Gesture Recognition For UCAV-N Flight Deck Operations

Problem Definition Final Report

Prepared For:
Van Olinger
DARPA UCAV Program Manager

Prepared By:
Naval Air Systems Command

23 January 2003

20040511 000

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

Best Available Copy

REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704-0188*

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE *(DD-MM-YYYY)*
23-01-2003

2. REPORT TYPE
Final Report

3. DATES COVERED *(From - To)*
Jul 02 to Dec 02

4. TITLE AND SUBTITLE
Gesture Recognition For UCAV-N Flight Deck Operations
Problem Definition Final Report

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Venetsky, Larry; Husni, Mark; Yager, Mark

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Naval Air Systems Command, Code 4.8T, Lakehurst, New Jersey 08733
DeveloSoft Corporation, 1630 30th St., Boulder, Colorado 80301

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Defense Advanced Research Projects Agency, Van Olinger

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
DISTRIBUTION STATEMENT A: Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT
How can unmanned air vehicles be transitioned seamlessly into carrier operations alongside legacy manned aircraft? One way could be through automatic recognition of the flight deck director's gestures to the aircraft. If a system could automatically recognize and respond to the director's signals (e.g. "come forward", "left", "right", "stop") just like a pilot would, then unmanned aircraft could be integrated onto the flight deck with minimal impact to operations or training. This report defines the problem of computer vision and gesture recognition on the carrier flight deck, and discusses technical issues and describes conditions in engineering terms (how fast, how far, where, how dark, min/max, variance) sufficient to select sensors and approaches with high likelihood of success.

15. SUBJECT TERMS
Gesture Recognition, Computer Vision, Carrier Flight Deck Operations, Aircraft Handling

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER *(Include area code)*

I. Table of Contents

II. Executive Summary

Flight deck operations are a "dance of chaos" with steam, constant motion, and crowding. Moreover, Fleet operations are continually subject to reduced manning and resistance to change. *Then how can UCAVs be transitioned seamlessly into carrier operations?*

If a system could be developed that **automatically recognized** the gestures that flight deck directors use in controlling aircraft, then UCAVs could be integrated with a minimum investment in additional manning or training cost.

We conducted a thorough **analysis** of carrier operations for directing unmanned aircraft. We interviewed flight deck directors, handlers, and carrier design personnel. We compared different directors, conducted ship visits – one while underway during heavy operations, and analyzed hours of video taken of flight director operations. We used in-house tools to take photogrammetric measurements and model illumination.

We defined **major challenges** to automated gesture recognition including:

- Low light
  Visibility during nighttime suffers.
- Pose
  Scaling and rotation of the director varies considerably which complicates the use of template matching.
- Occlusion
  People, equipment and steam can move in front of the director.
- Director variability
  Because of differences in technique and style, the same gesture can look different or different gestures can look the same.
- Scene clutter
  There are many objects, both moving and stationary, in the scene, including other directors.

We investigated the potential of a **machine vision** based recognition system aboard the UCAV to recognize flight deck director gestures. We applied expertise in three recent carrier automation projects for machine vision. We performed a limited test of machine vision on recorded gestures.

Our **recommendations** are that machine vision could be feasible, but only if:

- Visibility of the director could be **augmented** through LEDs in wands and cranials, IR illumination, or body-mounted sensors.
- Sensor(s) could be **placed high** on the aircraft (at least 8' off the deck) to see over objects and avoid occlusion problems.
- Director practices could be more **standardized**. Training the directors to face the aircraft more and eliminate non-standard cues would go a long way in enabling a machine vision system to successfully recognize director gestures for the UCAV.

III. Detail

A. Objective

The objective of this effort was to define the problem and quantify how directors guide aircraft in realistic conditions. We wanted to describe the environment in engineering terms (how fast, how far, where, how dark, min/max, variance) sufficient to select sensors and approaches with high likelihood of success.

We conducted a thorough analysis of carrier operations and the environment. We interviewed and videotaped numerous flight directors from Lakehurst Air Ops, Naval Air Force Atlantic Fleet, and the USS Harry S Truman (CVN-75) while in port. We observed and videotaped operations aboard the USS Nimitz (CVN-68) from positions on the flight deck and high in the island while she conducted fleet exercises with a full airwing. We took photogrammetric measurements of video from yet another ship's operations to determine director positioning. And we used in-house models to determine illuminations at night.

B. Gesture Lexicon

The Aircraft Signals NATOPS manual (see Appendix A) defines a standard set of gestures and is the guiding document for directing aircraft on the flight deck. There are 64 gestures listed in the NATOPS manual. Of these, 16 are applicable to fixed-wing UCAV operations. There were four additional gestures that were not in NATOPS. This brought the total set of gestures to 20.

C. Typical Scenario

Flight deck director activities could be separated into three major categories:

1) Taxiing and Parking: Directing aircraft from one parking spot to another.
2) Recovery: Directing the aircraft as it disengages the arresting cable, moves out of the landing area, and taxiies to an initial spot.
3) Launch: Directing the aircraft into the catapult and through the launch process.

Consider the following scenario for a UCAV-N as it taxiies out of spot, into the catapult for launch, and is recovered after its mission.

1. Taxiing

First, the director makes initial contact with the aircraft. When there is a pilot in the cockpit, this initial contact is not much of an issue; for an unmanned aircraft, this will require some thought. It is important that the UCAV correctly identify the director controlling it. The director will then issue commands (*come forward, move right, move* *left*) for the aircraft to begin taxiing out of the spot. A director will remain in place while the aircraft moves. As the aircraft is moving out of the director's range, he will hand control off to another director located further along the aircraft's path (*pass control* command). Multiple directors, as many as six, will control one aircraft before it reaches its destination.

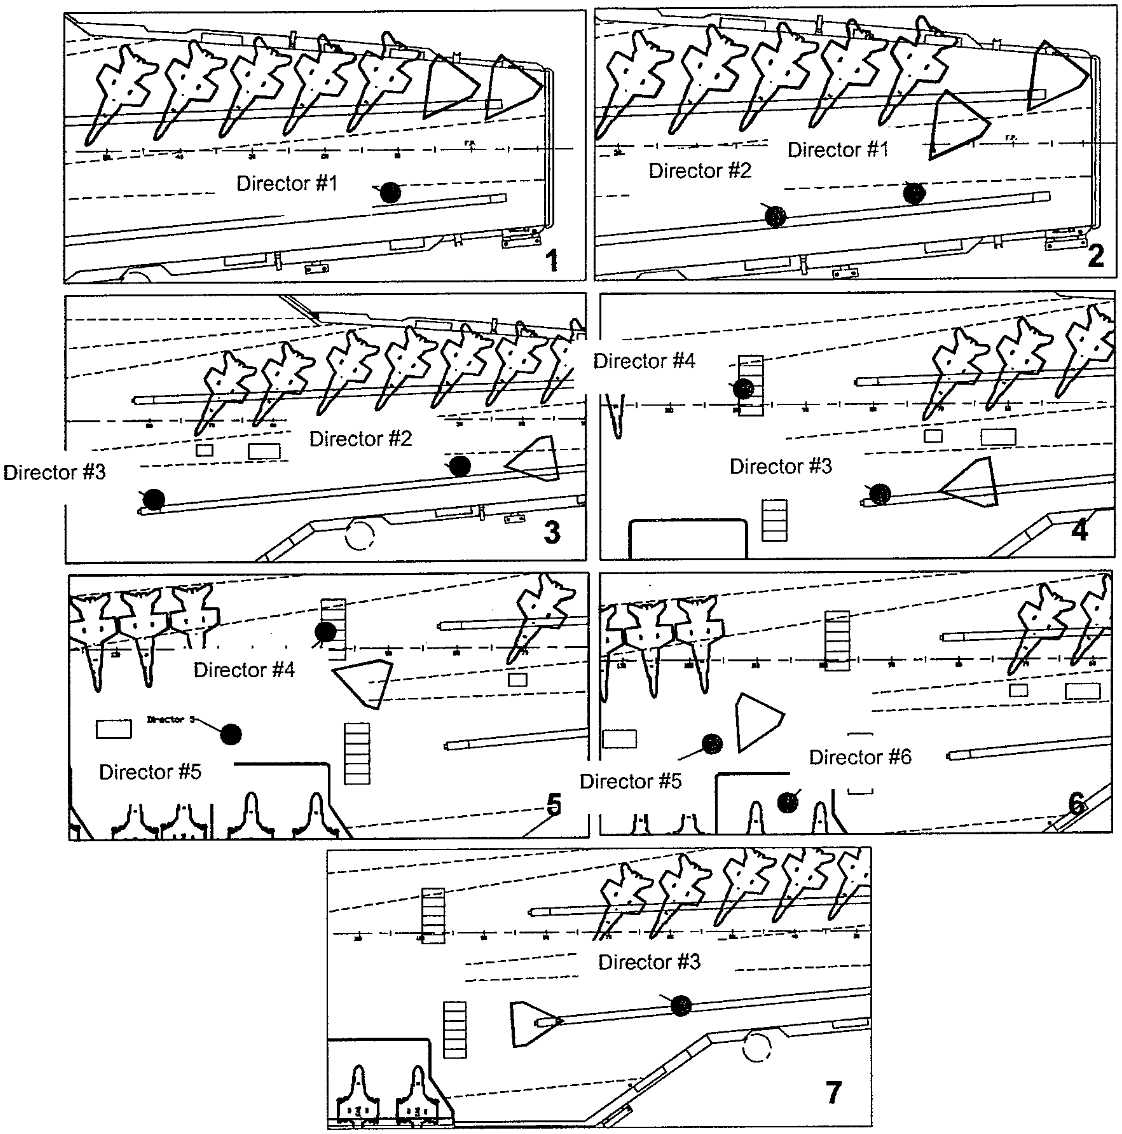

Typical taxiing sequence from parking spot on the bow to the catapult. The UCAV could see up to 6 directors in the path.

Once the aircraft is in the desired spot, the last director will give a *brakes on* command. Sometimes several deck hands will manually push the aircraft backwards until the *brakes on* command is given.

Commands during taxiing and parking are: *move forward, turn right/left, brakes on, brakes off, pass control* to the next director, *back-up, slow down, I have command,* and *slow down engine on indicated side.* Appendix A illustrates each gesture.

2. Launch

Now the UCAV is ready to be positioned into the catapult. This operation is very demanding on a director, in terms of the precision and tolerances required to get the launch bar to mate with the catapult.

The aircraft is positioned behind a jet blast deflector (JBD) with wings unfolded, waiting for its turn for the catapult. The *unfold wings* command was given *a priori.* The JBD is lowered and the aircraft taxiies to the catapult under director control. The aircraft is stopped at the entry area of the guide track (called the "Y"). At this point, the nose wheel may be off center. Therefore, the director must align it with the shuttle spreader by signaling the pilot to turn the nose gear to the right or left. The director then signals to the pilot to lower flaps and assume the take-off configuration. A "green shirt" installs a holdback bar for a specific catapult tension. The director then signals to the pilot to lower the nose gear launch bar. The director will then move to a position on the side of the aircraft and signal to the pilot to slowly taxi forward. The launch bar is guided by a track over the deck hardware. The holdback bar engages the buffer hooks and stops the aircraft when the launch bar is just forward of the shuttle. The director will give the command to the cat officer (not the aircraft) to take tension. The spreader will grab the launch bar. Settings are verified, safety checks are performed, and the aircraft is launched. If for some reason, the launch is suspended, the director (along with everyone else) will wave their hands over their hands as an indication that they all know the launch has been suspended. The director will then give the pilot the *launch bar up* command and the *throttle back* command. The aircraft will then taxi around and try again.

Commands during the launch process are: *move forward, unfold wings, brakes on, brakes off, slow down, flaps down, turn right/left, launch bar down, nose gear right/left, launch bar up, tension* (for the cat officer), and *I have command.* Appendix A illustrates each gesture.

The *flaps down* command may not be applicable to the UCAV if the aircraft can sense the launch sequence and can automatically assume the launch configuration.

3. Recovery

After the UCAV has completed its mission, it's time to come home. The aircraft will safely hit a wire, thanks to the Joint Precision Aircraft Landing System (JPALS). Once the aircraft comes to a stop after arrestment, a director will move into position in the landing area in front of the aircraft. A "green shirt" positions himself or herself close to the tailhook to observe the cable disengaging from the hook. If the arresting cable did not disengage the hook, the director will signal the aircraft to lower the hook again (*down hook* command) so that the aircraft can be quickly pushed back just enough to disengage the arresting cable. If the cable was loose, the "green shirt" sends a positive indication to the director. The director will then give the *brakes off* command and the *up hook* command to raise the tailhook. As the director commands the aircraft to taxi forward, he will issue the *wings fold* command. He will then hand off control to the next director. Again, the UCAV could "see" up to 6 directors on the way to its parking spot.

Relevant signals are: *up hook, move forward, turn right/left, fold wings, pass control* to the next director, *brakes on, brakes off, I have control, down hook,* and *move back.* Appendix A illustrates each gesture.

D. Shipboard Constraints

1. Stealth

Carriers do not want to be detected, especially in combat situations. Signatures of all kinds, including RF, heat/infrared, visual, and electromagnetic, are of concern. In order to reduce the probability of being detected, the carrier will at times operate under Emissions Control (EMCON) conditions. This means using low power, low probability of intercept for wireless communications, and darker decks. As the trend moves towards use of night vision devices on the flight deck, the decks will get darker.

2. Manning

Because of life cycle cost considerations, carriers will be operating with fewer people in the future. Approaches that require additional manning will be discouraged. For example, manually controlling the UCAV with joysticks could add up to 30 new people on the flight deck. Consider also the additional weight for each person. If taxi wands get heavy after 8 hours, imagine carrying a joystick, transmitter and battery.

3. Training

Implementing new equipment and operations requires a significant investment in training. New courses have to be developed and taught, and equipment needs to be deployed to the training command.

How malleable would the Fleet be in standardizing gestures or adopting new ones if the UCAV required them? Directors that we interviewed felt that two or three new gestures would probably be acceptable. Although NATOPS attempts to standardize the set of gestures, it was apparent early on that gestures could differ. One reason is that different ships and coasts train differently. Directors that we interviewed felt that more standardization was acceptable. But the system needs to be robust enough to account for the differences that naturally occur because people are different.

4. Legacy Aircraft

UCAVs will be operating side-by-side with legacy, manned aircraft. JSFs and F/A-18 Super Hornets will be on the flight deck for much of the UCAV's lifetime. Manned and unmanned aircraft will need to co-exist. Having dual, separate processes for each will not work on the carrier.

E. Machine Vision Based Approach

1. Trade-offs

How could the UCAV-N aircraft be handled on the flight deck? We considered approaches that would require no additional manning (see Appendix B). Simply hooking a tow tractor to the aircraft was considered too slow, especially since that aircraft are required to recover every 45 seconds and need to taxi out of the landing area for the next aircraft. Giving the director a joystick would tax his/her workload with an additional process and would negatively impact training and operations. Director-based sensors (*e.g.* data gloves, accelerometers) could be effective if they could be inexpensive, light-weight, and unbreakable. Our initial position was that a machine vision based approach would least impact operations and training and therefore held the most promise. Sensor(s) would be mounted on the aircraft. Image recognition software would reside on the onboard computer. These would provide inputs to the UCAV-N control system (see Appendix C).

2. Gesture Recognition State-of-the-Art

There has been much work in recent years in gesture recognition. These include hand motion, body motion (arm, leg, head), and facial gestures (for an in-depth survey, see [Cohen] and [Fisher]. Issues for any gesture recognition system include:

- Gesture lexicon – How large is the set of gestures?
- Punctuation – Do discrete start and stop points of a gesture need to be identified?
- Variance – How consistent are people in their use of gestures?
- 2D vs. 3D modeling
- Recognition rate
- Environment

Several research efforts are furthering the state-of-the-art in body motion gesture recognition. Bradski and Davis [Bradski 2002] use motion history images to recognize waving and overhead clapping motions for controlling a music synthesis program. Cybernet Corp (Dr. Charles Cohen) developed a system that recognizes arm motions of an Army squad leader in a training simulator environment. In KidsRoom, Bobick and Davis [Bobick 97] developed a system where virtual monsters dance with a child based on the child's dance moves. Starner and Pentland use hidden Markov models to recognize forty American Sign Language gestures [Starner 95].

While this is all important work, virtually all the current gesture recognition systems run inside in a controlled, lab environment. Gestures are constrained and performed at close range with nothing moving in the foreground or background.

3. Machine Vision Efforts on Carriers

In the past few years, we have built machine vision image recognition systems for different applications on carriers. Although still in the developmental stage, these efforts have given us experience with the carrier environment. The Embarked Aircraft Tracking System successfully identified and tracked aircraft on the flight deck. Under this effort, we developed algorithms to enhance poorly illuminated imagery, correct for distortion from jet exhaust, filter out irrelevant motion (*e.g.* from ocean waves, electromagnetic interference noise, and moving people/equipment), and algorithmically dampen vibration. Another effort aimed at recognizing the proper mating of launch bar to catapult spreader. This showed us that for short ranges, steam is primarily an illumination issue, and machine vision can work in the presence of steam if properly illuminated.

F. Challenges

1. Low light

During night operations, the deck is typically lighted using yellow and blue floodlights. These lights produce extremely weak light levels. We found that illumination levels at night were between 0.15 and 1.5 foot-candles (see below). This is analogous to reading a book at night with no light other than the moon and stars. This is a very low light level for many standard cameras.

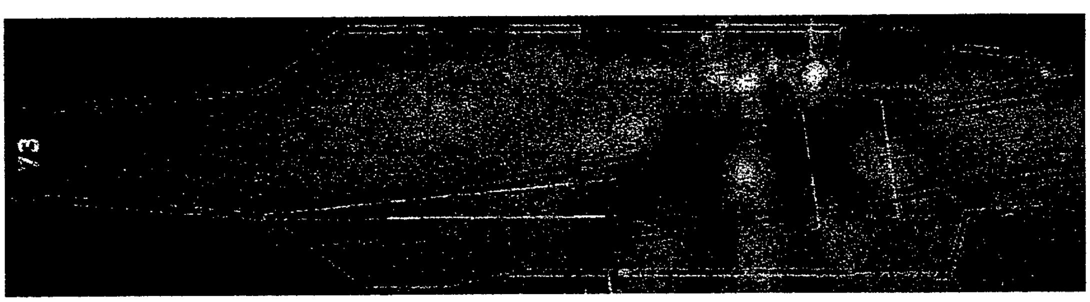

Illumination Contour Map from Overhead Floodlights, CVN-73

Over the past few years, NAVAIR Lakehurst has experimented with a number of off-the-shelf cameras and found that there are low light level CCDs with enough sensitivity to be adequate for the flight deck. One example is a non-intensified camera built by Watekc which has good low light response while minimizing blooming effects. Of course, a lot depends on moon phase and flight deck flood light settings. But we found these cameras can see out to a range of 635 feet on the flight deck.

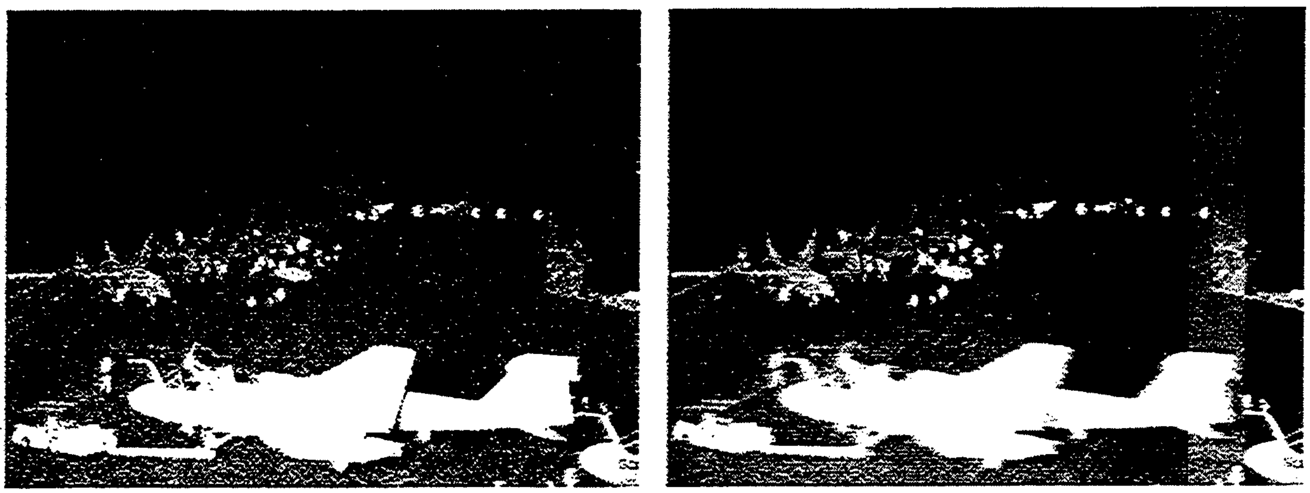

Watec imagery, non-integrated and integrated during a new moon.
The F18 at the top-of-the-bow is 635 feet from the camera.

Other options include infrared cameras, although their performance could degrade in the rain or a sunny hot day. Actively illuminating the flight deck in the near or mid IR can help but may compromise stealth. Through the use of optical filters, the cameras could be tuned to the particular frequency. Laser imagers are also an option, but not recommended due to eye safety and emissions concerns.

Directors operate holding lighted wands (standard flashlights with cones attached) that are colored yellow. In many cases only the wands can be seen. This problem is exacerbated when there are multiple directors in the scene. Who is directing whom? When faces and bodies are not visible, this becomes difficult.

Image of director with wands on CVN-68
using standard camcorder at range of 200 ft

2. Blooming

Blooming from bright light flashes such as engine exhaust or spot light sources is a constant problem for cameras. Filters are required to minimize the adverse effects of blooming on the image.

The effects without and with a hot mirror filter to suppress glare from the near-IR illuminators adjacent to the bow deck edge cameras

3. Sun Glare

Sun glare on the flight deck can be a big problem, especially if the sun is low in the sky. Whole sections of the scene can be washed out. Very little algorithmic enhancement can be done when pixel values are saturated.

It is difficult to see a director when the sun is directly behind them. Cameras have difficulty as well. The camera will degain and the director's image becomes black. Critical features, such as arms and legs, could be lost.

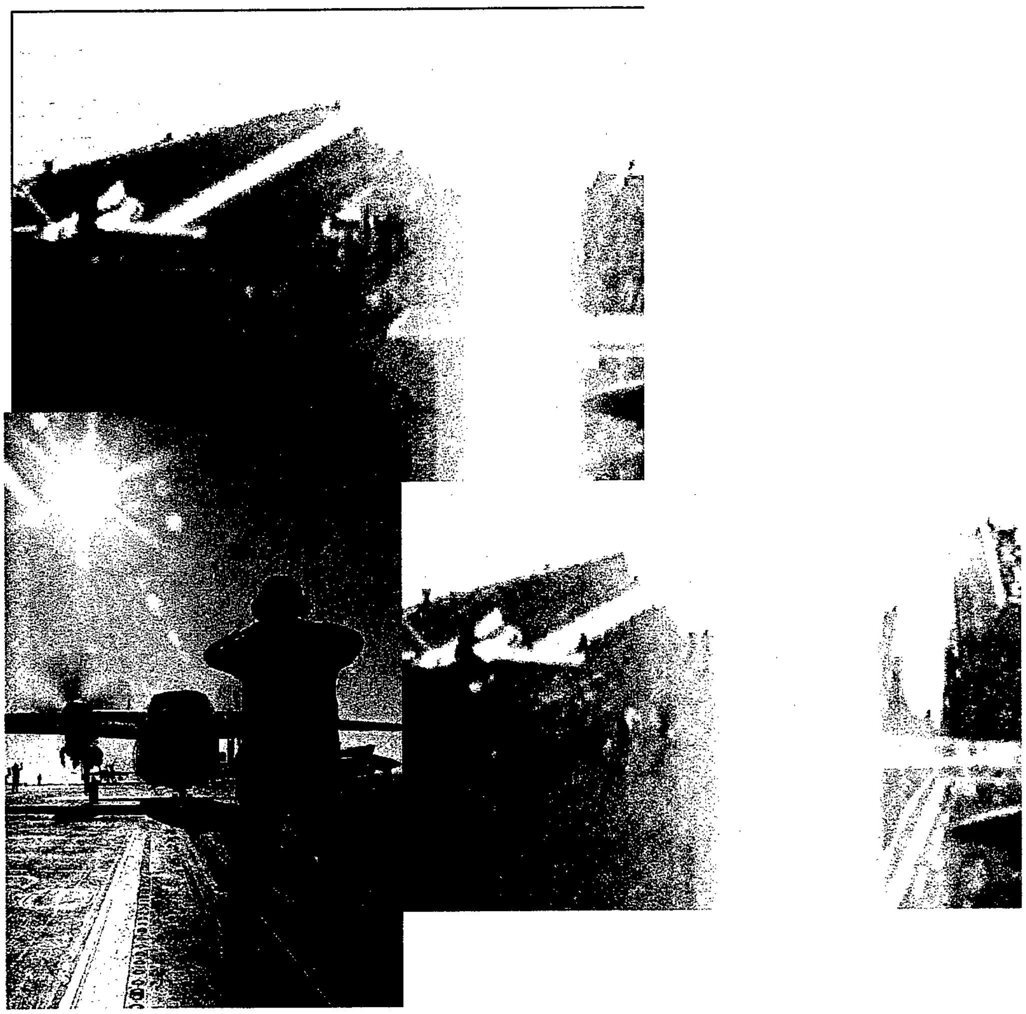

Examples of the effect of sun glare

4. Steam

After each launch, steam will billow from the catapult. Steam could occlude the director who is controlling the aircraft next in line for the catapult, although this is a problem only for short periods of time (under a sec). The wind over deck keeps the steam moving. Generally, the upper part of the director is more visible, and the pilot's position high off the deck helps him/her to better see the director. Steam is troublesome not only because it reduces the detail, but also because it creates huge motion fields in the imagery which any algorithms will have to suppress.

Effects of Steam

5. Occlusion from People and Objects

We observed various times when people or deck equipment would move in between a director and his/her aircraft, or as an aircraft would turn, people or equipment would temporarily be in the way. This is not a problem for the pilot who is positioned up height (~13 feet off the deck) and can see the director over any obstructions. However, for a camera mounted on the nose gear (~4-5 feet off the deck), this would be a problem. The obstruction could last for 3-5 seconds. Given a frame rate of 30 frames/sec, the image recognition system would have a blind spot equal to 90 to 150 frames. The system could miss up to 4 commands.

Therefore, the sensor would need to be mounted higher on the UCAV, at least 8 feet off the deck. Candidate locations could be the wing tips, canopy, or on the refueling probe which would have to be deployed while on the deck.

6. Background Scene Clutter

The flight deck is a very dynamic environment with many stationary and moving objects in the scene. Consider the following image.

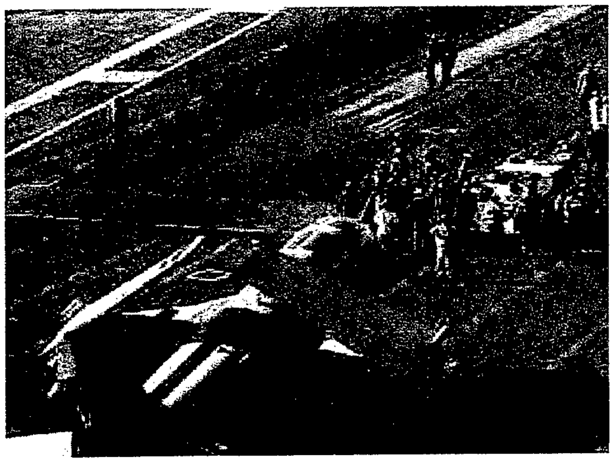

Where's the director?

Where's the director in all this? Humans are well suited to rapidly finding the target among a clutter of other objects. Computers have a harder time. In addition, there could easily be other yellow objects besides the director. There could be other directors in the scene, and deck equipment that is painted yellow. So to find the director, the system cannot simply cue on the yellow color.

Multiple directors controlling multiple aircraft

We counted objects in various scenes and found the maximum clutter to target ratio to be between 50:1 and 70:1. Objects could be people, aircraft or deck equipment, either moving or stationary.

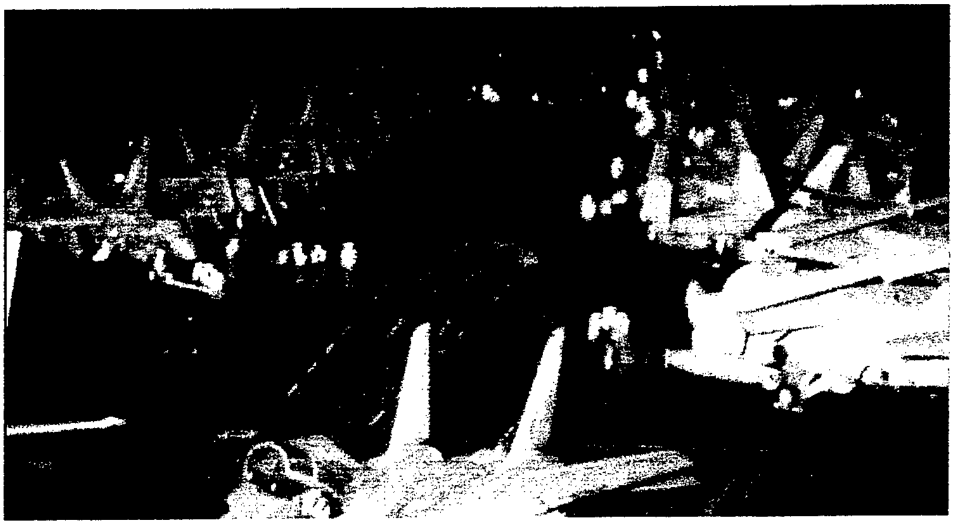

Scene clutter from aircraft

7. Color Variability

Again, if the system is cueing on color, it is important to note that the color can differ. Jerseys can get dirty and faded. Bright sun or a cloudy day could render different shades of yellow. Reflective tape on the float coat could be in different areas.

In addition, we observed different colored gloves (brown, black, green or yellow) or no gloves at all.

8. Jet Exhaust

Note the flight director in the exhaust plume of the forward F/A-18. Considerable image distortion may be present in cases like these because a UCAV in the aft F/A-18's position may have a distorted image from heat.

Light cat 1 aircraft exhaust as "seen" from the 010 level approximately 300 feet away

Image under heavy exhaust – Note that contrast is degraded

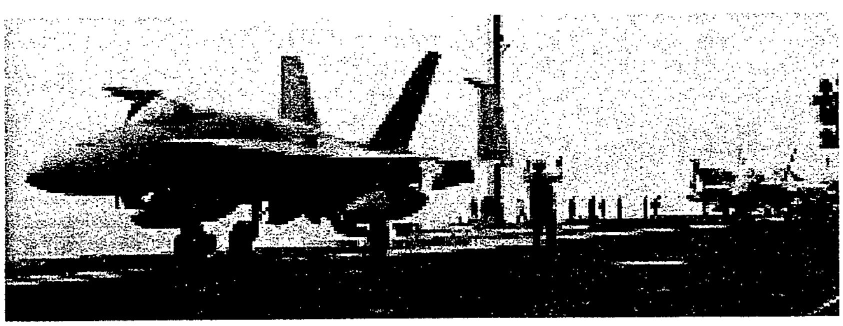

Note the flight director in the exhaust plume of the forward F/A-18. Considerable image distortion may be present in cases like these because a UCAV in the aft F/A-18's position may have a distorted image from heat.

9. Range

In order to provide requirements for the sensor, we needed to determine the range of director positions and orientations with respect to the aircraft. To accomplish this, we used an in-house system that we developed for tracking aircraft on the flight deck to track both aircraft and director positions and orientations on existing video. (Appendix E contains all the findings from this analysis.) From these photogrammetric measurements, we determined the worst-case distance between the aircraft and director to be 243 feet (or ~250 feet). This was confirmed through observing operations aboard CVN-68 and interviews with directors from the CVN-75 in Norfolk.

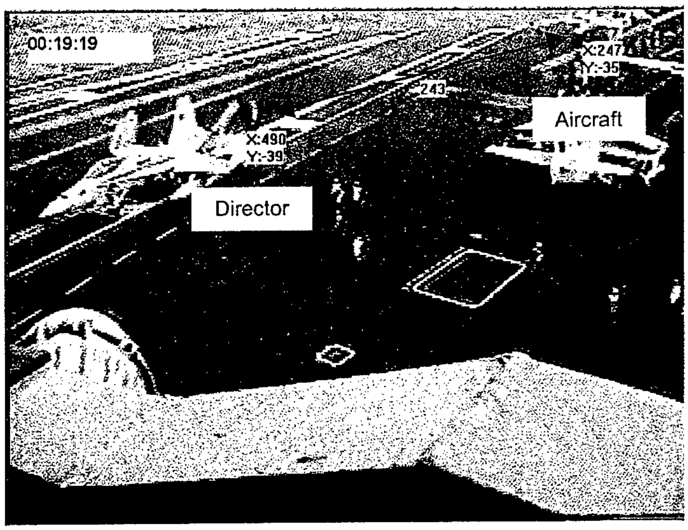

Director's position relative to the aircraft in this example is 243 feet

During night operations or poor visibility conditions, the director will stand closer to the aircraft. For example, the maximum distance for directors operating on the stern (the area known as "FLY-3") was 75 feet during the day, but only 35 feet during the night. This is a mitigating factor for a machine vision system. If the pilot has a hard time seeing the director, the director will generally move closer.

From the photogrammetric measurements and observations, we determined the minimum distance to be between 10 and 15 feet.

A distance-to-target of 10'-250' does present a significant scaling issue for machine vision. However, the pilot has a hard time seeing the director at 250' as well. At 250', the director will typically use only signals: *come forward* and *stop.* There won't be any intricate, detailed gestures, and there won't be any turning. This mitigates the scaling problem. We can envision two separate algorithms: one for shorter range, and one for longer range.

10. Orientation With Respect To Aircraft Nose

As long as he is confident that the pilot can see him, the director will stand anywhere. This could be up to +/- 100° with respect to the aircraft nose, a function of the canopy visibility and the limits of the pilot's head swivel. On the starboard side of two-seaters, such as the EA-6B, the director will stay closer to the nose. Remember that for safety purposes, the director will stay in one place as the aircraft moves. As the aircraft taxies past the director (past the ~100° limit), the director will pass control to another director further down the path. It is unlikely that a fixed sensor can be found to satisfy this field-of-view requirement. Current state-of-the-art panoramic camera systems package multiple cameras staring into a mirror. They can achieve up to 360° coverage, but are not nearly small enough to mount on an aircraft. This +/- 100° field-of-view requirement would drive a pan/tilt/zoom solution. Many miniature pan/tilt/zoom cameras exist off-the-shelf that would be small enough to mount on an aircraft.

11. Pose or Rotation With Respect to Aircraft

This refers to the rotation of the director's body on a vertical axis relative to the aircraft. This generally will not vary by more than +/- 20° as the director tries to remain facing the aircraft. However, there are many times when the director needs to turn, e.g. to see where the next director is, or to make sure another aircraft is not bearing down on him or about to fry him with hot exhaust. We have observed directors facing away from an aircraft, but continuing to wave his arms to direct the aircraft *behind* him!

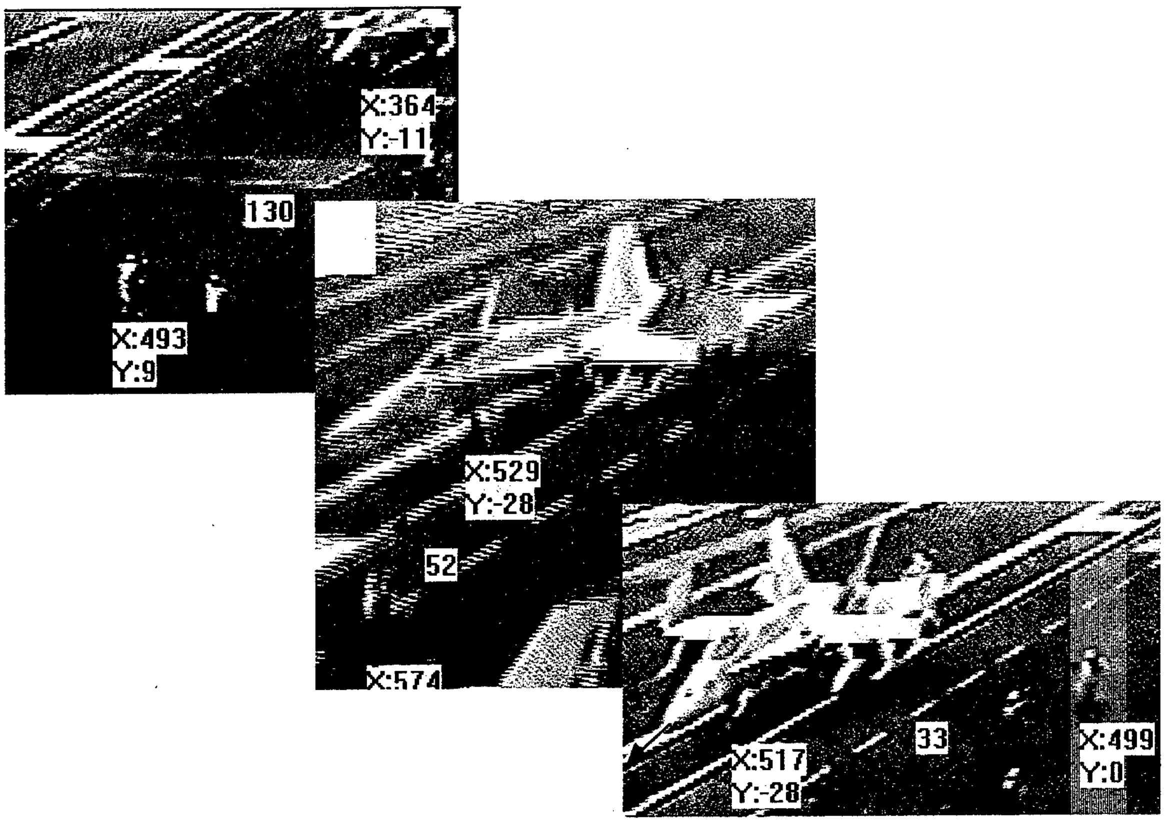

As the aircraft has overtaken this director, he continues to signal while facing away (middle figure) and sideways (rightmost figure) to the aircraft!

12. Gesture Similarity

A machine vision system will have to classify gestures that are similar to each other, either by definition or because of director practice. One example is *Emergency Stop* and *Brakes On*. These can be confused. Some directors perform *Emergency Stop* with their hands together, some with their hands apart, which could be confused with *Brakes On*.

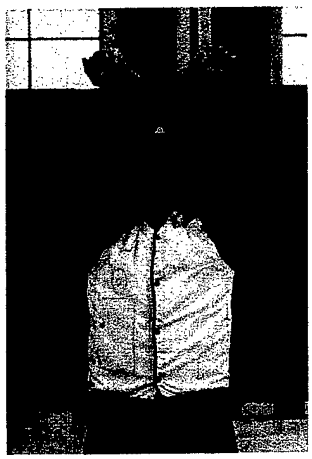

Emergency Stop

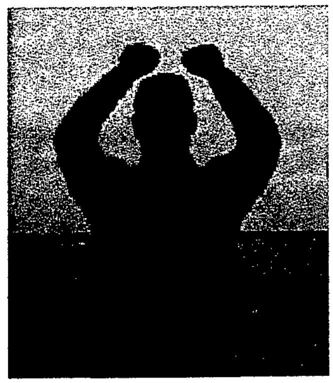

Brakes On

Another example is the *Slow Down* signal. The NATOPS manual describes *Slow Down* as an up-and-down motion with both hands. But we observed most directors to use more of a circular motion. A portion of this circular motion could certainly be misconstrued as the *Come Forward* motion. Machine vision algorithms would have to be carefully developed so as not to act on the subset of the gesture.

Some similar gestures are distinguished only by their sequence. *Brakes On* is an open hand to a fist. *Brakes Off* is a fist to an open hand (although the director will rapidly flash the *Brakes Off* signal a few times). At a distance, these gestures could be confused. Another example is *Wingfold* and *Wingspread.* The *Wingfold* signal is both arms straight out at sides, then swept forward to a hugging position. *Wingspread* is the same but in reverse. It is important to note that a machine vision algorithm would need to capture the sequence in order to properly classify the gesture.

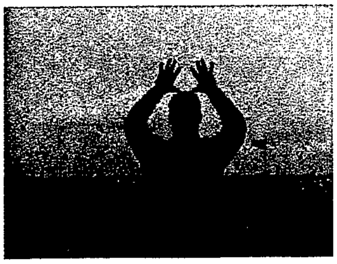

Brakes On
Open hand to fist

Brakes Off
Flashing fist to open hand

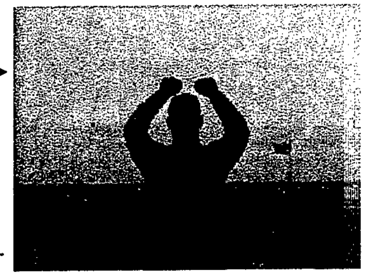

Wingfold
Arms straight out at sides then swept forward and hugged around shoulders

Wingspread
Arms hugged around shoulders, then swept straight out to the sides

Launch Bar Up and *Launch Bar Down* are also distinguished by their sequence, and have a 3-D component as well. If viewing the gestures head-on, it may be difficult to distinguish these gestures, even if given the proper sequence. A mitigating factor with these gestures is that they depend on a state of the aircraft, i.e. if the aircraft's launch bar is currently up, then a *Launch Bar Up* command does not make sense, and the command must be *Launch Bar Down.*

Small gestures are also confusing. It is hard to differentiate between *Hook Down* and *Hook Up*, especially at a distance. *Hook Up* could also be confused with *Emergency Stop* (see above).

Hook Down

Hook Up

Context is an important characteristic for distinguishing similar gestures. The *Turn Right* gesture looks a lot like *Final Tension.* The difference is that *Turn Right* is given during the taxiing process and *Final Tension* is given toward the end of the launch process. If the UCAV could sense when the launch bar is in the guide track and turn off the nose wheel steering, then this problem goes away.

Turn Right

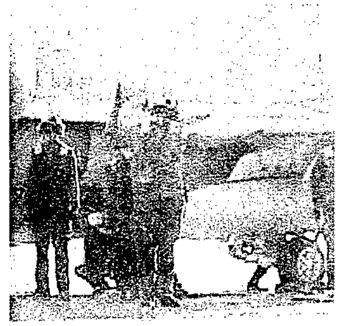

Final Tension

13. Subtle Cues

Directors will sometimes use subtle, non-standard cues when trying to convey information to the pilot. These would not be found in the NATOPS manual. They include:

- Head nod – commanding the pilot to nudge nose wheel steering on cat
- Head nod could also convey confidence – "Yeah, you've cleared the combing."
- Kick – "Director, get over there", or "Pilot, cut the turn faster"
- Point to eyes – "You're not looking at me" or "You're not responding"

14. 3D Gestures

One concern was that there would be three-dimensional components of gestures that would render the gesture unrecognizable by a two-dimensional camera system. This does not seem to be an issue with any of the gestures, other than differentiating between the *Launch Bar Up* and *Launch Bar Down* signals (discussed above).

15. Director Variances

Different directors can perform the same gesture differently. This could be attributed to:

- Director technique
- Fatigue
- Training differences

Algorithms will need to be robust enough to handle variations in a single gesture, correctly classifying the gesture even if its use is inconsistent.

Gestures that are especially prone to director variances are:

- *Come Forward*
- *Pass Control*
- *Move Right/Left*
- *Receive Control*
- *Slow Down Engine*

To give the *Come Forward* signal, some directors pivot at the elbows and swing their entire forearms. Others hold their forearms stationary, pivot at the wrist and wave their hands. Some directors wave over their heads, while others wave in front of their face. This signal is particularly vulnerable to director fatigue. As the director gets tired, the position of the hands and forearms starts sagging lower and lower.

Variances in the Come Forward *Signal*

For the *Move Right/Left* and the *Pass Control* commands, the height of the pointing arm can vary from a position pointing horizontally to a position pointing toward the sky. The director could point with one finger, two fingers or a flat hand.

The *Pass Control* command has been observed with either one hand or both hands pointing. Sometimes the director will tap his/her head before pointing, sometimes not.

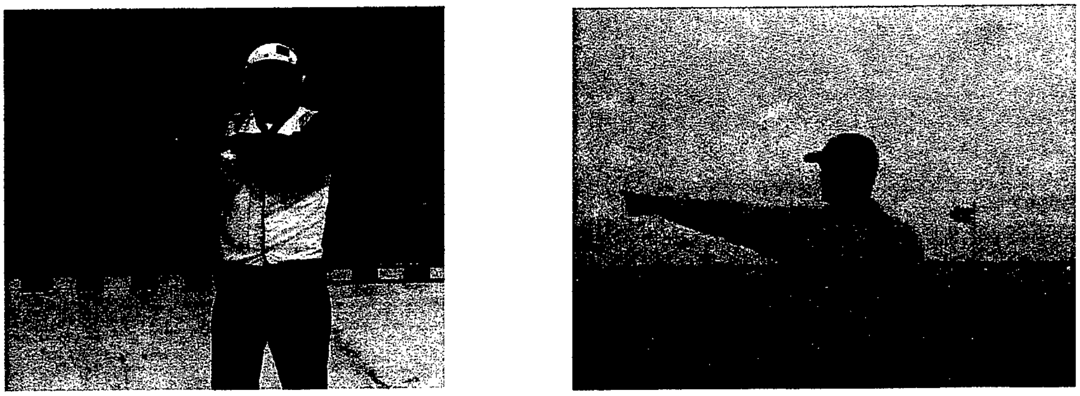

Variances in Pass Control *signal*

Another example of a difference based on training is the *slow down engine* gesture. We observed directors on the USS Truman indicating the gesture by grabbing the thumb and making a "throttle back" motion like a pilot would use with the stick. This is actually quite different than what is portrayed in NATOPS – a pushing down motion with a level hand. One director we interviewed, who was trained on the USS Kitty Hawk, used this gesture. This could be attributed to different coasts using different techniques.

The *Receive Control* gesture varied as well. This is the indication to the pilot that this new director is assuming control of the aircraft, and is an important signal to the UCAV. A major issue with automatic gesture recognition is ensuring the right director is controlling the right aircraft. Some touched their chest or shoulders with both hands, then held up a hand. Other just repeatedly touched their chest or shoulders. And others simply held up a hand. This seemed to be a product of differences in training.

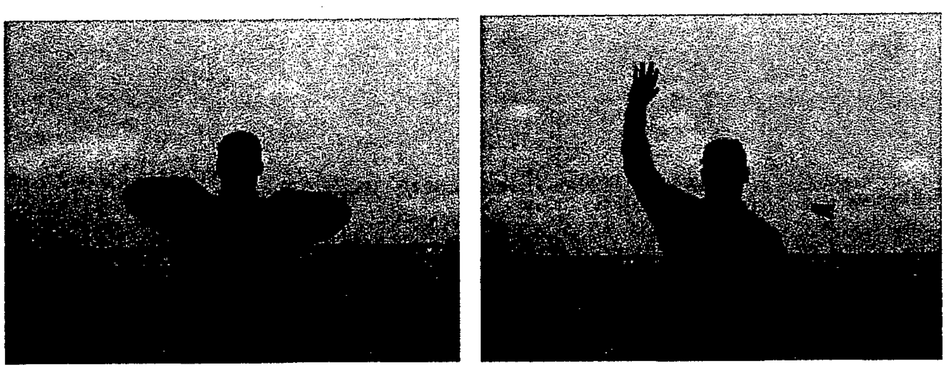

Variances in the Receive Control *or* I Have Command *signal*

Some gestures can be performed with either the right hand/arm or left hand/arm without changing the meaning (*e.g. Receive control, launch bar up/down, hook up/down*). A machine vision will have to account for both variations of these gestures, effectively expanding the gesture set.

16. Camera Retargeting

Consider the typical scenario: the aircraft is moving but the director stays in one spot. The system must retarget the aircraft camera(s) in real time. As the aircraft moves, the director shifts in the image. The algorithms need to pan the camera to keep the director's image centered.

An even larger issue is what happens when the director hands off to the next director. How do you find the next director? The human pilot sees that the first director is pointing left, for example. The pilot will then scan left until he sees a director with one hand up (*receive control* gesture). He assumes that this is his director and he will begin responding. A machine vision could mimic this process. Cameras would have to slew in the direction the yellow shirt is pointing and find the new director. This would be a major challenge given the carrier environment.

17. Temporal Resolution

A time budget for processing the algorithms is a direct function of how much latency the UCAV would be allowed during operations. The following questions are relevant:

- How fast do current manned aircraft respond to director commands?
- When would lag time start affecting tempo of operations?
- When would lag time become disconcerting to the directors?

During our visit to the CVN-68, we timed pilot reaction to director commands. We used a stopwatch with resolution to $1/100^{th}$ of a second. The stopwatch started when a director gave a new signal and stopped when the aircraft responded differently to the new signal. We found typical response time to be 7/10 sec. We understand that hand/eye response time for pushing the stopwatch button must be factored in. But this gave us an idea of the current response time for manned aircraft – nearly instantaneous. Given more sophisticated equipment, budget and time, we could come up with a more exact requirement.

Instantaneous may not be a realistic requirement for complex machine vision algorithms. If the requirement were tied to tempo of operations, then sortie generation rate could be the governing criterion. Sortie rate is defined by several factors, including aircraft availability, time to fuel the aircraft, time to loading ordnance, time to fix the aircraft, etc. All these times add up to all overall turnaround time of around 1 hour and 30 mins. It is hard to see that a gesture recognition time of perhaps 1 or 2 sec would have an appreciable effect on sortie rate, especially if UCAVs were already taxiing at a slower speed (~3 mph) than the other aircraft (~5 mph).

The real requirement may be tied to safety and what's acceptable to the directors. If the director is controlling a tight maneuver and the aircraft doesn't respond quickly enough to a change in signal, then the aircraft could end up over the side or into another aircraft.

More human factors analysis would be needed to determine just how much lag time is acceptable. When the aircraft is taxiing forward toward the deck edge, perhaps the directors would become conditioned to give the *move left* signal a split second earlier to account for latency, but we can't depend on that. Experience tells us that 1.5 sec may be the upper bound.

G. Feedback to the Directors

We interviewed a number of directors to determine the instances when they would want feedback from the UCAV. Their responses were:

- When handoff occurs
  Director wants confirmation that the aircraft correctly identifies him/her as the controlling director. This could be with some sort of flashing green signal from a landing gear light.

- When a signal changes
  Directors were divided as to whether a confirmation light was necessary for each signal change. Perhaps the change in UCAV response is enough. But some thought a flashing light would inspire more confidence in the director.

IV. Summary of Requirements

- Identify controlling director among other directors on deck
- Classify gestures in the lexicon (20 separate gestures), see Appendix A
- Spatial resolution: 6 ft tall man at 250 feet
- Color resolution: shades of yellow
- Temporal resolution: 1 to 1 ½ sec from command to UCAV response
  This is the time budget for all processing: retargeting, filtering out scene clutter, correcting for scaling and rotation, and gesture classification.
  60 Hz frame rate for distinguishing between 20 separate gestures and pace changes.
- Range of illumination: 10,000 foot-candles (direct sunlight) to 0.1 foot-candles (night)
- Scene clutter: object to target ratio of 50-70 to 1
- Scaling range: 10 to 250 feet
- Field of view: +/- 100°
- Rotation: Typically +/- 20°, can be up to 90° for short durations
- Gesture variances and inconsistency need to be taken into account
- No requirement for director to "punctuate," *i.e.* define gesture start and stop points
- Classification accuracy: near 100%

V. Recommendations

A. Cooperative Targeting

The scene clutter, lack of illumination at night, steam occlusion and weather issues make it challenging to identify the director and track his/her arm movements. We could make things a little easier by enhancing the visual cues through cooperative targeting. The director could be marked by a frequency modulated LED, on the gloves or jersey (to determine gesture trajectory) and the front of the cranial or float coat (to determine pose). The LED would be pulsing in the non-visible spectrum range (i.e. IR). The IR sensor would be tuned to the LED's particular frequency so that only the pertinent target would be tracked. IR is preferred so that pilots and flight deck personnel are not confused by signals in the visual range.

Another method would be to glue retro-reflective tape to the uniform arms and gloves that only reflects certain IR frequencies. Such tape is commercially available. An IR illuminator on the UCAV could illuminate the target (*i.e.* the director), providing clear gesture characteristics to the IR tuned optics on the UCAV sensor.

At night the directors operate with wands. These wands could be configured to function in tandem with the UCAV sensor. Yellow LEDs (for the humans) could be interlaced with IR LEDs that would be tuned to the sensor. The wands could be used during the day as an cheaper alternative to sewing sensors into gloves or clothing.

It is important to identify the right director in a group of other directors. Placing RF beacons on directors could be the answer. Each director would have a unique frequency, and each UCAV would need to "dial in" to the director's frequency (or vice versa). This is technically feasible, although systems engineering and human factors work would be required to determine how it would be done and if it could be done fast enough to keep up with the tempo of operations. Emissions control and stealth are issues, but distances are short enough for low power transmission, and spread spectrum can be used.

B. Associative Memory

There are several popular approaches to the recognition of human motion in real-time, including analysis of temporal trajectories and hidden Markov models. One promising approach proposed in [Bradski 2002] is associative memory using motion history gradients. This is a type of template matching where successive layering of image silhouettes of a person into a single template is used to represent and recognize patterns of human motion. Every time a new video frame arrives, the existing silhouettes are decreased in value subject to some threshold and the new silhouette (if any) is overlaid at maximal brightness. This layered motion image is termed a motion history image (MHI). MHI representations have the advantage that a range of times from frame to frame to several seconds may be encoded in a single image. In this way, MHIs span the time scales of human gestures.

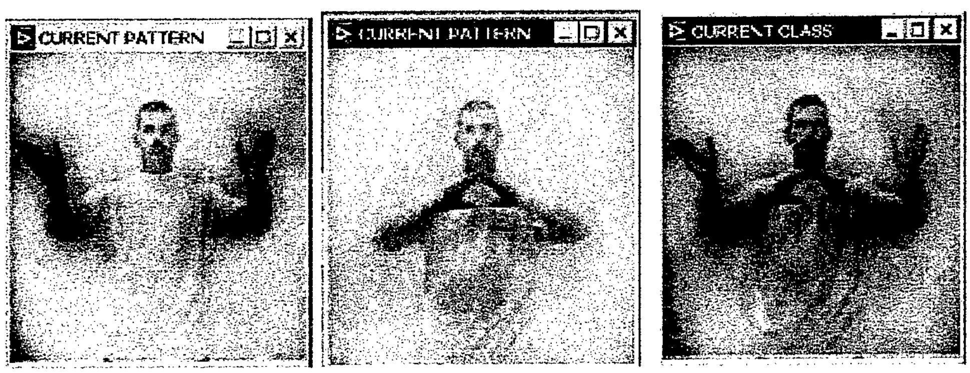

Images of Come Forward *gesture and their Motion History Image*

C. Rotation and Variance – Generalization

Image recognition is a pattern classification problem. We store a set of images as associative memory or templates. To classify the target image, we cull through the set of stored images until a match is found. If we are looking for exact images, we will never find the matching template.

Image rotation and variances (*e.g.* director-induced gesture variances) are an issue for template matching techniques. To classify the target image, we cull through the set of stored images until a match is found. If we are looking for exact images, we will never find the matching template.

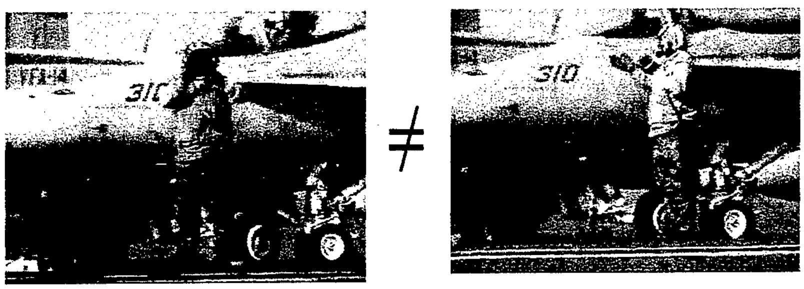

MEMORY IMAGERY

To take care of rotations and variances, we would use an algorithm with generalization properties, such as the Fuzzy ART algorithm (see Appendix D). A minimal set of templates representing the range of possible images is stored, and Fuzzy ART can classify the target image – this gesture *is closest to* X, so it must be X. If a higher degree of separability is required, then the set of templates could be expanded. The advantage of using Fuzzy ART with associative memory is its computational speed and the ability to generalize with inherited ability to reason with incomplete or imprecise information.

D. Scaling

The following approaches could be useful in attacking the problem of scaling:

- Measuring distance to target
  - Laser rangefinder
  - Stereoscopic vision
  - Using known landmarks
  - Approximating distance from pixel array
- Algorithmic method
  - Fourier transform
  - Log-log transform
  - Fuzzy ART

1. Laser rangefinder

A laser rangefinder could be mounted along with the electro-optic (EO) or infrared (IR) sensor to get real time range information. This approach is universally used for tracking objects and would easily get range. The problems are: (1) effects on flight deck personnel, and (2) clutter. Fleet operators we talked to were adamantly against multiple laser rangefinders shining on the deck. The objection is not so much with eye-safety, causing permanent blindness to personnel. But if a laser happened to shine dead-on into someone's eye, that person could be temporarily blinded – try staring at a light bulb, then reading a book. Dust, steam and exhaust would defract the laser beams, and you would see red lines all over the deck. Also, because of clutter, you are not sure if the laser bounced off the target (*i.e.* director) or another object. This approach is acceptable for tracking aircraft in the sky, but would be more difficult in a densely packed environment.

2. Stereoscopic vision

Stereoscopic vision would make the system much more robust than a single camera. Pixels would have not only intensity but range values associated – a flight director's arms and hands would stand out in more detail than just from a single intensity image. Using two cameras mimics the depth perception present by a human's offset eyes. Try holding a finger up in front of you and slowly blinking your eyes alternately. In each view you will notice your finger shifting by a certain amount. Moving your finger away from your face reduces the shift, while as it nears your face the shift becomes large. This shift is called *disparity* and is related to the distance between the viewer and the feature being viewed. So, for each pixel in the image, you can theoretically compute a range. This makes vision much easier in lots of cases because the flight director will "jump out" of the scene, and you can filter out the scene clutter to only that within a certain depth of field. The rule of thumb is that you want the distance between the cameras to be about ½ the distance to the object to provide enough depth resolution to do useful things. A maximum range of 250' means the separation between the cameras would be 125' – not likely on the UCAV. But stereoscopy would be very good for gesture recognition at shorter ranges (ranges of 75'-100' if cameras were mounted on the wing tips or leading edges) and obstacle avoidance, where the range to target would be short.

3. Using known landmarks

If finding places on the UCAV for two cameras is a problem, then another approach could be exploiting the known landmarks on the flight deck. These landmarks could be deck markings, equipment (such as jet blast deflectors or the island) or deck edges. Comparing the relationship of the target to known landmarks could yield a fairly accurate range. The issue becomes: how many landmarks would likely be in the scene. If the camera is mounted on the nose gear looking horizontally, it won't see be seeing much deck and landmarks would be limited. If the camera is mounted higher up on the aircraft and looking down, there may be more deck and more landmarks.

4. Approximating distance from pixel array

You can calculate range given the aspect ratio of the target (*e.g.* the typical man is 6 ft x 1 ½ ft) and the aspect ratio of the target image (*i.e.* the number of pixels in the array). This approach tends to be less accurate, though, especially for steeper staring angles.

5. Algorithmic approaches

Using purely algorithmic approaches have the advantage of requiring no additional equipment for measuring range. Fourier and Log-log transforms can be computationally expensive. The same Fuzzy ART algorithm used to generalize rotation and gesture variances could be used for scaling as well. This algorithm is much faster than Fourier or Log-log transforms. Note, though, that range would still be important for any obstacle avoidance or path planning.

E. Multiple Sensors

Body mounted sensors, such as accelerometers or data gloves, would be useful in lessening the technical issues associated with machine vision. These sensors would track the hand and arm motion trajectories in real-time and send the data back to the UCAV's computer via a wireless RF link. This would eliminate issues of illumination, scaling, rotation and shirt. This is a major advantage. The problem of gesture variance still remains as well as the introduction of emissions sources. We could solve variance using a generalization algorithm such as Fuzzy ART discussed above.

Issues of cost, ruggedness and equipment life would still need to be addressed before these sensors could be implemented. Sailors have been known to be rough on gear, especially delicate instruments. Weight should also be a consideration, especially since personnel will be wearing this gear for 12-14 hours on the flight deck.

A promising product is the "ShapeWrap" system, built by Measurand, Inc. (www.measurand.com). Fiber optic bundles are attached to arms, legs, wrists and fingers. Gestures are tracked by measuring the bending and twisting of joints. Processing power is minimized since only 2 degrees of freedom are tracked, rather than 6 for conventional accelerometers.

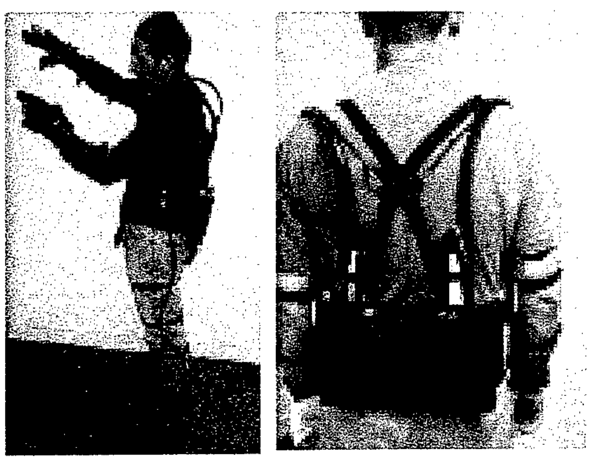

ShapeWrap system by Measurand, Inc.

VI. Summary

We have sought to characterize the carrier flight deck environment and the technical issues that would be associated with gesture recognition. The chaos on the deck with problems of illumination, occlusion, scaling, rotation and gesture variance are a challenge. We believe machine vision is feasible for this application, only if:

- Cooperative targeting is employed – the director's image is augmented with LED or active illumination
- The UCAV-N is configured with the sensor high off the ground (>8 ft)
- Multi-sensor modality can be used for scaling and/or obtaining range data
- Generalization technique is used with associative memory for rotation, gesture variance, and possibly scaling
- Separate classes are established where gesture sequence is important

Machine vision has the advantage of being passive, using less costly equipment, and lending itself to an obstacle avoidance/path planning solution for enhancing the system's safety. Body-mounted sensors hold promise for alleviating some of the issues and could be implemented as part of a total package or stand-alone system.

VII. Authors

Larry Venetsky, Naval Air Systems Command, larry.venetsky@navy.mil
Mark Husni, Naval Air Systems Command, mark.husni@navy.mil
Mark Yager, DeveloSoft Corporation, myager@develosoft.com

VIII. Acknowledgements

The authors wish to thank Warren Baker and Shaun Donnelly from NAVAIR for their advice during this project. We also appreciated the support of the Naval Air Force, Atlantic Fleet (CNAL) and Naval Air Force, Pacific Fleet (CNAP), specifically CDR Mike Yoast, CDR Chuck Shirley and ABHCM Stackler from CNAL, CDR Nel Tabinga from CNAP, and LT George Sharp and LT Szitta from the USS Nimitz. We also thank all the directors stationed at Lakehurst, including ABHCM Mike Bensinger and ABHCS Mike Fenton.

Appendix A – Gesture Lexicon

These 20 signals represent the set of gestures that would apply to controlling the UCAV. This is a subset of the signals found in the Aircraft Signals NATOPS Manual, NAVAIR 00-80T-113, with the addition of four signals not found in NATOPS: *Launch Bar Up, Launch Bar Down, Pivot to Left,* and *Pivot to Right.*

Taxiing Gestures	
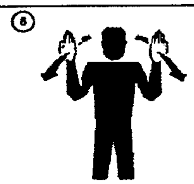 **MOVE AHEAD**	Arms extended from body and held horizontal to shoulders with hands upraised and above eye level, palms facing backwards. Execute beckoning arm motion angled backward. Rapidity indicated speed desired of aircraft.
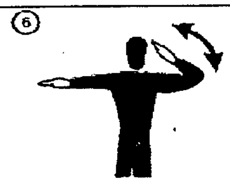 **TURN TO LEFT**	Extend right arm horizontally, left arm repeatedly moved upward. Speed of movement Indicates the rate of turn. At night is the same as day signal with addition of wands. A clenched fist (day) or down-turned wand (night) directs the pilot to lock the indicated brake.
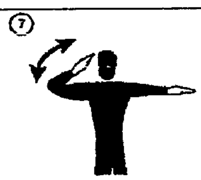 **TURN TO RIGHT**	Extend left arm horizontally, right arm repeatedly moved upward. Speed of movement indicates the rate of turn. At night is the same as day signal with addition of wands
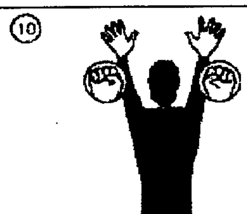 **BRAKES**	BRAKES ON -Arms above head, open palms and fingers raised with palms toward aircraft, then fist closed. BRAKES OFF - Reverse of above. Will repeat rapidly a few times. KEEP BRAKES ON – Hold up one fist and rotate fist back and forth. This is used, for example, if the aircraft needs to stay in place until something in front of him is cleared. At night: BRAKES ON - Arms above head, then wands crossed. BRAKES OFF - Crossed wands, then uncrossed.
 **PIVOT TO LEFT**	The director is telling the pilot to hold one main landing gear brake while turning. Point right closed fist to the port side brake. Repeated move left arm upward like *Turn to Left* signal.

 PIVOT TO RIGHT	The director is telling the pilot to hold one main landing gear brake while turning. Point left closed fist to the starboard side brake. Repeated move left arm upward like *Turn to Right* signal.
(63) 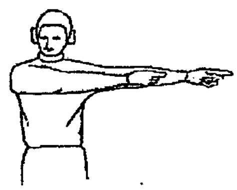 PASS CONTROL	With both arms shoulder height in direction of person receiving control. Will touch eyes or hand with both hands first before pointing. At night is the same as day signal except point amber wand. Will touch helmet with both wands before pointing.
(50)  I HAVE COMMAND	Hold one hand open motionless and high above head, with palm forward. At night is the same as day except with wand. NOTE: Another signal for "I have command" seen during operations is tapping the chest with both hands. Sometimes the director will tap the chest, then raise one hand.
(5) 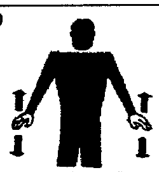 SLOW DOWN	Arms down with palms towards ground. Then moved up and down several times. At night is the same as day signal with addition of wands. NOTE: This gesture may not be relevant to UCAV operations, since the UCAV is already expected to taxi slowly, not much faster than 3 mph.
(9)  STOP	Arms crossed above the head, palms facing forward. Same as day signal with addition of wands. Emergency Stop (as opposed to "brakes on")
(11) 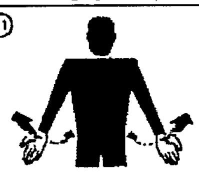 MOVE BACK (ALSO USED TO PULL BACK AIRCRAFT UTILIZING ARRESTING WIRE)	Arms by sides, palms facing forward, swept forward and upward repeatedly to shoulder height. At night is the same as day signal with addition of wands.
(23) 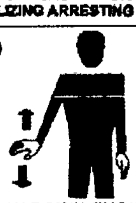 SLOW DOWN ENGINE(S) ON INDICATED SIDE	Arms down with palms toward ground, then either right or left arm waved up and down indicating that left or right side engines respectively should be slowed down. At night is the same as day signal with addition of wands. NOTE: This gesture has also been observed as: grab one thumb with the other hand and wiggle back and forth.

Launch Gestures
(Those used in addition to taxiing gestures)

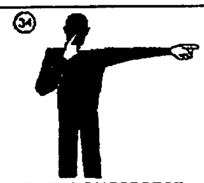 ENGAGE NOSEGEAR STEERING	Point to nose with index flinger while indicating direction of turn with other index finger. At night is the same as day signal with addition of wands. This command is given to the pilot to turn the nose wheel prior to entering the "Y".
 DISENGAGE NOSEGEAR STEERING	Point to nose with index finger, lateral wave with open palm of other hand at shoulder height At night is the same as day signal with addition of wands. This command is given to the pilot prior entering the "Y".
 LAUNCH BAR UP	One arm down and supported above the elbow. Pivot it up.
 LAUNCH BAR DOWN	Support one arm under the elbow in point up position. Pivot it down.
 SPREAD WINGS/ HELICOPTER BLADES	Arms hugged around shoulders, then swept straight out to the sides. At night is the same as day signal with addition of wands.

Recovery Gestures **(Those used in addition to taxiing gestures)**	
	Right fist, thumb extended upward, raised suddenly to meet horizontal palm of left hand. At night is the same as day signal with addition of wands..
	Right fist, thumb extended downward, lowered suddenly to meet horizontal palm of left hand. At night is the same as day signal with addition of wands.
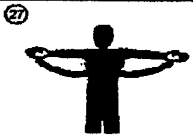	Arms straight out at sides then swept forward and hugged around shoulders. At night is the same as day signal with addition of wands.
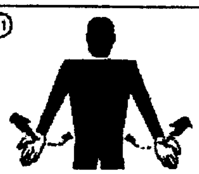	Arms by sides, palms facing forward, swept forward and upward repeatedly to shoulder height. At night is the same as day signal with addition of wands.

Appendix B – Trade-Off Analysis

The following was our initial trade-off analysis of concepts for UCAV-N deck handling that did not require additional manning on the flight deck.

Concept	Advantages	Issues
Hook up a tow tractor (vice taxiing)	• No technical risk	• Hookup times too long for recovery • Queue behind catapult too long • Additional tow tractors a negative impact to ship
Director controls UCAV via joystick	• Minimal technical risk • Hardware is cheap	• Large additional workload for director • Large negative impact to ops, training • Custom hardware • Something for sailors to break
"Data Glove" or accelerometers in wrist bands or wands	• Minimal impact to ops • Mature technology	• Expensive • Custom hardware • Something for sailors to break • Large reqt for available freqs
Autonomously follow deck markings	• Mature technology	• Fixed paths too inflexible for dynamic ops on flight deck
Sensor suite on ship	• Minimal impact to ops • No impact to aircraft	• Occlusion from aircraft, yellow gear, bodies • Infinite range of rotation of director • Weather, steam
Sensor on UCAV nose gear	• Minimal impact to ops • Manageable technical risk • Would work at airfield also	• Additional weight on aircraft • Weather, steam

Appendix C – Process for Gesture Recognition and UCAV Taxiing Control

The entire UCAV taxiing system would conform to the classic control configuration for intelligent control systems. This configuration is the Perception, Cognition, and Control loop as described in many intelligent control systems (see figure below).

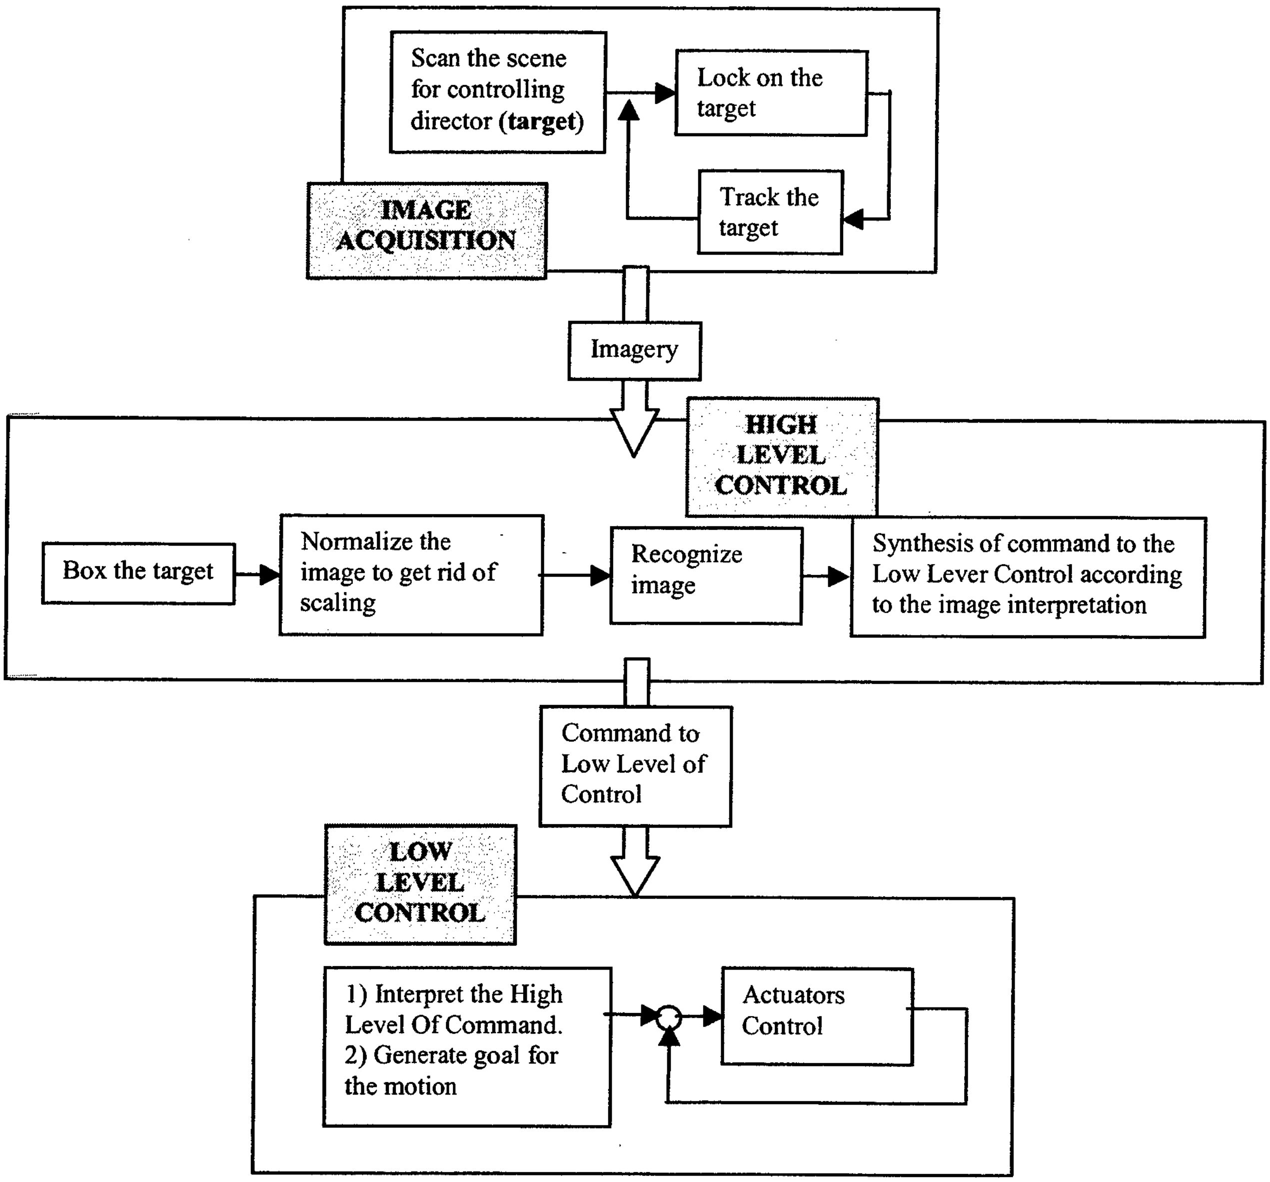

The Gesture recognition system as proposed should consist of three separate processes performed concurrently.

Image Acquisition: This subsystem should provide uninterrupted imagery stream to the High Level of Control. Primary functions of this subsystem are

1.1 Finding the target (director in control of taxiing) in the scene

1.2 Tracking the target.
1.3 Acquire and provide imagery of signaling director.

High Level of Control: This module should provide high level of decisions according to the imagery received from the acquisition unit.

1. Preprocess imagery for further image recognition
   1.1 Box the interest area
   1.2 Get rid of scaling problem
2. Recognize imagery. If the recognition is done by Fuzzy ART neural network than the rotational, noise and variability problems could be taken care of by the way associative memory is created.
3. Synthesis high level of control command. This command would contain necessary and sufficient information about the UCAV movement and actions during all stages of deck operations, be it arrestment, taxiing and parking or launch.

Low Level of Control: This module would interpret the High Level of command and compute necessary control for the UCAV actuators to perform taxiing.

Appendix D – Description of Fuzzy ART Algorithm

Fuzzy Art is a similar algorithm to ART 1 with fuzzy logic calculations to determine the relationship of memorized patterns to the current pattern presented for recognition. In ART 1 the weighted vectors stored in memory represent the patterns to be recognized. The intersection of the weighted patterns forms a region of uncertainty that presents difficulty for the standard classifiers to deal with. The region of uncertainty for Fuzzy Artmap can actually work to its advantage.

The basic principles behind fuzzy logic is the argument that a set A can intersect with its compliment A^c ($A \cap A^c \neq 0$) [Carpenter 1992]. The idea is that instead of unique existence for A and A^c, there exists a fuzziness at the boundaries between the two sets. The degree of fuzziness is a measure of this uncertainty. So it can be extended that if more than one set exist within this universe the measure of uncertainty is not a crisp result, but a measure of degree (fuzziness) that one set belongs to another within this region of uncertainty. Using this relationship, each set can be viewed as a fuzzy set. In these fuzzy sets the elements contain a measure of "membership" between all sets within their universe.

First, the input to the input layer (F_0) of Fuzzy Artmap normalizes the real number values within the range of [0,1] by the equation $a_i = IN_i / ||IN||$ where IN is the input vector supplied to the input layer and $||IN|| = \Sigma(INi)$. Normalization is performed to establish membership functions. The compliment of a_i is $a_i^c = 1 - a_i$, thus the input to the vigilance layer F1 becomes $I = (a_i, a_i^c)$. Notice the input set includes both the actual input and its compliment. Therefore, for each entry there exists a unique compliment with the exception when $a_i = 0.5$ ($a_i^c = 1 - a_i = 0.5$).

Network weights initially are set to 1 where $w_{j1}(0) = 1 \cdots w_{j2m}(0) = 1$; where m is number of inputs (notice, there are 2m weight for each row including the compliment). The output nodes of layer F_2 are uncommitted. Weights are determined by the test for vigilance $\frac{|\mathrm{I} \wedge \mathrm{w}_{\mathrm{J}i}|}{|I|} \geq \rho$, where ρ is the vigilance parameter. However, in this case ρ takes on a very special role that keeps the algorithm true to the fuzzy domain. Notice in the above relationship, If $\frac{|\mathrm{I} \wedge \mathrm{w}_{\mathrm{J}i}|}{|I|} < \rho$, it is said a mismatch occurs and the output for this node is set to zero so it is not chosen again. It forces the nodes above the selected category for the I being evaluated to be smaller than the node being updated. So, the weights are updated (resonance) if the vigilance criteria is met $\frac{|\mathrm{I} \wedge \mathrm{w}_{\mathrm{J}i}|}{|I|} \geq \rho$. This is also $|\mathrm{I} \wedge \mathrm{w}_{\mathrm{J}i}| \geq \rho |I|$. The weights are updated by: $w_J^{(new)} = \beta (I \wedge w_J^{(old)}) + (1-\beta)\, w_J^{(old)}$

β sets the speed of re-coding. Observe, however the weights must converge to $|\mathrm{I} \wedge \mathrm{w}_{\mathrm{J}i}| \geq \rho |I|$ to meet the stopping condition.

$w_J^{(new)} = \beta (I \wedge w_J^{(old)}) + (1-\beta)\, w_J^{(old)} = \beta (\rho |\mathrm{I}|) + (1-\beta)\, w_J^{(old)}$, and if β is set to 1, fast learning occurs giving the update values by $w_J^{(new)} = (\rho |\mathrm{I}|)$. As ρ is increased, the degree of fuzziness is decreased, and as ρ is decreased the opposite is true. Thus, defining the width of the membership between each training set.

$Tj(I) = \frac{|I \wedge w_J|}{\alpha + |w_J|}$, where $|I \wedge w_J|$ is the fuzzy "and" defining the intersection of the two fuzzy sets. min(I,w). (α=0.001 for the study presented in this paper). The output winning node is Y_J=max(T_j).

This is summarized in the figure below.

Fuzzy Adaptive Resonance Theory (Fuzzy ART) Block diagram

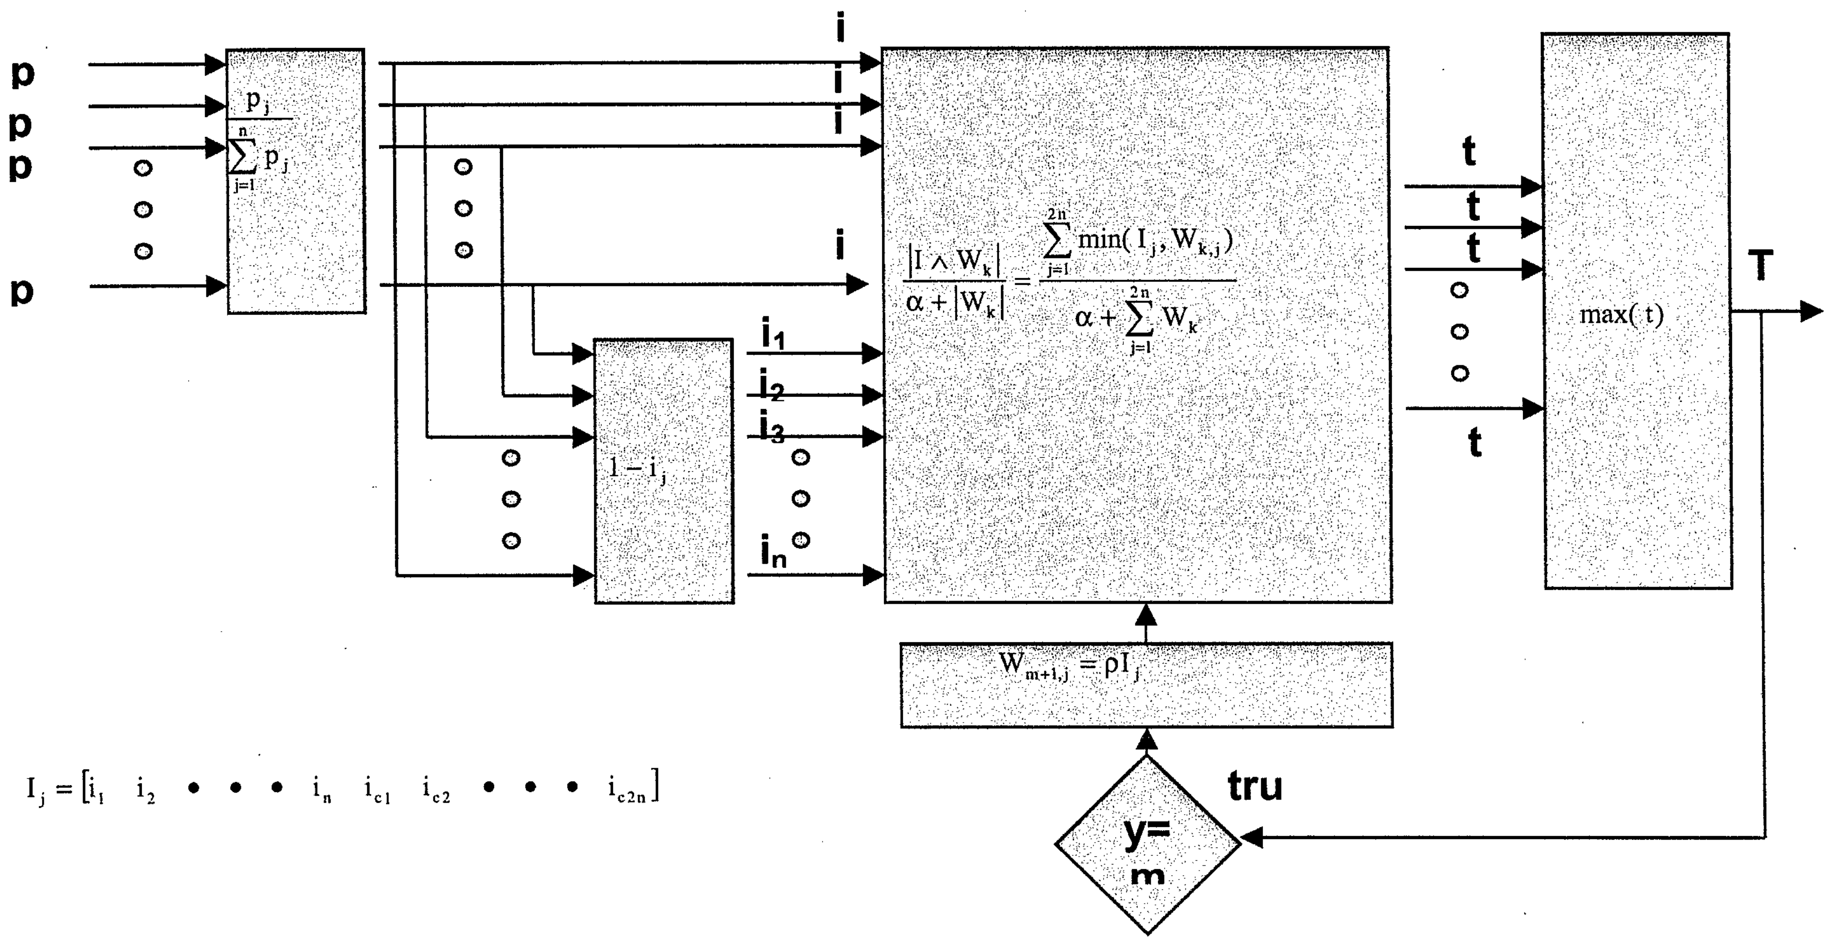

$$I_j = [i_1 \quad i_2 \quad \bullet \quad \bullet \quad \bullet \quad i_n \quad i_{c1} \quad i_{c2} \quad \bullet \quad \bullet \quad \bullet \quad i_{c2n}]$$

Appendix E – Analysis of Director Positions

The following are the results of photogrammetric analysis of in-house video to determine requirements for director positioning on the flight deck.

Descriptor	VHS Tape Index HH	MM	SS	Flight Director Coordinates X	Y	Aircraft Coordinates X	Y	Distance From Director To Aircraft (ft)	Orientation Of Director Relative To Aircraft (deg)	Orientation Of Aircraft Relative To Director (deg)	Handoff
devfltdir01	0	1	44	493	9	364	-11	130	0	90	N
devfltdir01	0	1	50	492	8	390	-13	104	-10	90	Y
devfltdir01	0	1	50	572	1	390	-13	182	0	0	N
devfltdir01	0	2	13	529	-28	574	-1	52	-30	180	Y
devfltdir01	0	3	1	490	0	405	-30	90	-30	10	N
devfltdir01	0	3	17	499	0	517	-28	33	-95	180	N
devfltdir01	0	4	9	555	-40	599	-7	55	-30	0	N
devfltdir01	0	4	19	554	-41	563	-5	37	-90	0	N
devfltdir01	0	4	26	450	18	542	-3	94	0	0	N
devfltdir01	0	4	44	452	21	449	-3	24	80	90	N
devfltdir01	0	6	32	414	72	470	-8	97	-10	0	N
devfltdir01	0	6	43	429	72	462	46	42	-20	-30	Y
devfltdir01	0	6	46	421	57	461	47	41	-80	-30	N
devfltdir01	0	6	50	423	55	460	51	37	-80	-90	Y
devfltdir01	0	6	50	306	-1	460	51	162	-110	0	N
devfltdir01	0	7	5	300	-2	331	13	34	-10	0	Y
devfltdir01	0	7	5	164	15	331	13	167	20	0	N
devfltdir01	0	7	32	317	17	436	57	125	-115	0	N
devfltdir01	0	7	59	288	-25	275	0	28	-90	0	Y
devfltdir01	0	7	59	226	-24	275	0	54	0	0	N
devfltdir01	0	8	20	421	47	447	-4	57	-10	80	N
devfltdir01	0	8	42	435	60	451	34	30	-5	0	Y
devfltdir01	0	8	42	531	66	454	34	83	80	0	N
devfltdir01	0	8	57	526	67	489	60	37	-15	20	Y
devfltdir01	0	9	8	419	51	457	-6	68	-10	70	N
devfltdir01	0	9	24	407	52	458	43	51	-70	0	Y
devfltdir01	0	9	24	274	14	458	43	186	-90	0	N
devfltdir01	0	9	53	278	12	290	17	13	0	0	Y
devfltdir01	0	9	53	154	0	290	17	137	0	0	N
devfltdir01	0	16	32	571	25	561	55	31	-70	20	N
devfltdir01	0	16	48	573	21	558	21	15	-90	0	N
devfltdir01	0	17	28	555	-39	513	-4	54	0	0	N
devfltdir01	0	17	44	552	-38	542	-19	21	20	0	Y
devfltdir01	0	17	44	561	51	542	-19	72	-90	0	N
devfltdir01	0	18	6	502	-62	381	-36	123	0	0	N
devfltdir01	0	18	43	498	-54	311	-24	189	15	0	N
devfltdir01	0	19	19	490	-39	247	-35	243	0	0	N
devfltdir01	0	26	22	541	68	506	57	36	5	0	N
devfltdir01	0	26	29	541	68	518	47	31	-50	0	Y
devfltdir01	0	26	29	597	3	518	47	90	30	0	N
devfltdir01	0	30	46	580	1	468	16	113	80	0	N

Appendix F - Fleet Discussions

9/24/02 Discussion with CDR Mike Yoast, CNAL Handling Team Officer

Multiple directors handling multiple aircraft are a problem, especially at night when it's hard to tell whether or not the director is facing you, the pilot. There have been numerous times at night when taxiing aircraft had to be stopped because pilots were responding to the wrong director.

CDR Yoast suggested that the plane captains could assume the role of handling the joystick. This would be a low technical risk option that would not require any additional flight deck manning. He noted, though, that there would be a training issue since the plane captain is typically junior and un-experienced.

CNAL Handling Team is involved with both UCAV contractors, and has been going to meetings on UCAV deck handling.

9/24/02 Comments from Shaun Donnelly, UCAV Deck Handling Lead

Re: Differences in director technique. A big difference is in the "come forward" signal. Some directors pivot at the elbows and swing their entire forearms. Other directors hold their forearms and wave their hands in front of their face.

Bright sun behind the director could be a big issue. Pilots frequently have a difficult time seeing the director in those conditions.

Trip to the USS Harry S Truman, CVN-75, 10/8/02

Interviewed and filmed 5 directors performing the various NATOPS signals that would apply to the UCAV. Purpose: To get an idea of the ways signals can differ over a range of directors, either because of director technique, fatigue, or training on a different ship.

The following are comments from the interviews:

<u>Orientation/Position:</u>

Worst case for director position relative to aircraft is +\-90° from the aircraft nose.
Differs for different aircraft based on pilot vision and canopy size.
E.g. F/A-18 has the least obstructions on the canopy – and the largest scope
For EA-6B and E-2, the copilot sits side by side with the pilot, blocking the pilot's view. In those cases the director will only go as far as 45° from the nose on the starboard side.

<u>Director Rotation:</u>

Director faces the aircraft most of the time.
Turns only during pass off. First to see if the next director is there, then to pass off.

<u>Distance from aircraft to director:</u>

During nighttime, the director will stand closer to the aircraft than during the day.

FLY-1: Max: From front of island to JBD
Min: 5 feet

FLY-3: Max: 50-75 ft Day
25-30 ft Night
Min: 10 ft

Signals that are similar to each other:

Engage, disengage nosegear – for towing (doesn't apply to UCAV)
Brakes on, off
Hook up, down
Wingfold, wingspread
Stop, brakes on are similar

Differences:

Wingspread

1. Start: hands on shoulders, or hands on elbows
2. Arc overhead, or
3. "Safe" sign

Turn Left/Right

Arm position could be pointed upward or horizontally
Open hand, 1 finger, or 2 fingers

Stop

Some people have hands together
Some people have hands apart

Brakes off

1 flash or multiple flashes

Slow down engine on indicated side

Two ways:
- Grasp thumb for "throttle back" indication
- Push down with open hand

Grasping thumb could be either hands on chest or hands away from chest

I Have Command

Receive from yellow shirt – hand up, either right or left
Receive from ordies – fists on chest

Anomalies – signals used on the ship but are not in NATOPS

1. Head nod – commanding the pilot to nudge nose wheel steering on cat
2. Kick – "Director, get over there", or "Pilot, cut the turn faster"
3. Touch nose and point to a direction for a push back
4. Turn on pivot – wave on one side, hold fist on the other (for holding the brake on this side)
5. Point to eyes – "You're not looking at me" or "You're not responding"

6. Head nod could also convey confidence – "Yeah, you've cleared the combing."

<u>Night time signals</u>

Brakes on (cross wands)

Brakes off (flick wands like snapping twigs)

Suspend launch (sweep horizontally at waist level)

Throttle back (hold wand parallel to deck, move up and down) – Could be right hand or left hand

Appendix G - References

Bobick A, Davis J, Intille S (1997) "The KidsRoom: an example application using a deep perceptual interface." In: Proceedings Perceptual User Interfaces, pp 1–4, October

Bradski J, and Davis J (2002) "Motion segmentation and pose recognition with motion history gradients." In Machine Vision and Applications 13: 174–184

Carpenter G, Grossberg S, Markuzon N, Reynolds J, and Rosen D (1992) "Fuzzy ARTMAP: A Neural Network Architecture for Incremental Supervised Learning of Analog Multidimensional Maps." In IEEE Transactions on Neural Networks, Vol 3 No 5, September

Cohen C, A Brief Overview of Gesture Recognition. Available online: www.dai.ed.ac.uk/CVonline/LOCAL_COPIES/COHEN/gesture_overview.html

Fisher R (ed). On-Line Compendium of Computer Vision. Available online: www.dai.ed.ac.uk/CVonline/

Kortenkamp D, Huber E, and Bonasso R P (1996) "Recognizing and interpreting gestures on a mobile robot." In Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI '96).

Starner T and Pentland A. "Visual recognition of American Sign Language using Hidden Markov Models." In IEEE International Symposium on Computer Vision, November 1995.