

GTGraffiti: Spray Painting Graffiti Art from Human Painting Motions with a Cable Driven Parallel Robot

Gerry Chen, Sereym Baek, Juan-Diego Florez, Wanli Qian, Sang-won Leigh, Seth Hutchinson, and Frank Dellaert

Abstract— We present GTGraffiti, a graffiti painting system from Georgia Tech that tackles challenges in art, hardware, and human-robot collaboration. The problem of painting graffiti in a human style is particularly challenging and requires a system-level approach because the robotics and art must be designed around each other. The robot must be highly dynamic over a large workspace while the artist must work within the robot’s limitations. Our approach consists of three stages: artwork capture, robot hardware, and planning & control. We use motion capture to capture collaborator painting motions which are then composed and processed into a time-varying linear feedback controller for a cable-driven parallel robot (CDPR) to execute. In this work, we will describe the capturing process, the design and construction of a purpose-built CDPR, and the software for turning an artist’s vision into control commands. Our work represents an important step towards faithfully recreating human graffiti artwork by demonstrating that we can reproduce artist motions up to 2m/s and 20m/s^2 within 9.3mm RMSE to paint artworks.

I. INTRODUCTION AND RELATED WORK

Spray painting graffiti art in a human style is an important, open problem that requires a systems approach. In this paper, we take the first step towards creating a system that can capture human graffiti artwork and collaborate with artists to create new and copied artworks in the public settings that define graffiti. In addition to the well-established sociological motivations for reproducing graffiti art [1], robot art is intrinsically motivating for its marriage of art and technology. By possessing physical abilities beyond those of its collaborating artists, a graffiti robot could reveal new artistic avenues highlighting human-robot collaboration for disabled [2] and able-bodied artists alike. To act as the hand of an artist poses inherently interconnected problems in art, hardware, control, and human-robot interaction. Furthermore, the technology required to produce the large-scale, dynamic motions required for graffiti has applications in warehouse/industrial logistics [3], agriculture [4], construction [5], [6], and motion simulation [7]. Creating graffiti art with a robot requires (1) capturing the motions of artists painting, (2) creating a robot that can achieve comparable motions to human artists, and (3) implementing algorithms that would allow the robot to execute on the artists’ visions. Despite considerable progress in each of these tasks, to our knowledge, no system has been demonstrated to achieve all three.

This work was supported by the NSF under Grant No. 2008302
Institute for Robotics and Intelligent Machines, College of Computing,
Georgia Institute of Technology, Atlanta, GA, {gchen328, sereymbaek,
jdfe3, wqian39, seth, fd27}@gatech.edu



Fig. 1. Our system captures artist painting motions of individual letter outlines which are composed and processed into controls for a cable robot to execute. Our system produced this painting of the letters “ATL” (Atlanta).

Prior work exists in capturing graffiti art, most notably the Graffiti Markup Language (GML) project [8]. Although the project has been successful in generating a large library of graffiti artwork, almost all the data was captured from digital interfaces (e.g. stylus) rather than full-body painting motions. This is problematic because an artist’s creative process may differ between virtual and physical mediums and because ignoring the physical painting motions neglects the challenge of generating robot trajectories. An exception is the GML Recording Machine [9], though it captures only 2 degree of freedom (DoF) planar motion.

Robots developed for spray painting have seen considerably more attention. Serial arm manipulators and gantry-based systems are precise and mature, but arms do not scale well to large workspaces [10], [11] and gantry-based systems exhibit a tradeoff between size and portability [12], [13]. Mobile manipulators address these issues, but are currently not as dynamic or precise as human artists [14]. Aerial robots are popular for their ability to paint otherwise inaccessible walls, but have been cited as being difficult to accurately control due to susceptibility to disturbances and comparatively limited acceleration capabilities [15]–[18]. Cable-based systems appear to be promising, but so far [2], [19] have only demonstrated raster- or stippling- style painting while [20] has not demonstrated the highly dynamic motions employed by human artists.

Finally, despite prior research in robot control and artistic composition, the software to enable graffiti painting does not currently exist. CDPR control (further discussed in Section IV-A.3) is relatively well understood, but has not been demonstrated for dynamic graffiti trajectories. Research on industrial painting robots has thoroughly studied paint dispersion and trajectory generation, but is primarily concerned with uniform coats on curved surfaces in contrast to graffiti

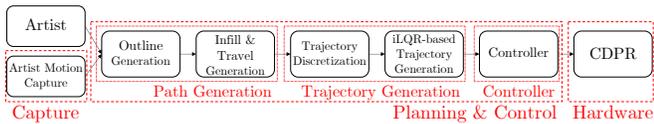


Fig. 2. This system overview depicts the capture, hardware, and planning & control components of our system.

art’s non-uniform coats on flat surfaces [21]–[24]. Berio is notable for his research in graffiti composition and stylization [11], [25]–[27], but focuses on digital rendering as opposed to producing trajectories.

We argue that the problem of creating graffiti artwork is sufficiently expansive and its components codependent that it requires a system-level approach. In this work, we propose a novel system towards creating graffiti artwork by improving and coordinating the capture, hardware, and software requirements. Fig. 1 depicts an example result of the GTGraffiti system summarized in Fig. 2. Our contributions include:

- **capturing** a library of 6DoF trajectories for creating graffiti artwork using motion capture (mocap),
- designing, building, and testing the **hardware** for a purpose-made robot platform to paint graffiti,
- proposing a **planning and control** pipeline to realize high-level artistic descriptions into motor torque commands, and
- demonstrating a system that can paint human-style graffiti artwork.

II. CAPTURE

The capture process is important for both learning the motions to produce graffiti art and establishing robot capability requirements. As such, our capture process is focused on obtaining the most artistically meaningful data while omitting less relevant data. In this work, we collect a library of simple, composable shape outlines.

A. Design Considerations

We first capture artwork using an OptiTrack™ mocap system for its simplicity and accuracy. Mocap systems have the advantage of directly outputting positions and orientations of rigid bodies with sub-millimeter accuracy which trivializes the process of obtaining the 6D trajectories of a spray paint can during painting. As we will discuss in Section IV, the processing and rendering components of our pipeline can optionally use other forms of captured artwork such as Scalable Vector Graphics (SVG) and GML files in addition to mocap data. Conversely, later sections will also discuss aspects of the artist motions captured with mocap which are beyond the robot’s capabilities (e.g. 6DoF vs planar motion).

We opt to capture only the outlines of shapes and omit the infills because, according to an artist collaborator, the particular pattern used to fill-in a shape is largely arbitrary and algorithmically generating one does not significantly detract from artistic value. Furthermore, the easiest infill path for a human may not be the easiest for a robot.



Fig. 3. Mocap setup for capturing painting motions of collaborating artists.

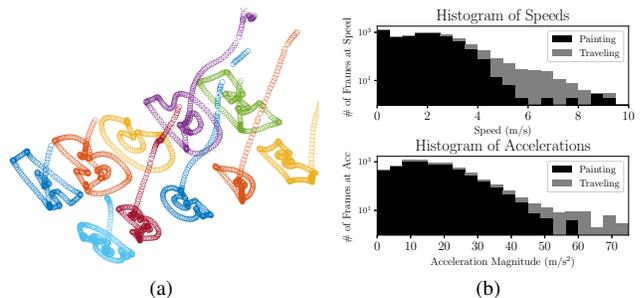


Fig. 4. (a) Captured shape outlines for the letters “ABCDE” by our two collaborating artists. (b) Speeds and accelerations of a human graffiti artist during painting which help inform minimum requirements for the robot.

B. Approach

We collected the full 6D trajectories of the spray cans and painting surfaces (plywood sheets) as two graffiti artist collaborators painted. Four mocap position markers were affixed each to the can (as shown in Fig. 3) and painting surface to extract the 6DoF poses for each time step at 120Hz. For each art collaborator, the 26 letters of the English alphabet were captured along with special symbols such as punctuation marks and small doodles (e.g. skull). We convert the pose of the spray can nozzle to the painting surface’s reference frame using coordinate transformations. Additional details can be found in our accompanying arXiv paper [28].

To determine when the spray nozzle is being depressed (painting vs traveling motions), we applied a number of heuristics for each candidate painting motion segment including distance between start and end points (assuming outlines are closed curves), maximum speed, arc length, non-maximum suppression, and manual annotations.

C. Results

Fig. 4a shows an example of our data by plotting the spray can nozzle translations. The fact that the data was collected while the artists were physically painting combined with the accuracy and 6DoF of motion capture gives our data the potential to better understand the nuanced motions of human graffiti painting e.g. biomechanically and with respect to can speeds, distances, and orientations.

We also study the speeds and accelerations reached by our collaborating artists during painting and traveling to inform the requirements of the cable robot, reasoning that a robot capable of these speeds and accelerations should be capable of reproducing the artworks. From Fig. 4b we estimate a speed of 6m/s and acceleration of 50m/s² to be sufficient to paint human-style graffiti artwork.

D. Discussion & Limitations

Although motion capture’s accuracy is unparalleled, there are several drawbacks. Most notably, cost and setup hinder the accessibility and mobility of capture systems. We were only able to capture motions in a controlled laboratory setting which limits the realism of the artwork. Additionally, human artists can move so fast that, even at 120Hz, our mocap system misses some detail.

The ability to capture nozzle actuation was also limited since we were unable to directly record actuation force which artists use to allow better paint control. During data collection, an additional marker was actually placed on the tip of the artist’s index finger to aid in identifying when the spray can nozzle was depressed. Upon analyzing the data, however, this was not found to be a consistent method of annotating binary nozzle actuation let alone actuation force. Even with various heuristics, manual annotation was needed to correct misclassifications in nozzle actuation.

Additional tags, characters, and full murals with photographs will be added to the library in the future.

III. ROBOT HARDWARE

Given the design requirements for painting graffiti based on human spray painting data, we believe that a CDPR is an ideal platform. In this section we detail our robot hardware.

A. Design Considerations

The primary design requirements for a graffiti painting system involve workspace size, maximum end-effector velocity, and maximum end-effector acceleration. We seek a platform which can be scaled to a workspace 20mx20m or larger, though in this work we only seek a demonstration sized at a few meters. Based on the analysis presented in Section II-C, we determined that we require 6m/s and 50m/s² of speed and acceleration, respectively. Assuming the mass of the spray can and actuating accessories do not exceed 2kg, including gravity the robot should be capable of exerting 120N upward and comparable forces in other directions.

Secondary design requirements include portability, accuracy, and stiffness. It should be feasible to disassemble and reassemble the robot on-site at the wall of a building. Accuracy and stiffness are considered secondary constraints because, compared to art forms such as brush painting or sculpture, graffiti is less sensitive to positional inaccuracies and experiences less reaction force. Based on the thickness of a line painted with a “needle” nozzle 5cm from the painting surface, we estimate 2.5cm of repeatability to be sufficient. We estimate an accuracy of 1% the size of the painting to be sufficient, based loosely on [29]. We estimate external disturbances to be negligible based on paint reaction forces and historical Atlanta wind speeds.

CDPRs present ideal platforms for graffiti painting given the aforementioned requirements. A CDPR is a robot whose end-effector is pulled by a set of cables which are driven by winches on a fixed base. Due to properties of cables, CDPRs can scale to extraordinary sizes and speeds [30], [31], albeit with reduced stiffness. These qualities make them

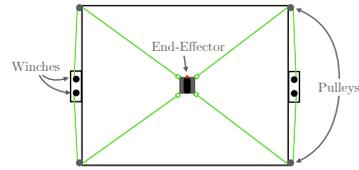


Fig. 5. Our planar CDPR has a 4-cable, rectangular configuration with the end-effector in the center carrying the spray paint can.



Fig. 6. Our cable robot (left) includes an end effector that carries the spray paint and actuator electronics (center) and 4 winch assemblies, each consisting of a shared motor controller, motor, and helical winch (right, x2).

ideally suited to the large but relatively undisturbed setting and modest accuracy requirement of graffiti painting.

CDPRs also have an active research community which has solved many challenges in workspace analysis [32]–[34], control, and estimation (further discussed in Section IV-A.3). Preliminarily, based on [3], we estimate a 1kHz update frequency to be necessary for real-time control.

Finally, we define the requirements to actuate the spray can nozzle. For a full can of Montana BLACK 400mL, the force required to depress the nozzle was measured to be 27N and the displacement was measured to be 2mm. Other 400mL spray cans by the brands Montana, Hardcore, and Kobra were found to have similar actuation forces and displacements.

B. Approach

Our CDPR uses 4 cables in a planar configuration to exert pulling forces on the end effector via 4 motor-driven winches (see Fig. 5). The end effector was built to be lightweight and carry the spray can and actuating electronics. It has 4 mounting points to connect to the 4 CDPR cables. The cables are low-stretch and lightweight (cable sag is negligible). The spray can nozzle actuating mechanism is wireless, battery-powered, and implemented using a servo with the lever-arm mechanism from [35]. Complete design details can be found in our accompanying arXiv paper [28].

C. Results

Our assembled robot is pictured in Figures 1 and 6. The winches satisfy our design requirements with each being capable of pulling a 2kg mass on the cable up to 7.6m/s and 94m/s² and bidirectionally communicating at 1kHz. The end-effector and spray can actuating mechanism are also pictured in Fig. 6. The total mass varied between 1006g and 1317g depending on the spray can. The spray can actuating mechanism was able to successfully depress the spray nozzle 100% of the time in a trial of 100, 1 second long actuations. The latency from commanding to dispensing paint was measured to be 400ms.

D. Discussion & Limitations

We are able to paint well, as in Fig. 1, despite not being able to use our 6DoF captured data to its full potential since we are limited to planar motion. Also, we will discuss in Section IV-C.2 that the paint limits us to a maximum speed far below what the hardware is capable of (and the speed suggested by Fig. 4b). A combination of hardware upgrades and intelligent paint modeling and optimization are likely necessary to leverage our system’s full potential, especially actuation to move the nozzle *closer to the canvas*.

IV. PLANNING AND CONTROL

From an artist’s input, we must control the robot to paint. We use a hierarchical approach with 3 levels:

- 1) Path Generation: turn the artist’s vision into a mathematical description
- 2) Trajectory Generation: find a trajectory within the robot’s capabilities while respecting the artist’s vision to the maximum extent possible
- 3) CDPR Control: execute the trajectory online

The interplay between the trajectory generation and CDPR control merits a summarized precursor explanation for clarity. During the trajectory generation phase, the optimal control problem of tracking a desired trajectory is solved *offline*. The iterative Linear Quadratic Regulator (iLQR) method [36] – iterating by applying the linearized system and control law forward in time to obtain a new linearized feedback law – is used to solve the optimal control problem. The feedback law from the final pass can then be used as the online controller.

A. Design Considerations

1) *Path Generation*: In this work, the artist composes artworks using the shapes in the shape library as a first step towards more general artistic descriptions. Given an artist’s specification for the placements of shapes from the library of captured art, we seek to generate the *paths*, in the form of Bézier curves, that the spray can must follow. This is a system-level problem because it requires suitable captured data, well-planned and modeled robot capabilities, and clear artist desires. We divide path generation into (1) outline, primarily a human-robot interaction problem, (2) infill, a coverage path planning problem, and (3) travel, a problem of filling in discontinuities. The latter two are unique to our system approach because, recalling the reasoning from Section II, only outlines are captured for the shape library.

2) *Trajectory Generation*: To create a physically realizable trajectory that is as similar as possible to the artist’s vision, we first discretize the path at 100Hz to obtain a direct-from-artist trajectory, \mathbf{x}_d (within speed and acceleration limits), then apply an offline iLQR-based optimization to obtain a smoothed reference trajectory, \mathbf{x}_{ff} , control signal, \mathbf{u}_{ff} , time-varying feedback gains, K , and paint timing.

Loosely inspired by [26], the iLQR-based optimization is used to strike a balance between the artist’s intent and the ease of controlling the robot. We express the iLQR problem

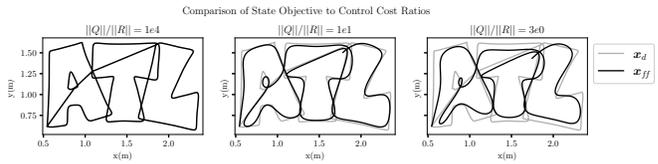


Fig. 7. Stylization from iLQR Q vs R also observed by [26]

in discretized form with time index k as:

$$\mathbf{x}_{ff}, \mathbf{u}_{ff} = \arg \min_{\mathbf{x}, \mathbf{u}} \sum_{k=0}^T \tilde{\mathbf{x}}[k]^T Q \tilde{\mathbf{x}}[k] + \tilde{\mathbf{u}}[k]^T R \tilde{\mathbf{u}}[k] \quad (1a)$$

$$\text{s.t.} \quad \mathbf{x}[k+1] = f(\mathbf{x}[k], \mathbf{u}[k]), \quad (1b)$$

$$\mathbf{x}[0] = \mathbf{x}_0 \quad (1c)$$

where $\tilde{\mathbf{x}}[k] := \mathbf{x}[k] - \mathbf{x}_d[k]$ is the deviation of the state $\mathbf{x}[k]$ from $\mathbf{x}_d[k]$: the artist’s intended trajectory, $\tilde{\mathbf{u}}[k] = \mathbf{u}[k] - \mathbf{u}_m$ is the deviation of the control $\mathbf{u}[k]$ from \mathbf{u}_m : the average of the minimum and maximum allowable torques [37], [38], \mathbf{x}_{ff} and \mathbf{u}_{ff} are the smoothed reference (nominal/feedforward) state and control signals, Q and R are the state objective and control cost matrices, $f(\mathbf{x}[k], \mathbf{u}[k])$ defines the nonlinear system dynamics, and \mathbf{x}_0 is the initial state. The state consists of the cartesian position and velocity: $\mathbf{x} = [p_x \ p_y \ \dot{p}_x \ \dot{p}_y]$, where \mathbf{p} denotes the position of the spray can’s nozzle. The control \mathbf{u} consists of the four motor torques. Section IV-C.3 experimentally justifies why orientation is omitted.

Intuitively, the state objective matrix, Q , advocates for the artist and the control cost matrix, R , penalizes being near torque limits. Interestingly, as discovered by [26], the relationship between Q and R can also be interpreted as an artistic parameter as visually depicted in Fig. 7.

3) *Control*: We seek a controller that can control the cable robot to achieve motions comparable to a graffiti artist (requirements are the same as in Section III-A).

Our cable robot controller is inspired by prior works. CDPR control places emphasis on “tension distribution” (TD) which is analogous to redundancy resolution in serial manipulators [3], [7], [37], [39]–[43]. These approaches typically use or assume a feedback controller whose control variable is a task space wrench. The TD algorithm then computes the optimal motor torques (or cable tensions) required to achieve the desired task space wrench.

However, since we are using an iLQR-based optimizer, which produces locally optimal control laws as described in Section IV-A.2, most aspects of control (including TD) have been shifted offline. Our online controller is then a simple linear feedback controller, as depicted in Fig. 8. Mathematically, our controller can be expressed as:

$$\boldsymbol{\tau} = K(t)(\mathbf{x}_{ff}(t) - \hat{\mathbf{x}}) + \mathbf{u}_{ff}(t) \quad (2)$$

where $\boldsymbol{\tau}$ is the 4-vector of motor torques, $\hat{\mathbf{x}}$ is the measured state, and $K(t)$ is the 4x4 time-varying gain matrix.

To compute $\hat{\mathbf{x}}$, we need to estimate the position and velocity of the spray can well enough to achieve our repeatability

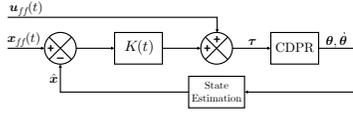


Fig. 8. CDPR controller block diagram, where u_{ff} , x_{ff} , and K are pre-computed offline using an iLQR-based optimizer implemented in GTSAM.

and accuracy requirements. Nonlinear least squares solvers are commonly used for CDPR state estimation, but we will show that a simpler solution is sufficient for our application.

B. Approach

1) *Path Generation*: When generating the outline, the largest challenge involves communicating the artist’s intent and robot’s capabilities between the artist and computer. In this work, we apply constraints to the artist when they are specifying their artistic vision. To constrain the canvas size, we use a rectangular approximation of the wrench feasible workspace (WFW) [33], [34] to define the space in which the artist may place library objects. To constrain layering specifications, we impose a strict layering of shapes such that each shape is either entirely on top of or entirely beneath another. Due to our robot’s planar limitation, we project the nozzle position for each frame from the mocap data onto the painting surface to form an ordered set of 2D line segments (allowing us to retain velocity information to be used during trajectory generation). Data sources other than our mocap library can be used but require velocity data.

To infill the shapes we apply an exact cellular decomposition and use a standard horizontal “zigzag” path within each cell [44]. Similar to the way each artist chooses a strategy for infilling based on personal preference (according to our artist collaborators), we choose this pattern for its ease of implementation and actuation: in most configurations our robot has the best control authority horizontally. Further details on the way we decompose the infill and compute the exhaustive walk to reduce nozzle actuation cycles are provided in our accompanying arXiv paper [28]. For each object, we paint the infill in the face color then the outline in the outline color. Instead of applying a hidden line removal algorithm [45], we simply finish painting each object before the next is started.

Finally, travel strokes must be added to make the path continuous in position. Although making the paths continuous in direction is an option (as in [20]), we opt to allow discontinuous directions but enforce continuous velocity in the trajectory generation stage. For every pair of strokes, we add a straight line from the end of the previous stroke to the start of the next stroke if they are not already coincident.

2) *Trajectory Generation*: Unique to the system-level approach, we discretize outlines and infill/travel strokes differently due to the different ways the paths were generated.

For outlines, we have velocities from mocap data so we need to apply time-scaling. We compute the original path’s speed and acceleration using finite differences assuming each line segment takes 1s, then apply the linear transformation,

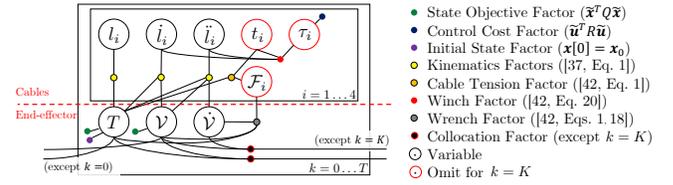


Fig. 9. Factor graph depicting the iLQR problem using plate notation. Circles represent variables to be solved while dots represent objectives or equations. l_i , \dot{l}_i , and \ddot{l}_i represent cable length, speed, and acceleration respectively. t_i , τ_i , and \mathcal{F}_i represent cable tension, motor torque, and the wrench on the end-effector caused by cable i , respectively. T , \mathcal{V} , and $\dot{\mathcal{V}}$ represent end-effector pose, twist, and twist-acceleration, respectively.

$t' = ct$, to match a predefined maximum speed and/or maximum acceleration and sample $x(t')$ from the path at 100Hz. We choose limits of 1.2m/s and 20m/s² based on the spray paint dispersion described further in Section IV-C.

For the infills and travel strokes, we need to generate rest-to-rest trajectories with continuous velocities. We choose trapezoidal velocity profiles for their popularity in industrial applications [46] with limits of 0.5m/s and 20m/s² (based on spray paint dispersion).

The iLQR-based optimization of (1) is performed offline using factor graphs and the GTSAM software library, but any iLQR implementation can be used. The system dynamics constraints (1b) are drawn from prior works in CDPR control including the standard equations for kinematics and cable tension/wrench equilibrium [37], winch model dynamics [3], [42], rigid body dynamics [42], and numerical integration [47], [48]. The iLQR problem is then expressed as the factor graph [49]–[51] shown in Fig. 9 and solved with the GTSAM software library using Levenberg-Marquardt and variable elimination. The solution gives x_{ff} and u_{ff} , while the Bayes Net obtained by the final iteration’s elimination step contains the locally optimal feedback gain matrix, K . Because we compute x_{ff} , u_{ff} offline, we do not need a full workspace analysis since we can check offline if u_{ff} is within the control limits. Full details on the equations and factor graph are available in our accompanying arXiv paper [28].

3) *Control*: First we interpolate K , x_{ff} , and u_{ff} since the trajectory generation phase was discretized at 100Hz while the controller runs online at 1kHz. K and u_{ff} are interpolated using a zero-order hold, while x_{ff} is interpolated using a first-order extrapolation from the most recent discrete x_{ff} .

To estimate position we discard redundant information and to estimate velocity we use a linear least squares solution. We discard the bottom two cable lengths then apply trigonometry using the top two cable lengths to estimate the spray can position assuming the spray can is always vertical. To estimate velocity, we solve the linear least squares problem:

$$\dot{p} = \arg \min_{\dot{p}} \left\| \dot{l} - W^T \dot{p} \right\|_2^2 = W^+ \dot{l} \text{ where } W \text{ is the wrench matrix [42] and } \cdot^+ \text{ is the Moore-Penrose left inverse.}$$

We also employ an *offline* calibration whereby a rectangular trajectory is run while collecting mocap and robot log data. A nonlinear least squares optimization is used to compute pulley locations and coefficients for cable length scaling which are hard-coded for subsequent runs.

C. Results

1) *Path Generation*: Fig. 10 shows an example result.

2) *Trajectory Generation*: The speed and acceleration limits were tuned for our painting distance of around 12cm. The outline limits of 1.2m/s and 20m/s² were tested using the Montana “Skinny Cap Beige” nozzle and the infill limits of 0.5m/s and 20m/s² were tested using the Montana “flat jet cap wide” nozzle. Faster speeds resulted in incomplete coverage while slower speeds resulted in “dripping”.

The offline iLQR-based optimization has $\mathcal{O}(n)$ complexity and runs at approximately half real-time (e.g. a 1-minute trajectory takes 2 minutes to optimize). We chose $Q = \text{diag}([1e4, 1e4, 0, 0])$ and $R = I_{4 \times 4}$ as a balance between tracking accuracy and stability.

3) *Control*: To evaluate our control stage, we use mocap for ground truth data and log \mathbf{x}_{ff} and $\hat{\mathbf{x}}$ from the robot at 100Hz for a challenging 2m/s and 20m/s² trajectory with sharp corners (Fig. 11). The mocap and robot coordinate frames were aligned using the 4 pulley locations and the mocap data was interpolated to match the 100Hz robot log frequency. The root mean square (RMS) error in control tracking, estimation, and repeatability across consecutive runs is 9.3mm, 3.4mm, and 3.3mm respectively (satisfying the requirements from Section III-A). We also validate our assumption that the end effector is always close to vertical, which is used for online control and estimation, by measuring the RMS rotation deviation to be 0.45°, 0.42°, and 1.57° in the horizontal, vertical, and normal directions respectively. We believe our proposed controller, which precomputes linear feedback gains offline, is accurate, easier to implement, and useful for CDPR applications other than graffiti as well.

In addition to the painting in Fig. 1, please refer to our supplemental video for additional painting results which qualitatively demonstrate our system’s capabilities. Also see our accompanying paper [28] for commentary on the quality of the supplemental video results as compared to Fig. 11.

D. Discussion & Limitations

When specifying an artist’s vision, the space of creative possibilities is immense. Non-flat overlap topologies [25], homographies, non-linear, and other 3D perspective transforms have artistic interest but are beyond the scope of this work. Graffiti stylization [26] and free-form inputs [52] are also beyond the scope of this work. In return, the artists get to explore the maxim “Creativity is born from limitations”.

Understanding the nuances of paint dispersion is another large area of study that is beyond the scope of this paper. For example, the fact that human artists consistently paint solid lines at 6m/s yet our robot’s lines begin losing complete coverage above 1.2m/s suggests a gap in our understanding. We believe moving closer to the canvas (artists were on average 3.0cm±0.1cm from the painting surface vs 12cm for our robot) and modeling special effects such as flares, blending, and intentional dripping are logical next steps.

Although our controller is generally reliable and robust to modeling inaccuracies, we do find some limitations on the Q matrix. For large Q matrices ($\|Q\|_2 / \|R\|_2 \geq 1e6N/m$),



Fig. 10. During path generation, we first produce the outline from artist inputs (left), then infill paths (center), and finally travel strokes (right).

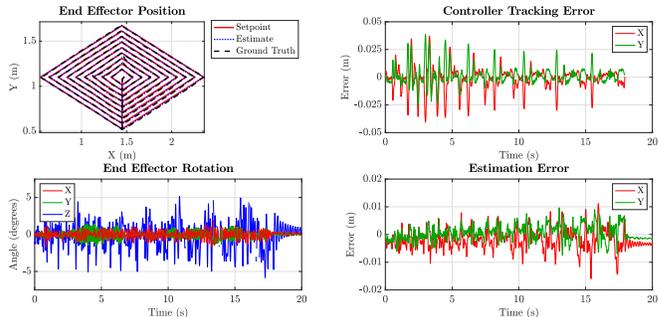


Fig. 11. Setpoint \mathbf{x}_{ff} , online estimate $\hat{\mathbf{x}}$, and ground-truth positions \mathbf{x} of the spray can (top left), end effector rotation (bottom left), tracking error $\mathbf{x}_{ff} - \hat{\mathbf{x}}$ (top right), and estimation error $\hat{\mathbf{x}} - \mathbf{x}$ (bottom right) for a challenging 2m/s, 20m/s² trajectory.

the controller resonates with the natural vibrations of the cable robot causing instability, while for small Q matrices ($\|Q\|_2 / \|R\|_2 \leq 1e2N/m$), the robot gets stuck for a few cm before overcoming the friction and returning to the setpoint trajectory. Still, compared to other methods such as [37], [38], we believe the iLQR method to be easier because it requires less tuning and modeling effort.

V. CONCLUSIONS & FUTURE WORK

In this paper, we presented a system for painting human-style graffiti art. Our work contributes to existing research by bridging three components in a systems approach: capturing artwork, building a graffiti robot, and planning and controlling the robot for painting graffiti. We illustrated the co-dependencies of various system design choices which suggest that future research in robot art should consider a more holistic approach. We also demonstrated that our system can successfully produce physical artworks from captured art. Our work can be applied to graffiti preservation by recreating captured artwork, to human-robot collaboration in art by enhancing the physical capabilities of artists, and to other fields through technology transfer for large-scale dynamic motion. Avenues for future work include a more portable graffiti capture device, better communication of robot limitations to the artist, style analysis and improvisation, paint dispersion analysis, real-time human-robot interaction, a larger workspace, and 6DoF robot motion.

ACKNOWLEDGEMENTS

We thank Max Ongena and Jules Dellaert for collaborating as graffiti artists, Prajwal Vedula and Zhangqi Luo for contributing to code, and Russel Gentry, Jake Thompkins, and Tristan Al-Haddad at Georgia Tech’s Digital Fabrication Lab for their hardware assistance and allowing us to use their space for painting.

REFERENCES

- [1] L. MacDowall, "In praise of 70k: Cultural heritage and graffiti style," *Continuum*, vol. 20, no. 4, pp. 471–484, 2006.
- [2] A. Liekens, Kenny, and L. Scheire, "Zet Kenny binnenkort zélf zijn eerste graffiti op een muur? — Team Scheire #9 [Will Kenny soon put his first graffiti on a wall? — Team Scheire #9]," <https://www.youtube.com/watch?v=jOpgvW7alhQ>, Sept 2020.
- [3] M. Gouttefarde, J. Lamaury, C. Reichert, and T. Bruckmann, "A versatile tension distribution algorithm for n -DOF parallel robots driven by $n + 2$ cables," *IEEE Transactions on Robotics*, vol. 31, no. 6, pp. 1444–1457, 2015.
- [4] J. M. Pagan, "Cable-suspended robot system with Real Time Kinematics GPS position correction for algae harvesting," Ph.D. dissertation, Ohio University, 2018.
- [5] F. Shahmiri and R. Gentry, "A survey of cable-suspended parallel robots and their applications in architecture and construction," *Blucher Design Proceedings*, vol. 3, no. 1, pp. 914 – 920, 2016.
- [6] R. Bostelman, J. Albus, N. Dagalakis, A. Jacoff, and J. Gross, "Applications of the NIST RoboCrane," in *Proceedings of the 5th International Symposium on Robotics and Manufacturing*, vol. 5, 1994.
- [7] P. Miermeister, M. Lächele, R. Boss, C. Masone, C. Schenk, J. Tesch, M. Kerger, H. Teufel, A. Pott, and H. H. Bühlhoff, "The CableRobot Simulator large scale motion platform based on cable robot technology," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016, pp. 3024–3029.
- [8] J. Wilkinson, E. Roth, T. Watson, C. Sugrue, and T. Vanderlin, "#000000book.com an open database for Graffiti Markup Language (GML) files." <https://000000book.com/>, 2021.
- [9] M. Yildirim and E. Roth, "GML recording machine," <http://fffff.at/gml-recording-machine/>, June 2011.
- [10] L. Scalera, E. Mazzon, P. Gallina, and A. Gasparetto, "Airbrush robotic painting system: Experimental validation of a colour spray model," in *Advances in Service and Industrial Robotics*, C. Ferraresi and G. Quaglia, Eds. Cham: Springer International Publishing, 2018, pp. 549–556.
- [11] D. Berio, S. Calinon, and F. F. Leymarie, "Learning dynamic graffiti strokes with a compliant robot," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 3981–3986.
- [12] L. Scalera, S. Seriani, A. Gasparetto, and P. Gallina, "Watercolour robotic painting: a novel automatic system for artistic rendering," *Journal of Intelligent and Robotic Systems*, pp. 1–16, 2018.
- [13] N. Roy, "Graffiti robot - train writing - planet256 / Niklas Roy / the fly - wall printer style machine," https://www.youtube.com/watch?v=zSIdvQsu27s&t=3s&ab_channel=ARTESANOBERLIN.
- [14] Y. Jun, G. Jang, B.-K. Cho, J. Trubatch, I. Kim, S.-D. Seo, and P. Y. Oh, "A humanoid doing an artistic work - graffiti on the wall," *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1538–1543, 2016.
- [15] A. Uryasheva, M. Kulbeda, N. Rodichenko, and D. Tsetserukou, "DroneGraffiti: Autonomous multi-UAV spray painting," in *ACM SIGGRAPH 2019 Studio*, ser. SIGGRAPH '19. New York, NY, USA: Association for Computing Machinery, 2019. [Online]. Available: <https://doi.org/10.1145/3306306.3328000>
- [16] A. S. Vempati, M. Kamel, N. Stulinovic, Q. Zhang, D. Reusser, I. Sa, J. Nieto, R. Siegwart, and P. Beardsley, "PaintCopter: An autonomous UAV for spray painting on three-dimensional surfaces," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 2862–2869, 2018.
- [17] B. Galea, E. Kia, N. Aird, and P. G. Kry, "Stippling with aerial robots," in *Computational Aesthetics (Expressive 2016)*, 2016, p. 10 pages.
- [18] TsuruRobotics, "Autonomous mural for sprite Ukraine." <https://tsuru.su/en/project/spritemural/>. [Online]. Available: <https://tsuru.su/en/project/spritemural/>
- [19] "Albert robot muralist," <https://www.robotmuralist.com/albert>. [Online]. Available: <https://www.robotmuralist.com/albert>
- [20] J. Lezni, "Soft monsters," *Perspecta*, vol. 40, pp. 22–27, 2008. [Online]. Available: <http://www.jstor.org/stable/40482274>
- [21] Y. Chen, W. Chen, B. Li, G. Zhang, and W. Zhang, "Paint thickness simulation for painting robot trajectory planning: a review," *Industrial Robot: An International Journal*, vol. 44, no. 5, pp. 629–638, 2021/09/04 2017. [Online]. Available: <https://doi.org/10.1108/IR-07-2016-0205>
- [22] H. Chen, T. Fuhlbrügge, and X. Li, "A review of CAD-based robot path planning for spray painting," *Industrial Robot: An International Journal*, vol. 36, no. 1, pp. 45–50, 2021/09/04 2009. [Online]. Available: <https://doi.org/10.1108/01439910910924666>
- [23] M. V. Andulkar, S. S. Chiddarwar, and A. S. Marathe, "Novel integrated offline trajectory generation approach for robot assisted spray painting operation," *Journal of Manufacturing Systems*, vol. 37, pp. 201–216, 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0278612515000229>
- [24] Y. Zeng, J. Gong, N. Xu, and N. Wu, "Tool trajectory optimization of spray painting robot for manytimes spray painting," *International Journal of Control and Automation*, vol. 7, pp. 193–208, 08 2014.
- [25] D. Berio, P. Asente, J. Echevarria, and F. F. Leymarie, "Sketching and layering graffiti primitives," in *8th ACM/Eurographics Expressive Symposium, Expressive 2019, Genoa, Italy, May 5-6, 2019, Proceedings*, C. S. Kaplan, A. G. Forbes, and S. DiVerdi, Eds. Eurographics Association, 2019, pp. 51–59. [Online]. Available: <https://doi.org/10.2312/exp.20191076>
- [26] D. Berio, S. Calinon, and F. F. Leymarie, "Dynamic graffiti stylisation with stochastic optimal control," in *Proceedings of the 4th International Conference on Movement Computing*. ACM, 2017, p. 18.
- [27] D. Berio and F. F. Leymarie, "Computational models for the analysis and synthesis of graffiti tag strokes," in *Proceedings of the Workshop on Computational Aesthetics*. Eurographics Association, 2015, pp. 35–47.
- [28] G. Chen, S. Baek, J.-D. Florez, W. Qian, S. won Leigh, S. Hutchinson, and F. Dellaert, "Extended version of GTGraffiti: Spray painting graffiti art from human painting motions with a cable driven parallel robot," 2021, arXiv:2109.06238 [cs.RO].
- [29] L. Della Valle, T. G. Andrews, and S. Ross, "Perceptual thresholds of curvilinearity and angularity as functions of line length." *J Exp Psychol*, vol. 51, no. 5, pp. 343–347, May 1956.
- [30] R. Nan, D. Li, C. Jin, Q. Wang, L. Zhu, W. Zhu, H. Zhang, Y. Yue, and L. Qian, "The five-hundred-meter aperture spherical radio telescope (FAST) project," *International Journal of Modern Physics D*, vol. 20, no. 06, pp. 989–1024, 2011.
- [31] S. Bandyopadhyay, "Lunar crater radio telescope (LCRT) on the far-side of the moon," April 2020. [Online]. Available: <https://www.nasa.gov/directorates/spacetech/niac/2020%5FPhase%5F1%5FPhase%5FII/lunar%5Fcrater%5Fradio%5Ftelescope/>
- [32] P. Bosscher, A. T. Riechel, and I. Ebert-Uphoff, "Wrench-feasible workspace generation for cable-driven robots," *IEEE Transactions on Robotics*, vol. 22, no. 5, pp. 890–902, 2006.
- [33] S. Bouchard, C. Gosselin, and B. Moore, "On the ability of a cable-driven robot to generate a prescribed set of wrenches," *Journal of Mechanisms and Robotics*, vol. 2, no. 1, 2/15/2021 2009.
- [34] M. Gouttefarde, D. Daney, and J. Merlet, "Interval-analysis-based determination of the wrench-feasible workspace of parallel cable-driven robots," *IEEE Transactions on Robotics*, vol. 27, no. 1, pp. 1–13, 2011.
- [35] A. Liekens, "Spray can servo mount," <https://www.thingiverse.com/thing:4622176>, Oct 2020.
- [36] E. Todorov and W. Li, "A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems," in *Proceedings of the 2005, American Control Conference, 2005.*, 2005, pp. 300–306.
- [37] A. Pott, T. Bruckmann, and L. Mikelsons, "Closed-form force distribution for parallel wire robots," in *Computational Kinematics*, A. Kecskeméthy and A. Müller, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 25–34.
- [38] P. Miermeister, A. Pott, and A. Verl, "Auto-calibration method for overconstrained cable-driven parallel robots," in *ROBOTIK 2012; 7th German Conference on Robotics*, 2012, pp. 1–6.
- [39] M. Agahi and L. Notash, "Redundancy resolution of wire-actuated parallel manipulators," *Transactions of the Canadian Society for Mechanical Engineering*, vol. 33, no. 4, pp. 561–573, 2009.
- [40] M. Hassan and A. Khajepour, "Analysis of bounded cable tensions in cable-actuated parallel manipulators," *IEEE Transactions on Robotics*, vol. 27, no. 5, pp. 891–900, 2011.
- [41] H. D. Taghirad and Y. B. Bedoustani, "An analytic-iterative redundancy resolution scheme for cable-driven redundant parallel manipulators," *IEEE Transactions on Robotics*, vol. 27, no. 6, pp. 1137–1143, 2011.
- [42] J. Lamaury and M. Gouttefarde, "Control of a large redundantly actuated cable-suspended parallel robot," in *2013 IEEE International Conference on Robotics and Automation*, 2013, pp. 4659–4664.

- [43] W. Shang, B. Zhang, S. Cong, and Y. Lou, "Dual-space adaptive synchronization control of redundantly-actuated cable-driven parallel robots," *Mechanism and Machine Theory*, vol. 152, p. 103954, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0094114X20301750>
- [44] E. Galceran and M. Carreras, "A survey on coverage path planning for robotics," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1258–1276, 2013. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S092188901300167X>
- [45] F. Devai, "Quadratic bounds for hidden line elimination," in *Proceedings of the Second Annual Symposium on Computational Geometry*, ser. SCG '86. New York, NY, USA: Association for Computing Machinery, 1986, pp. 269–275. [Online]. Available: <https://doi.org/10.1145/10515.10544>
- [46] B. Siciliano, L. Sciavicco, L. Villani, and G. Oriolo, *Robotics: modelling, planning and control*. Springer Science & Business Media, 2010, ch. 4.
- [47] D. Lau, J. Eden, Y. Tan, and D. Oetomo, "CASPR: A comprehensive cable-robot analysis and simulation platform for the research of cable-driven parallel robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016.
- [48] J. C. Butcher, *Numerical Differential Equation Methods*. John Wiley & Sons, Ltd, 2016, ch. 2, pp. 55–142. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781119121534.ch2>
- [49] F. Dellaert and M. Kaess, "Factor graphs for robot perception," *Foundations and Trends in Robotics*, vol. 6, pp. 1–139, 2017.
- [50] S. Yang, G. Chen, Y. Zhang, F. Dellaert, and H. Choset, "Equality constrained linear optimal control with factor graphs," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [51] G. Chen and Y. Zhang, "LQR control using factor graphs," <https://gtsam.org/2019/11/07/lqr-control.html>, Nov 2019, note: Superceding ICRA 2021 paper pending review.
- [52] D. Berio, F. F. Leymarie, and R. Plamondon, "Expressive Curve Editing with the Sigma Lognormal Model," *Eurographics 2018 - Short Papers*, 2018.