From virtual images to actions

Heinrich H. Bülthoff , Daniel Berger and Cengiz Terzibas

Max Planck Institute for Biological Cybernetics, Tübingen

Most experiments which study the mechanisms by which different senses interact in humans focus on perception. In most natural tasks, however, sensory signals are not ultimately used for perception, but rather for action. The effects of the action are sensed again by the sensory system, so that perception and action are complementary parts of a dynamic control system. To get a better understanding of how different senses interact in self-motion, we study the control of self-motion in a closed perception-action loop. Here we investigated how cues from different sensory modalities (visual cues and body cues) are used when humans stabilize a simulated helicopter at a target location.

Introduction

A helicopter is inherently unstable, and therefore flying a helicopter requires constant control input. If the helicopter is not perfectly upright, it will accelerate in roughly the direction to which it is tilted. If, for example, it leans slightly to the left, it will start accelerating to the left. To hover at a fixed position requires the pilot to continuously compensate such drifts, which can be caused also by external forces, e.g., wind gust.

Conventional helicopters are mainly controlled using four control input devices. The *cyclic stick* can be moved forward/backward and sideways, similar to a computer joystick, and controls pitch and roll rotations of the helicopter by changing each blade angle individually during a rotation cycle. Foot *pedals* control heading rotations by adjusting the blade angles of the tail rotor. The *collective* lever controls the angle of the main rotor's blades collectively, and with this the lift force of the helicopter. What makes helicopter flight control particularly difficult is the fact that the different degrees of freedom are coupled by the helicopter dynamics. For example, a change of the collective control changes the torque of the main rotor systems, and the resulting yaw rotation of the helicopter has to be compensated with the appropriate control input at the foot pedals. During training, pilots learn the dynamics of the helicopter, and with enough training pilots are able to hover the helicopter above a point with centimeter precision. Pilots handle this complex control problem without much cognitive effort, perhaps comparable to the (much simpler) control involved in riding a bicycle.

To date, it is still unclear which senses helicopter pilots' use for hover control. Common sense tells us that it is not possible to stabilize a helicopter without vision, since a blindfolded pilot has no means of distinguishing standing still from moving with a constant velocity[1]. Even when starting in a stabilized position, small disturbances will add up over time and the pilot will start drifting if he or she does not perceive and compensate those drifts in position and orientation. But vision is probably not the only cue for stabilization: pilots often report that a "seat of the pants" feeling, the sensation of accelerations from pressure on the skin and probably also from the vestibular system, is a particularly helpful cue for flight control.

There are several visual cues a pilot might use for stabilization. One important visual cue for the orientation of the helicopter in pitch and roll is provided by the *horizon*. When the helicopter tilts forward, backward, or from side to side, the horizon will move downwards, upwards, or roll in the pilot's view, respectively. Yaw (heading) rotations can also be noticed visually - and distinguished from lateral displacement - when objects which are far away (close to the horizon) move sideways. For nearby objects the perceived motion of a single object cannot be easily used to separate helicopter rotation from lateral drifts.

[1] Here we ignore the aircraft vibrations and flying noise under different flight conditions.

Another important visual cue is *optical flow*. When an observer moves, the visual features in the environment are perceived as moving, forming a particular motion field pattern (optical flow). This movement field contains all information necessary to reconstruct the motion of the observer (up to a scaling factor), if the visual features are not themselves moving and the shape of the terrain is known. However, it has been shown experimentally that humans are bad at detecting changes of forward self-motion velocity from optical flow (Monen and Brenner, 1994).

Apart from vision, pilots can also use force cues of self-motion. One important sensory modality to detect inertial forces on the head is the vestibular system in the inner ear, which can sense both rotations and accelerations of the head. There are also other sensors in the body, which can sense body accelerations, for example, pressure sensors in the skin.

The integration of multiple cues for vehicle action control has been extensively studied in driving (for review see Kemeny and Panerai, 2003). Motion cueing and force cueing for flight simulators have also been investigated in several studies (Heintzman, 1996; Telban and Cardullo, 2005). Chung et al. (2004) argue that the effectiveness of motion cueing depends on task, vehicle dynamics and the properties of the motion cueing. They suggest that future studies should carefully document the characteristics of the simulation and the algorithms used for motion cueing.

Experiments which specifically addressed sensori-motor control in helicopter stabilization, however, are comparatively rare. A study by Ricard and Parrish (1984) investigated the role of inertial motion cueing and visual delays on helicopter stabilization, with the simple methods available at the time. They found significantly better stabilization performance (measured in terms of mean-distance-to-target) when inertial motion cueing was available. For the visual delay, results were not so clear: most participants stabilized better with a shorter delay, but for some the opposite was the case, and others did not show any difference. An earlier study (Ringland et al. 1971, cited in Ricard and Parrish, 1984) found best control performance when only rotations, but not translations, were simulated with inertial motion cues. Hall (1978) also reported that inertial motion cueing can improve roll stabilization in a *Harrier GR Mk 3* flight simulation.

The experiment reported here focused on the effects of whole-body rotation cueing and two different visual cues of self-motion, a horizon and optical flow, on helicopter stabilization. The task was to stabilize the simulated helicopter on a target spot. Stabilization performance was measured with four different visual conditions, both with and without platform rotation.

Materials and Methods

Participants were seated on a Stewart motion platform (*CueSim MaxCue*), equipped with a projection screen (see Figure 1, left). The visual field was 70° horizontally and 54° vertically; viewing distance to screen was 1.19m. During the experiment the motion platform cabin was closed and the participant could not look outside. Noise-cancellation headphones playing noise and seat shakers were used to mask the sounds and vibrations of the motion platform. For a complete description of the motion platform set-up, see von der Heyde (2001).

A real-time simulation of the dynamics and aerodynamics of a small helicopter (similar to a *Robinson R-22*) was used for this experiment (Terzibas, 2004). The task for the participants was to stabilize the helicopter on a target spot. This was visualized using two spheres, one representing the target and the other marking the position of the helicopter (see Figure 1 right panel). While the participants stabilized, we continuously measured the distance of the position marker from the target in front-back and sideways directions, helicopter velocity, and pitch and roll angles of the helicopter. The latter are correlated with the accelerations of the helicopter, as the simulated helicopter accelerates in the direction to which it is tilted. Good stabilization is characterized by small distances to the target, low velocities and small tilt angles.

Participants used a realistic helicopter cyclic stick to control forwards-backwards and sideways drift, to stay as close as possible at a visible target. Height above ground and heading was automatically stabilized, so the *collective lever* and the *pedals* were not used. The cyclic stick was passive and did not provide any force cues.

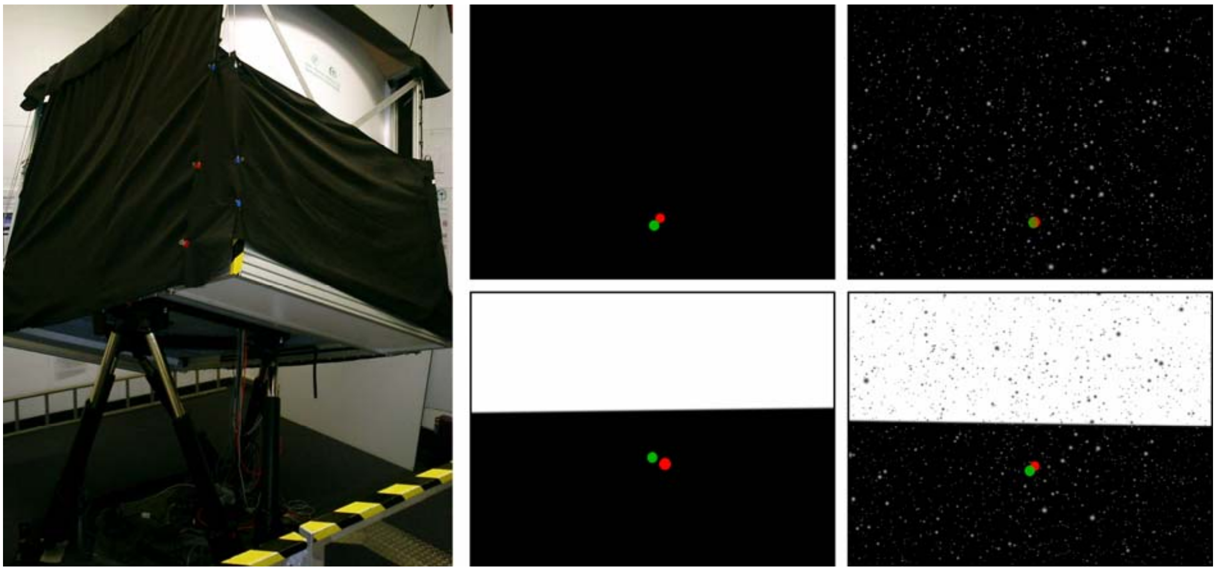

Figure 1: *Left:* Motion simulator set-up. *Right:* Four visual conditions: *black background* (B), random dot pattern to provide *optical flow* (OF), *horizon* (H), and both *horizon and optical flow* (H+OF). The red sphere is the target, the green sphere represents the position of the helicopter.

This experiment investigated the influence of the visual cues, *horizon* (H) and *optical flow* (OF), on the stabilization performance, and the interaction of those cues with whole-body motion cues. Five different visual scenes were used, four of which are reported here. In all of them, two spheres (0.6m diameter) were shown. The red sphere represented the target position and was fixed in the environment. The green sphere represented the position of the helicopter, and was placed at a simulated distance of 15 meters in front and 2 meters below the observer. The height of the helicopter was stabilized automatically at a height of 2 meters. The visual conditions presented together with the target and reference spheres were: *black background* (B), *horizon* (H), *optical flow* (OF), and both *horizon and optical flow* (H+OF) (see Figure 1, right). The random dot pattern used in the optical flow conditions were static in world coordinates and provided visual cues for self-motion, telling the participants about translations and rotations of the helicopter, but not about absolute position or orientation in space. In the H condition, participants were given visual information on helicopter pitch and roll orientation, and rotation velocity. The black background did not provide any visual cue, and served as a control condition to see how much information can be gained from the relative position or motion of the two spheres. The visual stimuli were designed so that they gave ambiguous information but not conflicting information. Optical flow does not provide any information about absolute pitch and roll because the random dot pattern is isotropic, and the horizon does not provide translation information because ground and sky are displayed without any texture.

All visual conditions were presented with and without platform rotation cueing, leading to 10 different stimulation conditions (8 of which are reported here). Pitch and roll body rotations were presented by tilting the platform exactly as the simulated helicopter tilted. The duration of each trial was 120 seconds. One experimental block contained one trial of each stimulation condition, with platform on and off conditions interleaved (20 minutes per block). Each participant ran 5 blocks with different pseudo-random trial orders, which amounts to a total of 10 minutes of stabilization per experimental condition.
Helicopter distance-to-target, velocity, and orientation were measured and analyzed for differences between the experimental conditions. Only results of distance-to-target are reported here.

Results

Six participants, who had already been trained on the task in another experiment, took part in this study (five male and one female). After re-training, each of them ran five blocks of 20 minutes each. Example trajectories produced by one of the participants (participant 1) are shown in Figure 2. It can be seen, that stabilization performance was better for trials in which platform rotations were presented (red trajectory), compared to trials in which the platform was off (blue trajectory). All participants showed this strong performance gain with platform rotations. Visual-only performance was best in the OF+H condition for participant 1.

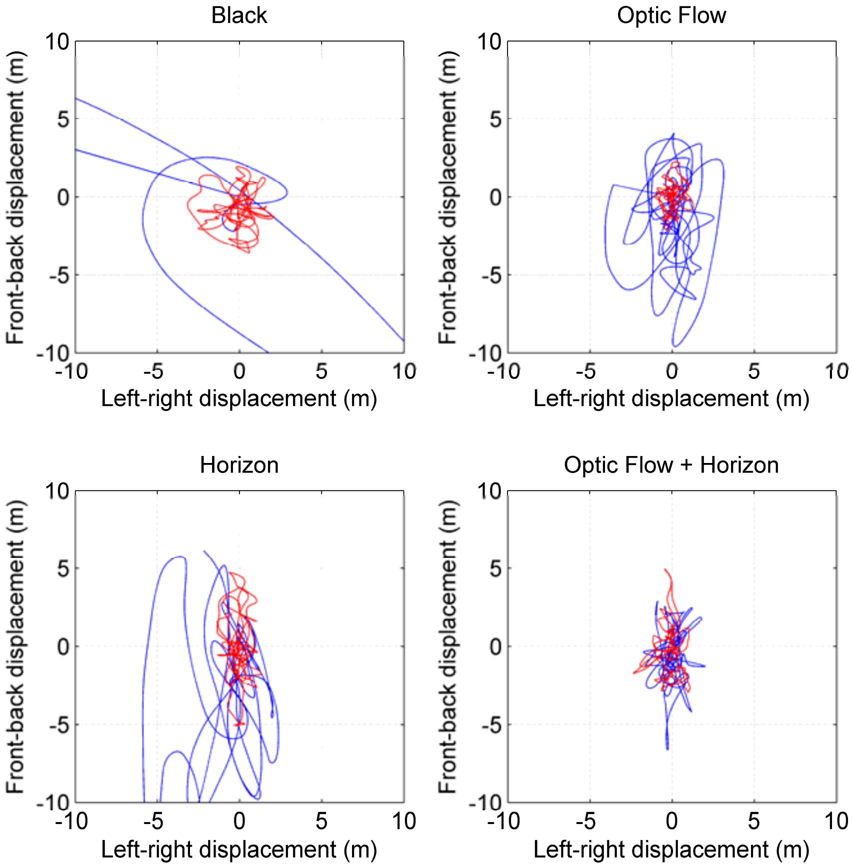

Figure 2: Representative trajectory examples of single trials of one participant (participant 1) in the different conditions. Blue: platform rotations off, red: platform rotations on. The target is at (0,0).

Figure 3 shows the resulting stabilization performance measures for all participants and all conditions. This data was analyzed with a four-way ANOVA for mean-distance-from-target; with direction (left / right vs. forward / backward), platform (rotation cueing on/off), horizon (on/off) and optical flow (on/off) as within-subject variables. Due to the skewed nature of the original data distributions, we used the (natural) logarithms of all measures for the ANOVAs, which made the distributions more Gaussian-like and the variances more similar for the different conditions (see Figure 3). Both are requirements for the ANOVA to work properly. The results of this ANOVA are shown in Table 1. Some of the effects are visualized separately in Figure 4.

In this analysis, all factors had significant main effects on distance-to-target. Also many interactions were significant. The B condition, without platform cueing, caused much larger response measures than all other conditions, which might be the reason so many interaction effects are significant. Participants stabilized significantly better in the left/right direction than in the forward/backward direction, as can be seen from mean distances in the two directions. Platform rotation improved stabilization performance significantly compared to visual-only stimulation for all measures (red vs. blue error bars in Figure 3). The improvement was particularly large in the B condition, and larger if the optical flow stimulus was shown visually compared to the H condition. In the B condition, visual feedback for control came only from the relative motion, position and sizes of the two spheres. If in this condition the platform motion was turned off, participants quickly lost control of the orientation of the helicopter. If, in contrast, platform rotation motion cueing was provided in the B condition, stabilization performance was almost as good as in the other visual conditions with the platform turned on.

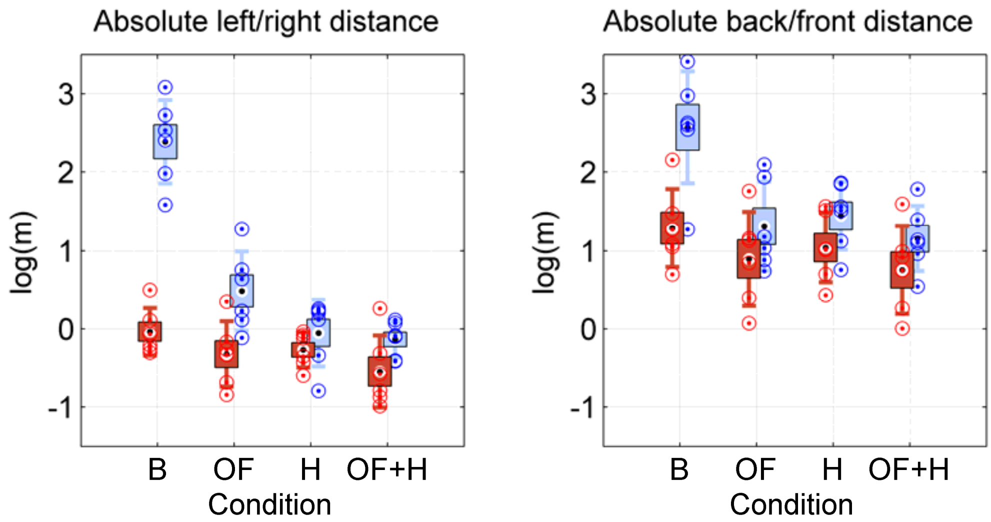

Figure 3: Results of all six participants. B: black background, OF: optical flow, H: horizon; Light blue: platform off, dark red: platform on. Measures are shown on logarithmic scale of helicopter target distance in meters. Error bars show standard deviations and boxes show standard error.

Mean	F=25.075	p=0.004**
D	F=251.609	p=0***
P	F=70.803	p=0***
D × P	F=7.124	p=0.044*
H	F=70.137	p=0***
D × H	F=19.144	p=0.007**
P × H	F=27.246	p=0.003**
D × P × H	F=55.93	p=0.001***
OF	F=47.582	p=0.001***
D × OF	F=1.678	p=0.252
P × OF	F=39.424	p=0.002**
D × P × OF	F=3.648	p=0.114
H × OF	F=21.786	p=0.005**
D × H × OF	F=10.165	p=0.024*
P × H × OF	F=14.962	p=0.012*
D × P × H × OF	F=18.144	p=0.008**

Table 1: Results of the four-way ANOVA with direction D (left/right vs. forward/backward), platform rotations P (on/off), horizon H (on/off), and optical flow OF (on/off) as within-subject variables.

The horizon improved stabilization performance significantly, and this effect was much stronger if the platform was off than if it was on (Figure 4, left panel). Presenting horizon and platform rotations together (P+H) still slightly improved stabilization performance for most measures, compared to the conditions in which either only horizon (H) or only platform rotations (P) were present. The horizon also improved stabilization performance more for the left/right direction than for the forward/backward direction. Horizon roll appears to be more useful to the observer than horizon pitch, possibly because we are more sensitive to horizon orientation (roll) than to absolute position in a large reference frame (display screen). Note however that in a real helicopter the artificial horizon cockpit instruments have reference markers both for orientation and vertical position of the horizon.

Optical flow also improved stabilization performance (Figure 4, right panel). The effect was again stronger if the platform was off than if it was on, probably because there was not much room for improvement in the platform-on condition.

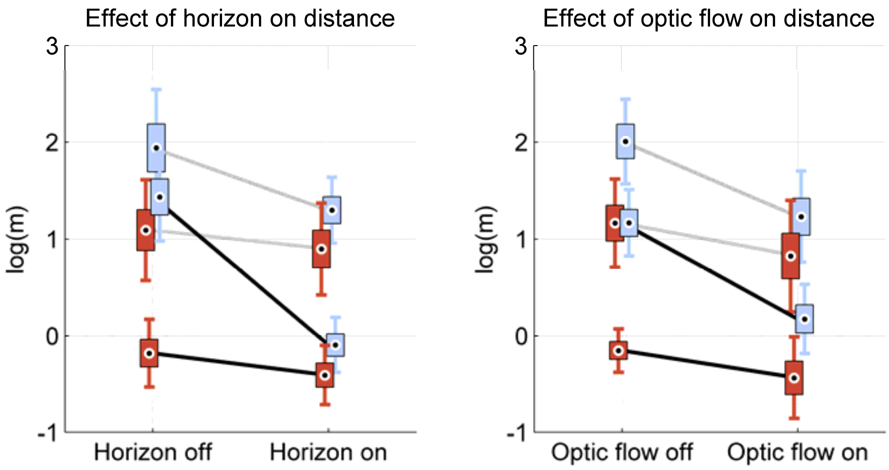

Figure 4: Visualisation of selected effects of different experimental parameters on the resulting distances from target. Dark red: platform on, light blue: platform off. Forward/backward distances are linked with grey lines, left/right distances are linked with black lines.

The overall best performance was reached if platform, horizon and optical flow cues were all available (P+H+OF). We found however large differences between participants for how much different cues helped stabilization. Example response patterns from two participants for the “distance-from-target” performance measure are shown in Figure 5. Most, but not all, participants produced response patterns similar to participant 1. Four participants stabilized better in the platform-on condition with optical flow information (P+OF) compared to the platform-on with horizon condition (P+H), but two subjects showed the opposite effect. Participant 3 was special in that he reached best stabilization for the horizon condition *without* platform cueing (see Figure 5 lower panel).

Is it possible to infer, from the data recorded, whether the three manipulated cues - platform rotations, horizon, and optical flow - are integrated by the participant in a statistically optimal, maximum-likelihood-like manner, as has been reported for multi-sensory integration with other modalities (Ernst and Banks, 2002; see Ernst and Bülthoff, 2004 for on overview of the MLE principle)?

Experimental tests of maximum likelihood estimation (MLE) in the integration of multiple cues are usually done by measuring psychometric functions in single-cue and combined-cue conditions (Ernst and Banks, 2002). Small conflicts between the cues are used in the combined-cue conditions, so that both variances and cue weights can be derived. From the response variances in the single-cue conditions, MLE predicts cue weights and variances in the combined-cue condition, which can then be compared to the measured responses.

In this experiment, we measured stabilization performance in single-cue conditions P, H and OF (platform-only, with neither horizon nor optical flow, horizon-only, and optical flow-only), and in combined-cue conditions P+H, P+OF, H+OF (platform-horizon, platform-optical flow, horizon-optical flow), and the triple-stimulation condition P+H+OF with platform, horizon and optical flow. However, since we did not introduce and manipulate small conflicts between the different cues, we do not know the cue weights in the combined conditions. Also, we cannot easily derive the perceptual variances (the reliabilities of the different sensory signals) from the performance measures (mean distance from target). Of course, the performance measures should be correlated with the perceptual variances, so that better perception leads to more accurate control, but how perceptual reliability and control accuracy are related exactly cannot be easily described, since the performance measures depend on the output of the complete pilot-helicopter system.

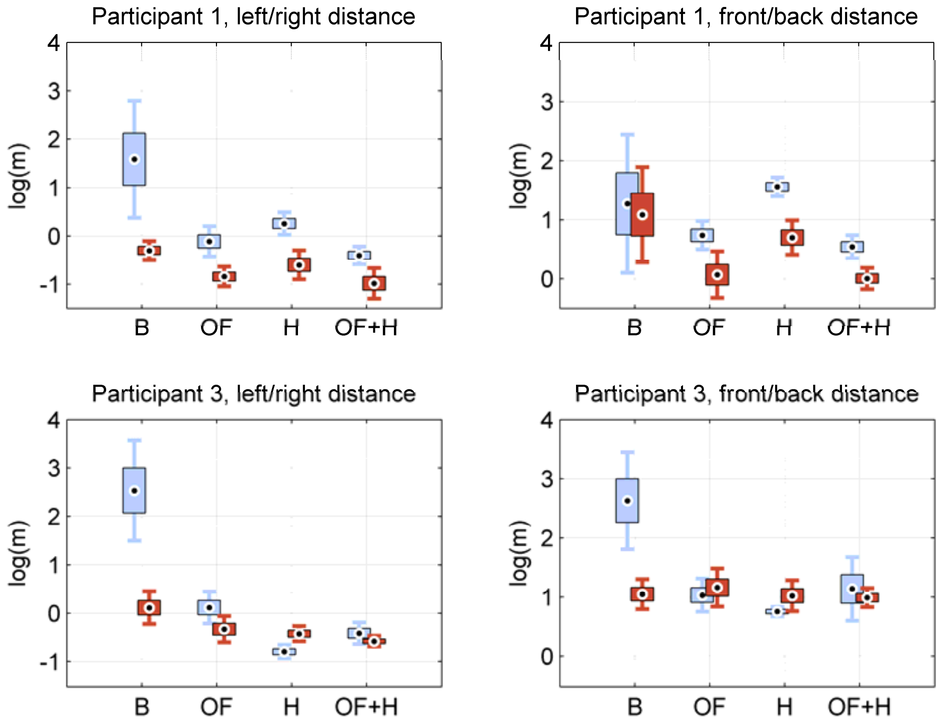

Figure 5: Individual differences between participants (here participants 1 and 3). B: black background, OF: optical flow, H: horizon; blue: platform off, red: platform on (P). Measures are shown on logarithmic scale of helicopter target distance in meters. Error bars show standard deviations and standard error.

We can still qualitatively describe the multi-sensory integration process. Stabilization performance and reliability of the sensory feedback should be correlated. Adding cues should in all conditions increase the quality of stabilization. Figure 5 shows that this is most often, but not always the case. For participant 3, for example, adding a platform cue (P), an optical flow cue (OF), or both (P+OF) to the horizon cue (H) *decreases* performance, which is evidence against MLE integration of these cues for this participant. The performance of participant 1, however, is in all cases qualitatively compatible with MLE integration: if a sensory cue is added, stabilization performance never decreases and gets better in most cases.

For participant 1, also the magnitude of the improvement seems to depend on how helpful a cue is alone. In the upper left panel of Figure 5, for example, platform alone (red error bar above 'B') is a better cue than the optical flow cue alone (blue error bar above 'OF'), and horizon alone is worse than both (blue error bar above 'H'). Adding platform movement to optical flow or horizon has a consistently stronger effect than adding a horizon cue to platform or optical flow. In the front/back direction (Figure 5, upper right panel), the platform cue is slightly worse than the optical flow cue, and indeed adding the platform cue to the horizon cue results in slightly worse performance than adding the optical flow cue to the horizon cue.

The single-subject data shows MLE-contradicting effects only between H and OF+H conditions. Participants 2, 5 and 6 showed the effect, but only when the platform was off; performance was worse when horizon and optical flow information were available (H+OF) than when only a horizon was shown (H). Participant 3 showed clear violation of MLE in all comparisons that involved the horizon condition without platform. Adding optical flow, platform motion, or both to the horizon cue degraded performance for this participant. Participants 1 and 4 did not show any MLE-contradicting performance. This suggests that some participants, in particular participant 3, might have used a different control strategy if the horizon was the only available cue, which allowed them to attain better performance than what the "standard control strategy" would have reached in this condition. Indeed, participant 3 reported that he tried to force himself to attend to horizon motion rather than target sphere motion in some conditions.

In all other cases, adding a cue does not reduce performance, but improves it significantly in many cases. This is consistent with true multi-sensory integration. Whether this integration follows the maximum likelihood principle remains to be shown.

Discussion

In this experiment, we investigated how different sensory cues are used by observers in a complex control task. We examined the influence of inertial motion cueing (platform rotations) and two visual cues (horizon and optical flow) on helicopter stabilization. We found that platform rotations helped stabilization significantly. Also, the horizon and the optical flow cues improved the helicopter stabilization significantly. Integration of the two visual cues (within modality) was not noticeably different from the integration of visual cues with inertial cue (across modalities).

The condition without platform movements, horizon, or optical flow shows that the information provided by the two spheres alone is not sufficient to stabilize the helicopter. As reported verbally, some subjects tried to infer the position and movement of a horizon from the relative position and size of the two spheres in this condition, but since this estimate is very unreliable, the performance was still a lot worse than in all other conditions. This shows that in our experiment, stabilization is indeed mostly based on the three cues which were manipulated, and not on the position or motion of the spheres.

The interaction of the three manipulated cues was qualitatively consistent with optimal cue integration (MLE) in most cases (adding a cue improves performance); however, in some cases we got MLE-contradicting results. The outlier was the horizon-only condition without platform motion (H), in which some participants performed better than when it was presented together with the optical flow stimulus, or, in participant 3, also together with the body rotations.

This may be explained by a control strategy switch, specifically for the horizon-only condition. In that condition, apparently a different strategy, which is optimized to use only the horizon cue for control, can lead to performance superior to the one which the standard strategy would attain in that condition. An interesting question is why the participants did not use this superior strategy any more when the horizon was presented together with platform rotations or the optical flow stimulus.

Different control strategies are an important issue for pilot training in a simulator. If trainees adopt a different control strategy in the simulator than in real flight, for example because no proper force cues are presented in the simulator, pilot training will not transfer well to real flight [Hall, 1978]. To make the trainee adopt a *natural control strategy*, inertial motion cueing are essential. The fact that in our experiment participants readily used inertial cues for stabilization supports this view.

Further work

The direction of the gravito-inertial force vector was not veridical in our simulation, since the platform rotations were determined only by the orientation of the simulated helicopter in space and not by the change of the gravito-inertial force vector induced by helicopter translations. However, since our participants were trained on this setup, they learned to use the platform rotation to stabilize the simulated helicopter. It would be interesting to test in further experiments what impact veridical force cueing versus rotation-only cueing has on the control performance.

With the Stewart motion platform, rotations can be simulated correctly, but the motion range was too limited for veridical translation motion cueing in this experiment. A recently acquired new setup, a simulator robot arm (*KUKA RoboCoaster*), will allow us to study helicopter hovering in realistic rotational and translational motion cueing situations.

In future studies we will also investigate the effect of other simulation parameters on stabilization performance, for example field of view, stereo projection, cue delays, and characteristics of both motion cueing and wash-out filters for translations and rotations.

Finally, a real helicopter has, of course, more degrees of freedom and more complex control dynamics than what was used in these experiments, in which we did not test experienced helicopter pilots. For a complete investigation of helicopter stabilization, the control task should incorporate all control devices (cyclic, collective and pedals) and all motion axes.

References

Chung, W. W., Bürki-Cohen, J., and Go, T. H. (2004). Task and vehicle dynamics based assessment of motion cueing requirements. In *Proceedings of the AIAA Modeling and Simulation Technologies Conference*, *August 16–19*, *Providence*, *Rhode Island*, paper number AIAA **2004-5154**, pages 1–10. American Institute of Aeronautics and Astronautics, Inc.

Ernst, M. O. and Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, **415**:429–433.

Ernst, M. O. and Bülthoff, H. H. (2004). Merging the senses into a robust percept. *Trends in Cognitive Sciences*, **8**(4):162–169.

Hall, J. (1978). Motion versus visual cues in piloted flight simulation. In *AGARD Conference Proceedings No. 249, Piloted aircraft environmental simulation techniques*, Meeting April 24-27, 1978, Brussels, pages 17/1–13. Controller HMSO, London.

Heintzman, R. J. (1996). Determination of force cueing requirements for tactical combat flight training devices. *ASC-TR-97*, **5001**:1–142.

Kemeny, A. and Panerai, F. (2003). Evaluating perception in driving simulation experiments. *Trends in Cognitive Sciences*, **7**(1):31–37.

Monen, J. and Brenner, E. (1994). Detecting changes in one's own velocity from the optical flow. *Perception*, **23**:681–690.

Ricard, G. L. and Parrish, R. V. (1984). Pilot differences and motion cueing effects on simulated helicopter hover. *Human Factors*, **26**(3):249–256.

Telban, R. J. and Cardullo, F. M. (2005). Motion cueing algorithm development: Human-centered linear and nonlinear approaches. *NASA/CR–2005*, **213747**:1–183.

Terzibas, C. (2004). Entwicklung eines Hubschraubersimulators zur Untersuchung der Lagestabilisierung mit visuellen und vestibulären Reizen. Master's thesis, Universität Tübingen, Abteilung Theoretische Astrophysik, Auf der Morgenstelle 10, 72076 Tübingen, Germany.

von der Heyde, M. (2001). A Distributed Virtual Reality System for Spatial Updating. PhD thesis, Universität Bielefeld; MPI Series in Biological Cybernetics, No. 02, Logos Verlag, Berlin, Germany.