The Virtual Sky is not the Limit: Ethics in Virtual Reality

Blay Whitby

February 1993

Introduction: What is VR?

Virtual reality (VR) is the name applied to one of the latest trends in high technology research. In essence it is the delivery to a human or several humans of the most convincing illusion possible that they are in another reality. This reality exists only in digital electronic form in the memory of a computer or several computers. Hence it is accurately described as 'virtual'. Its reality stems from the convincing nature of the illusion, and most importantly for moral considerations, the way in which human participants can interact with it. If one were to ask for a demonstration of VR, one would probably be asked to don a strange looking helmet. Inside this helmet would be a number of small screens on which pictures are projected immediately in front of the wearer's eyes. One might also be asked to wear one or more 'data gloves' or similar devices. Like the helmet, these would be generously connected by wires to the associated computing machinery. The function of a data glove is to transmit, as accurately as possible, the movements of the wearer's hand. These movements are fed into the computer, where they are translated into 'actions' (perhaps 'virtual actions') within the VR. Devices can also be attached to one's legs, though this is less common - users tend to 'float' in a given direction, rather than walk, in present day VR. The experience of being connected to all this high technology would be (fairly) close to entering another world. One sees this world via the screens within the helmet. Movement of the head and or the eyes is sensed and the pictures appropriately modified enabling one to 'look around' the world. One can (generally) move, pick up objects and interact with other characters within this world. These characters can be either computer-generated or other human participants, similarly connected via helmets, data gloves and so on. Within the world a large amount of activity may be possible. Because VR is highly interactive, what actually happens is determined by the user or users. On the other hand, what is possible is determined by the programmer or programmers. This is an important distinction to which we will return. Many readers may feel that the experience of present-day VR is probably not convincing enough to carry the label 'reality'. Against this argument it must be noted that humans have immense powers of imagination and a willingness to suspend disbelief. In other words the simulated reality does not need to be a perfect simulation in order to for users to come to believe that it is some form of reality. In addition, the main imperfection in simulation at present comes from the difficulties inherent in presenting a sufficiently convincing computer generated image. Presenting a convincing visual input

to a human being requires a computer which can handle a vast amount of information. This is both difficult and expensive with existing technology. It is now a familiar, but true, cliche to observe that the power and capability of computing machinery is increasing at a tremendous rate. In addition, a wide variety of techniques for producing more convincing VR are currently being researched. It is reasonable to expect, therefore, that in this crucial respect VR will become steadily more convincing. Of course, this steady improvement will depend on the availability of financial support. This, in turn depends on the existing and anticipated applications of VR. The question of what the applications might be is of central importance to any discussion of the ethical implications. While the technology remains a quaint experiment, it raises few, if any, moral problems. If it becomes generally applied, the moral implications will become much more important.

Implications: The Potential Applications of VR

Predictions of the future applications of new technologies are notorious for their inaccuracy. However there is little difficulty in making various predictions about the future of VR. In many respects VR represents a collection of developments of previous technologies, rather than complete innovation. Principal among these are simulators, interactive multimedia systems, and computer and arcade games. When a new technology, such as VR, becomes fashionable, its enthusiasts (and salesmen) may invest a great deal of energy in asserting that it is radically distinct from its technological antecedents. In the case of commercially available systems they have a clear vested interest in doing so. A more realistic approach would stress that most technological innovations, and VR is typical, are a relatively small development and improvement of pre-existing technologies. This approach may also help show the apparent novelty of the moral problems raised to be largely illusory. One of the major antecedents of VR was work on flight simulation for combat helicopter pilots. VR retains much in common with flight simulation. In particular, it can provide a training environment in which mistakes are less permanent and costly than they would be in reality. This feature will provide a wealth of application areas for VR, quite possibly assisted, by the availability of generous funding for potential military applications. Just as pilots, submarine captains, and tank commanders are today trained in complex simulators they will in the near future be trained in VR. One particular advantage of VR over a simulator in this application is the way in which it can incorporate multiple participants, for example, in competition or in combat with each other. VR is likely to find successful applications in many forms of combat training. The usefulness of VR in combat training is also accompanied by its usefulness in many forms of civilian training. In VR dangerous chemicals and machinery can be handled realistically without physical danger to the user or users. In many situations it is desirable to allow learners to make mistakes and yet protect them from the consequences of those mistakes. This is obviously the case with training for the management of nuclear power stations or dangerous chemical plants. What may not be so obvious might the advantages of using VR instead of a laboratory for teaching physics and chemistry in schools. Of

course, such a 'virtual laboratory' will be limited in its ability to give practical familiarity with the equipment and techniques. However in many areas, the handling of radioactive and other dangerous materials being a conspicuous example, it will have clear advantages. Even less obvious may be the usefulness of using VR to let users 'enter' a period of history. This would provide a useful way of teaching history, either in a school or a museum. The use of VR is already proving of interest to creative artists. Anne Barclay Morgan (Barclay Morgan 1992) has pointed out that VR (or cyberspace as it sometimes called in this area) is a medium which offers possibilities such as interactive paintings and sculptures and plots which can be changed by the audience. In addition to training and the acquiring of factual information VR will provide a useful tool for education. It is, in many ways, the ultimate development of Seymour Papert's 'microworlds' (Papert 1980). The attractions of being able to learn through doing, particularly in co-operation with others will soon be seen by educators. Another major area where VR will sooner or later find successful application is in the field of entertainment. This again, is an application area which is essentially an extension of existing technologies and practices. Just as the flight simulator can be seen as a training precursor of VR, so can the cinema,video game, and computer-game be seen as entertainment precursors of VR. Already arcade games are moving towards military VR in terms of the realism of their displays and the richness of interactions possible. As the technology improves there will be a strong market for VR in arcade and home entertainment. The development of the entertainment market depends on the technology achieving a sufficiently low cost and this may take slightly longer than the developments discussed in the last couple of paragraphs. However, there is every good reason to believe that VR used for entertainment will become commonplace within a few years. Two routes of development are possible. Firstly, (relatively) low-tech, low-cost VR could become available for use in the home. Secondly, (relatively) high-tech, high-cost VR might be set up in population centres and hired by the minute by large numbers of users. These two routes are not mutually exclusive and may well be able to co-exist, just as at present many people both visit the cinema and own a video player. Putting these two applications together produces the rather depressing conclusion that VR will also be used for advertising. The fact that there is an, almost literally, captive audience to which information can be delivered with convincing realism (yet with no test of truthfulness) will make VR just too good to miss for the advertiser. VR which claims to give the 'experience' of driving the latest sports car (complete, no doubt, with admiring crowds who all seem to know one's name) will abound. The need for ethical, and legislative, controls on this application of VR should be obvious. There can be little doubt that VR will become widespread in training, entertainment, and advertising applications. Probably most worries over their moral implications will be expressed in the context of entertainment, rather than the other application areas. However, there is a need for moral scrutiny in all three areas. In the case of training, those designing the VR will probably have clearer objectives and may well, therefore, limit the possibilities open to the user accordingly. There is also likely to be some monitoring, even if it is just some sort of final examination, of the users of training VR. This monitoring provides at least some control over possible misuse and bad outcomes for the users of training VR. In the case of advertising there is an obvious need to prevent the user of VR becoming a

victim of intrusive advertising and to ensure that the advertising does not give a totally false impression of the value of the product to him. This may prove somewhat more difficult in the case of VR than in in television or film because of the ability of a VR designer to tailor the advertising to the needs and desires of an individual. The present day advertiser tends to work with media that are aimed at groups of potential customers. A VR designer, by contrast, can easily take account of choices made by a user and use those choices to target more effective advertising techniques at a particular individual. However, there are codes of practice in place for television and film and there is no obvious reason why these should not be immediately extended to cover VR. In the light of the sort of problem discussed in the last paragraph, improvements to these codes may well prove desirable. It is important to recognize that this constitutes no argument against using what already exists now. Most public concern is likely to be voiced with respect to the use of VR in entertainment. The following section will therefore take the entertainment applications of VR as typical, though the conclusion will attempt to draw all three application areas together. It is the entertainment application area which will be most likely to be experienced by the general public. In addition, there are important ways in which the use of VR in entertainment is likely to be less tightly controlled than in training. On the other hand, it would seem that with VR as entertainment the sky is the limit.

The Ethical Implications of VR

Doubts have been voiced about the implications of the sort of freedom that can be provided by VR. In particular, there are worries about users having the freedom to commit rape and murder within VR. Before examining such worries in detail it is worth observing that this is an ethical rather than technical issue. It is technically possible to construct VR in such a way that almost every possibility of the user's imagination can be fulfilled. It is also possible for designers to place arbitrary limits on what is possible within a particular VR. They could, for example, simply set up a VR in such a way that killing, rape, and many other morally proscribed actions are impossible. They could even set up VR in which some form of punishment (virtual or real) was the consequence of attempting to commit a proscribed action within that VR. It is also worth observing that this issue needs to resolved in the immediate future, since the existence of VR which allows killing or similar morally reprehensible acts will itself constitute a powerful argument against those want to place some sort of restriction upon VR. These are exciting, intriguing, but dangerous possibilities. When new technologies raise ethical problems in this way, misunderstandings are often generated. It is easy, for example, to assume that the novelty of the technology is, or perhaps ought to be, reflected in the novelty of the moral issues surrounding it. This is unlikely to be the case, as will be argued in the following sections. A second group of misunderstandings stems from the fact that moral principles and beliefs are seen to be in a state of flux. The new possibilities opened up by technologies such as VR may heighten anxieties about this. In a philosophical overview of this area Colin Beardon (Beardon 1992) has argued that the emergence of VR is related to a contem-

porary crisis in philosophy. There is evidence that philosophical problems are raised by VR. In particular, VR (at least in its most hyped versions) closely resembles the philosophers' notion of 'the experience machine'(Nozick 1974, Glover 1984). Beardon is correct to point out that debates about VR can aggravate cultural and philosophical splits in contemporary society, however this paper takes his conclusion that the best response to this is a pragmatic one. It is sometimes even argued that morality itself no longer has any meaning with the rise of modern secular societies. To a certain extent it is simply the case that morality has always been in a state of flux. That is to say that there is a process of general debate on moral questions which probably rarely approaches consensus. This is not the place to attempt in any way to expand on this sort of debate. Instead it will be argued that there is an immediate need to resolve certain questions about what is morally acceptable in VR. These questions can be resolved by applying familiar principles. The doubts mentioned at the start of this section are about the impact of VR on human beings and the debate is therefore easiest to resolve when seen as a continuation of similar debates about the impact of older technologies on human beings. A further group of misunderstandings surrounds questions as to who should take responsibility for discussing and resolving the moral questions surrounding new technology. The (usually unjustified) belief on the part of laymen that they are incapable of understanding the technology makes them reluctant to enter the debate. The (sometimes unjustified) belief on the part of technologists that moral questions are something they neither know nor care about makes them reluctant to start such debates. The position is further complicated by the fact that those designers of VR who set themselves high ethical standards need support, preferably from the public at large. Without this sort of support they will not be able counter the arguments of customers or managers who demand morally dubious features in VR. A less satisfactory, but more practical alternative may be to form professional organizations and draw up codes, as has been done with many other forms of technology. However this takes time and there is a certain urgency to these matters.

Virtual Reality: The Case for Restrictions

VR is a technology which can offer significant benefits in training applications. In entertainment applications, it is probable that we could feel at least as positive about VR as we do about visual art or cinema, for example. In addition there are a number of arguments related to traditional views of the the freedom of the individual. These take as central the technical claim that what happens within a VR is truly private. (At least in the case where there is only a single user). If one believes that individuals should be free to do absolutely anything which does not affect the freedom of others, then a VR would seem to be the ideal place to do such things. It may not always be completely true that others' freedoms are unaffected by what one does within VR. In a multi-user competitive VR winning will involve someone else losing, for example. It is important not to confuse this issue. What one does within a single-user VR does not directly affect others and can therefore be regarded as private. (Indirect effects will be considered in the

following.) In a multi-user VR, one can carry out actions which directly affect other users, firstly in a 'virtual' sense. The 'virtual' nature of these actions clearly reduces their moral significance, but may not completely remove it. The degree to which a 'virtual' offence is morally reprehensible depends on (among other things) its believability to the user against whom it is committed. This would seem to be an area where empirical research is needed. Not all the offences which might be committed within a VR are necessarily 'virtual' in the above sense. The nature of interaction in a multi-user VR renders physical offences 'virtual' in this sense, but there is a whole range of non-physical offences such as slander, libel, and verbal degradation which is just as 'real' when committed within a VR. This is another area where the correct moral response is not difficult, but there is some urgency in ensuring that existing provisions are extended to cover the area of VR. A more difficult set of moral problems is raised by the case of the single-user VR. If we are to deal with these problems, we need a clear account of the moral status of immoral behaviour within a VR, even when no other person is directly affected by that behaviour. The first pragmatic step toward such an account must be to once again deny the claim that the technology makes any fundamental difference to the moral problems involved. The moral questions are to be resolved solely by consideration of the effects upon human beings of the technology. The complexity or novelty of the technology is of no concern, other than its tendency to obscure these effects. In considering the moral problems raised by VR, therefore, we are considering human problems only. Technological questions are merely a fog surrounding those problems. A further step in clarification is to recall that VR has important overlaps with older technologies. Discussion of its moral implications, therefore, can draw on many existing notions. There is no need to return to fundamental moral principles in order to deal with most of the issues raised by VR. There seem to be four main arguments for the restriction of certain types of activity within VR. Those clamouring for restrictions on VR may combine these arguments in practice. A clearer picture of their validity is likely to emerge from detailed and separate examination:

1) 'They might do it for real'

This argument suggests that people who regularly perform morally reprehensible acts such as rape and murder within VR are as a consequence more likely to perform such acts in reality. This is certainly not a new departure in the discussion of ethics. In fact the counter argument to this suggestion is as least as old as the third century BC. It is based on the Aristotelian notion of catharsis (Aristotle 1968). Essentially, this counter argument claims that performing morally reprehensible acts within VR would tend to reduce the need for the user to perform such acts in reality. The question as to which of these two arguments is correct is a purely empirical one. Unfortunately, it is not clear what sort of experiment could ever resolve the issue. A high correlation, for example, between those who perform rape and murder in VR and those who do it in reality does not establish any causal link. It may be that there is a level of motivation to perform morally reprehensible acts in some individuals which even the most effective catharsis cannot assuage. A high correlation of this sort can therefore be interpreted in two completely

different ways. On the one hand it might be seen as an indication that the use of VR had delayed the real performance of the morally reprehensible act. On this view the pressure to perform such acts might sometimes become too great for the cathartic effect of VR to work. On the other hand, the correlation might be interpreted as showing that performing events in VR often leads to performing them in reality. There is little prospect of resolving this debate in a scientific fashion. However, it would be extremely foolish to dismiss this argument simply because we can see no way of testing its major claim. With many Western societies showing both a rise in civil violence and crime and an increase in the portrayal of such actions by entertainment media, there is at least the possibility of a causal link. There is also a possibility that VR might pose more of problem than previous more 'passive' media. This is because it involves physically 'practising' ,in an important sense, the morally reprehensible acts which we would not wish performed in reality. It may well be the case that some of 'behavioural conditioning' can therefore more readily be produced by VR than by previous technologies. If there is such a process, there should already be reliable, but secret, data emerging from the area of military training. Perhaps a 'peace dividend' for psychological researchers could be in the form of unrestricted access to this data. The difficulty of resolving the empirical questions should not cause us to ignore the problem. Morally speaking, it behoves scientists to commit a vast research effort to devising some way of answering these empirical questions. In the absence of such hard evidence, many people will simply assume that the answer to the empirical question must be in line with their personal prejudices. A more realistic response to this argument is that, not only do we not know at present, but we are not sure how to find out. The present state of knowledge, therefore, entails that this argument, in isolation, will not justify restrictions on VR.

2) 'Some things are not acceptable even in private'

This argument rejects the traditional claims of personal freedom. According to the proponents of this argument, one simply does not have the right to perform morally reprehensible acts, even if no-one else will be affected by them. In other words, it is the sheer unpleasantness of an individual's actions which render them morally acceptable, even if they have no consequences whatsoever. Another way in which this might be interpreted is as having a moral duty to oneself. The counter to this argument is the libertarian tradition on which most Western secular societies are based. Its classical expression is in J.S.Mill's *On Liberty*:

> The only part of the conduct of anyone, for which he is amenable to society, is that which concerns others. In the part which merely concerns himself, his independence is, of right, absolute. Over himself, over his own body and mind, the individual is sovereign.
> (Mill 1859 p.14)

The widespread influence of views similar to Mill's is likely to form the basis of opposition to any restriction on VR based on this argument. It must be concluded, therefore, that this argument alone does not justify any restriction on the use of VR. However, it is worth noting that Mill and most authorities in the libertarian tradition specifically exclude children from the claim of individual authority. This exclusion is reflected in the existing censorship provisions for media such as television and film. Even if video games have unfortunately slipped through this censorship net, there is an immediate need to extend provisions for the protection of children to the technology of VR.

3) 'People will prefer the virtual to the real'

According to one version of this argument, many people will become so entranced by VR that they will avoid the less compliant and enjoyable real world. VR will therefore become the ultimate opiate. If one believes that this will apply to a significant number of people for a significant proportion of their time, VR will be a threat to the fabric of society. Depending on just how convincing one feels the technology will eventually become, this argument gains plausibility from the possibility of spending time in a 'world' where everything can be just as one wants it - rather than the way it is. One detailed treatment of this sort of argument (Frude 1983) has been made in relation to technologies which promise less than VR. This is discussed elsewhere (Whitby 1988) and need not be repeated here. Instead it is worthwhile to continue the theme of observing that, despite the excitement and hype, VR is not the first technology to offer this possibility. In fact, in relation to this possibility, VR is simply the latest development in a tradition which goes at least as far back as prehistoric cave-painting. Escapist literature, films, plays and even unaided fantasy are all capable of distracting our attention and interest from reality. VR is a (possibly) more effective way of doing this. Thus this argument also is unconvincing in isolation. Many new technologies, in particular television, have been cast as a threat to the fabric of society. Society, however, continues more or less successfully in spite of the amount of time which many people spend watching television. There is every good reason to believe that mature people can allocate their time between entertainment and work to the detriment of neither. A further strong counter to this argument is seen when we consider its policy implications. If some people consistently find VR more attractive than 'real reality' then can it be morally correct to act so as to deny them this alternative? Surely, the morally correct course of action is to pursue ways in which 'real reality' can be made more attractive to them. If their choice is rational, then any attempt to deny them their preference represents the infliction of unnecessary suffering.

4) 'The designers of VR can signal social approval and disapproval'

This argument takes note of the fact that what people generally do within a VR may come be seen as acceptable in some sense. Thus the designers of VR have the ability to provide a degree of social approval or disapproval for the categories of actions which they allow within a VR. Recall that what is possible within a VR is determined by its designers. They have it within their power to reward morally reprehensible behaviour; to prevent morally reprehensible behaviour; or to punish it. Many examples of the rewarding of morally reprehensible behaviour are provided by the current crop of arcade games. These require and encourage a level of simulated violence which would be unacceptable in reality. The designers of these games need to examine their consciences! The marketing success of these games, particularly among the young, is a gloomy portent of what might be expected from crude home-entertainment VR. The designer of VR could just as easily prevent a user from engaging in morally reprehensible acts by simply not allowing murder, rapes, and the like to happen with the VR, or by automatically ejecting the user who attempted to commit such acts. Similarly, there would be no technical problem in providing for a suitable 'virtual punishment' for users who attempt to commit various morally reprehensible acts within VR. The moral implications of these technical possibilities are discussed in the next section. This argument seems much more convincing than the previous three. Even in the absence of evidence that users are more likely to commit morally reprehensible acts in practice, drop out of 'real' interactions, and so on, there is no doubt that legitimacy is given to actions by one's being encouraged to perform them within VR. This argument is undoubtedly the most persuasive case for restriction upon VR. However, people may consider that it is not entirely practical. The existence of, and demand for, extreme violence in computer and arcade games may seem to put irresistible pressure on VR designers to allow a similar frequency of killing and maiming. The survival of the pornography industry in spite of legal restriction, suggests that there will inevitably be virtual pornography, even if underground. These practical problems are not an effective counter to this argument, since it is important to show approval and disapproval of certain activities, even if these standards are not always attained in practice. Ultimately, much depends on the attitude of VR designers. It should be clear that they carry a burden of moral responsibility and need to ensure that they set themselves the highest possible standards.

Conclusion: Ethical Virtual Reality

Irrespective of the problems (both practical and philosophical) involved there is a need for urgent action to discuss and identify the ethical issues involved with VR. The pace of development of modern technology can frequently prove too fast for a leisurely academic development of a philosophical and moral position. The proliferation of computer games which encourage extreme violence, for example, seems to have taken place in advance

of widespread social debate¸1¸. Urgency of itself need not entail poorly thought out responses. This paper makes two, fairly simple, practical suggestions. Firstly there is the immediate need to extend the age-based censorship on media such as film and television to all forseeable interactive media, including VR. Secondly it encourages public (which might entail legislative) support of VR designers in establishing high ethical standards in their work. More might well be needed, but there are greater dangers in procrastination than in partial action. Solving the immediate problems is, of course, a beginning rather than a conclusion to debate on the ethics of VR. That is an inevitable consequence of the nature of morality. Since morality entails unconstrained choices by human agents, the idea of a code (or any similar device) removing the need to think through the moral implications of our choices is impossible. This observation applies equally to the idea that VR designers can simply outlaw immoral behaviour within their system. Since the users are effectively denied the choice of whether or not to behave in a moral fashion, their behaviour cannot therefore be described as moral. Ideally, therefore VR should allow to users to behave in ways as wicked or as saintly as are possible in reality. Ideally again, the consequences of those behaviours should be as close as possible to reality. To constrain VR users is to deny them the chance to be moral within VR. It should be noted that this neutral position is most certainly not attained by the current crop of arcade games which allow a user to indulge in extreme violence against the person without experiencing any of its adverse consequences. So far it may seem that there has been a concentration on the detrimental moral implications of VR. It is necessary to redress this by pointing out there are many positive moral implications of VR. VR can and with any luck will be used to train and explore positive moral interaction. Moral philosophers have already begun to explore the use of simulated agents in a 'virtual world' to test the rationality of moral theories and behaviour (Danielson 1992). This could be developed in fascinating ways through the technical possibilities of VR. In particular, VR could be used as a medium in which to explore the consequences of various types of behaviour. Theorists of behaviour could use this an experimental technique to refine views on exactly how and when humans behave as they do. Theorists of morality could use them in a vast development of Danielson's work to explore the consequences of partial and widespread adoption of various ethical standards. Finally, and perhaps most importantly, the facility to have a reasonable simulation of another's experiences could refine the human ability to empathize. That is to say that the use of VR to give, for example, the experience of crime from the point of view of criminal, victim, legislator, and law-enforcer might give us all a technologically-based route to far greater moral sensibility.

1. This implies no criticism of those who spoke out against 'shoot-em-up' games in the early 1980's, many of whom may be readers of this journal. I know, for I was one of them!

Acknowledgements

I am indebted to Prof. Margaret Boden, Dr. Mike Sharples and Kieran Dale for their comments on an early draft of this paper, which I have freely incorporated into this version.

References

ARISTOTLE (1968) The Poetics, in Lucas D.W.(ed)Aristotle, OUP.

BARCLAY MORGAN, A. (1991) Interactivity: From Sound to Motion to Narrative, in Art Papers, Vol. 15 No.55. Sept./Oct. 1991.

BEARDON C. (1992) The Ethics of Virtual Reality, Intelligent Tutoring Media, Vol. 3, No. 1. pp.23-27.

DANIELSON, P. (1992) Artificial Morality, Routledge, London.

FRUDE, N. (1983) The Intimate Machine, New American Library, New York.

GLOVER, J. (1984) What Sort of People Should There Be?, Penguin, Harmondsworth, pp.92-113.

MILL, J.S., (1859) On Liberty, reprinted in John Stuart Mill A Selection of his Works, Robson J.M. (ed) (1966) Macmillan, Toronto, p14.

NOZICK, R. (1974) Anarchy, State and Utopia, Basic Books, New York, pp.42-45.

PAPERT, S. (1980) Mindstorms: Children , Computers, and Powerful Ideas, Basic Books, New York.

WHITBY, B. (1988) Artificial Intelligence: A Handbook of Professionalism, Ellis Horwood, Chichester.