

# Leonardo Da Robot

Authors: Eric Chang: Electrical and Computer Engineering, Carnegie Mellon University  
 Christopher Bayley: Electrical and Computer Engineering, Carnegie Mellon University  
 Harsh Yallapantula: Electrical and Computer Engineering, Carnegie Mellon University

**Abstract**—A system capable of taking a digital image and painting a watercolor version on a physical canvas. The aim of the project is to create an image which looks naturally painted, not a replica of the source image. At minimum, this product will be able to represent simple images or match the shape of the given image. The resulting image should be similar in terms of color gradients and general appearance and contrast. The system will consist of a 2-dimensional gantry system to physically draw the picture along with a software component that will handle image processing.

**Index Terms**— Gantry, Image Processing, Painting, Robot

## I. INTRODUCTION

ONE of the biggest goals in robotics is to create systems that behave more human-like. A large amount of research today is focused on creating robotic systems that replicate human tasks such as driving, speech comprehension and vision. Our intent is to break down this barrier between humans and robots in art. Our design is inspired by the gantry systems of 3D printers, which we combine with image processing and control systems to make a robot that can receive an image and paint it on a canvas. The most difficult part of this is to make the painting look natural and not robotic, while also maintaining accuracy and speed. This project does not aim to replicate how a printer makes an image pixel-by-pixel. To accomplish this task, we paint the picture using a paint brush and human-like strokes. Our approach takes the proven efficiency and accuracy of 3D printers to allow us to paint a pre-processed image using smooth and natural strokes

Critical to our design is the ability to paint an image which is an accurate representation of the source image. This is measured using the structural similarity (SSIM) index, where we aim for a score of at least 0.2. Additionally, the total time for painting an image should be reasonable to the scale and complexity of the image, limited to 8 hours in the worst case. These metrics are key for our project to meet our goals of connecting robotic systems with art through a design which produces a painted image of good quality and can do so in a reasonable amount of time.

## II. DESIGN REQUIREMENTS

The first requirement is that a digital input image of any size is capable of being rendered as the target image to paint and displayed back to the user. This step allows for an image to be rejected if the rendered painting is not of acceptable quality. This requirement is purely digital in nature and therefore can be

tested using an image bank of 10 images which vary in image size and complexity. The successful design will be able to create renders of consistent quality across all scales and complexities, and will scale images which exceed the bounds of the painting space. This requirement will simply be measured by the program's success at processing the given input image.

Our next requirement is that we effectively use the full range of the palette which is integrated into the system. When painting an image the palette color with the lowest difference to the desired color should be chosen, measured using the HSV (hue, saturation, value) color model. Specifically, this color difference will be measured using the following equation:

$$\Delta = (|H_0 - H_r| + f * 0.1 * |S_0 - S_r| + f * 0.15 * |V_0 - V_r|) \quad (1)$$

In this equation,  $\Delta$  is the color difference,  $H_0$  is the hue of the original color from the image,  $H_r$  is the hue of the color chosen by the robot,  $S_0$  is the saturation of the original color from the image,  $S_r$  is the saturation of the color chosen by the robot,  $V_0$  is the value of the original color from the image,  $V_r$  is the value of the color chosen by the robot, and  $f$  is the weighting factor for brown. In our experimentation of color picking, we found several difficulties and realized it was a much more complicated process than initially expected. Hue, saturation, and value can not be weighted equally when determining the closest color. As a result, we choose to weight saturation and value less than hue by multiplying them with the constants of 0.1 and 0.15 respectively. In addition, we introduced  $f$ , a weighting factor for brown. Unlike other colors, brown is not determined by a unique hue. Rather, its hue is orange, and it is defined by having a low value. As a result, when the normalized hue of the color is between 0 and 0.125, which is a region of orange-like hues, we define  $f$  to be 1.5, while normally it is 1, in order to more heavily weight value and saturation to more accurately pick out brown from orange. Another important consideration was that a hue of 0 and 1 both corresponded to red, as HSV values are depicted as a cylinder which wrap around. This was accounted for when writing the code that performed this algorithm and calculation. Finally, white, black, and gray are tested for initially before (1) is used, as those colors occur at extreme values of saturation and value.

This color requirement can be tested using an image of our own palette; the pigment in the palette image should always be painted with that pigment. The image of our palette we will use is shown in Fig. 1, which will be used to verify the accuracy of our color picking.



Fig. 1. An image of our palette which will be used to test our robot's color selection.

Our third requirement is that our product will create paintings that are visually similar to the original image. We require that our output painted image has a score of at least 0.2 according to the structural similarity (SSIM) index. This is a metric which is used to measure the perceived similarity of an image to the source and is commonly used for film and television [4]. We use this metric for this reason as it corresponds closely to the goal of our project. After testing the SSIM for various images, we found that professionally done water painting images of customer supplied images was roughly 0.4 on average. We will test our performance using an image bank of 10 images, shown in Fig. 2, which features images which grow in complexity. Using this bank, we will paint the images and test that the SSIM score is at least 0.2 for the first 8 images. The final two images are more complicated and are our stretch goals once we reach the scores for the first 8, so they will not necessarily be tested.

Our final requirement is that our design operates in a reasonable amount of time as a function of image size and complexity, which we define using the following equation:

$$t = (1 + \alpha) * (\beta * s) \tag{2}$$

In this equation,  $t$  is the estimated time,  $\alpha$  is a measure of image complexity, measured as  $4 * (\text{uncompressed image size} / \text{JPEG compressed size})$ ,  $s$  is the size of the painted image in square inches, and  $\beta$  is a constant scaling factor of 4 (translating to 4 minutes to paint each square inch). The equation for  $\alpha$  follows the logic that JPEG compression uses DCT-II coefficients to compress the image from a raw uncompressed size of the total image size and dimensions, meaning that the JPEG compressed size of an image gives rough estimate of how much information is contained in an image [5]. The image size painted can be variable, but is limited to 7.5x10 inches, and from experimental testing  $\alpha$  can be up to 0.70 for a very complex landscape image, and as low as 0.04 for a single line. Equation (2) gives a rough estimate of the upper limit of the testing time, and is designed to have a limit of roughly 8 hours for the largest and most complex image. We will test our design using the same image bank from Fig. 2, which contains variably sized images as well, and will ensure that they all print in a time corresponding to (2).

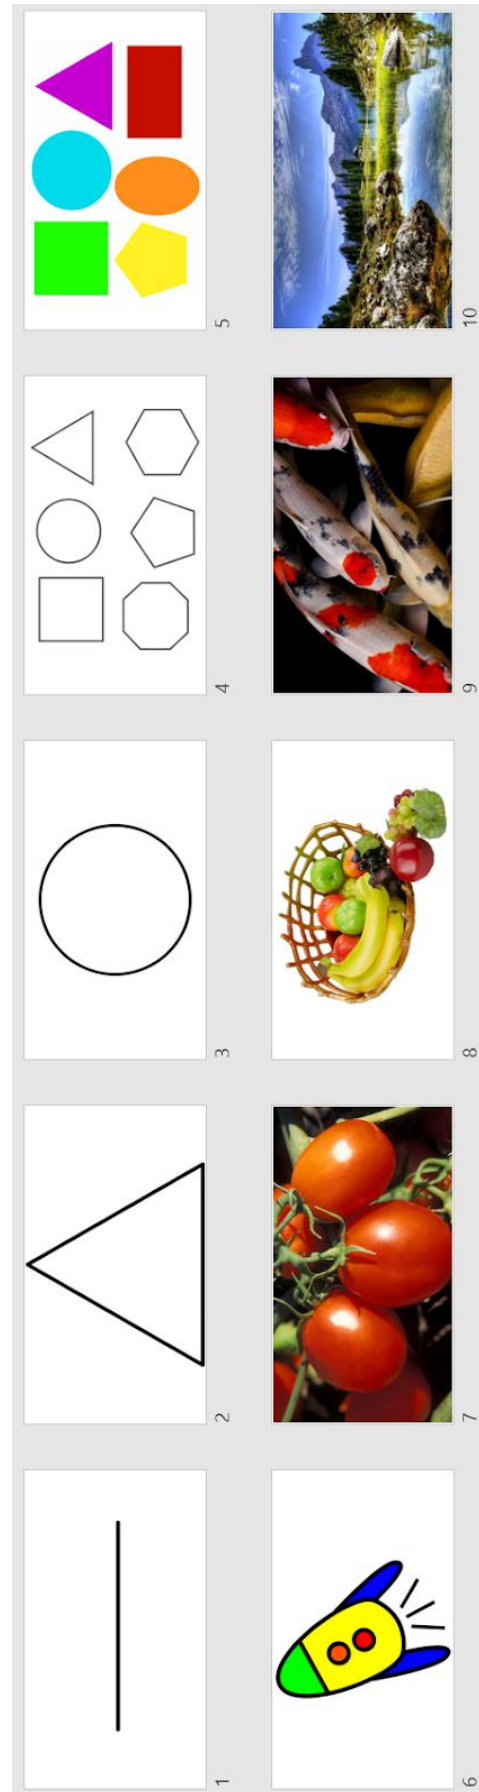


Fig. 2. The test image bank we will use in testing.

### III. ARCHITECTURE AND/OR PRINCIPLE OF OPERATION

Our project will primarily be split into a physical hardware component and a software image processing component. The overall architecture of our project is represented in the block diagram depicted in Fig. 3. The physical component will be the robot itself, which includes the base and frame that supports the 2D gantry system. The gantry system is supported by a frame that lies on the base of the robot. The gantry system’s purpose is to move a painting head around the XY coordinate space above the paper, water, and palette which are lying on the base. The painting head will control a paintbrush which is attached by a servomotor, and the head will be moved around by two stepper motors. These motors will be controlled by a Raspberry Pi. The servomotor is connected to the Raspberry Pi through a GPIO pin. In our original design, the stepper motors were controlled through a motor shield attached to the Pi, but in our final design these were controlled by stepper motor drivers in a cnc shield attached to an Arduino. The Pi sends commands to the Arduino which then controls the motors.

The software component of our project centers on using image processing to convert a digital input image into instructions for our motors in order to successfully paint an image. A monitor will be connected to the Raspberry Pi which will allow the user to submit the digital image they wish to be painted. The monitor will also be used to display the result of the image segmentation process to the user, in order to show the user what the approximate final result of the painting will be. The input image will go through our image processing algorithm, which will first modify the image to be easier to paint, in the process of image segmentation. This process will reduce details and cluster colors together to create a simpler image for painting, which is what will be displayed to the user on the monitor. The image segmentation process will output

data and information such as edges and colors, which will then be used in our stroke creation algorithm. This process will generate the sequence of strokes that our robot will need to perform in order to paint the image, with information such as color, length, and direction encoded.

This sequence of strokes will be given to the stroke routine, which will control the motors. A list of strokes with a specified color and location are provided to the stroke routine, and the routine will instruct the motors of our robot to dip the brush into the water, paint, and finally onto the paper. Our motor control routines are the bridge between the software containing information of what to paint and the hardware which is capable of painting. The interface from the lowest level of motor drivers allow the stepper motors to rotate a set number of steps in a set delay time between each step. This functionality is provided through the use of an open source CNC control software called GRBL. Communication between the Pi and the Arduino is over the Arduino’s serial port, and GCode formatted instructions are sent which contain destinations for the X and Y components along with a set speed. This is an improvement over our original design review, which involved controlling the motors through a motor shield using code provided by the manufacturer. This change in the system gave us greater control of the movement of the motors for smooth acceleration and deceleration, and allowed us to provide the necessary power at times to overcome friction in our physical system. The servomotor is controlled through PWM, and its driver has an interface that allows the motor to be set to a certain angle. These interfaces are used by the control routines to abstract to higher level concepts such as dipping the brush to grab pigment or performing a stroke at set locations. These control routines receive input from the painting and stroke generation algorithms.

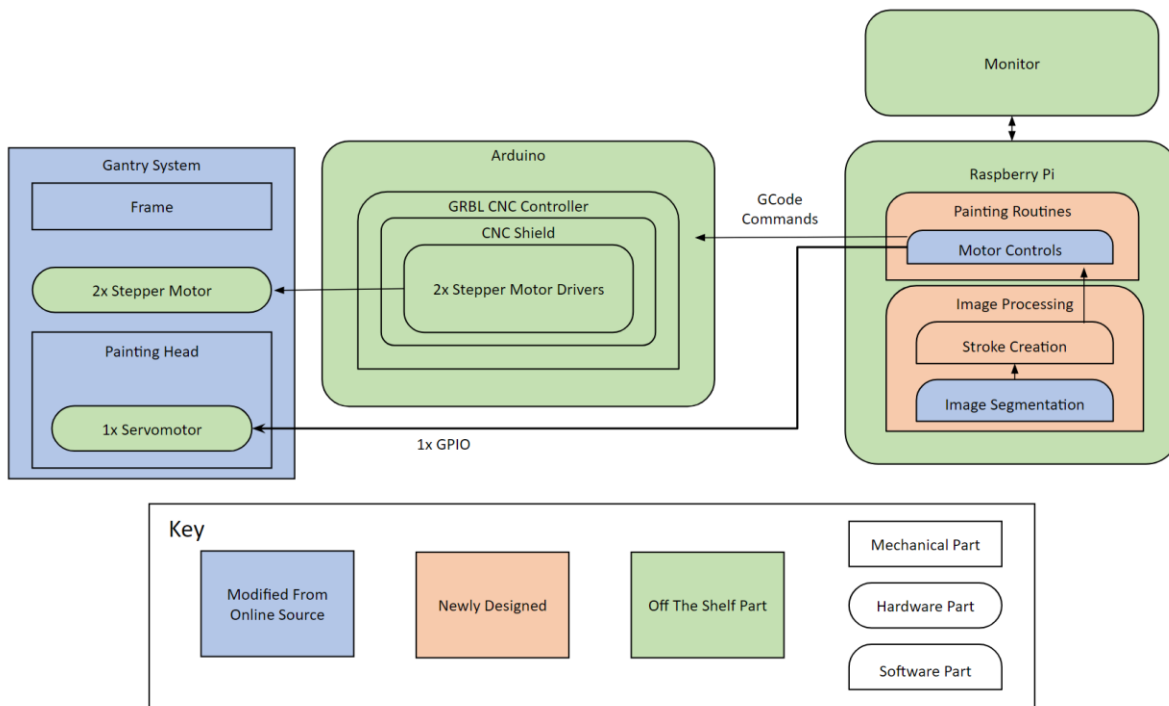


Fig. 3. The block diagram showing our project’s architecture and subsystems.

#### IV. DESIGN TRADE STUDIES

The four main requirements of our project were successful image rendering, accurate color picking from the available palette, visual similarity between the painted image and the original input, and a reasonable time taken for painting. We evaluated our design using the tests described earlier in the paper, and compared and analyzed these results against our expectations. These details will be outlined more comprehensively below in a section dedicated to each specification and metric.

In designing our project, there were numerous trade-offs and options to consider when finalizing decisions. Two of the most important aspects of our design were the 2D gantry system, which is the primary physical and hardware component of our robot, and the image segmentation algorithm, which is the first step of our software component. A significant amount of testing, research, and experimentation was conducted in order to choose the best option to proceed with in these two subsystems. More information about this is included below, in two sections centered around the design decisions we made when creating the gantry system and choosing the image segmentation algorithm.

##### A. Specification – Successful Rendering

Our first specification was that our robot could, given any input image of any size, successfully process and render the image and display the result back to the user. This specification was not explicitly tested with a strict methodology, as any issues arising from a failed render would be quickly fixed. However, by the end of the project, we have used numerous images as input to our program, including all ten images in our image test bank, and all can be successfully rendered. As a result, we believe our program’s ability to process any input is fairly robust and comprehensive.

##### B. Specification – Color Picking

The next requirement was accurately choosing colors for each segmented region in the input image. Table I shows the results of a test involving the image of our palette from Fig. 1. As Table I shows, our color picking algorithm was not completely accurate. This was due to several complications which were mentioned previously; black, gray, white, red, and brown have unusual or extreme values in their HSV representation, which is why (1) went through many revisions. Ultimately, the algorithm we created satisfied us by being accurate in many scenarios. Although there were some colors which were not fully correct, as in the exact same pigment from the palette was not chosen for itself, all the incorrect colors were close in terms of being the same hue. Because of the similarity in many of the hues on our palette, for example there are several reds and blues on our palette which all look very visually similar, we did not think it was necessary to dedicate more time and effort to refining the algorithm even more.

##### C. Specification – Visual Similarity

The next specification of our system was that each painted image is visually similar to the original input image. As previously stated, this was tested using the structural similarity

TABLE I. CORRECTNESS OF OUR COLOR PICKING ALGORITHM AS TESTED BY AN IMAGE OF OUR PALETTE

Category	Number of Colors
Correct	15
Close	9
Incorrect	0

(SSIM) index. This requirement was tested against our image bank of ten images, with the results shown in Table II. A small description of each image is included for reference. Images 9 and 10 were not tested due to concerns over complexity and time, as those two images were originally stretch goals as well. As can be seen from the results in the table, the first 8 images in our image bank all had an SSIM higher than 0.2, which was the threshold we determined when planning our tests and metrics. In fact, the results of the structural similarity tests were much higher than we expected, with every image in the first 8 having an SSIM of above 0.4. The results of some select images from the test bank are shown in Fig. 4, which displays our robot’s drawings of a rocket (image 6), tomatoes (image 7), and a fruit basket (image 8).

##### D. Specification – Time Taken

The final specification of our system was that it operated in a reasonable amount of time. Again referring to the results in Table II, the time it takes for the robot to paint an image increases as the complexity and size of the image increases. In addition, no image so far has taken over eight hours to paint, which was the maximum amount of time we wished our robot to perform in. However, images 9 and 10 are highly detailed and complex, and may possibly have taken over eight hours to paint. This was one of the reasons why we did not test images 9 and 10, due to the interest of time.

TABLE II. STRUCTURAL SIMILARITY AND TIME TAKEN OF IMAGES IN OUR TEST BANK

Image Description	Image Number	SSIM	Time Taken
Line	1	0.683	2 min 6 s
Triangle	2	0.655	5 min 14 s
Circle	3	0.691	5 min 3 s
Shapes	4	0.465	8 min 29 s
Colored Shapes	5	0.436	42 min 11 s
Rocket	6	0.404	37 min 40 s
Tomatoes	7	0.434	5 hrs 56 min 17 s
Fruit Basket	8	0.552	2 hrs 32 min 34 s



Fig. 4. The robot’s drawings of a rocket, tomatoes, and a fruit basket. Original images are images 6, 7, and 8 in the test image bank respectively.



### E. Subsystem - 2D Gantry System

There were several possible implementation paths available for designing the 2D gantry, and all would meet the major constraints of our project which are having a large enough workspace and a fine enough granularity. Because of this, the main considerations made when choosing the gantry design were complexity, cost, and risk. Our original plan was to use threaded rods which would support the carriage, allowing rotation of the threads to carry the device. We found that these designs were generally used for much larger constructions, as it can support a much heavier payload. For this reason, we chose to use a belt-based design as is commonly used in 3D printers, as these designs are commonly used for lighter loads and there are plenty of reference implementations available in the form of 3D printers.

When considering possible belt-based gantry designs, we found some of the most common designs are the Hbot system (depicted in Fig. 5), CoreXY (depicted in Fig. 6), and the Ultimaker style [2]. The Hbot uses a single very long belt arranged in the shape of an H, but the force applied is uneven and results in a moment created on the edge of the print head. The CoreXY system solves this problem using two belts arranged in an H shape which cross near the top. This design is fairly complicated, and the exact arrangement of the belts and pulleys required leaves a lot of room for error. Based on these factors, we chose the design used by the Ultimaker printers. This design uses several sets of belts which run in one direction and is almost entirely symmetric, making it simpler to implement. Additionally, Ultimaker is open source, allowing us to leverage their proven system with small modifications to work in our own design. We selected the gantry specification based on the simplicity of the Ultimaker style and the low risk due to the open-source nature of Ultimaker.

The final aspect of our design decision with respect to the gantry was to use existing parts or design our own. By designing and 3D printing our own parts, we allow for a more custom design, fast part availability, and an overall lower cost than purchasing the Ultimaker parts. The drawbacks are that this allows room for error by modifying an existing design which is known to work well. The Ultimaker parts were not available

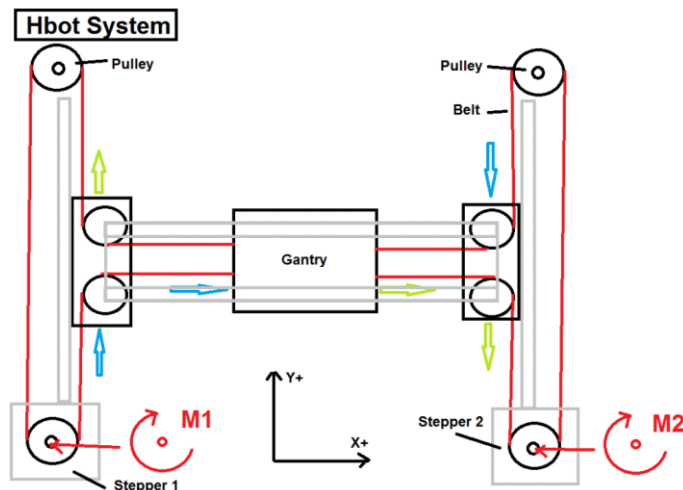


Fig. 5. The Hbot gantry design.

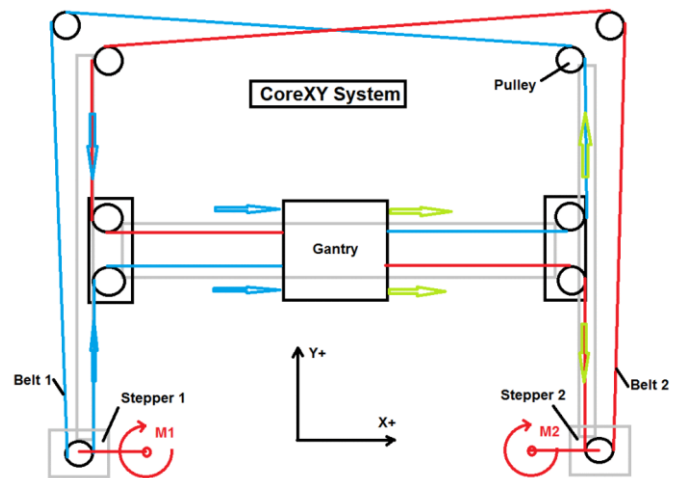


Fig. 6. The CoreXY gantry design.

from any major retailers in the US, and additionally the replacement part packs were expensive and contained many parts we wouldn't need. For these reasons we chose to design our own parts and have them 3D printed. This design choice has allowed us to save over \$100 of our budget, and allowed us to expand the working space of the head considerably. Inevitably the print head would have to be a custom part as well, and by modifying all of the other parts we reduced the number of constraints on the geometry of the print head which allowed us to pursue a clean and simple design.

Overall, our gantry design decision diverged first between a threaded rod or belt design, between several different belt designs, and finally on using custom or premade parts. We made the choice of a belt-based design for its prevalence in an area which is very similar to our goal, allowing us to keep risk low and manage complexity by using well documented existing designs. The Ultimaker design was chosen among other belt-based designs again for its simpler design and proven effectiveness, which greatly lower our risk. Finally, the decision to modify the parts came at the price of a small increase in risk for the benefit of a significantly less constrained design space and lower cost.

### F. Subsystem - Image Segmentation Algorithm

There were a few methods we considered for implementing an algorithm to pre-process the image. The reason that we needed to pre-process was because of the limitations of the physical system. The brush width must stay constant throughout the entire painting process, and thus we cannot draw something thinner than this width. The second limitation is that we have 24 colors available. This means that all colors have to be approximated to those 24 colors. The third limitation is time. Since this is a physical system where motors have to move everything around, there are time constraints. To tackle these constraints, we considered the following pre-processing algorithms: blurring, edge drawing, k-means image segmentation, and mean shift segmentation. After some consideration, we decided to discard the idea of blurring. The original image used to test all of these image segmentation algorithms, depicting a fruit basket, is shown in Fig. 7. The

result of blurring it is shown in Fig. 8.

Although blurring removes the detailed parts of the image, it still has gradients, which are very difficult to produce. It also becomes difficult to see where one object ends and another begins. Therefore, we decided to further explore drawing only the edges and image segmentation. To choose between these three methods, we calculated the accuracy of the output along with how complex the output image was. For accuracy, we looked at the output image and used the Structural Similarity Index to classify the methods. For complexity, we took the output JPEG file and compared it with the original JPEG file in terms of file size. As stated earlier, the JPEG file format compression is correlated to how complex an image is, so the more compressed the output file is, the less complex it is [5]. Table III shows the complexity and similarity for the three methods.

The conclusions we can draw are that the k-means clustering and mean shift clustering are far better than edge drawing in both accuracy and complexity. Mean shift clustering is 27% less complex than k-means clustering but still manages to beat out k-means in terms of accuracy. It also returns objects of defined edges and uniform color, which is much easier for the physical system to draw than what k-means returns. The following images show the output of edge detection, k-means clustering, and mean shift segmentation. Fig. 9 shows the output of the original image after running edge detection, Fig. 10 shows the result of running k-means, and Fig. 11 shows the result of running mean shift segmentation. All these images use the original fruit basket image in Fig. 7 as input.

TABLE III. STRUCTURAL SIMILARITY AND COMPLEXITY RATIO OF DIFFERENT IMAGE SEGMENTATION ALGORITHMS

	SSIM	Complexity Ratio
<b>Edge detection</b>	0.0335	1.8982
<b>Mean shift</b>	0.9432	0.7654
<b>k-means</b>	0.9352	1.0471



Fig. 7. The original image of a fruit basket used in our image processing experimentation.



Fig. 8. The result of blurring the image in Fig. 7.




Fig. 9. The result of running edge detection on the image in Fig. 7.



Fig. 10. The result of running k-means on the image in Fig. 7.



Fig. 11. The result of running mean shift segmentation on the image in Fig. 7.

## V. SYSTEM DESCRIPTION

Our project is split into three main subsystems: the 2D gantry system, which is the physical component that controls movement of the brush; the gantry control layer, which controls the motors of the gantry; and image processing and stroke generation, which is the software component of our project. An image of our robot as a whole can be seen in Fig. 12.

### A. 2D Gantry System

The gantry system forms the physical portion of the project. This is a system of shafts, pulleys, and motors which allow the robot to dip a paintbrush in water, collect pigment from a palette, and perform a stroke on the paper. The design for our gantry follows the gantry design used in the Ultimaker line of 3D printers, which is proven to be effective and accurate. Ultimaker is also entirely open source, and we were able to modify their CAD files as well as design our own which match our design more precisely. Our design does not use the Ultimaker parts, although they are available for purchase online through 3D printer repair sites. This is because by modifying the Ultimaker designs or creating our own we are able to keep costs low by 3D printing these parts, as well as make custom parts which match our design requirements more precisely. For example, the carriages were redesigned to fit an axes arrangement which was simpler for us to execute, as well as to carry the print head with the crossbar shafts in a different arrangement than the original design.

The core of the gantry design is outlined in Fig. 13. Two stepper motors are located at one corner of the gantry and can be positioned anywhere in the z plane below the axes. Two sets of parallel axes are positioned to form a square, mounted in a bearing allowing the shafts to rotate easily. Both the stepper motors and the axes have a pulley around them, and are connected by a timing belt, labelled purple. This translates the rotation of the stepper motor to the axes. Each set of parallel




Fig. 12. An image displaying the main portion of our robot.




Fig. 13. The Ultimaker gantry design.

axes are connected by two belts as well, labelled red and blue, which coordinate their rotations. A carriage rests on the axes and is attached to the belts, which allows the rotational motion of the opposite axes to be transferred into translational motion on the carriage. The carriages carry shafts which support a paint head in the center space, allowing the head to be coordinated in two dimensions. The paint head carries a servomotor which has a brush connected, allowing the brush to be raised from the page or lowered to make contact. This is suspended via a frame of T-slotted aluminum, allowing the paint head to travel a fixed distance above the base. The implementation of our design is shown in Fig. 14.

The major constraints on this system are that the brush head have a movable space which is large enough and that the brush can be controlled with sufficient precision. The working space must fit our canvas, 4x6 inches, our watercolor palette, about 3.5x8.5 inches, and a cup of water, about 2x2 inches. This totals a working space requirement of about 8.5x8.5 inches. Our design consists of a frame of aluminum with inner dimensions of 16x16 inches. The entire inner dimension is not usable however, due to the space required within the frame for




Fig. 14. A CAD model of our gantry system.

the axes to run uninterrupted. Considering the space required for the axes, as well as the pulleys and carriages mounted on the axes, the dimensions of the working space are 14x14 inches. Then this meets our requirement for the brush to be able to move over a sufficiently large area. This also leaves enough room for us to add a mechanism to stabilize the position of the paint head. By surrounding the shafts of the gantry with plastic tubing, we can further restrict the area the paint head can travel within. This allows us to place zero points near the components on the board which are frequently visited, and by driving the paint head into these tubes we can effectively guarantee the current position of the paint head. For our design with no feedback mechanisms, this design choice was critical to creating a system which could account for error. Additionally, the mounting height of the entire gantry is adjustable, but fixed during operation. This allows us to modify the height to fine tune the amount of contact the brush makes with the paper as it is rotated by the servomotor. This is the advantage of using T-slotted aluminum as well as our own designed mounts for the axes.

To address the precision allowed by this design, we must consider the precision of the stepper motors and the dimensions of the pulleys and timing belts. The stepper motors step size translates to a rotation of the axes by the same amount, as they are coupled by a timing belt. The rotation of the axes will result in a translational movement of the belt which carries the carriage according to the degree change and the radius of the pulley. From this we can form the equation  $x = 2\pi r * (\theta / 360)$ , where  $x$  is the translation of the carriage,  $r$  is the radius of the pulley, and  $\theta$  is the change in angle of the axes. The pulley used in the Ultimaker, which we use in our design with a small change to the inner radius, has a radius of 0.25 in. Our stepper motors have 200 steps / revolution, which translates to a 1.8 degree step size. Of all the cost-effective stepper motors we considered, this step size was standard. This gives us a minimum horizontal translation  $x = 2\pi * (0.25 \text{ in}) * (1.8 / 360) \approx 0.008 \text{ in.}$  or 0.2 mm. This means we can operate the paint head with a granularity of 0.008 inches, which is more than sufficient to carry out all the requisite strokes and operations.

The integration of the gantry into our overall design is at the motors. The two stepper motors are connected to CNC shield mounted on the Arduino, and the servomotor are connected directly to the Raspberry Pi. Low level drivers are used for basic control of these devices, allowing their rotations to be orchestrated into an operating gantry system.

The design of our gantry did not change much throughout our project, as we successfully leveraged the working designs of the Ultimaker. The major parts which were revised were the 3D printed parts, in order to allow for easier assembly and operation. The carriages were redesigned multiple times to create something which could connect easier than the Ultimaker design and be much more sturdy. Additionally, the pulleys were redesigned to allow for 4 screws rather than one for fixing onto the shafts. This allowed us to perfectly center the pulleys and reduce the amount of wobble in our gantry.

## B. Gantry Control Layer

The gantry control layer consists of the motors that will control the 2D gantry system as well as the software in the Raspberry Pi that controls the motors. The code for this subsystem will be written in Python, and will use libraries to help interface the motors with the Pi. The servomotor, which controls the painting head, will be wired to the Raspberry Pi through a GPIO pin. We will be using the library `gpiotermio` to interface with the servomotor, specifically the `AngularServo` class, which extends the `Servo` class [6]. This class allows control of a rotational PWM-controlled servomotor, and gives us the ability to set it to specific angles. For the library to work, we must set the servo to its maximum position and its minimum position and measure its angles, and input these angles as the constructor for an instance of the class. This setup will let us move the servomotor to any angle in between.

The two stepper motors are controlled by stepper motor drivers, A4988, which are mounted into a CNC shield on the Arduino. This allows for precise control of the motors through the open source CNC software GRBL. The drivers allow us to precisely control the power which goes to the motors, and have a sufficiently high current rating to allow for the friction in the system to be overcome. The painting routine algorithm generates GCode instructions which are sent to the Arduino over the serial port. The GCode sent is very simple, with lines consisting solely of X and Y positioning arguments and a speed. During initialization of the gantry control layer, GRBL is configured to match the rotations needed to translate a single inch.

The original design used a motor shield which connected directly to the Pi. While this design gave us access to the motors directly, it proved ineffective at supplying power as there was no way to set current limits to the motors. By using stepper motor drivers we were able to ensure that the motors would always receive sufficient power and not draw a dangerous amount. Additionally the GRBL software allowed smoother accelerations and decelerations which allowed for smoother control of the gantry.

With this, an interface was created for use by the painting routines. The gantry control layer will act as the intermediary between the software and hardware components of our project, allowing our code to easily call functions that will move the motors how we desire. This layer exposes a single interface which is the stroke routine, which can receive a list of xy coordinates which define a series of straight line segments and a number corresponding to one of the available colors in the palette. This function will then paint the entire stroke and return when completed. Doing so will require the paint head be moved to wash the brush, grab pigment from the palette, and trace the line segments on the page. It may be necessary to gather more water and more pigment while painting a single stroke, which is the responsibility of this layer to control. The layer above provides no information on how frequently to collect pigment. In addition to this single routine, a software model of how this layer should behave was created. This layer exposes an API which is exactly the same as the gantry control layer, but instead of sending the strokes to the gantry to be executed, they are



instead used to create an image of the lines in the stroke. This model is used to show an in-progress view of the painting, and also allowed for easy debugging as we had a comparison for exactly how the gantry should behave.

In creating the stroke routine which is exposed to the layer above, routines for washing the brush and collecting a specific pigment are required. These routines are also a part of the motor control layer but are not exposed to the layer above. These are used internally by the stroke routine when it is necessary to wet the brush and collect pigment. As the locations of all of the needed objects to carry out these routines are fixed, much of the control routines will be moving to hard-coded locations, not operating by some feedback informing the head where the objects are. For this to work this layer must also always keep track of the precise location of where the head is. An additional challenge is that the coordinate systems used by the image must match the coordinate system used by the control layer. To address this the layer also exposes an initialization routine which will receive information from the above layer regarding the image size and desired output size, in inches. This information allows the future routines to normalize the pixels received and translate to its own coordinate system which is rooted in the physical space rather than a digital space.

This layer bridges the gap between the hardware and software in our design. Built up from the provided device drivers for the motors, this layer offers a single control routine for painting a desired stroke with a specified color. This abstraction allows the painting to be easily performed from the above layer after a list of strokes representing the image has been created.

### C. Image Processing and Stroke Generation

The Image Processing and Stroke Generation algorithms form the software component of this project. The first part is the Image Processing algorithm. This converts the original digital image into something that the physical apparatus can draw. The reason this is required is that the physical part of the robot has 3 main constraints: brush width, number of colors, and time. The width of the brush is constant from the beginning to the end of the painting process, since we won't swap out brushes during a painting. We are using 24 colors, so the robot will have to select the closest color to the 24 colors. There is also a time limitation. Therefore an image clustering algorithm called Mean Shift Segmentation was used. This algorithm takes regions of similar color that are close together and turns them into regions of uniform color. Three parameters are used to toggle the properties of the segmentation algorithm. The first parameter is the distance from the center that a point can be such that it is still able to be part of that segment. This is called the spatial radius. The second parameter is the range radius, which specifies the range of colors that can be in the segment. The third parameter is the minimum density of points, which indicates the number of pixels that can be inside one segment, thereby dictating the size of the largest segment. These parameters ensure that the objects are wide enough to be drawn by the brush. This algorithm reduces the complexity of the image significantly (about 24%) while keeping the accuracy

close to the original. The second benefit is that the algorithm outputs a map with the labeled regions for each pixel, which allows us to have objects with defined edges and uniform color. This makes it simple for the next part of the software component, namely the Stroke Generation algorithm.

The Mean Shift Segmentation algorithm is an open source function implemented in the `pymeanshift` library [8]. By changing the parameters, we are able to segment an image in whichever way the problem requires. Fig. 15 shows the result of running light segmentation while Fig. 16 shows the result of running stronger segmentation.

The second part of the software component is Stroke Generation. Once the segmented image has been created, the segments are separated into lists of strokes. There are two different kinds of strokes: perimeter strokes and fill-in strokes. Each object in the segmented image is made up of one perimeter stroke and one fill-in stroke. The perimeter strokes trace the outline of an object with the color of that object. The fill-in strokes fill in the object with horizontal straight lines of that color. First the perimeter stroke is drawn for one object, and then that object is filled in with fill-in strokes. The order of objects to be drawn is based on the length of the perimeter stroke, to make sure that the lowest detailed objects are drawn first. A perimeter stroke is made up of several very small line segments which trace the outline of the object. A fill-in stroke is made up of several horizontal line segments which go from the top of the object to the bottom.

Once the strokes have been created and ordered, they are broken up into line segments which are defined by a starting



Fig. 15. The result of light segmentation on the image in Fig. 7.



Fig. 16. The result of stronger segmentation on the image in Fig. 7.

coordinate, ending coordinate, and color.

The color chosen from a region uses a weighted cost function of HSV values. Some special colors, white, gray and black, have special logic related to the saturation and the value HSV fields. A weighted sum allows us to place more emphasis on the hue of the color when relevant, as this is the most important factor in matching the color. The color choosing algorithm also considers the neighbors of a region, and is biased towards choosing colors which are different than its neighbors. This allows the rendered image to maintain the contrast present in the original image. The algorithm will consider using the second best color match over the best color match when the best match is used by a neighbor. The second best color will only be used if the loss of color accuracy is below some threshold.

The order the regions are drawn in mostly follow the order that they are created by default using the library. However, we made one small optimization. When testing, we noticed that some black outlines would cover up certain colored segments due to it being such a dark color. As a result, we make the program paint all the black regions of the painting first before painting the rest.

## VI. PROJECT MANAGEMENT

### A. *Schedule*

The schedule for the project is shown on page 13. Each task is color coded based on which team member primarily worked on it, and lines between tasks indicate which tasks depended on the completion of others. Our schedule did not change significantly from the version presented in the design report. The major changes are that some tasks were broken up into smaller parts, and the time taken for certain tasks was longer than initially predicted. Nevertheless, we met all of our major deadlines and remained mostly on track for the duration of the project. We had plenty of time remaining in the days up to the final demo to test, debug, improve, and modify the system.

### B. *Team Member Responsibilities*

In the first few weeks, all three of us focused on the physical portion of the design. Since none of us have had any mechanical design experience before, we have decided to frontload that part of the project. Chris has taken the lead on the mechanical design and oversaw building the physical system. Harsh focused on the software portion which involved the image processing. Eric has taken charge of the control layer of the robot, which includes the hardware and control of the motors. Although each focused on an individual subsystem, each team member also helped out and discussed decisions regarding other portions of the project whenever applicable and relevant. As the schedule shows, we all sharing some responsibilities at the end.

### C. *Budget*

The budget consisted of ordered parts and 3D printed parts. Refer to the Bill of Materials on page 14.

### D. *Risk Management*

The major risks to our project were being able to construct a reliable gantry system, being able to accurately control the system, being able to paint images in a reasonable amount of time, and consistent paint quality for all colors used.

In order to mitigate the risk of an inconsistent gantry design we drew heavily from the Ultimaker design. We were able to replicate the principles of operation from their design, and modify it to fit our design. These modifications consisted of custom length rods, custom frames, custom paint head, custom pulleys, and custom carriages. After construction we saw that our design was not as reliable as we would have liked, and we were able to limit this through iterative redesigns of all of our custom parts and a very fine attention to detail to reduce any sources of friction or non-uniform movements. This problem was also minimized by front-loading the construction of the gantry, giving us ample time to improve the physical design. From our experience, hardware problems are much more difficult to identify and fix than software issues, and this strategy gave us lots of time to fix all of our hardware issues.

We were able to limit the risk of how we operated our physical system by using GRBL, an open source software library for controlling CNCs. This library takes the responsibility of tracking the current head location and smooth motor controls. By using this library rather than creating this

software ourselves we limit the error in operation of the gantry, so that the only source of error comes from inconsistencies in the physical gantry itself.

In order to reduce the amount of time taken we were able to set different operating modes of the gantry. When painting the image it moves slower and more carefully to ensure the lines look correct in the image, but moves much faster when moving to gather pigment or rinse the brush. This allows the system to operate in a reasonable amount of time without sacrificing quality of the painted image. Additionally, the strokes which are used to paint the image have been optimized to reduce the distance travelled by the paint head when performing infill strokes, which greatly reduces the total operation time.

Finally to ensure consistent paint quality across different colors we allowed for a custom number of palette swishes based on color. For lighter colors, like yellow or orange, the system will swish the paint much more than darker colors like black or gray. This ensures that the amount of pigment placed on the paper is roughly equal across all colors and that no color will dominate the resulting image.

## VII. RELATED WORK

The goal of having an independent robot independently paint an image is not unique to our project; this is the same goal of the Robot Art Competition and Exhibition [1]. This is a competition held annually for robotics enthusiasts to submit images which were painted by their creations. From this competition and the gallery of designs they provide we were able to see how many of the best robots performed and how they were designed. A stand out in the competition is a man named Pindar Van Arman and his CloudPainter, which is an AI driven robot which has both a robotic arm based design and a gantry based design. His design is so effective that he has built a brand and sells the paintings his robot creates for a large sum. Most other designs featured on the Robot Art site use gantry-based designs, including another watercolor painting robot. These designs were our main source of confirmation that our idea for a gantry-based painting robot could be a success.

## VIII. SUMMARY

Overall, our system was able to successfully meet all our design specifications. Visual similarity, which was the specification with the biggest weight when dealing with painting images, was passed with flying colors with SSIM much higher than the threshold that we originally set. We were able to successfully render all input images and the painting operated in a reasonable amount of time. Color selection was not perfectly accurate, but we felt it was accurate enough and also valued contrast and aesthetic appeal along with having accurate colors.

### A. Future work

One thing we want to work on is removing the downward lines that are created after each stroke. We think that this is a purely software problem, and we would be able to reduce this effect given more time. The further optimizations require making physical changes. One major problem we faced during the entire project was that the gantry system had large non-uniform resistance during movements. We pinpointed the problem to the fact that the pulleys were too large for the shaft, leading to them not being centered on the shaft which caused the belts to stretch and contract with each rotation. To mitigate this problem in the limited time we had, we printed new pulleys with screws on all 4 sides to be able to center the pulley on the shaft. A final solution would involve obtaining pulleys which fit perfectly on the shaft. These would preferably be CNC'd, since metal parts would be more accurate and wear down slower. We also want to replace our bearings since they are low quality and don't eliminate friction well. Another problem we had was that we didn't have any feedback. We had no idea whether the painting head went to the right position or if it drew the right stroke. We got around this by zeroing after every movement. However, a final solution would involve using a camera to track the painting head, and encoders to track motor movement. Color blending could also be added to more perfectly match the input image, but for our project during this semester this was out of scope.


### B. Lessons Learned

The biggest lesson we learned was that mechanical problems are much harder to fix than software problems. Software bugs can be fixed on a computer in a few hours, whereas mechanical problems require making the right measurements, waiting days for parts or 3D prints, and then fitting them onto the entire setup. Another lesson we anticipated and confirmed was that one should never try to recreate something that has already been created. To create our gantry system, we used a design inspired by 3D printers. We looked at the Ultimaker printer to understand its design and modeled a large part of our 2D gantry around it. Since we knew that this design already worked, we were confident that our gantry would function as well. This was highly helpful because none of our team members had much experience in mechanical design. We also front loaded the tasks we weren't familiar with (mechanical design) to make sure we would have enough time for debugging. This turned out to be a good scheduling decision as we had enough time to learn from our mistakes and fix our design. This would definitely be a good lesson for us in future projects. One last thing was that we had several unexpected problems towards the end of the project, and we had to think quickly to come up with solutions. One should always expect the unexpected and account for it in design decisions. We had decided to 3D print most of our parts at the beginning, so when parts failed or were causing problems, we were able to reprint parts quickly. This was very helpful the weekend before demo day since we needed new pulleys.

## REFERENCES

- [1] <https://robotart.org/>
- [2] <https://maxdesign1990.wordpress.com/2016/05/22/gmtech-printer-motion-platform-research/> - gantry designs
- [3] <https://imiloainf.wordpress.com/2012/06/13/mean-shift-segmentation/> - Mean Shift Segmentation
- [4] <http://www.imatest.com/docs/ssim/> - SSIM info
- [5] <https://en.wikipedia.org/wiki/JPEG> - JPEG Compressions algorithm
- [6] [https://gpiozero.readthedocs.io/en/stable/api\\_output.html](https://gpiozero.readthedocs.io/en/stable/api_output.html) - gpiozero servo library
- [7] <https://github.com/sbcshop/MotorShield> - MotorShield library
- [8] <https://github.com/fjean/pymeanshift> - Pymeanshift library





Bill of Materials

Part name	Quantity	Extra Quantity	Total Quantity	Price/unit	Base Price	Shipping/Tax	Total Price	Notes
RPI	1	0	1	\$34.49	\$34.49	\$0.00	\$34.49	
RPI Motor Shield	1	1	2	\$14.99	\$29.98	\$0.00	\$29.98	
RPI Power Adapter	1	0	1	\$9.99	\$9.99	\$0.00	\$9.99	
MicroSD with Adapter	1	0	1	\$7.99	\$7.99	\$0.00	\$7.99	
T-slotted Aluminum Beams	96	32	128	\$0.23	\$29.44	\$7.80	\$37.24	Price per inch, shipping/tax price is for 4 cuts at \$1.95 each
Beam Caps	8	2	10	\$1.20	\$12.00		\$12.00	Shipping/tax included in Beam T-Nut and Screw
Cap Push-In Fastener	8	2	10	\$0.13	\$1.30		\$1.30	Shipping/tax included in Beam T-Nut and Screw
Beam Corner Brackets	8	2	10	\$2.75	\$27.50		\$27.50	Shipping/tax included in Beam T-Nut and Screw
Beam T-Nut and Screw	8	2	10	\$0.46	\$4.60	\$22.57	\$27.17	Shipping/tax is for all of the 80/20 parts
Beam T-Nut and Screw (for bearing mounts)	16	0	16	\$0.46	\$7.36	\$0.00	\$7.36	Identical as above part, ordered separately
Beam T-Nut and Screw (for bearing mounts)	1	0	1	\$8.99	\$8.99	\$0.00	\$8.99	Part is same as above, incorrect amount ordered first time
Stepper Motor	2	1	3	\$9.63	\$28.88	\$0.00	\$28.88	Nema 17 (from Ultimaker 2) (pack of 3 + cables)
Servomotor	1	1	2	\$3.95	\$7.89	\$0.00	\$7.89	Pack of 2
Servomotor extension cable	1	0	1	\$6.99	\$6.99	\$0.00	\$6.99	Pack of 12, 320mm
Jumper Cables	1	0	1	\$5.99	\$5.99		\$5.99	Pack of 80
12v 1A DC Power Supply	1	0	1	\$8.95	\$8.95	\$8.94	\$17.89	Shipping and tax includes cost for following part
Female DC Power Adapter	1	0	1	\$2.00	\$2.00		\$2.00	
Paint Set	1	0	1	\$10.29	\$10.29	\$0.00	\$10.29	
Base Board	1	0	1	\$8.29	\$8.29	\$6.85	\$15.14	
Stepper Mounter Mounts	1	0	1	\$10.99	\$10.99	\$0.00	\$10.99	One pack has 3 mounts and screws
Shafts	3	1	4	\$4.67	\$18.68	\$7.12	\$25.80	Each shaft will be cut into 2
Pack of Bearings	1	0	1	\$24.95	\$24.95	\$4.98	\$29.93	Contains 10 Bearings. 8 needed 2 extra
Motor Belts	1	0	1	\$6.99	\$6.99	\$0.00	\$6.99	2 belts per pack
Belts	1	0	1	\$8.88	\$8.88	\$0.00	\$8.88	5m of belt
Belt torsion springs	1	0	1	\$5.00	\$5.00		\$5.00	Pack of 10
Watercolor paper	1	1	2	\$7.19	\$14.38		\$14.38	Pack of 25 sheets
Arduino	1	1	2	\$20.69	\$20.69		\$20.69	
Stepper motor drivers	1	1	2	\$8.11	\$8.11		\$8.11	
CNC Shield	1	1	2	\$6.79	\$6.79		\$6.79	
<b>Total Cost</b>							<b>\$436.64</b>	

Total Cost

<b>Hardware Parts</b>	<b>\$436.64</b>
<b>3D Printed Parts</b>	<b>\$66.00</b>
<b>Total Cost</b>	<b>\$502.64</b>
<b>Budget Remaining</b>	<b>\$97.36</b>