Visual Features Involving Motion Seen from Airport Control Towers

Stephen R. Ellis*. Dorion Liston.**

**NASA Ames Research Center Moffett Field, CA 940355*
USA (Tel: +1 650 604 6147; e-mail: Stephen R. Ellis@nasa.gov).
***San Jose State University Foundation, NASA Ames Research Center, CA 94035*
(e-mail: dorion.b.liston@nasa.gov)

Abstract: Visual motion cues are used by tower controllers to support both *visual* and *anticipated* separation. Some of these cues are tabulated as part of the overall set of visual features used in towers to separate aircraft. An initial analyses of one motion cue, landing deceleration, is provided as a basis for evaluating how controllers detect and use it for spacing aircraft on or near the surface. Understanding cues like it will help determine if they can be safely used in a remote/virtual tower in which their presentation may be visually degraded.

Keywords: Human vision, air traffic control, visual motion, human factors, perception.

1. INTRODUCTION

The visual cues necessary to fly and land an aircraft have been well studied over many decades (e.g. Gibson *et al*, 1955; Grunwald & Kohn, 1994). In particular, the degradation in piloting performance and the consequent need to reduce airport capacity due to bad weather is fairly well understood. (FAA 71010.65R, 2006). The present report outlines an approach to a complementary side of the airport capacity-safety trade-off quantifying the visual features and properties used by tower controllers. These features are now particularly interesting due to recent proposals for technology and procedures in which controllers lose visual contact with their controlled space as in a "virtual tower" physically removed from the airport (JPDO, 2007).

What important visual information would be lost in a virtual or remote tower and how may it be replaced by new aircraft/airport sensors or managed by new procedures or displays? Could controllers return to some form of visual procedures were these new sensors and associated displays and procedures to fail? These are some of the questions ultimately needing answers.

The following discussion of these questions will first point out relevant visual elements of the control task facing the tower evident in previous task analyses (see below). However, this earlier work appears to only provide very general descriptions of the specific visual features to which that the controllers attend. To the extent the visual functions important to the controllers are discussed with precision, they are generally limited to questions of detection, recognition and identification. The following discussion will consider other visual features, in particular motion of the controlled aircraft. The preliminary conclusion of the discussion is that tower controllers use visual features to provide predictive position information allowing them to use *anticipated separation* to effectively and safely merge and space aircraft to maximize airport capacity.

The visual cues used by controllers are important for several reasons. In the first place, there is FAA interest in increasing airport capacity so that current operations under non-visual flight rules with reduced capacity may be modified to allow higher capacity visual operations. For this purpose the currently used visual information needs to be provided by alternative means. Such "Equivalent Visual Operations" described FAA/NASA planning documents may be achieved with synthetic visual system, i.e., (Kramer, Williams, Wilz & Arthur, 2008) visual replacement with visualized electronic position data. But these replacements will not be fully successful, and may even be tragically misleading, if the useful visual affordances provided by the real scene are not appropriately provided. Although Equivalent Visual Operations has primarily been thought of from the pilot's viewpoint and discussed in terms of flight displays using new sensor data for synthetic vision, it has a flip-side for which synthetic vision or camera-based displays could present useful visual features in a remote or virtual tower.

Significantly, this information need not be provided in the form of an image, but could be provided in a more map-like plan view format and conceivably could even come along non-visual sensory channels, e.g. auditory or haptic. In fact, it could be based on data directly down-linked to ground displays from the aircraft indicating its state, e.g., spoilers deployed (Hannon, et al., 2008).

The visual environment in an airport tower may be illustrated by considering the view from a tower such as that of San Francisco International Airport (SFO). Such tower views show significant perspective compression at the ~1 n.mi. range to runways and taxiways, making commercial aircraft subtend small visual angles, and posing viewing difficulties due to background visual clutter. Interestingly, during low visibility CAT III operations at SFO, airport operations may be conducted with the controllers never actually seeing the aircraft. Thus, since it is already possible for the controllers

to continue many of their control tasks without visual contact, the idea of a remote tower may have merit. But without visual contact, they must inform the pilot and those monitoring their communications that visual contact has been lost. At the SFO tower continued operation without visual contact is associated with a significant loss (~50%) of airport capacity. In contrast at an airport such as Arlanda, Sweden (ARN) with the parallel runways ~1 km apart, total loss of visual contact can have virtually no impact on capacity when ground radar is properly functioning.

2. SFO OPERATIONS

A more specific analysis of the role of visual features in tower control can be developed from a more concrete discussion of operations for a particular airport, SFO. A sense of the overall strategy for some aspects of usual airport operation at SFO is best gotten from plan-views maps (Figures 1 and 2). Aircraft are taxied from their gates to the south ends of runways 1L and 1R and launched in staggered pairs that are interleaved between approximately paired aircraft landing on Runways (RW) 28L and 28R which taxi to their gates crossing to the southwest. (Figure 1.). Current winds, weather, and special operational requirements, of course, can significantly alter this pattern. Sometimes the longer 28 runways are needed, for example, for departing heavy transpacific aircraft. Detailed descriptions of the approach and departure procedures can be found in the Standard Instrument Departures (SID) and Standard Terminal Arrival Routes (STARS) associated with the airport, but the local controller's responsibility generally begins with radio contact somewhat before the aircraft crosses the San Mateo Bridge and ends 1 n.mi. beyond the departure runway. By FAA rules, the local controller is generally responsible for aircraft entering and leaving the runways whereas the ground controllers handle most of the taxing to and from the gate. These two positions in addition to that of the supervisor are the ones that make the most use of the out-the-window information. The Flight Data and Clearance Delivery controllers primarily use inside-the-tower information sources and voice communications.

3. VISUAL INFORMATION IN THE TOWER

The primary responsibility of the control tower is to ensure sufficient runway separation between landing and departing aircraft. (FAA, 2006) A back propagating process may be used to understand the visual requirements supporting the tower controller's primary responsibility.

This process first identifies the visual affordances that the controllers' tasks involve. Affordances are the higher-level behavioural capacities that vision must support. Controllers, for example, must identify the aircraft type, company and flight status. They must control and recognize aircraft speed, direction and position. They must establish a movement plan involving a succession of spatial goals. They must communicate this plan to the aircraft, coordinate it with other controllers and pilots as necessary, establish whether aircraft comply appropriately, and recognize and resolve spatial and other conflicts that may arise. These higher-level elements are supported visually by a number of visual functions: detection, recognition, and perception of the static and dynamic state of the aircraft. These functions are supported by still lower level visual mechanisms underlie luminance, color, control, position, and movement pro-cessing. These three levels of analysis provide a basis for describing the controllers visual task.

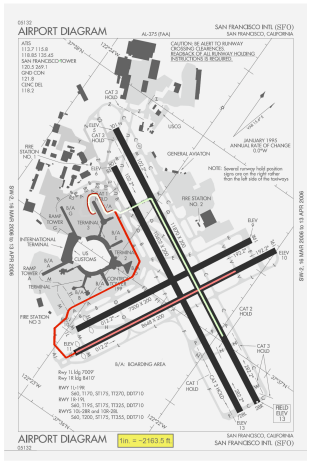

Fig. 1 *SFO airport diagram showing typical movement paths for United Airlines, departures (dark/red), arrivals (light/ green)*

Fig. 2. *The first and last positions where SFO controllers report useful visual information w/r to landing (RW 28) and departing aircraft RW 1).*

The tower controller's overall task has, of course, been analyzed within and outside of the FAA. It may be broken down to six different job subtasks: separation, coordination, control judgment, methods/procedures, equipment, and communication. The five of these subtasks which involving vision have been identified by boldface type in Table 1.

The assurance and maintenance of spatial separation is, of course, a visual task since regardless whether separation is determined by radar or direct view, it is definitely recognized visually. Handoffs and point-outs clearly are also intrinsically dependent upon vision, though the need for the controller to adopt the pilot's spatial frame of reference to direct attention toward objects and aircraft is also a significant cognitive task. Control judgment, being essentially a mental and cognitive issue, does not have an intrinsically visual component. But its connection with maintenance of effective and efficient traffic flow does emphasize the critical importance of time in traffic control. Three general methods and procedures directly involve vision. These include establishment and maintenance of aircraft identify, posting and correct annotation of flight strips, and continual scanning of the entire control environment. Associated with these methods is the admonition to work quickly and rapidly recover from errors or off nominal conditions. Because each tower's environment is to some extent unique, the specifics of their procedures differ from tower to tower. All control techniques are, of course, consistent with the regulations cited and described in the FAA air traffic control, *Order 7110.65R,* but unique procedures and heuristics are passed on to future controllers by onsite training.

The overall tower control process has been formally analyzed and modelled including visual and nonvisual components (Alexander *et al.,* 1989, Werther, 2006). For example, the MANTEA notation (Zografos & Hesselink, 2000) has been applied to analyze controller activity in the tower. Some of the elements identified in the MANTEA analyses are, in fact, visual, but the visual components are only described in very general terms such as "visualize runway," "visualize meteo," etc. These descriptions really only identify the sensory modality used to gather the information and a general description of the content of the visual information, but they say nothing specific about the actual visual viewing conditions or about the specific visual stimuli. This feature is, in fact, common in other more recent and more sophisticated task analyses of visual features seen from the tower. Even the recent modelling done with Petri nets (Werther, 2006) does not identify specific visual stimuli but is more connected to estimates of time required for the precision with which various visual subfunctions maybe executed and to the logical conditions and consequences associated with the functions.

The FAA has done some analysis of the specific visual performance expected from Tower controllers. The work primarily focuses on the controller's surveillance function and has been based on visual performance models developed for the military by CERDEC at Ft. Belvoir (e.g., Vollmerhausen & Jacobs, 2004). These models primarily are intended to predict the probability of visual detection, recognition, and identification of known targets. "Detection" refers to users' ability to notice the presence of a particular object. "Recognition" refers to their ability to categorize the object into a general class such as a tank, light aircraft, or truck. "Identification" refers to their ability to determine the specific type of object, i.e., an Abrams tank, a Cessena 172, or a Ford refueling tanker. More modern similar visual performance models do not require specific calibration techniques to determine model parameters for specific targets (Watson, Ramirez, & Salud, 2009).

The CERDEC analysis predicts specific object perception from towers of various heights, during a variety of atmospheric conditions and object distances has been incorporated into a web tool to help tower designer ensure that specific architectural and sitting decisions for new towers will meet FAA requirements. Significantly, this tool also just focuses on the surveillance function and does not address the aspects of visual motion that tower controllers use for the information, separation and safety tasks.

In order to understand the details of the visual features used in tower control it is first necessary to identify the range within which controllers use visual information. We can use the example of SFO. Informal voluntary discussions with ten active controllers who work at this tower were analyzed for the physical locations identified as points where various types of visual references are taken from approaching or departing aircraft (Figure 2.). These points include positions where visual contact with the aircraft is first or last considered to provide useful information. These positions include those for which they come under or leave tower control, where they pass important ground references, or where visual contact provides other useful information. The points were determined independently from each of the controllers in response to the question, "When you are in the Local controller position, where are the aircraft when you usefully observe them visually, what visual aspects of the aircraft do you observe and why?" Controllers could designate more then one point of interest for departing and more than one for arriving traffic, only two controllers took this option. One point represents nine controllers' overlapping responses identifying approximately the same location about 1 n.mi. beyond the end of the departure Runway 1.

In general it is apparent from the distribution of points that controllers' visual attention is much more spatially distributed to the aircraft approaching the 28LR runways and rather abruptly drops off about 1 mile off the end of the usual departure runways 1LR. These observations refer to the most common aircraft flow at SFO but suggest the generalization that the local controllers' visual attention to approaching aircraft is distributed over a much large area than that corresponding to departing aircraft. A likely reason for this is that departing traffic is handed off to approach/departure control and generally not thereafter of concern to the tower.

A significant aspect of the controllers' remarks concerning when they first start paying visual attention, or when they last pay attention, to aircraft is that they rarely mentioned the aircraft's visual motion[1]. One reason is that for the viewing angles and distances to the aircraft approaching SFO, this motion is very small in terms of degrees per second, often the azimuth rate is on the order of much less than 0.25°/s and rarely more than 0.5°/s. The visual accelerations are even much smaller and difficult to see because of atmospheric haze, thermal effects, and the visual range beyond 5 miles. Visual rates of motion are more important for closer aircraft on or near the runways and taxiways.

Probably the most obvious need for visual contact by controllers in the tower is to immediately note unusual events that are not detected by electronic sensors such as radar. Examples could be heavy bird activity or an aircraft leaking fuel onto a taxiway. But there are a wide variety of other visual features that controllers use on a more regular basis when aircraft are close enough for the visual motion to be more easily noticed. Discussions with controllers as provided a list of some that are used.

A tabulation (Table 1) of the visual features mentioned in the discussions with ten SFO controllers shows the relative frequencies with which different features were mentioned. These discussions used a "cognitive walk-through technique in which the controllers were asked to imagine representative approaching, departing, and taxiing aircraft under a variety of visual conditions and to report what they looked for visually to assist their control tasks. The most frequently mentioned features were relative motion between landing or taking off aircraft and obstacles that could be on the runway. The first

[1] Visual motion is defined as the angular rate of change of the line of sight angle to an aircraft from the tower.

of these features is probably prominent because SFO has intersecting runways commonly used for takeoffs and landings. An assessment of all of the features mentioned, however, shows what may be a more general element. Seven of the 13 features identified in the interviews note that the feature helps the controller anticipate future activity. This information provides insight into pilot intent, knowledge, and likelihood of aberrant behaviour. These predictive cues help the controller with the short term planning needed for *anticipated separation*.

Table 1. Counts of visual features identified in interviews with 10 SFO controllers.

Feature	#	Comments
1. Relative visual motion used to interleave departures and arrivals	5	Controllers verify predicted separation using relative motion w/r stationary references.
2. Visual check for obstacles or other a/c to verify a clearance	5	Obstacle checks include ground vehicles, aircraft, birds, people
3. Taxing "With authority" helps attention allocation	4	Fast and confident a/c motion allows controllers to distribute attention to pilots who maneuver hesitantly allowing **anticipation** of future problems
4. Aircraft attitude/ altitude predicts a "Go Around"	4	Controllers anticipate "Go Around" by checking a/c passage through various approach gates defined by altitude and attitude
5. Visually apparent acceleration, speed or turn rates	4	Controllers mentally integrate motion features to anticipate taxiway and ground route selection
6. Visual and radar position and speed are cross-checked	4	Specific visual land-marks are selected to cross check radar
7. Visible wing dip predicts coming turn	3	Visible banking confirms initial conformance to turn clearance
8. "Mike and a mile" rule for interleaved take offs and landings	3	Predictive rule: Departing A/C must be rolling across taxiway Mike on RW1 when matched landing A/C on RW28 is at least 1 mi out for required separation.
9. Engine smoke or heat confirms take off start	3	Modern engines don't smoke much and have cooler exhaust
10. Onset of navigation lights precedes call to tower	2	Controllers can anticipate coming workload
11. Visual resolution of motion and position is better at airport than radar	1	1-2 nmi. From the tower the "visual display" of the real world has more "pixels" than associated radar displays
12. Visual double check on a/c tail to verify company	1	
13. Check landing gear	1	This is an isolated comment probably because it is done automatically and is an infrequent is issue for major airlines.

Many of the predictive cues, particularly the motion-based cues, are available because of the high dynamic fidelity of direct visual contact. Because of noise, sampling artefacts, or undue delay these features may not easily be seen on electronically mediated information such as that on radar displays. A better understanding of exactly what these continuous cues are can be developed by examining one them quantitatively. An example of such analysis is presented below with respect to landing deceleration at SFO.

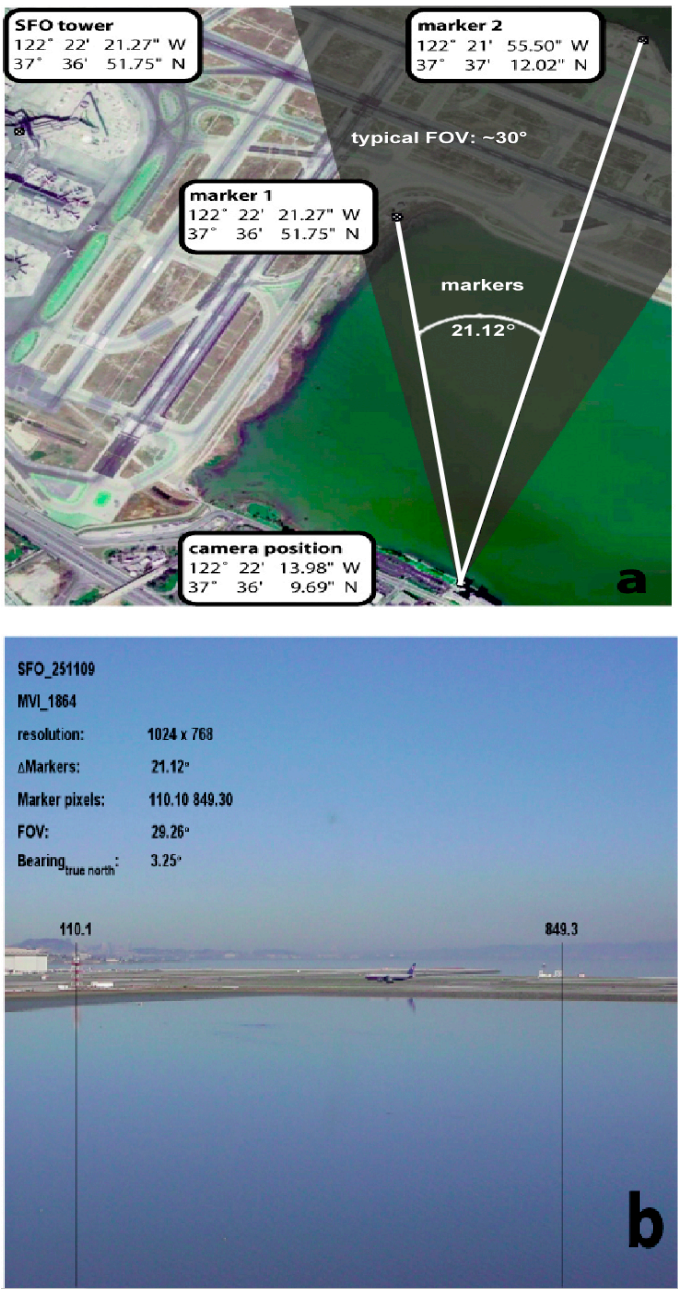

Figure 3. The video field of view used for motion tracking during the braking phase of landing.

4. DECELERATION DURING LANDING AT SFO

The following analysis begins to determine the magnitude of this visually sensed deceleration and how it could be used by controllers. In doing so we will identify one of the dynamic visual features used in traffic control from the airport

tower: the change in speed evident during a single glance a controller might make towards a decelerating landing aircraft[2]. In thinking about what specific aspects of the visual stimulus to which the controllers might be attending, it is helpful to remember that perceptual discriminations of commonly experienced magnitudes of sensory quantities such as velocity are fairly well described by Weber's Law, which states that the just noticeable difference is a constant proportion of the quantity's magnitude. This so-called Weber fraction generally had a midrange value around 3-5% (Goldstein, 2007) for common perceptual discriminations such as those that could be made from the control tower for landing and departing aircraft.

It is also important to understand that controllers may not be directly sensing the visual velocities *per se* but may develop viewing strategies allowing them to translate speed into displacement during relatively fixed time intervals, thus making the detection of unusual rates of change easier. Additionally, alterative visual cues to quantities such as deceleration, but which are correlated with them, could be used. Aircraft pitch, for example, could be a clue to braking.

It is not so much the visual aspect of the visual information that is important as it is the fact that the information revealed by vision is relevant, real, direct, unmediated, immediate and continuous that makes it possible for the best possible anticipation of future action. This is why the visual input could be critical. Replacements for it need to capture the same predictive, informational features.

In order to begin to analyze the visual features actually present in real landing trajectories we have initially focused on the deceleration profile of aircraft landing on the 28L and 28R runways at SFO. Controllers report that they use their sense of degree and timing of this deceleration to anticipate which taxiway would be needed for the aircraft to exit the active runway. This decision is time critical during heavy runway use since landing aircraft are staggered in pairs and interleaved with departures on crossing runways 1R/1L.

We have made 15 frame/s video recordings at 1024 X 768 resolution of the braking phase of 45 aircraft landing on 28L and 28R and processed the recordings to measure changes in visual velocity. We used a custom MatLab image processing technique that isolated the moving contours across a set of two frames and averaged them to localize the aircraft and provide their screen velocity in degrees per second. Using the viewing geometry described in Figure 3, we have recovered the aircraft braking profile and computed the changes in visual velocity as viewed from the control tower by re-projecting the movement, as it would have been seen from the tower. Thirty of these velocity profiles (low pass filtered with a 1Hz cut-off) are shown in Figures 4.

Because of the noise present in our current recording technique, we were unable to obtain velocity and acceleration values with acceptable noise levels. We were, however, able to obtain a braking deceleration profile for an A319 aircraft landing on Runway 28 from the same company, comparably loaded and flying in the same wind and weather conditions as one of the aircraft we had recorded visually. Since we knew the touchdown points for these two A319 landings, we've combined the two trajectories producing what we believe is a fairly accurate landing profile as seen from the tower (Fig. 5).

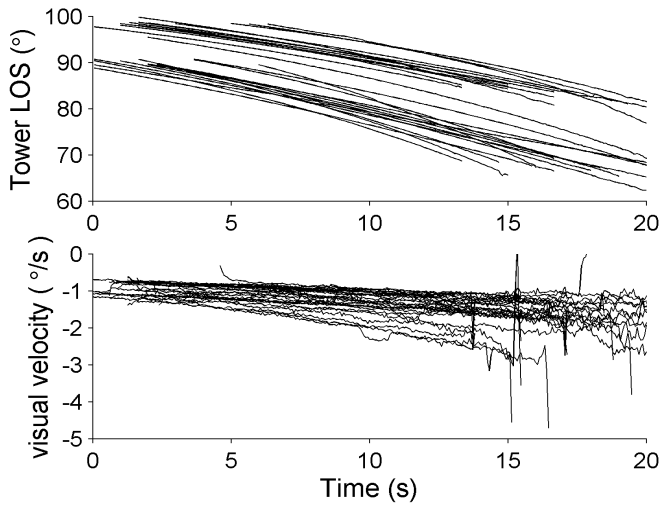

Figure 4. Line of sight (LOS) position and visual velocity of 23 aircraft landings at SFO showing growing tracking noise after ~14 seconds.

Laboratory research has examined the perceptual mechanisms humans use to perceive 2D acceleration in stimuli like our LOS data, and have quantified perceptual thresholds for acceleration as low as 30 deg/s^2 for integration times of 200 ms or more (Calderone and Kaiser, 1989; Stone and Ersheid, 2006). From the point of view of the control tower, the visual deceleration signals (e.g., the deceleration of the visual image of the plane) are much lower (e.g., <1 deg/s^2) than these threshold values, and highlight the clear distinction that must be made between deceleration of the visual image of the aircraft (e.g., the slope of visual velocity traces in Figs. 4 and 5) and the higher-order inference that the ground speed of the plane on the runway has decreased. Although these data allow us to rule out the idea that the former supports the latter, the higher-order mechanisms of inference supporting percepts of deceleration in the 3D world likely "re-project" velocity signals on the retina onto a mental model of the runway, allowing the construction of a deceleration percept. The present data do not speak to how this might occur, although design concepts for remote towers will have to identify the types of visual information that will allow controllers a robust percept of motion within the 3D world model.

[2]During normal vision, people make from 3-5 fixations per second (Rayner & Castelhano, 2007). However, when studying some aspect of an ATC image, fixations duration can increase but rarely grow longer than approximately 1.3 s (e.g. Remington, Lee, Ravinder, Matessa, 2004). Consequently, a reasonable constraint for modeling the duration of a controller's glance would be to insure that they are 1.3 s or less.

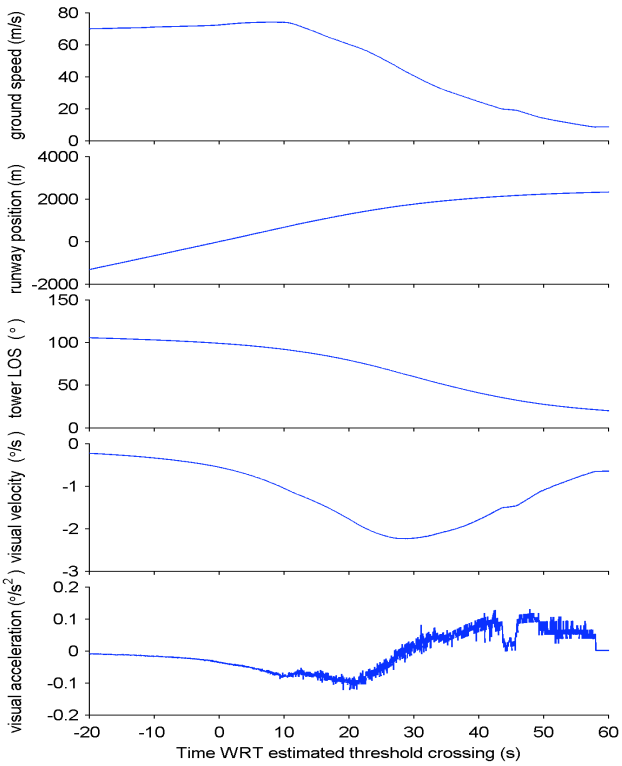

Figure 5. Line of sight changes for a single medium jet transport landing at SFO.

The deceleration profile in Figure 5 shows the aircraft approaching and passing the tower as it decelerates. In fact, during the approach the visual velocity actually increases during the deceleration because of the decreasing distance between the aircraft and the tower. It is clear from the deceleration profile that there are several phases of braking due to deployment of the thrust reversers, spoilers, and mechanical brakes and further data collection and processing needs to be done to more precisely identify these periods. However, the very smooth velocity plot in Figure 5 (third panel from top) already shows that the amounts of velocity change in the braking within any short time window 2 s or less are well less than the 3-5% Weber fraction for a just noticeable difference. This level is defined by convention to be that difference in a sensory quantity that can be detected correctly 75% of the time and is therefore not evidence of a very strong sensory stimulus.

This observation leads to some skepticism that the controllers are detecting velocity change per se, but rather may have developed a strategy to detect speed change by some other means, perhaps by comparing displacement for approximately equal time periods. Such a timing strategy might be evident in eye tracking records of controllers judging aircraft deceleration. Of particular interest will be future analyses and experimental to determine how well the controller's sense of aircraft deceleration can be maintained with airport imagery spatially degraded by pixelation and sensor noise, and temporally degraded by low sampling rate. These issues will be addressed by future research.

Acknowledgements: Discussions with L. S. Stone, NSF Award 09-250-01 and NASA/FAA institutional support.

REFERENCES

Alexander, J. R., Alley,V. I., Ammerman, W. S., Fairhurst, W. S. , Hostetler, C.M., Jones, G.W., and Rainey, C.L. (1989) FAA air traffic control operations concepts, Volume VII, ACTC Tower controllers. DOT/FAA/AP-87-01, FAA/DOT,Washington, D.C.

Calderone, J. B., & Kaiser, M. K. (1989). Visual acceleration detection: effect of sign and motion orientation. *Perception & Psychophysics*, *45*(5), 391-394.

Dorighi, N., and Sullivan, B., Future Flight Central: A Revolutionary Air Traffic Control Tower Simulation Facility, Proc. American Institute of Aeronautics and Astronautics (AIAA), Sep. 2003.

FAA Task order 7110.65R (2006) Air traffic control February 16, 2006, http://www.faa.gov/atpubs

Fürstenau, N., Möhlenbrink, C., Rudolph, M., Schmidt, M., and Halle, W. (2008) Augmented Vision Videopanorama System for Remote Airport Tower Operation. Proc. 26th Int. Congress of the Aeronautical Sciences, I. Grant (Ed.), Anchorage, Sept. 14-19 2008, ISBN 0-9533991-9-2.

Gibson, J. J., Olum, P, and Rosenblatt, F (1955), 'Parallax and perspective during aircraft landings', *American Journal of Psychology,* vol. 68, pp. 372-385.

Goldstein, E.B. (2008) *Sensation and perception, 7th Edition,* Wadsworth, N.Y. p.15.

Grunwald, A. J. and Kohn, S. (1994) Visual field information in low-altitude visual flight by line-of-sight slaved head mounted displays. *IEEE Systems Man and Cybernetics, 24*, 1, 120-134.

Hannon, D., Lee, J., Geyer, T., M., Sheridan, T.,Francis, M., Woods, S., Malonson, M. (2008) Feasibility evaluation of a staffed virtual tower. *Journal of Air Traffic Control,* Winter , 27-39

JPDO (Joint Planning and Development Office) (2007) Concept of operations for the Next Generation Air Transportation System Version 2.0, 13 June 2007, pp. 2-27--2-36.

Kramer, L. J., Williams, S. P., Wilz, S. J. and Arthur, J. J. III (2008) Simulation evaluation of equivalent vision technologies for aerospace operations. SPIE Proceedings, Vol 6557, 69570K (2008); doi: 10.1117/12.772775

Paul, S., Zografos, K., and Hesselink, H., (2000). MANTEA Final Report, MANTEA/ISR DOC-D83-137-R1, TR 1036, Telematics Application Programme, (Transport/Air).

Rayner, K. and Castelhano, M. (2007), Scholarpedia, 2(10):3649. doi:10.4249/scholarpedia.3649 revision #54795.

Remington, R. W., Lee, S. M., Ravinder, U., and Matessa, M. (2004) Observations on human performance in air traffic control operations: Preliminaries to a cognitive model. in *Proc. 2004 Behavioral Representation in Modeling & Simulation*, Arlington, VA.

Ruffner, J., Deaver, D. M., and Henry, D. J. (2003) Requirements analysis for an air traffic control tower surface surveillance enhanced vision system. SPIE Conference on Enhanced and synthetic vision 2003. Orlando FL, 9-10 September 2003 SPIE Proceedings vol. 5081, pp. 124-135, Bellingham, WA.

Stone, L.S., and Ersheid, R., (2006) Timecourse of human sensitivity to visual acceleration. Soc Neurosci Abstr. 32:641.5.

Vollmerhausen, R. H., and Jacobs, E. (2004) The Targeting Task Performance (TTP) metric: a new model for predicting target acquisition performance," Technical Report AMSEL-NV-TR-230, U. S. Army CERDEC, Ft. Belvoir, VA.

Watson, A.B., Ramirez, C.V. and Salud, E. (2009) Predicting visibility of aircraft. PLoS ONE. 2009; 4(5): e5594. Published online 2009 May 20. doi: 10.1371/ journal.pone.0005594.

Werther, B., (2006) Colored Petri net based modeling of airport control process. Computation. Intelligence for Modeling, Control and Automation (CIMCA) Sydney, 0-7695-2731-0