An Investigation into the design of an Interface for Interaction with a Virtual Environment representing a four-dimensional object

R Wellard and S.C Chapman

Physics Department, University of Warwick,
Coventry CV4 7AL
{wellard, sandrac}@astro.warwick.ac.uk

Abstract. Five methods of mapping the input from a six-degrees-of-freedom input device to the set of rotations available in four-dimensional space are evaluated with respect to how well they can be used to perform four-dimensional target aquisition tasks. Also evaluated are the user's opinions of the methods of interaction. Two of the five interaction methods perform significantly better than the rest allowing some general results to be drawn. ...

1 Introduction

4-Space is the theoretical four-dimensional vector space spanned by four orthogonal vectors. The vector space most closely associated with the universe we inhabit is three-dimensional, that is, it is spanned by three orthogonal vectors. Any position in 3-space can be described by its distance from a convenient point of origin along each of three orthogonal vectors, and illustrated in the familiar format *(x,y,z)*. For a position in 4-space to be similarly described a new variable must be introduced, the distance from the origin along the fourth orthogonal vector. The position can now be referenced as the point *(x,y,z,w)*. The properties of four-dimensional objects are very different to those of three-dimensional objects, and are well explained in works by Banchoff [3], Manning [12], and Abbott [1]. This work is only interested in controlling the extended set of rotations available in 4-space, and in particular which of the five methods of interaction presented (Section 3) form the most effective link between user and four-dimensional rotations. Although the aim of this work is to classify these methods of interaction according to how well they enable the user to control four-dimensional rotations, it is hoped that some general results for four-dimensional manipulation design can be drawn. In order to evaluate the five methods of interaction (which shall now be referred to as the five *control mappings*) an experiment was designed to test user performance under each of the control mappings, a complete description of which is provided in Section 2.

Wellard & Chapman [14] contains a brief explaination of 4-D rotations, projections through a point, and their effects on the hypercube; specifically aimed at explaining these issues in the context of this work. However, as an aid to understanding this paper, we reiterate some of the basic concepts here. Rotations in 4-Space are about a plane (in comparison with rotations in 3-space, which are about a line), and all orientations of a 4-D object can be realised by rotations about three perpendicular planes. The hypercube [2] is the four-dimensional analogy of the cube, and consists of eight cube-faces attached at their faces (cf a cube is constructed by attaching six square-faces together at their edges). With a hypercube in a certain orientation, it's 3-D projection through a point forms a cube-inside-a-cube shape ([3] p120-122).

Lessons can be learnt from work on control mappings between user and a 3-D virtual environment [13], [5], [7], and studies on input devices and their effectiveness with respect to the user [10], [11], [16]. However, the work presented here addresses an original research problem since the abstract nature of the task to be performed, and the fact that interaction is with a four-dimensional object and not with a familiar three-dimensional one, means that not all the previous results can be extended to this work.

Within this paper, Section 2 describes the experimental method; Section 3 introduces the control mappings ([14] contains a discussion of their theorized affordances and weaknesses); Section 4 presents the significant results, which are discussed in Section 5 and collated into the conclusions of Section 6.

2 Experimental Method

A within subjects 5x7 repeated measures experiment was carried out in which subjects performed seven target aquisition tasks using different control mappings. The dependent variable was *level of closeness to target achieved*. The independent variables were *control mapping*, at five levels (mappings 1 to 5), and *target orientation*, at seven levels (seven different targets presented in the same order). After completing the seven tasks using a particular mapping the subjects answered a questionnaire; resulting in a separate within subjects, repeated measures data set, where the dependent variable was *questionnaire-score*, and the independent variable was *control mapping*.

Eighteen (14M 4F) subjects were randomly selected from a set of 24 (19M 5F) physics postgraduate volunteers; as with the within subjects design, each subject completed all the tasks using every control mapping. The order in which the subjects used the control mappings was randomly[1] assigned, however the target aquisition tasks were presented in the same order for each control mapping, and for each subject. The equipment used was a FakeSpace ImmersaDesk package, which consists of a five foot by four foot stereo-projection screen, a pair of tracked CrystalEyes shutter glasses, and a tracked FakeSpace WorkWand input device with three digital buttons and a two-degrees-of-freedom analogue joystick. Together with code written specifically for this experiment; and a questionnaire

---

[1] a problem with the equipment meant that there was a tendency for control mappings 4 and 5 to be used after the others

constructed using the presence questionnaire developed by Witmer and Singer [15] as a guideline.

Visually the subjects were presented with the projected image of a hypercube drawn in transparent blue four feet above the ground and approximately three feet in front of them, with it's orientation in 4-space such that it's projection takes the cube-inside-a-cube form. The 'inside' cube-face of the hypercube is drawn opaque red, and it's sides labeled one to six. Also presented is the projection of an opaque green target cube that also has it's sides labeled one to six. This target cube exists in the same 4-space as the hypercube. The target aquisition task is for the subject to rotate the hypercube in 4-space so that the red cube-face of the hypercube occupies the same space (in 4-space and therefore in the 3-D projection space) as the green target cube, and the numbers on the square faces of the two cubes match up. Once this task has been completed the orientation of the hypercube is re-set, and a new target is presented.

Subjects were met individually by the experimenter and asked to read an introduction to the experiment, and then a brief set of instructions for the control mapping they were about to use. They then put on the CrystalEyes shutter glasses, were given the Workwand, and the code was started. The subjects were given two minutes to get used to the control mapping during which they were presented with a practice target so that they knew what they would have to do during the trial. After the two minutes the orientation of the hypercube was re-set and the first target was presented. Once either the subject was satisfied that the target orientation had been achieved, or they had given up, or four minutes had elapsed, the hypercube's orientation was re-set, and the next target presented. All measurements were automatically recorded by the computer. Distance from the target was calculated as the average of the linear distances, in 4-space, between each vertex of the red cube-face and the appropriate vertex of the green target cube. The minimum of this distance value during the subject's attempt to match the target was then converted into the interval scale, between 0 (distance greater than 0.24) and 8 (distance less than 0.03), *level of closeness to target.*

3 Control mappings

This section describes each of the five proposed control mappings using the following conventions and notation. The coordinate system for the three-dimensional space the user moves in consists of three perpendicular vectors, x_3, y_3 and z_3. With the user facing the ImmersaDesk, x_3 is in the direction of the user's right, y_3 increases upwards, and z_3 increases in the opposite direction to the way the user is facing. This set of axes has as its origin the centre of the Workwand; thus the Workwand itself forms a vector through the origin, and all rotations of the Workwand are about a line through the origin.

The coordinate system for the four-dimensional space, within which the four-dimensional objects are manipulated, consists of four perpendicular vectors, x_4, y_4, z_4 and w_4. The origin of this coordinate system is at the centre of the four-dimensional object. Thus all rotations of the four-dimensional object about planes through the origin do not translate the object's centre. The four-dimensional object is projected through a point on the w_4-axis onto a three-dimensional hyperplane. This projection is displayed in such a fashion that movement of a four-dimensional object in the positive x_4-direction would cause the image to move to the user's right, movement in the positive y_4-direction would cause the image to move upwards, movement in the positive z_4-direction would cause the image to move towards the user, and movement in the positive w_4-direction would cause the image to get larger. This obvious choice of projection and display-orientation, means that rotations of the four-dimensional object affecting only the x_4, y_4, and z_4 dimensional coordinates of its points will cause the projected image to rotate as a three-dimensional object (in Section 5 these rotations are referred to as *non-morphing* rotations). For example, rotating the four-dimensional object about the x_4-w_4-plane would cause the image to rotate about a line through its centre parallel to the x_3-axis.

A virtual linear slider is a one degree of freedom input device, which comprises of a handle that can be moved one-dimensionally between two end points. In the presented control mappings the position of the handle on the linear slider represents the rotation of the four-dimensional object about a particular plane. Which plane the virtual linear slider is associated to, and how it is manipulated, depends upon the control mapping.

Rotations of the Workwand are separated into two categories: rotations altering the Workwand's directional axis (a straight line through the Workwand in the direction the Workwand is pointing in), and rotations about the Workwand's directional axis. This second form of rotations shall be called twists about the Workwand's axis. For all the mappings presented, except mapping 3, twisting the Workwand will perform the expected action; that is rotating the Workwand about a certain axis has the same effect as twisting the Workwand, while that axis is the directional axis of the Workwand. With all the control mappings presented in this paper which utilise the movement of the Workwand, the trigger button will act as a ratchet button ([6], p70-71) to the movement of the four-dimensional object; hold the button down to perform a rotation, release and movement of the Workwand has no effect. This makes rotating the four-dimensional object about large angles easier, and enables the user to have the Workwand permanently in a comfortable position.

3.1 Mapping 1: Three-dimensional wand rotations with the ability to change the dimensions the rotations are mapped to

This control method maps the three-dimensional rotations of the Workwand to rotations altering three-dimensional coordinates of the four-dimensional object. For the user to be able to rotate the object fully, they are able to select which three of the four dimensions their movement affects. For example the user may rotate the Workwand in 3-space, and cause a rotation of the four-dimensional object changing only its x_4, y_4 and z_4 values; then push a button to change the dimensions their movements are mapped to, and now rotations of the wand will cause rotations of the four-dimensional object changing only its w_4, y_4 and z_4 values. Specifically, one of the four dimensions is fixed, and the other three have a one to one correspondence with the three dimensions the user is moving the Workwand in. Denoting the 4-space axes a_4, b_4, c_4, d_4; if d_4 is the fixed axis and the other three are mapped such that x_3=a_4, y_3=b_4, and z_3=c_4, then a rotation of the Workwand about a vector in x_3-y_3-z_3-space is mapped to a rotation of the four-dimensional object about the plane formed by the corresponding vector in a_4-b_4-c_4-space and a vector in the d_4 direction. The sub-mappings available to the user are:

(1) x_3=x_4 y_3=y_4 z_3=z_4 w_4 fixed
(2) x_3=-w_4 y_3=y_4 z_3=z_4 x_4 fixed
(3) x_3=x_4 y_3=-w_4 z_3=z_4 y_4 fixed
(4) x_3=x_4 y_3=y_4 z_3=-w_4 z_4 fixed

Mapping the 3-space axes to -w_4 is done so as to keep the other two 3-space axes mapped to their 4-space counterparts in the same direction as they are in sub-mapping 1. For example, whilst sub-mapping 1 is selected, a rotation of the Workwand about the x_3-axis corresponds to a rotation of the four-dimensional object about the x_4-w_4-plane. And whilst sub-mapping 4 is selected, a rotation of the Workwand about the y_3-axis corresponds to a rotation of the four-dimensional object about the y_4-z_4-plane. The user is able to cycle through these sub-mappings in one direction by pressing the left button on the Workwand, and cycle back through by pressing the right button. A set of axes is displayed in the lower left corner of the ImmersaDesk to indicate to the user the non-fixed dimensions, and their appropriate directions.

3.2 Mapping 2: Three-dimensional wand rotation together with change in position of wand

Under this mapping, rotation of the Workwand about a vector in x_3-y_3-z_3-space corresponds to a rotation of the four-dimensional object about the plane formed by the same vector, in x_4-y_4-z_4-space, and the w_4 vector. This is the same manipulation scheme as in sub-mapping 1 of mapping 1. But in addition to this, spatial movement of the Workwand whilst the ratchet button (trigger button) is held down also causes a rotation of the four-dimensional object thus: The change in position of the Workwand between tracker position updates forms a vector. Perpendicular to this vector is a unique plane (through the origin) in x_3-y_3-z_3-space, which corresponds to a plane in x_4-y_4-z_4-w_4-space (under the mapping x_3=x_4, y_3=y_4, z_3=z_4, w_4=0), and it is about this plane that the four-dimensional object is rotated. The length of the vector formed by the change in position of the Workwand is proportional to the angle the four-dimensional object is rotated through. For example, if the Workwand is moved along the x_3-axis (i.e. to the right), the four-dimensional object is rotated about the y_4-z_4-plane (as the y_3-z_3-plane is perpendicular to the x_3-axis).

3.3 Mapping 3: Three-dimensional rotation together with an alternative action for twisting the wand

This mapping separates the action of twisting the Workwand about its directional axis from rotating the Workwand and causing its directional axis to change. This mapping is a variation of mapping 2, described in section 3.2; instead of the change in position of the Workwand describing a vector in 3-space, the directional axis of the Workwand describes the vector. With the angle the Workwand is twisted, being equal to the angle through which the four-dimensional object is rotated, about the plane perpendicular to the described vector. All other (non-twisting) rotations of the Workwand have the same effect as those in mapping 2, and sub-mapping 1 of mapping 1.

3.4 Mapping 4: Three virtual linear sliders

As mentioned in section 1, all orientations of a four-dimensional object are achievable via rotations about three perpendicular planes. This mapping maps the position of the handles of three virtual linear sliders to the rotations of the four-dimensional object about three perpendicular planes; the z_4-w_4-plane, the x_4-w_4-plane, and the y_4-z_4-plane. The limits of the sliders represent 180° and -180° rotations of the four-dimensional object, from its original orientation, about each of the planes. For example, the virtual linear sliders are positioned horizontally, and moving the handle from its original central position, to the extreme right causes the four-dimensional object to rotate 180° in one direction about the appropriate plane. The letters ”z-w”, ”x-w” and ”y-z” are displayed above the sliders indicating which plane that slider rotates the four-dimensional object about. The user can move the handle on the selected slider by pushing the analogue joystick on the Workwand, left or right; and pressing one of the three buttons on the wand selects one of the sliders, one button per slider.

3.5 Mapping 5: Three-dimensional Workwand rotation together with one virtual linear slider

This control mapping uses the same mapping for the rotations (and twisting) of the Workwand as mapping 2, and sub-mapping 1 of mapping 1. Together with this manipulation, the user is given control over a virtual linear slider which controls the rotation of the four-dimensional object about the y_4-z_4-plane. This slider is controlled in the same way, and has the same limits, as those of mapping 4, except that it is always selected.

4 Results

Subjects *level of closeness* results were submitted to a 5 (control mappings) × 7 (target orientation) repeated measures analysis of variance. This analysis revealed highly significant effects of control mapping [$F(4,17)=15.75$, $p<0.00001$]

and target orientation [F(6,17)=33.59, p<0.00001]; with a lesser, but still highly significant control mapping × target orientation interaction [F(24,17)=5.63, p<0.0003]. Figure 1 shows the mean level of closeness together with the 95% confidence interval for each control mapping and target orientation. It is clear from Figure 1b that not all the targets were equally difficult to achieve. It is possible to justifiably discard one of the targets and split the remaining six into two equal sized catagories, *easy* and *difficult* targets. Targets 1, 3, 5 and 6 can be achieved by rotating the hypercube through 90° about the planes formed only by the axes, as a result the shape of the projected image of the target cube is either a cube or a truncated triangular based pyramid. Target 5 is to be discarded, as the target cube occludes the hypercube in the visual feedback, making manipulation of the hypercube unfairly difficult; the rest (targets 1, 3 and 6) form the *easy-target* catagory. Target orientations 2, 4 and 7 are more general, and the projected images of the target cubes for these orientations do not form regular 3-D shapes; these targets form the *difficult-target* catagory. Subsequently the data was collapsed across the easy and difficult target catagories and subjected to a 5 (control mapping) × 2 (target difficulty) repeated measures analysis of variance. Revealed by this analysis were strong dependencies on control mapping [F(4,17)=16.79, p<0.00001], and task difficulty [F(1,17)=146.02, p<0.00001]; together with a highly significant control mapping × task difficulty interaction [F(4,17)=18.124, p<0.00001]. Figure 2a shows the mean level of closeness together with the 95% confidence interval for each control mapping, catagorized into easy and difficult target orientations.

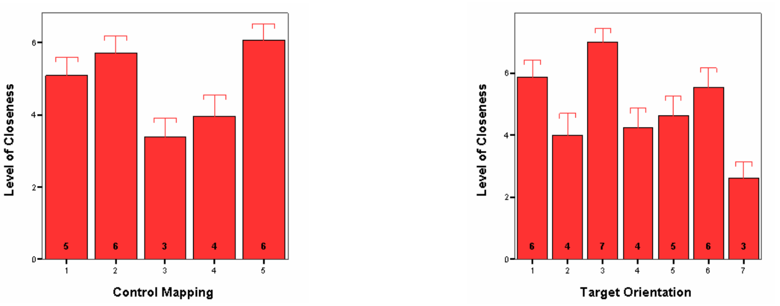

Fig. 1. Graphs showing the mean level of closeness for each (**a**) control mapping (left) and (**b**) target orientation (right). Error bars indicate 95% confidence interval

Questionnaire results were collated by equating an ordinal answer-score between 0 and 6 to each answer depending upon which box was ticked (0 for an answer implying that the control mapping was difficult or awkward to use). The *questionnaire-score* given by a particular subject for a particular mapping was then calculated as the mean of the individual answer-scores. Subject's questionnaire-scores were then subjected to a repeated measures analysis of variance. This analysis revealed that the questionnaire-scores were strongly dependent on the control mapping [$F(4,17)=15.68$, $p<0.00002$]. Figure 2b shows the mean questionnaire-score together with the 95% confidence interval for each control mapping.

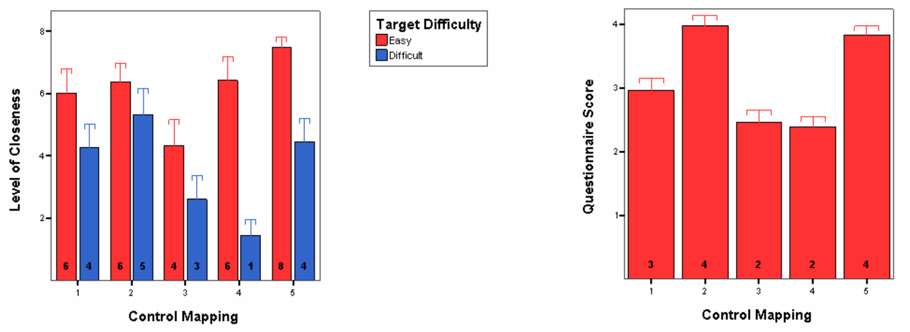

Fig. 2. **a**(left) Bars show mean level of closeness for each control mapping, catagorized into easy and difficult target orientations. **b**(right) Bars show mean questionnaire score for each control mapping. For both graphs, error bars indicate 95% confidence interval

5 Discussion

It is clear from these results that control mappings 2 and 5 form the most effective (Figure 1a), and most popular (Figure 2b), link between user and space of four-dimensional rotations; mappings 3 and 4 form the weakest and least popular link; and control mapping 1 lies somewhere in between. According to Wellard & Chapman [14] the only attributes mappings 2 and 5 have in common are that they do not treat all four-dimensional rotations equally, but rather classify them into two sets, rotations which affect the shape of the 3-D projection, *morphing* rotations, and those which do not, *non-morphing* rotations. Evident from observations whilst conducting the experiment and points of view written by the subjects on the back of the questionnaire (although no statistical evidence can be presented here), was that the subjects also performed this classification. Thus, for a control mapping to perform well it must form an effective yet separate link between manipulation of the input device and both these sets of rotations. Non-morphing rotations cause the projection to rotate as a 3-D object, therefore the most obvious way to perform these rotations is in the same way that it is done in everyday life; the rotation performed with your hand is the rotation the object undergoes. Mappings 2 and 5 have this method of controlling this set of rotations, as does mapping 1 under sub-mapping 1; but mapping 1 is weakened by the amount the user has to learn, each sub-mapping has new rules defining how movements of the workwand effect changes in the visual feedback.

Most users only learnt the effects of sub-mapping 1, and just cycled through the others until a useful effect was noticed (if one was noticed at all).

For morphing rotations the slider control was very useful, subjects had no pre-conceived ideas about the effect of moving a slider. Mappings 4 and 5 used a slider for morphing rotations; however, mapping 4 had a very cumbersome [4] method of performing non-morphing rotations. Translating the workwand under mapping 2 causes the image of the red cube-face of the hypercube to translate, from it's initial position, in the direction the workwand has been moved. This effect is reversed once the hypercube has been rotated through about 110°, but this direct link between movement of the workwand and movement of the visual feedback was very easy for the subjects to learn. Also the dimension of the space of morphing rotations available under mapping 2 (and mapping 3) is greater than that offered by any finite number of sliders. Irrespective of target orientation, the target shape of the projection can be formed by one translational movement under mapping 2, and therefore movement can always be performed towards the target. This is not the case for mappings 1, 4, and 5 where, due to the reduced set of rotations available, movement away from the target is sometimes necessary in order for the target to be achieved. Mappings 4 and 5 have the smallest set of morphing rotations available, and Figure 2a shows that these mappings have the largest difference in performance for easy and difficult target orientations, suggesting that reducing this set makes achieving difficult orientations harder.

Many subjects were not able to separate their movements of the workwand into twisting and non-twisting rotations, and therefore were not able to separate the morphing and non-morphing rotations under mapping 3. Separating these two types of 4-D rotations was easy under mappings 4 and 5, but caused a little difficulty under mapping 2, as translating the workwand often resulted in a slight rotation, and vise-versa.

The most interesting result from this experiment is the lower dependency task difficulty had on the level of closeness results for control mapping 2 (Figure 2a). For complete control over 4-D rotations, target orientation should not be a factor; if a user is performing rotation tasks in 4-space there should not be preferred orientations, as 4-space does not have preferred orientations. This low dependency on task difficulty, together with the high results for both difficulty levels, indicates that control mapping 2 forms a very strong, functional link between user and the space of 4-D rotations.

6 Conclusions

No rotary-to-rotary [10] or integrality and separability [8] [9] dependencies, as hypothesised by Wellard & Chapman [14], were found. However, some general results can be extracted from this investigation. Firstly, the projected image is three-dimensional, yet it's shape cannot be easily recognised without the ability to rotate and examine it. This, together with the separate nature of the conceptual model users form for the space of 4-D rotations, means that they must be able to separate morphing rotations from non-morphing rotations. Secondly, performing the non-morphing rotations in the manner of mappings 2 and 5 appears to be the most intuitive way. Finally, the degrees of freedom available for morphing rotations influences the ability of the user to achieve 'difficult' orientations. If the user can only perform morphing rotations about one plane (mapping 5), then they must guess the shape of the projection needed, rotate the projection to see if the shape is correct, and if it is not then repeat this procedure; ballistic movements. If they can perform morphing rotations about any plane (mapping 2), then the user can always move towards the desired shape of the projection; the shape is continuously steered towards it's target.

References

1. Abbott EA (1944) Flatland (5 ed), Basil Blackwell, Oxford
2. Strauss C and Banchoff TF (1978) The Hypercube - Projections and slicing, Currently unavailable, see http://www.math.brown.edu/ banchoff for details
3. Banchoff TF (1996) Beyond the Third Dimension, Scientific American Library, New York
4. Chen M, Mountford SJ, Sellen A (1988) A Study in Interactive 3-D Rotation Using 2-D Control Devices, Computer Graphics, 22(4), 121-129
5. England R (1995) Sensory-motor systems in virtual manipulation, in Simulated and Virtual Realities, Carr and England, ch 7, 131-117
6. Gabbard JL and Hix D (1997) A Taxonomy of Usability Characteristics in Virtual Environments, available at http://csgrad.cs.vt.edu/ jgabbard/ve/taxonomy/
7. Houde S (1992) Iterative Design of an Interface for Easy 3-D Direct Manipulation, CHI '92 Conference of Human Factors in Computing Systems, 3 - 7 May 1992, 135-142
8. Jacob RJK and Sibert LE (1992) The Perceptual Structure of Multidimensional Input Device Selection, CHI '92 Conference of Human Factors in Computing Systems, 3 - 7 May 1992, 211-218
9. Jacob et al (1994) Integrality and Separability of Input Devices, ACM Transactions on Computer-Human Interaction, 1(1), 3-26
10. MacKenzie IS (1995) Input Devices and Interaction Techniques for Advanced Computing, in Virtual Environments and Advanced Interface Design, Barfield and Furness, Oxford University Press, ch 11, 437-470
11. MacKinlay JD, Card SK, Robertson GG (1990) A semantic analysis of the design space of input devices, Human Computer Interaction, 5, 145-190
12. Manning HP (1960) The Fourth Dimension Simply Explained, Dover Publications, New York
13. Smets et al (1995) Designing in Virtual Reality: perception-action coupling and affordances, in Simulated and Virtual Realities, Carr and England, ch 9, 189-208
14. Wellard & Chapman (2000) An Investigation into Control Mappings for the Manipulation of Four-Dimensional Objects Represented within a Virtual Environment, EGUK 2000 Conference Proceedings, 4-6 April, 159-170
15. Witmer et al (1998) Measuring Presence in Virtual Environments: A Presence Questionnaire, Presence: Teleoperators and Virtual Environments, 7(3), 225-240
16. Zhai S and Milgram P (1993) Human Performance Evaluation of Manipulation Schemes in Virtual Environments, IEEE Virtual Reality Annual International Symposium Proceedings, 155-161, available at http://vered.rose.toronto.edu/people/shumin_dir/papers/VRAIS93.html