Closing Keynote Speech: Reconfigured Self as Basis for Humanistic Intelligence *

Prof. Steve Mann
University of Toronto, Department of Electrical and Computer Engineering
10 King's College Road, Toronto, Canada
mann@eecg.toronto.edu   http://wearcomp.org/steve.html

USENIX-98, New Orleans June 15-19, 1998

Abstract

John S. Quarterman asks the questions: "Will the U.S. continue to contain more than half of the Internet? Will one company make all the software?"[1]. One could also ask related questions like "Will one university or educational organization fabricate the future by manufacturing consensus of ideas and ideals?". That the U.S. government, together with a single company, might control the world's information flow and software, and thus indirectly, the world's thoughts, ideas, etc., is a troubling thought to many. One possible solution to this problem is in the tradition of science, and thus the notion of disclosure and open peer review, as a basis for allowing anyone the option of acquiring, and thus advancing the world's knowledge base. A further construct called "Humanistic Intelligence (HI)", motivated by the philosophy of science, is proposed. HI provides a new synergy between humans and machines that seeks to involve the human rather than having computers emulate human thought or replace humans. Particular goals of HI are human involvement at the individual level, and providing individuals with tools to challenge society's pre–conceived notions of human–computer relationships. An emphasis in this article is on computational frameworks surrounding "visual intelligence" devices, such as video cameras interfaced to computer systems.

1 Problem statement

Software hegemony, seamlessness of thought, and the building of computer science upon a foundation of secrecy, present a series of problems to the individual. Advanced computer systems is one area where a single individual can make a tremendous contribution to the advancement of human knowledge. A system that excludes any individual from exploring it fully, may prevent that individual from "thinking outside the box" and therefore advancing the state–of–the–art.

Some of the new directions in Human–Computer Interaction (HCI) suggest bringing advanced computing into all aspects of life. Computers everywhere, constantly monitoring our activities, and responding "intelligently" have the potential to make matters worse, from the above perspective, because of the possibility of excluding the individual user from knowledge not only of certain aspects of the computer upon his or her desk, but also of the principle of operation and the function of everyday things.

1.1 Spreading of sterility

The U.S. Department of Agriculture, working together with Delta and Pine Land Company, has engaged in genetic engineering research to develop seeds that will produce sterile plants[2]. This effort, they claim, will boost sales of new seeds by preventing farmers from using seeds obtained from their previous year's plants. Preventing farmers from replanting from a previous year's crops will force them to buy new seeds each year, thereby increasing Delta and Pine Land Company's revenue.

This proposed invention has created outrage in the community, and critics have pointed out that the invention could threaten the world's food supply through the spreading of sterility. Moreover, there is concern that this single corporation would control the world's food supply. Furthermore, this would end the centuries-old practice of farmers doing cross–breeding themselves to adapt plants to local soil and climate conditions.

*Thanks to Kodak and Digital Equipment Corporation (DEC) for assistance with the Personal Imaging and Humanistic Intelligence projects

1.2 Hegemony of the software barons

What do seeds have to do with advanced computing systems? Delta and Pine's ability to genetically turn the reproductive system of seeds on and off is analogous to a compiler's ability to produce executable programs in an intellectually–encrypted form, so that the software vendor can decide whether or not to include source–code and other conceptual disclosure.

Programs distributed with source code are analogous to the traditional seeds that contribute to a healthy and diverse world food supply. Programs that are operationally disclosed (with source code, etc.) will be adapted to local requirements, through modification and sometimes cross–breeding with other program source code. Each disclosed program contributes to the knowledge base of individual users. Not necessarily all users will modify the software, but even a small fraction of the users doing this can and often will improve or contribute to the knowledge base.

While there may be an API (applications program interface) provided for software development, and this may satisfy 99% of the users, it is the very 1% of users it doesn't satisfy, who are the ones "thinking outside the box" and therefore perhaps most apt to advance the state–of–the–art. By shutting off access to that very small percentage of the user–population, software barons have contributed to a situation of intellectual stagnation outside their regimes. In this way they have set the stage for shutting out a certain amount of intellectual diversity.

1.3 Seamlessness of thought as intellectual monocropping

A varied supply of ideas, like biodiversity, leads to advancement in thought, and new paradigms in computing. Thus a single entity controlling all the world's supply of software, and in some sense, thought (e.g. a single entity providing all of the world's WWW browsers, and therefore, in some sense influencing the world's Internet content, or a single entity providing all of the world's dictionaries or automatic spell checking programs, in some sense influencing the use of language), could reduce the world's diversity of thought. Seamlessness of thought may reduce intellectual diversity, in the same way that monocropping reduces biodiversity.

1.4 Computer Science *or* Computer Secrecy

Science provides us with ever–changing schools of thought, opinions, ideas, and the like, while all building upon a foundation of verifiable (and sometimes evolving) truth. The foundations, laws, and theories of science, although true by assumption may at any time be called into question as new experimental results unfold. Thus when doing an experiment, we may begin by making certain assumptions, but at any time, these assumptions may be verified. In particular, a scientific experiment is a form of investigation that leads wherever the evidence takes us. In many cases, the evidence takes us back to questioning the very assumptions and foundations we had previously taken as truth, and in some cases, instead of making a new discovery along the lines anticipated by previous scientists, we discover that another previous discovery was false or inaccurate. Sometimes these are the biggest and most important discoveries — things that are discovered by accident. Thus any scientific system that tries to anticipate "what 99% of the users of our result will need" may be constructing a thought prison for the very people who would otherwise be the most likely to advance human knowledge. In many ways their entire user–base is in this thought prison, but many would never know it, since their own explorations do not take them to its outermost walls.

Thus a situation in which one or more of the foundation elements are held in secret is contrary to the principles of science. Although many results in science are treated as a "black box", for operational simplicity, there is always the possibility that the evidence may want to lead us inside that box.

Imagine, for example, conducting an experiment on a chemical reaction between a proprietary solution "A", mixed with a secret powder "B", brought to a temperature of 325 degrees T. (Top secret temperature scale which you are not allowed to convert to other units). (See Figure 1.)

Now it is quite likely that one could make some new discoveries about the chemical reaction between A and B, without knowing what A and B are, and one might even be able to complete a doctoral dissertation and obtain a PhD for the study of the reaction between A and B (assuming a large enough quantity of A and B were available). However, one might ask where one would publish these findings, except maybe in the Journal of Irreproducible Results.

Results in Computer Science that are based, in part, on undisclosed matters, inhibit the ability of

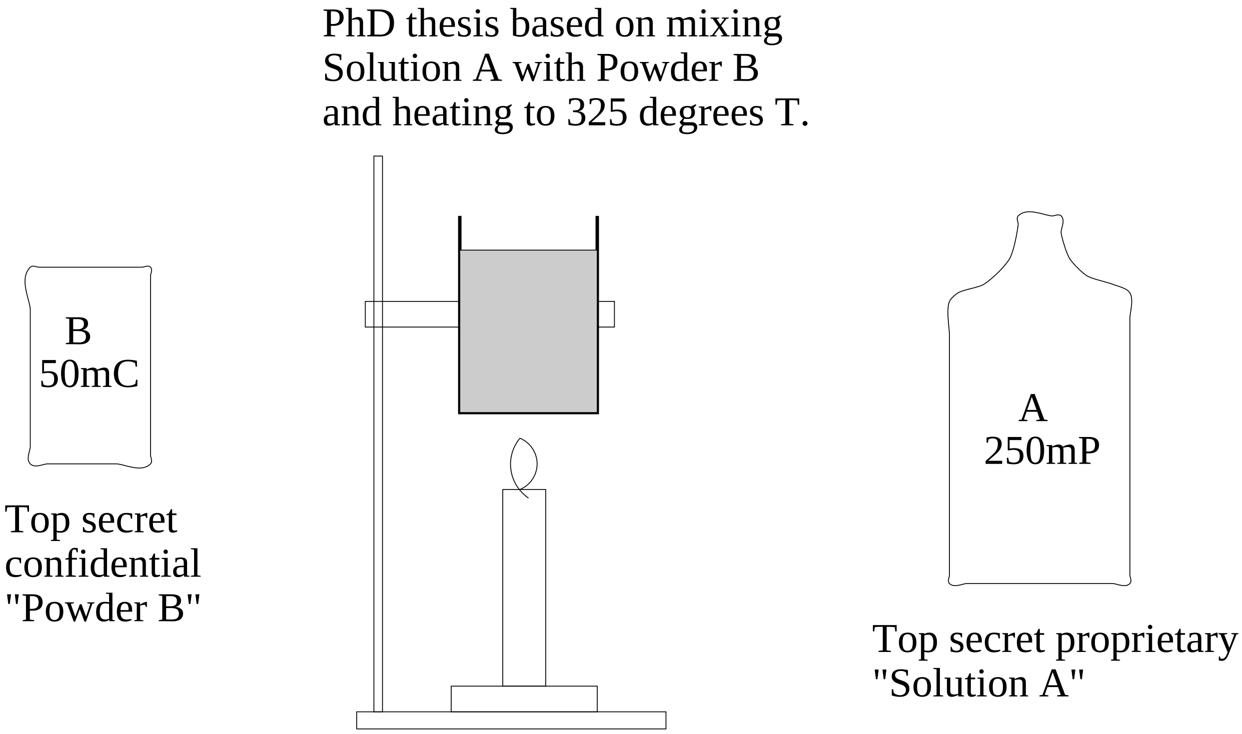

Figure 1: Imagine, for example, conducting an experiment on a chemical reaction between 250mP (proprietary units of volume) of proprietary solution "A", and 50mC (confidential units of mass) of a secret powder "B", where the mixture is raised to 325 degrees T (top secret temperature scale). It is hard to imagine where one might publish results of such an experiment, except, perhaps, in the Journal of Irreproducible Results.

the scientist to follow the evidence wherever it may lead. Even in a situation where the evidence does not lead inside one of the secret "black boxes", science conducted in this manner is irresponsible in the sense that another scientist in the future may wish to build upon the result, and may, in fact, conduct an experiment that leads backwards, as well as forwards. Should the new scientist follow evidence that leads backwards, inside one of these secret "black boxes", then the first scientist will have created a foundation contaminated by secrecy. In the interest of academic integrity, better science would result if all the foundations upon which it were built were such as to be subject to full examination by any scientist who might, at some time in the future, wish to build upon a given scientific discovery.

Thus, although many computer scientists may work at a high–level, there would be great merit in a computational foundation open to examination by others, even if the particular scientist using the computational foundation does not wish to examine it. For example, the designer of a high level numerical algorithm, who uses a computer with a fully disclosed operating system (such as Linux), does other scientists a great service, even if he or she himself or herself only uses it at the API level and never intends to look at its source code or that of the Linux operating system underneath it.

1.5 Obvious or obfuscated?

Imagine a clock that was designed so that when the cover was lifted off, all the gears would fly out in different directions, so that it would be more difficult for a young child to open up his or her parents' clock and determine how it works. Alternatively, imagine the clock were loaded with explosives, so that it would completely self–destruct upon opening.

Assuming a child survived such a bad experience, it is still doubtful that devices made in this manner would be good for society, in particular, for the growth and development of young engineers and scientists with a natural curiosity about the world around them.

As the boundary between software and hardware blurs, devices are becoming more and more difficult to understand. This difficulty arises in part, as a result of deliberate obfuscation on the part of product manufacturers. More and more devices contain general–purpose microprocessors, so that their function depends on software. Specificity of function is achieved through specificity of software rather than specificity of physical form. By manufacturing everyday devices in which there is provided only executable code, without source code, manufacturers have provided a first level of obfuscation. Furthermore, additional obfuscation tools are often used in order to make the executable task image more difficult to understand. These tools include strippers that remove object link names, etc., and even tools for building encrypted executables which contain dynamic decryption function that generates a narrow sliding window of unencrypted executable, so that only a small fragment of the executable is decrypted at any given time. In this way, not only is the end–user deprived of source–code, but the executable code itself is encrypted, making it difficult or impossible to look at the code even at the machine code level.

Moreover, devices such as Complex Programmable Logic Devices (CPLDs), such as the Alterra 7000 series, often have provisions to permanently destroy the data and address lines leading into a device, so that a single chip device can operate as a finite–state machine yet conceal even its machine–level contents from examination. (An excellent tutorial on FPGAs and CPLDs may be found in[3].) Devices such as Clipper chips go a step further by incorporating fluorine atoms, so that if the user attempts to put the device into a milling machine, to mill off layer–by–layer for examination under an electron microscope, the device will self–destruct in a drastic manner that destroys structure. Thus the Clipper phones could contain a "trojan horse", or some other kind of "back door", and we might never be able to determine whether or not this is the case. This is yet another example of deliberate obfuscation of the operational principles of everyday things.

Thus we have a growing number of general–purpose devices whose function or purpose depends on software, downloaded code or microcode. Because this code is intellectually encrypted, so is the purpose and function of the device. In this way, manufacturers may provide us with a stated function or purpose, but the actual function or purpose may differ, or may include extra features, of which we are not aware.

1.6 "Environmental intelligence" or intelligence–gathering systems

There are a number of researchers who have been proposing new computer user–interfaces based on environmental sensors. Buxton, who did much of the early pioneering research into intelligent environments (smart rooms, etc.), was inspired by automatic flush urinals, (as described, for example, in U.S. Pat. 4309781, 5170514, etc.), and formulated, designed, and built a human–computer interaction system called the "Reactive Room"[4] [5]. This system consisted of various sensors, including optical sensors (such as video cameras), and processing, so that the room would respond to the user's movement and activity.

Increasingly we are witnessing the emergence of "intelligent highways", "smart rooms", "smart floors", "smart ceilings", "smart toilets", "smart elevators", "smart lightswitches", etc.. However, a typical attribute of these "smart spaces" is that they were architected by someone other than the occupant. Thus the end–user of the space often does not have a full disclosure of the operational characteristics of the sensory apparatus and the flow of intelligence data from the sensory apparatus.

In addition to the intellectual encryption described in the previous section, where manufacturers could make it difficult, or perhaps impossible for the end user to disassemble such sensory units in order to determine their actual function, there is also the growth of hidden intelligence, in which the user may not even be aware of the sensory apparatus. For example, U.S. Pat. 4309781 (for a urinal flushing device) describes:

> ...sensor...hidden from view and thus discourage tampering with the sensor...when the body moves away from the viewing area...located such that an adult user of average
> height will not see it...sensing means, will be behind other components...positioned below the solenoid to allow light in and out. But the solenoid acts in the nature of a hood or canopy to shield the sensing means from the normal line of sight of most users.... Thus most users will not be aware of the sensing means. This will aid in discouraging tampering with the sensing means. A possible alternate arrangement would be to place the sensing means below and behind the inlet pipe.

U.S. Pat. 4998673 describes a viewing window concealed inside the nozzle of a shower head, where a fiber optics system is disclosed as a means of making the sensor remote, concealment to prevent users from being aware of its presence. U.S. Pat. 5199639 describes a more advanced system where the beam pattern of the nozzle is adapted to one or more characteristics of the user, while U.S. Pat. 3576277 discloses a similar system based on an array of sensing elements.

Means of creating viewing windows to observe the occupants of a space while, at the same time, making it difficult for the occupants to know if and when they are observed, are proposed in U.S. Pat. 4225881 and U.S. Pat. 5726706.

In addition to concealing the sensory apparatus, a goal of many visual observation systems is to serve the needs of the system architect rather than the occupants. For example, U.S. Pat. 5202666 discloses a system for monitoring employees within a restroom environment, in order to enforce hygiene (washing of hands after use of toilet).

Other forms of "intelligence", such as "intelligent highways" often have additional uses, beyond those purported by those installing the systems. For example, traffic monitoring cameras were used to round–up, detain, and execute peaceful protesters in China's Tiananmen Square.

U.S. Pat. 4614968 discloses a system where a video camera is used to detect smoke by virtue of the fact that smoke reduces the contrast of a fixed pattern opposite the video camera. However, the patent also notes that the camera can be also used for other functions such as visual surveillance of an area, since only one segment or line of the camera is needed for smoke detection. Again, the camera may thus be justified for one use, and additional uses, not disclosed to occupants of the space, may then evolve. U.S. Pat. 5061977 and 4924416 disclose the use of video cameras to monitor crowds and automatically control lighting, in response to the absorption of light by the crowds. While this form of environmental intelligence is purportedly for the benefit of the occupants (to provide them with improved lighting), there are obvious other uses.

U.S. Pat. 5387768 discloses the use of visual inspection of users in and around an automated elevator. Again, these provide simple examples of environmental intelligence, in which there are other uses, such as security and surveillance. Although even those other uses (security and surveillance) are purportedly for the benefits of the occupants, and it is often even argued that concealing operational aspects of the system from the occupants is also for their benefit, it is an object of this paper to challenge these assumptions and to provide an alternate form of intelligence.

When the operational characteristics, function, data flow, and even the very existence of sensory apparatus is concealed from the end–user, such as behind the grille of a smoke detector, environmental intelligence does not necessarily represent the best form of human–machine relationship for all concerned. Even when the sensors are visible, there must be the constant question as to whether or not the interests of the occupant are identical to those who control the intelligence–gathering infrastructure.

The need for personal space, free from monitoring, has also been recognized[6] as essential to a healthy life. As more and more personal space is stolen from us, we may need to architect an alternative space of our own.

2 A proposed solution: Humanistic Intelligence for the re–configured self

What is proposed is a computational framework for individual personal empowerment. This framework involves the architecting of a new kind of personal space, through an apparatus that is owned, operated, and controlled by the occupant of that space. In some sense, it is like a building, built for one occupant, and collapsed down around that one occupant. This computational framework for HI is called "WearComp", and will now be described.

2.1 WearComp as basis for HI

WearComp, which originated in the 1970s as a tool for the visual arts[7], is a new form of human–computer interaction comprising a computer that is subsumed into the personal space of the user (e.g. the computer may be worn, hence the term "user" and "wearer" of the computer may be used interchangeably), controlled by the wearer, and has both operational and interactional constancy, eg. is always on and always ready and accessible [8]. Typical embodiments of WearComp comprise a body–worn computer system, a visual display over one or both eyes with text and graphics display capability, and an input device consisting of typically five or more pushbutton switches that may be operated by one hand. Other input devices typically include a microphone and video camera positioned such that it provides a view of the same subject matter the wearer sees. While this apparatus may sound, at first, unwieldy, it has evolved, over the last 20 years or so, into a normal–looking (e.g. unobtrusive) form of clothing. (See Fig 2.)

(a) (b) (c)

Figure 2: Typical embodiments of the WearComp "photographer's assistant" for of Personal Imaging. (a) The "photographer's assistant" project of the 1970s eventually gave rise to cumbersome backpack systems of the early 1980s. (b) These systems evolved into a series of late 1980s clothing-based signal processing and personal imaging systems, which, although still somewhat cumbersome, were more comfortable to use due to better distribution of lesser weight. Here, for example, is a system with left eye display and beam splitter, which was an experiment in partially–mediated reality. Separate antennas facilitated simultaneous voice, video, and data communication. (c) Current (1990s) embodiments of WearComp are unobtrusive. These are more suitable for everyday use.

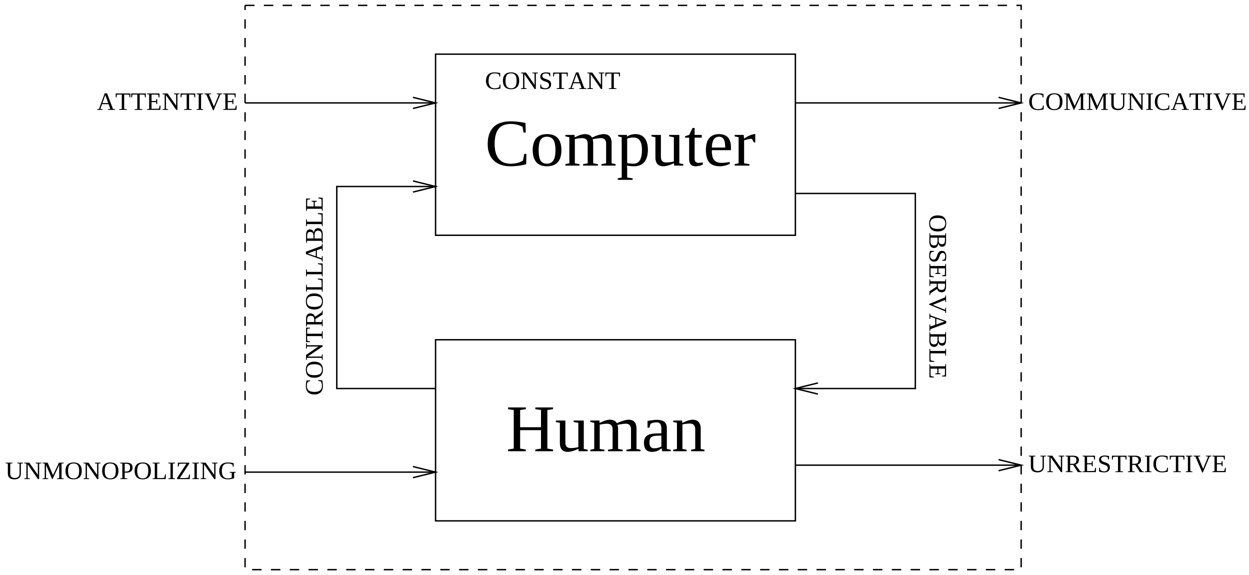

Figure 3: **A new form of Human–Computer Interaction**: WearComp is constantly attentive to its surroundings, is inextricably intertwined with its user/host, yet in a way which neither wholly monopolizes the user's attention nor appreciably restricts the user's ability to simultaneously do other things (computing being a secondary task as opposed to the primary task that it normally is in its usual office or desktop context).

2.2 Definition of 'WearComp'

'WearComp', is with reference to Fig 3 in terms of its eight attributes, as follows:

1. **CONSTANT**: Always ready. May have "sleep modes" but never "dead". Unlike a laptop computer which must be opened up, switched on, and booted up before use, it is always on and always running.

2. **UNRESTRICTIVE to the user**: ambulatory, mobile, roving, "you can do other things while using it", e.g. you can type while jogging, etc.

3. **UNMONOPOLIZING of the user's attention**: it does not cut you off from the outside world like a virtual reality game or the like. You can attend to other matters while using the apparatus. It is built with the assumption that computing will be a secondary activity, rather than a primary focus of attention. In fact, ideally, it will provide enhanced sensory capabilities. It may, however, mediate (augment, alter, or deliberately diminish) the sensory capabilities.

4. **OBSERVABLE by the user**: It can get your attention continuously if you want it to. Almost–always–observable: within reasonable limitations (e.g. that you might not see the screen while you blink or look away momentarily) the output medium is constantly perceptible by the wearer.

5. **CONTROLLABLE by the user**: Responsive. You can grab control of it at any time you wish. Even in automated processes you can manually override to break open the control loop and become part of the loop at any time you want to (example: "a big Halt button you want as an application mindlessly opens all 50 documents that were highlighted when you accidently pressed Enter" would make a computer more CONTROLLABLE. Infinitely–often–controllable: the constancy of user–interface results from almost–always observability and infinitely–often controllability in the sense that there is always a potential for manual override which need not be always exercised.

6. **ATTENTIVE to the environment**: Environmentally aware, multimodal, multisensory. (As a result this ultimately gives the user increased situational awareness).

7. **COMMUNICATIVE to others**: Can be used as a communications medium when you want it to. Expressive: allows the wearer to be expressive through the medium, whether as a direct communications medium to others, or as means of assisting the production of expressive media (artistic or otherwise).

8. **PERSONAL**: Human and computer are inextricably intertwined.

2.3 Wearable, tetherless computer–mediated reality

Such a computational framework allows one to subsume all of the personal electronics devices that one might normally carry, such as:

- cellular phone;
- pager;
- wrist watch;
- heart monitor;
- camera;
- video camera

into a single device. Obviously, since it is a fully featured computer, it is possible to respond to email, plan events on a calendar, type a report, or the like, while walking, standing in line at the bank, or the like. It this way WearComp anticipated the later arrival of the so–called "laptop computer", but has advantages over the laptop computer in the sense that it can be used while walking around doing other things. However, the real power of WearComp is in its ability to serve as a basis for Personal Imaging and Humanistic Intelligence.

2.4 Personal Imaging

Personal Imaging is a camera–based computational framework in which the camera behaves as a true extension of the mind and body, after a period of long–term adaptation, as described in [9]. In this framework the computer becomes a device that allows the wearer to augment, diminish, or otherwise alter his or her visual perception of reality. Moreover, it allows the wearer to allow others to alter his or her visual perception of reality, and therefore becomes a communications device.

Mediated reality arises, for example, through the eyeglass–based version of WearComp (called "WearCam", as described in http://wearcam.org/research.html) in which the glasses absorb and quantify incoming rays of light, process this visual information, and then reconstruct new rays of light in response to the processing. In a fully–mediated reality environment, the wearer of the glasses would be blind to the outside world, but for the processing that is inserted between the analysis ("camera") and synthesis ("display/viewfinder") portions of the glasses.

2.5 Reality User Interface

One application of computer–mediated reality is to create, for each user of the apparatus, a possibly different interpretation of the same visual reality. Since the apparatus shares the same first–person perspective as does the user (and in fact since the apparatus is what enables the user to see at all!!!), then of course the apparatus provides the processing system (WearComp) with a view of how the user is interacting with the world. In this way, each user may built his or her own user–interface within the real world. For example, one user may decide to have the computer automatically run a telephone directory program whenever it sees the user pick up a telephone. This example is similar to hypertext, in the sense that picking up the telephone is like clicking on the telephone with a mouse, if it were in an HTML document. "Clicking" on real objects is done by simply touching them.

Outlining objects with the fingertip is another example of a reality user interface (RUI).

2.6 Reality Window Manager

When windows are used together with the RUI, a new kind of window manager results. For example, while waiting in lounge or other waiting area, a user might define walls around the lounge as various windows. In this way screen real–estate is essentially infinite. Although not all screens are visible at any one time, portions of them become visible through the WearComp glasses, when they are looked at. Others in the lounge need not be able to see them, unless they are wearing similar glasses and the user has also permitted them access to these windows (as when two users are planning upon the same calendar space).

There are no specific boundaries in this form of window manager. For example, if the user runs out of space in the lounge, he or she can walk out into the hall and create more windows on the walls of the hallway leading into the lounge. It is also easier to remember where all the windows are when they are associated with the real world. Part of this ease of memory comes from having to walk around the space, or at least turn one's head around in the space.

This window manager, called RWM, also provides means of making the back of the head "transparent", in a sense, so that one can see windows in the front as rightside–up and windows behind as upside–down. This scheme simply obeys the laws of projective geometry. Rearview windows may be turned on and off, since they are distracting for concentration, but they are useful for quick navigation around a room.

2.7 The VideoOrbits head–tracker

Of course one cannot expect a head–tracking device to be provided in all possible environments, so head–tracking is done by the reality–mediator itself, using the VideoOrbits[10] tracking algorithm. The VideoOrbits head–tracker does head–tracking based on a visually observed environment, yet works without the need for high–level object recognition.

VideoOrbits builds upon the tradition of image processing [11] [12] combined with the Horn and Schunk equations [13] and some new ideas in algebraic projective geometry and homometric imaging, using a spatiotonal model, $\tilde{p}$, that works in the neighbourhood of the identity:

$$\left(\sum_{x,y} (\phi(x,y)\phi^T(x,y)) \right) \tilde{\mathbf{p}} = -\sum_{x,y} F_t \phi(x,y) \quad (1)$$

where $\phi^T = [F_x(xy,x,y,1), F_y(xy,x,y,1), F, 1]$, $F(x,t) = f(q(x))$ at time t, $F_x(x,t) = (df/dq)(dq(x)/dx)$, at time t, and $F_t(x,t)$ is the difference of adjacent frames. This "approximate model" is used in the innermost loop of a repetitive process, and then related to the parameters of an exact projectivity and gain group of transformations, so that the true group structure is preserved throughout. In this way, virtual objects inserted into the "reality stream" of the person wearing the glasses, follow the orbit of this group of transformations, hence the name 'VideoOrbits'.

A quantagraphic version of VideoOrbits is also based on the fact that the unknown nonlinearity of the camera, f, can be obtained from differently exposed images $f(q)$ and $f(kq)$, etc., and that these can be combined to estimate the actual quantity of light entering the imaging system:

$$\hat{q}(\mathbf{x}) = \frac{\sum_i c_i(\frac{\mathbf{Ax}+\mathbf{b}}{\mathbf{cx}+1}) \frac{1}{k_i} f^{-1}(F_i(\frac{\mathbf{Ax}+\mathbf{b}}{\mathbf{cx}+1}))}{\sum_i c_i(\frac{\mathbf{Ax}+\mathbf{b}}{\mathbf{cx}+1})} \quad (2)$$

where c_i is the derivative of the recovered nonlinear response function of the camera, f, and A, b, and c are the parameters of the true projective coordinate transformation of the light falling on the image sensor. This method allows the actual quantity of light entering the reality–mediator to be determined. In this way, the reality–mediator absorbs and truly quantifies the rays of light entering it. Moreover, light entering the eye due to the real and virtual objects are therefore placed on an equal footing.

Other researchers, such as Feiner[14][15][16][17], propose augmented–reality environments with some similar characteristics, although these are tethered to a specific location, in part, because of the non–vision–based tracking. Feiner's work is seminal, and of great value, in the context of location–specific augmented reality. However, an object of Humanistic Intelligence is to be able to architect a personal visual reality that does no rely on specific environmental provisions.

2.8 Personal Safety Device

The RUI and RWM are both examples of how WearComp not only subsumes the function of the laptop computer, but goes beyond it.

Another area in which WearComp provides a truly new form of user–interface not found on laptop computers and PDAs is in its constancy of user–interface, and constancy of operation. This characteristic perhaps becomes most evident in its use as a personal safety/security camera. Imagine, perhaps as you walk down some quiet street late at night, an assailant wielding a sawn–off shotgun, demanding cash from you. You would not like have time or opportunity to pull out a camcorder to record the experience, but since the eyeglasses are worn constantly, you would have captured the experience.

2.9 Camera of the future

Less extreme examples of WearComp as a new user–interface include the ability to construct a personal documentary video without conscious thought or effort. For example, in a fully–mediated reality, all light entering the eyes, in effect, passes through the computer, and may therefore be recorded (and possibly transmitted to remote locations). Wearable Wireless Webcam[18] was one example of a personal documentary video recorded using a reality mediator.

In the future, we may well have the capability to capture and recall our own personal experiences, and to have photo albums generated automatically for us. We will never miss baby's first steps, because we will have a retroactive record feature that lets us, for example, "begin recording from 5 minutes ago". Photo albums, in addition to being generated automatically, may also be exhibited while they are being generated. Rather than sending postcards to friends and relatives, or showing them an album after you come back from vacation, you may just put on your sunglasses and have the album sent to them automatically, as was done with the Wearable Wireless Webcam experiment in which video was transmitted, and still images were also automatically selected from the video.

2.10 The "Intelligence Arms Race"

While there will no doubt be more environmental intelligence than personal intelligence, there is at least the hope that there might be an end to the drastic imbalance between personal intelligence and environmental intelligence. The individual making a purchase in a department store may have several cameras pointing at him to make sure that if he removed merchandise without payment that there would be evidence that he did not pay for the item. However, in the future, he will have a means of collecting evidence that he did pay for the item, or a recorded statement of a clerk about the refund policy. More extreme examples, such as the case of Latasha Harlins[1] also come to mind.

In this sense, the camera–based reality–mediator becomes an equalizer much like the Colt45 in the Wild West. When there's a standoff, it doesn't matter whether one person has a big gun and the other has a small gun, so long as there is enough ammunition for mutually assured destruction.

In the WearCam case, it is simply a matter of mutually assured accountability.

3 Conclusions

3.1 Intelligence: Real or Artificial?

A goal of Artificial Intelligence (AI), is to emulate human intelligence. This goal has evolved toward Environmental Intelligence in which cameras and microphones everywhere monitor us, purportedly to make our lives better (Fig 4). In many ways, the goals of AI and EI are similar, in the sense that both attempt to emulate human intelligence in a machine. HI is presented as an alternate direction, in which

[1] A customer falsely accused of shoplifting, and fatally shot in the back by a shopkeeper as she attempted to walk out of the shop.

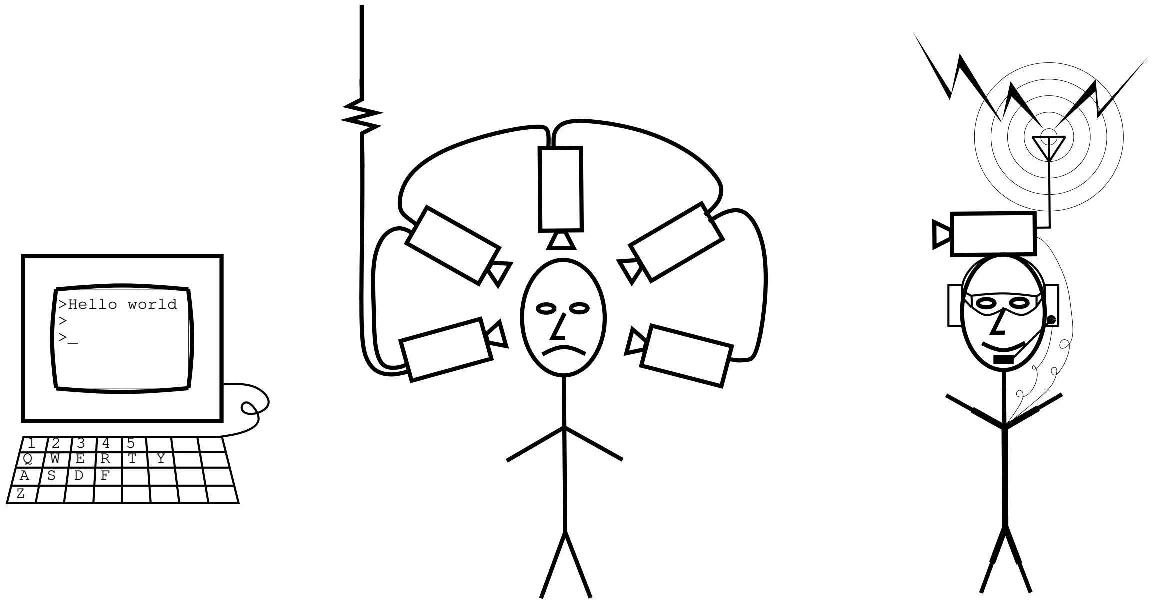

Artificial Intelligence (AI) **Environmental Intelligence (EI)** **Humanistic Intelligence (HI)**

Figure 4: A goal of Artificial Intelligence (EI) is to emulate human thought. A goal of Environmental Intelligence (EI) is to have an environment that monitors the user and "tries to be helpful", although it then opens up the question as to whether or not the architect of the space is serving his or her own needs which may be different than, or in addition to the needs of the occupants. Humanistic Intelligence (HI) encourages individuals to architect their own space. Rather than attempt to emulate human thought, HI recognizes the synergy between human and machine that arises from having human intelligence in the feedback loop of a computational process. HI also builds, to some extent, on the existential principle of self–determination and mastery over one's own destiny.

the human becomes the machine. While we already have smart rooms, smart floors, smart toilets, smart elevators, smart light switches, etc., we may how have "smart people" — people equipped with devices that have traditionally only been part of the architecture or environment.

3.2 The new business opportunity: System Administration and consulting

HI sets forth a new computational framework in which each computer is extremely customized to the needs of one and only one individual. The computer becomes very much like a prosthetic device, or like prescription eyeglasses. And just as you would not want to wear another person's undergarments, or another person's mouthguard, you may not likely find yourself wearing another person's computer.

And the traditional paradigm of one world-wide software vendor providing to everyone identical copies of an executable–only distribution no longer applies. Instead, complete reconfigurability is needed. Instead, each user will customize his or her own environment. Since many layperson users are not well–versed in operating system kernel source code, there will be a growing need for system administrators and consultants.

In the future, software will be free, and users will buy support. There will be little problem with software piracy, both because the software will be free, and because a version of the software customized for one person will be of lesser use to another person with different needs. Each person will likely have different user preferences. And because the computer will function as a true extension of the user's mind and body, it would not do the user well to ingest software owned by somebody else. The computer will function much like a "second brain", and In the true spirit of freedom of thought, it would be preferable that any commercial interests in the customization and contents of one's "second brain" be a work for hire (e.g. an interaction in which the end–user owns the rights) rather than a software purchase. Thus there will be an exponentially growing need for personal system administrators as new people are born into the community of connected collective humanistic intelligence.

References

[1] John S. Quarterman. The future of the internet. In *Proceedings of USENIX 98*, page New Orleans. USENIX, June 15–19 1998.

[2] Leora Broydo. A seedy business. *The Mojo Wire*, April 7–13 1998. http://wearcam.org/seedy.html.

[3] Stephen Brown and Jonathan Rose. Fpga and cpld architectures: A tutorial. *IEEE Design and Test of Computers*, 13(2), 1996. see also http://wearcam.org/jayarpubs.html.

[4] Jeremy Cooperstock and Bill Buxton. Reactive room, 1995. http://www.dgp.utoronto.ca:80/people/rroom/rroom.html.

[5] J. R. Cooperstock, S. S. Fels, W. Buxton, and K. C. Smith. Reactive environments: Throwing away your keyboard and mouse, 1997. http://www.csl.sony.co.jp/person/jer/pub/cacm/cacm.html.

[6] Erving Goffman. *The Presentation of Self in Everyday Life*. Doubleday, Garden City, New York, 1959.

[7] Steve Mann. An historical account of the 'WearComp' and 'WearCam' projects developed for 'personal imaging'. In *International Symposium on Wearable Computing*, Cambridge, Massachusetts, October 13-14 1997. IEEE.

[8] Steve Mann. Smart clothing: The wearable computer and wearcam. *Personal Technologies*, March 1997. Volume 1, Issue 1.

[9] S. Mann. 'mediated reality'. TR 260, M.I.T. Media Lab Perceptual Computing Section, Cambridge, Massachusetts, http://wearcam.org/mediated-reality/index.html, 1994.

[10] S. Mann and R. W. Picard. Video orbits of the projective group; a simple approach to featureless estimation of parameters. TR 338, Massachusetts Institute of Technology, Cambridge, Massachusetts, See http://n1nlf-1.eecg.toronto.edu/tip.ps.gz 1995. Also appears IEEE Trans. Image Proc., Sept 1997, Vol. 6 No. 9.

[11] Anastasios N. Venetsanopoulos. Nonlinear mean filters in image processing. *IEEE Trans. ASSP*, 34(3):573–584, June 1986.

[12] T. G. Stockham, Jr. Image processing in the context of a visual model. *Proc. IEEE*, 60(7):828–842, July 1972.

[13] B. Horn and B. Schunk. Determining Optical Flow. *Artificial Intelligence*, 1981.

[14] S. Feiner, B. MacIntyre, and D. Seligmann. Karma (knowledge-based augmented reality for maintenance assistance), 1993. http://www.cs.columbia.edu/graphics/projects/karma/karma.html.

[15] S. Feiner, B. MacIntyre, and D. Seligmann. Knowledge-based augmented reality, Jul 1993. Communications of the ACM, 36(7).

[16] S. Feiner, Webster, Krueger, B. MacIntyre, and Keller.

[17] S. Feiner, Webster, Krueger, B. MacIntyre, and Keller. Architectural anatomy, 1995. Presence, 4(3), 318-325.

[18] S. Mann. Wearable Wireless Webcam, 1994. http://wearcam.org.