P. Cisek, T. Drew & J.F. Kalaska (Eds.) Progress in Brain Research, Vol. 165 ISSN 0079-6123 Copyright © 2007 Elsevier B.V. All rights reserved

CHAPTER 11

Sensorimotor optimization in higher dimensions

Douglas Tweed*

Departments of Physiology and Medicine, University of Toronto, Centre for Vision Research, York University, Toronto, ON, Canada

Abstract: Most studies of neural control have looked at constrained tasks, with only a few degrees of freedom, but real sensorimotor systems are high dimensional — e.g. gaze-control systems that coordinate the head and two eyes have to work with 12 degrees of freedom in all. These extra degrees of freedom matter, because they bring with them new issues and questions, which make it hard to translate low-dimensional findings into theories of real neural control. Here I show that it is possible to predict high-dimensional behavior if we apply the optimization principles introduced by 19th-century neuroscientists like Helmholtz, Listing, and Wundt. Using three examples — the vestibulo-ocular reflex, saccadic eye movements, and depth vision — I show how simple optimization theories can predict complex, unexpected behaviors and reveal fundamental features of sensorimotor control, e.g. that neural circuits perform non-commutative algebra; that in rapid gaze shifts the eye controllers deliver commands with three degrees of freedom, not two; and that the eyes roll about their lines of sight in a way that may simplify stereopsis.

Keywords: sensorimotor; optimization; degrees of freedom; control; oculomotor; vision; vestibulo-ocular reflex (VOR); saccades; stereopsis; computational; behavioral

In sensorimotor control, as in science fiction, strange things happen in higher dimensions. For simplicity, most studies of neural control have focused on low-dimensional tasks, meaning ones with few degrees of freedom, such as purely horizontal movements of an eye or flexions of a single joint. But real sensorimotor systems are high dimensional. An arm, for instance, has 7 degrees of freedom — 3 for the shoulder, 2 for the elbow, and 2 for the wrist. A single eye rotates with 3 degrees of freedom — horizontal, vertical, and torsional. The head moves with 6 degrees of freedom, and so gaze-control systems that coordinate the head and two eyes have to work with 12 dimensions in all. In these cases and others, the hope has been that if we

*Corresponding author. Tel.: +1-416-978-2603;

can first manage to understand the system in a simple, constrained setting, we can then extrapolate to higher dimensions. But the extrapolation has often proved difficult. The key problem is that fundamentally new issues arise in higher dimensions, making it hard to generalize from lowdimensional findings. Here I will give examples of new concepts emerging in this way, but I will also argue that it is possible to predict high-dimensional behavior if we extrapolate in the right way, based on optimization.

Optimization theories of the brain go back to Helmholtz, Listing, Wundt, and other oculomotor pioneers of the 19th century (Helmholtz, 1867). To analyze a neural system by this approach, you first figure out what it is trying to do and state your guess precisely, in the form of a cost function. In an eye or arm movement, for instance, the cost

Fax: +1-416-978-4373; E-mail: douglas.tweed@utoronto.ca

might be the time to reach the target. To test the theory, you devise a controller that minimizes that cost function, and compare it to real human or animal behavior. Here I will show how simple optimization theories have predicted complex, high-dimensional behaviors that might otherwise have been inscrutable, or might never have been found. For instance, they predicted that certain random-dot stereograms are perceived as threedimensional (3-D) only when they are viewed looking up, not down; and that in some gaze shifts the eveballs twirl about their lines of sight at up to 200°/s for a fraction of a second and then unwind again. More importantly, these theories have revealed fundamental features of sensorimotor control, e.g. that in depth vision the brain searches for matching images in the two eyes over fixed rather than mobile patches of retina and that the two eyes are coordinated so as to shrink these patches; that eye control during gaze shifts is 3-D even though the line of sight has just 2 degrees of freedom; and that there is noncommutative computation in the sensorimotor circuitry of the

Noncommutativity in the brain

brain.

One example of a concept that emerges in higher dimensions is noncommutativity. A process is said to be noncommutative if order matters when things combine; if order makes no difference, the process is commutative (Hamilton, 1853). For instance ordinary multiplication of numbers is commutative because the order of factors is irrelevant, e.g. $5 \times 7 = 7 \times 5$. 1-D rotations are also commutative — turning first 10° right and then 20° right yields the same outcome as turning first 20° and then 10° . But 3-D rotations are noncommutative: the same two rotations applied in different orders can yield different overall rotations (Westheimer, 1957; Tweed and Vilis, 1987; Carpenter, 1988; McCarthy, 1990). This point is illustrated in Fig. 1, where the chess knight, starting from the same orientation, undergoes the same two rotations, 90° right and 90° down, in different orders and winds up in different final positions (in this figure the motions are defined in a knight-fixed frame, but



Fig. 1. Three-dimensional rotations do not commute. (Adapted with permission from Tweed (2003).)

rotations defined in a space-fixed frame do not commute either).

Why does noncommutativity of rotations matter for the brain? Because many brain processes have to deal with rotations, e.g. processes such as spatial perception, navigation, and the control of rotary joints. If they are to do their jobs even half decently, these systems have to represent and compute rotations, and for this they need noncommutative algebra (Westheimer, 1957; Tweed and Vilis, 1987, 1990; Crawford and Vilis, 1991; Minken et al., 1993; Hestenes, 1994a, b; Tweed et al., 1994; Tweed 1997a; Henriques et al., 1998; Smith and Crawford 1998), though for a long time this idea was controversial (van Opstal et al., 1991; Tweed et al., 1994, 1999; Straumann et al., 1995; Raphan, 1997, 1998; Quaia and Optican, 1998; Smith and Crawford, 1998; Schnabolk and Raphan, 1994).

Of all the neural systems that deal with rotations, maybe the simplest is the vestibulo-ocular reflex, or VOR. This reflex acts like a Steadicam for the eyeballs, stabilizing the retinal images when the head moves. Sense organs in the inner ear measure head velocity and send commands to the eye muscles, moving the eyes in the opposite direction when the head turns, so as to prevent the eyeballs rotating relative to space (Carpenter, 1988).

That the VOR needs noncommutative computation is illustrated in Fig. 2. Here a subject sits in a rotary capsule and looks out through a viewing screen at a space-fixed target, the black disk. In



Fig. 2. An optimal VOR must be noncommutative. In both the upper and lower series, the subject sits in a rotary capsule viewing a space-fixed object (the black disk), and in both series the target's initial location relative to the subject is the same: 30° directly to the left. Then the lights go out, and the subject tries to keep looking at the unseen disk while undergoing two rotations. In the two series the rotations are identical but are applied in opposite orders. Because of noncommutativity, the target's final locations relative to the subject are different. (Adapted with permission from Tweed (2003).)

both the upper and lower sequences, the subject starts out in the same position relative to the target: looking at it 30° directly to his left. In the upper series the subject turns first 10° counterclockwise (CCW) and then 60° left, so to keep his eyes on the target; he has to end up looking 30° right and 5° up. In the lower series, the subject undergoes the same two rotations in reverse order and winds up looking right and *down*. In other words the VOR must compute different final eyeposition commands when the subject goes through the same rotations in different orders (Tweed et al., 1999).

The motion of the eye in the head is plotted in Fig. 3A. If the subject turns first CCW then left, the eyes counter-rotate first clockwise (CW) then right, winding up at Position 3. When the subject rotates in the reverse order the eyes turn right and then CW, winding up at Position 3'. These trajectories are simulations of a theory of the 3-D VOR (Tweed, 1997b; Tweed et al., 1999) that was extrapolated from earlier, 1-D theories where the eye moved purely horizontally. There are many ways to extrapolate from low dimensions to high, but here the extrapolation preserved the optimization



Fig. 3. The real VOR is noncommutative. (A) Motion of the eye in the head, during the task from Fig. 2, as predicted by a theory of the VOR where the retinal image is perfectly stabilized. The system is noncommutative, yielding different final eye positions depending on the order of head rotations. (B) Performance, on the same task, of a VOR model in which all neural processing is commutative. The final eye positions do not depend on the sequence of body rotations. (C) A real human subject shows noncommutativity: final eye Positions 3 and 3' differ by about 10°, as predicted by the optimization theory. (Adapted with permission from Tweed et al. (1999).)

principle that the VOR acts to minimize retinalimage slip. So noncommutativity emerged as a necessary feature.

This optimal behavior was not predicted by previous 3-D models of the VOR, because they were extrapolated from 1-D theories in a different way, by preserving the 1-D principle that eyeposition commands are integrals of eye-velocity commands. But integration is commutative in the sense that the final value of an integral does not depend on the temporal order of its inputs, and therefore models based on this principle neglect noncommutativity and are incompatible with optimal image stabilization in 3-D. For instance, Fig. 3B shows one such commutative model (Raphan, 1997): regardless of the order of rotations, it brings the eye to the same final orientation relative to the head, and so relative to space the eye is incorrectly positioned, off the target.

On this same task, real human subjects closely matched the optimization theory, adopting different final eye positions that depended on the order of rotations. For the subject shown in Fig. 3C, the difference between Positions 3 and 3' (averaged over several trials) was 9.0° vertically, as compared to the optimal value of 10° . Averaged across all five subjects, the difference was 10.3° (range 7.4–12.6), and it was significant for each individual subject.

These findings established that there is noncommutativity in the VOR: the reflex correctly computes different final eye-position commands when put through identical rotations in different orders. And the broader point is that a simple optimization theory, based on minimizing retinal slip, predicted a fundamental feature of eye control that was absent in 1-D and was missed by other approaches. This theory (for details, see Tweed et al., 1994; Tweed, 1997a, b) has predicted many features of ocular control (e.g. Tweed, 1997b; Tweed et al., 1999; Misslisch and Tweed, 2000) and continues to find experimental support; e.g. Klier et al. (2006) recently showed that stimulating the abducens nerve rotates the eyeball around an axis that tilts as a function of eye position, in the pattern predicted by this theory.

Optimizing gaze control in three dimensions

In this section we focus on another high-dimensional concept, kinematic redundancy. We say a system is kinematically redundant if it has more degrees of freedom that it needs for some job. For example an arm has 7 degrees of freedom, but it needs only 6 to place the hand in any possible position (within a reasonable range near the shoulder joint). And in 3-D, an eye also has kinematic redundancy: it rotates with 3 degrees of freedom, but the line of sight has just 2, so there are infinitely many different eye positions that are all compatible with any one gaze direction (Fig. 4A shows three possible eye positions for straightahead gaze).



Fig. 4. (A) Listing's law. The three heads show three of the infinitely many different possible eye orientations for the same gaze direction, but Listing' law says the eye always chooses the orientation where torsion is zero. (B) A twisting-saccade task shows that a time-optimal saccadic system must control ocular torsion independently of horizontal and vertical eye position (see text). (C) Experimental data confirm that the eye controller for saccades has independent control of torsion. In the top row, a time-optimal controller performs the twisting-saccade task (number labels 1-4 correspond to Stages 1-4 in the illustration above). The controller does not move the eye directly (in the head frame) from its initial Position 1 to its final Position 3, but instead drives the eye out to 15° of torsion, which is the maximal allowable in this simulation, and then unwinds to its destination. This way, the eye stays near zero torsion between movements, and during gaze shifts it takes the fastest possible path to its final position in space. In the bottom row, a real human subject shows the same pattern. (Adapted with permission from Tweed et al. (1998).)

From this infinity of potential positions, the brain consistently chooses just one — the orientation in which the eye's torsional angle is zero, as shown by the central face in Fig. 4A. This zerotorsion rule is known as Listing's law, and it holds to within a degree or two during fixation and in the rapid gaze shifts known as saccades, as long as the head stays still (Helmholtz, 1867; Tweed and Vilis, 1990; Minken et al., 1993; Straumann et al., 1995).

Listing's law has been taken to mean that the eye controller for gaze shifts is 2-D, generating only horizontal and vertical commands. No torsional commands are needed, in this view, because torsion just stays at zero. But from an optimization viewpoint, there is reason to suspect that torsion is under separate, active neural control, and that the torsional command might be revealed by looking at saccades involving eye and head. More precisely, we need torsional control if the saccadic system is even roughly time-optimal, bringing the eye to its target position quickly (Tweed et al., 1998).

The crucial thought experiment is shown in Fig. 4B. The subject starts out in Position 1, with the head tilted 30° left-ear-down, looking at a target light which is 20° straight down relative to the head, and 1 m away. Then the target jumps sideways and the subject makes a twisting eye-head gaze shift to refixate it, passing through Position 2 in mid-saccade and ending up in Position 3. The interesting stage is 2: the eye is quicker than the head - it reorients more swiftly when an interesting object appears in the visual periphery (Roucoux et al., 1980; Laurutis and Robinson, 1986; Guitton and Volle, 1987; Tweed et al., 1995) — so a time-optimal controller would exploit that speed, flicking the eye to the target and locking on while the head catches up. The eye should move quickly to its final 3-D orientation in space, turning not just horizontally and vertically but also torsionally, so that midway through the head movement, the eye is twisted in its socket in the CW direction (from the subject's viewpoint), as shown in Stage 2. It should then hold still in space as the head completes its motion. If the target then jumps back to its original location, we should expect a similar return trajectory, this time with a strong CCW twist in mid-saccade, as in Stage 4.

This is the time-optimal strategy, and clearly it requires a torsional controller that can twist the eye rapidly in its orbit.

The top row of Fig. 4C shows a simulated timeoptimal controller (Tweed, 1997a; Tweed et al., 1998) performing this task. Again, the interesting thing is the eye's path relative to the head: it does not simply jump from Position 1 to 3, but takes a wide horizontal and torsional detour through Position 2, twisting through about 15° and then unwinding back to near-zero torsion (and similarly on the return trip through Position 4).

Faced with this same task, real human subjects behave like the time-optimal model. The subject in Fig. 4C showed the predicted torsional loops, ranging from 17° CW to 15° CCW. Across all four subjects, the torsional range averaged 29°, which far exceeds the 2-4° seen during head-fixed gaze shifts (Helmholtz, 1867; Tweed and Vilis, 1990; Minken et al., 1993: Straumann et al., 1995). And these huge torsional excursions really were visually evoked gaze shifts, not vestibular reactions to head motion, because they usually began 20-60 ms before the head started moving (Tweed et al., 1998). The eye spun about its line of sight at up to $200^{\circ}/s$ for 80 ms and then unwound to near-zero torsion over 200 ms (Tweed et al., 1998), so that Listing's law was in force at the end of the movement.

Obeying Listing's law brings advantages: it likely requires less muscle force to hold torsion near zero; and the eve, at the center of its torsional range, is optimally placed for the next gaze shift, which may go either CW or CCW (Hepp, 1990; Tweed, 1997c). So why does the eye break Listing's law during the gaze shift? As shown in Fig. 4B, the eye twists to anticipate the impending torsional motion of the head. This way, it reaches its final position in space while the head is still in midmovement. From then on, the eve holds a stable orientation in space, so the visual world remains stationary on the retina, blur is reduced, and visual analysis is simplified in other ways as well (Tweed et al., 1998). So there is more to torsional eye movement than simply holding at zero, and this study shows how an optimization model led to the discovery of an independent torsion-control system that helps drive saccades and underlies Listing's law.

The motor side of depth vision

For our final example we turn to stereopsis, where the visual system computes the 3-D locations of objects based on their images in the two eyes. The first step is to identify corresponding image features on the two retinas (Julesz, 1960). Figure 5 illustrates the problem: the eyes view a cloud of 21 dots, which cast 21 identical images on each retina. How does the brain know which dot on the right retina corresponds to which one on the left? We know the brain can find these matches, even when the images are thousands of identical dots, as in random-dot stereograms.

How does it manage? Geometry may help: as shown in Fig. 5, the optics of the situation restrict matching images to what are called epipolar lines (Ogle, 1950; Rogers and Bradshaw, 1996). So if it could locate these epipolar lines, the brain could simplify its quest for matching images: it would not have to search the entire retina for a match, but could carry out a 1-D search along the epipolar line, like looking for lost hikers along a single trail rather than combing the whole forest.

Most theories of stereopsis have proposed that the brain searches along epipolar lines. But these theories were worked out assuming stationary eyes. When we consider that the eyes move, the theories hits a snag: the epipolar lines migrate on the retinas (Garding et al., 1995; Stevenson and Schor, 1997; Tweed, 1997c). As shown in Fig. 5, the same point on one retina corresponds to different epipolar lines on the other retina, depending on the configuration of the eyes (in the figure, the eyes rotate about their own lines of sight, but other sorts of rotations also shift the epipolar lines). Again this is a problem of dimensionality: earlier theories neglected all three dimensions of eve rotation (or all six, counting both eyes), and new issues arise when we consider these extra degrees of freedom.

In light of this complication, there are two ways the brain might find matching images in mobile eyes (Schreiber et al., 2001). The options are illustrated in Fig. 6. Given an image falling on some locus in one retina, the brain could use eye-position information to locate its epipolar line on the other retina. The other option is to forget about finding epipolar lines and instead search a 2-D patch of retina large enough to encompass all possible locations of the epipolar line in any likely eye configuration. This way, the stereoptic system would not have to monitor eye position, but it would lose the advantage of a 1-D search. So the question is: Does the brain search for matches along epipolar lines, or over retina-fixed 2-D zones?

We can answer this question using rotated stereograms, as shown in Fig. 7. We construct a random-dot stereogram in the usual way and then rotate the disks. If the disk viewed by the right eye is turned CCW, and the other CW, as in Fig. 7, the stereogram is incyclorotated. If the rotations are reversed, it is excyclorotated. Why are these stereograms useful? We know that when people converge their eyes and look up - when they look at something close to their forehead — they excycloverge, rotating the upper poles of both eyeballs outward (Allen, 1954; Mok et al., 1992; Van Rijn and Van den Berg, 1993; Minken and Van Gisbergen, 1994; Tweed, 1997c; Kapoula et al., 1999; Steffen et al., 2000; Schreiber et al., 2001); and when they converge and look down, they incycloverge. So the prediction is this: if our stereo search zones are retina-fixed, we should be better able to see excyclorotated stereograms on upgaze, incyclorotated on downgaze.

For example, suppose you view a stereogram that is excyclorotated by 5°. When your eyes are also excycloverged 5°, just like the stereogram disks, then the optical correspondence should be normal, just as if you were viewing a normal, nonrotated stereogram with zero cyclovergence, so the image should be easy to see. But when your eyes are cycloverged 0° then corresponding dots in the two excyclorotated disks will project onto odd locations on your two retinas, making the stereo-image hard to see. And again, this is the prediction *if* stereo search zones are retina-fixed; if instead the search zones move with the epipolar lines then eye position should not affect visibility.

Figure 8A confirms that the search zones move with the retinas, not with the epipolar lines. It plots the probability of stereoptic vision versus cyclorotation of the stereogram at three eye elevations for a typical subject. For instance when a



Fig. 5. Any animal with stereopsis must solve the stereo-matching problem, deducing which images on the right retina correspond to which ones on the left. The task can be simplified using epipolar lines, but when the eyes move, the epipolar lines migrate on the retinas.



Fig. 6. There are two ways the visual system could look for matching images in mobile eyes. The nine small circles are nine images projected onto the right retina, one foveal and the others 15° eccentric (the large gray disk is the region within 22.5° of the fovea). Corresponding images on the left retina must lie somewhere on the line segments, which are pieces of epipolar lines, but the lines are in different places depending on the positions of the eyes. White patches cover the ranges of motion of the epipolar segments when the eyes move over a realistic range. (Adapted with permission from Schreiber et al. (2001).)



Fig. 7. Top: cyclorotated stereograms are visible only in certain eye positions. Cross-fuse the disks from 30 cm away and depress your gaze as far as possible, holding the paper orthogonal to the plane of your sight lines. You should see a depth image (a triangle) in this position, but not when you do the same on upward gaze. If the image never disappears, your search zones are too large for this stereogram; try the examples in Schreiber et al. (2001).

stereogram is incyclorotated by 6° then it is perceived with probability 1 when the eyes are directed 30° down (dotted line) and with probability 0 when the eyes are 30° up (thin gray line) — i.e. this stereogram is visible on downgaze but not on upgaze.

Figure 8B plots perception thresholds — the angles of cyclorotation at which stereograms were



Fig. 8. Stereopsis depends on gaze elevation. (A) For this typical subject, the small black dots show the probability of seeing the stereoimage as a function of the cyclorotation of the stereogram when the elevation of the eyes is 0°. Curves plot performance at three gaze elevations. Larger circles are perceptual thresholds — cyclorotation angles at which stereoimages were perceived with probability 0.5. (B) Stereoptic thresholds depend on gaze elevation and cyclovergence. (C) For all five subjects, these thresholds varied significantly with gaze elevation, shifting toward incyclorotation on downgaze (leftmost symbol in each string of three) and toward excyclorotation on upgaze (rightmost symbol). The average slope is 1.06, very close to the slope of 1 predicted if stereo search zones are retina-fixed. (Adapted with permission from Schreiber et al. (2001).)

perceived 50% of the time — versus ocular cyclovergence, for the same subject as in Fig. 8A, for the same three eye elevations. And as in Fig. 8A, white symbols indicate data collected on downgaze, grav means upgaze, and black means level. So for instance, of the nine plotted points in this panel, the small white one at the lower left corner means that this subject, when looking 30° down (and converging 30°) had about 4.5° of incyclovergence, and an incvclorotation threshold of 10° ; i.e. the subject perceived stereoimages with probability 0.5 when the stereogram was incyclorotated by that amount. Similarly, the leftmost point on the upper line of the plot means this subject's excyclorotation threshold under these conditions was about 4° . The large dot halfway between the in- and excyclorotation thresholds is the average of the two thresholds.

Figure 8C shows thresholds and midpoints for all five subjects. If stereo search zones were perfectly fixed on the retinas then the line of midpoints would have a slope of 1 (because the cyclorotation thresholds would rotate exactly as far as the eyes), as indicated by the dashed line. The actual slope, averaged over all subjects, was 1.06, and not significantly different from 1. So the data indicate that stereo search zones are retinafixed.

This finding suggests that eye control plays a central part in stereopsis. An optimized controller could coordinate the eyes so as to minimize the motion of the epipolar lines, allowing stereopsis to get by with the smallest possible search zones. The normal pattern of eye control when viewing distant objects is Listing's law (Helmholtz, 1867; Carpenter, 1988), but on near gaze the law is broken (Allen, 1954; Mok et al., 1992; Van Rijn and Van den Berg, 1993; Minken and Van Gisbergen, 1994: Tweed, 1997c: Kapoula et al., 1999: Steffen et al., 2000; Schreiber et al., 2001) and it can be shown that the deviations from Listing's law shrink the required search zones (Schreiber et al., 2001). The zones are not precisely minimized the eye's deviations from Listing's law are not large enough for that — but the reason may be that the controller is balancing the benefits of small zones against the advantages of Listing's law (Carpenter, 1988; Hepp, 1990; Tweed, 1997c).

Conclusion

Most studies of neural control have focused on low-dimensional tasks, with few degrees of freedom, but real sensorimotor systems are high dimensional. I have argued that new issues arise in higher dimensions, but I have also shown, with three examples, that it is nevertheless possible to extrapolate usefully from low-dimensional findings if we do it based on optimization principles. Each of these three examples suggests further questions and generalizations. I have shown that there is noncommutative computation in the circuitry of the VOR, and by similar reasoning, one would expect noncommutativity also in many other brain systems that deal with rotations, such as those for head and limb control, auditory and visual localization, space constancy, and mental rotation of objects (Hestenes, 1994b; Tweed, 1997a; Henriques et al., 1998). Optimization ideas clarified the implementation of Listing's and Donders' laws, and there are doubtless, waiting to be discovered, many higher-dimensional analogs of these laws, constraining the motions of the eyes, head, and limbs in complex tasks. An optimization model clarified the relation between stereopsis and eve control in six dimensions, and this model, too, leads to further predictions, for instance that the layout of stereo search zones on the retinas should resemble the optimal pattern in Fig. 6 (Schreiber et al., 2001). And optimization methods have been applied with great success to many other sensorimotor problems besides my specific examples. From the pioneering work of Helmholtz to the present day, probably no other approach has been so successful at illuminating the complex control systems of the brain.

Acknowledgments

I thank my co-authors on the original studies, J. D. Crawford, M. Fetter, T. Haslwanter, V. Happe, and K. Schreiber. For comments and technical help I thank M. Abdelghani, K. Beykirch, D. Broussard, L. Chinta Venkataswararao, J. Dichgans, S. Ferber, K. Fortney, D. Goche Montes, P. Hallett, C. Hawkins, D. Henriques, I. Howard, E. Klier, H. Misslisch, P. Nguyen, M. Niemeier, J. Sharpe, T. Vilis, H. Wang, A. Wong, and J. Zacher. This work was supported by the Canadian Institutes for Health Research and the Deutsche Forschungsgemeinschaft.

References

- Allen, M.J. (1954) The dependence of cyclophoria on convergence, elevation and the system of axes. Am. J. Optom. Arch. Am. Acad. Optom., 31: 297–307.
- Carpenter, R.H.S. (1988) Movements of the Eyes. Pion, London.
- Crawford, J.D. and Vilis, T. (1991) Axes of eye rotation and Listing's law during rotations of the head. J. Neurophysiol., 65: 407–423.
- Garding, J., Porrill, J., Mayhew, J.E. and Frisby, J.P. (1995) Stereopsis, vertical disparity and relief transformations. Vision Res., 35: 703–722.
- Guitton, D. and Volle, M. (1987) Gaze control in humans: eyehead coordination during orienting movements to targets within and beyond the oculomotor range. J. Neurophysiol., 58: 427–459.
- Hamilton, W. (1853) Lectures on Quaternions: Containing a Systematic Statement of a New Mathematical Method. Hodges and Smith, Dublin.
- Helmholtz, H.von. (1867) Handbuch der Physiologischen Optik. L. Voss, Leipzig.
- Henriques, D.Y., Klier, E.M., Smith, M.A., Lowy, D. and Crawford, J.D. (1998) Gaze-centered remapping of remembered visual space in an open-loop pointing task. J. Neurosci., 18: 1583–1594.
- Hepp, K. (1990) On Listing's law. Commun. Math. Phys. V, 132: 285–292.
- Hestenes, D. (1994a) Invariant body kinematics: I. Saccadic and compensatory eye movements. Neural Netw., 7: 65–77.
- Hestenes, D. (1994b) Invariant body kinematics: II. Reaching and neurogeometry. Neural Netw., 7: 79–88.
- Julesz, B. (1960) Binocular depth perception of computergenerated patterns. BSTJ, 39: 1125–1162.
- Kapoula, Z., Bernotas, M. and Haslwanter, T. (1999) Listing's plane rotation with convergence: role of disparity, accommodation, and depth perception. Exp. Brain Res., 126: 175–186.
- Klier, E.M., Meng, H. and Angelaki, D.E. (2006) Three-dimensional kinematics at the level of the oculomotor plant. J. Neurosci., 26: 2732–2737.
- Laurutis, V.P. and Robinson, D.A. (1986) The vestibulo-ocular reflex during human saccadic eye movements. J. Physiol., 373: 209–233.
- McCarthy, J.M. (1990) An Introduction to Theoretical Kinematics. MIT Press, Cambridge, MA.
- Minken, A.W. and Van Gisbergen, J.A. (1994) A three-dimensional analysis of vergence movements at various levels of elevation. Exp. Brain Res., 101: 331–345.

- Minken, A.W., Van Opstal, A.J. and Van Gisbergen, J.A. (1993) Three-dimensional analysis of strongly curved saccades elicited by double-step stimuli. Exp. Brain Res., 93: 521–533.
- Misslisch, H. and Tweed, D. (2000) Torsional dynamics and cross-coupling in the human vestibulo-ocular reflex during active head rotation. J. Vestib. Res., 10: 119–125.
- Mok, D., Crawford, J.D. and Vilis, T. (1992) Rotation of Listing's plane during vergence. Vision Res., 32: 2055–2064.
- Ogle, K.N. (1950) Researches in Binocular Vision. W.B. Saunders, Philadelphia, PA.
- van Opstal, A.J., Hepp, K., Hess, B.J., Straumann, D. and Henn, V. (1991) Two- rather than three-dimensional representation of saccades in monkey superior colliculus. Science, 252: 1313–1315.
- Quaia, C. and Optican, L.M. (1998) Commutative saccadic generator is sufficient to control a 3-D ocular plant with pulleys. J. Neurophysiol., 79: 3197–3215.
- Raphan, T. (1997) Modelling control of eye orientation in three dimensions. In: Fetter M., Haslwanter T., Misslisch H. and Tweed D. (Eds.), Three-dimensional Kinematics of Eye, Head and Limb Movements. Harwood Academic Publishers, Amsterdam, The Netherlands, pp. 359–376.
- Raphan, T. (1998) Modeling control of eye orientation in three dimensions. I. Role of muscle pulleys in determining saccadic trajectory. J. Neurophysiol., 79: 2653–2667.
- Rogers, B.J. and Bradshaw, M.F. (1996) Does the visual system use the epipolar constraint for matching binocular images? Invest Ophthalmol. Vis. Sci., 37(Suppl.): 31–25.
- Roucoux, A., Crommelinck, M., Guerit, J.M. and Meulders, M. (1980) In: Fuchs A.F. and Becker W.E. (Eds.), Progress in Oculomotor Research. Elsevier, Amsterdam, pp. 309–315.
- Schnabolk, C. and Raphan, T. (1994) Modeling three-dimensional velocity-to-position transformation in oculomotor control. J. Neurophysiol., 71: 623–638.
- Schreiber, K., Crawford, J.D., Fetter, M. and Tweed, D. (2001) The motor side of depth vision. Nature, 410: 819–822.
- Smith, M.A. and Crawford, J.D. (1998) Neural control of rotational kinematics within realistic vestibuloocular coordinate systems. J. Neurophysiol., 80: 2295–2315.
- Steffen, H., Walker, M.F. and Zee, D.S. (2000) Rotation of Listing's plane with convergence: independence from eye position. Invest. Ophthalmol. Vis. Sci., 41: 715–721.
- Stevenson, S.B. and Schor, C.M. (1997) Human stereo matching is not restricted to epipolar lines. Vision Res., 37: 2717–2723.
- Straumann, D., Zee, D.S., Solomon, D., Lasker, A.G. and Roberts, D.C. (1995) Transient torsion during and after saccades. Vision Res., 35: 3321–3334.
- Tweed, D. (1997a) Three-dimensional model of the human eyehead saccadic system. J. Neurophysiol., 77: 654–666.
- Tweed, D. (1997b) Velocity-to-position transformations in the VOR and the saccadic system. In: Fetter M., Haslwanter T., Misslisch H. and Tweed D. (Eds.), Three-dimensional Kinematics of Eye, Head and Limb movements. Harwood Academic Publishers, Amsterdam, The Netherlands, pp. 375–386.

- Tweed, D. (1997c) Visual-motor optimization in binocular control. Vision Res., 37: 1939–1951.
- Tweed, D. (2003) Microcosms of the Brain. Oxford University Press, Oxford, UK.
- Tweed, D. and Vilis, T. (1987) Implications of rotational kinematics for the oculomotor system in three dimensions. J. Neurophysiol., 58: 832–849.
- Tweed, D. and Vilis, T. (1990) Geometric relations of eye position and velocity vectors during saccades. Vision Res., 30: 111–127.
- Tweed, D., Glenn, B. and Vilis, T. (1995) Eye-head coordination during large gaze shifts. J. Neurophysiol., 73: 766–779.

- Tweed, D., Haslwanter, T. and Fetter, M. (1998) Optimizing gaze control in three dimensions. Science, 281: 1363–1366.
- Tweed, D., Misslisch, H. and Fetter, M. (1994) Testing models of the oculomotor velocity-to-position transformation. J. Neurophysiol., 72: 1425–1429.
- Tweed, D.B., Haslwanter, T.P., Happe, V. and Fetter, M. (1999) Non-commutativity in the brain. Nature, 399: 261–263.
- Van Rijn, L.J. and Van den Berg, A.V. (1993) Binocular eye orientation during fixations: Listing's law extended to include eye vergence. Vision Res., 33: 691–708.
- Westheimer, G. (1957) Kinematics of the eye. J. Opt. Soc. Am., 47: 967–974.