Synthetic Experience: A Taxonomy, Survey of Earlier Thought, and Speculations on the Future

Warren Robinett
Department of Computer Science
University of North Carolina, Chapel Hill, NC 27599-3175

Abstract

A taxonomy is proposed to classify all varieties of technologically-mediated experience. This includes virtual reality and teleoperation, as well as earlier devices such as the microscope and telephone. The model of mediated interaction assumes a sensor-display link from the world to the human, and an action-actuator link going back from the human to the world, with the mediating technology transforming the transmitted experience in some way. The taxonomy is used to classify a number of sample systems.

Earlier visions of technologically-based experience are compared with the ideas presented in this paper. Then the long-term prospects of this field are speculated upon, ignoring constraints of cost, effort, or time to develop. Finally, the ultimate limits of synthetic experience are discussed, derived from properties of the physical universe and the human neural apparatus.

1. Introduction

The Head-Mounted Display (HMD) has been used in two distinctly different kinds of applications: teleoperation, in which a human operator's senses are projected into a remote robot body, and virtual environments, in which the human can move through and interact with a three-dimensional computer-generated virtual world. New uses for the HMD are currently being discovered, such as flight simulation, night vision goggles, micro-teleoperation, and augmented reality.

This paper proposes a taxonomy for classifying systems which incorporate an HMD. Systems are classified according to eight independent dimensions, each of which can take on a number of discrete values. The domain of this classification method is broad enough to include technological precursors to the HMD, such as the telescope, microscope, television, and telephone.

This taxonomy attempts to impose some sense onto a very broad and very new area which is pregnant with unexplored possibilities. The method of making sense is to cleave the set of possible systems into a small number of disjoint sets by imposing distinctions. Each of these distinctions corresponds to a dimension of the taxonomy. For example, the dimension called "causality" distinguishes between teleoperation, in which the operator's actions affect the real world, and virtual environments, in which the operator's actions affect only a simulated world. Other dimensions have to do with the sensory modalities used by the systems (vision, hearing, touch, and others), the nature of the representations (or models) of the environments surrounding the user, and displacements or scaling in time or space between the user's true position and the environment with which the user interacts.

The attempt to classify, or even talk about, devices which produce reproductions of sensory experience immediately brings up many difficult issues in philosophy, psychology, and other fields. What is experience? What is reality? What is a representation of the world or of an object? Is perfect reproduction of human sensory experience possible? Many more questions and issues of this sort could be listed here. Rather than be scared off by these difficult and complex issues, I have tried to mention the issues that I think are relevant to the discussion, and leave it to others to correct and clarify errors and omissions, if they desire. I hope this taxonomy can serve as a point of departure for us collectively to understand and develop Head-Mounted Displays into useful and widely-used tools.

The discussion in this paper is somewhat biased by my own experience in designing computer-simulated virtual worlds, and my lack of detailed knowledge of the work done over the last several decades in teleoperation, and perhaps other related fields. All of us in this diverse field have our specialties and our blind spots. My experience tells me that the distinctions that I put forward in this paper are important ones, and rather than waiting until I achieve broad knowledge of all the fields touching on virtual worlds, I put the ideas forward now to serve as a starting point for discussion.

The common theme of all these devices and systems is *technologically-mediated experience*. The older systems use optics or analog electronics to mediate and transform the user's experience, whereas the more recent systems rely heavily on computers and digital electronics. In both cases, the general model of technologically-mediated experience is the same, as shown in Figure 1.

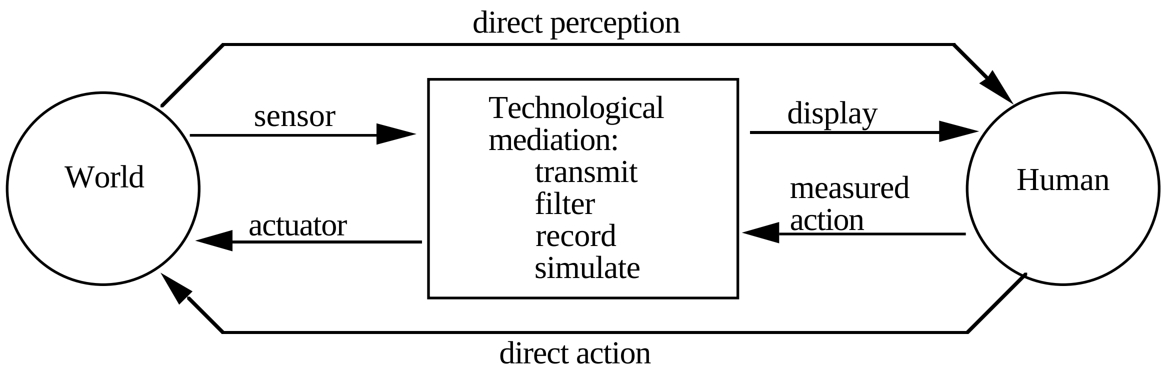

Figure 1. Technologically-mediated experience.

The new devices incorporating the HMD did not come out of nowhere, but are extensions and refinements of earlier devices and media. Media began to evolve thousands of years ago when prehistoric man created visual representations of the world using paint. Painting was followed by the telescope, the microscope, photography, the phonograph, the telephone, film, television, and video games. Each of these devices derives its usefulness from being able to modify, record, or transmit some aspect of human sensory experience. For each of these devices, sensory experience is captured, processed, and then displayed to a human user.

The HMD is one further step along this evolutionary path. It improves over earlier visual media in being able to give the user a perception of a surrounding three-dimensional space, rather than just a look into a space from a fixed viewpoint. It is not simply a visual display technique, but rather a multi-sensory display technique involving the visual, the vestibular, and the proprioceptive systems, in which the visuals depicting the surrounding 3D virtual world are generated so as to match the user's voluntary head movements.

We offer some definitions:

natural experience:
directly perceiving the properties or behavior of something physically present before the perceiver.

synthetic experience:
perceiving a representation or simulacrum of something physically real rather than the thing itself.

The dictionary definitions (Webster's Ninth New Collegiate Dictionary) help to clarify the term "synthetic experience":

experience 1a: direct observation of or participation in events as a basis of knowledge.
3a: the conscious events that make up an individual's life.
5: the act or process of directly perceiving events or reality.

synthetic 4b: devised, arranged, or fabricated for special situations to imitate or replace usual realities.
5: something resulting from synthesis rather than occurring naturally.

synthesis 1a: the composition or combination of parts or elements so as to form a whole.

The term *synthetic experience* encompasses virtual environments, teleoperation, other uses of the HMD, film, the telephone, video games, and most earlier media. It is meant to be synonymous with the term *technologically-mediated experience,* used earlier. We limit the scope of synthetic experience to reproductions of sensory experience. We exclude verbal descriptions such as novels and oral story-telling.

We also exclude theater from this classification system, though there is clearly a common thread running from story-telling to theater to film. That theater is in one sense a natural experience of watching human actors and at the same time a recreation of a (hypothetical) earlier action shows that it may be difficult to draw a distinct boundary between natural and synthetic experience. In the broadest sense, a student's mimicking of a tennis pro's serve is a reproduction of an earlier action, and role-playing in group therapy is a simulated experience. However, we limit the scope of synthetic experience to technologically-mediated reproductions of sensory experience.

1.1 Examples of synthetic experience

Some of the most important current applications of the HMD are landmarks that help to map out the scope of synthetic experience. These examples are meant to illustrate the breadth of synthetic experience and are not meant to be definitions.

Virtual reality uses a stereoscopic, wide-angle HMD to create the illusion of a three-dimensional surrounding fantasy world, a 3D video game that allows one or more players to get inside and interact with one another.

Flight simulation also defines a simulated 3D world in which actions have effects, but in this case the simulation is intended to accurately model the behavior of a real aircraft so as to give the pilot experience in dangerous situations without mistakes being fatal.

Teleoperation uses a HMD and force-feedback handgrip that are electronically linked to a distant robot body with a robot arm and a pair of video cameras on the robot head. The robot head turns to mimic the operator's head motions and the robot arm mimics hand motions, so that the operator's eyes and hands are effectively projected into the remote environment, and the operator can look around and do things through the robot body. The remote environment may be a dangerous one, such as the bottom of the ocean, inside a nuclear power plant, or in space.

Micro-teleoperation replaces the human-scale anthropomorphic robot of ordinary teleoperation with a microscope and micromanipulator so as to give the operator the sense of presence and the ability to act in the microscopic environment. The Scanning-Tunneling Microscope (STM) is well-suited to micro-teleoperation since it uses a tiny probe scanned over the sample surface to capture a 3D image of the surface (at atomic resolution), and the probe tip can also be used as a micromanipulator to interact with the sample material (Robinett, Taylor, Chi, Wright, Brooks, Williams & Snyder, 1992).

Telecommunication is familiar to us through daily use of the telephone, and video teleconferencing extends this remote communication with other human beings to include the sense of sight, and to allow communication among groups of people rather than just two at a time. The operator of a tele-controlled robot is able to speak and listen to another human being in front of the robot.

Technological masquerade has been used to study intra-species communication in animals by using recordings and sophisticated puppets to fool the animals into behaving as if they were interacting with another member of their species. This has been done extensively with recorded bird calls (Bright, 1984), and also with a computer-controlled robot bee, which was able to direct real bees to specific locations far from the hive by moving in the patterned "dance" that bees use to communicate and then dispensing a sweet liquid "sample" of the (pretended) distant pollen (Weiss, 1989). The connection with synthetic experience is that a human could potentially teleoperate a robot bee to attempt to communicate in real-time with real bees.

Augmented reality uses a see-through HMD, in which half-silvered mirrors allow the user to see through directly to the real world and at the same time spatially superimpose the virtual world on top of the real world. The superimposed virtual world may be labels or diagrams located at specific points in the real world (Caudell & Mizell, 1992). It may also be information derived from sensors which is superimposed onto the user's direct view of the real world, as with helicopter pilots flying at night through canyons using Forward-Looking Infrared (FLIR) sensors, who also have a direct view out into the darkness in case there is anything bright enough to see.

A *synthetic sense* is created when a sensor for a phenomenon that is imperceptible to human senses is linked to a display device. This gives the user the ability to perceive phenomena that are invisible, silent, and intangible without technological augmentation. Night vision goggles are an example of this. Another example, currently being prototyped, is medical "X-ray vision," in which an obstetrician uses a see-through HMD to view data from a hand-held ultrasound scanner. The doctor can see and touch the abdomen of a pregnant woman and sees the data from the ultrasound scanner superimposed at the location from which it came, giving the perception of seeing into the living tissue (Robinett, 1991a).
A *sensory prosthesis* corrects, amplifies, or otherwise improves the fidelity of an ordinary "built-in" human sense. Examples are corrective spectacles, sunglasses, and hearing aids. For people with defective or non-functional senses, *cross-sensory substitution* can compensate for the disability. For example, the Opticon (Linvill, 1973) is an optical-to-tactile transducer array which allows blind people, after some training, to read from ordinary printed books by, in effect, running their finger over the printed text and feeling the black marks as raised bumps.

1.2 On taxonomies in general

Taxonomy can help to think about a group of ideas. A taxonomy characterizes ideas along several more or less independent axis so as to form a matrix. The taxonomy has two important parts, the axes and the entries in the matrix.

The items listed on each axis are related by whatever characteristic is represented by that axis. If the variable chosen is continuous, the axis will represent a spectrum of values which may be more or less well defined. If the variable chosen has only a few fixed values, then those values divide the space of all ideas considered into categories. For example, a taxonomy of people might include eye color as a variable, characterizing eye color into a small number of common values, for example, blue, brown, hazel, etc. The several axes of the taxonomy divide the intellectual space of interest into a number of boxes, each of which can be examined in turn.

The boxes in a taxonomy are used to enumerate ideas that have common characteristics in the axes chosen. One can list in each such box a set of ideas that are related to each other by their position on the axis.

The number of boxes in a taxonomy is its volume, that is, the product of the number of values permitted to each variable. The most useful taxonomies generally have fewer than 100 boxes. With more boxes it is difficult to think about each one separately. Some useful taxonomies have only a few boxes; as few as four are sometimes useful. The desire to limit the number of boxes limits the number of values permitted to each variable. Of course if there are more variables, each must be permitted fewer values. A useful discrete variable usually provides 2 to 5 values; a useful continuous variable generally approximates continuity with small, medium, and large.

There are two parts of building a taxonomy that make building it interesting. First there is the selection of axes and values along them. In forming the axes, one is forced to think of orthogonal dimensions. Which variables are related and which are independent. It often happens that as the taxonomy takes shape one finds that the variables chosen are not really independent.

The second interesting part of building a taxonomy is filling in the boxes. Here the really interesting results come out by examining combinations of variables that may hitherto not have been considered together.

2. Dimensions of synthetic experience

The proposed taxonomy for classifying types of synthetic experience is shown below in Table 1. The eight dimensions of the classification system are largely independent of one another, so the space of all possible types of synthetic experience should be conceived as a matrix (with eight dimensions) rather than a hierarchy. The first four dimensions describe the basic nature of the technological mediation in a synthetic experience device, whereas the last four dimensions have to do with which sensory channels and motor channels are employed.

Dimension	Possibilities	Contrasting examples
causality	simulated recorded transmitted	flight simulator film teleoperation
model source	scanned constructed computed edited	night vision goggles video game computational fluid dynamics film
time	1-to-1 accelerated (or retarded) frozen distorted	film time-lapse photography photograph edited video recording of event
space	registered remote miniaturized (or enlarged) distorted	night vision goggles teleoperation micro-teleoperation (STM) STM with heights exaggerated
superposition	merged isolated	augmented reality virtual reality
display type	HMD screen speaker (many more -- see Table 3)	virtual reality video game recorded music
sensor type	photomultiplier STM ultrasound scanner (many more -- see Table 4)	night vision goggles micro-teleoperation medical “X-ray vision”
action measurement type	tracker & glove joystick force feedback arm (many more -- see Table 5)	virtual reality video game teleoperation
actuator type	robot arm STM tip aircraft flaps (many more -- see Table 6)	teleoperation micro-teleoperation remote piloted aircraft

Table 1. Classification system for types of synthetic experience

2.1 Causality

The first dimension of the classification system, *causality*, makes the most fundamental distinctions among types of synthetic experience. The three possibilities are to *transmit*, *record*, or *simulate* experience. These three categories correspond to the way that we experience the world -- not only do we experience the present, but we also remember the past and imagine the future. Replaying a recording has similarities with remembering: it is re-experiencing past events. Participating in an interactive simulation has similarities with imagining; it is trying out courses of action on an imaginary stage, perhaps to see what the consequences might be. Engaging in real-time transmitted experience through, for example, a teleoperator system, has similarities with normal active experience in the present:; the operator's actions affect the world (Robinett, 1991b).

The effect of voluntary actions is different in each of these these cases. In a simulated virtual world (for example, in a flight simulator), actions have effects within that simulated world but not in the real world.

(There is no plane to crash; nobody will die.) In a virtual world which is a real-time reproduction of some part of the real world (for example, a pilot flying a remote-piloted aircraft), actions do affect the real world. (The plane can crash and burn.) In a recording of past events (for example, the “black box” recording of what happened in an airliner crash), what happened was recorded, and actions by the user cannot change what happened. This dimension is called “causality” because, for the three cases of simulated, transmitted and recorded experience, actions by the user cause either effects in the simulated world, effects in the real world, or no effects at all.

This dimension might possibly have been called “time,” since recordings replay past events, transmitted experience takes place in real-time in the present, and simulations sometimes are used to predict future events. However, simulations are not necessarily of the future (for example, a simulation of continental drift), so it is best to name this dimension “causality” to capture the real differences between transmitted, recorded, and simulated experience.

Figure 2 shows diagrams of the primary data flow for transmitted, recorded, and simulated experience, with each shown as a special case of the diagram for technologically-mediated experience in Figure 1.

In the case of transmitted experience, the diagram of Figure 1 is a good model for the data flow -- the user observes the world through the sensor-display data path and performs actions that affect the world through the action-actuator data path. In the case of recorded experience, the sensor data is stored in some kind of memory device (such as magnetic tape), and at a later time this data is replayed through the display to the user. An actuator is not needed in this activity, and user actions are only needed to control the replay process itself. In the case of simulated experience, the primary data path is from the measured actions of the user, through the simulation, and back through the display to the user. Again, the actuator is not needed, and the sensor channel is needed only if the simulated virtual world is based at least partly on scanned-in data from the real world.

The fourth diagram shown in Figure 2 is a variety of transmitted experience, with a data path introduced to allow autonomous actions by a telerobot, under supervision of a human operator. An operator could alternate between passive real-time observation of the telerobot’s actions and taking direct control of the telerobot’s actions as in normal transmitted experience.

In a system in which all of these data paths are present, all of these modes of operation are possible. In the UNC Nanomanipulator Project (Robinett, Taylor, Chi, Wright, Brooks, Williams & Snyder, 1992), for example, in which a HMD and force-feedback arm control a Scanning Tunneling Microscope (STM), transmitted experience, recorded experience, simulated experience, and supervisory control are all possible. The user may directly control the STM tip through the force-feedback arm and modify the sample surface (transmitted experience). The user may record a snapshot or sequence of images of the surface and view them through the HMD at a later time (recorded experience). The user may manipulate simulated molecules through the force-feedback arm and HMD with no connection to the microscope (simulated experience). We plan later to allow the user to initiate algorithmically-controlled modifications of the sample surface, with the possibility of intervening (transmitted experience with supervisory control).

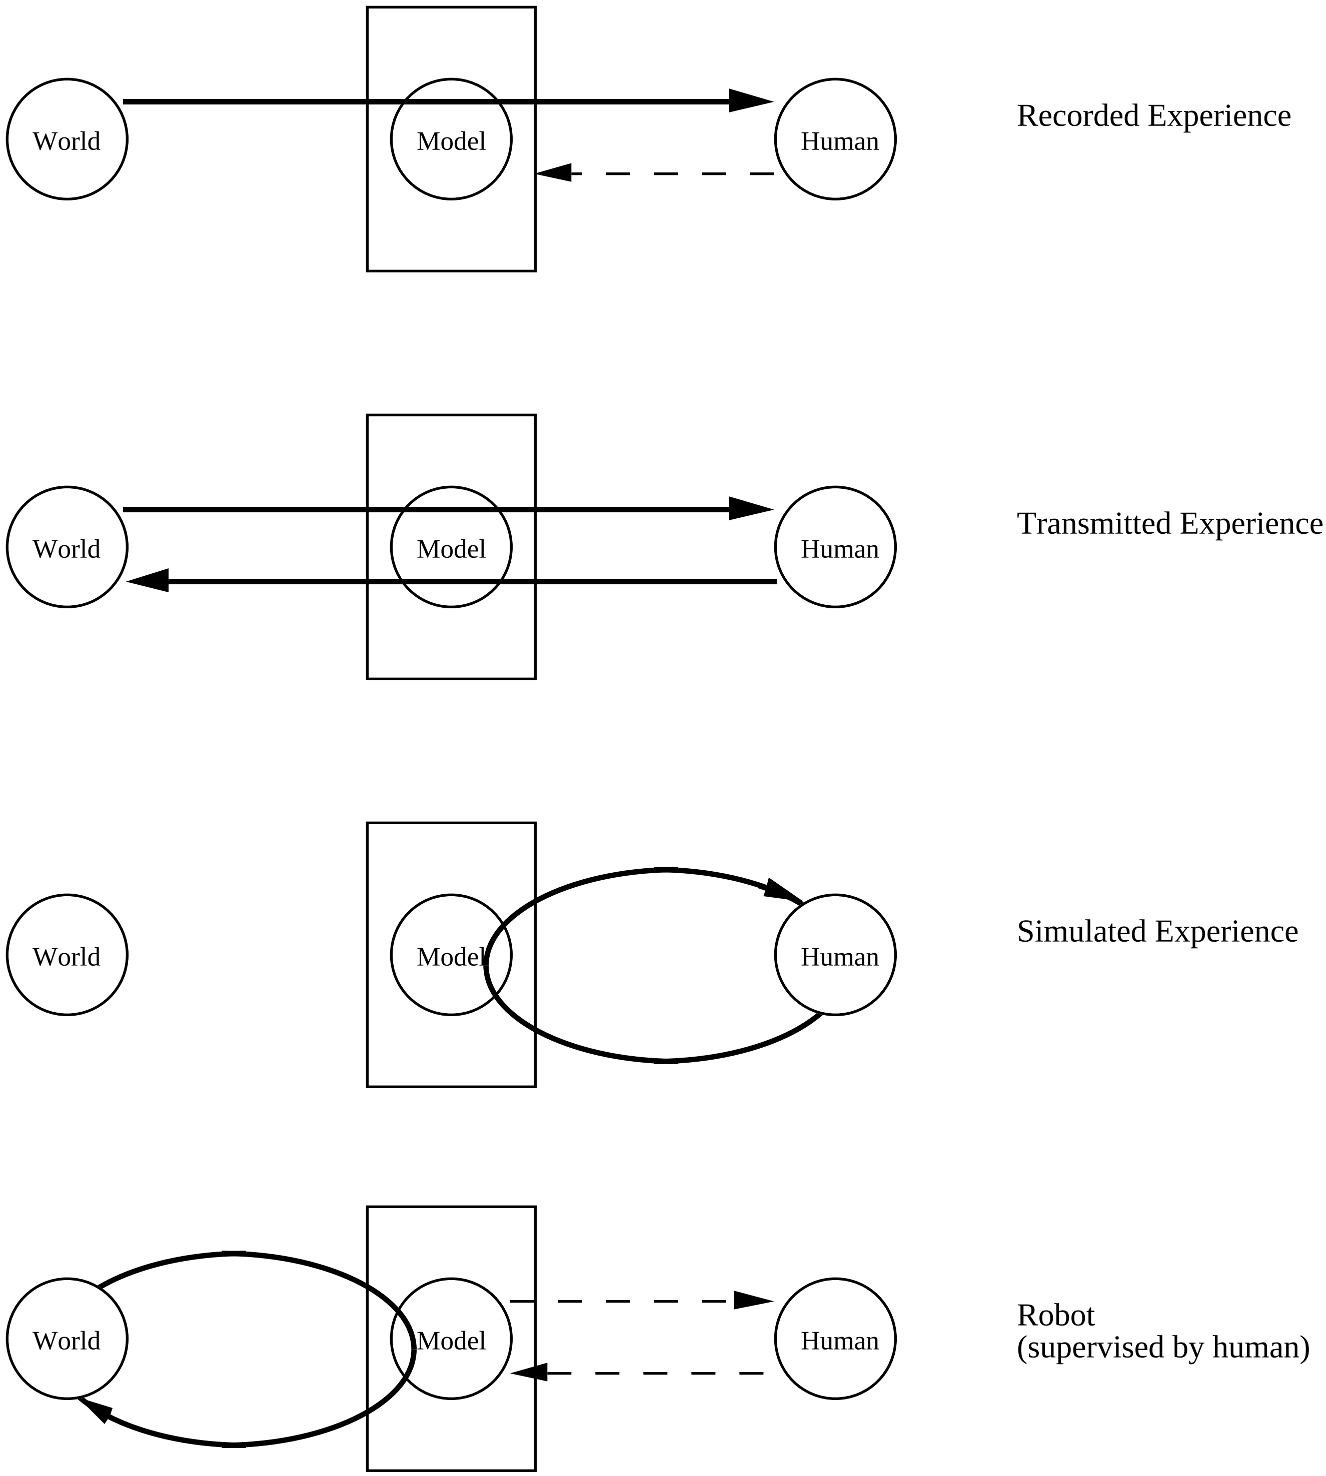

Figure 2. Data flow for types of mediated experience.

2.2 Model source

In a synthetic experience, the human user perceives a virtual world which is defined by a (possibly changing) database called a *model*.. This model is stored, at least transitorily, in some kind of memory device. The model defines what the virtual world looks like, sounds like, and feels like, according to which display devices are available.

There are three main sources for this model data. A sensor can *scan* the real world to produce a model for later display to the user, a human artist or craftsman can laboriously *construct* a model, piece by piece, or a dynamic model can be *computed* on the fly by a computational model. Examples of these three cases are live television, with the world scanned by the video camera, Disney-style animated cartoons, with each animation frame drawn by an artist, and computational fluid dynamics, where the simulation code generates new model data as needed. However, these cases are not exclusive, and a scanned-in model can be chopped up and *edited* to construct a model that is partly based on the real world but is different. A good example of this is film, in which raw footage from the initial shooting is heavily edited and some animated special effects are thrown in to produce the final movie.

2.3 Time and space

For data scanned in from the real world, in some cases (such as night vision goggles), the data will be displayed in exactly the location from which it was derived, whereas in other cases (such as teleoperation), the scan space is displaced from the display space. The display space may also differ in scale from the scan space, as in micro-teleoperation. The mapping from scan space to display space may include a spatial deformation.

Furthermore, the scan and display may be aligned or displaced in time (transmitted experience versus recorded experience). Scan time and display time may also differ in time-scale, as with time-lapse photography. Display time could be related to scan time by a non-linear distortion mapping, for example, in the replay of an explosion where initial events occur more rapidly than later events.

These possibilities may be summed up by saying that, for both time and space, the scan and display may be either aligned, displaced, differ in scale, or be related by a distortion mapping, as shown in Table 2. The relationship of Table 2 to the overall taxonomy of Table 1 is that Table 2 emphasizes the similarity of the values which can be assigned to the two dimensions "time" and "space" of the taxonomy.

	Time	Space
Aligned	transmitted in real-time, 1-to-1 time scale *live television*	registered, 1-to-1 scale *night vision goggles*
Displaced	recorded earlier, 1-to-1 time scale *TV rebroadcast of live event*	remote, 1-to-1 scale *teleoperation*
Scaled	recorded earlier, accelerated (or retarded) time *slow motion instant replay on TV*	remote, expanded (or miniaturized) *micro-teleoperation*
Distorted	recorded earlier, distorted time *TV event with dull parts edited out*	remote, distorted space *micro-teleop, exaggerated height*

Table 2. Relative scale and displacement in time and space for sensor and display.

Since this is a comparison of the time and space coordinates of the sensor and display, this comparison only makes sense when there is a sensor involved. A constructed model comes out of nothingness and therefore has no real world coordinates with which the display might align. Likewise, an edited model may have pieces that come from specific locations in the real world, but there is no way to match the whole of the model to the real world.

If we imagine two clocks displaying Greenwich Mean Time, one being scanned by the sensor and the other with the user beside the display, we may ask if the two clocks are displaying the same time and if they are running at the same rate. For transmitted experience, the two clocks must match from moment to moment, so 1-to-1 time scale is required. In replaying a recording, the clocks might run at the same rate but display different times. However, as with a VCR, the recording might also be played back in slow motion, faster than normal, in reverse, or paused with the action frozen. All these varieties of time-progression are possible for any recording technique.

Since the distinction between transmitted and recorded experience is already covered by the causality dimension of the classification system, the time dimension of the classification system focuses on time-scale, with the possibilities being *1-to-1* time-scale, *accelerated (or retarded)* time, *frozen* time and *distorted* time.

In the same way that we used two clocks to judge the time-offset and time-scale differences between scan time and display time, we may also use two spatial markers to judge the offset and scale difference between scan space and display space. For this we use a pair of three-dimensional coordinate axes, one being scanned by the sensor and the other measuring the space the user occupies. For concreteness, let us imagine that the user wears a see-through HMD which uses half-silvered mirrors to spatially superimpose the real and virtual worlds. In this case, the coordinate axes that are actually present in front of the user can be seen by the user with image of the scan-space coordinate axes optically superimposed. Both coordinate axes are ruled in centimeters. We can now ask whether the two axes are aligned or displaced from one another, and whether they appear to be the same or different sizes.

The main possibilities are that, relative to scan space, display space is *registered*, *displaced*, or *expanded (or miniaturized)*. It is also possible to introduce various *distortion* mappings between scan space and display space. An example is in micro-teleoperation using the STM, where we wish to exaggerate the height variations of the sample being scanned so as to make very slight height steps more obvious.

2.4 Superposition

A virtual world may be *merged*, perhaps using half-silvered mirrors, with the surrounding real world. It may be convenient to combine the real and virtual worlds by using a video camera to capture an image of the real world and then doing a more sophisticated merge than is possible with optics. This example of using cameras to capture and merge the real world with a virtual one shows that the surrounding real world itself may be thought of as model on an equal footing with the virtual world model, and the two models may be edited together if desired. A powerful technique is to spatially superimpose two models of the same region of space, creating a sort of three-dimensional Rosetta Stone through the spatial correspondence of pairs of matching points in the two models.

Data from multiple sensors may be fused into a single virtual world model. An example is a conference telephone call. This may be thought of as automatic real-time editing, in which the data from three or more sensors (microphones) is integrated and then displayed through the speaker in each user's handset.

On the other hand, a display, particularly an HMD, may block out the real world and *isolate* the user within the virtual world.

2.5. Senses and sensors

Display type and sensor type are two of the dimensions of the classification system, and a few examples are given in Table 1 where the classification system is defined. However, the complete range of displays can in principle cover every phenomenon that human beings possess sensory organs to detect. Likewise, the complete range of sensors encompasses all measurable or detectable phenomena. Table 3 lists human sensory channels and corresponding display devices, and Table 4 lists some sensors and the phenomena to which they are sensitive. Neither the list of display devices nor the list of sensors is exhaustive.

Sensed phenomena	Sensory system	Display device
visible light (400 - 700 nm)	vision	display screen (CRT, LCD, or other) head-mounted display (HMD) individual lights, dials and gauges
vibrations in the air (20 Hz - 20 kHz)	hearing	speaker, headphones headphones with spatialized sound
force vibration surface texture temperature	“sense of touch”	force-feedback device buzzer tactor array, air bladders heater, cooler, fan
chemical composition of air	smell	Sensorama smell display
chemical composition of food	taste	
acceleration of body	vestibular system	motion platform
limb & body position	proprioception	exoskeleton with forced movements
internal state of body (hunger, thirst, fatigue, etc.)	interoceptors	intravenous medical device to monitor and control contents of bloodstream
damage to body	pain	

Table 3. Human sensory channels and display devices

Sensed phenomena	Sensor
visible light	still camera
visible light	video camera
sound	microphone
position of moving objects	radar
distance to object	range-finder
position & orientation of moving object	tracker (6 DOF)
inside of human body	ultrasound scanner computer-aided tomography (CAT) scan nuclear magnetic resonance (NMR) imaging
infrared light	night vision goggles
ultraviolet light	UV detector
Xrays	fluoroscope
magnetism	electronic compass
radiation	Geiger counter
3D surface shape	3D laser scanner
3D topography of the earth	aerial photography & photogrammetry
3D surface of microscopic sample	scanning-tunneling microscope
image of distant object	telescope
chemical composition	gas chromatograph, spectrograph
movement & vibration	accelerometer gyroscope
gravitation field variations	mass detector

Table 4. Sensors and what they sense

Many of the sensors listed detect phenomena that are imperceptible to human senses, and by linking such sensors to display devices, these imperceptible phenomena can be rendered visible, audible, touchable, or otherwise perceptible to a human being. A sensor-display linkage of this sort creates a synthetic sense, an apparatus that extends human perception and awareness (Robinett, 1991a).

Using a HMD as a display device offers the possibility of mapping sensed phenomena to specific locations relative to the body of the user. This is what vision and hearing do, and for those sensors which are able to establish the location or direction of the phenomena they sense, this positional information can be interpreted through the visual (and auditory) channels of the HMD to depict the sensor data as emanating from specific locations in space.

There are a considerable number of imperceptible phenomena. Every one of them can be given a visible form, or sound, or tactile representation. Every detectable phenomenon can be given a perceptible representation, regardless of its remoteness in space, time, scale, or time-scale, and regardless of what form of energy or matter is being detected. By linking sensors and displays to create synthetic senses, every phenomenon that exists can be rendered directly perceptible to the human senses.

What do these imperceptible things look like? Since they are imperceptible, they don't look like anything. A representation must be invented, and choices present themselves. In general, many representations are possible for a given phenomenon, and different representations may be useful at different times. For example, a number of visual representations of molecules are used in different situations: touching sphere model, ball and stick model, solvent-accessible surface model, ribbon following the backbone of a protein. What these invisible things should look like is a graphic design problem that, in time, can be expected to settle out on the basis of informativeness, aesthetics, convention, and accidents of history.

2.6. Actions and actuators

In the same way that sensors and human senses can be linked to cover all detectable phenomena, a linkage from manual and other input devices to actuators should be able to control any device or system designed to be controlled. This is a relatively unexplored area, with the main work so far having been done in teleoperation and remote piloted vehicles. Most other human tools, vehicles, environments, and instruments each have their own idiosyncratic locally-operated control panels.

In a few years visual telepresence may be widely available, so that a person can move by virtual travel instantly to distant locations, just as is now possible with the telephone for hearing only. If, at that time, most controllable devices are linked to the communications network, then it will be possible for a person to project by virtual travel to a distant location and initiate actions there through the actuators available at that site. For safety and security reasons, remote access will probably not be allowed for some types of devices, but for many devices it may make sense. Another issue is who has permission to control which devices. In spite of these probable limitations, we can still imagine a future world in which an enormous traffic of ghostly presences leap about the planet, manipulating distant parts of the world through briefly occupied robot bodies.

Table 5 lists human motor channels and some devices available to measure human actions. Table 6 lists some actuators that have so far been used in remote presence systems, and, of course, there are many more devices and systems that could potentially be controlled over the communication network.

Motor channel	Behavior measurement device
hands	hand-tracker (6 DOF) hand-held pushbuttons instrumented glove keyboard, mouse, joystick
feet	foot-pedal
eyes (gaze-direction, blinking)	gaze-tracker
head position	head-tracker (6 DOF)
body posture	instrumented body-suit
voice	speech recognition
breath	breath controller
heartbeat	EKG machine

Table 5. Human motor channels and measurement devices

Actuator
robot arm
STM tip
remote piloted vehicle

Table 6. Actuators used in teleoperated systems

2.7. Classification of some specific systems

Table 7 below shows the synthetic experience types of a number of specific systems and devices.

	Causality	Model source	Time	Space	Super-position	Display	Sensor	Action measure	Actuator
Tele-operation	transmit	scan	1-to-1	remote	isolated	HMD, force feed-back arm	camera on robot head	force feed-back arm	robot arm
Micro-tele-operation	transmit	scan	1-to-1	expanded	isolated	HMD, force feed-back arm	STM	force feed-back arm	STM tip
Remote piloted aircraft	transmit	scan	1-to-1	remote	isolated	screen	video camera	joystick	flap actuators in aircarft
Flight simulation	simulate	computed	1-to-1	remote	isolated	HMD, motion base	satellite photo-graphy	cockpit controls	---
Virtual reality game	simulate	construct	---	---	isolated	HMD	---	tracker, glove	---
Video game	simulate	construct	---	---	isolated	screen, speaker	---	joystick	---
Augmented reality -- helicopter	transmit	scan	1-to-1	registered	merge	HMD	FLIR	---	---
Night vision goggles	transmit	scan	1-to-1	registered	merge	HMD	photo-multiplier	---	---
Medical "X-ray vision"	transmit	scan	1-to-1	registered	merge	HMD	ultrasound scanner	hand tracker	---
Telephone	transmit	scan	1-to-1	remote	merge	speaker	micro-phone	keypad	---
Live television	transmit	scan	1-to-1	remote	---	screen, speaker	camera, micro-phone	---	---
Film	record	edit	1-to-1	remote	---	screen, speaker	camera, micro-phone	---	---
Video cassette recorder	record	edit	1-to-1, fast, slow, frozen	remote	---	screen, speaker	camera, micro-phone	keypad	---
Time-lapse photo-graphy	record	scan	accelerated	remote	---	screen	camera	---	---
Photo-graphy	record	scan	frozen	remote	---	print	still camera	---	---
Painting	record	construct	frozen	---	---	canvas	---	---	---

Table 7. Examples of how specific systems are classified

Comparing different lines in the table suggests variations and extensions for some of the systems. For example, a hybrid of film and virtual reality would give us 3D recording of earlier actions that the user could fly through to observe from any viewpoint.

A hybrid of micro-teleoperation with the STM and the video cassette recorder would allow rapid events occuring at the microscopic scale to be rapidly scanned as they occur, and then played back at a slower speed later, pausing and backing up to observe interesting events.

The dimensions of the classification system are largely independent of one another, so it is possible to go into the line for a given system and ask what kind of system would result by changing it along one dimension. The main dependencies between the dimensions are that transmitted experience requires 1-to-1 time scale, recorded experience needs no actuator or action measurement (except to control the replay), and simulated experience needs no actuator. Also, the dimensions of sensor-to-display relative time and space only apply to models scanned in by sensors from the real world.

3. Comparisons with earlier visions of technologically-based experience

In this section, we compare the idea of synthetic experience proposed in this paper with the visions of earlier thinkers who have explored the use of technology to augment human perceptual powers. We have attempted to put these thinkers into chronological order to reflect the development of ideas in the field, but there are overlaps because each person's work spans many years. We have attempted to assign a starting year to each person based on either the publication of a paper or the completion of system which was relevant to synthetic experience.

Because the proposed taxonomy spans such a broad intellectual territory, it is difficult to give a complete survey of all the thinkers in all the fields which contributed to synthetic experience. The selection of people and ideas described in this section is biased by my own experience and familiarity with computer-based simulations of virtual worlds. The field of teleoperation developed pretty much independently of the computer graphics field, and the first practical uses of HMDs were for teleoperation. Because of limited knowledge of the field of teleoperation, I can only give a few scattered references (Goertz, 1961; Goertz, 1965; Johnsen & Corliss, 1967; Corliss & Johnsen, 1968; Sheridan, 1992).

There are psychologists, such as J. J. Gibson, whose work was important to this field, and likewise some philosophers who did important groundwork. These fields are not covered either in this section.

Vannevar Bush, 1945

In his article "As We May Think," Vannevar Bush, who headed the American scientific effort during World War II, proposed the "memex," a storage and access device for human knowledge (Bush, 1945). The idea was basically to create an automated library, containing text, mathematical models, and images. By automating formerly manual functions, he foresaw keyword searches, easily followed links among related items (what we now call hypertext), assistance with the mechanics of symbolic manipulation in mathematics, optical character recognition, and speech recognition. His vision was somewhat colored by the technology which he proposed to use: photography, microfilm, electronics, and photocells.

Bush also proposed that people would capture images at will using a small camera mounted on the forehead. He mentioned photographing the invisible and visualizing microsecond-scale events on an oscilloscope. He finished by mentioning the possibility of direct neural connection to machines.

Mort Heilig, 1955

In his article "Cinema of the Future," Mort Heilig, a film-maker, proposed creating multi-sensory films that covered the entire human sensorium (Heilig, 1955). He asserted that the relative importance of the senses of vision, hearing, smell, touch, and taste were in the proportion 70%, 20%, 5%, 4%, and 1%. He contrasted the field of view, acuity, and stereoscopic vision of natural human experience with the limits of the cinema of the day. He saw the goal of film as the reproduction of experience, with new less literal art forms following in the wake of successful reproduction, just as abstract painting followed realistic painting and Disney-style animated cartoons followed motion picture photography.

Heilig considered mainly recorded experience but did mention transmitted experience (television). He also touched on the ability to expand the senses with the microscope, slow-motion photography, and other such devices.

Heilig built a prototype multi-sensory film machine, the Sensorama, illustrating his ideas. It used a multi-track photographic and magnetic film to record the experiences to be replayed, which were displayed to a single user through a color stereoscopic viewer, stereophonic headphones, vibrating seat and handlebars, wind generated by a fan, and a smell display. In the smell display, wax pellets impregnated with the odors for a given film were placed in small chambers which could be opened under control of the smell-track on the film, with air forced through the chamber to emerge near the user's nose.

Doug Englebart, 1963

In his paper "A Conceptual Framework for the Augmentation of Man's Intellect," Doug Englebart proposed that computers could be used, not just for calculation, but to increase human understanding and problem-solving ability (Englebart, 1963). He foresaw this as being an interactive process, involving computer-controlled changing images on cathode-ray tubes and controlled by the human with interactive manual input devices. He described the key elements of augmentation as a trained human, with artifacts, language, and methodology. "Artifacts" referred to the display and input devices. He foresaw in increase in intellectual power as humans moved successively from mental symbol manipulation (language), to manual external symbol manipulation (writing and mathematics), and finally to automatic external symbol manipulation.

Under his leadership, a prototype system was built in the late Sixties at Stanford Research Institute. In this laboratory, the first prototypes were created for the mouse, chord keyboard, word processing, electronic mail, teleconferencing, and hypertext.

Ivan Sutherland, 1965

In his paper "The Ultimate Display," Ivan Sutherland proposed that a computer-based, multi-sensory, interactive simulated world could be created with the ultimate goal that it "look real, feel real, and act real" (Sutherland, 1965). He proposed that a mathematical simulation, such as one of forces on charged particles, could give a user an experience not possible in the physical world. He considered covering all the senses, and discussed visual, auditory, and kinesthetic displays but dismissed smell and taste displays. His discussion was of simulated experience and did not consider transmitted or recorded experience.

A few years later, Sutherland built the first head-mounted display for displaying computer-simulated 3D virtual worlds (Sutherland, 1968). This prototype was monochromatic, eventually stereoscopic, and could display simple wire-frame images in correct perspective. The position and orientation of the user's head, which was allowed to move within certain limits, was measured with a mechanical linkage. A manual input device was later added which could also move in 3D (Vickers, 1974).

Sutherland went on to found Evans & Sutherland Corporation, which pioneered using computer graphics and computer-simulated virtual worlds for training aircraft pilots.

Fred Brooks, 1968

Beginning in the late Sixties, Fred Brooks and his students constructed and tested force-feedback devices for display of force to the hand, which were used in conjunction with graphic displays of interactive virtual worlds (Steelman, 1968)(Batter & Brooks, 1971)(Ouh-Young, Beard & Brooks, 1989)(Brooks, Ouh-Young, Batter & Kilpatrick, 1990). They also explored interactive 3D graphical simulations, with particular emphasis on using computer graphics to represent and simulate molecules (Wright, 1972) and to represent the interior of buildings (Brooks, 1986)(Airey, Rohlf & Brooks, 1990). A see-through HMD was built and tested for various applications (Chung, Harris, Brooks, Fuchs, Kelley, Hughes, Ouh-Young, Cheung, Holloway & Pique, 1989).

In his paper "Grasping Reality Through Illusion: Interactive Graphics Serving Science," Brooks made a case for interactive computer graphic simulation as a "potent tool ... in man's on-going scientific enterprise -- the understanding of the physical universe" (Brooks, 1988). He argued, based on the experience of building a dozen interactive graphics systems for various scientific applications, that graphics were useful for visualizing numerical simulation data, that 3D graphics were essential because the real world is 3D, that making displays multi-sensory made them more effective, that interactivity is crucial for exploring behavior in simulations, and that interactive steering of the simulation computation itself was a powerful technique.

Brooks observed that the massive computing power now available enables scientists to build sophisticated models of complex natural phenomena. These simulation-based mathematical models are qualitatively different from the closed-form algebraic models traditionally used by scientists since the days of Isaac Newton. Some of the characteristics of these new models are that they may be: massive, detailed, 3D, non-linear, discontinuous, discrete, indeterminate, parametric, symbolic. He listed examples of complex natural phenomena which could be profitably explored through such simulation: species population dynamics, protein structures, magnetohydrodynamics, plate tektonics of the earth, shock-wave propagation, black holes, quarks, metereology, oil-field geology, blood flow in the body. He noted that for simulations of phenomena which are not directly observable, such as individual molecules, there is no single correct visualization, but rather different graphic representations of the same phenomena can produce different insights.

Brooks also observed that human beings and computers each had distinct capabilities that the other lacked. Humans surpass computers in the ability to recognize patterns in noisy data, to utilize context, and to form judgements in complicated situations. Computers surpass humans in the ability to remember without forgetting, to calculate accurately, and to rapidly and accurately carry out repetitive tasks. Brooks asserted that at any given point in the development of computer technology, the combination of a human and computer could be more effective than an autonomous "artificial intelligence," and referred to the strategy of interactively linking human and computer with emphasis on their respective strengths as "intelligence amplification."

Myron Krueger, 1974

Myron Krueger conceived the idea of an electronically-mediated responsive environment, and built several systems as art pieces in the Seventies at the University of Wisconsin. Krueger's vision was not of an apparatus strapped to the face, but rather a room filled with devices that monitored and responded to the user's behavior. He called this "artificial reality."

His early pieces used pressure sensors in floor and other sensors which were linked to actuators in the room to cause things to happen, and he also used video cameras linked to monitors, sometimes cross-coupled with one another. His later work used video cameras to capture images of the user, and real-time image processing on the user's silhouette to allow the user's gestures to control events in a computer graphic world. The computer graphics were merged with the video image and projected onto a large screen so that the user could see both him- or herself and the simulated graphic world being controlled.

In 1983, Krueger published a book, "Artificial Reality", which described responsive environments which he had built and laid out his vision of the possibilities (Krueger, 1983).

Tom Furness, 1977

For two decades, Tom Furness directed a program at Wright-Patterson Air Force Base which was aimed at creating a "Super Cockpit". In this new kind of cockpit for aircraft pilots, the aircraft sensor data was transformed for display to the pilot into 3D visual and auditory representations (Buchroeder, Seeley & Vukobradatovich, 1981) (Furness, 1986a) (Kocian, 1988). This was, in essence, creating synthetic senses for the pilot to allow, for example, seeing compass headings superimposed on the horizon, seeing maps superimposed onto the world for night flying, seeing spatial envelopes representing dangerous areas to avoid, and hearing the direction of approaching aircraft. The Visually-Coupled Airborne Systems Simulator (VCASS) was built in 1977, and was a ground-based simulator intended to prove concepts for later incorporation into cockpits of real planes.

Many of the components needed for VCASS did not exist initially, so Furness's lab funded development of see-through HMDs, magnetic tracking of head and hand position, and spatialized sound. VCASS also included speech synthesis and speech recognition subsystems. Furness also proposed using gaze-direction as a means of control. Furness also sponsored considerable human factors research, resulting in a large reference work, the *Handbook of Human Perception and Performance* (Boff, Kaufman & Thomas, 1986).

Furness's work with the Air Force focussed on synthetic senses for aircraft pilots and possible remote piloting. He later wrote about applications of HMDs in other fields, for example, computer-aided design, sensory prostheses, medical imaging, and the control of microscopes (Furness, 1986b).

Henry Fuchs, 1977

Henry Fuchs was a student at the University of Utah where Sutherland and his students experimented with the first HMD prototype. Tracking was a major problem for that early HMD, and Fuchs built a one-dimensional CCD-based optical tracker that measured lateral head movements (Fuchs, Duran & Johnson, 1977). For his Ph.D. dissertation in 1977, Fuchs built a laser-based 3D scanner which could capture the 3D profile of a surface.

Fuchs went to the University of North Carolina, where, in 1986, he co-directed the construction of a see-through HMD system which used half-silvered mirrors and small color LCD displays (Holloway, 1987). Throughout the Eighties, Fuchs and his colleagues worked to design special-purpose real-time 3D graphics computers, and this work culminated in Pixel-Planes 4 (Fuchs, Goldfeather, Hultquist, Spach, Austin, Brooks, Eyles & Poulton, 1985) (Eyles, Austin, Fuchs, Greer & Poulton, 1987) and Pixel-Planes 5 (Fuchs, Poulton, Eyles, Greer, Goldfeather, Ellsworth, Molnar, Turk, Tebbs & Israel, 1989). These high-powered graphics engines were used to explore applications for the HMD in medical imaging (Fuchs, Levoy & Pizer, 1989) (Mills & Fuchs, 1990) (Ohbuchi & Fuchs, 1990) (Ohbuchi & Fuchs, 1991), molecular graphics, and architecture. An optical tracker was also developed (Wang, Azuma, Bishop, Chi, Eyles & Fuchs, 1990).

Micheal Naimark, 1980

Micheal Naimark did the cinematography for the Aspen "movie-map," the interactive videodisc-based simulation of driving around Aspen, Colorado (Lippman, 1980). He later directed other movie-maps using interactive videodiscs, including one for the Exploratorium in San Francisco that allows the user to pan in two dimensions through aerial views of the Bay Area with the direction of view always centered on the Golden Gate Bridge.

In his paper "Elements of Realspace Imaging," Naimark proposed a taxonomy of six successively more complete methods for recording and reproducing experience (Naimark, 1991). These are: monoscopic imaging, stereoscopic imaging, multiscopic imaging, panoramics, surrogate travel, and "real-time" imaging. The first five categories cover techniques which allow successively greater freedom in achieving during replay arbitrary viewpoints in the 3D scene which was recorded. He defines the last category, real-time imaging, as the "process of recording and displaying temporal sensory information indistinguishable from unmediated reality," in other words, high-fidelity motion pictures. He gives a long list of techniques and problems in recording and replaying experience. Some of the issues and problems are orthoscopy, resolution, color, dynamic range, brightness, spatial consistency, accomodation, image stabilization, match-cuts, and storage capacity. Display techniques mentioned are mirrors, relief projection, holography, viewpoint-dependent imaging, and projection onto surrounding screens. Naimark primarily surveyed visual recording techniques in this paper, but he did also mention auditory, force, and vestibular (motion platform) displays.

Scott Fisher, 1982

At MIT in 1982, Scott Fisher built a viewpoint-dependent imaging system which used a stereoscopic display, head tracker, and videodisc to give the user head-motion parallax while looking through the window of the screen at a pre-recorded static scene (Fisher, 1982). The videodisc stored a 2D array of images of the same scene from varying viewpoints, and the head tracking allowed selecting the displayed image to match the user's head position as it changed from moment to moment. The effect was similar to looking at a hologram.

In 1985, Fisher went to the NASA Ames Research Center to work on an HMD being developed for teleoperation of robots in space. In his paper "Virtual Environment Display System," he describes the HMD system developed at Ames (Fisher, McGreevy, Humphries & Robinett, 1986). The intended applications were control of autonomous and semi-autonomous telerobots, display of abstract "data-spaces," and human factors research. The system consisted of a relatively inexpensive HMD made from monochrome LCD displays and wide-angle optics (Howlett, 1983), a magnetic tracker for the head and hand, an instrumented glove which measured the bending angles of the fingers (Foley, 1987), a 3D sound subsystem (Wenzel, Wightman & Foster, 1988), and a speech recognizer. A freely-movable stereoscopic display mounted on a counter-balanced mechanical linkage was also developed.

Teleoperation experiments were done linking the HMD to a pair of video cameras mounted on a motorized gimbal able to mimic the operator's head motions in pitch, roll and yaw. The NASA HMD system was developed to be useful for both simulated virtual worlds and telerobotic control, thus covering both simulated and transmitted experience.

Warren Robinett, 1982

Warren Robinett ,(the author of this paper), designed video games at Atari starting in 1977. He went on to design computer games to teach mathematics to children. *Rocky's Boots* , published in 1982, was an interactive graphical simulation which taught digital logic design to children by allowing them to manually construct circuits on the screen and see signals flowing through the circuitry as it operated (Robinett & Grimm, 1982)(Robinett, 1983).

In 1986, Robinett joined the team at NASA Ames with McGreevy and Fisher and wrote much of the initial software for the NASA HMD system (Fisher, McGreevy, Humphries & Robinett, 1986)(Rheingold, 1991).

In 1989, he went to the University of North Carolina where he directed the Head-Mounted Display Project (Robinett, 1990) (Robinett & Rolland, 1991) (Robinett, 1991a) (Robinett, 1991b) (Holloway, Fuchs & Robinett, 1991) (Robinett & Holloway, 1992) and, later co-directed the Nanomanipulator Project

(Robinett, Taylor, Chi, Wright, Brooks, Williams & Snyder, 1992).

Mike McGreevy, 1984

At NASA Ames Research Center, starting in 1984, Mike McGreevy assembled LCD screens from pocket television sets, stereoscopic optics, a head tracker, and a motorcycle helmet to make the prototype HMD which was later refined and extended by Fisher, Robinett, and others (McGreevy, 1984)(Fisher, McGreevy, Humphries & Robinett, 1986).

McGreevy's interests were in space applications of the HMD, and he tapped the image data of Mars returned by the Viking spacecraft to make a prototype system allowing simulated presence on the real topography of the Martian landscape as it was sensed by the spacecraft (McGreevy, 1991). One potential use for this, given that lightspeed delays made teleoperation impossible on Mars, was to plan the exploration route for a semi-autonomous unmanned Mars rover vehicle.

Jaron Lanier, 1986

The instrumented glove used in the NASA Ames HMD system was supplied to NASA in 1986 by VPL Research, a small company founded by Jaron Lanier. The glove was invented by Tom Zimmerman of VPL. In 1989, VPL began selling a commercial HMD, the EyePhone, and the hardware and software to link together two people wearing HMDs in the same virtual world. This system was called "Reality Built for Two" or RB2 (Blanchard, Burgess, Harvill, Lanier, Lasko, Oberman & Teitel, 1990).

Lanier coined the term “virtual reality.” He asserted that this computer-simulated world could have any desired properties, and described it as an "intentional dream” (Barlow, 1990). He talked about "jacking into lobsters," that is, using your own body to control simulations of creatures with more limbs than you possess. He asserted that it would become easy to create complex objects with manual control gestures, so that easily-created visual representations of objects in the world would take over the function of words, making language unnecessary. He called this "post-symbolic communication."

William Bricken, 1990

Starting in 1987, William Bricken worked on the Cyberspace project at AutoDesk, Inc., and subsequently went to the Human Interface Technology Laboratory at the University of Washington, working at both places on computer-simulated virtual worlds viewed with HMDs.

In his 1990 papers, "A Vision of Virtual Reality" and "Virtual Reality: Directions for Growth," Bricken defines virtual reality broadly as "the body of techniques that apply computation to the generation of experientially valid realities” (Bricken, 1990a)(Bricken, 1990b). He discussed a great range of topics, including multi-sensory display, behavior transducers, telepresence, interactive fiction, collaborative work, and simulated creatures (which he called "entities"). He discussed many applications, including therapy, art, travel, education, and simulation. He offered many aphoristic observations, for example, "Psychology is the physics of virtual reality. Our body is our interface. Knowledge is in experience. Data is in the environment. Scale and time are explorable dimensions. One experience is worth a trillion bits. Realism is not necessary."

David Zeltzer, 1991

In his 1991 paper "Autonomy, Interaction, and Presence," David Zeltzer proposed "a taxonomy of graphical simulation systems” (Zeltzer, 1992). The taxonomy consisted of three independent scalar dimensions which defined a space of possibilities, the "AIP cube". The dimension *autonomy* described the sophistication and dynamics of the model defining the virtual world. *Interaction* measured the degree to which user actions could affect what happened in the virtual world. *Presence* measured the sensory fidelity and breadth (the number of senses to which displays were aimed). The three dimensions were presented as rough lumped measures of the the sophistication of the three key components of a simulated virtual world: the model defining the virtual world (autonomy), the input devices which let the user affect what happens in the virtual world (interaction), and the displays which let the user perceive the virtual world (presence).

Each dimension was measured by a scalar running from zero to one. Zeltzer commented that it was "not clear how to rigorously quantify" these dimensions, so assigning values for a given system would seem to be more a matter of judgement than measurement. He gave examples for each dimension. Autonomy ranged from a static model (autonomy = 0) to a fully autonomous agent (autonomy = 1). Interaction ranged from batch (interaction = 0) to real-time access to all model paramaters (interaction = 1). Presence ranged from static graphics (presence = 0) to sensory stimulation indistinguishable from the real world (presence = 1).

Zeltzer gave examples of where specific systems fell within the AIP cube. He mapped the point (autonomy = 0, interaction = 0, presence = 0) to typical computing in the early Sixties: batch processing of simple graphical models with output on a plotter. The point (autonomy = 1, interaction = 1, presence = 1) mapped to "fully autonomous agents and objects which act and react according to the state of the simulation, and which are equally responsive to the actions of the human participant(s). In addition, sensory stimulation provided to the participant(s) in the virtual environment is indistinguishable from what would be expected in a physical setting". He stated that the (1,1,1) point was probably not achievable without direct neural connection. He described existing graphics systems that lacked one thing or another and therefore mapped to other corners of the cube. An interesting (hypothetical) system was "digital Shakespeare" at the point (autonomy =1, interaction = 0, presence = 1). The user could view the action of the play from any viewpoint, and could rewind or fast-forward but would be unable to affect what happened in the play.

Zeltzer observed that "it is not possible to simulate the world in all its detail and complexity, so for a given task we need to carefully identify the sensory cues that must be provided in order for a human to accomplish the task, and match as closely as possible the human perceptual and motor performance required for the task". This technique is called *selective fidelity*.

The AIP cube is more useful conceptually for suggesting untried possibilities than for actually classifying systems. The problem is that the possibilities lumped into each dimension do not map in any obvious fashion to a single linear scale. It is not clear how the autonomy, interaction, or presence of a system would be measured.

Howard Rheingold, 1991

Howard Rheingold's 1991 book *Virtual Reality* surveys the field, describing people, giving anecdotes, and sketching the historical sequence of ideas (Rheingold, 1991). In the last section of the book, Rheingold considers teledildonics, "electronic LSD," and remotely-operated warplanes, which lead him to questions about ethics and about human nature.

Brenda Laurel, 1991

Brenda Laurel's book, *Computers As Theater* ,uses uses a theatrical point of view to analyze computer-mediated experience (Laurel, 1991a). *The Art of Human-Computer Interface Design*, edited by Laurel, also has several articles on virtual environments (Laurel, 1991b).

4. Prospects and Limits of Synthetic Experience

The developments now underway in the various subclasses of synthetic experience are far from mature, and it seems clear that further exploitation of their inherent possibilities offers to humanity great increases in awareness, power, and the ability to effectively use our vast corpus of information and knowledge.

With this power comes danger, and also the likelihood that these tools will become so pervasive and symbiotic with us that they change the very nature of being human. This is not necessarily a bad thing. Agriculture has changed humanity, as have writing, mathematics, and science. But for an enterprise that portends great changes for humanity, it is sensible to ask: How far will it go? How far *can* it go? What is most likely to happen? What are the ultimate limits?

In this section of the paper, we will discuss some likely paths of development for various strands of synthetic experience, as best we can anticipate from our present knowledge of the physical universe and the human neural apparatus. We will try to extrapolate as far as we can, without regard to cost or effort or time to develop, but limited by what ultimately seems possible. Our model in this constrained speculation is Arthur C. Clarke's *Profiles of the Future* (Clarke, 1962).

We first imagine synthetic experience developed as far as we can foresee, and then consider the physical and psychological limits that form the ultimate boundary of the possible.

4.1 A vision of the potential long-term development of synthetic experience

4.1.1 Perfect fidelity of synthetic experience

Reproducing various aspects of experience has been the goal of most media from painting onward. The ultimate development of this aspect of synthetic experience is coverage of all human senses and perfect fidelity for each sensory channel, so that the human is unable to tell the difference between synthetic and natural experience. This criterion applies mainly to transmitted and simulated experience -- you feel like you're really there. For recorded experience, since you have no ability to act and must merely observe the action, you cannot be fooled into believing that you are having a natural experience.

4.1.2 Synthetic senses spanning all detectable phenomena

For every known detectable phenomenon involving energy or matter, and for every kind of sensor that exists, a mapping may be made from the phenomenon to the built-in human senses. This implies that all these imperceptible phenomena have been given visual representations, or representations matched to other human senses.

4.1.3 Instant travel at the speed of light

Transmitted experience permits experience at a distance, and with multiple teleoperated robots the human operator could switch his or her presence from one site to another as easily as people today call around the world to various sites on the telephone. With many sensors scanning the world in real-time, an integrated 3D global database could be maintained in real-time, so that the virtual travel from one site to another could be continuous motion with a changing viewpoint, rather than teleporting from site to site, as with the telephone. The speed of travel can be as fast as the user desires.

4.1.4 Apparent magic

With all vehicles, tools, factories, libraries, and other controllable systems connected to the world-wide communications grid, a virtually present operator will be able to control devices or systems at any location. The input devices and control gestures will be arbitrary and independent of the devices being controlled. Arbitrary gestures activating, moving, creating and destroying objects of the physical world will be similar in appearance and capabilities to the mythical idea of magic, for example, like the wizard Dr. Strange of the comic books.

4.1.5 Adventures in microworlds

Micro-teleoperation will permit human operators to perceive and manipulate things in microscopic worlds ranging in scale from the merely tiny (where a bee is your size), down to the microscopic (where a bacterium is your size), and on down to atomic scale (where an atom is your size). Micromanipulators will allow actions in these microworlds. Among the things to do down there are to build things (nanotechnology), to explore and probe (biological and physical science), and to interact with microscopic creatures (entertainment).

4.1.6 Global experience database and time-travel

In order to augment their native memories, all people may in the future come to wear multi-sensory recording devices, much as people today wear eyeglasses. In addition there will be many other sensors that are simultaneously scanning the real-world and recording this data. The data from the sensors at these many locations may be integrated into a global 3D database, spanning the experience of all of humanity since people began to wear the recorders. It will be possible, using this database, to share the experience of a distant person in real-time, or to relive the experience of any person from any time in the past so long as the experience is recorded in the database. This is, in effect, time-travel into the recorded past. Time-travel into the future is possible to the extent that a simulation can predict what will occur, but there are fundamental theoretical limitations about how well we will ever be able to precisely predict the future. To the same degree that future events can be accurately predicted by simulation, it should be possible to travel into the simulated past and be given the ability to perform actions, and thus to experience what might have been.

4.1.7 Simulating the real world and fantasy worlds

There will be two main kinds of simulated experience: simulations that attempt to accurately mimic or predict what can happen in the real world; and convincing fantasy worlds that dispense with the constraints and physical laws of reality. The accurate simulations will be useful for education, for training, and for exploring the consequences of contemplated actions. The fantasy worlds will entertain, and perhaps delude and addict. Many human participants will be able to simultaneously inhabit these simulated worlds, seeing and interacting with each other. In addition, simulated creatures will also inhabit these worlds. Simulated creatures, also known as autonomous agents or artificial intelligences (AIs), are today quite limited in capability as compared with human beings or even the most primitive animals. At the simplest, a creature is an object in the virtual world that moves about on its own, initiating actions. More sophisticated creatures would include in their behaviors simulations of the abilities to recognize, to remember, and to plan.

4.1.8 Overlays onto the real world

There will be many databases registered with the real world and able to be superimposed onto it, for example: labels, maps, notes to specific people, diagrams, paths, graffiti, as well as the actions from earlier times recorded in the experience database. It will be a matter of choice which, if any, of these overlays are viewed by each human at any given moment.

4.1.9 Shared virtual worlds

Many people will be able to enter simultaneously into these virtual worlds, including real-time transmitted virtual worlds (by virtual travel to the same location), and recorded experience (by traveling to the same place and time in the experience database). For two people simultaneously reliving a particular trace through recorded space and time, it will be a matter of choice whether they see one other. They could equally well relive the same experience independently, or see representations of one another as they observe the action from separate locations. It would also be possible to observe the traces of earlier observers of a given recorded action.

4.1.10 Animals in virtual worlds

Humans will be able to masquerade in the real world as animals of any species for which teleoperated robot manikins can be built. To interact effectively with these animals, an understanding of how they communicate is necessary, and this understanding may be facilitated by the very existence of such robots. A particularly effective way to make, for example, a teleoperated cockroach, would be implant remote controls into its neural tissue. This would be a teleoperated but biologically real creature -- a zombie cockroach. Zombies would probably be better at fooling animals than the most sophisticated puppet mechanism. However, if cockroaches, frogs and squirrels can be wired as zombies, then most likely humans can be also.

It will also be possible for higher animals, such as mammals, to enter into virtual worlds in the same way that humans do. If we can reproduce experience with perfect fidelity for humans, then we should be able to do the same for a cat or a snake, with the displays and input devices tailored for their particular senses and physiognomies. Thus, an ape or dog or cat could perceive a surrounding 3D world, move through it, and interact with it, more or less like they do in the real world. It would be interesting to see if an ape, given the ability to capture and replay representations of its experience for other apes, could make use of this ability for communication.

4.1.11 Direct neural connection

Since all perception and action is accomplished through the human sensory and motor nerves, display devices could stimulate the nerve fibers directly rather than the sensory organs, and in principle achieve the same perception. Likewise, motor nerve impulses could bypass the muscles and directly trigger actuators which manipulate the world.

4.2 Ultimate limits

Having imagined many grandiose extrapolations to our present capabilities for synthetic experience, we turn to the question of ultimate limits. What aspects and laws of the physical universe and the structure of the human body and brain will constrain what kinds of synthetic experiences are achievable?

4.2.1 Fidelity

For the fidelity of reproduced experience, it seems possible in principle both to cover all human senses simultaneously with display devices, and to achieve arbitrarily high fidelity for each sensory channel. For hearing, CD-quality stereophonic sound has already reached the point of indistinguishability from natural sound. For vision, improvements in video can be posited that would reach the limits of human visual acuity, field of view, color palette, motion detection, and depth perception. This will not be easy, but neither is it impossible.

Taste has four dimensions, (salty, sour, sweet, and bitter), and arbitrary tastes may be synthesized with combinations of these primaries, just as arbitrary colors are synthesized from the red, green, and blue phosphor dots of the television screen. Similarly, smell appears to have seven dimensions (Kandel & Schwartz, 1985), although there is some scientific dispute about this. Thus, smells could probably be synthesized from primary components also.

Displaying to the vestibular system presents more of a problem -- motion platforms can tilt and can apply strong but brief accelerations, within the limits of their travel, and imperceptibly drift back to center position. However, displaying a zero-G experience or a two minute-long 3G Apollo blast-off would seem to require more than a motion platform, perhaps a simulation chamber in high Earth orbit that could be left floating or accelerated at arbitrary rates for long periods.

Displaying with perfect fidelity to the haptic and tactile senses presents such a daunting engineering challenge that it seems nearly impossible. Some sort of whole body exoskeleton would be needed, with integrated arrays of pressure, vibration, and temperature displays covering the entire body surface. However, pieces of this have been done already. The Jacobsen Arm, developed by Steve Jacobsen at the University of Utah, is an exoskeleton with force feedback that covers all the joints of the arm and hand including shoulder, elbow, wrist, thumb, and fingers. A full body force-feedback exoskeleton can therefore be imagined. Tactile arrays that can display texture and vibration to the finger or other surfaces of the skin exist (Linvill, 1973)(Rheingold, 1991). Building a flexible, body-covering tactile array would be very difficult but not impossible.

This completes the list of external sensory organs, leading to the conclusion that experience can potentially be reproduced with a fidelity that is indistinguishable from natural experience. However, a great deal of effort and expense would be required to develop some of the required display devices.

4.2.2 Transmitted experience

The speed of light limits real-time transmitted experience. In teleoperating a distant robot, there is a speed of light time lag proportional to the remoteness of the robot. This time lag applies both to the transmission of the sensor data to the operator's display and the transmission of the operator's actions to the remote actuators. Humans are sensitive to such time lags, with a lag of 1 second causing enough confusion to the operator to make many manual tasks impossible. A lag of 100 milliseconds is quite noticeable, and some experiments show sensitivity down to 5 milliseconds. A light-speed circumnavigation of the Earth takes roughly 150 milliseconds, so teleoperation is feasible anywhere on the planet, or in low Earth orbit. However, the lag will begin to be noticeable for operator-robot distances of about 1000 kilometers.

With a 3-second round trip speed of light delay from the Earth to the Moon, teleoperation is not feasible for dextrous manual tasks that require force-feedback or continuous guidance. However, some things can be accomplished by greatly slowing down the operations performed. To Mars, the minimum round trip time lag is about 10 minutes, so real-time teleoperation from Earth to Mars is impossible.

4.2.3 Recorded experience

The fundamental limit on recorded experience is storage capacity. We can estimate the data rate for human experience by using the standard NTSC video data rate for comparison. A hand-held videotape can currently store an hour's worth of visual and auditory experience.

In the future, fidelity and therefore data rate will increase for recorded experience, but at the same time storage density will increase. Let's explore the likely changes in these two parameters. NTSC video nominally transmits an image frame of 640 x 480 pixels at a rate of 30 frames per second with the color of each pixel encoded by 8 bits for each of the primary colors red, green, and blue. This is roughly 200 million bits per second. To be conservative, we will increase this by a factor of 250 to allow for the greater resolution, field of view, and so forth, needed for perfect visual fidelity. Since humans are primarily visual creatures, another factor of 2 should be sufficient to record all the other senses. This gives us a data rate for human experience of 10^{11} bits per second, which is probably much more than is needed.

Current common storage techniques, such as music compact discs, can store roughly a billion bits per cubic centimeter. However, we can expect storage density to continue to increase until it hits some sort of physical limit. It should ultimately be possible to encode information in the arrangement of matter on the atomic scale (Feynman, 1960)(Drexler, 1991). Assuming a nanotechnological storage device that stores 1 bit for every 1000 atoms gives us a storage density of 10^{20} bits per cubic centimeter. At this density, a lifetime (100 years) of human experience can fit into the volume of a large grape (3 cm^3), and the experiences of all of humanity (10 billion people) for 10,000 years would fit into a cubic mile of nano-storage.

Storing a continuous recording of every human's egocentric experience is thus possible in principle. But what about storing data from the many other sensors scattered throughout the world so as to be able at a later time to review what happened at a location which no one observed at the time of interest? This is possible to a certain degree, but it is impossible to record everything that happened to an arbitrary degree of accuracy. Reality is too complex -- even if it were possible to sense the positions of all the atoms in a scene from nanosecond to nanosecond, this would overwhelm even the nano-storage postulated above. It does not seem possible to record everything that one might want to later experience. Recording will necessarily be selective.

4.2.4 Simulated experience

The fundamental limit on the accuracy of simulations of reality is the inherent complexity of reality itself. Also, our current understanding is that reality is fundamentally unpredictable at the quantum level and therefore at all higher levels. We can expect simulations that attempt to predict and model reality to achieve increasing fidelity within restricted domains (such as vehicle simulators, weather prediction, or orbital mechanics), but there can be no expectation of being able to predict the *exact* behavior of macroscopic quantities of matter under arbitrary conditions. There are several reasons for this. Chaotic systems existing in nature are sensitive to tiny perturbations and therefore defy prediction. The uncertainty principle bars accurate knowledge of the parameters needed for prediction of behavior on the atomic scale. Certain quantum events appear to be truly unpredictable. Even without all the foregoing reasons, the sheer number of atoms involved makes such an atom-by-atom simulation unfeasible.

It would seem that it will never be possible to predict the exact behavior of an ensemble of atoms because of the reasons of chaos, uncertainty, and quantum unpredictability. And even without those objections, it would also seem that it would take a simulation computer enormously larger in physical size than the object of the simulation, and that the simulation would proceed much slower than the atomic interactions being simulated. But to be cautious, we might phrase it as a question: Does it always take more atoms and more time to simulate, and predict the outcome of, a physical process than are involved in the process itself?

This limit on simulation accuracy is a fundamental limit, not just a temporary obstacle than can be overcome with effort and time. How can simulated experience deal with this? Several techniques have proved useful so far. Selective fidelity looks at the limited domain being simulated and the task to be done in order to select which aspects of experience to simulate accuractely. Interactive steering of the display computation dynamically allocates computational resources to where they will be most effective, as in, for example, multiple levels of detail for objects in flight simulators which are chosen according to the user's current location.

For fantasy worlds that don't attempt to accurately model the real world, it is hard to see a fundamental limit. With the same display apparatus that can reproduce transmitted experience with perfect fidelity, the display should not impose any limit. Extrapolating into the future, we can expect enormous increases in computational power available for running the simulations underlying these fantasy worlds. Perhaps the limitation will simply be a noticeable difference between the way things happen in simulated worlds and the way things happen in reality which cannot be simulated in detail.

5. Conclusions

The taxonomy presented in this paper, with its eight dimensions, offers a way to classify devices that use technology to transmit, filter, record, or simulate experience. The taxonomy also helps to understand the relationships among existing synthetic experience devices and to suggest as-yet-untried possibilities.

It appears that the new capabilities offered by the various kinds of synthetic experience can be developed rather far before we hit fundamental limits. It seems plausible that in the near future, we will be able to transmit experience from a distance with good fidelity and coverage of most senses, see the invisible through sensors linked to HMDs, travel visually to distant places as easily as we make telephone calls, record 3D scenes and actions for later replay, and enter simulated 3D worlds. These simulated worlds may be either fantasy worlds or may try to accurately model some aspect of the real world.

The speed of light limits the distance which experience may be transmitted without perceptible time lags. The accuracy with which the real world can be simulated is limited by several factors which make it unlikely and probably impossible to predict by simulation the *exact* behavior of many aspects of the real world. The recording of experience is limited by storage capacity. If storage of information can be realized in the arrangement of matter at the atomic scale, continuous recording of human experience may be feasible.

6. Acknowledgements

I would like to thank my colleagues in the UNC Computer Science Department for their help and support, particularly the Head-Mounted Display team, the Pixel-Planes team, and the Nanomanipulator team. Fred Brooks and Ivan Sutherland suggested modifications to earlier versions of the classification system presented in this paper.

Fred Brooks, Henry Fuchs, and Gary Bishop helped to guide our HMD-related research at UNC with good technical judgement. Rich Holloway implemented the core software library, with help from Jim Chung, Drew Davidson, Erik Erikson, and Mark Mine. Jannick Rolland, Vern Chi, and Doug Holmgren designed and built prototype optical systems. Stan Williams and his team at UCLA developed the scanning-tunneling microscope which was brought to UNC and interfaced to the UNC HMD system by Russ Taylor and Rick Snyder (of UCLA), with guidance from Stan Williams and Bill Wright. Henry Fuchs, John Poulton, John Eyles, Trey Greer, Laura Weaver, and others designed and built the Pixel-Planes 4 and Pixel-Planes 5 graphics engines which were crucial to our ability to explore computer-generated virtual worlds at UNC. The Pixel-Planes software team did a good job of exploiting the latent power of the Pixel-Planes hardware.

I would also like to thank my colleagues from my days at NASA Ames Research Center, The Learning Company, and Atari, where my ideas about computer-simulated worlds evolved. I would like to thank Julius Smith, Scott Kim, Fred Lakin, Ken Harrenstien, and Ken Shoemake for many years of technical discussions, critiques, and support. I would like to thank the Banff Centre for the Arts in Canada for providing an environment in which to study synthetic experience as a new artistic medium and thereby gain a non-technical perspective on this new phenomenon.

This research was supported by the following grants: Defense Advanced Project Research Agency #DAEA 18-90-C-0044, National Science Foundation Cooperative Agreement (with DARPA) #ASC-8920219, Office of Naval Research #N00014-86-K-0680, and National Institutes of Health #5-R24-RR-02170.

7. References

Airey, J. , Rohlf, J., Brooks, F. P. (1990). Towards Image Realism with Interactive Update Rates in Complex Virtual Building Environments, *Computer Graphics: Proceedings of 1990 Symposium on Interactive 3D Graphics,* Snowbird, Utah, 41-50.

Barlow, J. P. (1990). Life in the DataCloud: Scratching Your Eyes Back In, *Mondo 2000*, Summer.

Batter, J.J. and F.P. Brooks, Jr. 1971. GROPE-I: A computer display to the sense of feel. *Information Processing: Proceedings of IFIP Congress 71.* 759-763.

Blanchard, C., S. Burgess, Y. Harvill, J. Lanier, A. Lasko, M. Oberman, M. Teitel. Reality Built for Two: A Virtual Reality Tool. *Proc. 1990 Workshop on Interactive 3D Graphics.* 35-36.

Boff, K. R., Kaufman, L., Thomas, J. P. (1986). *Handbook of Perception and Human Performance* . New York: John Wiley & Sons.

Bright, M. (1984). *Animal Language*, 55-108. London: British Broadcasting Corporation.

Brooks, F. P. (1986). Walkthrough—A dynamic graphics system for simulating virtual buildings. In *Proceedings of 1986 Workshop on Interactive 3D Graphics.* UNC-Chapel Hill. 9-21.

Brooks, F.P., Jr., M. Ouh-Young, J.J. Batter, and P.J. Kilpatrick. 1990. Project GROPE—Haptic displays for scientific visualization. *Computer Graphics: Proceedings of SIGGRAPH '90*. Dallas,TX. 177-185.

Brooks, F. P. Jr. (1988). Grasping reality through illusion. *Proceedings of CHI '88*. Washington, D.C.

Bricken, W. (1990a). *A Vision of Virtual Reality,* Tech. Report HITL-P-90-1, Human Interface Technology Lab, University of Washington.

Bricken, W. (1990b). *Virtual Reality: Directions of Growth,* Tech. Report HITL-R-90-1, Human Interface Technology Lab, University of Washington.

Buchroeder, R. A., Seeley, G. W., & Vukobradatovich, D. (1981). *Design of a Catadioptric VCASS Helmet-Mounted Display,* Optical Sciences Center, University of Arizona, under contract to the U.S. Air Force Armstrong Aerospace Medical Research Laboratory, Wright-Patterson Air Force Base, Dayton, Ohio, AFAMRL-TR-81-133.

Bush, V. (1945). As We May Think, *Atlantic Monthly*, July.

Chung, J. C., Harris, M. R., Brooks, F. P., Jr., Fuchs, H., Kelley, M. T., Hughes, J., Ouh-Young, M., Cheung, C., Holloway, R. L., & Pique, M. (1989). Exploring Virtual Worlds with Head-Mounted Displays, *Non-Holographic True 3-Dimensional Display Technologies*, SPIE Proc. Vol. 1083.

Caudell, T. P., Mizell, D. W. (1992). Augmented Reality: an application of heads-up display technology to manual manufacturing processes, HICSS conference, January 1992.

Clarke, A. C. (1962). *Profiles of the Future*. New York: Holt, Reinhart & Winston.

Corliss & Johnsen (1968). NASA Special Publication 5070.

Drexler, K. E. (1991). *Unbounding the Future*. New York: William Morrow & Co.

Englebart, D. (1963). A Conceptual Framework for the Augmentation of Man's Intellect, *Vistas in Information-Handling, 1,* 1-29, Howerton, P. W., ed. Washington DC: Spartan Books.

Eyles, J., J. Austin, H. Fuchs, T. Greer and J. Poulton. Pixel-Planes 4: A summary. *Proceedings of the Eurographics '87 Second Workshop on Graphics Hardware: Advances in Computer Graphics Hardware II*. 183-208.

Feynman, R. P. (1960). There's Plenty of Room at the Bottom, reprinted in *Journal of Micromechanical Systems 1*(1), January 1992.

Fisher, S. S. (1982). Viewpoint Dependent Imaging: An Interactive Stereoscopic Display, *Proceedings SPIE* 367.

Fisher, S.S., McGreevy, M., Humphries, J., & Robinett, W. (1986). Virtual Environment Display System, *Proc. 1986 Workshop on Interactive 3D Graphics*, 77-87.

Foley, J. D. (1987). Interfaces for Advanced Computing, *Scientific American* 257(4) 126-135, October.

Fuchs, H., Duran, J., & Johnson, B. (1977). A system for automatic acquisition of three-dimensional data. In *Proceedings of the 1977 National Computer Conference*. 46:49-53.

Fuchs, H., J. Goldfeather, J.P. Hultquist, S. Spach, J.D. Austin, F.P. Brooks, Jr., J.G. Eyles and J. Poulton. 1985. Fast spheres, shadows, textures, transparencies, and image enhancements in Pixel-Planes. *Computer Graphics*: *SIGGRAPH '85 Conference Proceedings.* 19:3:111-120. Reprinted in *Advances in Computer Graphics I*. G. Enderle, M. Grave, and F. Lillehagen, eds. Springer-Verlag. 1986. 169-187. Also reprinted in *Tutorial: Computer Graphics Hardware—Image Generation and Display.* Hassan K. Reghbati and Anson Y.C. Lee, eds. IEEE Computer Society Press. 1988. 222-231.

Fuchs, H., J. Poulton, J. Eyles, T. Greer, J. Goldfeather, D. Ellsworth, S. Molnar, G. Turk, B. Tebbs and L. Israel. 1989. A heterogeneous multiprocessor graphics system using processor-enhanced memories. *Computer Graphics: Proceedings of SIGGRAPH '89*. 23:4:79-88.

Fuchs, H., M. Levoy, S. M. Pizer, "Interactive Visualization of 3D Medical Data," *IEEE Computer*, Vol. 22, No. 8, August 1989, pp. 46-51.

Furness, T. (1986a). *The Super Cockpit and Human Factors Challenges*, Tech. Report HITL-M-86-1, Human Interface Technology Lab, University of Washington.

Furness, T. (1986b). *Putting Humans into Virtual Space*, Tech. Report HITL-R-86-1, Human Interface Technology Lab, University of Washington.

Goertz, R. C. (1961). The ANL Model 3 Master-Slave Electric Manipulator -- Its Design and Use in a Cave, *Proc. 9th Conf. on Hot Laboratories and Equipment*, Washington D.C., United States Atomic Energy Commission.

Goertz, R. C. (1965). An experimental head-controlled television system to provide viewing for a manipulator operator, Proc. 13th RSTD Conf. 57.

Heilig, M. (1955). Cinema of the Future, *Espacios* (Mexico City), January.

Holloway, R. L. (1987). *Head-Mounted Display Technical Report*, #TR87-015, Dept. of Computer Science, University of North Carolina at Chapel Hill

Holloway, R., H. Fuchs, W. Robinett. 1991. Virtual-Worlds Research at the University of North Carolina at Chapel Hill. *Proc. Computer Graphics '91*. London, England.

Howlett, E. M. (1983, September 27). *Wide Angle Color Photography Method and System*, U.S. Patent Number 4,406,532.

Johnsen & Corliss (1967). NASA Special Publication 5047.

Kandel, E. R. & Schwartz, J. H. (1985). Principles of Neural Science, 418-419. New York: Elsevier Science Publishing Co.

Kocian, D. F. (1988). *Design Considerations for Virtual Panoramic Display (VPD) Helmet Systems*, Armstrong Aerospace Medical Research Laboratory, Visual Display Systems Branch, Wright-Patterson Air Force Base, Dayton, Ohio 45433-6573.

Krueger, M. (1983). *Artificial Reality*. Reading MA: Addison-Wesley.

Laurel, B. (1991a). *Computers As Theater*. New York: Addison-Wesley.

Laurel, B. (1991b). *The Art of Human-Computer Interface Design*. New York: Addison-Wesley.

Linvill, J. G. (1973). *Research and Development of Tactile Facsimile Reading Aid for the Blind (the Opticon)*, report to U.S. Dept. of Health, Education, and Welfare, Stanford Electronics Laboratory, Stanford University.

Lippman, A. (1980). Movie-Maps: An Application of the Optical Videodisc to Computer Graphics, Computer Graphics *14*(3).

McGreevy, M. W. (1984). NASA Ames Virtual Environment Display: Applications Requirements [internal technical document], Aerospace Human Factors Research Division, NASA Ames Research Center.

McGreevy, M. W. (1991). Virtual Reality and Planetary Exploration, *29th AAS Goddard Memorial Symposium*, Washington DC.

Mills, Peter H., Henry Fuchs. 1990. 3D ultrasound display using optical tracking. *Proceedings of the First Conference on Visualization in Biomedical Computing.* Atlanta, GA. 490-497.

Naimark, M. (1991). Elements of Realspace Imaging: A Proposed Taxonomy, *Proceedings of SPIE: Electronic Imaging, 1457.*

Ohbuchi, R. and H. Fuchs. 1990. Incremental 3D ultrasound imaging from a 2D scanner. *Proceedings of the First Conference on Visualization in Biomedical Computing.* Atlanta, GA. 360-367.

Ohbuchi, R., and H. Fuchs. 1991. Incremental volume rendering algorithm for interactive 3D ultrasound imaging. *Proceedings of the Information Processing in Medical Imaging (IPMI) Conference XII.* 486-500.

Ouh-Young, M., D.V. Beard and F.P. Brooks, Jr. 1989. Force display performs better than visual display in a simple 6-D docking task. *Proceedings of IEEE 1989 Robotics and Automation Conference.* Scottsdale, AZ. 3:1462-1466.

Rheingold, H.. 1991. *Virtual Reality.* New York. Summit.

Robinett, W., Grimm, L. (1982). *Rocky's Boots* [software product]. Fremont CA: The Learning Company.

Robinett, W. (1983). *Imaginary Worlds* (unpublished book manuscript).

Robinett, W. (1990). Artificial Reality at UNC Chapel Hill [videotape], 10 minutes, *SIGGRAPH Video Review*.

Robinett, W., and J. Rolland. 1991. A computational model for the stereoscopic optics of a head-mounted display. To appear in *Presence. 1:1.* Also in *Proceedings of SPIE: Electronic Imaging*, Vol. 1457, Santa Clara, California, February 1991.

Robinett, W. (1991a). Electronic expansion of human perception. *Whole Earth Review.* 16-21.

Robinett, W. (1991b). Technological Augmentation of Memory, Perception, and Imagination, *Virtual Seminar on the Bioapparatus*, 17. Banff, Alberta, Canada: The Banff Centre for the Arts.

Robinett, W., and R. Holloway. 1992. Flying, grabbing and scaling in virtual worlds. Proceedings 1992 ACM Symposium on Interactive 3D Graphics, Cambridge MA.

Robinett, W., Taylor, R., Chi, V., Wright. W. V., Brooks, F. P. Jr., Williams, R. S., & Snyder, E. J. (1992). *The Nanomanipulator Project: An Atomic Scale Teleoperator* (to appear in 1992 SIGGRAPH course notes for the course "Implementation of Immersive Virtual Worlds").

Sheridan, T. B. (1992). Musings on Telepresence and Virtual Presence, *Presence 1*(1).

Steelman, H.S. 1968. The GROPE-I system: An analysis of friction and backlash problems. MS thesis. University of North Carolina, Chapel Hill.

Sutherland, I. E. (1965). The ultimate display, 1965 *Proceedings of the IFIPS Congress*, *2*, 506-508.

Sutherland, I. E. (1968). A head-mounted three-dimensional display, 1968 *Fall Joint Computer Conference, AFIPS Conference Proceedings*, 33, 757-764.

Vickers, D. L. (1974). *Sorceror's Apprentice: head mounted display and wand*, Ph.D. dissertation, Dept. of Computer Science, Univ. of Utah, Salt Lake City.

Wang, J., R. Azuma, G. Bishop, V. Chi, J. Eyles, and H. Fuchs. 1990. Tracking a head-mounted display in a room-sized environment with head-mounted cameras. *Proceedings: SPIE '90 Technical Symposium on Optical Engineering & Photonics in Aerospace Sensing.* Orlando, FL.

Weiss, R. (1989). New Dancer in the Hive, *Science News 136*(18), 282-3

Wenzel, E. M., Wightman, F. L., Foster, S. H. (1988). A Virtual Acoustic Display for Conveying Three-Dimensional Information, *Proceedings of the Human Factors Society*.

Wright, W. V. (1972). *An Interactive Computer Graphics System for Molecular Studies*, Ph.D. dissertation, Dept. of Computer Science, University of North Carolina, Chapel Hill.

Zeltzer, D. (1992). Autonomy, Interaction, and Presence, *Presence 1*(1).