

Design and Evaluation of a Hair Combing System Using a General-Purpose Robotic Arm

Nathaniel Dennler¹, Eura Shin², Maja Mataric¹ and Stefanos Nikolaidis¹

Abstract—This work introduces an approach for automatic hair combing by a lightweight robot. For people living with limited mobility, dexterity, or chronic fatigue, combing hair is often a difficult task that negatively impacts personal routines. We propose a modular system for enabling general robot manipulators to assist with a hair-combing task. The system consists of three main components. The first component is the segmentation module, which segments the location of hair in space. The second component is the path planning module that proposes automatically-generated paths through hair based on user input. The final component creates a trajectory for the robot to execute. We quantitatively evaluate the effectiveness of the paths planned by the system with 48 users and qualitatively evaluate the system with 30 users watching videos of the robot performing a hair-combing task in the physical world. The system is shown to effectively comb different hairstyles.

I. INTRODUCTION

There are 40.7 million people in United States living with some form of reduced mobility [1] resulting from a variety of factors (e.g., traumatic injury, stroke, genetic causes, etc.), according to the National Health Interview Survey. Most items used in activities of daily living (ADLs) are not explicitly designed for people living with reduced mobility and may be completely unusable. This requires people living with limited mobility to either purchase specialized items designed for mobility limitations or acquire assistance from trained professionals. However, for many people, neither option may be accessible nor affordable.

The high costs of formal caregivers has necessitated their replacement with family members who give care informally. In 2015, 43.5 million people in the US reported providing informal care in the previous 12 months. Informal caregiving requires providing companionship, personal care, scheduling health services, providing transportation, and more, while at the same time managing one's own job, family finances, household tasks, and other daily activities [2]. Due to the increased caregiver load, informal caregiving can be a source of stress and can lead to health problems [3]. However, assisting people with limited mobility with ADLs is necessary to decrease their risk of depression and, in some forms of limited mobility, can aid in their rehabilitation processes [4].

Assistive robots could provide a solution to this growing issue by enabling people with limited mobility to independently perform necessary daily tasks such as eating [5], drinking [6], and grabbing objects [7]. Areas of self-care




Fig. 1: System setup and trajectories for different hairstyles.

such as shaving [8], cleaning [9], and dressing [10], [11] have similarly shown potential for assistive robotics for non-critical tasks. However, fully autonomous systems are not always desirable, as it limits the user's choice and leads to lower acceptance of the system in the event of errors [5]. A key design goal is thus to include user input to the system in a way that reduces the effort of the user yet simultaneously maintains the user's own autonomy.

We propose to expand the abilities of general-purpose assistive robotic arms to the domain of hair-combing, an important and under-explored self-care ADL. In this work, we provide a minimal system to comb different kinds of hairstyles. The system consists of three modules: an image segmentation module, a path-planning module, and a trajectory generation module. The main contributions of this paper are: 1) the formulation of a path planning method for hair combing, and 2) insights from naïve users on the efficacy of that approach to hair combing. Together, the system provides users or remote caretakers a way to automatically generate paths through a user's hair and have a robot comb along the generated paths. The system provides functionality using a robotic arm and an RGB-D camera. This work presents a physical implementation of the system on a lightweight Kinova Gen2 robot arm and an evaluation of its performance. Our results demonstrate that the system successfully combs a variety of hairstyles.

II. RELATED WORK

Assistive Robotics: Advances have made it possible to provide robot-assisted ADLs [12]. Many of the existing systems are high degree-of-freedom arms mounted on tables [13], [14], wheelchairs [15]–[17], or mobile platforms [18], [19] capable of accomplishing a wide variety of tasks. Several are

¹Nathan Dennler, Maja Mataric, and Stefanos Nikolaidis are from the Computer Science Department at the University of Southern California, Los Angeles, CA {dennler, mataric, nikolaid}@usc.edu

²Eura Shin is from the Computer Science Department at Harvard University, Cambridge, MA eurashin@g.harvard.edu

commercially available [20]–[22]. However, these systems are often difficult to control due to the high number of controllable joints, creating a need for some autonomy in order to perform ADLs in an efficient and low-effort manner [23]. In addition, many of these systems require specialized or wearable sensors that can have unintended negative physical and psychological consequences for the user [24]. To mitigate those issues, our work aims to be minimally invasive and uses a minimal setup, while additionally utilizing a general-purpose robot arm that is useful in other capacities.

Hair Modeling: Current works in hair modeling focus on recreating and rendering entire hair structures in 3D from 2D images [25]. In general, hair can be modeled as a collection of curves in 3D space that follow an underlying orientation field that can be inferred from the surface of the hair [26], [27]. Some approaches utilize a database of 3D models, which can be combined to form new styles that are found in single-view images [28]. Other approaches [29] use a pipeline for editing hairstyles in portraits by finding orientation fields from images to generate additional strands [29]. Our system takes inspiration from these works, and models hair as an orientation field. We emphasize being able to generate paths in real time on hair that becomes neater as the system combs it.

III. SYSTEM DESCRIPTION

Our goal is to design a system for people living with limited mobility so they can comb their own hair. To do this, we propose a system consisting of three modules: the *segmentation module* localizes the hair in space, the *path planning* module proposes brush strokes from a starting point input by the user, and the *trajectory generation module* provides the robot with a trajectory through space to move along the user-selected path.

A. Hardware

The hardware for this system consists of a depth camera and a robot manipulator. Our implementation uses a Microsoft Kinect v1 depth sensor which is commonly used in various robotics domains. We used the Kinova JACO Gen 2 lightweight robotic manipulator, which is also used effectively in several other assistive settings [23], [30] because it is safe for physical interaction. While a comb is also required, our system is designed to work with any comb that can be grasped by the robot arm. We chose a wide-toothed comb, but the comb should be chosen with consideration for the user’s hair type. A diagram of the system setup for different hairstyles is shown in Fig. 1.

B. Segmentation Module

Hair is difficult to locate in images due to its high variability in color, length, and style. We use a learning-based approach to segment hair in RGB images and consider this mask a noisy estimate of the hair. We refine this estimate through signal processing techniques as outlined in Fig. 2




Fig. 2: Diagram of the image segmentation module.

1) *Segmentation Formulation:* Segmentation is a well-established field of study in computer vision [31], [32]. The overall approach for neural architectures is to take a high dimensional image and learn a low-dimensional representation of the image with an encoder network. The low-dimensional representation is then up-sampled to form an image of the same size as the original image that contains the mask. The loss is quantified as the binary cross-entropy loss between a ground-truth mask and the predicted mask.

Due to the success and efficiency of fine-tuning pre-trained models, we use SqueezeNet [33] as the encoder; it is a small-sized network that achieves high accuracy on Image Net Classification tasks. The smaller model provides comparable accuracy with fewer parameters and thus has a faster run time, which is essential for real-time applications like ours. The decoder has a similar architecture to the encoder, but uses deconvolutional layers instead. We used the Figaro 1k dataset [34] for training; it contains a wide variety of labelled hair segments from unconstrained camera angles.

2) *Refining the segmented region:* To improve the fidelity of the mask, we extract the points in XYZ space that correspond with the hair pixels. Two main sources of error are of concern in accurately segmenting hair from RGB images: *temporal* errors and *spatial* errors.

Temporal errors can result from sensor noise, causing the estimated hair mask to form discontinuities or disappear completely for a single frame. To reduce the impact of such high-frequency noise, we use an exponential smoothing filter to average values from the stream of masks. This filter averages masks over time, but gives exponentially higher weights to images nearby in the sequence. This reduces the effects of high-frequency noise (i.e., sensor noise) but allows low-frequency patterns (i.e., user movement) to pass through. The filter is described by the following iterative function:

$$mask(t, \hat{Y}_t) = \begin{cases} \hat{Y}_t & t = 0 \\ \alpha \cdot \hat{Y}_t + (1 - \alpha) \cdot mask(t - 1, \hat{Y}_{t-1}) & t > 0 \end{cases} \quad (1)$$

where t is the timestep, \hat{Y}_t is the output of the segmentation network for frame t , and α is the smoothing parameter. Higher values of α allow better tracking of user movement and lower values reduce the effect of sensor noise. Our system uses $\alpha = 0.9$.

To address spatial errors (i.e., false-positives, other people in the frame, etc.) we first select the largest contiguous area in

the mask. We then find the median value of pixels selected as representing the depth of the hair, an estimate that is robust to outliers. Points that are within two standard deviations of this median value and have hues within the range of the pixels in the mask are added to the mask to reduce the number of false negatives.

C. Path Planning Module

We developed a path planning module that takes the current hair state from the segmentation module and a user-specified point as input and produces a path through the hair that aligns with the natural curvature of the hair. For this work, we use a mouse to make the selection, but this pipeline is not dependent on the input device. Following works from Computer Graphics on hair rendering, we model the flow of hair as a differential field [29]. To obtain the differential field, we find the orientation of the hair in the image. We consider a path that follows the flow of the hair as a solution to the differential field, with initial conditions defined as the point a user selects as the path's starting point. For the following sections, we leverage the common assumption that an image is a smooth continuous function that maps XY coordinates to intensity values. The function is discretized into pixels and the intensity value at each pixel is noisily sampled from the underlying intensity function. This assumption allows us to take derivatives over the image to construct the orientation field.

1) *Coherence-Enhancing Shock Filters*: We first use the method described by Weickert et al. [35] to make the natural flows of the hair more coherent, reducing the effect that stray strands of hair on the calculation of orientation. To do this, we iteratively refine an image by increasing contrast in the direction of the greatest change of the intensity gradient (directions perpendicular to the flow of hair), and decreasing contrast in the direction of the lowest change of the intensity gradient.

First, we calculate the dominant eigenvector e for each pixel's local neighborhood, which provide the direction of greatest change in intensity. The convexity in the direction of e is determined by $e^T H e$, where H is the Hessian matrix in each pixel's neighborhood. If the convexity is positive, the algorithm intensifies the pixel to match its surroundings. If the convexity is negative, the algorithm reduces the pixels intensity. The net effect of this process is to increase the contrast in the direction of the greatest change in intensity. The process is shown in Algorithm 1, and the effect on images is seen in Fig. 3.

2) *Calculating Orientation*: Once the image has been filtered to enhance the coherence of the hair flows, the system calculates the direction of the hair flows in the image. The gradient structure tensor provides this information, which is defined continuously for a point p as:

$$S_{K_E} = \begin{bmatrix} \int_r K_E(r) \cdot (I_x(p-r))^2 & \int_r K_E(r) \cdot I_x(p-r)I_y(p-r) \\ \int_r K_E(r) \cdot I_y(p-r)I_x(p-r) & \int_r K_E(r) \cdot (I_y(p-r))^2 \end{bmatrix} \quad (2)$$

Algorithm 1 Coherence Filter

- 1: **Input:** Image I , kernel size K_δ for approximating derivative, kernel size K_e for calculating eigenvectors, kernel size K_m for calculating max and min values, Constant C_{blend} the blending rate for each iteration, and T the total number of iterations.
- 2: **for** $t = 1, 2, \dots, T$ **do**
- 3: Calculate first (normalized) eigenvector $[e_x, e_y]$ for each sub image of size $K_e \times K_e$ in I_t
- 4: Approximate $\frac{\delta^2 I}{\delta x^2}$, $\frac{\delta^2 I}{\delta x \delta y}$, and $\frac{\delta^2 I}{\delta y^2}$ with a sobel filter of size K_δ .
- 5: Compute $I_{vv} \leftarrow e_x^2 * \frac{\delta^2 I}{\delta x^2} + 2 * e_x * e_y * \frac{\delta^2 I}{\delta x \delta y} + e_y^2 * \frac{\delta^2 I}{\delta y^2}$
- 6: Create a new image I'_t , where for each sub image $K_m \times K_m$ in I , we take the max value of the sub image if $I_{vv} > 0$ and the min if $I_{vv} < 0$
- 7: $I_{t+1} \leftarrow I_t * C_{blend} + I'_t * (1 - C_{blend})$
- 8: **return** image
- 9: *Note: our implementation uses $K_\delta = 7$, $K_e = 11$, $K_m = 3$, $C_{blend} = 0.9$, and $T = 3$*

Algorithm 2 Find Orientations

- 1: **Input:** Image I , kernel size K_δ for approximating derivative, kernel size K_E for approximating expectation
- 2: Approximate $\frac{\delta I}{\delta x}$ and $\frac{\delta I}{\delta y}$ (with a sobel filter of size K_δ).
- 3: Construct: $\begin{bmatrix} I_x I_x & I_x I_y \\ I_y I_x & I_y I_y \end{bmatrix} = \nabla I^T \nabla I = S_0$
- 4: Locally average each element in S_0 with kernel size K_E to calculate Structure Tensor: $\begin{bmatrix} J_{11} & J_{12} \\ J_{21} & J_{22} \end{bmatrix} = S_{K_E}$
- 5: Calculate local orientation: $\Theta = \frac{1}{2} \text{atan2}(\frac{J_{12} + J_{21}}{J_{22} - J_{11}}) + \frac{\pi}{2}$
- 6: **return** Θ
- 7: *Note: our implementation uses $K_\delta = 3$ and $K_E = 5$*




Fig. 3: Overview of path-planning module.

where K_E represents a window of the part of the image to consider. The integral corresponds to the convolution operation over a continuous image. The pseudo-code for computing S_{K_E} discretely for all pixels is detailed in Algorithm 2 on lines 2-4. Orientations from this structure tensor following the equation on line 5, as detailed by Yang [36].

3) *Generating the path*: Once the orientation field is constructed, the user selects a point on the image for the starting point for the comb. Once the point is selected, the path is iteratively generated by adding new points that

Algorithm 3 Calculate Path

- 1: **Input:** Orientation image Θ , selected pixel for the start of the path $p_0 = [p_x, p_y]$, step size k (in pixels).
- 2: initialize $path$ with v_0
- 3: **while** p_t is within bounds of hair **do**
- 4: $\theta_{p_t} = \Theta[p_t]$
- 5: $p_{t+1} \leftarrow p_t + k \cdot [\cos(\theta_{p_t}), \sin(\theta_{p_t})]$
- 6: append p_{t+1} to the path
- 7: **return** $path$
- 8: *Note: our implementation uses $k = 6$*




Fig. 4: Creating task-space poses from points of stroke.

follow the direction of the orientation field, as described in Algorithm 3.

D. Trajectory Generation Module


The trajectory generation module takes in the path in the form of pixel coordinates and maps them to task-space poses. While we used a stationary camera for the mapping, this is not required, since we only need the transformation between the camera and robot frames when the image was taken.

We retrieve the XYZ position of each point in the path from the point cloud, where we index points in the point cloud by the pixels in the path. We assign orientations to each point, so that the teeth of the comb (z-axis of the end-effector, following the Denavit–Hartenberg convention) point orthogonally to the user’s hair plane H , formed by the visible portion of the hair. Its normal vector is \vec{h} .

The Y axis of the end effector is aligned with the vector tangent to the path that the robot is following, projected onto the hair plane. This is approximated with the centered difference of the previous and next points in the path. Figure 4 shows this graphically, where the direction of the end effector’s y coordinate frame before normalization (measured from the camera’s reference frame) at point p_t is:

$$EE_y = (\vec{p}_{t-1} - \vec{p}_{t+1}) - \frac{(\vec{p}_{t-1} - \vec{p}_{t+1}) \cdot \vec{h}}{\|\vec{h}\|} \vec{h} \quad (3)$$

A motion plan for the robot end effector is planned in task space, following joint velocity and acceleration constraints. A preview of the resulting motion plan is shown to the user. If the user accepts the preview, the robot executes the motion plan, otherwise the system prompts the user to select a new start point in the hair and generates a new path and trajectory.



(a) Mesh-based Algorithm. (b) Image-based Algorithm.

Fig. 5: Illustrative hairstyle where the image-based algorithm performs differently than the mesh-based for corresponding starting points.

IV. PATH PLANNING EVALUATION

To evaluate the path planning module’s design choice of following curves of the hair as an appropriate combing strategy, we propose a baseline alternative, called the mesh-based method. In this approach, the strategy for combing is instead based on selecting paths that follow the shortest path to the bottom of the hair from an initial starting position.

A. Mesh-Based Method

To create paths that move directly downward, we plan directly on a mesh created from the point cloud of the hair, taken by the Kinect camera. Next we segment the image based on color and depth to find candidate point clouds of the head. A human operator selects the subset of point clouds that correspond to the user’s hair. The union of these point clouds represent the entirety of the hair structure to be brushed. The resulting point cloud is then converted to a mesh by using a greedy surface triangulation algorithm [37] from the PointCloud Library [38]. Once the mesh is created, paths are formed by sampling mesh vertices of the hair mesh model as starting points. The mesh is a graph with nodes corresponding to mesh vertices and edges corresponding to the edges of the triangles that define the mesh. We define the goal as any point in the bottom portion of the hair mesh. We then formulate path planning as finding a low-cost path from the start vertex to one of the vertices in the goal set. We find these paths using the A* algorithm [39] with Euclidean distance as the distance metric and the straight-line distance to the bottom of the mesh as the heuristic.

To generate a variety of paths, the algorithm is run several times using a variety of sampled start states. Examples of paths generated by this algorithm and how it compared to the image-based algorithm is shown in Figure 5. The key difference between these algorithms is that the mesh-based algorithm combs in a downward manner, whereas the image-based algorithm is able to additionally brush sideways paths.

B. Study Setup

To evaluate the strengths and weaknesses of the these approaches, we measured user perception of the methods through an on-line study using Amazon Mechanical Turk. A



Fig. 6: Examples of generated paths randomly sampled from the set of solutions on the thirty hairstyles. Each hairstyle came in five different colors and from one of three camera viewpoints used to evaluate the algorithms in the user study. Red paths represent the mesh-based method, orange paths represent the image-based method, and blue paths represent the human-drawn paths.

variety of hairstyles were selected from the USC hair salon database [28]. Ten short, ten medium, and ten long hairstyles were selected at random, and rendered in one of five hair colors (black, brown, auburn, blond(e), and grey) from one of three camera angles.

Each MTurk participant saw 10 of the 30 hairstyles at random, and was shown three types of paths: the mesh-based method, the image-based method, and a set of human-drawn paths as a high-quality reference point. Human-drawn paths were generated by an independent researcher not affiliated with the project. For each algorithm, participants were shown three paths randomly selected from a set of solutions at different starting points to provide a holistic view of the types of paths the algorithms generated. For each algorithm, participants rated the generated paths on: 1) whether the paths *are able* to brush hair (referred to as **completeness**), and 2) whether the paths *effectively* brush the hair (referred to as **effectiveness**).

C. Hypotheses

H1: Participants will prefer paths that vary in direction, and thus rate the image-based method as generating paths that are more complete than the mesh-based method.

H2: Participants will prefer paths that follow the flow of the hair, and thus rate the image-based method as generating paths that are more effective than the mesh-based method.

D. Results

The study received 48 respondents (27 women, 21 men; median age range 25-34). Responses were measured on a scale of -3 to 3, with 0 indicating neutral responses. A repeated-measures ANOVA test found significant differences for the main effect of algorithm on reported **completeness** of planned path, with $F(2,94) = 24.12$, $p < .001$. Pairwise t-tests were used for post-hoc analysis, revealing that all pairwise comparisons were significant after Bonferroni correction.

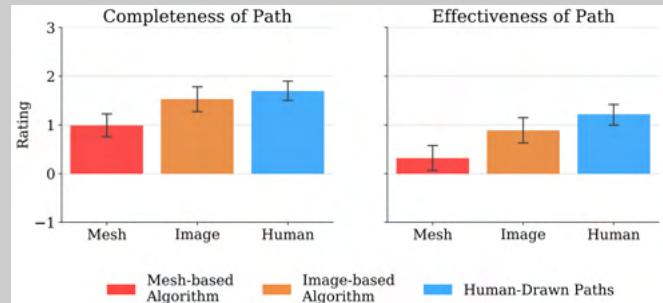


Fig. 7: Means of the three different algorithms. Error bars represent 95% Confidence Intervals of mean ratings. All differences are significant.

The mesh-based method of planning paths ($M=.99$) received significantly lower ratings than the image-based method ($M=1.53$), $p < .001$, $\eta^2 = .09$ which was in turn rated as less complete than the human-drawn paths ($M=1.70$), $p = .01$, $\eta^2 = .01$. The difference between the mesh-based method and the human-drawn paths was also significant, $p < .001$, $\eta^2 = .17$. This supports **H1**.

We performed the same analysis for **effectiveness** ratings. A repeated-measures ANOVA test shows significant differences by algorithm type, with $F(2,94) = 21.69$, $p < .001$. Pairwise t-tests show significance between all comparisons with $p < .001$, where the mesh-based method ($M=.32$) was lower than the image-based method ($M=.89$), $\eta^2 = .08$, and the human-drawn paths ($M=1.22$) were rated the highest, $\eta^2 = .04$. This supports **H2**.

Qualitative notes from participants in the study revealed that they found the human paths to be rated the most highly because they tended to “follow the hair flow” but were also “long and straight”. The main concern with the image-based method was that it followed the path of the hair too closely. This is evident for the hairstyles that have waves, where combing in a wavy path may be inefficient. Because the

Theme	Positive Example	Negative Example
Effectiveness	"...the hair [that] the robot worked on looked much better..."	"...there is no effect on the hair that the robot is combing..."
Force	"It did seem gentle, so that is good."	"...it looked like it was pushing rather rough at the hair."
Movement	"This robot appeared to have finesse and sensitive movements."	"...there is no grace or nuance to the motion."

TABLE I: Examples of positive and negative qualitative responses from participants for each theme.




Fig. 8: Sample frames of the video shown to the participants and forces measured by the arm for different strokes of each hairstyle. Orange lines represent mean force values across 25 strokes with the orange region illustrating the first and third quartiles. Blue lines represent individual stroke force values. All force readings were measured at 10Hz and were post-processed with a sliding average of 9 timesteps for visualization.

hair can deform as the comb passes through it, moving along the waves is less direct than moving straight. However, moving along the waves is still preferable to moving directly downward without following the general direction of the hair, as in the mesh-based method.

V. TRAJECTORY EXECUTION

A. On-line User Study

To evaluate the performance of the physical robot implementation of the system, we recorded three videos of the robot combing three different wigs.¹ The wigs varied in color and hairstyle and were representative of the hairstyles in the first study, consisting of a short brown wig, a medium-length pink wig, and a long blonde wig. The wigs were teased so that the hair was in a state that would warrant using a hair-combing device. Each video showed the wig and robot arm from the side so that study participants could see both the end effector and wig at all times, shown in Fig. 8. We used a plain white background for all videos. A human operator selected the start points for the robot system off-screen. The participants watched the wig being combed by the robot, where the first and final strokes were filmed in real time and the middle strokes were filmed at 20x speed to show the progress that the robot makes over the course of the hair-combing interaction to reduce fatigue effects. Participants were then asked to provide qualitative feedback on the effectiveness of the combing as well as their willingness to use the system.

¹We are unable to perform in-person studies due to the ongoing pandemic.

B. Qualitative Analysis

The study received 30 responses (17 men, 13 women; median age range 25-34). We conducted a qualitative analysis of the open-ended survey responses utilizing a grounded-theory approach [40] where we categorized responses by theme in a data-driven manner. The qualitative analysis revealed three main themes in participants' evaluations of the system. The first theme was related to the effectiveness of the task of hair combing. The second theme was the perception of the amount of force that was used throughout the interaction. The third theme was related to how the robot moved while combing hair. Representative positive and negative examples of each theme are shown in Table I.

1) *Effectiveness*: The theme of effectiveness had several subcomponents. The top three were: final alignment of hair, depth of combing, and time it takes the robot to comb.

The final alignment of the hair was mentioned in 31 of the 90 responses; 25 of these responses indicated that respondents thought that the hair ended up "...straighter than before it started combing" or some variant highlighting that the hair looked neater after the combing interaction. This subcomponent received the highest number of mentions across all categories, indicating that the robot is functionally capable of combing hair.

Related to depth of combing, the majority of negative responses came from the short brown hairstyle due to the fact that the hair lies very close to the scalp. 19 responses mentioned that the robot appears to comb the hair on the outermost layer, but 4 of the responses indicate that the robot also penetrates the deeper layers of the hair. All 4 positive responses were for the medium and long hairstyles.

The subcomponent of timing was noted in 24 responses.

19 of these responses indicated that the robot took a long time to move when combing, indicating that they were willing to have a robot that has higher velocity limits for this task. Interestingly, 4 responses positively described the slow movement of the robot as cautious and gentle.

2) *Use of Force*: The theme of force was concerned with the amount of force the robot used during the combing interaction. The most common comment was that the robot appeared to use too much force when combing the wig, present in 15 of the responses. This was either stated generally, or quantified as the amount of movement that the mannequin head made during the combing execution. This brought up concerns for a human user's safety when using such a system in 6 of the participants' responses. Conversely, the amount of force used was perceived as gentle and appropriate by 6 of the responses, who mentioned that the robot appeared to move carefully. Participants responded that the robot did not appear to have an understanding of the sensation of touch, which worried them as the robot may not realized how hard it is pulling on knots or tangles. Physically, the maximum force exerted by the system for any hairstyle did not exceed 15N, with the majority of strokes exerting peaks of less than 5N for each hair style as measured by joint efforts of the arm (see Fig. 8). These values are similar in magnitude to sensitive physical interaction tasks such as shaving (7N peak magnitude) and wiping cheeks (14N peak magnitude) [41].

3) *Characteristics of Movement*: The responses that regarded the characteristics of movement mentioned both auditory and visual components. Participants mentioned the lack of sound during the movement as uncomfortable. Visually, some participants described the static orientation of the hand during the approach as lacking nuance, whereas others indicated that the robot following the curves and shape of the hair as having finesse and being smooth. This indicates a need to adapt motion to the specific user of the system.

VI. DISCUSSION AND CONCLUSION

Results from the qualitative survey indicate that participants found the proposed system to be effective in hair combing, but they raised interesting points to improve the user's experience during the interaction.

A key limitation of this study is that it was conducted through online videos and not in person, limiting participants to only visual and auditory modalities of perception. This allowed them to accurately make judgements of the sound that the robot makes while moving, but prevented them from directly experiencing tactile information such as forces used by the system. While the robot in the videos implemented position control, our system provides only the set of goal positions and their associated times. Implementing hybrid position-force controllers [42] would allow the system to move between these same goal positions while simultaneously enforcing force constraints, allowing users to specify the force that they are comfortable with. Being controller-agnostic additionally allows the force to be estimated through different kinds of sensors depending on the robot (e.g.,

current sensors, end effector force sensors, etc.). Further work should critically evaluate controller design and user acceptance in a real-world setting.

Participants indicated concern that the silence of the interaction was unsettling, reflecting findings of robots in social settings [41], [43], [44], but contrary to findings in other impersonal and intimate interaction settings [45], [46]. Future work should explore the implementation of pleasant auditory elements that allow the anticipation of the robot's motion.

Some participants additionally indicated that the qualitative aspects of the robot's movement lacked nuance. Previous works have linked user perception of robot movement to the orientation and speed of the robot [47], [48], warranting future work in modulating the end effector's orientation and speed during its approach. For use in intended assistive services, the robot's physical appearance can also be improved with best embodiment design practices, as discussed in Deng et al. [49], e.g., using robotic arms with a traditionally less mechanical appearances for these social tasks.

While our algorithm performs well on the hairstyles selected from the database, the evaluation does not cover all possible hairstyles, such as very curly hair or sparse hair which are realistic to encounter in naturalistic settings. However, it is important to consider that not all hairstyles should be dry-combed daily; for example, combing very curly hair may damage coils, and certain styles such as buzz cuts or braided hair would not benefit from utilizing this system. Nevertheless, validating this approach on other relevant hairstyles is an important direction for future work.

Overall, this work brings about a better understanding of how general-purpose robotic arms can effectively comb hair by adapting to different hairstyles. We presented a modular system for robotic hair combing using a minimal setup that can work in conjunction with other assistive applications. We found that users' experience is related to the robot's movement, sound, and perceived applied force, thus opening up a range of human-robot interaction research questions.

REFERENCES

- [1] NHIS, "Table a-10. difficulties in physical functioning among adults aged 18 and over, by selected characteristics: United states, 2018," Nov 2019.
- [2] L. Feinberg, S. C. Reinhard, A. Houser, R. Choula, *et al.*, "Valuing the invaluable: 2011 update, the growing contributions and costs of family caregiving." *Washington, DC: AARP Public Policy Institute*, vol. 32, p. 2011, 2011.
- [3] M. Pinquart and S. Sørensen, "Differences between caregivers and noncaregivers in psychological health and physical health: a meta-analysis." *Psychology and aging*, vol. 18, no. 2, p. 250, 2003.
- [4] I.-F. Lin and H.-S. Wu, "Does informal care attenuate the cycle of adl/iadl disability and depressive symptoms in late life?," *Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, vol. 66, no. 5, pp. 585–594, 2011.
- [5] T. Bhattacharjee, E. K. Gordon, R. Scalise, M. E. Cabrera, A. Caspi, M. Cakmak, and S. S. Srinivasa, "Is more autonomy always better? exploring preferences of users with mobility impairments in robot-assisted feeding," in *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 181–190, 2020.
- [6] F. F. Goldau, T. K. Shastha, M. Kyrarini, and A. Gräser, "Autonomous multi-sensory robotic assistant for a drinking task," in *2019 IEEE 16th International Conference on Rehabilitation Robotics (ICORR)*, pp. 210–216, IEEE, 2019.

- [7] A. Jain and C. C. Kemp, "El-e: an assistive mobile manipulator that autonomously fetches objects from flat surfaces," *Autonomous Robots*, vol. 28, no. 1, p. 45, 2010.
- [8] K. P. Hawkins, P. M. Grice, T. L. Chen, C.-H. King, and C. C. Kemp, "Assistive mobile manipulation for self-care tasks around the head," in *2014 IEEE Symposium on computational intelligence in robotic rehabilitation and assistive technologies (CIR2AT)*, pp. 16–25, IEEE, 2014.
- [9] C.-H. King, T. L. Chen, A. Jain, and C. C. Kemp, "Towards an assistive robot that autonomously performs bed baths for patient hygiene," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 319–324, IEEE, 2010.
- [10] Z. Erickson, H. M. Clever, G. Turk, C. K. Liu, and C. C. Kemp, "Deep haptic model predictive control for robot-assisted dressing," in *2018 IEEE international conference on robotics and automation (ICRA)*, pp. 1–8, IEEE, 2018.
- [11] A. Kapusta, Z. Erickson, H. M. Clever, W. Yu, C. K. Liu, G. Turk, and C. C. Kemp, "Personalized collaborative plans for robot-assisted dressing via optimization and simulation," *Autonomous Robots*, vol. 43, no. 8, pp. 2183–2207, 2019.
- [12] S. W. Brose, D. J. Weber, B. A. Salatin, G. G. Grindle, H. Wang, J. J. Vazquez, and R. A. Cooper, "The role of assistive robotics in the lives of persons with disability," *American Journal of Physical Medicine & Rehabilitation*, vol. 89, no. 6, pp. 509–521, 2010.
- [13] A. Candeias, T. Rhodes, M. Marques, M. Veloso, et al., "Vision augmented robot feeding," in *PROCEEDINGS of the European Conference on Computer Vision (ECCV)*, pp. 0–0, 2018.
- [14] A. Jardón, C. A. Monje, and C. Balaguer, "Functional evaluation of asibot: A new approach on portable robotic system for disabled people," *Applied Bionics and Biomechanics*, vol. 9, no. 1, pp. 85–97, 2012.
- [15] R. M. Alqasemi, E. J. McCaffrey, K. D. Edwards, and R. V. Dubey, "Wheelchair-mounted robotic arms: Analysis, evaluation and development," in *Proceedings, 2005 IEEE/ASME International Conference on Advanced Intelligent Mechatronics.*, pp. 1164–1169, IEEE, 2005.
- [16] P. Schrock, F. Farello, R. Alqasemi, and R. Dubey, "Design, simulation and testing of a new modular wheelchair mounted robotic arm to perform activities of daily living," in *2009 IEEE International Conference on Rehabilitation Robotics*, pp. 518–523, IEEE, 2009.
- [17] D. Gallenberger, T. Bhattacharjee, Y. Kim, and S. S. Srinivasa, "Transfer depends on acquisition: Analyzing manipulation strategies for robotic feeding," in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 267–276, IEEE, 2019.
- [18] D. Park, Y. K. Kim, Z. M. Erickson, and C. C. Kemp, "Towards assistive feeding with a general-purpose mobile manipulator," *arXiv preprint arXiv:1605.07996*, 2016.
- [19] T. Rhodes and M. Veloso, "Robot-driven trajectory improvement for feeding tasks," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2991–2996, IEEE, 2018.
- [20] Kinova, "Assistive solutions."
- [21] P. Health, "Obi robotic feeder."
- [22] A. Innovations, "Assistive innovations - iarm: Robotic arm for humans, mountable on powered wheelchair."
- [23] L. V. Herlant, R. M. Holladay, and S. S. Srinivasa, "Assistive teleoperation of robot arms via automatic time-optimal mode switching," in *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 35–42, IEEE, 2016.
- [24] M. Schukat, D. McCaldin, K. Wang, G. Schreier, N. H. Lovell, M. Marschollek, and S. J. Redmond, "Unintended consequences of wearable sensor use in healthcare: Contribution of the imia wearable sensors in healthcare wg," *Yearbook of medical informatics*, no. 1, p. 73, 2016.
- [25] K. Ward, F. Bertails, T.-Y. Kim, S. R. Marschner, M.-P. Cani, and M. C. Lin, "A survey on hair modeling: Styling, simulation, and rendering," *IEEE transactions on visualization and computer graphics*, vol. 13, no. 2, pp. 213–234, 2007.
- [26] S. Saito, L. Hu, C. Ma, H. Ibayashi, L. Luo, and H. Li, "3d hair synthesis using volumetric variational autoencoders," *ACM Transactions on Graphics (TOG)*, vol. 37, no. 6, pp. 1–12, 2018.
- [27] M. Zhang and Y. Zheng, "Hair-gan: Recovering 3d hair structure from a single image using generative adversarial networks," *Visual Informatics*, vol. 3, no. 2, pp. 102–112, 2019.
- [28] L. Hu, C. Ma, L. Luo, and H. Li, "Single-view hair modeling using a hairstyle database," *ACM Transactions on Graphics (ToG)*, vol. 34, no. 4, pp. 1–9, 2015.
- [29] M. Chai, L. Wang, Y. Weng, Y. Yu, B. Guo, and K. Zhou, "Single-view hair modeling for portrait manipulation," *ACM Transactions on Graphics (TOG)*, vol. 31, no. 4, pp. 1–8, 2012.
- [30] V. Maheu, P. S. Archambault, J. Frappier, and F. Routhier, "Evaluation of the jaco robotic arm: Clinico-economic study for powered wheelchair users with upper-extremity disabilities," in *2011 IEEE International Conference on Rehabilitation Robotics*, pp. 1–5, IEEE, 2011.
- [31] R. M. Haralick and L. G. Shapiro, "Image segmentation techniques," *Computer vision, graphics, and image processing*, vol. 29, no. 1, pp. 100–132, 1985.
- [32] S. Ghosh, N. Das, I. Das, and U. Maulik, "Understanding deep learning techniques for image segmentation," *ACM Computing Surveys (CSUR)*, vol. 52, no. 4, pp. 1–35, 2019.
- [33] F. N. Iandola, S. Han, M. W. Moskewicz, K. Ashraf, W. J. Dally, and K. Keutzer, "Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 0.5 mb model size," *arXiv preprint arXiv:1602.07360*, 2016.
- [34] U. R. Muhammad, M. Svanera, R. Leonardi, and S. Benini, "Hair detection, segmentation, and hairstyle classification in the wild," *Image and Vision Computing*, 2018.
- [35] J. Weickert, "Coherence-enhancing shock filters," in *Joint Pattern Recognition Symposium*, pp. 1–8, Springer, 2003.
- [36] G.-Z. Yang, P. Burger, D. N. Firmin, and S. Underwood, "Structure adaptive anisotropic image filtering," *Image and Vision Computing*, vol. 14, no. 2, pp. 135–145, 1996.
- [37] Z. C. Marton, R. B. Rusu, and M. Beetz, "On Fast Surface Reconstruction Methods for Large and Noisy Datasets," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, (Kobe, Japan), May 12–17 2009.
- [38] R. B. Rusu and S. Cousins, "3D is here: Point Cloud Library (PCL)," in *IEEE International Conference on Robotics and Automation (ICRA)*, (Shanghai, China), May 9–13 2011.
- [39] S. Russell and P. Norvig, "Artificial intelligence: a modern approach," 2002.
- [40] H. Heath and S. Cowley, "Developing a grounded theory approach: a comparison of glaser and Strauss," *International journal of nursing studies*, vol. 41, no. 2, pp. 141–150, 2004.
- [41] K. P. Hawkins, C.-H. King, T. L. Chen, and C. C. Kemp, "Informing assistive robots with models of contact forces from able-bodied face wiping and shaving," in *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication*, pp. 251–258, IEEE, 2012.
- [42] M. H. Raibert and J. J. Craig, "Hybrid position/force control of manipulators," 1981.
- [43] T. L. Chen, M. Ciocarlie, S. Cousins, P. M. Grice, K. Hawkins, K. Hsiao, C. C. Kemp, C.-H. King, D. A. Lazewatsky, A. E. Leeper, et al., "Robots for humanity: using assistive robotics to empower people with disabilities," *IEEE Robotics & Automation Magazine*, vol. 20, no. 1, pp. 30–39, 2013.
- [44] E. Cha, A. D. Dragan, and S. S. Srinivasa, "Perceived robot capability," in *2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pp. 541–548, IEEE, 2015.
- [45] T. L. Chen, C.-H. King, A. L. Thomaz, and C. C. Kemp, "Touched by a robot: An investigation of subjective responses to robot-initiated touch," in *2011 6th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 457–464, IEEE, 2011.
- [46] G. Trovato, R. Paredes, J. Balvin, F. Cuellar, N. B. Thomsen, S. Bech, and Z.-H. Tan, "The sound or silence: investigating the influence of robot noise on proxemics," in *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pp. 713–718, IEEE, 2018.
- [47] H. Knight, R. Thielstrom, and R. Simmons, "Expressive path shape (swagger): Simple features that illustrate a robot's attitude toward its goal in real time," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1475–1482, IEEE, 2016.
- [48] J. Rond, A. Sanchez, J. Berger, and H. Knight, "Improv with robots: Creativity, inspiration, co-performance," in *2019 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pp. 1–8, IEEE, 2019.
- [49] E. Deng, B. Mutlu, and M. Mataric, "Embodiment in socially interactive robots," *arXiv preprint arXiv:1912.00312*, 2019.