between March and August of 2020.

1.1.1. Division of Labor

The Mytide thesis project was divided into three projects, each to be completed by an individual team member of the Mytide thesis group. Each project took a unique approach by focusing on improving a different part of the process that contributed towards the overall goal of improving the workflow of the system. Liudi focused on synthesis anomaly detection and peptide purity prediction, which can have a significant effect on eliminating wasted time processing faulty product. Ben worked on predictive modeling of liquid chromatography retention time, which was

used to optimize solvent gradients in the purification process to greatly reduce purification time. Finally, the author developed a machine vision system for the robotic platform for in-process inspection of the product and real-time fault detection.

1.2. Background on peptides and peptide applications

Bioactive and naturally occurring peptides are pervasive throughout all life [1] and play a central role in several biological processes [2]. Peptides have recently become of interest in the pharmaceutical community due to their strong potential as novel therapeutic treatments. Defined as polymers consisting of links of amino acid monomers, peptides have shown promising results in treating a multitude of ailments, such as type 2 diabetes, Acromegaly, prostate and breast cancer, HIV, and more [3]. Over 7,000 naturally occurring peptides have been discovered, with 60 currently approved for clinical use and 500 derivatives in research and developmental phases [3]. Peptides have a high degree of customizability, making them a versatile and tunable therapeutic agent. Benefits of peptide-based therapeutics include their high potency and target selectivity, broad range of potential biological targets, low tissue accumulation, diversity, and predictable metabolism. These characteristics make them attractive to researchers and clinicians. There are a few disadvantages, such as poor metabolic stability and membrane permeability, high production costs, and low biostability [3]. The goal of peptide research today is to overcome their disadvantages and bring peptides into wide clinical usage, which has resulted in an acceleration of investment into peptide research.

Several categories of therapeutic peptides are on the market and in research today. These include antimicrobial peptides, cancer therapeutics, self-assembling, and treatment-specific research in areas such as type 2 diabetes, hypertension, and anticoagulation [3]. Antimicrobial peptides are hypothesized to have a lower risk of generating resistance, due to the several demonstrated modes

of combating pathogens, both intracellular and extracellular [3]. Cancer therapeutics is one of the most promising fields within peptide research. Some short peptides, known as tumor-specific antigens (TAAs), help the immune system target tumor cells exhibiting those antigens. Certain peptides have recently been introduced for therapeutic cancer vaccines, which are used in conjunction with traditional cancer treatments [3].

1.3. Background on Peptide Synthesis

Natural peptides can be isolated using various methods, but only chemical peptide synthesis methods can create synthetic peptides. Chemical peptide synthesis is also capable of mass mass-producing natural and synthetic peptides [2]. The general modern method of manufacturing peptides and small proteins uses an insoluble polymeric support resin as a base on which the peptide is built. The primary difference between these two types of molecules are peptides range from 2-100 amino acids in length while proteins begin at a length of 100 amino acids. A protector for the α-amino group of the amino acid is attached to avoid any unwanted side reactions, which is removed before the coupling of the next amino acid. The process is repeated until the desired sequence is obtained. This became known as the solid-phase peptide synthesis (SPPS) method.

The discovery of SPPS in 1963 significantly decreased production time by moving an otherwise iterated isolation and purification step performed after each amino acid addition to a single process at the end of production [4]. Only a wash step to remove the reagents and byproducts is needed between the addition of amino acids. Since the process of adding amino acids and removing the protective group is a stepwise procedure, incomplete coupling and deprotection steps can lead to high quantities of unwanted side products. To avoid this issue, which becomes more prominent as the amino chain gets longer, a new process called chemical ligation was discovered in 1992 to make the synthesis of proteins possible [2]. More recently, microwave synthesis has become

prevalent in the peptide synthesis industry since its targeted and rapid heating allows for increased reaction rates up to 1000-fold [5].

SPPS was traditionally carried out in batches, but the introduction of continuous flow technology allows for the production of peptides with any sequence using stock amino acids [6]. It is more efficient than batch methodologies and allows for standard operating procedures to be adapted to any peptide. The use of flow chemistry opened the doors to the automation of the peptide synthesis field and is a core technology used in the automated platform at Mytide.

Several post-synthesis steps are completed to produce the final product. After the amino acids have been coupled, the peptides produced are cleaved off of the support resin with a cleavage cocktail. The resin is collected, and the cleavage cocktail is vaporized with forced air. Once the peptide has been separated from the resin, the product is then reconstituted in solvent. A first step quality check is completed with liquid chromatography and mass spectroscopy to determine the amount of the pure product, byproducts, and secondary structure of the produced peptides. The purity of the product is estimated using an integration of the liquid chromatography peak that represents the correct product, as identified by mass spectroscopy, over the integration of the entire output. The product is then isolated from side products by chromatographic purification, utilizing their differences in hydrophobicity.

1.4. Motivation for Peptide Synthesis Automation

Laboratory scale synthesis of peptides is straightforward, as explained in the previous section. However, due to the sequential building of the peptide, the percentage yield can drop significantly as the peptide chain increases in length. This is especially prevalent in peptide chains longer than 50 amino acids, which presents significant challenges when scaling manufacturing processes to

longer chains and larger batches [3]. The ability to synthesize organic compounds, such as peptides automatically and on-demand, can make this process faster, more efficient, and produce higher purity peptides that are currently unobtainable through traditional manual synthesis [6]. As peptide and protein research moves forward, it will become increasingly important to design and deliver new variants quickly to iterate and optimize the designs for potency, stability, and selectivity [7]. Automation of pharmaceutical manufacturing processes leads to improved quality and yield of the product through techniques such as process control, sequencing, data analysis [8]. By introducing these techniques into peptide synthesis, customizable peptide chains will become readily available at a higher purity and shorter yield time than the current industry standard, further enabling research progress and iteration of peptide variants.

A fully-automated platform for dedicated peptide synthesis was developed in 2017 [7] that uses a flow-based approach to solid-phase polypeptide synthesis and can form an amide bond in 7 seconds, with total synthesis times of 40 seconds per amino acid residue. It is integrated into a robotic platform that can execute all the necessary steps in production without intervention. The peptides produced using this system were comparable to, if not better than, peptides produced through other methods. This work served as the basis and minimum viable product for the Mytide manufacturing platform, which has seen continuous development since. In 2019, a technique combining Artificial Intelligence (AI) and robotics was used to evaluate the likely success of desired reactions using a library of data based on previously published reactions. Using this library, the likely success of desired reactions was evaluated to determine the optimal synthesis parameters [6].

1.5. Background on Machine Vision

In 1963, Larry Roberts published "Machine Perception of Three-Dimensional Solids," where he discussed methods for extracting information about 3D objects from a 2D image [9]. This paper was the beginning of a field known today as computer vision. Computer vision, in an engineering sense, was developed to build autonomous systems that could perform tasks that humans often perform visually [10]. Today, computer vision is used in applications such as facial recognition, autonomous vehicles, classifying medical conditions, identifying healthy crops in a field, and manufacturing [11]. In industrial applications, a subfield of computer vision, known as machine vision, is applied to manufacturing processes in which computer vision and image analysis are combined with automation to inspect parts and ensure that a process runs smoothly [12]. A key component to Industry 4.0, the vision system uses software and imaging hardware to identify certain pre-programmed conditions to trigger a set of tasks based on the image [12]. It is often used to aid in manual assembly, vision-guide robots, inspect parts, and extract information about processes. It can also be used in conjunction with statistical process control to check measurements and analyze trends [13].

1.6. Background on Mytide Therapeutics

Mytide Therapeutics is a biotech startup based in Boston, Massachusetts, that is addressing the peptide market with an automated, proprietary peptide solid-phase slug flow (SPSF) platform. Founded in 2018, Mytide has been using this technology, alongside predictive analytics and machine learning, to manufacture customized peptide chains that can be delivered to customers faster, and at a higher purity. Utilizing this technology for research, development, and ultimately large-scale manufacture, Mytide's goal is to accelerate the discovery, manufacture, and delivery of

peptide therapeutics, ultimately making customized treatment options more accessible and economically feasible. Customers can currently submit a custom chain of up to 40 amino acids to be manufactured and delivered in 5 days with a purity greater than 90%. Mytide operates as both a business to business (B2B) and a business to consumer (B2C) company, providing peptides to research labs, institutions, and pharmaceutical companies [14].

1.7. Thesis Outline

Chapter 2 will give an overview of the manufacturing process at Mytide, and Chapter 3 will outline the problem statement and project goals to be addressed throughout the thesis. Chapter 4 will review image processing and machine history and methods that informed the subsequent methodologies for improvements. Chapter 5 covers the equipment used, methods of image analysis, and their implementation into the Mytide system. Chapter 6 will show the results of the implementation of the machine vision system and the value added to the synthesis platform and manufacturing process as a whole. Finally, chapter 7 will conclude the thesis and make recommendations for future work.

Chapter 2: Mytide Process Overview

At Mytide, the process for synthesizing peptides has been automated utilizing their robotic peptide synthesis platform, “Rapid Automated Computation, Coupling, Cleavage, and Chromatography Execution” (RAC4E). Utilizing artificial intelligence and robotics, this platform is able to rapidly produce peptide molecules.

Figure 1: *The robotic peptide manufacturing platform at Mytide*

The platform utilizes reactor vials to contain the resin, peptide, and other chemicals used throughout the process. To prepare a reactor vial for synthesis, a frit is pressed to the bottom of the vial. The frit is a porous plastic wafer that acts as a filter to allow fluids to flow past it but retains the resin inside of the reactor. The resin is then loaded into the vial on top of the frit, and then the vial queued up on the platform.

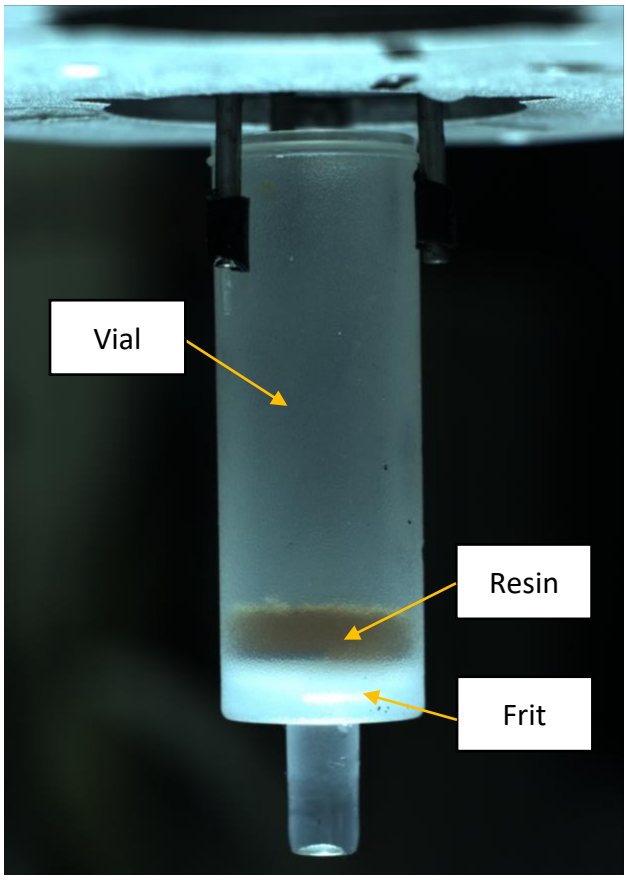

Figure 2: *Reactor vial that has been prepped for synthesis by placing resin in the vial on top of the frit.*

When a peptide's synthesis is scheduled to begin, the vial is picked up by the robot and moved to the synthesizer. The vial is secured on the synthesizer's radial seal, which contacts the inside diameter of the vial, then filled with Dimethylformamide (DMF), which serves as a solvent for the coupling reactions. The vial is then lowered into a reactor where all the necessary amino acids are coupled to form the desired peptide. Once synthesis has been completed, the vial is lifted from the reactor, then removed from the seal by the robot. The vial is transferred to the drier, where the bulk of the liquid in the reactor is removed. Once the peptide has been dried, the robot transfers the reactor vial to the cleaver.

The vial is secured to the cleaver seal, then lowered into another reactor, where the peptide is cleaved from the support resin with a cleavage cocktail. Once cleavage has been completed, the

vial is lifted out of the reactor and removed from the cleaver seal by the robot. The product is then isolated from the cleavage cocktail, reconstituted in liquid, and then isolated from other side products. This is done by a chromatographic purification in which the physical separation of the desired product is achieved through exploiting the varying hydrophobicity of different products. The purified samples are then combined, and the product is packaged for delivery to the customer.

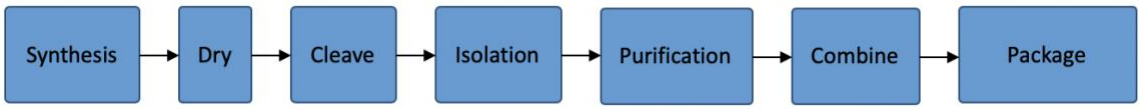

Figure 3: *Process flow diagram of the Mytide peptide manufacturing system.*

Chapter 3: Problem Statement

This thesis investigates machine vision techniques to aid in the automated manufacture of peptide-based therapeutics. The robot on the platform is currently running "blind," meaning that positioning and movement is controlled without sensor feedback. This leaves the system susceptible to obstructions, crashes, and misplacements. The implementation of a machine vision system will enable the robot to "see" the tasks as they are performed, and allows for checks along the way to ensure that all parts are present and in their correct positions. It will also allow the system to monitor the peptide as it moves throughout the manufacturing process by ensuring the vials are gripped, placed, and transferred properly between process steps.

3.1. Project Outline

The primary goal for this system is to detect specific and well-known fault conditions within the system in order to trigger the recovery steps necessary to overcome the identified faults. Throughout the development and use of the automated synthesis platform, several consistent fault conditions have been identified that decrease the platform's overall mean time to failure (MTTF). The mean time to repair (MTTR) can vary wildly, depending mainly on when an operator is available to repair. If a crash occurs during an overnight run, then the system will lose several hours of production time before it can be repaired. By integrating the platform with a machine vision system, these fault conditions can be detected early in the process and mitigated before any damage is done to the system or peptide and prevent machine downtime. Additionally, the machine vision system will be used to gather data relevant to peptide synthesis that could lead to future process improvement.

1.1.2. Fault Condition Detection

The first condition that will be monitored is the presence or absence of a vial at different stages of the manufacturing process. As a peptide progresses through the manufacturing process, its vial is moved from station to station by the robot. Occasionally, the robot will fail to grip a vial and unknowingly leave it behind in a tray. At the synthesizer, it will sometimes fail to lift the vial out of the synthesizer reactor, leaving it behind. Failing to move a vial as expected halts that peptide's progress and impedes the progress of the next peptide. When the robot attempts to place the next vial in a machine with the previous vial still present, the two vials will collide, which causes a crash and halts production. By placing cameras at critical locations where this fault is likely to occur, the system will be able to "see" when a vial is present in the machine or robot gripper or and when it is clear and ready for the placement of the next vial.

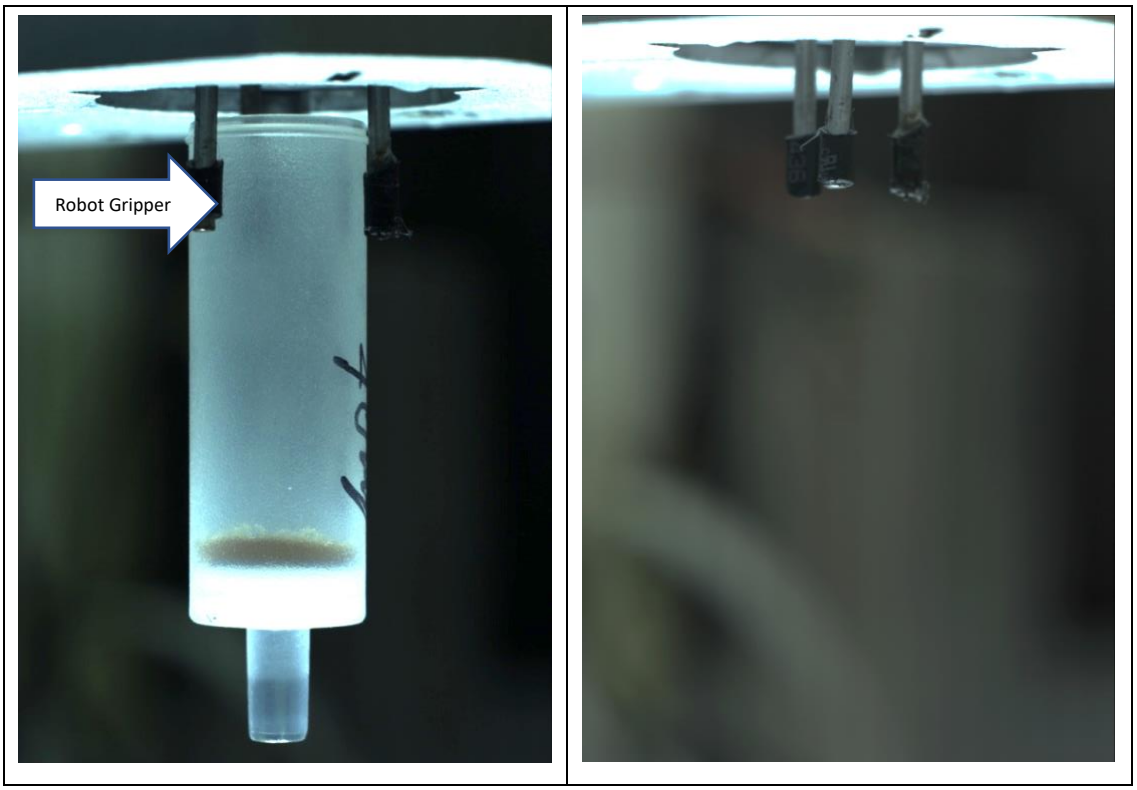

Figure 4: *Example of images taken by the central inspection camera to check for vial presence*

Left: Image of the vial when successfully gripped by the robot
Right: Image of the empty robot gripper after failing to pick up a vial

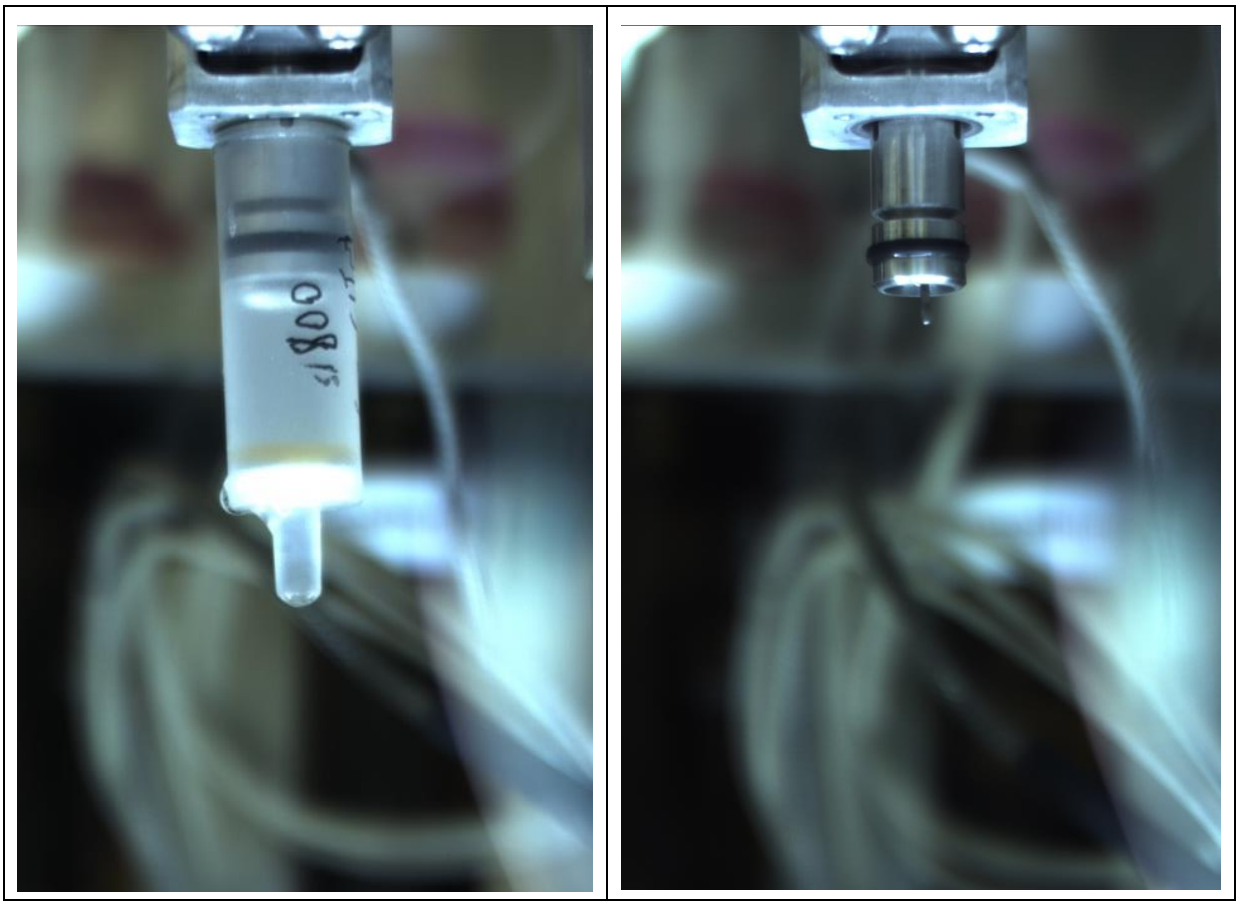

Figure 5: *Example of images taken by the synthesizer camera to check for vial presence*

Left: Example image of the synthesizer when the vial has been successfully placed by the robot or lifted out of the reactor
Right: Example of an empty synthesizer seal after failing to either place a vial correctly or lift a vial out of the reactor.

As the vial is transferred by the robot from station to station, there is the possibility for the vial to be gripped by the robot or by one of the machines slightly out of place. If the vial is gripped at the wrong location, it could cause the robot or machine to either crash the vial into its next placement or fail to engage the vial with the gripping mechanism properly, causing it to fall. By identifying the location of the vial relative to known landmarks and specific regions where it is expected to be, the system will be able to account for the improper placement and take measures to correct it.

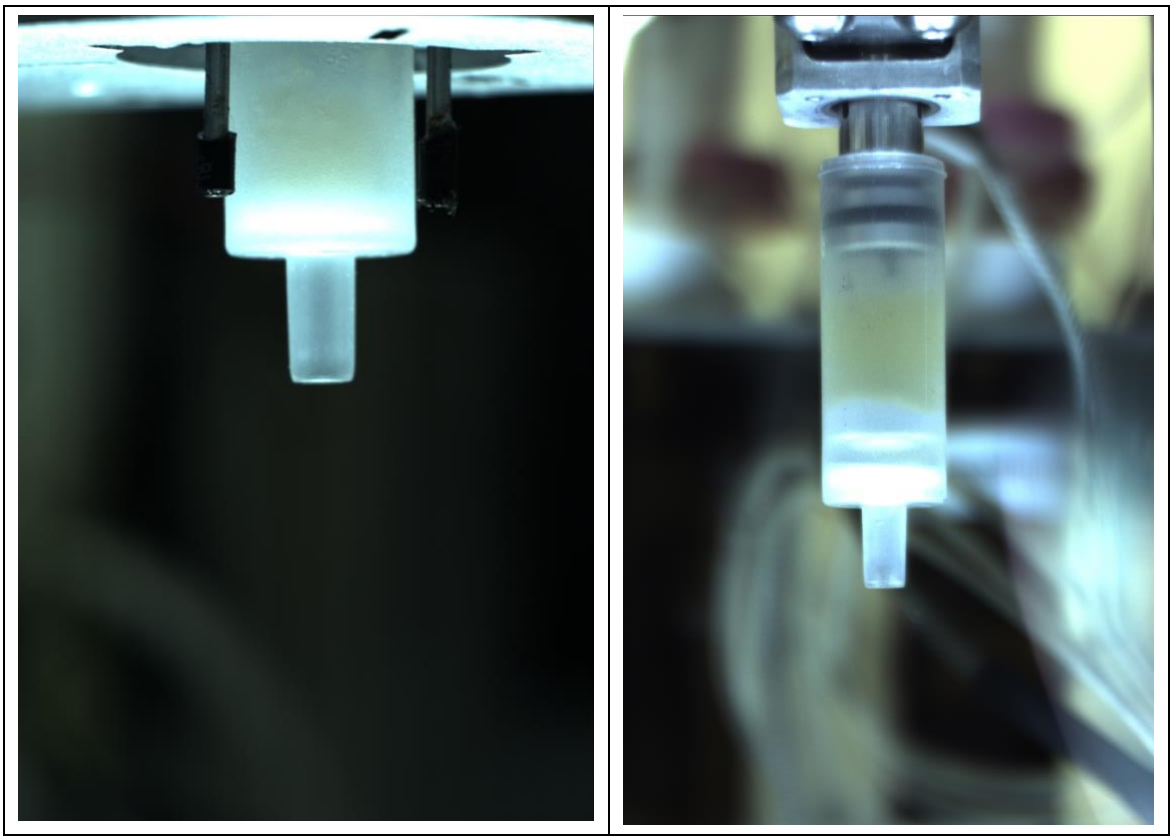

Figure 6: *Example of images taken to check that the vial has been gripped properly*

Left: Image of the vial taken by the central inspection camera when gripped incorrectly by the robot gripper.
Right: Image of the vial taken by the synthesizer camera when gripped incorrectly by the synthesizer seal.

Additionally, the robot may pick up the vial in a way that results in the vial being tilted within the robot grip, causing the tip of the vial no longer to be in line with the top. As the placement for these vials are machined trays with arrays of holes slightly larger than the vial outer diameter, there is potential for failed placement of the vial into said machine or tray. By measuring the degree of tilt of the vial, the robot will be able to make any adjustments needed to account for the resulting displacement of the tip when placing the vial in its next location.

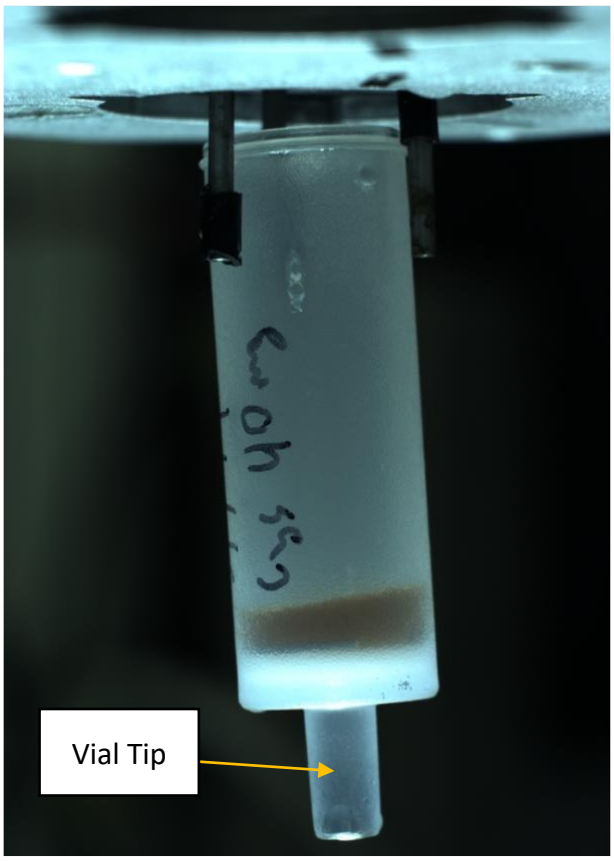

Figure 7: *Example of a vial that is gripped by the robot at an angle*
This causes the vial to be tilted and the tip to be unaligned with the center of the vial.

Another fault condition that the machine vision system will monitor is the dislodgement of the frit in the vial during synthesis. The frit is a small disk that is placed at the bottom of the vial that prevents the resin from draining out the open tip, while still allowing for the flow of solvent and reagent. If the frit becomes dislodged during synthesis, it partially uncovers an open channel in the tip of the vial, potentially making the drying and cleave steps dangerous. It could also allow for resin beads to be sucked out of the vial while being dried, or allow resin to leak out of the reactor vial and into the collection of the cleaver.

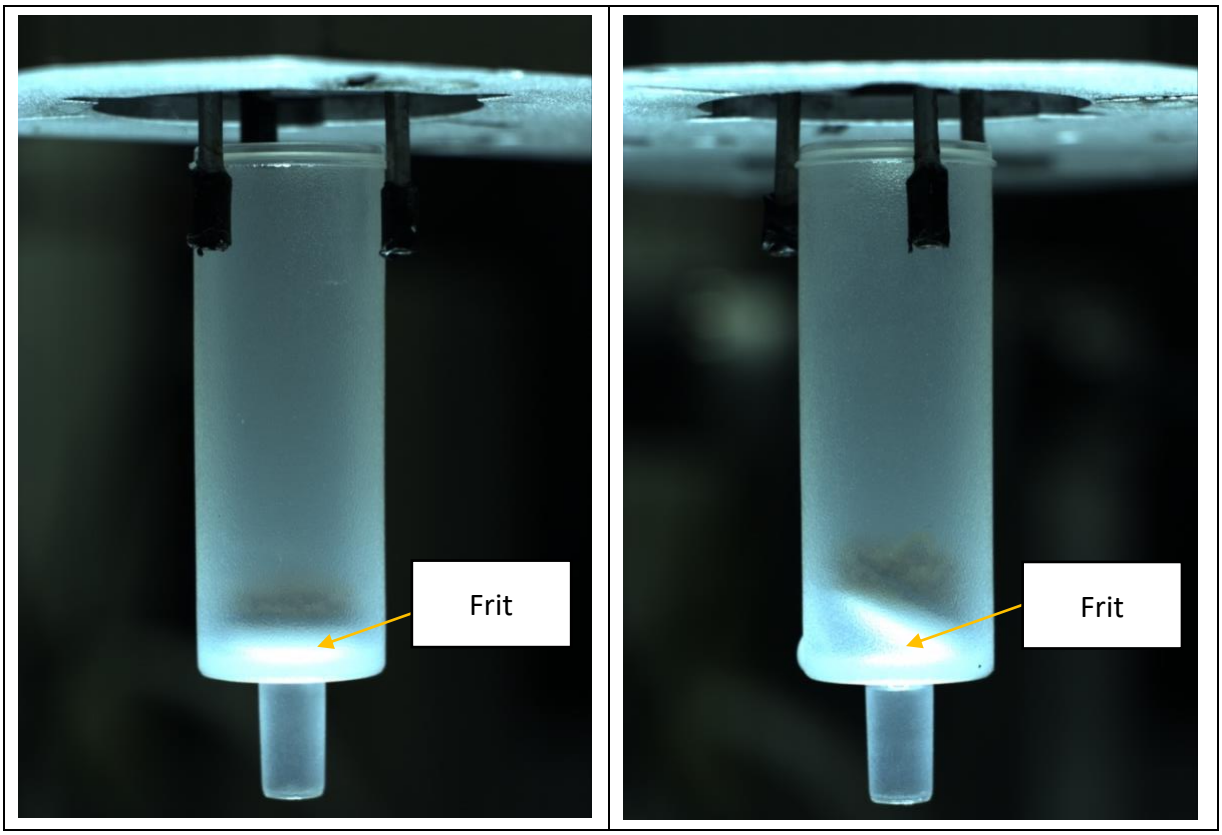

Figure 8: *Example of images taken by the central inspection camera to check for frit tilt*

Left: Example of a vial post-synthesis where the frit has remained in place.
Right: Example of a vial post-synthesis where the frit has been dislodged, resulting in the frit being tilted

Finally, the cameras will be used to determine whether the vial has been adequately removed from the radial seal at the synthesizer. To remove a vial from either machine, it is gripped by the robot, released from the seal, then lowered by the robot before retreating from the machine and moving to the next station. If something goes wrong during removal, such as the robot slipping down the vial, the top of the vial may not be low enough to clear the seal as it is being pulled from the machine. This would cause the vial to crash into the seal, causing it to dislodge the vial or critical components of the machine. By checking the height of the vial relative to the seal before it is moved, the system can avoid such crashes.

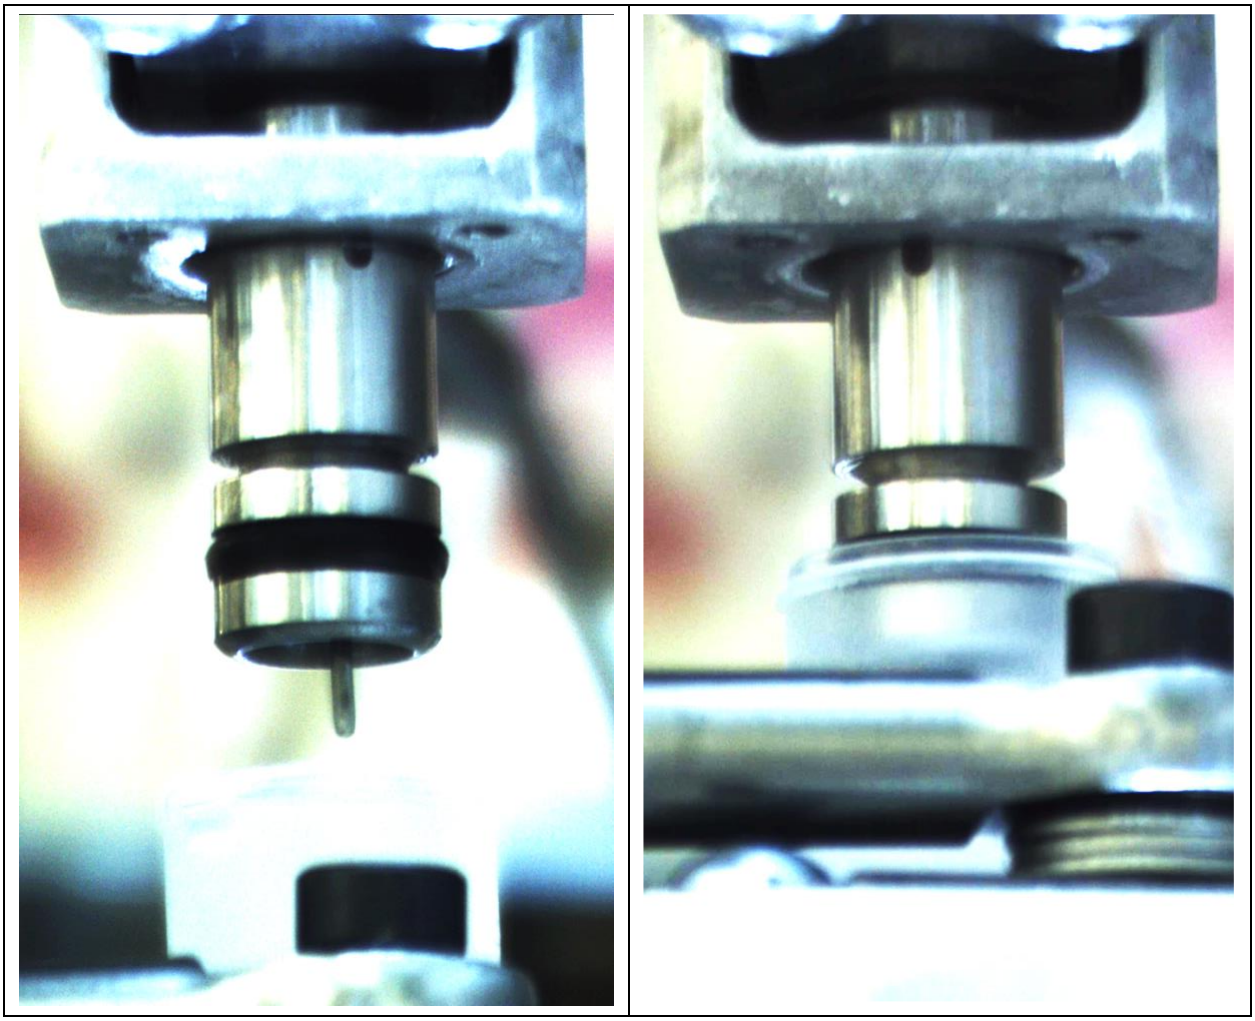

Figure 9: *Examples of images taken by the synthesizer camera during vial removal*

Left: Example of a vial being removed that has cleared the synthesizer seal.
Right: Example of a vial being removed that has not cleared the synthesizer seal.

3.1.2. Data Collection

The machine vision system will also be used to collect data relevant to peptide synthesis that will be used to understand the process better and lead to potential future improvements. Primarily, it will collect data pertaining to the resin during peptide synthesis. To prep a vial for synthesis, chemists weigh the appropriate resin that has been prepared for synthesis by undergoing a series of steps involving swelling the resin in Dichloromethane (DCM) and, if necessary, preloading it

with an amino acid. The resin is loaded into a vial that is then placed in a tray on the platform to be queued up for synthesis. During synthesis, the resin acts as a base or scaffolding for the amino acids. As more amino acids are coupled, the resin swells, causing the volume of the resin to increase. Understanding the percent change of volume of the resins during synthesis of specific peptides will provide a greater understanding of the process and potential optimization opportunities.

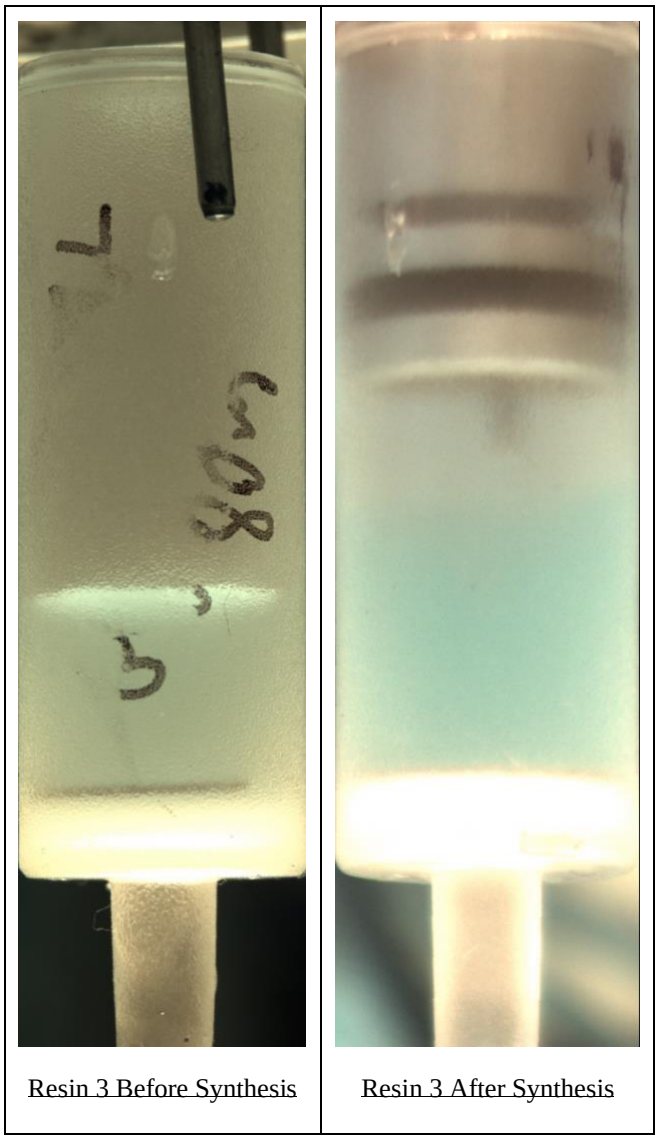

Resin 3 Before Synthesis | Resin 3 After Synthesis

Figure 10: *Example images of Resin 2 taken before and after synthesis*

Left: Example of resin in a vial before synthesis.
Right: Example of resin in a vial after synthesis.

Chapter 4: Literature Review

4.1. Image Processing

Digital image processing can trace its origins back to the 1960s at the Jet Propulsion Laboratory in Pasadena, California when the Ranger 7 space probe transmitted the first images of the lunar surface back to Earth. These images were processed by a computer to correct distortion inherent to the on-board cameras. Many of the lessons learned from Ranger 7 would serve as the basis for the development of many modern imaging process techniques. Further advances in satellite imagery and medical imaging helped propel the processing techniques developed in the 1960s and 70s into a vast and diverse range of applications such as archaeology, defense, autonomous vehicles, meteorology, and manufacturing [15]. As computers became faster, cheaper, and more widely available, image processing began to take on more complex analyses at greater speeds. This led to real-time image analyses and the development of specialized processes. Two such processes, Canny edge detection and image segmentation, were employed during the course of this project.

This thesis utilizes OpenCV, an open-source and cross-platform imaging library originally developed by Intel. OpenCV is a library of functions developed for real-time computer vision applications [16]. Included in this library are functions used perform the image processing techniques that are described in this thesis.

4.1.1. Canny Edge Detection

Canny edge detection is a popular tool used as a pre-processing step in many computer and machine vision algorithms. Developed by John Canny in 1986, this method takes a computational approach to locating the edges in an image by taking into account specific criteria for detection

and localization, and then adding in another criterion to ensure that each edge is only detected once [17]. The algorithm performs several steps, which include smoothing, filtering, suppressing non-maxima, then performs a connected component analysis that detects the edges of objects within an image [18]. OpenCV contains a Canny edge detection method, cv2.Canny() that goes through four steps. The first step is to reduce the amount of noise in the image, since edge detection can be highly susceptible to noise, and thus uses a 5x5 Gaussian filter to filter out as much noise as possible. The second step filters the image using a Sobel kernel to find the gradient intensity of the image. By using the kernel in both the horizontal (G_x) and vertical(G_y) direction, the function can find the gradient direction (θ) of the edge.

$$G = \sqrt{G_x^2 + G_y^2}$$

Equation 1: Edge Gradient

$$\theta = tan^{-1}\left(\frac{G_y}{G_x}\right)$$

Equation 2: Gradient Direction

The third step in Canny Edge Detection scans the image to locate any pixels that do not represent a local maximum and suppresses them (sets their pixel value equal to 0). This results in a binary image with clear, unclouded edges. Finally, the image goes through Hysteresis Thresholding, where the function determines which are edges and which are not, based primarily on their connectivity to each other [19]. The result is a binary image showing the edges of the original images in white pixels, with all other elements suppressed to black pixels.

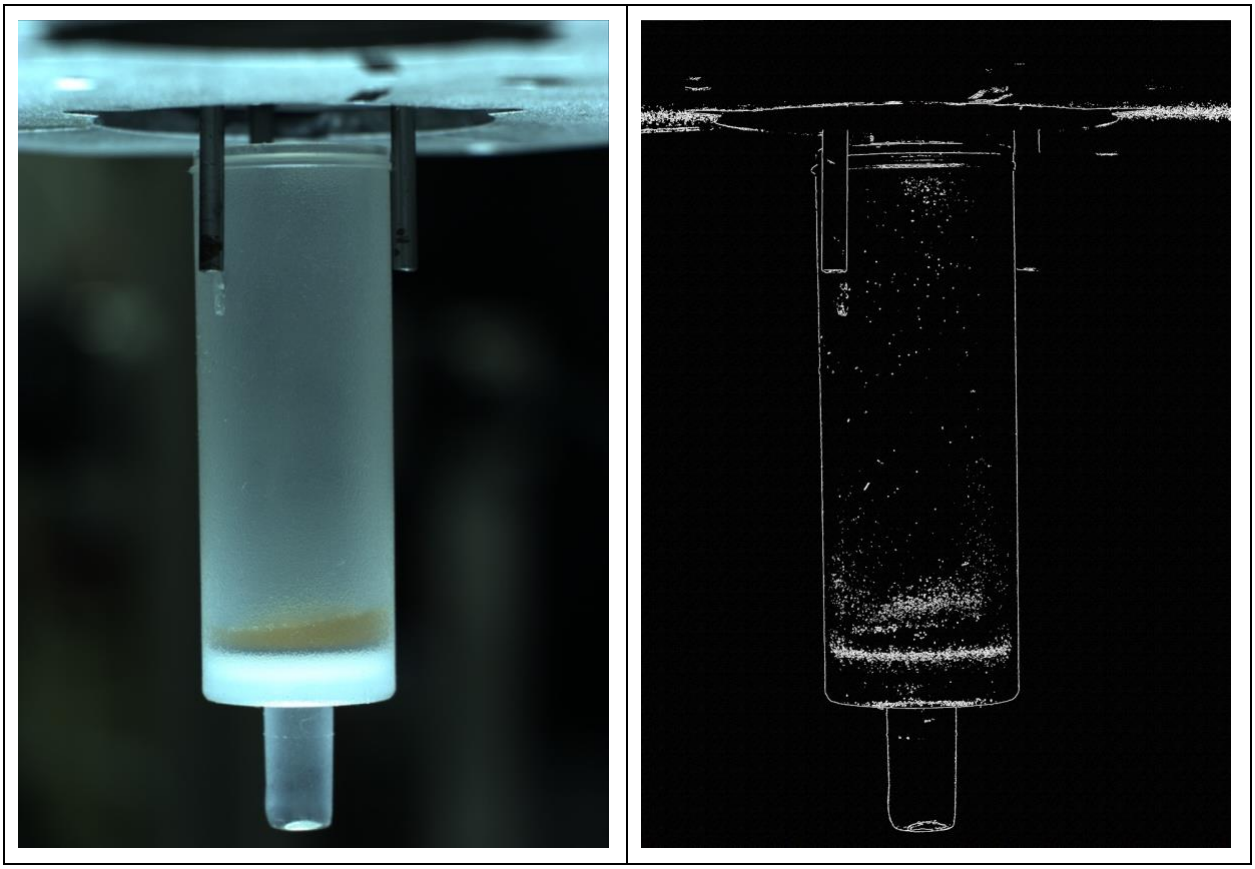

Figure 11: *Example of the Canny Edge Detection algorithm performed on a vial image*

Left: Original image used as the input into the algorithm.
Right: Output image showing edges of the input image.

4.1.2. Image Segmentation

Image Segmentation is the process of segmenting regions of pixels into separate partitions in order to simplify the analysis of the image or make it more meaningful. The pixels in a segmented region are similar to each other with respect to certain characteristics, such as color or intensity [20]. Several techniques of image segmentation have been developed, such as thresholding, clustering, edge detection, and statistical region merging. The simplest version of image segmentation is thresholding. This method works by selecting a threshold value to determine which pixels are set to black and which are set to white, resulting in a binary image. The threshold can be set manually

or determined automatically. The methods for setting thresholds can be divided into six groups, histogram, clustering, entropy, object-attribute, spatial, and local-based methods [21].

OpenCV contains a function called cv2.inRange() that checks an array of elements to determine if they lie within a specific range [22]. In image processing, this function is used as a form of image segmentation to determine which pixels in an image lie within a specified color range and then separate them from the rest of the image. The function takes in an image, as well as lower and upper thresholds (set by the user) in the form of two sets of pixel values in the same color profile as the image. Defining these ranges of color spaces allows the function is able to identify the pixels that lie within the color range and produce a binary image containing only those pixels whose color lies within the specified range. This technique allows specific, uniform regions of an image to be analyzed separately from the rest of the image. This approach works for a wide range of color profiles, including RGB (red, green, blue), BGR (blue, green, red), and HSV (hue, saturation, value) [23].

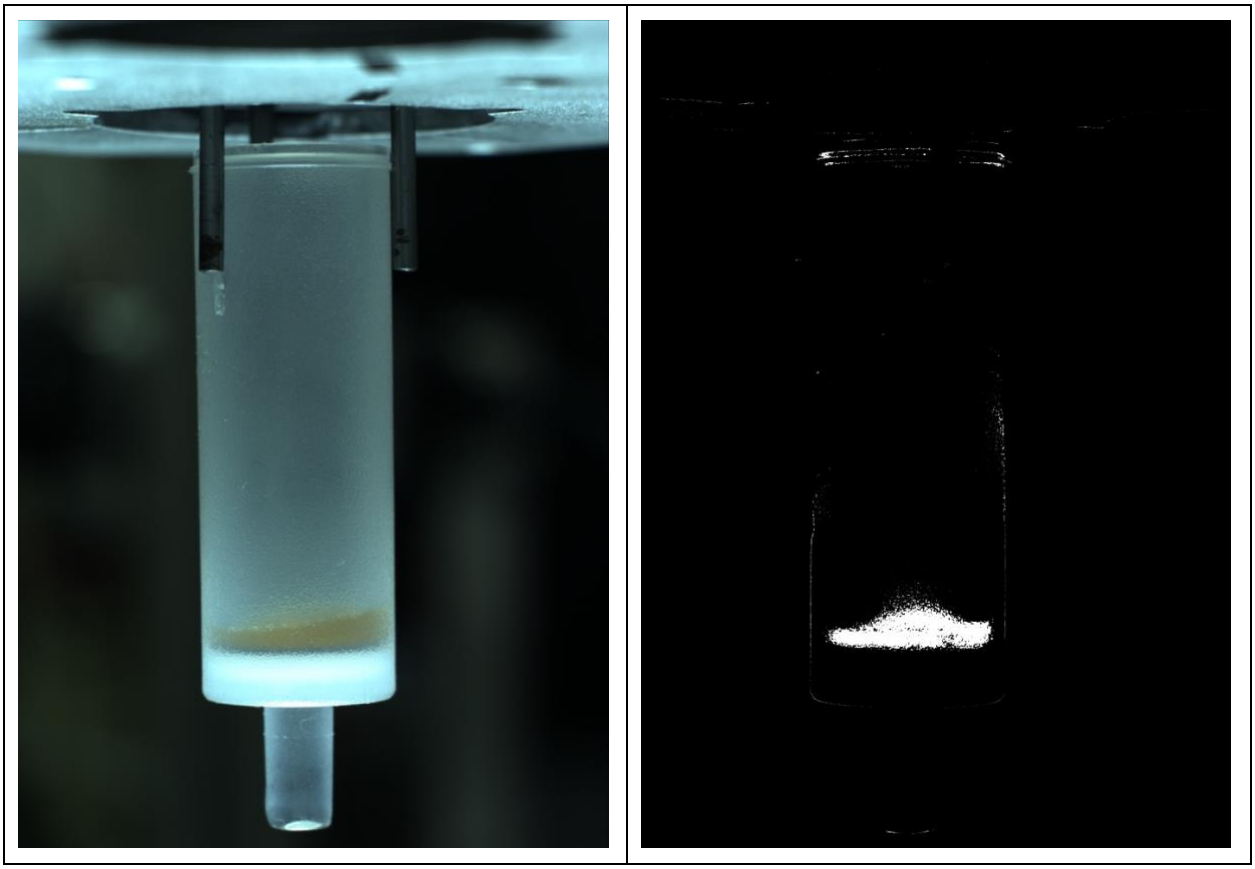

Figure 12: *Example of Image Segmentation used to separate the resin pixels from the rest of the image.*
Left: Original image input into the system in the standard RGB (red, green, blue) color space.
Right: Output image of the image segmentation algorithm, showing the pixels associated with the resin highlighted in white.

4.2. Machine Vision

Advancements in image processing have led to the introduction of several new fields of research and numerous practical applications, such as computer and machine vision. The ultimate goal of computer vision is to emulate human vision and make judgments and decisions in place of a human. As a subdivision of Artificial Intelligence (AI), this includes learning from and making decisions based on the results of the analyzed images [15]. Machine vision is a subfield of computer vision that utilizes strategically placed cameras and image analysis for automated image capture and processing in an industrial environment. The use of machine vision in place of human

inspectors allows for inspection processes to happen faster, continuously, and objectively. Machine vision has proven to be highly effective in quality control of products and helped push the manufacturing industry towards the automation and digitization goals of Industry 4.0.

Typically, in machine vision systems, a camera is placed with a view of an illuminated section of interest in a manufacturing line. The photos or videos taken of that section during the manufacturing process are then promptly processed to obtain useful visual information about the system or product. That information is then used to inform a decision-making system or a scheduling system that will determine subsequent moves in the manufacturing process [24].

As a critical component of Industry 4.0, machine vision allows for the collection and analysis of a high volume of visual data. This data can be used to identify defective products, understand their deficiencies, and enable fast and effective intervention when things go wrong [25]. Applications of machine vision in manufacturing include identifying surface cracks and defects, volume calculations, quality detection, separation of good and bad product, and reading labels [26]. In the pharmaceutical industry, these machine vision applications have been used to determine print regions of drug capsules [27], count and classify pills and tablets, inspect packaging and bottle caps [28], determine liquid levels, and much more.

Chapter 5: Methodology

5.1. Equipment

For this project, two cameras were placed on the peptide synthesis platform, one at the synthesizer, and one at a central inspection station. The cameras mounted on the platform were Basler acA4600-7gc GigE, 14-megapixel cameras. Each camera was equipped with a Megapixel Monofocal Lens with a manual iris from Computar. The lens on the central inspection camera is an M2514-MP2 25mm monofocal lens, while the synthesizer camera uses an M1614-MP2 16mm monofocal lens. Additionally, ring lights were placed below each station with upward illumination.

Each camera was connected to the system via ethernet and controlled through the system's scheduler. The scheduler is responsible for communicating with each machine on the platform in order to coordinate the processing of several peptides. Once the process has reached a state in the process that requires an image to be taken, the scheduler turns on the ring light and takes a picture with the appropriate camera. The image is analyzed, and a copy is saved temporarily in the system's memory. For permanent storage, a copy of the image is also sent to the manufacturing backend server.

5.2. Camera Mounting and Placement

The central inspection and synthesizer cameras were placed strategically in order to obtain a clear and unobstructed view of the desired machine or station while staying out of the robot's path. The analysis performed on the images taken at each location is specific to that camera, due to the differing views, backgrounds, and landmarks in each image. While each camera takes unique

pictures distinct from the other, there is some overlap in the desired information to be acquired from the images, such as vial presence and location.

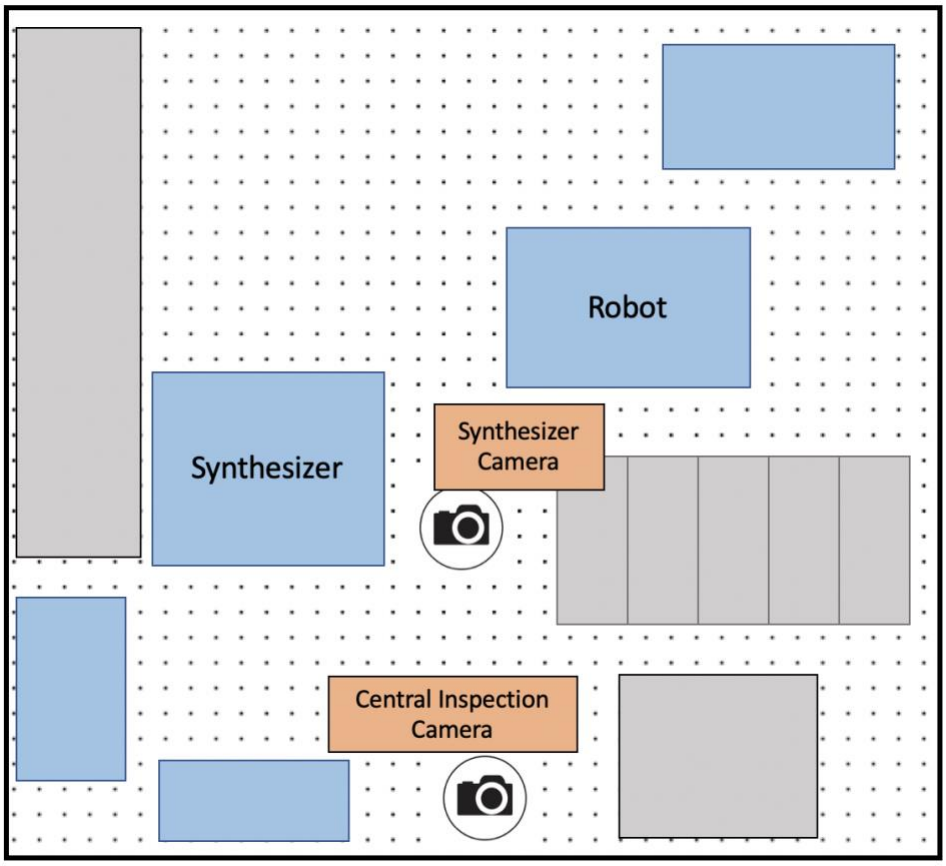

Figure 13: *Diagram showing the placement of each camera within the robot platform*

Images are taken at five points in the process. Before synthesis, images of the vial are taken both at the central inspection camera and after placement of the vial into the synthesizer by the synthesizer camera. Another image of the vial is then taken after synthesis by the synthesizer camera, and again during the removal of the vial from the synthesizer. Finally, the central inspection camera takes a series of images of the vial after it has been dried.

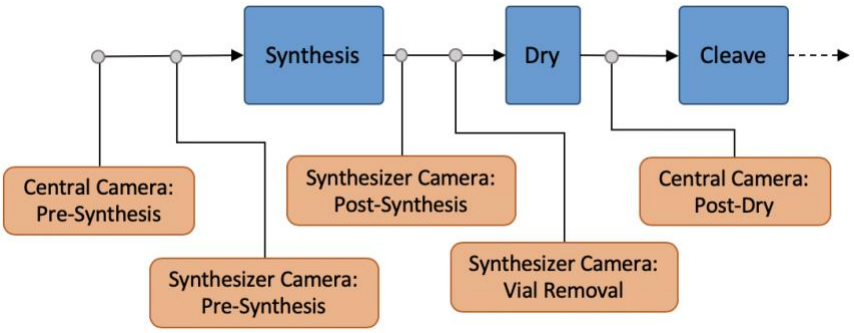

Figure 14: *Flow chart showing the steps in the process where images are taken of the vial. The blue boxes represent the production step, and orange boxes indicate the images that are taken at the step.*

5.2.1. Central Inspection Camera

The central inspection camera is centered at the front of the robot platform and is used as an inspection point for vials as they are being transferred between machines. The robot presents the vial to the camera, at which point a ring light below the designated inspection area is turned on to illuminate the vial and its contents. The camera takes four images as the robot rotates the vial by 60 degrees between each to obtain images of the vial at 90, 150, 210, and 270 degrees. The combination of these four images provides a more accurate analysis of the vial, as it allows for a more complete view of the vial and its contents.

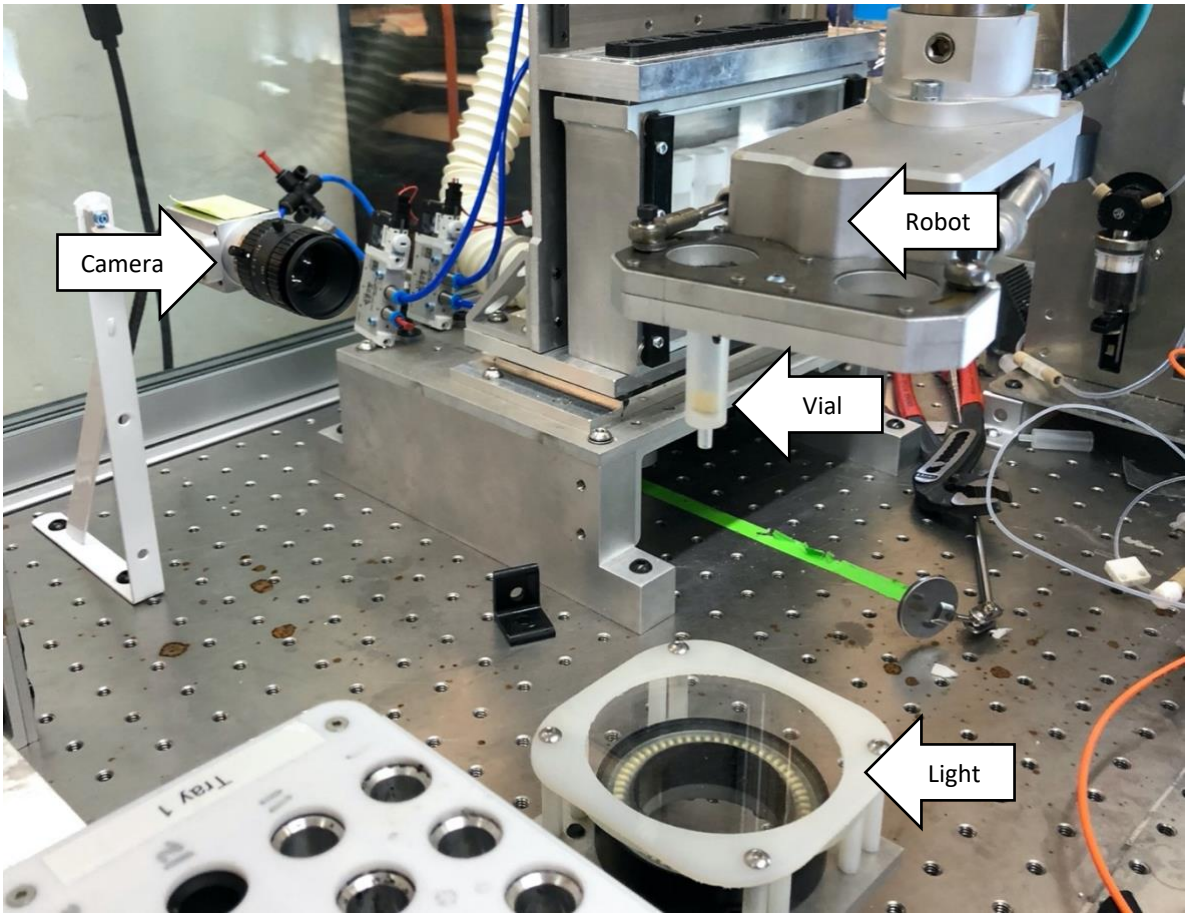

Figure 15: *The central inspection camera mounted at the center front of the platform. A vial is presented to the camera by the robot over the ring light.*

The central inspection camera takes images at two points in the process, before the vial is placed in the synthesizer, and after it is removed from the drier. Both sets of images are used to verify that a vial is present in the robot gripper, determine the height of the vial within the gripper, and measure the tilt of the vial. Additionally, the "pre-synthesis" images are used to estimate the volume of the resin with the vial, while the "post-dry" images are analyzed to determine if the frit has been dislodged during synthesis.

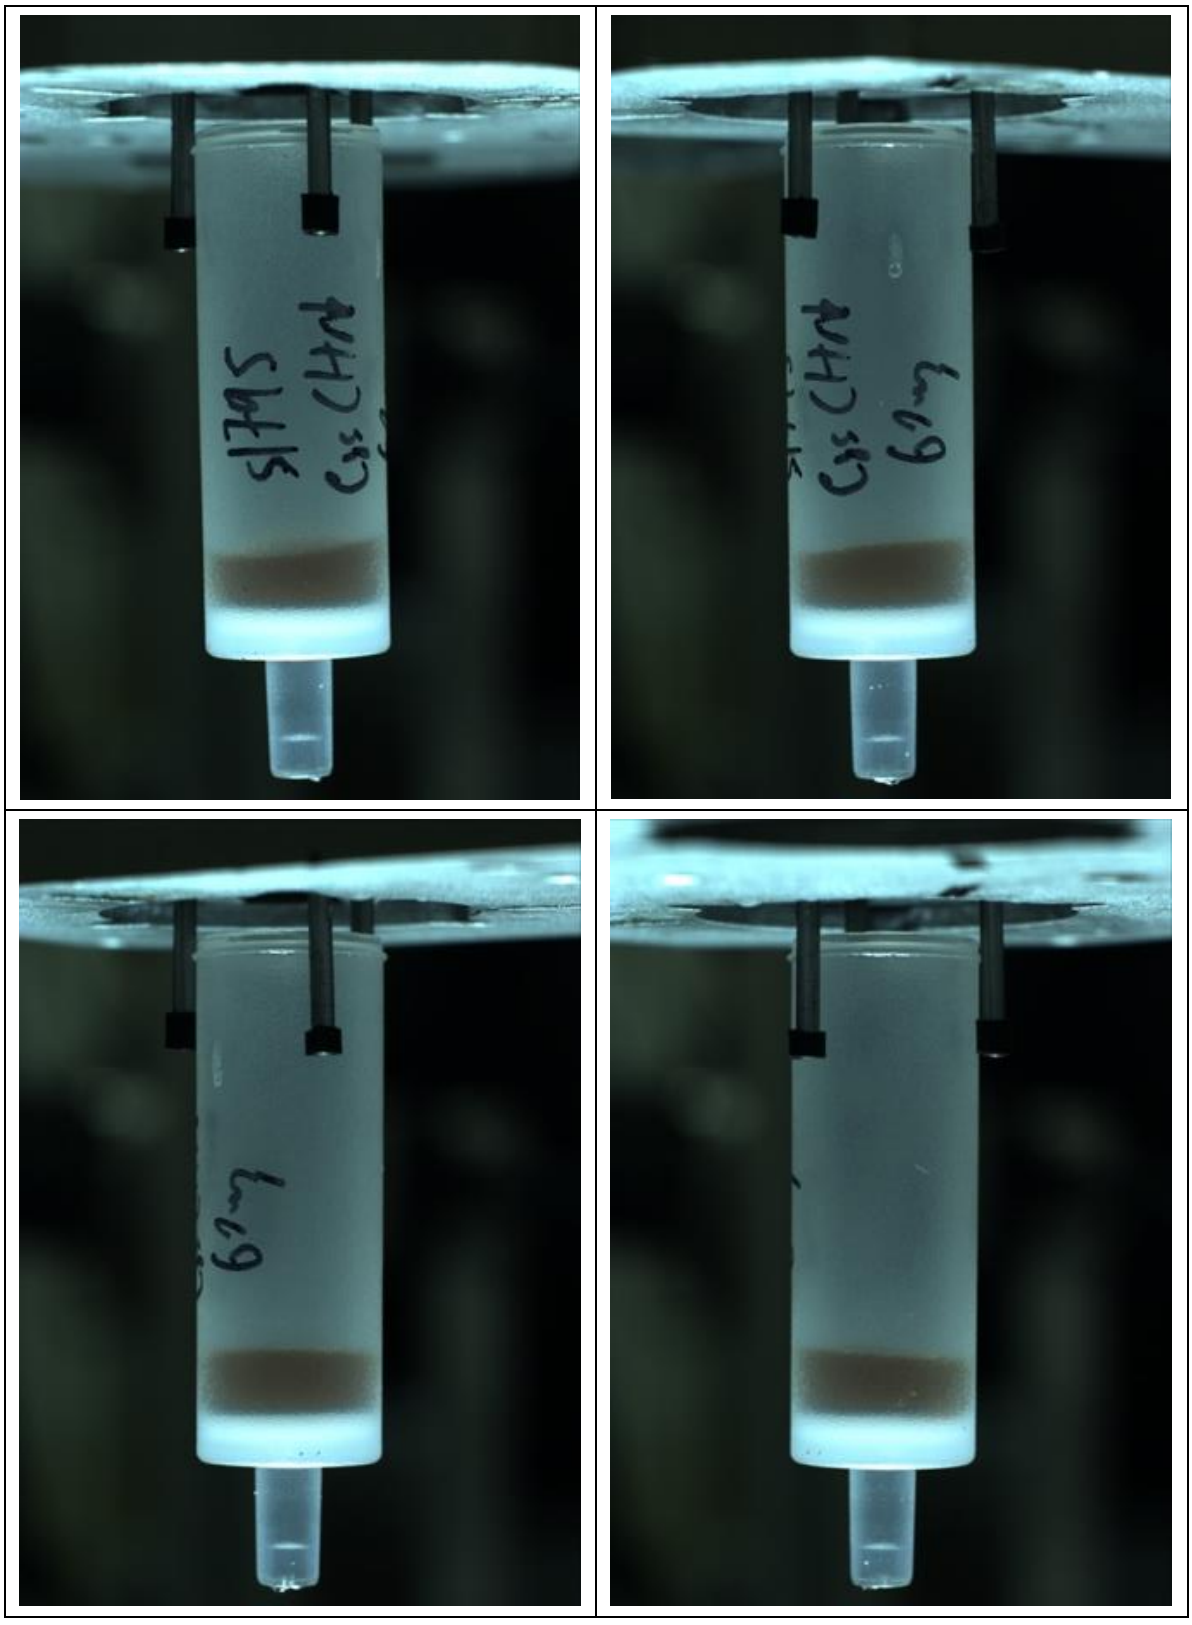

Figure 16: *Example set of four images taken by the central inspection camera before peptide synthesis. Images taken at negative 90 (top left), 150 (top right), 210 (bottom left), and 270 (bottom right) degrees.*

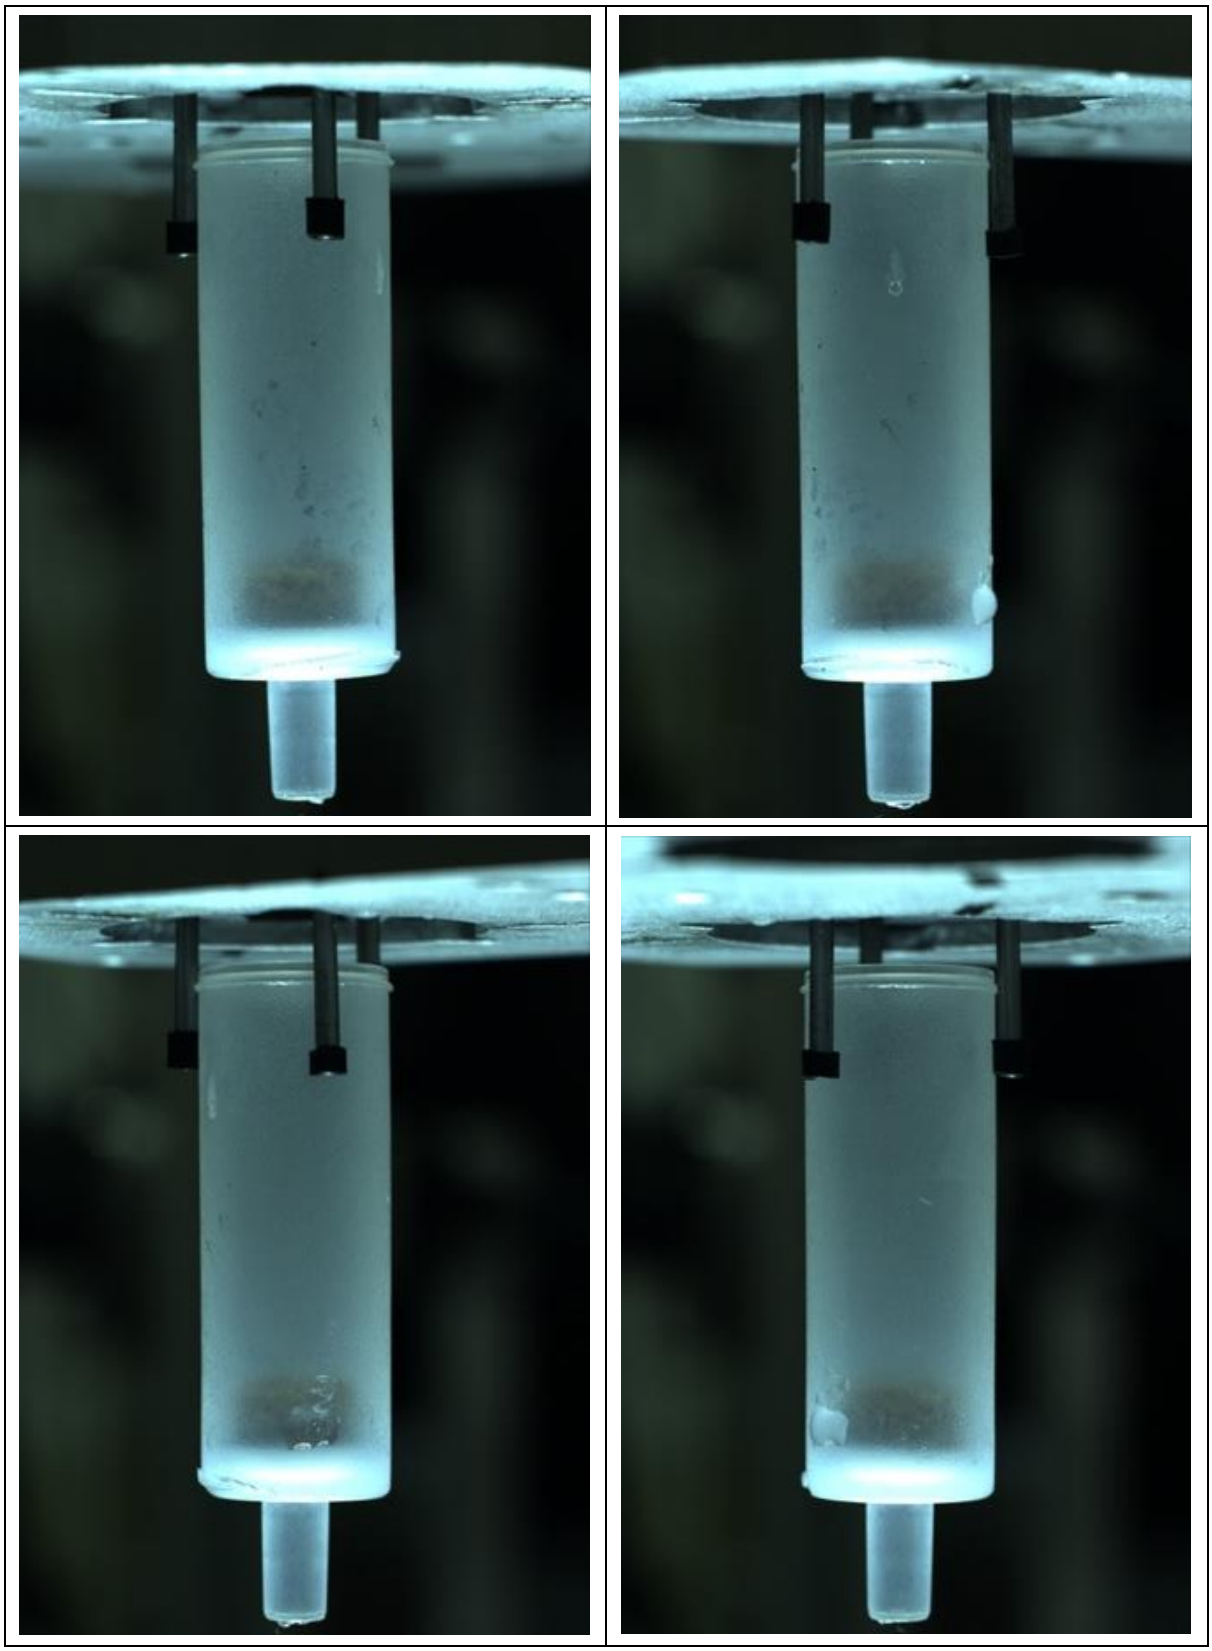

Figure 17: *Example set of four images taken by the central inspection camera before peptide synthesis. Images taken at negative 90 (top left), 150 (top right), 210 (bottom left), and 270 (bottom right) degrees.*

5.2.2. Synthesizer Camera

The synthesizer camera is mounted directly to the right of the synthesizer, towards the center of the platform. It takes three pictures of the synthesizer seal at critical points of transition in the process and is illuminated by a ring light placed below the reactor. These three images are taken after the vial has been placed in the synthesizer seal and filled with DMF, immediately after synthesis, and after it has been removed from the synthesizer seal by the robot.

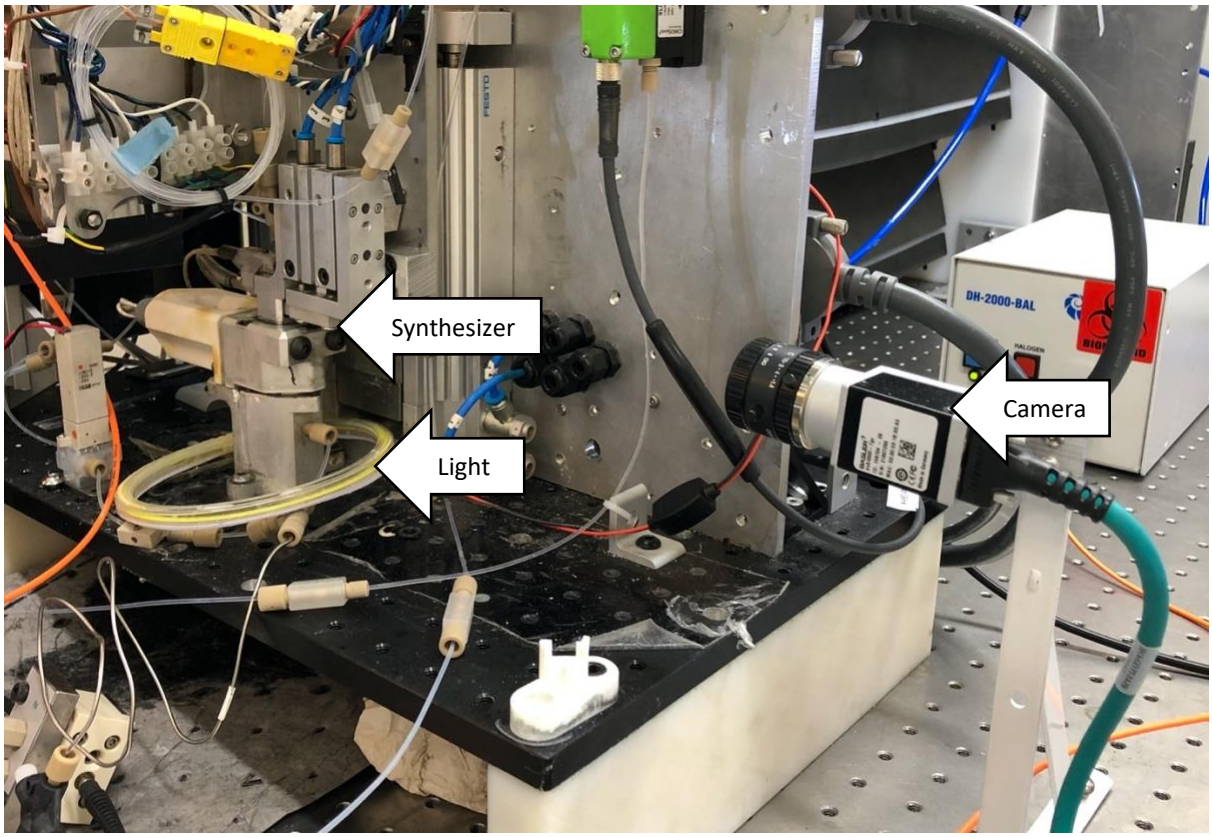

Figure 18: *The synthesizer camera mounted to the right of the synthesizer*

The first and second images taken by the camera are analyzed to determine both the presence of a vial in the gripper as well as the placement of the vial on the gripper. The second image, taken immediately after synthesis, is also analyzed to determine if the frit was dislodged during synthesis and to estimate the volume of the resin in the vial. The third and final image is taken after the robot

has removed the vial from the seal, but before it retreats from the synthesizer. This photo is analyzed to ensure that the vial has been removed completely before it is taken away in order to prevent the vial from crashing into the synthesizer.

Before Synthesis

After Synthesis

Vial Removal

Figure 19: *Example set of the three images taken by synthesizer camera*
Images taken before synthesis (left), after synthesis (center), and during vial removal (right)

5.3. Image Analysis Functions

Each image taken by the cameras is analyzed to extract specific information about the process and peptide. Several functions for performing the analyses were written and applied to each image of interest. All image analysis code was written in Python 3.7.6. Several functions applied to images taken at multiple cameras and required slight adjustments in the code to account for that camera's specific lighting and frame. The results of these functions are outputted and stored in a cloud database as part of the peptide's metadata. When a fault condition is detected, the image analysis code will send a signal to the robot, indicating that recovery measures need to be taken.

5.3.1. Vial Presence Detection

The first task required by the vision system is to determine if there is a vial present in the relevant gripper. Each image taken undergoes a series of image processing steps that are used to verify the presence or absence of a vial. To start, Canny edge detection is performed on the image to isolate prominent edges from the rest of the image. The code then looks for columns in the image with a significant number of white pixels present. The nearly empty background for each image ensures that the only columns identified as having a significant (greater than 60) number of white pixels are there due to the presence of a vial. If there are no such columns present, there is no vial present. If there are, then the first and last columns identified coincide with the two outer edges of the vial. When a vial is present, the function outputs a Boolean value "True," as well as the identified vial edge locations. If no vial is present, the function outputs a Boolean value "False," and a "None" data type. This function serves as the first check performed on every image, as all other functions require a vial to be present to perform their checks.

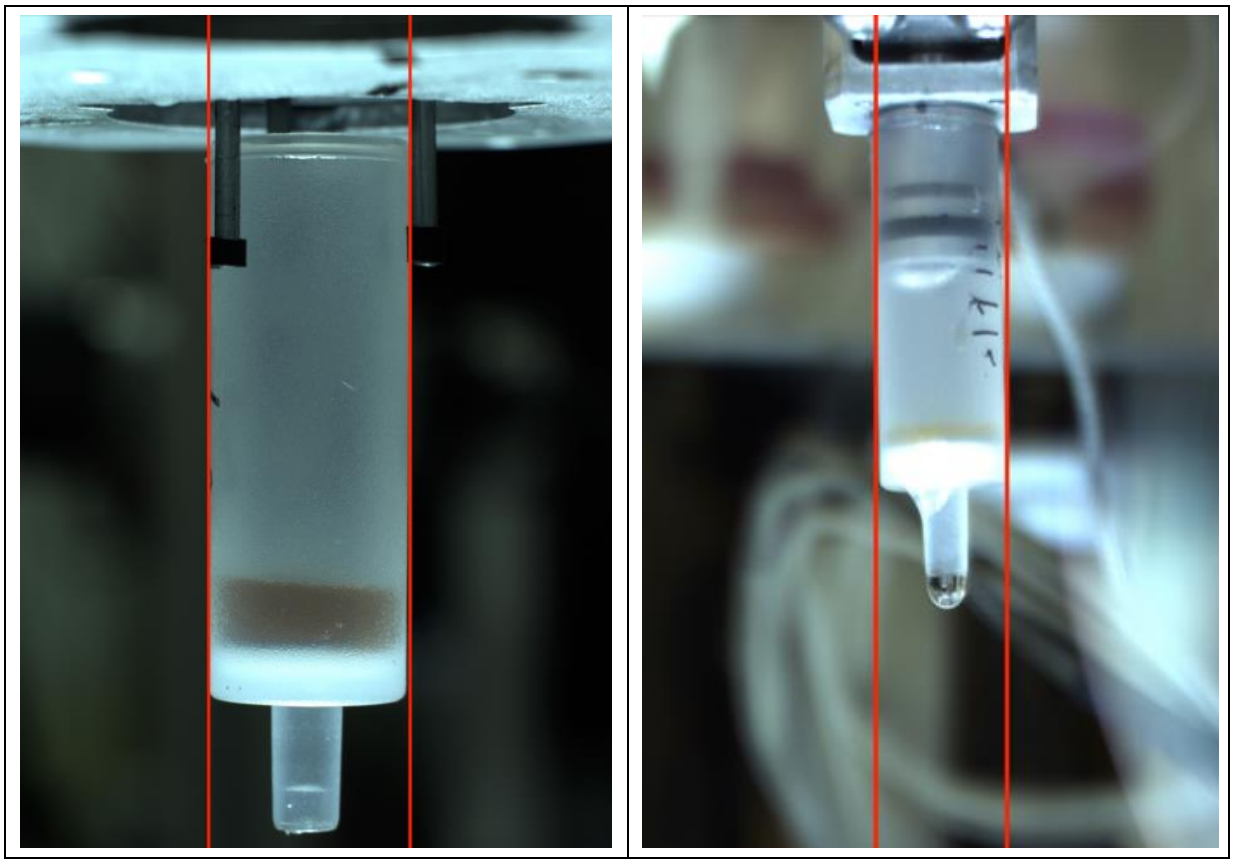

Figure 20: *Example set of images after being processed by the vial presence function*

Left: Image of vial at the central inspection station with two red lines indicating the vial edges found using the vial presence detection function.
Right: Image of a vial in the synthesizer with two red lines indicating the vial edges found using the vial presence detection function.

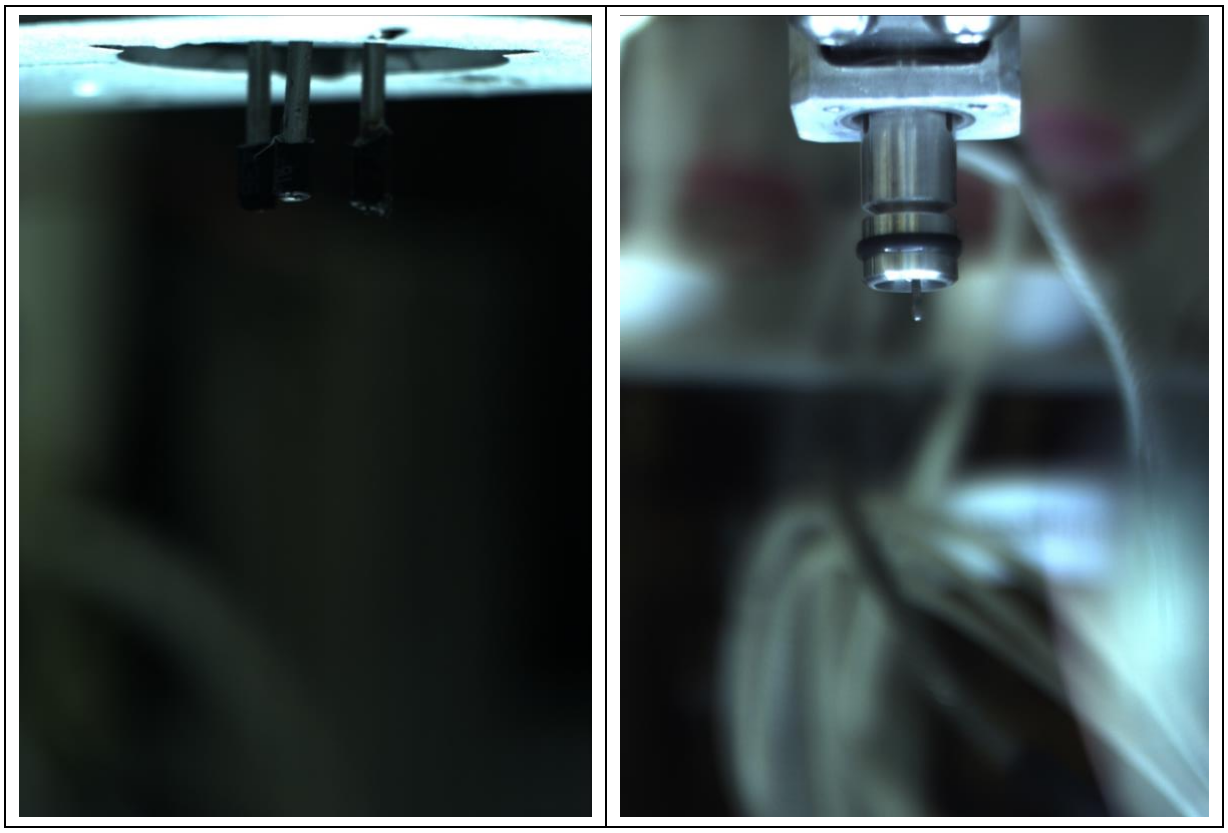

Figure 21: *Example set of images when no vial is present*

Left: Image of a vial presence fail case at the central inspection.
Right: Image of a vial presence fail case at the synthesizer.

5.3.2. Vial Grip Analysis

Another critical piece of information obtained from the images is the location of the tip of the vial relative to the gripper, which will indicate whether or not the vial has been gripped properly. The process of locating the tip of the vial is similar to locating the vial edges, however, instead of looking at the columns of pixels in the image, it looks at the rows. The image is again run through Canny edge detection, after which the program determines all the rows in the image that contain a significant number of white pixels. For images taken by the central inspection camera, the last row containing white pixels is associated with the tip of the vial. Images of the vial at the synthesizer, however, require additional filtering to identify the tip. Before synthesis begins, the vial is filled

with DMF. During synthesis, it is common for DMF to collect on the outside of the vial, and collect at the tip when lifted out of the reactor, making it challenging to determine the actual location of the tip. To filter out the edges created by this drop of DMF, the program calculates the width of the tip by finding the locations of the first and last white pixel in each row. When the distance between these two pixels begins to decrease at a rate greater than 5 pixels per row, the program determines that it has reached the tip of the vial.

Once the vial tip has been located, it is then compared to the region of rows in the image in which the tip is expected to be relative to the gripper. This region is defined using the top of the gripper and known dimensions of the vial. If the distance from the top of the gripper to the tip of the vial is greater than the length of the vial by more than one millimeter, the function outputs a Boolean value of "False" to let the robot know that the vial is not gripped correctly. If it lies within the expected region, the function outputs a Boolean value of "True," indicating that the vial is gripped correctly. The function also outputs the location of the tip.

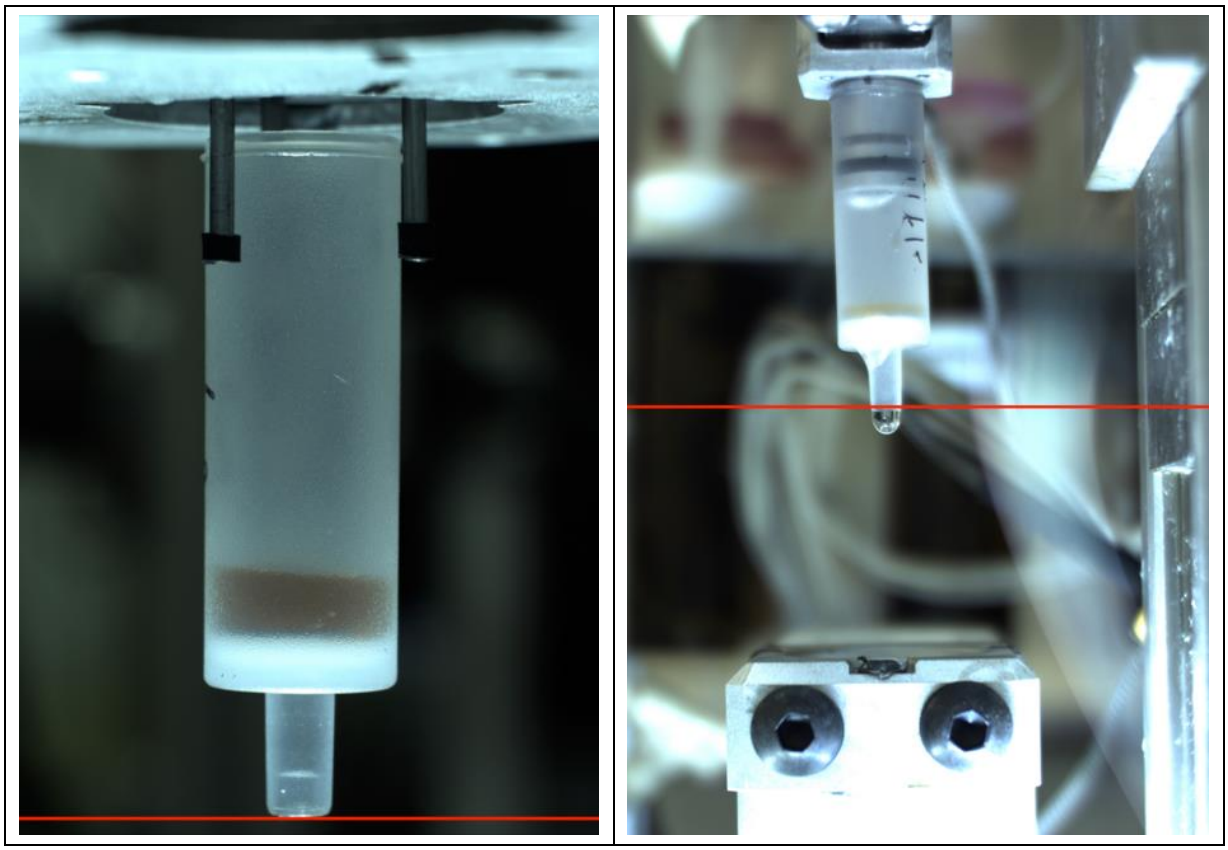

Figure 22: *Example images of a correctly gripped vial after being processed by the vial location function*

Left: Image of the vial at the central inspection station with a red line indicating the tip of the vial found using the tip location function.

Right: Image of the vial in the synthesizer with a red line indicating the tip of the vial found using the tip location function.

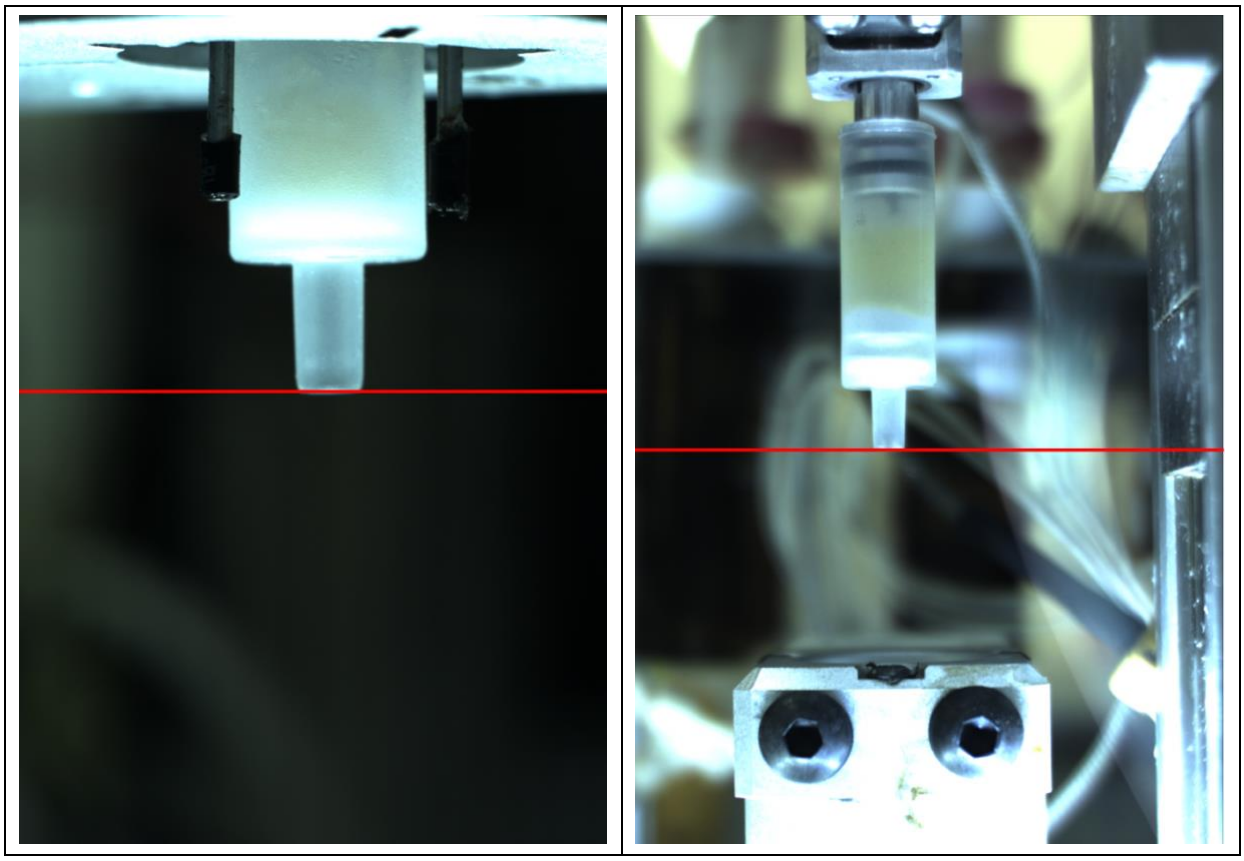

Figure 23: *Example images of an incorrectly gripped vial after being processed by the vial location function*

Left: Image of the vial at the central inspection station with a red line indicating the tip of the vial found using the tip location function.
Right: Image of the vial in the synthesizer with a red line indicating the tip of the vial found using the tip location function.

5.3.3. Vial Tilt Measurement

The third function that is used to determine when the vial has been gripped properly is the vial tilt measurement function. To measure the degree of tilt of the vial within the gripper, the function restricts the region of the image to be analyzed based on the edges and tip location found in the previous two functions. It then focuses specifically around the location of the left vial edge. Next, the program needs to locate points along the edge of the vial with which to measure the tilt. First, it identifies two rows in the constrained region containing pixels associated with the edge of the vial. Once those two rows have been selected, the program finds the first white pixel in that row and records its location.

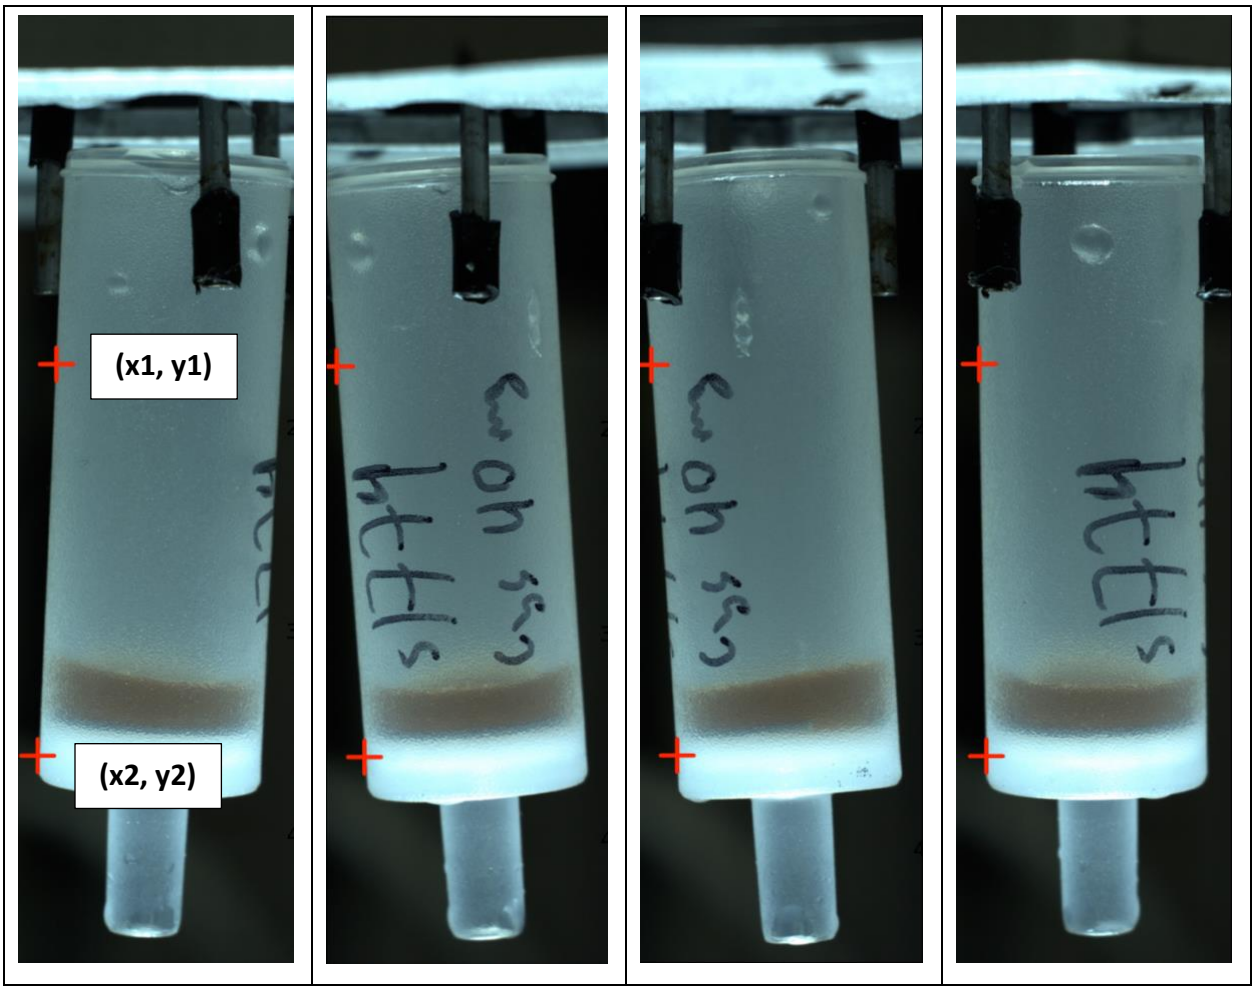

Figure 24: *Example set of images showing the two points on the vial from which the tilt is measured.*

Using these two pixels as reference points, the program measures the tilt of the vial in the gripper. It does this by calculating the inverse tangent of the difference between the pixel rows $(x_2 - x_1)$, divided by the difference between the pixel rows $(y_2 - y_1)$, then converting from radians to degrees.

$$\theta_{vial,\ degrees} = \text{atan}\left(\frac{x_2 - x_1}{y_2 - y_1}\right) * \frac{180}{\pi}$$

Equation 3: Angle of vial tilt in degrees

Once the tilt of the vial is known, the program determines the resulting displacement of the tip from the vertical of the vial by converting the difference between the two previously identified pixel rows to millimeters. The scale for this conversion is determined using the known diameter of the vial as a scale reference. The outputs of the function include the measured tilt of the vial in degrees, the displacement of the tip in millimeters, as well as a Boolean value of "True" if the displacement is greater than a minimum threshold, and "False" if it is below the threshold. The threshold was set to 1.5 mm, as any greater displacement would cause the tip of the vial to be unable to be placed in its next tray or station.

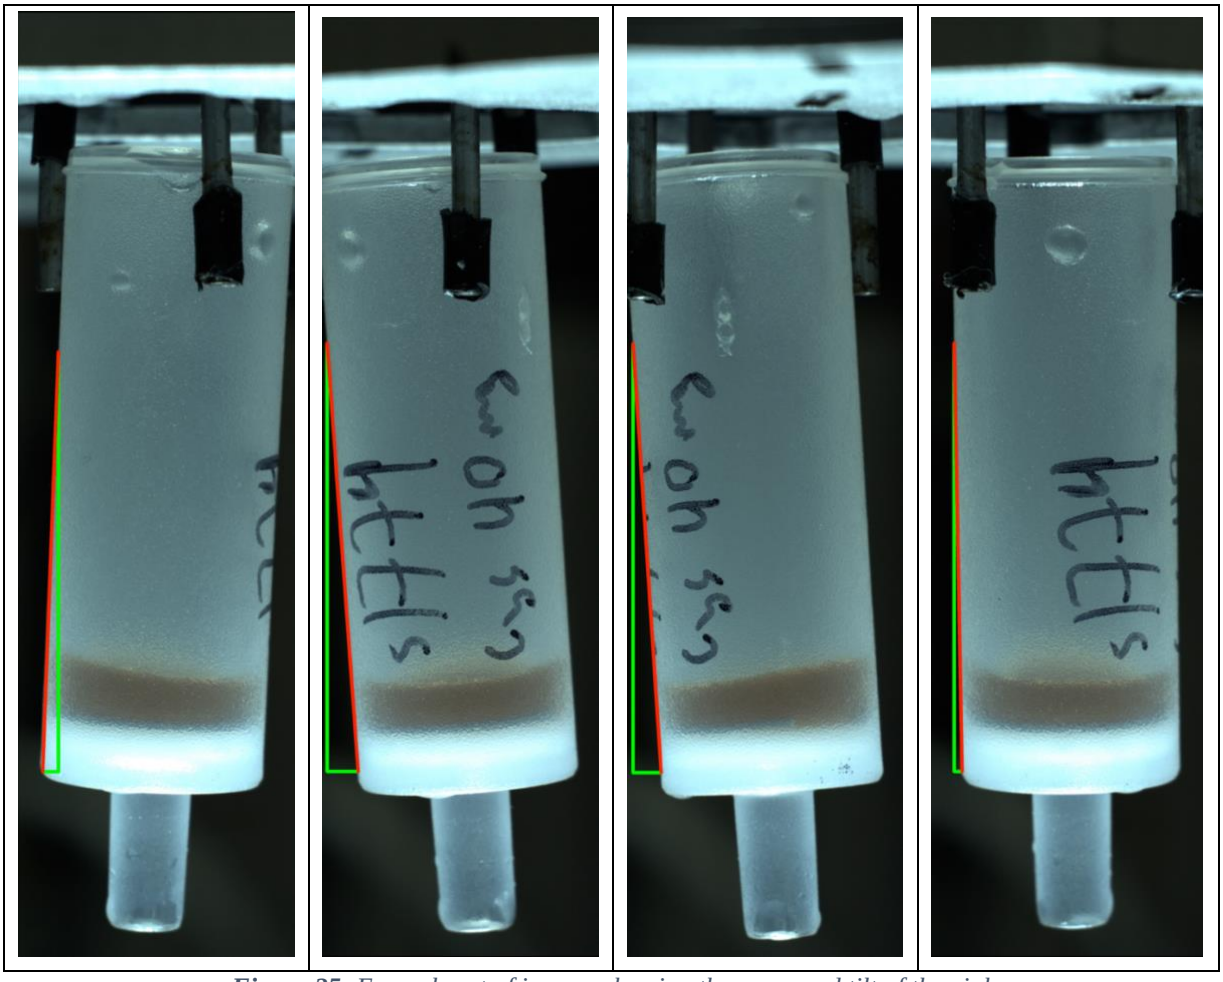

Figure 25: *Example set of images showing the measured tilt of the vial*

The red line follows the edge of the vial, while the green lines indicate the magnitude of the vertical and horizontal components of the tilt.

5.3.4. Identification of Frit Tilt Post Synthesis

The function to identify if the frit has been dislodged during synthesis takes in an image that has been cropped to the region of the image containing only the vial based on the vial edges, tip location, and vial tilt previously determined. The frit can be isolated from the rest of the cropped image by identifying the pixels that lie within the color range associated with the frit. With the ring light placed below the vial, the frit appears as a bright white region near the bottom of the image, making it easy to separate from the rest of the image. The pixels associated with the frit are found through image segmentation using the InRange function from the OpenCV package in Python. A frit that is sitting correctly at the bottom of the vial will be represented by a slim grouping of white pixels at the bottom of the vial. By contrast, a dislodged frit will appear as a larger, white oval, affecting a much larger region. If the shape of the segmented frit pixels is larger than is expected for a frit that is sitting correctly, the function outputs a Boolean value of "True." If the region of pixels is within the expected range, the function outputs a Boolean value of "False." This function also identifies the location of the top of the frit by looking for the first row to contain a significant number of white pixels. This location is output from the frit tilt function and used as an input into the resin volume function.

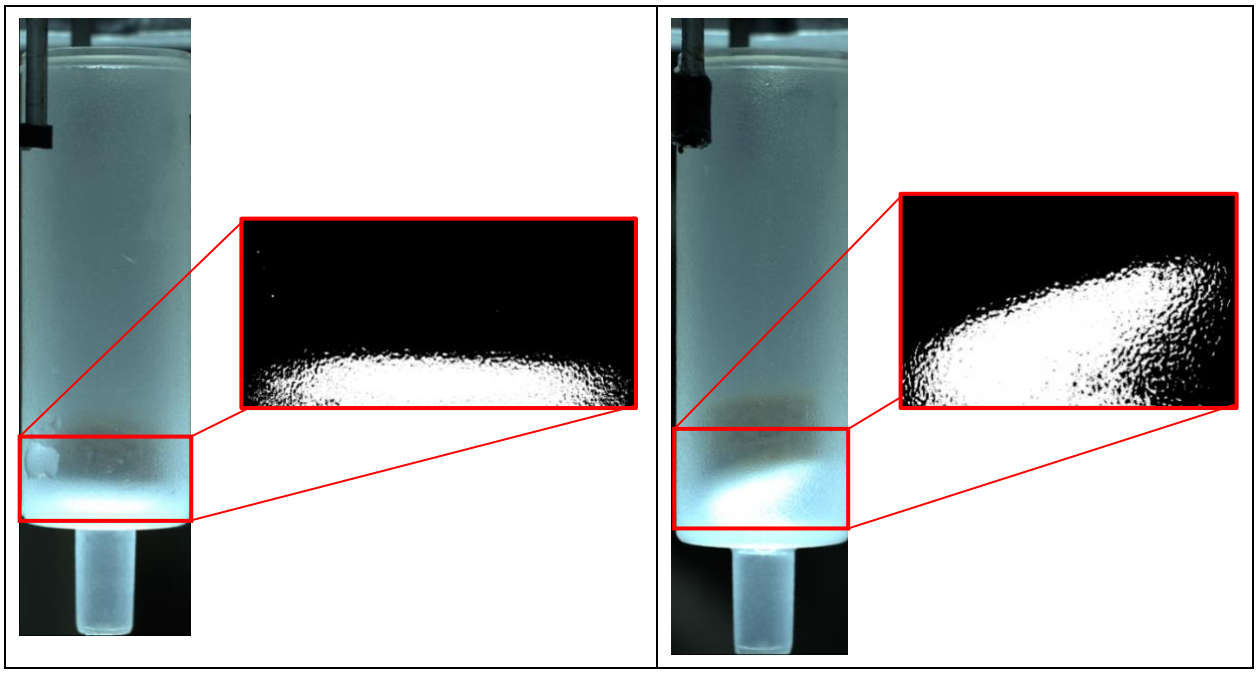

Figure 26: *Images taken by the central camera showing the isolated frit found using the frit tilt function.*

Left: Isolated image of a frit when sitting correctly in the vial.
Right: Isolated image of a frit that has been dislodged from its correct sitting position during the synthesis process

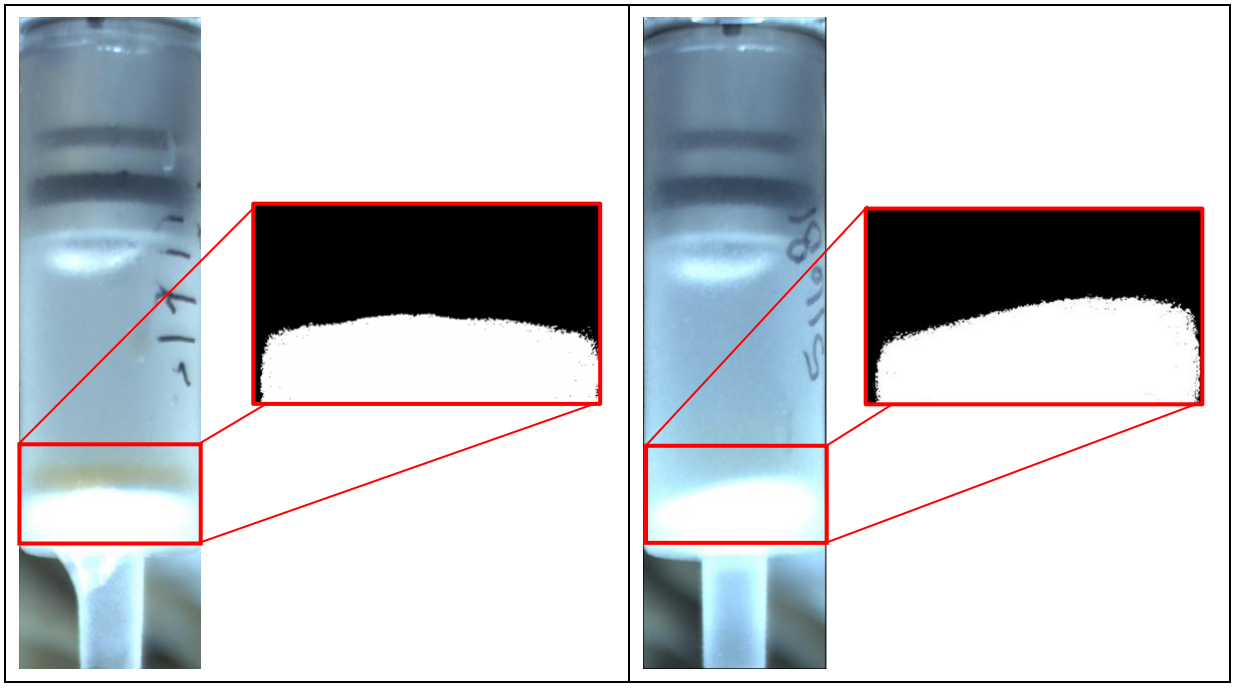

Figure 27: *Images taken by the synthesizer camera showing the isolated frit found using the frit tilt function.*

Left: Isolated image of a frit when sitting correctly in the vial.
Right: Isolated image of a frit that has been dislodged from its correct sitting position during the synthesis process

5.3.5. Determining Vial Clearance of Gripper for Safe Removal

Determining if the vial has cleared the seal of the synthesizer requires the successful identification and classification of three objects in the image: the seal, the thermocouple, and the vial.

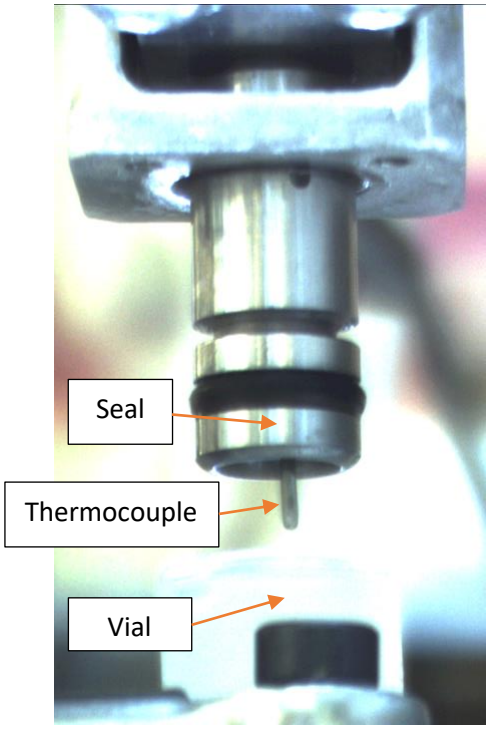

Figure 28: *Landmarks located by the vial clearance function, the seal, thermocouple, and vial*

First, Canny edge detection is performed on the original image and is then cropped based on the location of the seal within the image. To identify the gripper, the program looks for two vertical lines of white pixels that are separated by the number of columns equivalent to the diameter of the gripper. The last row of these two columns is labeled as the end of the visible seal. Next, the program looks for the thermocouple, starting immediately below the end of the gripper. It identifies the thermocouple using the same method used to identify the seal. However, it adjusts the number of columns expected between the two vertical lines of white pixels to be equivalent to the diameter of the thermocouple. If it is unable to locate the thermocouple, the program assumes it is covered by the vial, indicating that the vial is too high on the gripper and needs to be lowered for safe

removal. If the thermocouple is identified, the program identifies the last row of pixels that are the correct distance apart and labels that row as the end of the thermocouple. The final object located by the function is the vial itself. Starting immediately below the end of the thermocouple, the function starts counting rows of pixels that are entirely black. These rows represent the open space between the vial and the thermocouple. When the top of the vial is encountered, the number of empty rows is converted to a distance in mm. If this distance is below a set minimum threshold, then the vial is determined to be too close to the gripper and unsafe to retract.

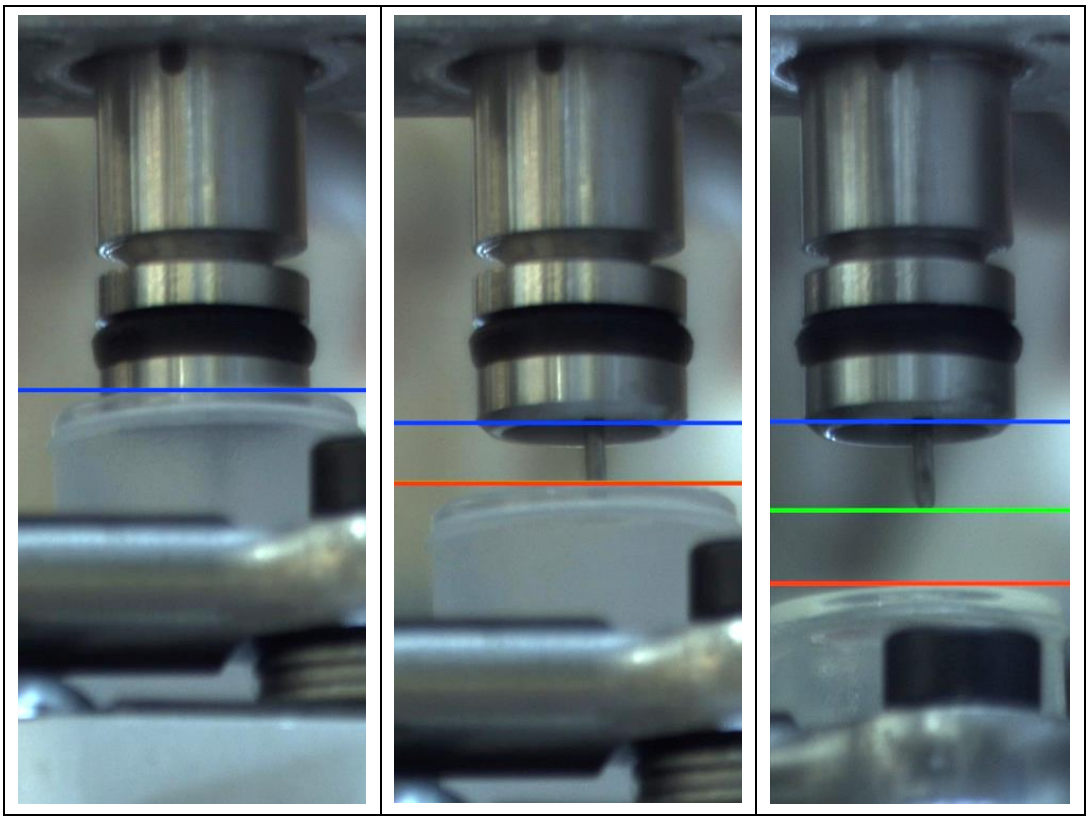

Figure 29: *Processed images taken by the synthesizer camera during removal of the vial from the synthesizer*

Left: Image showing the vial when not cleared from the synthesizer seal. The line indicates the point where the seal and the vial meet.
Center: Image showing the vial when not cleared from the thermocouple. The first line indicates the bottom of the gripper. The second line indicates the meeting point between the thermocouple and the vial.
Right: Image showing the vial when it has cleared both the synthesizer seal and thermocouple. The top line indicates the end of the gripper, the center line indicates the end of the thermocouple, and the bottom line indicates the top of the vial.

5.3.6. Estimating Resin Volume Pre and Post Synthesis

The volume of resin in the vial is estimated before and after synthesis. Estimating the resin of the volume requires two inputs, both dependent on the information obtained during the vial analysis function. The first is an image that has been cropped to the region containing only the vial based on the vial edges, tip location, and vial tilt data found in the previous functions. The second is the location of the top of the frit, found through the frit tilt function. This function isolates the resin in the vial in much the same way the frit tilt function isolates the frit by finding the pixels in the image that lie within a color range associated with the desired object, in this case, the resin. Several color profiles were considered, but it was determined that the blue, green, and red (BGR) color profile best distinguished the resin color from the rest of the image. This particular color profile switches the blue and red pixel intensities from the image's original red, green, and blue (RGB) color profile. This color profile was selected as it is the default color profile used by OpenCV and exhibits stark contrast of color between the resin and the vial. When analyzing the color distribution of the pixels in an image, the BGR color profile shows a much clearer distinction between the pixels associated with the resin and the pixels associated with the vial. Using the color distinction between the resin and the vial, the function separates the resin pixels from the rest of the image using the same InRange function that was used to identify frit tilt.

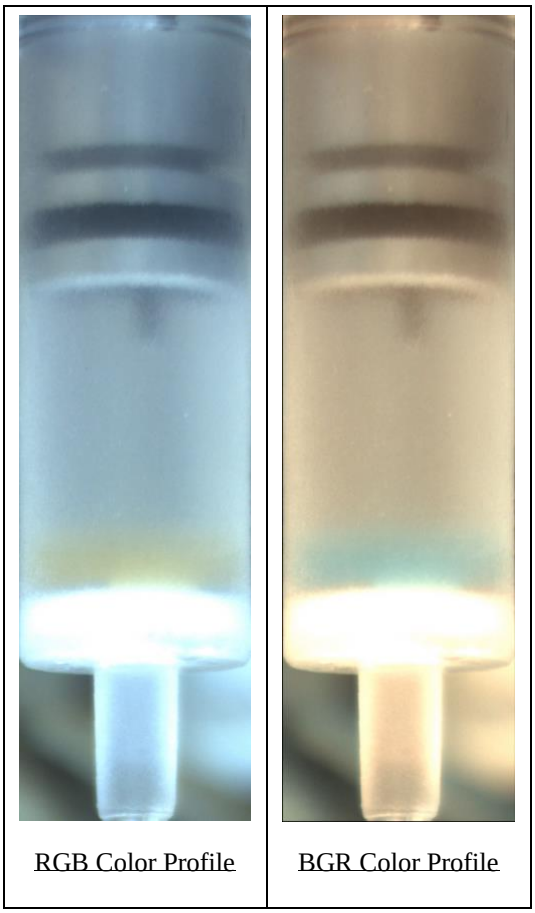

Figure 30: *Example image of a vial that has been converted from the RGB color profile (left) to BGR(right).*

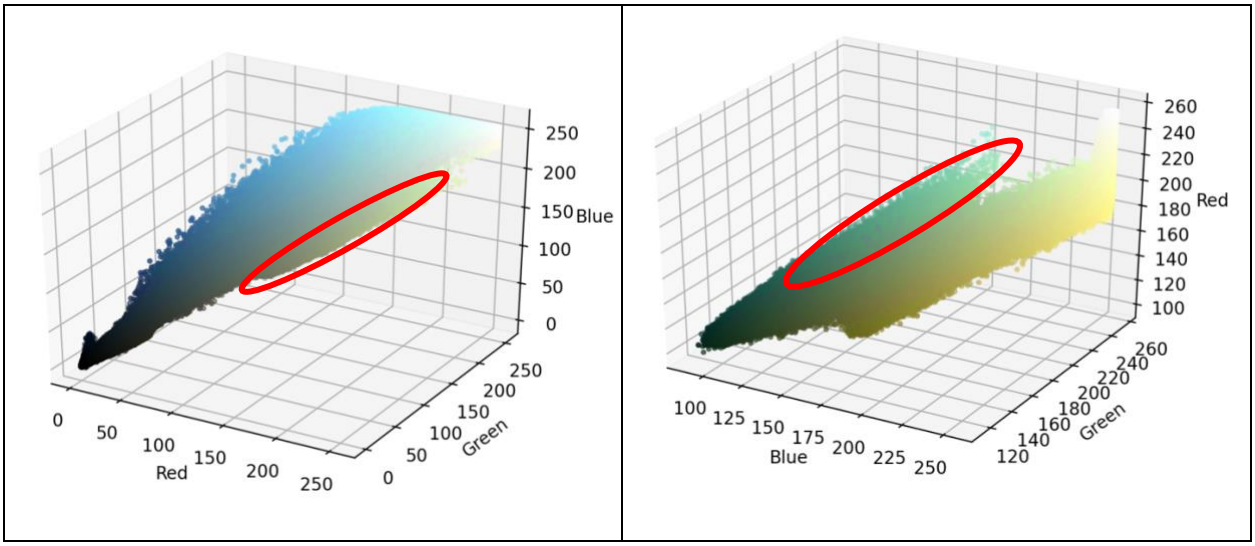

Figure 31*: Color plots showcasing the color distribution of pixels of a standard vial image.*

The red circle indicates the region of pixels pertaining to the resin.
Left: Pixel color distribution in the red, green, blue (RGB) color profile. Resin pixels
Right: Pixel color distribution in the blue, green, red (BGR) color profile.

Several types of resin can be used to serve as scaffolding for the peptide during synthesis. Mytide uses three types of resin, which will be referred to as Resins 1, 2, and 3 in this document. It is also often necessary to preload the resin with the first amino acid to be coupled. This helps increase the efficiency of the subsequent amino acid coupling reactions. To isolate the resin from the rest of the image, the program looks for a specific pixel color intensity range associated with that resin. However, this color range differs by both resin type and by amino acid preload. Additionally, the resins are prepped in batches, which are then used for multiple peptides. The color variation between these batches is also significant. This makes it difficult to predict the color range of pixels to separate, which necessitates the automatic generation of the color profile for each image. Finally, because of the differences in the appearance of the resin before and after synthesis, separate methods were developed to isolate the resin and estimate the volume for each state.

5.3.6.1. Pre-Synthesis Resin Isolation and Volume Estimation

The three types of resin used by Mytide vary in color, contour, and mass, making generating a color profile for each difficult. For example, Resin 1 is dark in color and is more likely to have an uneven distribution of resin throughout the vial. Resins 2 and 3 are lighter in color but have a more even distribution compared to Resin 1.

***Table 1**: Examples of different types of resins and preloads used at Mytide*

Resin 1 - Gly	Resin 1 - Cys	Resin 1 - Phe	Resin 1 - Gly	Reisn 1 - Gly
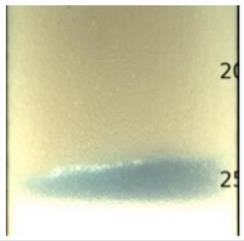	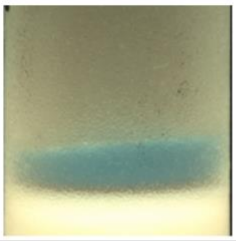	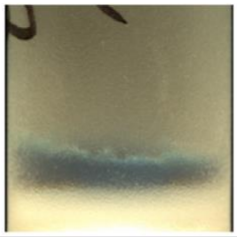	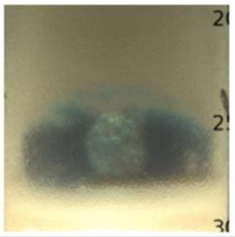	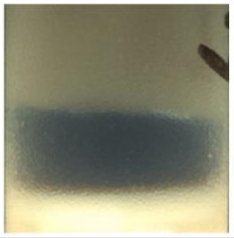
Resin 2	**Resin 2**	**Resin 2**	**Resin 3 - Gly**	**Resin 3 - Gly**
	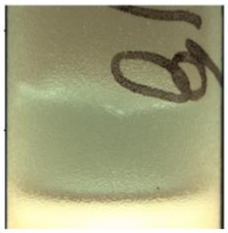		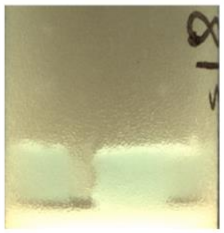	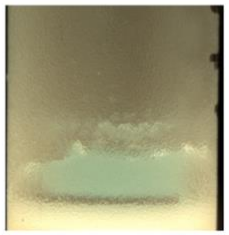

To account for this variation, the function takes a tiered approach when determining the range of pixel intensities to look for. Three tiers of pixel ranges are defined; the first uses a wide set of parameters when looking for pixels associated with the resin. Each tier below it narrows the parameters until an acceptable isolation is obtained. If the final tier is unable to isolate the resin, then the function returns “None”.

The first step in determining these parameters is identifying the region of the image where the resin is expected to be located. The primary point of reference is the frit location found in the frit tilt function. Because the resin is always expected to lie directly on top of the frit, the function assumes that the region directly above this location is going to consist of pixels corresponding to the resin. The function takes a sample of pixels just above the frit and calculates the average and standard deviation of the three color intensities assigned to each pixel. These numbers are then used as the basis for generating a color profile for the resin. The lower pixel color value is calculated as the average (μ) minus x number of standard deviations (σ), while the upper pixel

color value is calculated as the average pixel intensity plus x number of standard deviations. The value of the factor x is different for each search tier. For the first tier, x = 3, the second x = 2, and the final tier x = 1.

$$p_{low} = \mu - (x * \sigma)$$

Equation 4: Low pixel value threshold

$$p_{high} = \mu + (x * \sigma)$$

Equation 5: High pixel value threshold

This results in a range of blue, green, and red pixel intensities associated with the resin specific to each image.

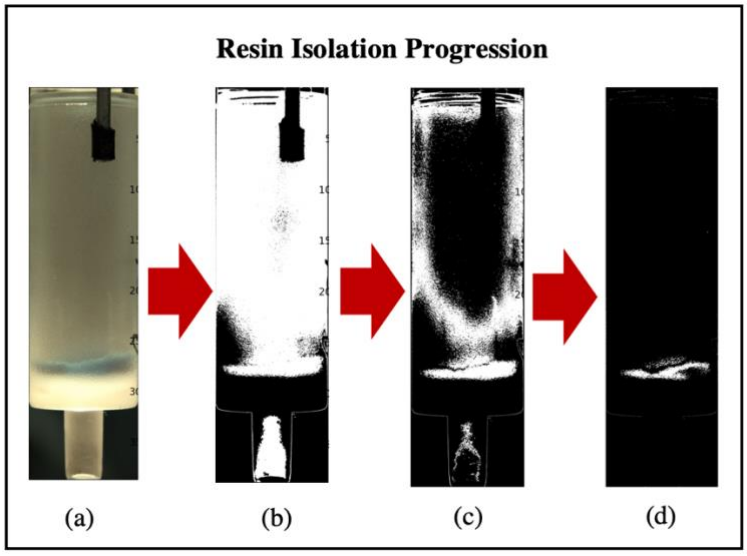

***Figure 32**: Progression of an image of resin in a vial as it moves through each tier in the resin isolation code.*

(a) The original image that has been cropped and rotated based on the vial edges, tip location, and tilt found in the vial presence, vial location, and vial tilt functions.
(b) The image after being processed through the first tier in the resin isolation function.
(c) The image after being processed through the second tier in the resin isolation function.
(d) The image after being processed through the third tier in the resin isolation function.

Once a color range has been determined, the function pulls out any pixel in the image that lies within that range. It then loops through each row of pixels from the top of the image to the top of the frit, separating rows containing a significant number of pixels from those that do not. The

function then identifies regions of the image where significant rows of pixels are grouped together and selects the grouping closest to the frit as the pixels most likely to be the resin.

Once the function has made an estimate of the resin location, it must verify that the resin has been properly isolated. It looks at the region of pixels just above this resin line for a "cloud" of pixels. If the isolation is clean, then this region will not contain many white pixels. If the isolation is still cloudy, then the function narrows the search range and tries again.

After the resin has been isolated, the function can estimate its volume. The function discretizes the resin volume by determining the approximate volume of each column of pixels within the resin. To determine the height of resin in each column, it calculates the difference between the top and bottom resin pixels in that column, centered around the resin line identified when isolating the resin. It then determines the column's distance from the center of the resin and uses that distance as the outer radius and sets the inner radius as one pixel closer. The column of pixels is revolved around the center of the resin 180 degrees and finds its volume based on the height of pixels in the column and distance from the center. It repeats this process for all resin columns and sums them up to get an estimate of the total volume. This number then serves as the primary output of this function.

$$scale = \frac{d_{vial}}{n_{cols}}$$

Equation 6: Image scale, number of millimeters per pixel

$$h_i = p_{bottom_i} - p_{top_i}$$

Equation 7: Height of a column of resin

$$r_{outer_i} = abs(col_{center} - col_i)$$

Equation 8: Outer radius of a column of resin

$$r_{inner_i} = r_{outer_i} - 1$$

Equation 9: Inner radius of a column of resin

$$A_i = \frac{\pi * {r_{outer_i}}^2 - {r_{inner_i}}^2}{2}$$

Equation 10: Area corresponding to a column of resin after a rotation of 180 degrees about the center of the resin

$$v_i = h_i * A_i * scale^3 * 0.001$$

Equation 11: Volume corresponding to a column of resin after a rotation of 180 degrees about the center of the resin

$$V = \sum_{i=1}^{n_{col}} v_i$$

Equation 12: Estimated volume of resin, pre-synthesis

Table 2: *Examples of resin volume estimation of Resin 1*

Examples have been preloaded with Glycine (Gly), Cysteine (Cys), and Phenylalanine (Phe).
The first image is the original resin image, the second is the resin isolation image. The third image depicts the height of the resin in each column used to calculate the volume.

Resin	Resin 1 - Gly	Resin 1 - Cys	Resin 1 - Phe	Resin 1 - Gly	Resin 1 - Gly
Original Image (BGR Color Profile)	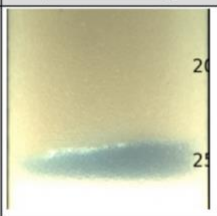	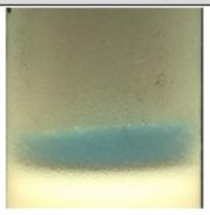	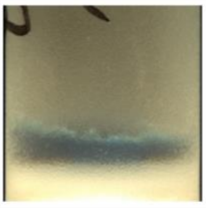	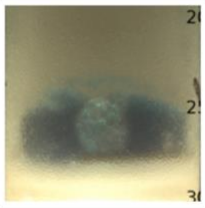	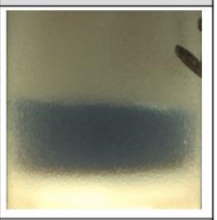
Isolated Resin				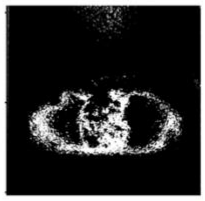	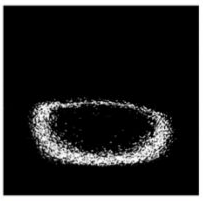
Resin Contour Profile	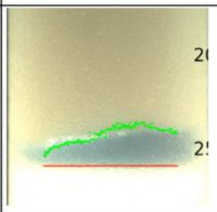	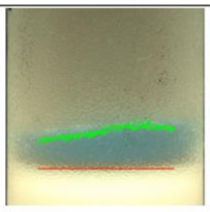	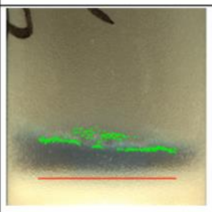	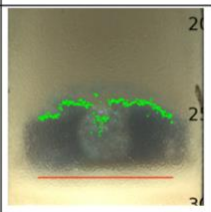	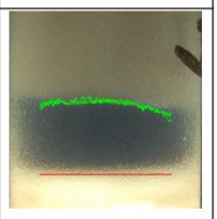
Volume	0.083 mL	0.084 mL	0.085 mL	0.166 mL	0.184 mL

***Table 3**: Examples of resin volume estimation of Resins 2 and 3.*

The first image is the original resin image, the second is the resin isolation image. The third image depicts the height of the resin in each column used to calculate the volume.

Resin	Resin 2	Resin 2	Resin 2	Resin 3	Resin 3
Original Image (BGR Color Profile)	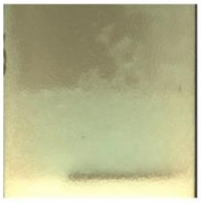	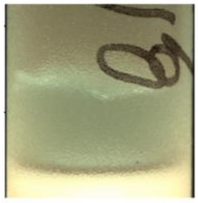		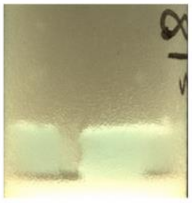	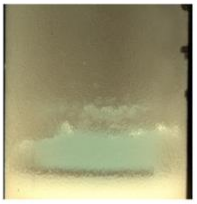
Isolated Resin		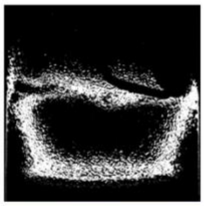			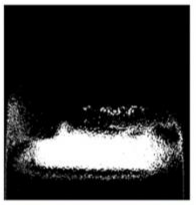
Resin Contour Profile	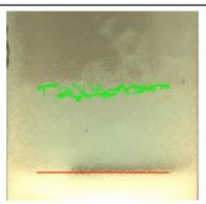	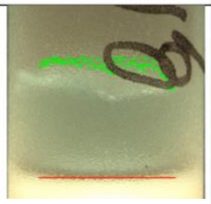	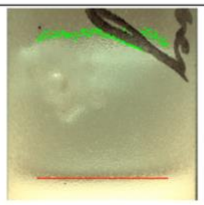	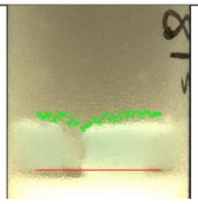	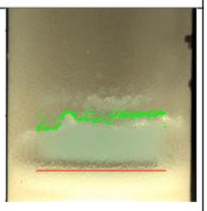
Volume	0.250 mL	0.311 mL	0.433 mL	0.202 mL	0.179 mL

5.3.6.2. Post-Synthesis Resin Isolation and Volume Estimation

During synthesis, the resin is suspended in DMF. This causes the resin to lose its defined edges and appears as a color gradient in the vial rather than the distinct mass of resin observed before synthesis.

***Table 4:** Examples of the different types of resins used at Mytide after synthesis*

Resin 1	Resin 1	Resin 2	Resin 3	Resin 3
				

The color range for isolating resin after synthesis is determined in a similar way as the color range for the pre-synthesis resin. A region of pixels just above the frit is selected as a sample of resin pixels and a color range is calculated based on the average and standard deviation of the blue, green and red pixel intensity values using equations 4 and 5. However, because the resin distribution is uniform, only two tiers of parameter settings are required to detect the level of resin in the vial.

The first tier is used to isolate Resin 1, which appears darker in color and exhibits a more distinct border between it and the DMF. The same equation is used to calculate the color range in this step, assigning $x = 3$ to equations 4 and 5. Because the resin color is uniform and evenly distributed, it is easily isolated. The function then takes the resulting binary image containing the resin pixels and searches for a region of the image just above the frit line. The first row containing a significant (greater than 80) number of white pixels is recorded as the location of the top of the resin.

In contrast, the Resins 2 and 3 resins are lighter in color and greater in volume than Resin 1. When images of these resins are passed through the first tier, the function only picks up the glare of the vial that is present due to the vial's stippling. Because of this, the image containing the isolated pixels will exhibit a large number of white pixels in the outer two-thirds of the image, and relatively few pixels in the center third. When this condition is detected, the function passes the image to the next tier.

To determine the color range, the second tier looks at a region of pixels further above the frit than the first tier to avoid the glare caused by light reflecting off of the frit, then assigns $x = 5$ in equations 4 and 5 to determine the color range. The resulting isolation still contains a large amount of glare in the outer two-thirds of the image, however, the center third of the image

contains a more accurate representation of the level of resin within the vial. When looking for rows containing significant numbers of white pixels, the function excludes the outer two-thirds of the image and only counts pixels in the center third. The level of resin is estimated to be the first significant row (containing more than 130 white pixels) identified as it moves from the top of the image to the bottom of the image.

Because the height of the resin is uniform after synthesis, it is not necessary to discretize the volume calculations as was done for estimating the volume before synthesis. Instead, the height of the resin is set to the difference between the top of the resin that was identified during the isolation step and the top of the frit. The volume is then calculated using the known dimensions of the vial.

$$V = h_{resin} * \pi * r_{i,vial}^2 * scale$$

Equation 12: Estimated volume of resin post-dry

Table 5: *Examples of resin volume estimation of Resins 1, 2, and 3.*

The first image is the original resin image, the second is the resin isolation image. The center third of this image is used to find the level of the resin in the vial. The third image depicts the height of the resin in each column used to calculate the volume found when analyzing the resin isolation image.

Resin	Resin 1	Resin 1	Resin 2	Resin 3	Resin 3
Original Image (BGR Color Profile)					
Isolated Resin					
Resin Contour Profile					
Volume	0.141 mL	0.160 mL	0.423 mL	0.350mL	0.486 mL

5.4. Function Architecture

Each camera is assigned a function containing a series of sub-functions for each image or set of images it produces. Each sub-function contains a series of calls to the image analysis functions required to obtain the information needed for those specific images. The central inspection camera function contains sub-functions for the pre-synthesis and post-dry image sets. The synthesizer function contains sub-functions pre-synthesis, post-synthesis, and vial removal images.

The image analysis functions contain sub-functions pertaining to each camera. While each sub-function is similar in goal, the different views and images taken at each camera vary enough to warrant specialized functions specific to each camera to handle these differences.

5.5. Central Inspection Camera Image Analysis

The central inspection camera takes a set of four images at two points in the process, just before the vial is placed in the synthesizer and just after it has been removed from the drier. Each individual photo has its own metadata attached to it containing the results of the analysis. The individual data will then be consolidated into a single set of results describing the status of the vial at the relevant step.

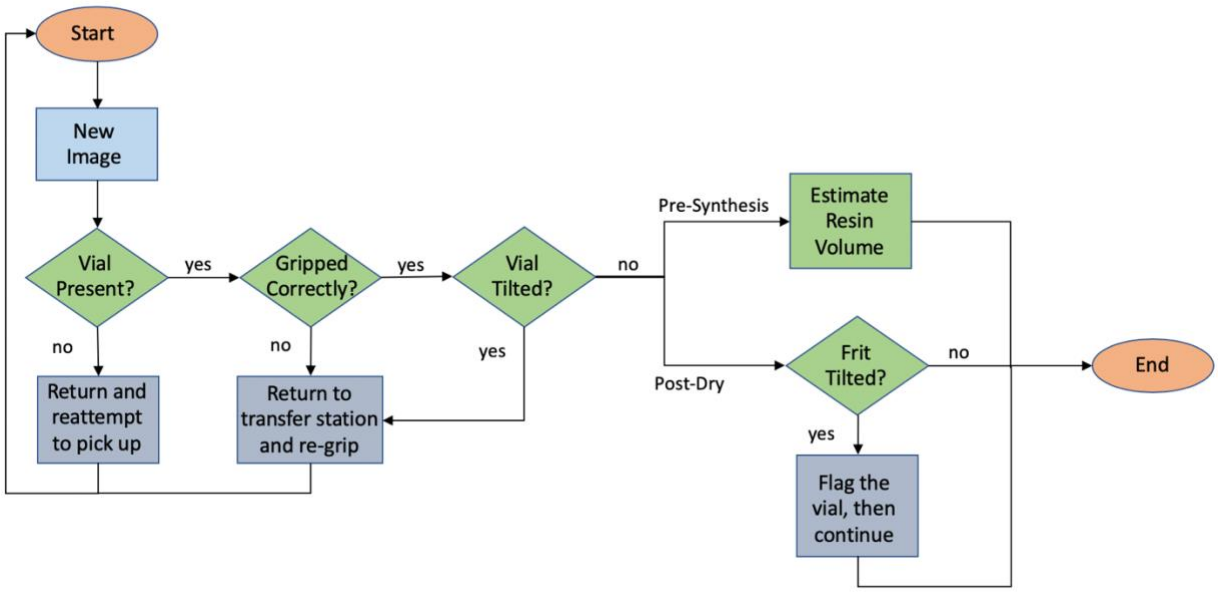

Figure 33: *Process flow diagram of the central inspection camera image analysis.*

5.5.1. Central Inspection Camera: Pre-Synthesis

The first images taken of the peptide vial are taken by the central inspection camera just before the vial is placed in the synthesizer. Four images are taken and analyzed. First, the function

confirms that a vial is present in all four images. If one or more images do not contain a vial, then the robot returns to the tray where the vials are initially queued up for synthesis, and reattempts to pick the vial up. After confirming the presence of a vial in the robot gripper, the function determines whether or not it has been gripped correctly. If the vial is found to be gripped incorrectly or tilted in one or more images, the robot returns the vial to the transfer station (a station used throughout the process for purposes of adjusting the robot's grip on the vial). The vial is regripped by the robot and then returns to the central inspection station again to re-analyze the grip.

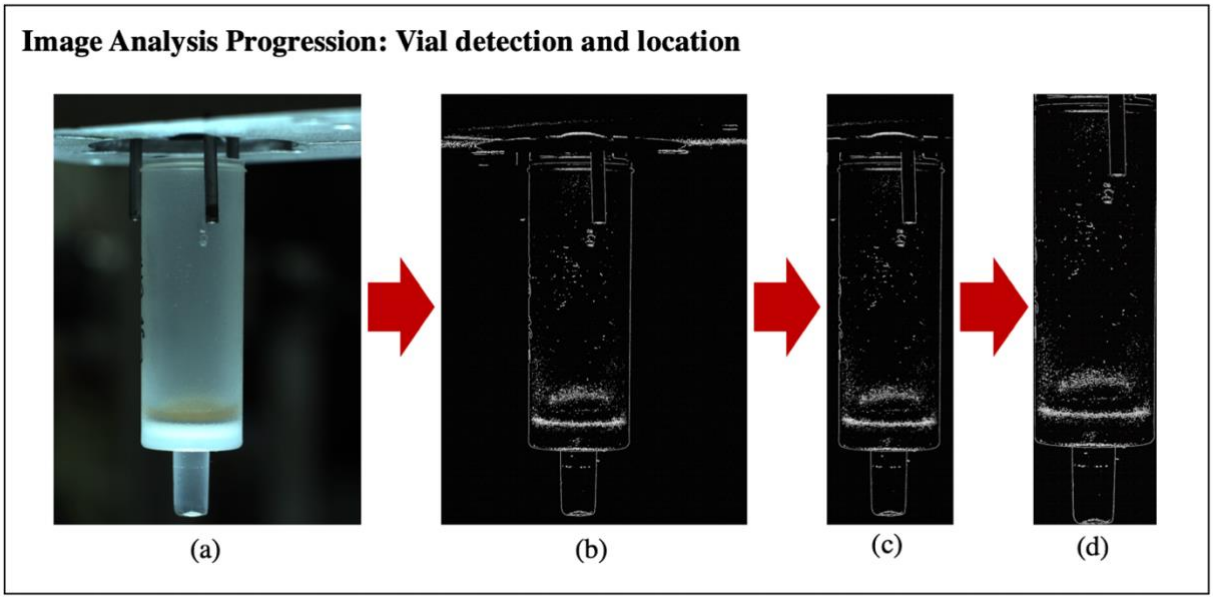

Figure 34: *Progression of an image taken by the central camera before synthesis as it is analyzed for vial presence, location, and tilt.*

(a) The original image
(b) Edge detection is performed on the original image
(c) The image is cropped based on the vial edges found in the vial presence function.
(d) The image is cropped again based on the tip location found in the vial location function and tilt measured in the vial tilt function.

Finally, the function estimates the volume of the resin for all four images. It then takes an average of the four estimations and records that value as part of the peptide's metadata.

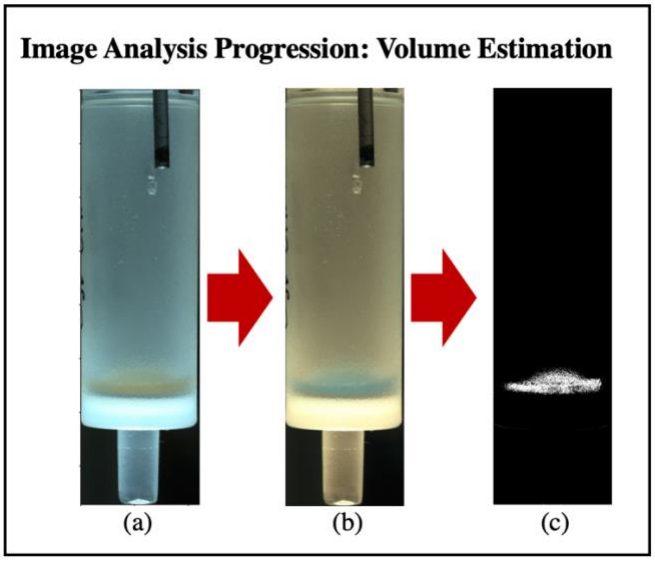

***Figure 35**: Progression of an image taken by the central camera before synthesis as it is analyzed for resin volume.*

(e) The original image that has been cropped and rotated based on the vial edges, tip location, and tilt found in the vial presence, vial location, and vial tilt functions.

(f) The image is switched from the RGB color profile to the BGR color profile.

(g) The pixels of the resin are isolated based on the resin color range automatically detected by the function.

5.5.2. Central Inspection Camera: Post-Dry

The set of images taken by the central inspection camera go through much the same analyses as the images taken before synthesis. The images are checked first for vial presence. If a vial is not found to be present, the robot returns to the drier to reattempt to grip the vial. After a vial is confirmed to be present in the robot gripper, the function confirms that the vial is gripped correctly by the robot and that it is not tilted, returning to the transfer station to regrip the vial if necessary.

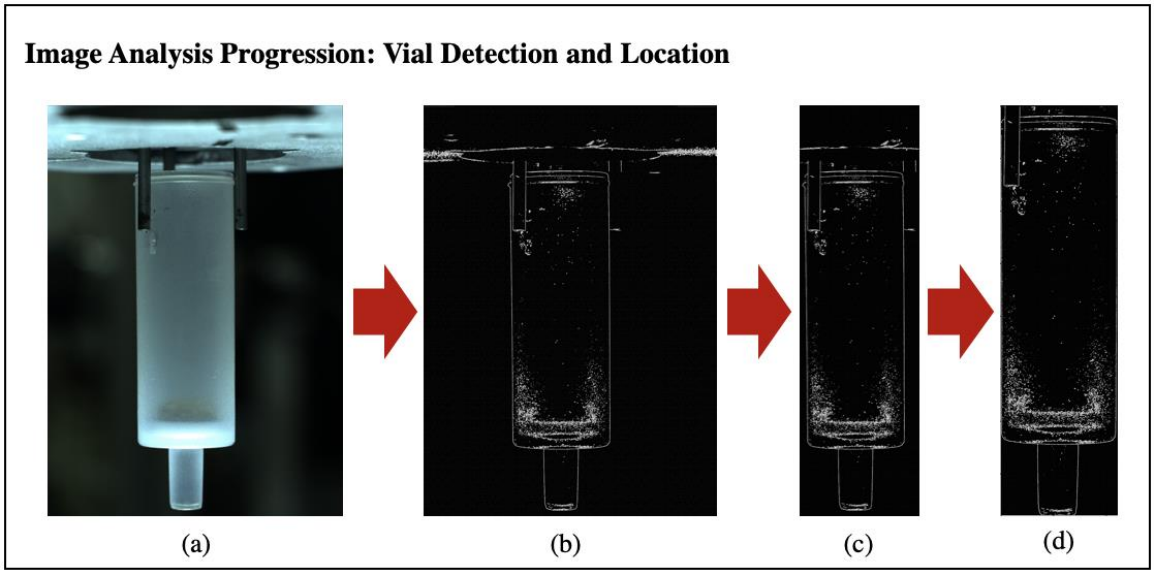

Figure 36: *The progression of an image taken by the central inspection camera after drying as it is analyzed for vial presence, location, and tilt.*

(a) The original image
(b) Edge detection is performed on the original image
(c) The image is cropped based on the vial edges found in the vial presence function.
(d) The image is cropped again based on the tip location found in the vial location function and tilt measured in the vial tilt function.

Lastly, instead of estimating resin volume, the function checks if the frit was dislodged during synthesis. If one or more images read that the frit is dislodged, the vial is flagged then moves on to the next stage of the manufacturing process.

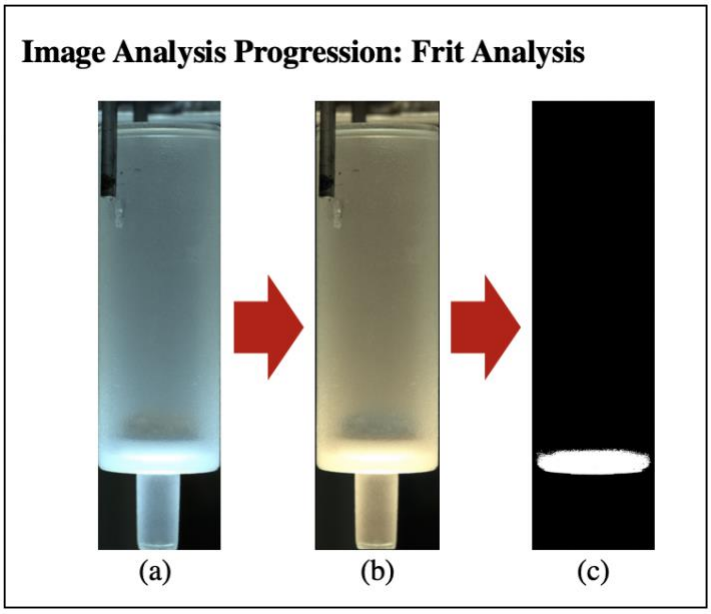

Figure 37: *The progression of an image taken by the central inspection camera after drying as it is analyzed for frit tilt.*

The original image that has been cropped and rotated based on the vial edges, tip location, and tilt found in the vial presence, vial location, and vial tilt functions.

(a) The image is switched from the RGB color profile to the BGR color profile.
(b) The pixels of the frit are isolated based on the resin color range automatically detected by the function.

5.6. Synthesis Camera Image Analysis

The synthesizer camera takes a total of three pictures, one before synthesis, one after synthesis, and one before removal of the vial. Each individual photo has its own metadata attached to it containing the results of the analysis.

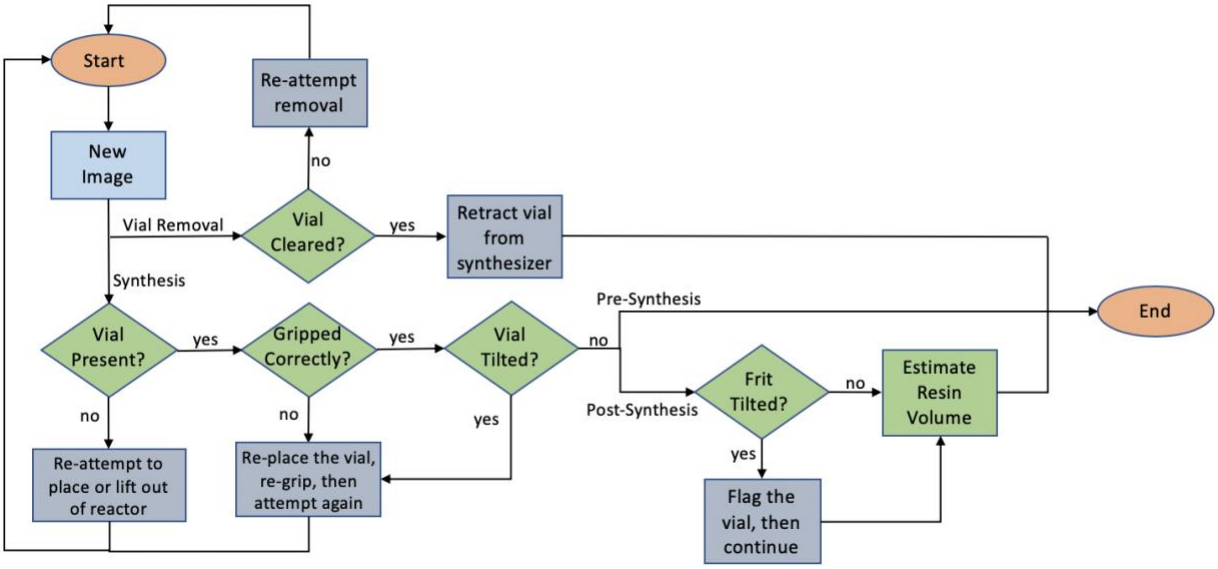

***Figure 38**: Process flow diagram of the synthesizer camera image analysis.*

5.6.1. Synthesizer Camera: Pre-Synthesis

Prior to synthesis, the main objective of the image analysis is to ensure that the vial is present and securely gripped before being lowered into the reactor. First, the function inputs the image into the "vial presence" function. If a vial is not found to be present, it sends a message to the robot to reattempt to place the vial on the synthesizer seal. After re-placing the vial, another photo will be taken, and the new photo will be checked for vial presence. If it fails a second time, production will be halted, and an alert will be sent out to the team. If a vial is present, the next step is to determine whether or not the vial is gripped correctly by the seal. If not, the robot will regrip the vial, release it from the synthesizer seal, then reattempt to place it. Once all checks have passed, the synthesizer is cleared to begin synthesis, and the vial is lowered into the reactor.

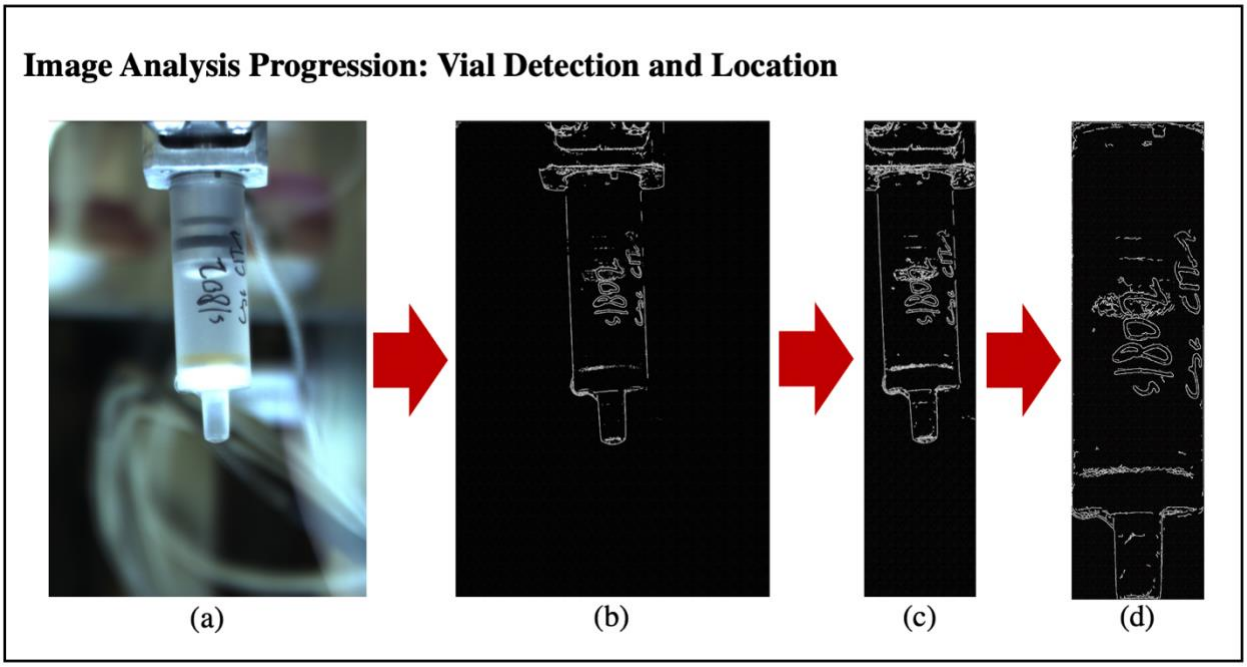

Figure 39: *Progression of an image taken by the synthesizer camera before synthesis as it is analyzed for vial presence, location, and tilt.*

(a) The original image
(b) Edge detection is performed on the original image
(c) The image is cropped based on the vial edges found in the vial presence function.
(d) The image is cropped again based on the tip location found in the vial location function and tilt measured in the vial tilt function.

5.6.2. Synthesizer Camera: Post-Synthesis

After synthesis, the vial is lifted up out of the reactor and another picture is taken. The image is analyzed for vial presence, location, and tilt. When the image is analyzed for vial presence, if a vial is not found, then it assumes that the vial was not lifted out of the reactor. The synthesizer is then instructed to reattempt lifting the vial, and another photo is taken. After a vial is confirmed present, the function checks for vial location, then vial tilt, and instructs the robot to adjust the grip if necessary.

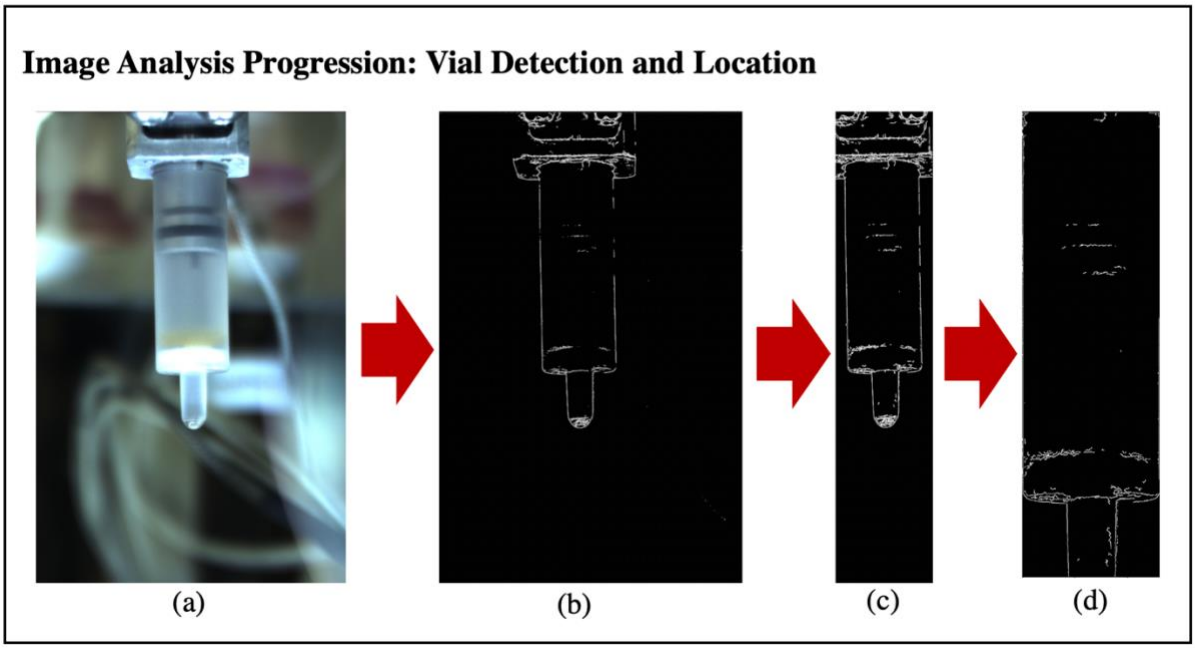

Figure 40: *Progression of an image taken by the synthesizer camera after synthesis as it is analyzed for vial presence, location, and tilt.*

(a) The original image
(b) Edge detection is performed on the original image
(c) The image is cropped based on the vial edges found in the vial presence function.
(d) The image is cropped again based on the tip location found in the vial location function and tilt measured in the vial tilt function.

Once the vial is properly gripped by the synthesizer seal, two additional checks are made. The first determines if the frit was dislodged during synthesis. If the frit was dislodged, it attaches a flag to that vial for review later. Lastly, the volume of the resin in the vial is estimated and recorded as part of the peptide's metadata.

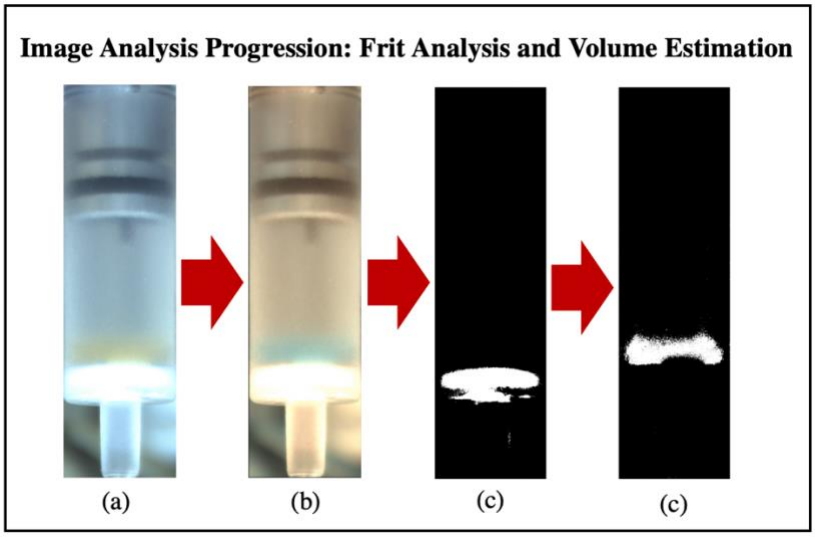

Figure 41: *Progression of an image taken by the synthesizer camera before synthesis as it is analyzed for frit tilt and resin volume*

(a) The original image that has been cropped and rotated based on the vial edges, tip location, and tilt found in the vial presence, vial location, and vial tilt functions.
(b) The image is switched from the RGB color profile to the BGR color profile.
(c) The pixels of the frit are isolated based on the resin color range automatically detected by the function.
(d) The pixels of the resin are isolated based on the resin color range automatically detected by the function.

5.6.3. Synthesizer Camera: Vial Removal

The final image is taken just after the vial has been removed and lowered from the synthesizer seal. Only one check is performed on this image, determining whether or not the vial has been cleared of the synthesizer seal before the robot attempts to retract it from the synthesizer. If the vial has not cleared the seal, the robot reattempts to remove the vial by re-placing it on the seal, then starting the removal process over. Another photo is then taken and checked. Once the vial is clear of the seal, it is retracted from the synthesizer and carries on through production.

Chapter 6: Results

In total, over 1400 images were taken and processed over 148 peptide synthesis runs in the development and validation of this machine vision system. Each image was processed using the methods described in Chapter 5. The analysis of these images provides insight into the most common fault conditions experienced by the system and serves as a base from which to develop a database of images and image analyses to inform future process improvements.

6.1. Fault Condition Detection Results

All 148 runs took pictures on the vial at the central inspection station before synthesis. Of these 148 runs, there were four instances where the vial was not present in the robot gripper when it was expected, resulting in a failure rate of 2.7%. Additionally, there were three cases of the vial tilt exceeding the maximum allowable tip displacement of 1.5 mm, resulting in a 2% failure rate. There were zero cases of the robot gripping the vial at the wrong height.

There were 62 peptide runs that included images of the vial taken by the synthesizer camera before synthesis. Out of these 62 runs, there was only one case of the vial not being present in the synthesizer when expected.

The synthesizer camera took images of the vial after synthesis during 58 peptide runs. There were 5 instances where the synthesizer failed to pull the vial out of the reactor, resulting in an 8.6% failure rate detected by the vial presence detection function. There were four additional runs where the vial was gripped insufficiently by the synthesizer as it was removed from the reactor, resulting in a 6.9% failure rate for the vial location function, given a vial was present. Finally, 7 of these runs resulted in a dislodged frit, exhibiting a 12% failure rate.

There were 35 images taken by the synthesizer during vial removal. All 35 of these images passed the clearance function, indicating that the vial was clear for removal from the synthesizer.

Finally, images of the vial post drying were taken by the central inspection camera 114 out of the 148 runs. In all runs, the vial was present and gripped at the correct height. However, there were two runs where the tilt of the vial resulted in a displacement of the tip greater than 1.5 mm. Additionally, there were 17 cases of frit dislodgement, a rate of 15%.

Table 6: *Summary of results of the vial analysis by the central inspection and synthesizer cameras*

	Central: Pre	**Synthesizer: Pre**	**Synthesizer: Post**	**Synthesizer: Removal**	**Central: Post**
Not present	4/148	1/62	5/58	N/A	0/114
Not gripped	0/148	0/62	4/58	N/A	0/114
Vial tilted	3/148	N/A	N/A	N/A	2/114
Frit dislodged	0/148	0/62	7/58	N/A	17/114
Not clear	N/A	N/A	N/A	0/22	N/A

The results of these runs indicate the most common fault conditions to occur on the system. At the central camera, there is more likely to be a problem grabbing the vial coming out of the initial queueing trays than there is taking the vial from the dryer. However, both are equally likely to grip the vial at a large enough tilt to cause the tip of the vial to be displaced enough to be a concern.

If a vial is not present in the synthesizer gripper prior to synthesis, it likely corresponds to a vial not present in the robot gripper prior to placement in the synthesizer. This indicates that if the machine vision system is able to catch the lack of a vial present in the robot gripper, then there

should be fewer instances of a vial not present in the synthesizer seal before synthesis when expected.

The step in the manufacturing process most likely to experience issues is removing the vial from the synthesizer reactor after synthesis. The primary issues observed at this step can be attributed to the synthesizer seal failing to maintain a solid grip on the vial as it is lifted out of the reactor. This results either in an improper grip on the vial or leaving the vial behind entirely. The most common fault condition observed was the dislodgement of the frit during synthesis, which can be observed both post synthesis and post dry.

There were fewer pictures of the vial removal from the synthesizer than at other points in the process. While the absence of any detected fault conditions indicates that there is not a profound issue in vial removal, more images are needed to be analyzed before any significant conclusions about the prevalence of this condition are made.

In total, the fault conditions identified correspond 14 total or 9.5% of runs where a system crash or failure could have been avoided with the implementation of the machine vision system. Worst case scenario, a system crash results in eight hours of downtime until an engineer is able to address the issue and get the system back up and running. At the current fault condition detection rate, there are 76 hours per 100 peptide runs that can be saved when the recovery steps are implemented into the system.

6.2. Volume Estimation Results

Of the 148 peptides that ran through the system, 144 had usable images taken of the resin and vial before synthesis. These runs were grouped by resin type, then resin mass. Resin 1 is used in masses of 30, 40, and 60 mg, Resin 2 in 30 and 40 mg, and Resin 3 in 30, 40, and 80 mg. The

average and standard deviation of each set of resins were taken to set a baseline for the volume of resin that can be expected as well as gain insight into any variation in the process with regards to the weighing, loading, and swelling for each type of resin.

Table 7: *Summary of results of the volume analysis by the central inspection and synthesizer cameras pre and post synthesis.*

Resin	**Resin Type**	Resin 1			Resin 2		Resin 3		
	Resin Mass (mg)	30	40	60	30	40	30	40	80
Pre-Synthesis	**n samples**	15	44	14	55	9	3	2	2
	Avg. Vol (mL)	0.084	0.107	0.175	0.232	0.369	0.209	0.221	0.437
	Std. Dev	0.028	0.056	0.014	0.059	0.039	0.002	0.048	0.003
Post-Synthesis	**n samples**	7	0	0	25	0	3	0	2
	Avg. Vol (mL)	0.150	-	-	0.418	-	0.395	-	0.461
	Std. Dev	0.006	-	-	0.052	-	0.058	-	0.025

The resin groups with the most amount of variation within the volume calculation were Resin 1 at 30 and 40 mg, and Resin 2 at 30 mg. The most significant source of this variation is the amount of DMF used by the chemists to prep the resin prior to synthesis. The exact amount of DMF used at this step does not affect synthesis and is not controlled as tightly as other parameters. This results in the volume of the resin to appear greater in cases where more DMF was used, and lower in cases where less was used.

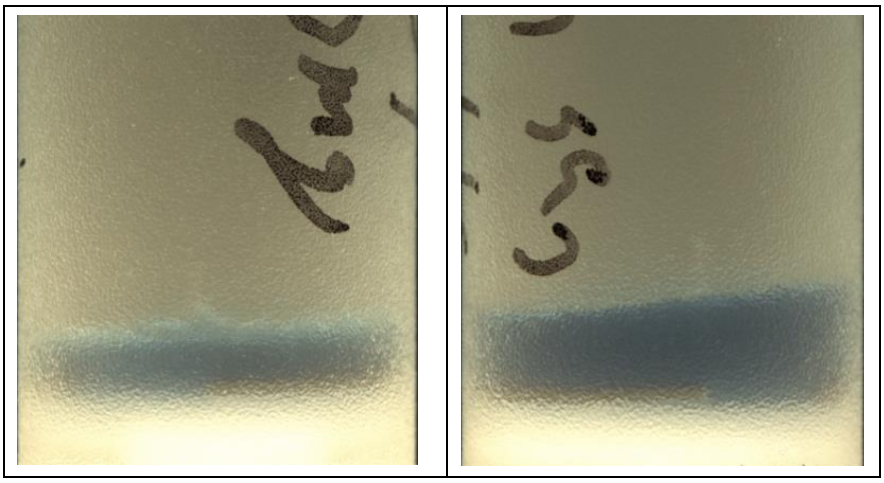

Figure 42: *Example of two different resins, both 40 mg of Resin 1, with differing volumes. Volumes differ due to variation in the amount of DMF used to swell the resin beads prior to synthesis. The resin on the right has a volume of 0.06 and the resin on the left has a volume of 0.13 mL*

There were, however, a few shortcomings that were identified throughout validation that are another source of variation within the measurements. First, undefined edges of the resin pre-synthesis made it difficult for the function to determine where the resin ends and vial begins. Inconsistent color distribution in the resin creates issues in the color range selection step during resin isolation as not all resin pixels will lie within the determined color range. Additionally, there were a couple of cases that caused the function to over or undershoot when identifying the top level of the resin. This is likely due to glare from the vial getting mixed with the edge of the resin during isolation. This can make the resin level to appear higher or cause the image to be passed down to the next tier of isolation parameters, where an excess number of resin pixels will be filtered out of the isolation image. This condition affects Resin 2 at 30 mg the most. While uncommon, these cases illustrate the effect that variation in resin color and distribution has on the accuracy of the volume measurements, and presents opportunities for improvement in both the resin loading process and volume estimation code.

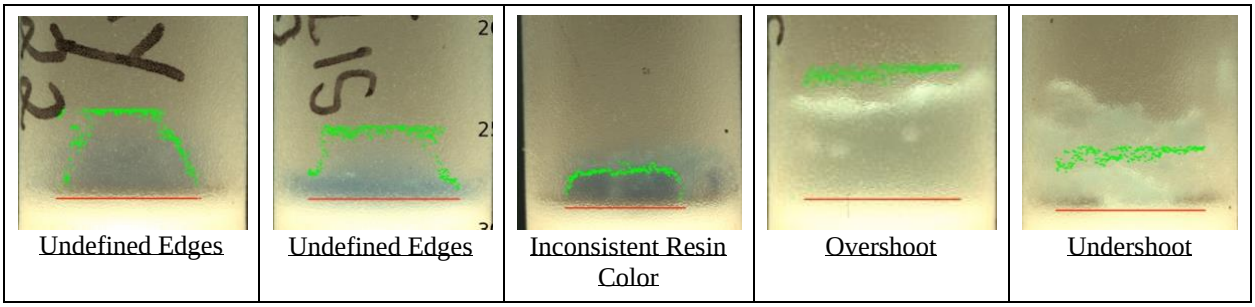

Figure 43: *Examples of inaccurate resin contour detection as a result of unsuccessful resin isolation.*

Due to the limited number of runs possible during the duration of this project, the current analyses of these resin groups are greatly limited by sample size. The volume measurements detailed in this thesis serve as a base from which a database of resin data can be built on as more peptides are manufactured. Moving forward, this database will help Mytide gain insight into the behavior of each type of resin during synthesis and better understand sources of variation within the process.

Chapter 7: Conclusions and Future Work

7.1. Conclusions

In summary, this thesis used several image processing methods to develop a machine vision system for process monitoring and data collection on an automated peptide manufacturing platform at Mytide Therapeutics. This system is designed to detect known fault conditions and designate when recovery steps need to be taken in order to prevent system crashes, decreasing downtime of the system, and increasing total throughput. Additionally, the fault conditions tracked by the machine vision system will give insight to the steps in the manufacturing process that have the greatest potential for improvement. Understanding where the system fails is a vital component to making these improvements to the process that will enable the system to reach the targeted production rates of peptides. The volume data collected through the machine vision system will provide insight into the behavior of the resins during synthesis and inform future process improvements.

The implementation and further development of this machine vision system will advance the automation of the peptide manufacturing platform, accelerating Mytide on their path to providing customers with peptides at a higher purity and faster delivery rate and advancing the field of peptide manufacturing.

7.2. Future Work

This thesis outlines the implementation of a machine vision system on the platform. There are many opportunities available to expand its functionality and capabilities as the company grows.

The following actions are suggested to further develop the machine vision system and increase its overall functionality, robustness, and accuracy.

1. The implementation of the recovery states defined in Chapter 5 for each camera station and image analysis function. This involves creating new recovery states in the robot state path that are triggered by the machine vision system.
2. Adding a camera to the cleaver will provide more steps within the manufacturing process that can monitored in real-time. Ensuring that a vial is present and gripped correctly while in the cleaver seal will not only ensure that the process runs smoothly, but can make the process safer overall by preventing leaking of the cleavage cocktail, which is dangerous to humans.
3. Implementing a final product inspection station to perform a color analysis of the peptide to ensure it meets industry color standards before being packaged and shipped to the customer.
4. Implementing a label reading functionality to each camera to keep track of each vial and its data as it moves through the system.

Bibliography

[1] S. Daffre, P. Bulet, A. Spisni, L. Ehret-Sabatier, E. G. Rodrigues and L. R. Travassos, "Bioactive Natural Peptides," *Studies in Natural Chemistry Products,* vol. 35, pp. 597-691, 2008.

[2] T. Kimmerlin and D. Seebach, "'100 years of peptide synthesis': ligation methods for peptide and protien synthesis with applications to b-peptide assemblies," *The Journal of Peptide Research,* vol. 65, no. 2, pp. 229-260, 2005.

[3] G. Laverty, J. Rafferty, H. Nagaraj, A. P. McCloskey, R. Huwaitat, S. Porter and A. Albadr, "Peptide Therapeutics and the Pharmaceutical Industry: Barriers Encountered Translating fro Laboratory to Patients," *Current Medicinal Chemistry,* vol. 23, no. 35, 2016.

[4] R. B. Merrifield, "Solid PHase Peptide Synthesis. I. The Synthesis of a Tetrapeptide," *Journal of the American Chemical Society,* vol. 85, pp. 2149-2154, 1963.

[5] S. A. Rahman, A. El-Kafrawy, A. Hattaba and M. F. Anwer, "Optimization of Solid_Phase Synthesis of Difficult Peptide Sequences via Comparison between Different Improved Approaches," *Amino Acids,* vol. 33, pp. 531-536, 2007.

[6] C. W. Coley, D. A. Thomass III, J. A. M. Lummiss, J. N. Jaworski, C. P. Breen, V. Schultz, T. Hart, J. S. Fishman, L. Rogers, H. Gao, R. W. Hicklin, P. P. Plehiers, J. Byington and J. Piotti, "A robotic platform for flow synthesis of organic compouns informed by AI Planning," *Science,* vol. 365, 2019.

[7] A. J. Mijalis, D. A. Thomas III, M. D. Simon, A. Adamo, R. Beaumont, K. F. Jensen and B. L. Pentelute, "A fully automated flow-based approach for accelearted peptide synthesis," *Nature Chemical Biology,* vol. 13, pp. 464-468, 2017.

[8] D. J. Alder, J. A. Herkap, J. R. Wiesler and S. B. Williams, "Life cycle cost and benefits of process automation in bulk pharmaceuticals," *ISA Transaction,* vol. 34, pp. 133-139, 1995.

[9] L. G. Roberts, *Machine Perception of Three-Dimensional Objects,* Massachusetts Institute of Technology, 1963.

[10] T. S. Huang, "Computer Vision: Evolution and Promise," 1996.

[11] D. Faggella, "Computber Vision Applications - Shopping Driving, and More," Emerj, 14 March 2020. [Online]. Available: https://emerj.com/ai-sector-overviews/computer-vision-applications-shopping-driving-and-more/. [Accessed 4 June 2020].

[12] AIA, "Computer Vision vs. Machine Vision," Vision Online, 16 Jaunary 2014. [Online]. Available: https://www.visiononline.org/vision-resources-details.cfm/vision-resources/Computer-Vision-vs-Machine-Vision/content_id/4585. [Accessed 4 June 2020].

[13] Christopher, "The role of Machine Vision in manufacturing," Medium, 7 January 2020. [Online]. Available: https://medium.com/technology-innovations-insights/the-role-of-machine-vision-in-manufacturing-97a0a0ad81df. [Accessed 4 June 2020].

[14] "Mytide Therapeutics," Mytide Therapeutics, 2020. [Online]. Available: https://www.mytide.io. [Accessed 4 June 2020].

[15] R. C. Gonzalez and R. E. Woods, Digital image processing, Upper Saddle River, New Jersey: Prentics Hall, 2008.

[16] K. Pulli, A. Baksheev, K. Kornyakov and V. Eruhimov, "Realtime Computer Vision with OpenCV," *Queue,* vol. 10, no. 4, pp. 40-56, 2012.

[17] J. Canny, "A Comutational Approach to Edge Detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence,* Vols. PAMI-8, no. 6, pp. 679-698, 1986.

[18] Y. Luo and R. Duraiswami, "Canny edge detecion on NVIDIA CUDA," in *IEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, Anchorage, AK, 2008.

[19] OpenCV Library, "Canny Edge Detection," [Online]. Available: https://docs.opencv.org/trunk/da/d22/tutorial_py_canny.html. [Accessed 14 June 2020].

[20] L. G. Shapiro and G. C. Stockman, in *Computer Vision*, New Jersey, Prentice-Hall, 2001, pp. 279-325.

[21] M. Sezgin and B. Sankur, "Survey over image thresholding techniques and quantitative performance evaluation," *Journal of Electronic Imaging,* vol. 1, no. 13, pp. 146-165, 2004.

[22] "Operations in Arrays," OpenCV Library, [Online]. Available: https://docs.opencv.org/3.4/d2/de8/group__core__array.html#ga48af0ab51e36436c5d04340e036ce981. [Accessed 14 June 2020].

[23] "Thresholding Operations using inRange," OpenCV Library, [Online]. Available: https://docs.opencv.org/3.4/da/d97/tutorial_threshold_inRange.html. [Accessed 14 June 2020].

[24] V. Karathanassi, C. Iossifidis and D. Rokos, "Application of machine vision techniques in the quality control of pharmaceutical solutions," *Computers in Industry,* vol. 32, pp. 169-179, 1996.

[25] Cognex, "Industry 4.0 and Machien Vision: Industry 4.0, or "The Industrial Internet of Things," will rely on machine visionto revolutionize industrial automation.," Cognex, [Online]. Available: https://www.cognex.com/what-is/industry-4-0-machine-vision. [Accessed 17 July 2020].

[26] S. Anand and L. Priya, A Guide for Machine Vision in Quality Control, New York, NY: CRC Press, 2019.

[27] A. Mehle, M. Bukovec, B. Likar and D. Tomazevic, "Print registration for automated visual inspection of trasnparent pharmaceutical capsules," *Machine Vision and Applications,* vol. 2017, pp. 1087-1102, 2016.

[28] M. Bahaghighat, M. Mirfattahi, L. Akbari and M. Babaie, "Designing Quality Control System Based on Vision Inspection in Pharmaceutical Product Lines," in *International Converence on Computing, Mathematics, and Engineering Technologies*, Sukkur, Pakistan, 2018.