Ready—Aim—Fly! Hands-Free Face-Based HRI for 3D Trajectory Control of UAVs

Jake Bruce, Jacob Perron, and Richard Vaughan
Autonomy Laboratory, School of Computing Science
Simon Fraser University
Burnaby, Canada
{jakeb, jperron, vaughan}@sfu.ca

***Abstract*—We present a novel user interface for aiming and launching flying robots on user-defined trajectories. The method requires no user instrumentation and is easy to learn by analogy to a slingshot. With a few minutes of practice users can send robots along a desired 3D trajectory and place them in 3D space, including at high altitude and beyond line-of-sight.**

With the robot hovering in front of the user, the robot tracks the user's face to estimate its relative pose. The azimuth, elevation and distance of this pose control the parameters of the robot's subsequent trajectory. The user triggers the robot to fly the trajectory by making a distinct pre-trained facial expression. We propose three different trajectory types for different applications: straight-line, parabola, and circling.

We also describe a simple training/startup interaction to select a trajectory type and train the aiming and triggering faces. In real-world experiments we demonstrate and evaluate the method. We also show that the face-recognition system is resistant to input from unauthorized users.

***Keywords*-human robot interaction; unmanned aerial vehicle; face recognition;**

I. Introduction

Recent innovations have brought us very capable, small-size, and low-cost unmanned aerial vehicles (UAVs). These have many applications and new industries forming around them. Current commercial operator control interfaces for UAVs use either a dedicated hardware controller or functionally similar software running on a tablet or smartphone. As a long-term research program we are interested in methods for robots to work together with humans outdoors in large environments. We and other researchers have begun to conduct user studies that investigate how untrained users choose to interact with robots using only 'natural' interfaces where the human participant is entirely uninstrumented, i.e. they carry no equipment, their appearance is unaltered, and little training is required [1], [2]. The long term goal of this work is to enable people to interact with robots and AIs as we now interact with people and trained animals, just as long imagined in science fiction. Further, the main interaction we propose here is hands-free, which is valuable for applications where the user's hands are busy.

This paper proposes a novel human-robot interaction (HRI) system for controlling the flight trajectory of a UAV by direct face-to-face engagement. The interaction system only requires the commodity sensors available onboard vehicles that cost a few hundred dollars today. Our image processing is currently done off-board over a wireless network, but the necessary computation power will be available onboard even low-cost UAVs in the very near future.

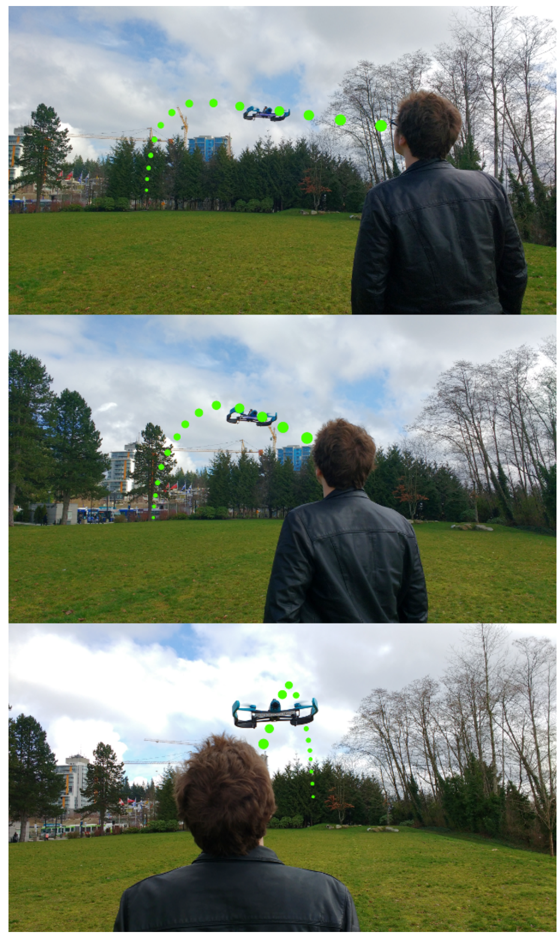

Fig. 1: Parabolic UAV trajectories being aimed by the user's face as described in this paper. The vector between the face and the robot determines the launch angle, and the size of the face determines the distance.

The proposed HRI system uses visual and orientation sensors onboard a quadrotor UAV to *1)* learn the identity and facial expressions of a user, *2)* accept user input through touch interaction, and *3)* aim the flight trajectory of the robot based on the relative pose of the user and the robot. This interaction is modelled on the act of drawing a bow or slingshot, in which the launch azimuth and elevation angles are lined up by eye, and the launch power is set by the magnitude of the draw. The 'slingshot' interaction allows the user to send the robot into predictable trajectories in 3D space without a hardware controller. We suggest three different trajectory types: straight-line, parabola, and circling, that have different applications.

Training interactions provide adaptation to a specific user and add robustness to environment and lighting conditions that may otherwise defeat current face recognition systems.

The contributions of this paper are the slingshot interaction and related recipes for face-based UAV interaction. In particular, we describe the benefits of a runtime-determined facial expression as a signalling method, and demonstrate these control signals to launch the robot at specific target locations on desired trajectories. Although we demonstrate with quadrotor UAVs, the method can be broadly applied to control robot trajectories in other domains.

II. RELATED WORK

Existing work on uninstrumented HRI with UAVs has focused on gesture-based interfaces [3], [4], [5], [6], although voice [7], [8], [9] and touch-based [10] interactions have also been proposed. Face detection has also been used to interact with UAVs: in [9], selection between multiple robots is accomplished with face engagement as an attentional cue; in [5], the face determines where to look for human gesture motion. Davis et al. [11] investigate face recognition for security applications involving UAVs. Bold et al. [12] investigate the performance of face recognition on consumer UAVs as a function of distance and angle to the human. Oreifej et al. [13] consider human identity recognition from the air, but they are concerned with scales at which face recognition is not feasible.

The method proposed in this paper uses the line-of-sight vector from the user's face to the robot to indicate a direction (see Figure 1). This has been investigated in related work involving gaze direction [14][15], and is similar to pointing gestures [16] in which a line is drawn between the user's eye and fingertip, communicating a pointing vector to a robot. However, to our knowledge, ours is the first demonstration of a pointing-like gesture to control a UAV using only monocular camera sensing.

III. SYSTEM

The proposed system involves three main phases:

- *Ready*— the user's identity and facial expressions are learned and input is provided through touch-based interaction
- *Aim*— the robot starts flying and keeps its user centered in its camera view, while the user lines up the trajectory and chooses its power by "drawing back" analogous to firing a bow or slingshot
- *Fly!*— the user signals to the robot to begin a preset parameterized trajectory. The robot executes the trajectory with parameters observed at the end of the *Aim* phase. Below we demonstrate three possible trajectory types.

Fig. 2: Example learned faces taken from our experiments. The top row shows the *neutral* face, and the bottom row shows the *trigger* face that is used to send the start signal. The two columns on the left were captured with the robot in the user's hand, while the two on the right were captured in flight. Note the low image quality: the off-the-shelf face recognition software was able to handle the poor imagery up to a distance of several meters, generalizing successfully between scales.

In this section we describe each phase of the system in detail and the components involved at each step.

A. Facial Expressions—Ready

Face recognition techniques require training to identify users. Training can be done in a separate phase and identity models can be assumed *a priori*, but there are advantages to including the training phase at runtime. In particular, the system will learn a model that is automatically calibrated to the current environment, lighting conditions, and transient appearance details of the user such as clothing, hairstyle, and worn accessories. In our demonstration system, the interaction begins with the robot in the user's hand and the user's identity is learned to sufficient confidence (100% precision over a 3 second window in our case) before the robot takes flight. This enables a degree of confidence that the robot will be able to find its user again once it has launched. Training time is typically less than one minute, and can be faster if it is the same user from a previous session.

We use the ROS *face_recognition* package [17] to learn the user's identity with a conventional camera sensor onboard the robot. As face recognition software is not yet meeting human performance, false positives can be expected that could cause the robot to take commands from other humans or human-like objects. In order to mitigate this weakness in the current state of face recognition technology, we borrow the concept of two-factor authentication where face identity alone is not sufficient to control the robot. Although a sensor modality independent from vision would be ideal for this purpose, once in flight, vision is the most feasible conventional sensor for detecting and interacting with humans. To address dual problems of signalling to a flying robot and providing robustness to false positives, we introduce the concept of a runtime-determined *trigger expression*.

Fig. 3: The two stages of the *Ready* phase: (left) learning the *neutral expression*; (right) learning the *trigger expression*.

A convenient side effect of current off-the-shelf face recognition is that the same user showing dramatically different facial expressions (Figure 2) can be detected as distinct face identities. We exploit this by training first on the user's *neutral expression* and, upon reaching 100% precision over a rolling time window, informing the user to choose a *trigger expression* (Figure 3). The *trigger expression* functions as a signal to launch the flight behavior, and incidentally functions as a security measure to ensure that other humans in the area who do not know the chosen expression cannot command and recover the robot.[1]

The user rotates the robot in their hand to signal that their *trigger expression* is being displayed and the system trains to recognize this new expression to 100% precision. The direction of rotation is remembered during the *Fly!* phase, where it can determine parameters of the flight path (such as direction to circle). Once training is complete, the robot takes flight and the *Aim* phase begins.

B. Face Position—Aim

Once the robot has begun flying, it waits for the learned user to appear in its view. On seeing the user's face, the robot controls its yaw heading to keep the human centered in its field of view. Adjusting its yaw heading helps keep the user visible and most importantly, allows the user to create a line-of-sight vector between the center of their face and the robot. This vector defines the direction of flight in the next phase. In addition to the location of the face, we use an analogy to shooting a bow or slingshot for distance control: in shooting a bow, the power of the shot is determined by how much energy is put into the arrow, usually controlled by how far the string is drawn back. We borrow this concept by determining the magnitude of the flight path with the size of the user's face in the image: a smaller face means the user is farther away and more power is put into the shot, so the robot travels farther in the *Fly!* phase.

The *trigger expression* learned in the first phase acts as the launch signal. Once the robot detects this expression, it launches on the trajectory defined by the line-of-sight vector, controlled by the *Fly!* phase.

C. Trajectory Execution—Fly!

Once the trigger signal has been received, the UAV executes the flight path. We demonstrate three different trajectories by analogy to different familiar shooting/throwing modalities: the *beam*, the *slingshot*, and the *boomerang*.

1) Beam: a straight path along a specified azimuth and elevation. The distance along this path is determined by modulating a predefined base distance with the size of the user's face on launch signal, a smaller face indicating more power and a greater distance. Once reaching the location the robot can perform a specific behavior such as video capture or mapping, but this is application-dependent and is outside the scope of this paper. The utility of *Beam* is that it allows the robot to be sent to a location in 3D space at arbitrary altitude, for example above a distant building.

2) Slingshot: a ballistic trajectory launched at the angle specified by the user's line-of-sight vector to the robot. The robot's vertical velocity is decreased at a constant rate during the flight path, thereby following a parabolic arc analogous to a thrown projectile, until reaching the original launch altitude at which point the robot executes an application specific behaviour. The size of the user's face on launch signal determines the distance covered by this ballistic arc, with smaller faces indicating a more powerful launch. The utility of *Slingshot* is that it can send a robot over an arbitrary-sized vertical obstacle and into an area out of line of sight, perhaps to perform video capture.

3) Boomerang: a circular path tangent to the line-of-sight vector between the user and the robot, analogous to a boomerang that is thrown forward and curves around to return to the thrower. The direction of the circle (curving leftward or rightward) is determined by the direction of rotation during the touch interaction in the *Ready* phase, and the radius of the circle is determined by the size of the user's face on launch signal. Smaller faces indicate a more powerful throw, and therefore a larger circle. The utility of *Boomerang* is to define a survey path, perhaps to obtain a 360-degree scan of an object such as a building or statue.

IV. EVALUATION

We performed simple experiments with a commercially available $800 UAV platform to test the three proposed behaviors in addition to the security aspect of the face recognition. In this section we describe our implementation and demonstrate and evaluate several aspects of its performance. Video of the demonstrations can be seen at https://youtu.be/sHkcVIJt2_Y.

[1]The interaction method described in this paper generalizes straightforwardly to other types of gestures: the system can be modified and extended with a variety of other physical interactions as gesture recognition technology improves.

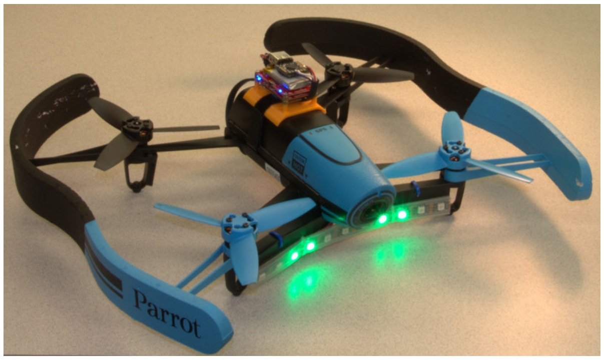

Fig. 4: Parrot Bebop quadrotor used in our experiments, equipped with a programmable colored LED strip for visual feedback to the user.

Participant	Successes	Mean Time (s)	Std. dev. (s)
1	20/20	3.49	1.61
2	18/20	8.63	9.61
3	20/20	3.57	0.71
Overall	58/60	5.23	6.13

TABLE I: Results of face recognition trials.

Participant	Successes	Mean Error (°)	Std. dev. (°)
1	6/8	12.78	9.20
*2	8/8	6.78	1.87
3	7/8	15.21	8.39
4	7/8	20.09	17.25
Overall	28/32	13.71	11.70

TABLE II: Results of *beam* trials. The participant marked with the asterisk (*) was an expert user—the developer of the system.

Robot

The robot for these demonstrations is a Parrot Bebop quadrotor (Figure 4), which has a well-stabilized frontal camera with a resolution of 640x368 at 30 frames per second and high-quality onboard orientation and velocity estimation. To control the robot and receive sensor data we used the open-source `bebop_autonomy`[2] ROS package. Due to the built-in image stabilization, the roll rotation during the *Ready* phase is almost imperceptible in the video feed, greatly simplifying the image processing and the use of orientation as a user signal. We augmented the robot with a programmable colored light strip for user feedback and debugging. Specifically, the light strip indicates training progress and request for the *trigger expression* during the first phase. Video processing and control computation was performed off board over the built-in wireless network of the Bebop on a consumer-quality laptop with a quad-core Intel processor at 2 GHz and 8 GB of RAM.

Evaluation—Face Recognition

We first evaluated the security aspect by determining whether the robot would reject face command attempts from a human who was not the learned user. We performed 20 trials with each of 3 users, in which the robot would go through the training phase and then hover, and only land when presented with the correct user's face. Half of the trials included only the learned user, and the other half had distractors attempting to command the robot to land for 30 seconds before the user entered the frame. We considered a trial to be successful if the robot accepted a command from only the authorized learned user, and a failure if it ever accepted a command from a non-trained (unauthorized) user. Success rates and time taken are reported in Table I. We observed two failures in 60 trials.

Evaluation—Beam

To evaluate the straight-line trajectory, we had the user attempt to launch the UAV through a 0.8 m diameter hoop located 8 meters from the user in our indoor motion-capture lab. Target hoops at four different locations were used, each at a different bearing and azimuth. 8 trials were performed per user, twice for each target hoop. We compared the resulting robot trajectory measured using an external Vicon motion capture system with the ideal trajectory. The results are reported in terms of the error in azimuth and elevation against the true angle to the target hoop. The robot did not fly perfectly straight due to the inevitable errors of real-world robotics, so we compute angle error based on a least-squares line through the trajectory as determined by the first principal component of a singular value decomposition (SVD) on the array of positions for each trajectory.

We repeated these trials with 4 users. Although trajectories were not perfect, they were quite good, with most having a root-mean-square deviation of less than 10 centimeters from the line of best fit. We considered a trial successful if the angle error to the target was the lowest among the four hoops, and we also report the error values. Results are summarized in Table II, and two sample trajectories are shown relative to the four hoops in Figure 5.

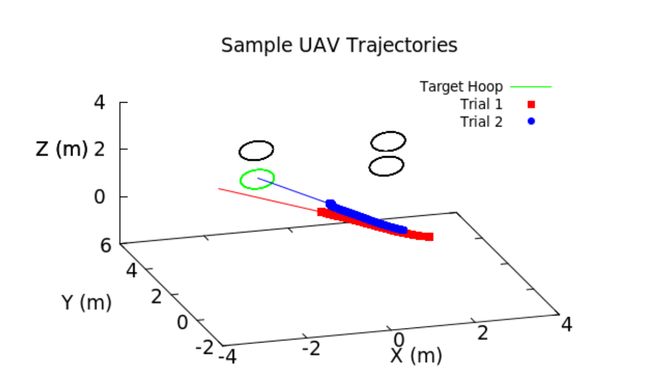

Fig. 5: Two sample trajectories from the *beam* experiments, with the target hoop in green. Best fit lines show that the robot would have flown very close to the target, although trajectory lengths were limited in these trials for safety reasons. The trajectories shown are taken from the first two trials of Participant 1.

[2]http://github.com/AutonomyLab/bebop_autonomy

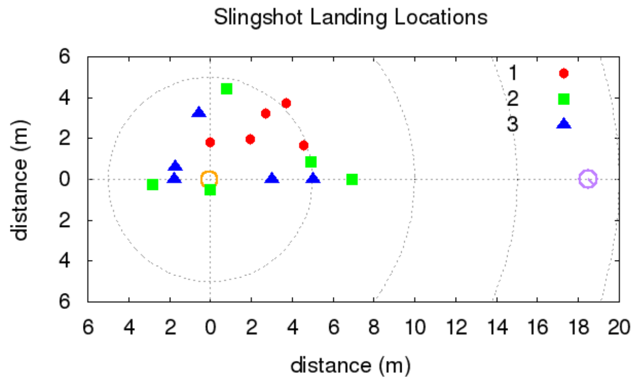

Fig. 6: Landing locations of the robot in the *slingshot* trials. The initial pose of the robot is shown on the right in purple (note the orientation of 45° away from the target hoop), and the orange circle shows the location and diameter of the target hoop. The three different glyphs identify the different users.

Evaluation—Slingshot

We evaluated the ballistic trajectory in an outdoor environment, by attempting to launch the UAV in an arc to land in a 0.8 meter diameter hoop located on the ground 18.5 meters from the robot. The user stood 1.5 meters in front of the robot. The base distance of the system was calibrated such that the largest face size would result in a flight distance of 3 meters, and the smallest resulting in a distance of 45 meters. Due to the small size of the target relative to the traversal distance, we considered a trial successful if the robot landed with 5 meters of the center of the hoop. We conducted 5 trials with each of 3 users and measured the accuracy of the robot's end location in meters from the center of the hoop. The robot began each trial with a yaw of 45° to the target, to ensure that users had to aim both the angle of the trajectory and the distance. Results are shown in Table III, and landing locations are shown in Figure 6.

Participant	Successes	Mean Error (m)	Std. dev. (m)
1	4/5	3.77	1.30
2	4/5	3.96	2.15
3	4/5	2.99	1.17
Overall	12/15	3.57	1.65

TABLE III: Results of *slingshot* trials.

We also performed *slingshot* trajectories outdoors over distances of up to 60 meters and out of line-of-sight. Without accurate ground-truth (GPS data were not accurate enough), we are not able to provide quantitative evaluations. As an example, the system was tested in the short-range scenario depicted in Figure 7a with the goal of taking a photo of a target object that is occluded from the user's view. The user signals to the robot to fly a *slingshot* trajectory over a dirt pile (7b), the robot takes a picture (7c), and returns to the user.

Demonstration—Boomerang

For the *boomerang* trajectory, we conducted simple outdoor demonstration trials in which the user attempted to send the UAV on a circular path in a particular direction, ending the trajectory back near the user. We had no clear evaluation metric for these demonstrations, but footage is included in the accompanying video, and Figure 8 illustrates a sample trajectory successfully orbiting a statue at close range.

V. DISCUSSION

Although this work is preliminary and should be considered primarily a set of recipes for interaction with UAVs, the results are encouraging. In the security evaluation we found that once trained to 100% precision in the *Ready* phase, the face recognition false positive rate during flight was very low. This is a good result, because although false negatives can delay a successful interaction, false positives can be catastrophic due to the initiation of an arbitrary, undesired trajectory. It is also an advantage for a UAV to accept commands only from its known user since this provides a degree of security, and with only 2 failures out of 60, we can be relatively confident in this aspect of the system.

In the *beam* experiments, we found that users were able to send the UAV toward the correct target hoop in 87.5% of the trials performed, with a mean angle error to the target of 13.7 (+/- 11.7) degrees. This is sufficient for positioning the robot at a coarse location in 3D space, and represents a new option for a hands-free natural interface to control a UAV for video capture or to initiate mapping. We also observed that expert users such as the system developer performed significantly better than untrained users, which indicates that users get better through experience with the system. This suggests the possibility of future work evaluating the performance gains for new users during the training process.

The *slingshot* experiments show agreement with the angular predictions of the *beam* experiments based on the spread of landing locations around the target hoop and demonstrate that users are able to use the slingshot metaphor to produce trajectories that land in the vicinity of the target. In only one trial, the robot came very close to landing inside of the hoop. However, given a range of possible distances from 3 to 45 meters depending on the size of the face, given the starting distance of 18.5 meters from the target, and given the lack of any robot-to-user distance feedback during the *Aim* phase, an average position error of 3.5 meters from the target is another encouraging result. The users were essentially "shooting blind": in video games, trajectories are often indicated by overlaying a virtual arc on the screen while aiming (similar to Figure 1), and it seems reasonable to expect that a similar augmented-reality aspect in a system like this could improve performance. In the absence of feedback, hours of practice may be required before developing a sense of the distances produced (as in shooting a real slingshot), and the other users in these trials had never used the system before.

For the demonstration of the *boomerang* trajectory, we successfully orbited a statue with the UAV using the face-

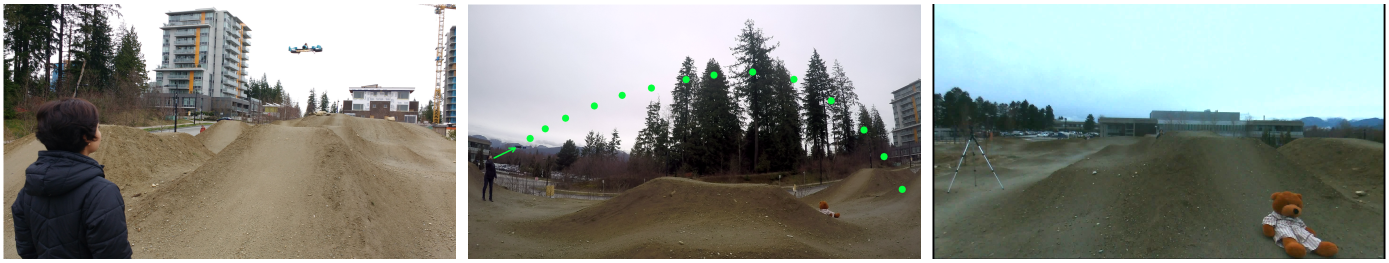

(a) Target is occluded from the user by a dirt pile. (b) *Slingshot* trajectory commanded by user. (c) An image acquired by the robot.

Fig. 7: A example scenario where the user commands the robot to inspect an occluded area.

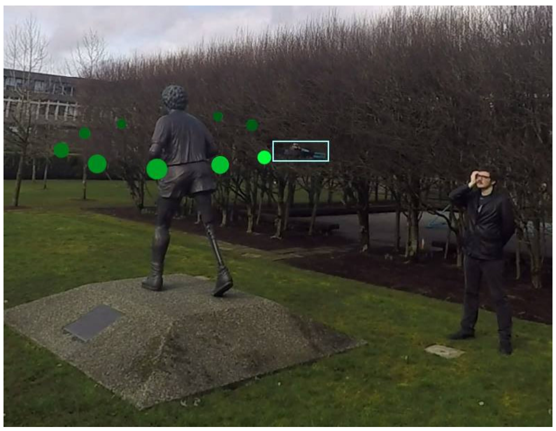

Fig. 8: User signalling the robot to begin the *boomerang* trajectory orbiting a statue, with robot highlighted in blue.

based control method in this paper to specify the direction and size of the circle, and using the orientation of the robot in the hand to determine the direction of revolution. Although we were interested only in investigating the behavior of the robot, the camera data from the robot during such an orbit would be useful for such purposes as creating a 3D reconstruction of the orbited statue.

VI. CONCLUSIONS AND FUTURE WORK

We propose and demonstrate a face-based system for uninstrumented HRI with a UAV in flight: the first system to our knowledge that uses only face recognition to send commands to a flying robot. Once trained on two facial expressions, the system uses one as a trigger, and the location of the face determines the angle and power of the flight. In addition to the signalling function of the facial expression system, it provides a degree of security to prevent false positives in face identity recognition from triggering behavior: the trigger expression is required before initiating motion, and this expression is known only by the primary user.

We plan to explore the utility of the *boomerang* trajectory for mapping a desired area, and for orbiting a target object or person in order to create a 3D model of the target. We would also like to conduct a user study to determine whether users would prefer the proposed slingshot-analogue method for setting trajectory power over the inverse design where smaller faces indicate less power, and as mentioned above, determine also how well user performance increases with experience.

Other potential improvements include using a dedicated facial expression detector as opposed to our dual-identity face recognition method. In the future, face recognition software is likely to be increasingly invariant to facial expression, but powerful software will no doubt continue to be developed for recognizing facial expressions. As mentioned in the discussion, we expect that including a form of augmented-reality feedback to the user during the *Aim* phase would greatly improve user's performance in judging the distance of a given trajectory.

This interaction design has the UAV flying very close to the user, which has potentially serious safety implications. For user safety in these experiments we rely on the inherent safety of the lightweight Bebop vehicle. But the general problem of having UAVs actively maintain safety when working closely around people is important for future work.

While the demonstrations in the paper have sent the robot on flights of 10 meters indoors (the range of our motion capture system) and 45 meters outdoors (the size of our local field), these interactions scale to hundreds of meters without modification. If the UAV was able to visually servo to a target of interest after reaching the peak of its trajectory (for example another person, as described in another paper under review) we might be able to "throw" the UAV from one person to another over a kilometer or more.

Finally, and informally, we assert that using the robot in this way is *fun*, so this interaction could have applications in entertainment.

ACKNOWLEDGMENT

This work was supported by the NSERC Canadian Field Robotics Network (NCFRN).

REFERENCES

[1] J. R. Cauchard, L. Jane, K. Y. Zhai, and J. A. Landay, "Drone & Me: an exploration into natural human-drone interaction," in *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 2015, pp. 361–365.

[2] S. Pourmehr, J. Thomas, and R. Vaughan, "What untrained people do when asked âĂIJmake the robot come to youâĂİ (late-breaking abstract)," in *Proceedings of the 2016 ACM/IEEE International Conference on Human-robot Interaction*, ser. HRI '16, 2016.

[3] W. S. Ng and E. Sharlin, "Collocated interaction with flying robots," in *RO-MAN, 2011 IEEE*. IEEE, 2011, pp. 143–149.

[4] K. Pfeil, S. L. Koh, and J. LaViola, “Exploring 3d gesture metaphors for interaction with unmanned aerial vehicles,” in *Proceedings of the 2013 international conference on Intelligent user interfaces*. ACM, 2013, pp. 257–266.

[5] M. Monajjemi, J. Wawerla, R. Vaughan, and G. Mori, “HRI in the Sky: Creating and commanding teams of UAVs with a vision-mediated gestural interface,” in *Intelligent Robots and Systems, 2013 Int. Conf. on*, Nov 2013, pp. 617–623.

[6] M. Monajjemi, J. Bruce, S. A. Sadat, J. Wawerla, and R. Vaughan, “UAV, Do You See Me? establishing mutual attention between an uninstrumented human and an outdoor UAV in flight,” in *Intelligent Robots and Systems (under review), 2015 Int. Conf. on*. IEEE, 2015.

[7] M. Draper, G. Calhoun, H. Ruff, D. Williamson, and T. Barry, “Manual versus speech input for unmanned aerial vehicle control station operations,” in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 47, no. 1. SAGE Publications, 2003, pp. 109–113.

[8] R. W. Beard, D. Kingston, M. Quigley, D. Snyder, R. Christiansen, W. Johnson, T. McLain, and M. Goodrich, “Autonomous vehicle technologies for small fixed-wing uavs,” *Journal of Aerospace Computing, Information, and Communication*, vol. 2, no. 1, pp. 92–108, 2005.

[9] S. Pourmehr, V. M. Monajjemi, S. A. Sadat, F. Zhan, J. Wawerla, G. Mori, and R. Vaughan, “"You Are Green": A touch-to-name interaction in an integrated multi-modal multi-robot HRI system (late-breaking abstract),” in *Proceedings of the 2014 ACM/IEEE International Conference on Human-robot Interaction*, ser. HRI ’14. New York, NY, USA: ACM, 2014, pp. 266–267. [Online]. Available: http://doi.acm.org/10.1145/2559636.2559806

[10] S. Pourmehr, V. M. Monajjemi, R. T. Vaughan, and G. Mori, “"You two! Take off!": Creating, modifying and commanding groups of robots using face engagement and indirect speech in voice commands,” in *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS’13)*, Tokyo, Japan, November 2013.

[11] N. Davis, F. Pittaluga, and K. Panetta, “Facial recognition using human visual system algorithms for robotic and uav platforms,” in *Technologies for Practical Robot Applications (TePRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 1–5.

[12] S. Bold, B. Sosorbaram, and S. R. Lee, “Implementation of autonomous unmanned aerial vehicle with moving-object detection and face recognition,” in *Information Science and Applications (ICISA) 2016*. Springer, 2016, pp. 361–370.

[13] O. Oreifej, R. Mehran, and M. Shah, “Human identity recognition in aerial images,” in *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. IEEE, 2010, pp. 709–716.

[14] A. Couture-Beil, R. T. Vaughan, and G. Mori, “Selecting and commanding individual robots in a vision-based multi-robot system,” in *Proceedings of the 5th ACM/IEEE international conference on Human-robot interaction*. IEEE Press, 2010, pp. 355–356.

[15] C. Rich, B. Ponsler, A. Holroyd, and C. L. Sidner, “Recognizing engagement in human-robot interaction,” in *Human-Robot Interaction (HRI), 2010 5th ACM/IEEE International Conference on*. IEEE, 2010, pp. 375–382.

[16] S. Pourmehr, M. Monajjemi, J. Wawerla, R. T. Vaughan, and G. Mori, “A robust integrated system for selecting and commanding multiple mobile robots,” in *Proceedings of the IEEE International Conference on Robotics and Automation*, May 2013.

[17] S. Emami and V. P. Suciu, “Facial recognition using OpenCV,” *Journal of Mobile, Embedded and Distributed Systems*, vol. 4, no. 1, pp. 38–43, 2012.