Ownership and Copyright

Virtual Reality (2002) 6:86–95

Simulating Self-Motion II: A Virtual Reality Tricycle

R. S. Allison[1], L. R. Harris[2,3], A. R. Hogue[1], U. T. Jasiobedzka[1], H. L. Jenkin[2], M. R. Jenkin[1], P. Jaekl[2], J. R. Laurence, G. Pentile[1], F. Redlick[3], J. Zacher[4], D. Zikovitz[3]

Centre for Vision Research, [1]Department of Computer Science, [2]Department of Psychology and [3]Department of Biology, York University; [4]Centre for Research in Earth and Space Technology, Toronto, Ontario, Canada

Abstract: When simulating self-motion, virtual reality designers ignore non-visual cues at their peril. But providing non-visual cues presents significant challenges. One approach is to accompany visual displays with corresponding real physical motion to stimulate the non-visual, motion-detecting sensory systems in a natural way. However, allowing real movement requires real space. Technologies such as Head Mounted Displays (HMDs) and CAVE™s can be used to provide large immersive visual displays within small physical spaces. It is difficult, however, to provide virtual environments that are as large physically as they are visually. A fundamental problem is that tracking technologies that work well in a small, enclosed environment do not function well over longer distances. Here we describe Trike – a 'rideable' computer system that can be used to present large virtual spaces both visually and physically, and thus provide appropriately matched stimulation to both visual and non-visual sensory systems.

Keyords: Self-motion simulation; Visual and vestibular egomotion cues

Introduction

Generating the appropriate visual view during Virtual Reality (VR) simulation requires knowing the viewer's instantaneous location. Although tracking technologies have seen considerable advances, fundamental problems remain, especially in applications where the user moves over a range of more than a few feet. With few exceptions, existing tracking systems are designed around technologies that limit the user to moving with a small physical area. For example, mechanical tracking systems such as Fake Space's BOOM and Puppetworks' trackers, physically connect the wearer to a fixed position, and the user is literally tethered to a point in space. Wireless electromagnetic systems such as Ascension Technology's Flock of Birds tracking system and Polhemus' Fastrak remove the mechanical link, but have a limited operational range. For example, the Flock of Birds has an operational range of ±4 feet with the standard transmitter and ±10 ft with the extended range unit. Vision-based trackers such as Northern Digital's POLARIS have the potential for a longer operating range, but their accuracy degrades with distance due to the use of triangulation to measure position. Acoustic tracking systems are limited by the decay in the audio signal as a function of distance and the sound absorption properties of air. Existing VR systems may emulate a large visual space, but the bulk of existing tracking technologies limit the user to moving in a relatively small physical world.

Various alternatives have been proposed to extend the range over which an immersive VR display can operate. Perhaps the most extreme version of these has been the development of wearable computers (e.g. [1,2]). These systems provide a computer with an interface that is 'worn' by the user and that augments their normal senses as they move. Although most existing wearable systems are designed to provide two-dimensional visual overlays on the wearer's normal visual field, some wearable computer systems provide more immersive visual displays. For example, Lewis et al. [3] describe a wearable computer system that in part mimics a pair of binoculars, but with identical views in both eyes. The wearer looks through the 'binoculars' and is presented with a computer-enhanced version of the environment. Feiner et al. [4] describe a backpack-based system where the user views the world through a HMD in which synthetic elements are superimposed over the normal view. The lack of a static base station and some mechanism for measuring displacement relative to an earth-fixed point make estimating absolute position very difficult in Feiner et al.'s system: only relative orientation information is available for the generation of their visual display.

Technologies which track over shorter ranges can be patchworked together to form a mosaic of trackable regions. For example, a network of Flock of Bird transmitters can be used to provide coverage over a large space, but at significant financial cost, and the system still requires a connection from the receiver to the base station. One successful application of a patchwork approach is Northern Digital's OPTOTRAK system. This system senses active infrared markers placed on the subject with a number of camera trackers to provide long-range tracking. Even with an unlimited budget, this type of tracking technology only provides tracking over a physical space the size of the environment in which the tracker is installed.

Although the Global Positioning System (GPS) – especially in its differential mode – can theoretically be used to provide tracking information over larger ranges, it is not without its problems. For example, it cannot be used in environments where no clear line of sight exists to the satellites, such as indoors, near mountain ranges, in forests, in built-up urban areas. It also has a relatively slow update rate (see [5] for a description of GPS and its limitations). That being said, in appropriate outdoor environments, GPS can be used to correct drift in other sensors.

2. Simulating Medium- and Long-range Motion

To provide a virtual environment over a large physical space and to permit the experiments of the type described in the previous paper [6], we have developed a rideable computer system (Trike), based on a commercially available tricycle. Trike can be used to provide a full range of naturalistic, motion-related cues while users work within a large scale virtual environment. Standard VR technology is used to provide an immersive visual display relative to Trike itself. The effort of pedalling Trike around provides cognitive and proprioceptive cues consistent with normal navigation and the physical motion of Trike generates the normal spectrum of non-visual cues. Trike is instrumented so that its motion relative to some initial base frame can be computed and Trike can be tracked using a ground-based visual system to deal with drift from the vehicle-based tracking system.

In addition to providing a natural mechanism for navigation in a large-scale virtual world, Trike has been used as a powerful research tool for investigating the relative importance of different visual and non-visual cues to the perception of self-motion in virtual environments (see [6]).

The complexity of the perceptual cues that an operator experiences when they move, presents a challenge to those attempting to simulate long-distance motion without actually moving the operator through the corresponding distance. A number of technologies have been developed and deployed to try to simulate these movements by stimulating one or more non-visual sensory systems while ignoring the others. Here we review some of the technologies employed to simulate medium- and long-range motions. We define short-range motions as those that occur within a radius of about three feet and can be tracked by a variety of conventional technologies. Medium-range motions are those that are possible within a controlled environment up to the size of a hall, ice rink or baseball arena. Long-range motions are those that are not usually restricted to a defined area such as the movements of cars, aircraft, spacecraft or other vehicles.

Static Automotive/Aircraft Simulators

Static (non-moving) simulators have found a wide range of applications in everything from aircraft pilot training to training student drivers. The devices replicate the appearance and functionality of only that part of the device needed for the training to be performed. The device is static, but the visual displays and instruments are driven by a computer-based simulation of the vehicle's systems to provide realistic responses to operator inputs. The sophistication of the simulation and displays runs the gamut from accurate, large aircraft simulations, such as those manufactured by Atlantis Aerospace, to a simple PC with a steering wheel and brake and accelerator pedals.

The sensation of motion in these devices is generated primarily by visual cues, and from the expectation of motion from the state of the instruments and controls.

VR Bicycles

A number of laboratories have adapted stationary bicycles (typically a road bicycle mounted on a trainer stand) to simulate long-range motion. The bicycle is mounted on a stand, and the steering angle and drive wheel are instrumented. Coupled with an appropriate kinematic model, this permits the simulation of motion through large-scale space (see [7] for one approach to constructing a VR Bicycle). It is even possible to provide computer control of the resistance that the operator experiences when cycling. For example the Computrainer system by RacerMate Inc. incorporates a computer-controlled device that opposes motion of the pedals. This can then realistically simulate the resistance associated with pedalling uphill.

A VR bicycle simulates many of the non-visual cues associated with long-range motion. However, as the bicycle does not physically move relative to the ground, neither the rotation-sensing nor translation-sensing parts of the vestibular system are activated by the motion of the bicycle itself. It is also important to note that VR bicycles are usually fixed in place, and do not lean like real bicycles. Tilting a static bicycle to the side as it goes round a corner is possible; a low-slung version has been developed by Heinrich Bülthoff and his group (see [8]). Tilting the bicycle makes the steering feel more natural, but does not by any means simulate all the physical forces (e.g. centripetal forces and linear and angular accelerations) that a cyclist normally experiences when turning.

VR Hang-glider

VR hang-gliders attempt to provide a realistic way to control a hang-glider. The Dreamglider system marketed by Dreamality Technologies is typical. The operator is strapped into a harness that mimics the harness worn normally by a hang-glider operator, and either views a fixed monitor or wears a head-mounted display. Operators shift their weight in the harness to control the flight of the device, similar to when flying a real hang-glider. NTT's Virtual Hang-glider adds simulated wind to the experience. Although these devices may be great toys, they cannot hope to present the physical motions associated with real hang-gliders.

VR Treadmills

Standard exercise treadmills can be instrumented to permit an operator to walk in one dimension [9]. The treadmill can either be electrically controlled to centre the operator within its operational range or a non-motorised treadmill can be used with the subject walking to provide treadmill motion. This second option is not as desirable as it requires rather unnatural motion on behalf of the operator. An obvious restriction is that only one-dimensional motion is possible. To address this issue, two-dimensional treadmills, such as the Omni-Directional Treadmill (ODT) manufactured by Virtual Space Devices and the Torus Treadmill at the University of Tskuba, have been developed (see [10,9]). An individual walking on the ODT can move in any direction, and the device maintains the user's position near the device's centre. The ODT is essentially two continuous belts, one for the X direction, one for the Y. The X belt is supported by the Y belt and is mechanically transparent so that the Y motion is conducted up through the X belt. Each belt is made up of series of rollers woven together. Each belt is motorised and motion of the operator relative to the treadmill is tracked. The belts are rotated so as to maintain the operator within the physical extent of the device.

Although the ODT can provide many non-visual cues to self-motion, the limited physical extent of the device requires that the device has to keep moving the operator back to the centre of its operating range. For relatively short-duration motions, or motions followed by stationary periods, these recovery motions can be of short duration use and very low magnitude accelerations in an attempt to keep them below the operator's perceptual threshold. For continuous motion, however, these re-centring motions are likely to confound the normal non-visual sensations experienced by the operator.

Physical Motion Bases

Physical motion bases have been used for aircraft pilot training since the 1920s. One of the earliest successful physical simulators was the Link trainer. The Link trainer originally provided the operator with a control

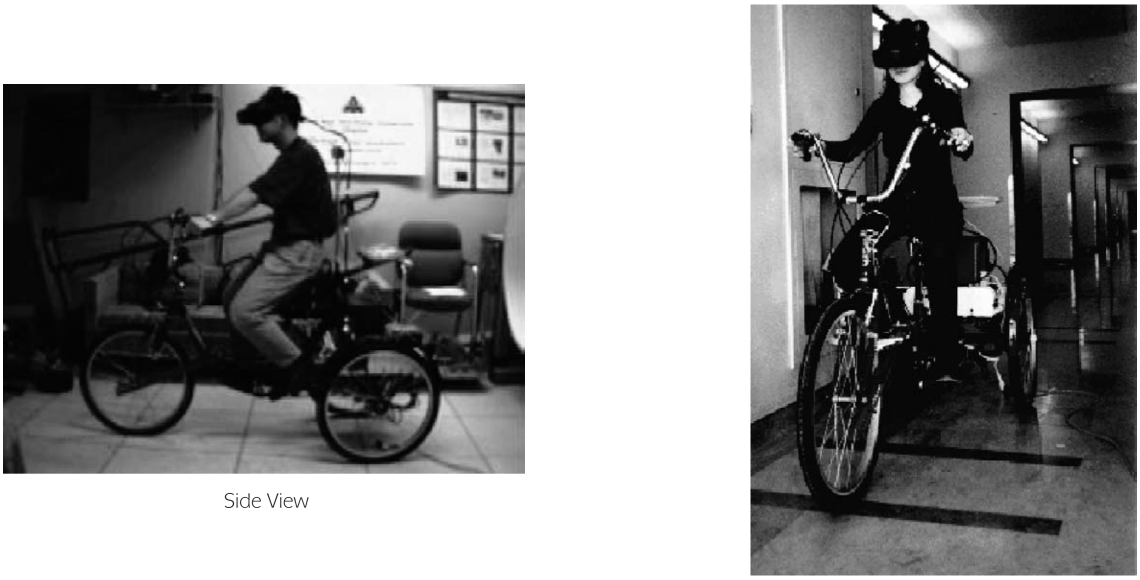

Side View

Front View

Fig. 1. The Virtual Reality Tricycle (Trike). Trike is based on an adult-sized tricycle that has been instrumented to measure the rotation of one of the rear drive shaft and the steering angle. A Polhemus Tracker provides a measurement of the rider's head relative to Trike. Visual display is presented via a V8 binocular HMD. Display is generated onboard via a SGI O2. Power, and optionally network access, is provided via a tether.

column and control wheel, two-foot pedals, and various flight and navigation instruments. The device sat on four pneumatic bellows and used air pressure to simulate flight motion such as climbing, diving and banking. The device responded to commands from the control column, and through a series of linkages, used the air to adjust the pose of the simulated cockpit.

Modern aircraft simulators such as those manufactured by CAE or TDI, utilise six degree of freedom motion bases coupled with complex visual displays and simulation software. The physical motion base typically has a very restricted motion envelope, which means that large physical motions must be simulated using only small-range physical motions.

For translation, motion bases can provide the initial linear acceleration and then (as for the 2D treadmill) attempt to return users to the centre of the operating range at subliminal accelerations (a process known as 'washout') during periods of constant velocity or no movement. They can also provide tilt to mimic an aircraft banking, or to simulate linear accelerations by directing an appropriate component of gravity (controlled by the amount of tilting) along the direction that a linear acceleration is expected.

Summary

These ingenious systems provide some of the natural *outgoing* cues, such as those generated by pedalling and steering (bicycle), walking (treadmill) or body orientation changes (hang-gliding). However, with the exception of motion bases, they do not provide or simulate the non-visual sensory cues that would naturally be associated with the motion of their mobile counterparts.

3. Building a Rideable Computer

As our goal is to allow physical motion over a large area, it is important that the device be tethered no more than necessary to a base station. To accomplish this we have constructed a VR tricycle (Trike), which is an almost self-contained device permitting a wide range of perceptual cues concerning self-motion to be generated and controlled. To make the device as mobile as possible, it is tethered only for power and optionally networking, and even these links could be removed by using self-carried batteries and a wireless network link. As Trike has a relatively high payload capacity and can be tethered for power, weight and power consumption are not significant constraints on Trike's design. This permits standard, off-the-shelf components to be used on the vehicle. Trike (Fig. 1) is based around a standard size adult tricycle. The drive train of the stock vehicle has been modified so that there is only a single gear and so that pedalling backwards causes the vehicle to go backwards. The slack in the drive chain, which is normally present in a bicycle or tricycle, has been minimised to reduce backlash.

To estimate the vehicle's motion and position relative to some initial position, the vehicle has been instrumented with a potentiometer connected to the steering axis that senses the steering direction, and the drive wheel has been instrumented with an optical encoder so that its rotation can be measured. A small onboard microprocessor monitors these sensors as well as two buttons mounted on the handlebars. The microprocessor communicates with the main onboard computer via a standard serial link.

The main onboard computer is an SGI O2 with dual display (the 'two-headed option'). The two video outputs of the machine are fed to a binocular HMD equipped with a Polhemus head tracker. The tracker's reference transmitter is mounted on the bicycle. The Polhemus tracker therefore reports the position of the helmet relative to Trike, and Trike's instrumentation yields motion and direction information of the machine relative to the ground. Thus the rider's position can be calculated from the combination of these two tracking systems and the appropriate visual display generated.

Tricycle Kinematics

To update the rider's position with respect to the world, it is necessary to construct a kinematic model of the vehicle. Fortunately, bicycles and tricycles have a straightforward kinematic model. Full details of the kinematic models of these and other wheeled vehicles can be found in [5], but the basic concept is sketched here.

Tricycles are steered vehicles in which changes in wheel orientation or steering direction are used to change the trajectory of the vehicle. For a wheeled vehicle to move without slippage, each wheel must follow a circular course around the vehicle's Instantaneous Centre of Curvature (ICC), and must roll on the ground with a velocity that is consistent with the geometry of the wheel placement. As power is provided to the wheels, the entire vehicle will then rotate about the ICC. The task of determining the kinematics of any wheeled vehicle, and the tricycle in particular, reduces to the task of determining the ICC of the vehicle.

The ICC must lie at the intersection of lines drawn through and perpendicular to the rotational axis of each wheel (Fig. 2). Thus for a tricycle, the ICC must lie on a line passing through the rotational axis of the rear wheels which is perpendicular to the body of the tricycle. The front wheel can be steered, and thus the ICC lies on that part of the line passing through the rear wheels that intersects the line drawn along the axis of the front wheel. Given a maximum steering angle of the front wheel, a tricycle has a minimum turning radius and rotates about a point on the line passing through the rear axle.

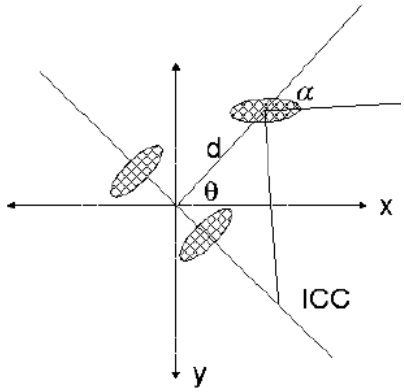

Fig. 2. Tricycle/tricycle kinematics. The shaded ovals represent the three wheels of Trike. The Instantaneous Centre of Curvature (ICC) must lie at the intersection of lines orthogonal to, and passing through the rotation axis of each wheel.

If the steered front wheel is set at an angle α from the straight-ahead direction, and moves with ground contact velocity ν, the tricycle will rotate with angular velocity ω about a point lying a distance R along the line perpendicular to and passing through the rear wheels, where R and ω are given by

$$R = d\tan(\pi/2 - \alpha) \qquad \omega = \nu/(d^2 + R^2)^{0.5}$$

and d is the distance from the front to the rear axle as shown in Fig. 2.

Suppose that the tricycle is at some position (x,y) and *facing* along a line making an angle θ (see Fig. 2) with the x-axis at time t, i.e. it has pose $[x \;\; y \;\; \theta]^T$. Then the ICC is given by

$$ICC = (x - R\sin(\theta), y + R\cos(\theta))$$

and after a short interval Δt the pose of Trike is given by

$$\begin{bmatrix} x' \\ y' \\ \theta' \end{bmatrix} = \begin{bmatrix} \cos(\omega\Delta t) & -\sin(\omega\Delta t) & 0 \\ \sin(\omega\Delta t) & \cos(\omega\Delta t) & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x - ICC_x \\ y - ICC_y \\ \theta \end{bmatrix} + \begin{bmatrix} ICC_x \\ ICC_y \\ \omega\Delta t \end{bmatrix} \tag{1}$$

This describes the motion of a bicycle rotating at distance R about its ICC with an angular velocity ω.

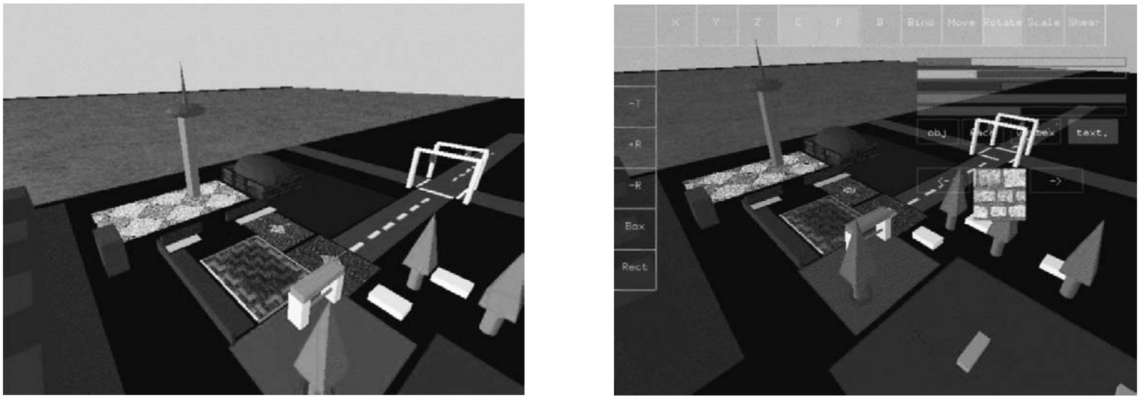

(a) User display

(b) User display with interface overlaid

Fig. 3. A sample rideable environment. The environment is generated using the current state of the head tracker and Trike kinematic model (see text). The user interface can be overlaid over the rendered 3D environment. User interaction is via three buttons which can be used to manipulate existing objects in the environment and to initiate the creation of new objects. The 'brick' pattern in the middle of the image is the currently selected texture for texturing surfaces. The interface permits both a user's eye view as well as the bird's eye view shown here.

The actual implementation on Trike is complicated by the fact that the encoder is mounted on one of the rear wheels, though the basic concept is as described here.

Constructing a Visual World

Given the kinematic model of Trike and the state of the head tracker, the next step is to generate a visual world around which to ride. Figure 3a shows a sample virtual environment used with Trike. The environment was constructed using a point-and-click interface (described below) and consists of polygonal structures which can be texture mapped. The entire environment is shaded and the operator's view is updated based on the state of Trike and the operator's head position and orientation.

Test environments are constructed through the 3D graphical interface shown in Fig. 3b. The user interface was designed to allow the user to interactively create a virtual environment encompassing all three spatial dimensions. It allows the creation of simple geometrical objects such as boxes and rectangles through a point-and-click mechanism. Objects can also be imported following a VRML-like file format. The interface allows a user to specify operations such as move, rotate, scale and shear to act upon existing objects. For many operations, the construction interface manages the parameters necessary to allow such implicitly 3D operations to be performed as sequences of 2D tasks. For example, an object may be moved on the XY, YZ or XZ planes, according to the settings on the interface. The interface also allows setting object properties such as colour, transparency and specularity. Operations through the interface can also take the form of assigning and manipulating object textures. Textures can be set on polygonal faces and then scaled, moved or rotated.

4. Extending Trike's Operational Range

Although Trike can be used 'as is' over short distances, the internal kinematic modelling of the device is not sufficiently accurate for the device to be used for medium- or long-range experiments into human performance in tasks such as cognitive mapping. To study cognitive mapping, it is necessary for Trike's internal model to be accurate with respect to the external physical world. This is required so that the experimenters know what non-visual stimuli the rider is experiencing. It is also important to calibrate the path with respect to a possibly restricted workplace such as an arena, to maintain Trike within the boundaries of the real physical space.

Inaccuracies in Trike's internal model arise due to the fact that its kinematic model is based only on Trike-mounted sensors. Small errors in these sensors and in the kinematic model of the vehicle build up over time resulting in a disassociation between the vehicle's true physical position and its internal model. This is related to the problem of pose maintenance in mobile robotics (see [5]), in which the ongoing pose of a mobile robot cannot be estimated using internal sensors only, and some reference to an external frame of reference is

(a) Trike with the addition of the visual target (large box)

(b) Stereo tracking system

Fig. 4. Hardware to permit long-range tracking of Trike. (a) Trike is enhanced through the addition of a large visual target which is colour coded to aid tracking. (b) Tracking is performed using a stereo camera system.

required to maintain the pose. For Trike, the internal model of the vehicle's pose is inaccurate due to a number of small but cumulative factors including wheel slippage and errors in the steering angle and kinematic model.

Wheel Slippage

Trike's sensor measures rotation of one of the vehicle's rear wheels. As the wheel is not in perfect contact with the ground, and as the wheel deforms as it rolls, the motion of the wheel does not correspond exactly to the motion of the vehicle on the ground.

Steering Angle Error

The potentiometer, which is used to measure the angle of the steering wheel, is calibrated against two hard stops. The orientation of these two hard stops is measured by hand, and this introduces a small error. Although the steering angle can be varied continuously and the resulting resistance of the potentiometer changes continuously, this value is then quantised and assumed to change in direct proportion with the steering angle.

Kinematic Modelling Error

The kinematic model of Trike described above is quite simple. It assumes perfect wheel-ground contact, idealised wheels, etc. It is unlikely that these assumptions will hold in practice. The accumulation of these errors makes the registration of the internal model of Trike drift relative to the real world over time. If Trike is to be used in experiments which involve registering its position with landmarks in the world, a mechanism must be provided to correct the model whenever possible to diminish, or at least bound, these errors.

To address these issues, a stereo vision system has been developed which is able to track Trike as it moves through the world. This stereo system obtains independent estimates of the vehicle's position in a global coordinate system, and sends the proper world coordinates through a network connection to Trike, where they can be integrated with the internal pose estimate of the vehicle. Although this presently represents a second tether, the network link could be established in the future by a radio link. Trike's stereo vision system is composed of two 'off-the-shelf' Logitech QuickCam Pro cameras which are used to track Trike (see Fig. 4b). They are separated with a baseline of 1 m, and are fixed in place so their image planes are as coplanar as possible. Each camera is calibrated independently for the internal parameters with the use of the freely available Intel® Camera Calibration Software. Trike is augmented with a large visual target which can be easily segmented from the background, and the position of the vehicle estimated.

The vehicle is tracked by colouring the visual target distinctively. Standard colour tracking techniques (described in [11]) are used to segment the stereo images obtained by the cameras into regions that contain the target with high probability. Instead of segmenting the images in RGB space, segmentation is

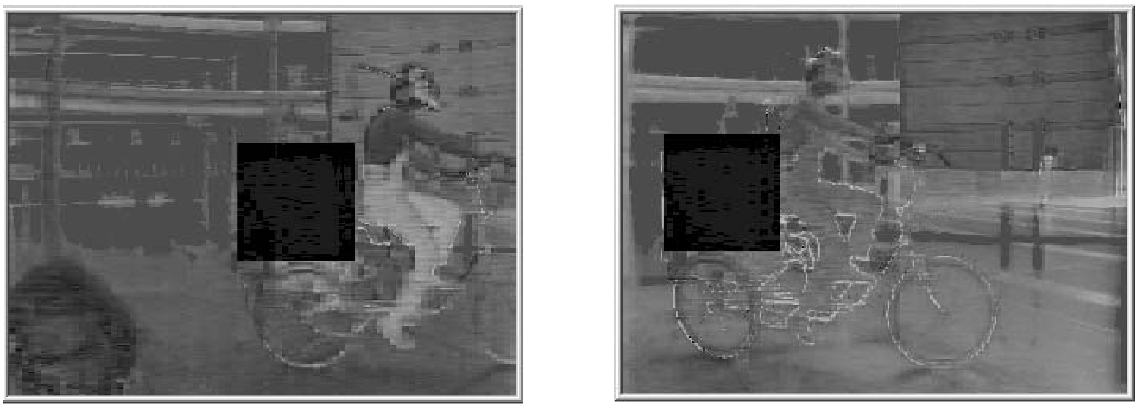

Fig. 5. HSV coded view from each of the stereo cameras. The visual target is easily identified in the images as a dark box. The left camera's image is shown on the left and the right camera's image on the right so this stereopair can be fused into a three-dimensional image by uncrossed fusion.

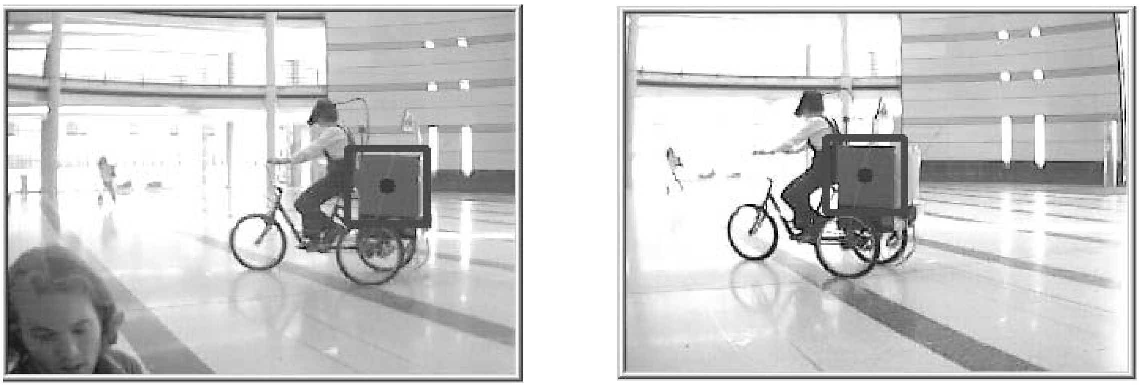

Fig. 6. Stereo view of Trike showing the tracked region marked by a box with a dot in the middle. The left camera's image is shown on the left and the right camera's image on the right so this stereopair can be fused into a three-dimensional image by uncrossed fusion.

based on the hue value of individual pixels (see Fig. 5). Segmentation is approached as a pixel classification process based on Bayesian probabilities. Conditional probabilities are used to determine the probability that a given pixel (hue) value corresponds to the visual target on Trike.

Samples of images of Trike in its operating environment are collected and conditional probabilities $P(Hue \mid non\text{-}object)$ and $P(Hue \mid non\text{-}object)$ are constructed as:

$$P(Hue|object) = \frac{object[Hue]}{T_{obj}} \quad (2)$$

$$P(Hue|non\text{-}object) = \frac{non\text{-}object[Hue]}{T_{non-obj}} \quad (3)$$

In the above, T_{obj} and $T_{non\text{-}obj}$ are the total number of elements in the corresponding Object or Non-object histograms obtained from the sample images. Here *object*[*Hue*] and *non-object*[*Hue*] are the bin counts of the corresponding histogram. Given a particular Hue value, it is then possible to compute the conditional probability that the hue value corresponds to the object using Baye's rule:

$$P(object|Hue) = \frac{P(Hue|object)P(object)}{P(Hue|object)P(object) + P(Hue|non\text{-}object)P(non\text{-}object)} \quad (4)$$

where:

$$P(object) = \frac{T_{obj}}{T_{obj} + T_{non\_obj}} \text{ and } P(non\text{-}object) = \frac{T_{non\_obj}}{T_{obj} + T_{non\_obj}} \quad (5)$$

Once the probability of a pixel being part of the object has been calculated, pixels are classified as being part of the visual target as $P(object \mid Hue)\ \Delta$, where Δ is a predefined threshold. Given the classification of pixels in a given image, the 0th and 1st order moments are computed to identify the centroid of the largest region of classified pixels.

The visual target is tracked independently in both images, and an epipolar constraint is applied to ensure that the same target is identified in both images. This is accomplished by calculating the tracked centre of the object and checking to see if the left and right centres are within a given distance threshold both vertically and horizontally. If the tracking regions deviate past these thresholds, the left image enforces the epipolar constraint on the right image by setting the search region to be along the same horizontal line in the image. When the target is acquired once again, the constraint is lifted and independent tracking resumes. Figure 6 shows Trike being tracked by the visual system. The large squares indicate the tracked target, with the centre of the target marked with a large dot.

In subsequent frames, the search region is moved in each image by means of a predictive Alpha-Beta Filter [12], which also helps to smooth the final data. The update equation of the Alpha-Beta Filter is defined as

$$\hat{x}_{(k+1|k+1)} = \hat{x}_{(k+1|k)} + \begin{bmatrix} \alpha \\ \beta / \Delta t \end{bmatrix} \left[\hat{z}_{(k+1)} - \hat{z}_{(k+1|k)} \right] \tag{6}$$

where α and β are optimally defined in [12].

From frame to frame, the location of the search region is updated as

$$x_{(k+1)} = x_{(k)} + v_{(k)} \Delta t \tag{7}$$

where $x_{(k)}$ is the current state vector (u,v image coordinates) and $v_{(k)}$ is the velocity of each component of the state. The prediction allows us to follow the object with higher accuracy, and to avoid getting confused with other objects that may have a similar hue.

The tracking data is sent to Trike through a peer-to-peer network connection. Trike is thus able to use the data as it moves around the physical world. At present, the model and tracking data are not integrated using a strategy such as Kalman filtering [13]; rather, Trike operates in either Internal Mode, using only its internal sensors or Tracker Mode, using the external video tracker. A Kalman filter-based approach to merging the data in an ongoing manner is currently under development.

Figure 7 illustrates the effectiveness (and need) for a tracking system which makes reference to the external world. Relying only on internal tracking leads to substantial drift (Fig. 7b). The external visual tracking keeps the location of Trike constrained to close to its true physical position (Fig. 7c).

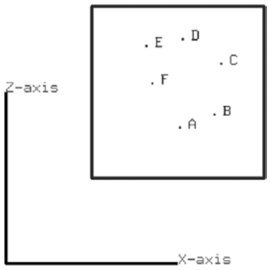

(a) Position of known landmarks in the environment

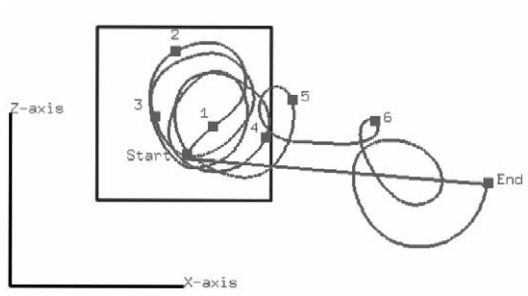

(b) Internal bike model position

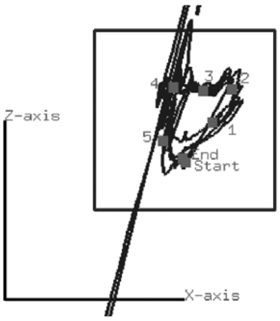

(c) External sensor model position

Fig. 7. Trike is ridden in a tight circle through landmarks A–F marked on the floor (a). (b) Position of Trike measured by the *internal sensors*. The numbered points are shown only to explain the direction Trike moves through the world. (c) Raw tracker data from the stereo tracker. Again, the numbers indicate only the direction in which Trike moves and do not correspond with those in (b). Some gating of the tracking data is still required – note the large error to the top and bottom of (c) when the tracking data are temporarily lost.

5. Discussion

During self-motion we receive a number of different cues from a variety of sensory systems. These cues include optic flow, auditory motion, vestibularly sensed angular and linear acceleration, proprioceptive and somatosensory information about the position and movements of the limbs, and knowledge that instructions have been sent to the various muscles. Each sensory system has different response characteristics, and the systems interact in complex and subtle ways to generate our perception of self-motion and orientation in the world. It might be thought that VR primarily seeks to mislead the visual sense, but, in fact, the process of providing an imaginary space that a person is able to explore, seeks to fool the totality of the various systems that contribute to spatial awareness.

VR systems which augment or modify one or more of these sensory systems may confound our overall sensation of motion in various ways. A common result of this confusion is nausea (cybersickness) and an associated degradation of performance. Designers of immersive visual systems must take great care that their augmentation of vision or other sensory inputs does not interfere unpredictably with the normal perceptual processes. Trike utilises standard VR technologies to generate an immersive visual display while utilizing real motion to generate compatible non-visual cues. This permits operators a wider operational range than is found in more traditional VR systems. Limited only by the length of the power cord and the available free space, subjects can explore large virtual environments and obtain appropriate visual and non-visual cues to their motion.

Acknowledgements

Financial support from the Centre for Research in Earth and Space Technology (CRESTech) and NSERC Canada are gratefully acknowledged.

References

1. Matias E, MacKenzie IS, Buxton W (1996) A wearable computer for use in microgravity space and other non-desktop environments. In: Proc Companion of the CHI'96 Conference on Human Factors in Computing Systems, New York, NY 69–70
2. Ockerman JJ, Pritchett AR (1998) Preliminary investigation of wearable computers for task guidance in aviation inspection. In: Proc IEEE 2nd Int Symp on Wearable Computers, Pittsburgh, PA
3. Lewis SA, Havey GD, Hanzal B (1998) Handheld and bodyworn graphical displays. In: Proc IEEE 2nd Int Symp on Wearable Computers, Pittsburgh, PA 102–107
4. Feiner S, McIntyre B, Höllerer T, Webster A (1997) A touring machine: prototyping 3d mobile augmented reality systems for exploring the urban environment. In: Proc IEEE 1st Int Symp on Wearable Computers, Cambridge, MA 74–81
5. Dudek G, Jenkin M (2000) Computational principles of mobile robotics. Cambridge University Press: New York, NY
6. Harris LR, Jenkin M, Zikovitz D, Redlick F, Jaekl P, Jasiobedzka U, Jenkin H, Allison RS (2002) Simulating self motion I: cues for the perception of motion. Virtual Reality 6(2): 75–85
7. Gradecki J (1994) The virtual reality construction kit. Wiley: New York, NY
8. Bülthoff H, van Veen HAC (1999) Vision and action in virtual environments. In: Jenkin M, Harris L (Eds). Vision and attention. New York, NY: Springer-Verlag
9. Iwata H (1999) Walking about virtual environments on infinite floor. In: Proc IEEE Int Symp on Virtual Reality 286–293
10. Darken RP, Cockayne WR, Carmein D (1997) The omni-directional treadmill: a locomotion device for virtual world. In: Proc UIST'97 213–221
11. Jones MJ, Rehg JM (1998) Statistical color models with applications to skin detection. Technical Report CRL 98/11, Compaq Computer Corp.: Cambridge, MA
12. Brown C Ed (1994) Tutorial on Filtering, Restoration, and State Estimation. Technical Report 534. Department of Computer Science, University of Rochester: Rochester, NY
13. Bozic SM (1979) Digital and Kalman filtering. Edward Arnold: London

---

Correspondence and offprint requests to: *M. R. Jenkin, Centre for Vision Research and Department of Computer Science, York University, 4700 Keele St, Ontario, M3J 1P3 Canada. Email: jenkin@cs.yorku.ca*