AN INTELLIGENT TOOTHBRUSH:
Machines for Smart Brushing

Anna Flagg,* Jennifer Boger, and Alex Mihailidis
Intelligent Assistive Technology and Systems Lab
Department of Occupational Science and Occupational Therapy, University of Toronto
*me@annaflagg.com

ABSTRACT

The elderly is the fastest growing portion of our population, worldwide. As our society ages, the need for support for the increasing numbers of those suffering from cognitive impairments increases too. This research examines the essential daily living task of brushing teeth, which can be difficult or impossible for older adults with dementia to complete on their own. Studies have shown, however, that these difficulties can often be overcome given a few simple guiding prompts. This paper presents a vision-based toothbrush tracker, a first step in the development of an intelligent system for automated, context-aware prompting.

INTRODUCTION

The last century has seen an increase in life expectancy of over 50%, in large part a happy result of modern medical advances [1]. At the same time, however, we witness an accompanying strain on the healthcare system, caregivers, patients and families to handle the increase in older adults and the health problems that come with aging. This has set in motion a race to develop an environment to support quality of life of an aging population.

People with dementia often encounter difficulties remembering how to perform essential daily living tasks, leading to loss of independence, and possible relocation to long-term care facilities. Studies have shown that in some cases the use of simple prompts can help alleviate these difficulties for some tasks, such as handwashing [2]. This suggests that automated prompting could be an effective way to increase the independence of these older adults. Such an approach would also support aging-in-place, which is beneficial in many respects: people tend to be happier, the familiar environment acts as a valuable tool for memory and cognition, and automation for simpler tasks eases caregiver burden. It is the goal of this project to support aging-in-place with the development of an intelligent aid for brushing teeth.

Relevant work in assistance for brushing teeth include Chang et al.'s Playful Toothbrush [3], and Newsome's Toothbrush Timer [4]. Both are game-like teaching aids for children, the first a vision-based tool and the second an iPad application. Similar to the purposes of this work, the Playful Toothbrush infers real-time human behavior during brushing. However, it requires a specialized brush with a large box of LEDs attached to its end. This would be inappropriate for this application because older adults with dementia tend to have difficulty recognizing and using unfamiliar objects, and therefore a toothbrush for this population needs to fit gracefully into the person's long-term memory of what a toothbrush conventionally looks like. The Toothbrush Timer is of interest because although it is not interactive, it is an effective teaching tool for autistic children, and ideally the work presented here could eventually be extended to the autism community as well.

AN ASSISTIVE PROMPTER

The goal of this work is the development of an intelligent prompting application for brushing teeth. The final system will model human behavior through vision-based analysis of motion and interaction of important points of interest in the scene, including facial features, toothbrush, toothpaste, towel, etc. The perceived behavior will then be fed into a Partially Observable Markov Decision Process (POMDP) to create a context-aware prompter. This paper presents the preliminary work completed towards the first step in the system, the object (i.e., toothbrush) tracking problem.

A TOOTHBRUSH TRACKER

Design approach

A key challenge in designing for persons with cognitive impairments is to avoid intrusiveness. Ideally, the design should be zero-effort technology that is built into the natural home environment, remaining invisible until needed, and disrupting one's normal routine as little as possible. With these principles in mind, we designed a color-based algorithm that works via an inexpensive webcam, requiring only that the toothbrush be striped with two alternating colors that run circularly, perpendicular to its length (see Figure 1).

Figure 1: Prototype of the smart toothbrush

At all times the application watches for movement, requiring no interaction from the person brushing, and no intrusive markers or sensors of any kind are worn on the person's body. The bi-color striping method maintains the recognizability of the toothbrush and provides a level of protection against distractors and occlusions, important features for a toothbrush tracker because it means different hand positions during brushing are not a problem. Additionally, this approach is low-tech, inexpensive, and easy to implement in real time, face and feature tracking included.

The algorithm

As shown in Figure 2, the toothbrush tracking application consists of three sub-trackers. The first detects and then tracks the person's face. The second locates the person's mouth, which is necessary so that the system knows that the person is likely to be actually brushing her or his teeth, as opposed to just moving the toothbrush around in space. The third sub-tracker locates and follows the actual toothbrush.

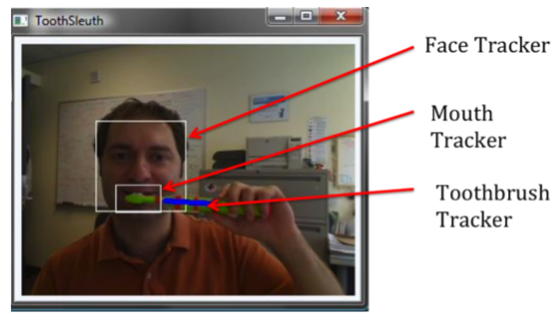

Figure 2: The overall tracking module consists of three sub-trackers: 1) face sub-tracker; 2) mouth sub-tracker; and 3) toothbrush sub-tracker.

The first step in detecting the striped toothbrush was to find all instances of the brush's two colors in the image. Then, because of the toothbrush's elongated shape, a line fitting approach was used on these colored points to make some plausible guesses at what might be a toothbrush. This was accomplished with the Hough transform, an image analysis technique for generalized feature extraction [5]. Stripes were chosen because they are seen as several disjointed, roughly collinear blobs by the computer, which is exactly the type of input that the line fitting routine is attracted to. It is also an unlikely pattern to exist elsewhere in the image, and is robust to a hand occluding part of the brush. This is because no matter how the brush is held, several stripes of both colors are generally still visible somewhere along the brush's length. Then all that remained was to perform face and feature detection with existing support from OpenCV, an open source computer vision library.

Color tracking

At present the tracker supports red, green and blue colors for toothbrushes (though it could be extended to others). The general routine to detect a color is to check if its corresponding channel is significantly bigger than the others. For example, to find the color red the tracker finds all pixels that have significantly larger values in the red channel than in green and blue combined.

Simplifying and fitting lines

The Hough transform was used here to find lines. Since it is designed to work on points, any connected component of detected color was simplified to a single point. This reduction also helped to reduce false positives because most distracters are single amorphous blobs of color, rather than several disjointed, collinear shapes. This meant that most distracters were reduced to just one point, and were therefore weak candidates for line fitting. Figure 3 shows an example of fitting lines to points extracted by color matching using a Hough transform.

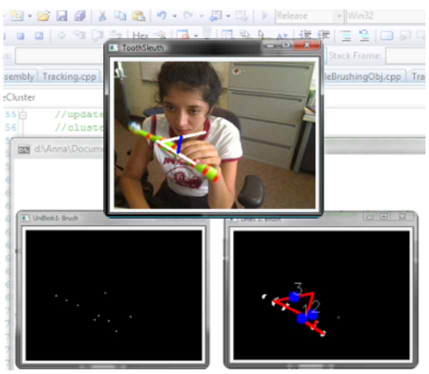

Figure 3: Hough transform line fitting in two colors

Finding a match

A toothbrush detection depended on the tracker's confidence that it had indeed found a pair of lines in the corresponding colors that described the same object, and that this was the same object that was seen in the previous frame (if any). For each line pair a similarity score was calculated based on if the two lines had similar angles, small distance between midpoints, small perpendicular distance and small difference in length. A pair was considered a good match if its similarity measure was below a certain threshold value, which was selected simply by observing which value maximized tracker performance in lab trials.

Face and feature tracking

OpenCV has methods for mouth detection, but the mouth is a small feature, and often turns up lots of false positives in a large image. Thus face detection was performed first as a means to improve results by narrowing down the search space for the mouth. At the same time, any positive detection of a face was tested by requiring that it contain facial features, including mouth and eyes. This test in turn helped reduce false positives in the face tracking results. A positive face detection confirmed by eye and mouth detections is shown in Figure 4.

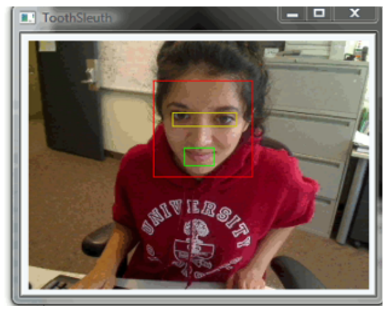

Figure 4: OpenCV face, eyes and mouth detection working together

PRELIMINARY RESULTS

In laboratory tests, the tracker seemed to be able to correctly find the prototype toothbrush the majority of the time, as well as identify when there was no toothbrush visible in the scene. It could handle even quite extensive hand occlusions, and rarely accidentally tracked similarly colored objects. Further tests are needed to define formal error figures, but informal observations of many lab trials suggested that the tracker could estimate the angle and position of the brush quite accurately, to within about 15 degrees and 10 pixels, respectively.

However the tracker operates in only two dimensions, and attempts no depth measurements. Therefore it could not, for instance, differentiate between proper brushing action or if the brush was simply moved within the mouth area. It often failed to measure the toothbrush length accurately, usually underestimating, especially if part of the brush was inside the mouth. The tracker results could also be somewhat unstable-it tended to jump around between slightly different midpoints on the brush due to small lighting changes. Another problem manifested when the brush was pointing directly at the camera, because the stripes were not visible, and so the algorithm could not see the brush. Finally, selection of the threshold value mentioned above, as well as other constants in the code, would often have a significant effect on the results. Different locations and lighting conditions often required manual calibration of these values.

Further work to address these failings could include integrating stabilizing components like a Kalman filter, or adding more constraints to the existing similarity function such as a prior model that "remembers" how the brush has been moving in the last several frames (right now it only considers the previous frame). More improvements could include adding support for automatic calibration, coupling the algorithm with another unobtrusive sensor such as an accelerometer, and possibly extending to 3D space.

CONCLUSIONS

Preliminary tests in a lab setting showed the above approach to be a robust way to track a toothbrush during brushing. The next step is to test our toothbrush tracker outside of laboratory conditions. We can then move on to the problem of using an artificially intelligent agent to perform customized action recognition for the steps involved in brushing teeth. Once the system is intelligent enough to recognize human behaviors, it could be integrated into a model such as a POMDP to create complete task recognition and guidance.

To conclude, we hope that by supporting the essential daily living task of brushing teeth, this work can help increase the independence and quality of life of the elderly.

REFERENCES

[1] Centers for Disease Control and Prevention, "United States Life Tables," National Vital Statistics Report, 2007.

[2] Mihailidis, Alex, Boger, Jennifer, Craig, Tammy, and Hoey, Jesse, "The COACH prompting system to assist older adults with dementia through handwashing: An efficacy study," BMC Geriatrics, 2008.

[3] Chang, Yu-chen, Huang, Chao-ju, Lo, Jin-ling, Chu, Hao-hua, "A Playful Toothbrush to Motivate Proper Brushing for Young Children," National Taiwan University, 2007.

[4] Newsome, Ryan, "Toothbrush Timer," itunes.apple.com/us/app/toothbrush-timer/id307084494?mt=8, 2010.

[5] Ballard, D.H., "Generalizing the Hough transform to detect arbitrary shapes," Pattern Recognition, Vol. 13, Issue 2, pg. 111-122, 1981.