VUI - Virtual Environments as the Next Generation User Interface

Real. Cool.

James Overton
CS 208 Canon of Computer Science
Professor Koltun
Spring 2011

Imagine you're sitting in your home office when you decide to do some reading. Reaching for some papers on the desk in front of you, you grab a stack of documents – some articles, some magazines, perhaps a book or two. Thumbing through them you sort them into piles, arranging them around you. Then, choosing one, you open it and start reading. Once you're finished you say "Computer – go to: entertainment center, browse videos"; suddenly the office setting – desk, papers, books and décor – vanishes before your eyes, phasing cleanly into an entertainment center replete with your entire DVD collection where documents on your 'desk' (now your 'coffee table') had been only moments before. Picking one up, you say "Computer – play this" and the movie begins.

While user interfaces that incorporate augmented reality holographic projections, voice commands and gesture control may not be technologically and commercially feasible for many years, the potential for virtual user interfaces (VUIs – the interface to a virtual environment) to become mainstream computer interfaces has never been greater.

Virtual environments are already superior application settings in many domains, including gaming, design, simulation and learning. While their core appeal and use will remain context-specific, virtual user interfaces (VUIs) may never fully supplant WIMP as the dominant operating system interface. Nevertheless, VUIs will achieve ubiquity as highly compelling complements and alternatives to the traditional WIMP user interface.

Domain and Existing Work

"The proper study of exemplars demands humility of approach… the task of the student is to find and master the excellence that has gone before, even if his muse or his new circumstances then drives him in a totally different direction" – Frederick Brooks, The Design of Design

A user interface is a communication channel between man and machine. And so of necessity, UIs have a proud tradition of being central to the field of computer science. From the first computers, which used batch interfaces, to Vannevar Bush's "Memex" thought experiment (which described the first desktop metaphor) of 1945, to command line interfaces, to Ivan Sutherland's "Sketchpad" (the first GUI) of 1963, to Engelbart's On-Line System which incorporated windows and a mouse-controlled pointer, to the first WIMP (window-icon-menu-pointer) GUI operating system developed at Xerox PARC in 1981 for the Xerox Star Workstation, to the now touch-screen, gesture-based WIMP interfaces that are all but ubiquitous in mobile and tablet computing, advances in user interface design have reflected and embodied great moments both technologically and experientially in the history of computer science.

The quality of a user interface might be tersely stated as the synergy of the following attributes:

1. Usability – ease of use, intuitive feel, etc.
2. Efficiency – speed + power
3. X-factor – how cool or fun it is to use

Ultimately, the goal of a UI might be stated as maximizing the computer's potential to augment human capabilities. A somewhat logical corollary is that augmenting human capabilities is more easily

facilitated if the computer presents a more natural, familiar setting. Even as early as 1965, in his now famous lecture "The Ultimate Display", Ivan Sutherland hinted at the impending shift toward more real, immersive, virtual interfaces: "Don't think of that thing as a *screen*, think of it as a *window*, a window through which one looks into a *virtual world*. The challenge to computer graphics is to make that virtual world look real, sound real, move and respond to interaction in real time, and even feel real" (Brooks 1999).

Virtual environments are not new, and so neither is the concept of a virtual user interface. In fact, researchers in computer science have been developing natural, intuitive 3D user interfaces[1] since the late 1980s. Despite having been successful in a research sense, they did not enjoy popular or commercial success. These 3D UIs were highly compelling, but were very much ahead of their time both inasmuch as the requisite hardware support and the proposed divergence from the prevailing user experience.

The first and perhaps most famous 3D UI was the Information Visualizer, completed by a research group led by Stuart Card at Xerox PARC in 1991. Developed on the heels of the first WIMP-based OS (also a product of Xerox PARC), this UI was decades ahead of its time. I had the honor of speaking with Professor Card on the matter (Card 2011), who indicated that two things were required for 3D UIs to attain popular success:

1) Technological readiness – the software and hardware that comprise the system must be both capable of supporting the feature set and inexpensive enough for commodity production

2) Dominance of design – the new UI must demonstrate a significant advantage to displace/supplant the entrenched technology

Technologically, it made use of custom SGI equipment (like the geometry engine) that were largely "unavailable on commodity systems at that time" (Heer 2011) – costing somewhere on the order of $50k (incidentally, with less computing power than that found in many mobile phones today). Also, by the time it was feasible, the inherent user experience represented too great a departure from the WIMP GUI that came to dominate mainstream computer/operating system UI.

But despite its inability to "achieve commercialization" (Winograd 2011), the Information Visualizer encompassed many compelling features – beyond mere aesthetic appeal – that underscored the benefits of a 3D UI. Among them were:

[1] While virtual environments might be more traditionally associated with virtual reality or virtual worlds (online communities populated by avatars), I use the term more loosely to encompass any virtual (computer) space in which the setting is meant to replicate or model a real one. And so inasmuch as their use in this paper – where both are semantically identical in that they model a real setting or entity – I use the terms "virtual environment" and "3D UI" interchangeably.

- "A paradigm for [new user] interfaces – the *information workspace…* a special environment in which the cost structure of the needed materials is tuned to the requirements of the work process using them"
- A characterization of the cost structure of information, with regard to "information-based work processes", as well as a way to optimize it, by "organizing the parts of the system hierarchically"
- A task-based focus for grouping programs logically, both for user visual-organization, and for advantaging locality of reference and reference clustering
- The role of information visualization as facilitating spatial memory in a UI

(Card 1991).

In addition, this early work brought to light some of the limitations and problems inherent to a 3D UI. Professor Jeffrey Heer, with whom I was also privileged to discuss the matter, enumerated the three key problems as:

1) Occlusion – entities blocking other entities visually
2) Greater effort & time – navigating 3D space is physically and cognitively more complicated
3) Perspective distortion – foreshortening, alignment, distance judgments may be unintentionally misleading

Professor Heer also cautioned that when the data is not inherently 3D, that "in most cases studies have shown that 2D displays – if you can do them well – will be more effective" (Heer 2011). This notion was voiced to some degree even as far back as 1996, when the 2D vs. 3D debate raged to the tune of "a tendency to conflate style with substance" (Mullet, Schiano 1996).

But what about Bob? Microsoft Bob was (in 1995) Microsoft's virtual home interface to a desktop computing environment. While it is widely regarded as one of the greatest technology failures of all time – one of Time Magazine's "50 worst inventions" (Time 2010) and CNET dot com's "worst product of the decade" (CNET 2007) – it was nonetheless a warning against (many things, but most relevantly) the risks involved in attempting to make mainstream OS UI more natural and familiar alone, without a more significant value proposition.

Influential research in the field continued at Xerox PARC and Microsoft, largely inspired by the Information Visualizer. The WebBook and Web Forager, projects also led by Professor Card, extended the Information Visualizer to the activity of web browsing, providing "an information workspace that enable[d] rapid interaction with materials gleaned from the Web" (Card 1996). A few years later at Microsoft, The Task Gallery was completed; developed by a group led by George Robertson, one of Professor Card's colleagues in 3D UI technology at Xerox PARC, it replaced the desktop metaphor with a gallery motif, with intent to "solve two problems with the current desktop metaphor: task management and comparison of multiple windows" (Robertson, 2000). While it was an innovative evolution of the

Information Visualizer, and had the benefit of its requisite technology being readily available, the necessary user support was still lacking.

But that may no longer be the case. The gaming industry has grown to a multi-billion dollar phenomena in which immersive sound and graphics and realistic, lifelike gameplay is merely 'par for the course'. Most users below the age of fifty were still in elementary school either during or after gaming's so-called 'Golden Age' (1978-86), and so have lived most of their lives exposed to the virtual environment, 3D UIs that have become a necessity for most games. Traditional virtual worlds like those embodied by Playstation Home have spawned massive followings, to the degree that these virtual communities claim a following of over one *billion* users (Nicole 2010).

If the combined gaming and online virtual world communities do not possess the inertia necessary to legitimize virtual environments as the worthy alternative/counterpart to a WIMP interface, then the proliferation of touchscreen and gesture controlled devices might provide a firm foothold upon which to move the masses. Specifically, users of mobile phones and tablet computers have grown accustomed to using touch, gesture, and in many instances voice as natural, intuitive, fun alternatives or supplements to the point and click mouse. Immersive apps in the genres of gaming and augmented reality may very well subconsciously warm this large user set to the appeal of virtual environments, effecting a subtle, yet significant paradigm shift even if the 3D UI is at best only an application-specific or even OS-level supplement to the traditional WIMP interface. According to Professor Heer: "I don't see WIMP going away – I just see other models of interaction being simultaneously realized" (Heer 2011).

Consider Microsoft Surface, initially released in 2008, as a cutting edge example of a device with a natural, intuitive interface. With support for up to 52-point multi-touch, object recognition and gesture control, it is perhaps the most technologically advanced commercially available touchscreen device. A similar example is embodied by the Microsoft Kinect, a 'body controller' extension to the Xbox 360 gaming platform. It includes a motion sensor, depth sensor and voice recognition, along with a suite of games that map natural and realistic motions to in-game events. The resulting gameplay enables augmented virtuality for immersive realism potentially never experienced in a mainstream gaming platform. As described by the Principal Design Director for Microsoft Surface, these products reflect a profound shift in the UI paradigm: "A new inflection point in human-computer interaction is upon us. Along with other technologies, Microsoft Surface [facilitates] designing for a new era in which emotional intent and intuitive interaction are the imperative" (de los Reyes, 2008). These sophisticated, natural UIs may serve to raise the baseline 'cool factor' for a user interface, leaving a void – an opportunity – for 3D UI to fill.

A perhaps more subtle example of virtual environments entering mainstream (sub)consciousness is in the domain of website design. The vast majority of websites model traditional printed media: newspapers, magazines and posters. While this may be appropriate for a media organization's website, in which the site is actually modeling their printed real world analog, or for a news organization like CNN, it is not the ideal representation for many websites on the Internet. And so many sites are breaking the traditional mold by modeling virtual spaces consistent with the site's identity, vision and/or value proposition. Two sites that exemplify this paradigm shift are Chiquita dot

com and Fashion Point (Appendix A). The user experience surpasses merely 'going to a website' – the feeling is more akin to entering a virtual world. Beyond merely immersive sound and graphics and compelling animation, the sites evoke a *feeling* in the user; the experience is compelling on many sensory levels beyond the purely visual. Fundamentally, what was appropriate and technologically feasible 20 years ago is no longer the case – highly expressive, interactive platforms like Flash, Silverlight and Unity3D are all but ubiquitous. While largely isolated examples of websites as virtual spaces have surfaced over the past several years, to this day they embody a great deal of the award-winning innovation and appeal in the industry.

Metaphor in VUI Design

"Metaphor plus magic" – Alan Kay on UI Design

A key challenge in virtual interface design, and in user interface design in general, is that of modeling. This notion is embodied by the so-called Principle of Metaphor of User Interface Design – "borrow behaviors from systems familiar to your users" (Talin 2008). Therefore it is beneficial to consider the real entity we are emulating with a UI element. In a computer system, that element will almost always represent or fundamentally contain some form of data, but how best to represent that information is far from trivial. For instance a book or article contains knowledge, while a pdf file models that printed volume.

The benefits of virtualization (translation from the real to the computer space) are quite clear: we retain the essence that most matters – if modeling a book, the textual information contained in its pages – and streamline processes of interest. For example, minimizing its representation for efficient storage and fast retrieval, serializing it so that it is searchable/indexable, etc. But in 'computerizing' a book into pdf form, we are stripping it of some of its essence – we are extracting from it some of its inherent expressive potential. Irrevocably, there is some intrinsic value that is 'lost in translation'. And therein lies a key challenge and perhaps the fundamental tradeoff in designing a VUI element – how are we to represent a real entity in virtual form (though it contain potentially vast amounts of information), retain its semantic value and expressive potential, and yet still not lose in overall value – usability, efficiency, or 'cool' factor?

One somewhat obvious, trivial solution might be termed a cop-out: simply present the user with a toggle or combination view – the ability to view the object in either 'real' (albeit through a virtual lens) and/or computerized form. While in this fashion the old and new can coexist, and therefore we can retain the benefits of both worlds while (hopefully) negating their drawbacks, the problem would benefit from a more careful analysis.

In comparing a virtualized entity with its real counterpart, we may more easily frame the comparison as such:

- In the computer model:
    - what are we gaining by virtualization?
    - what are we losing?
- In the real entity:

- what are important aspects that we might want to keep?
- what aspects are meaningful and/or useful to the user, and why?

Ultimately, our goal is to maintain the familiarity and semantic richness of the real world entity while still gaining the efficiency of a virtual representation, essentially combining the benefits of both worlds:

1) Real world semantics and familiarity ("Metaphor")
2) Computer efficiency and power ("Magic")

For example, in a simple WIMP directory, the typical 2D tree hierarchy representation lacks a visual-spatial notion of size and location; if we maintain efficiency of search and selection, we may gain semantic richness by representing the directory as a virtual library. A digitized version of an analog clock gains a smaller representation and potentially faster reading of time, for which it sacrifices a spatial representation of time itself. As before, the typical website is modeled after printed media, but a much more compelling alternative is to present it as a space consistent with the site (virtual store, virtual office, virtual dojo, etc.) itself, while still retaining informational properties. Some examples of common UI metaphors are:

Real Entity	Virtual Counterpart	Result of Virtualization Process
Frame (Television/ Photo)	Window	+/- Neutral
Photo (Miniaturized)	Icon	+ Efficiency of storage & access
Menu	Menu	+ flexibility and power + expressive potential
Hand (Gesture Control)	Pointer/Mouse	- expressive potential (from essentially infinite gestures in time and space to very limited 2D 'point and click' interface) + maintain (one to one) impedance mapping from arm to mouse
Document/ Paper	File	- no ability to manipulate non-trivially (annotate/highlight) without (arguably awkward UX in) MS Word/Adobe Reader-like program or (potentially costly) tablet + lose semantic (and immediately recognizable)notion of size + facilitates edit + Efficiency of storage/retrieval/search
File Cabinet	Folder/Directory	- lose spatial representation - lose potentially meaningful (semantically rich) metadata – postit notes, physical notes, color-coordination

		+ Efficiency of storage/retrieval/search
Newspaper/ magazine/ poster	Website	+ Efficiency of storage/retrieval/search - Semantic richness lost

In general, some major benefits of a digital representation are:

1) Efficiency of storage/retrieval/search
2) Flexibility & power

And some benefits of a real entity with respect to its virtual counterpart are:

1) Sensory feedback
    a. Visual feedback
    b. Haptic feedback –ability to *hold* something in your hand, feel its weight, texture etc.
2) Semantic purity – in general, beyond the obvious primary or secondary attributes, a real entity has inherently different meaning(s) that its virtual counterpart

Note that while there is value in modeling, it is simply not enough to motivate a design. As proposed by Hollan and Stornetta in "Beyond Being There", we shouldn't constrain our thinking such that we are only using a computer to replicate the properties of real world entities; we should think further, and advantage a computer's ability to "go beyond" what is possible in reality.

The Future

"I saw the angel in the marble and carved until I set him free" – Michelangelo

The need for a divergent UI is imminent. WIMP was okay in the 1980s, but (all kidding aside) it alone is too wimpy for the next generation of users. I believe that VUI is the logical counterpart to WIMP for a generation in which immersive sound and graphics, 3D acceleration and touch-screen/gesture capability will soon be an assumption of the past, rather than a hope for the future.

Now the operating system interface of the future will very likely present – by default – a basic WIMP interface. But beyond point and click, it will provide touch screen, gesture control and voice control interface alternatives. And as a visual alternative, it will provide a virtual environment operating system interface; the ability to switch between different Views of the same underlying data representation – from the 2D desktop to/from:

- A 3D desktop (Bumptop) or workspace (Information Visualizer)
- A 3D space such as a gallery (Task Gallery), home (Daddy.fr), or solar system (Helsinki)
- A user configured space consistent with his desires

It will include a 3D representation of programs, very likely as 2.5D windows, logically grouped into tasks. It will afford gesture capable, physics-based organization of passive and running programs. It will allow

augmented reality and/or augmented virtuality for the aforementioned ('configured space') customization; consider the following thought experiment enumerating the steps that may be necessary to realize such a feature:

1) Training – the user inputs (via image(s)/video) a real environment, such as different rooms within his home, to the OS
2) Calibrating – the user classifies entities in his home to act as UI elements in his customized OS interface (augmented reality)
3) Processing – the OS processes the training and calibration phases, creating a virtual model of the user's home (augmented virtuality)
4) Use – the OS allows the user to navigate through this virtual environment, and to organize his system (files, folders, passive and running programs) within his customized space

This fictional application does require training and calibration steps, but nonetheless has the advantage that – in customization – it may provide more meaning to individual users.

In any case, the fundamental question still remains to be answered – when do 3D UIs provide more value than 2D alternatives? Professors Card and Heer agree that the most important aspect is:

1) "When the data is inherently 3D – you're not creating or designing an artificial mapping from an abstract data set to space; rather, the spatial coordinates are given, and you need to reason about them" (Heer 2011).

Professor Card continued, to add:

2) When 3D provides semantic enrichment
    a. Perspective emphasis
    b. Spatial relationships to indicate logical separation/organization
3) When 3D enables physics-like interactions

(Card 2011). And so it follows from the first condition that the relative merits of 3D UI are very context specific. Therefore, some domains in which 3D UIs will continue to provide the most appealing application interface are:

- Modeling and design (CAD, 3D modeling, architecture, drafting, etc.)
- Simulation (Wargaming, as a historical example)
- Learning and training

Therefore the future will continue to have, albeit to a larger degree: artists using tablets as VUIs to a virtual drawing, painting or modeling environment; architects, draftsman and designers of all disciplines using touch and gesture as a complement to the mouse for designing entities in their respective

disciplines; virtual environments for simulation, learning or training in which performing the real activity would be prohibitively dangerous or expensive.

With the Internet as the medium, and sites as the vehicle, websites will continue to be a manifestation of the VUI imperative – an expression of the immersive and enriching potential of the computing medium. Others might agree, one day, that this is true Web 3.0.

As an example, consider an online store. While the site models a virtual Online Store – existing examples of real, natural navigation (Fashion Point site) and highly customizable shopping (Soccer Teddy site) foreshadow a more user friendly, 'real' UX that models a real shopping experience (without the waiting in line). While the site loads, you have the option of uploading either a 3D avatar or a 2D image (which might then be UV-mapped onto a 3D model) so that 'you' can actually 'try-on' clothes in real time. You might similarly go on virtual tours to intended travel destinations. User photo data will be mined and aggregated (using a tool such as Make3D) intro fly-through three-dimensional models, allowing for instant navigation of virtual spaces that model of any destination in the world. Users may similarly aggregate photos of their home into virtual spaces, inviting other friends over for a virtual – or holographically projected augmented reality – visit. In short, the virtual realm will continue to both emulate and transcend the real world in form, function and creative capacity.

What will present an even greater divergence in UI, however, is the inclusion of near-sentient automata as helpers in our everyday lives – virtual personal assistants (as exemplified by Tony Stark's "Jarvis" super-futuristic-computer interface in the movie Iron Man, or by Apple's 1987 Knowledge Navigator demo of future UI) that possess a sophistication of AI well beyond the capabilities of today's computers. While a full discussion is beyond the scope of this paper, they will be able to carry out a broad set of mundane or even complex tasks, have insight into our mood by detecting changes in our vocal tones or heart rate, be trainable so as to be optimized to our likes, dislikes and habits, and in general help us like loyal, helpful friends. That is, until Skynet.

References

"1 Billion Virtual World Users. And They're Mostly Pre-Teen Girls", Kristen Nicole, Silicon Angle, http://siliconangle.com/blog/2010/10/01/1-billion-virtual-world-users-and-theyre-mostly-pre-teen-girls/.

"3D or Not 3D: "More is Better" or "Less is More"?" Kevin Mullet, Diane J. Schiano, SIGCHI 1995.

"A Summary of Principles for User-Interface Design." Talin, ACM, 1998.

"Beyond Being There", Jim Hollan, Scott Stornetta, CHI 1992.

"CityWall Interactive Multitouch Display Now Has a Glorious 3D Interface", Adrian Covert, http://gizmodo.com/5062005/citywall-interactive-multitouch-display-now-has-a-glorious-3d-interface.

"CNET Top 5: Worst products in a decade." CNET.com, Tom Merrit, 4/30/2007.

"Distributed Systems: CS 244B course notes" (48-50). David R. Cheriton, 3/25/2011.

"Effective Visual Communication for Graphical User Interfaces." Suzanne Martin.

"Playstation Home", US site, http://us.playstation.com/.

"Predicting the Past." August de los Reyes, Web Directions South 2008, Sydney Convention Centre, 8/25/2008.

"The 50 Worst Inventions." Time Magazine, 5/27/2010.

"The Design of Design – Essays from a Computer Scientist." Frederick P. Brooks, Jr. Addison-Wesley, 2010.

"The Golden Age of Video Arcade Games", Wikipedia. http://en.wikipedia.org/wiki/Golden_age_of_video_arcade_games

"The Information Visualizer, An Information Workspace", S. K. Card, G. G. Robertson, J. D. Mackinlay, 1991.

"The Task Gallery: A 3D Window Manager", George Robertson et al., 2000.

"The WebBook and Web Forager: An Information Workspace for the World-Wide Web", S.K. Card, G. G. Robertson, W. York, Proceedings of CHI 1996.

"What's Real about Virtual Reality?" Frederick P. Brooks, Jr. IEEE Computer Graphics and Applications, November/December 1999.

Bumptop, http://bumptop.com/.

Email Correspondence, Professor Terry Winograd, Stanford University, 6/9/2011.

Email Correspondence, Professor Vladlen Koltun, Stanford University, 5/18/2011.

Fashion Point website, http://fashionpoint.ru/start_eng.html.

Knowledge Navigator (video demonstration), Apple Corp., 1987, http://video.google.com/videoplay?docid=-5144094928842683632#.

Office Hours Discussion, Professor Jeffrey Heer, Stanford University, 6/10/2011.

Office Hours Discussion, Professor Stuart Card, Stanford University, 6/8/2011.

Soccer Teddy, http://www.soccerted.co.uk.

Appendix A – 'Web 3.0' (Virtual Space) Samples

Virtual Space

- Come into the Closet Let's Dance (Ikea)
    - http://demo.fb.se/e/ikea/comeintothecloset2/site/default.html
    - An ikea (non-)shopping site that is a tour through several closet/room shopping areas. A wacky, eccentric, highly appealing and innovative concept site; merely for entertainment purposes, as it does not integrate with the rest of the Ikea site (shopping cart for items you like, for instance) or provide seamless navigation freedom, but is nonetheless highly interactive and just plain fun!
- Agency Net dot com
    - At the outset, a site that is very alive. Provides on-click zoom-navigation for an interesting shift to divergent content that does not depart from the site's look, but perhaps a little bit from its feel. Contains at least one sub-site that is itself a virtual space - http://agencynet.com/#/clients/ - which allows you to navigate downtown NYC to view some of the agency's esteemed clientele.
- Media Boom
    - http://mediaboom.com/flash/
    - A site that is alive! The viewing area is a postcard that sits atop other typical traveling documents (other postcards, a map), and on-click of the site's various pages the user is zoomed through downtown NYC to a location analogous in some way with the page content.
- Inbred Boy
    - http://www.s16504.gridserver.com/projects/inbredboy/main.htm
    - A portfolio site in the swamp motif, no less. While a somewhat crude, coarse example of a virtual space, and without truly smooth transitions between pages, it does nonetheless provide an immersive experience.

Virtual World

- Chiquita dot com
    - Tropical island
- Sean Kingston dot com
    - Excellent example of the kind of graphic appeal and feel that is possible; lacks a bit of the natural flow that would be present in a real space (must navigate 'back' in the browser to return to the 'world' view).
- Adobe Creative Mind

    - http://unit9.com/creativemind/
    - Purposed with illustrating the creative capabilities of the Abobe Creative Suite, this site is mostly just for entertainment purposes. Nonetheless, it is awesome! Entertaining, immersive, appealing, and (naturally) highly creative...

Virtual Gallery

- Tim Burton dot com
    - Wacky, weird site where the user actually walks through a museum/gallery-like setting (versus the typical corollary of viewing images in a mock/static photo album, like Facebook, for instance)

Virtual Desk

- World of Mr. Cogito
    - http://herbert2008.pl/cogito/
    - A very compelling, elegant example of what every scrapbook/photo/article site could be; a virtual desk, highly interactive with seamless, natural navigation – highly creative and compelling take on what most sites present as static, comparatively flaccid content. Excellent.
- Calamity Physics
    - http://calamityphysics.com
    - A site by Freedom Interactive Design
    - Very innovative, appealing desk virtual environment; very natural navigation around the site itself, with interactive elements in the environment that force the user to (happily) discover it; very cool and very fun.

Virtual Home

- Daddy's House
    - http://www.daddy.fr/
    - The future of operating systems? (Look past the pink...) Potentially one 'View' that might be available in OSs of the future; a more natural and visually appealing alternative to the directories & windows (GUI) model, with rooms representing different categories (Entertainment, Office, Kitchen, etc.) of the computer system/user experience.
    - Interactive, visually appealing, good, near-seamless navigation to different 'pages' (which are represented as 'rooms') ... a good example of the new imperative.

Virtual Store

- Fashion Point

    - http://fashionpoint.ru/start_eng.html
    - Very cool concept site for a virtual store; navigation is fast and seamless, as if you were traveling through an actual store (at walking, running or teleportation speeds) ... immersive, but also efficient enough so that the user can enjoy the experience while not feeling wasteful. All in all a very cool example of the sort of feel, and at the very least navigation, to which online stores could aspire.
- Soccer Teddy
    - http://www.soccerted.co.uk
    - The site allows for the customization and purchase of a single item – a silver teddy keychain. The teddy bear is highly customizable in real time, allowing the user to visualize their choices. This site is (in small part) an indication of the future UX of online stores.

Virtual Album/Scrapbook

- I'm a Cyborg, But That's Okay
    - http://cyborg.d-o-e-s.com/
    - A promotional site for a South Korean movie where the user essentially navigates through a highly interactive scrap-book/album. The navigation from 'page' to 'page' is seamless, animated, natural and (essentially) perfect. A very fun, immersive, highly-interactive site.

Virtual Portfolio

- Jay Z Greatest Hits
    - http://www.jayzhitscollection.com

A site completed by Agency Net, a web design & marketing firm who themselves appear on this list, this site is highly immersive – providing music throughout, video segues between pages, and a consistent look and feel. While arguably not a virtual space in the traditional sense (being the most static site on this list), it is nonetheless a highly appealing representation of an artist's portfolio.

Appendix B – Meeting w/ Professor Card
June 9th, 2011

Work was, and actually still is, very much before it's time.

Web design works against 3D; standard GL for 3D

Why didn't 3D UIs take on?
We thought market was going to explode for 3D; market wouldn't be ready for about twenty years.
Blitting aka bit blip aka bit block transfer - moves pixel bits around directly in memory.
Invented by a secretary at PARC!
Jim Clark - made 12 step pipeline so that clipping, (others?) could be done simultaneously. Made real time interactive 3D graphics possible (late 1980s) on a relatively cheap machine (50k)
GUI based bit blip, could get the next generation of UI based on the geometry engine. So the question was that would be?
Invented the basic GUI at Xerox PARC, then Apple started putting that out. They then expanded (leapfrogging to the next generation) this - trying to build the next major UI by adding the concept of multiple workspaces; the first one was called Rooms, and Apple made a bastardized version of it called Spaces, but they made 3 fundamental errors:
1. Application centered; must be task-based instead!
   Xerox Star and Apple Lisa were task-based
2. Overlap windows (overhead moving windows around, resizing - versus artfully tiling them); windows exhibit locality of reference, just like cells in a virtual memory OS

Many benefits to a task-based rather than an application-based focus; can better advantage locality of reference, reduced thrashing, can restore state with better organization and semantics (visual representation helps in recall for intermittent activities),

When working set exceeds the size of physical memory; if you do nothing, you risk going into thrashing hell (spend more time looking for pages than actually paging).

Doors go from Task to another; not as useful as you might think, because difficult to predict what door you're going to go to.

Need some theory about what is going on, so that you can reason about what is going on, and predict the "sweet spots of the design space."

How to present information so that it is more semantically meaningful/valuable?
Information visualization - hard to measure that in the abstract, but the value of representation is relative to the task - it must cue you to action. How to increase relative mapping, situational awareness, etc.

Bertan - french cartographer who had a system, a graphics design theory for representing information. Quantitative
Jacques (did thesis) - formalized that; two rules:
1) Expressible. The representation has to not necessarily be efficient, but basically not anti-alias (not mislead with your representation of data); visualization has to be in some order, but data may not have order, so you

potentially cause a semantic aliasing by adding information that really isn't there, essentially. Represent what is in the data, and don't represent what's not in the data.
2) Efficient. Some representations are more visual than others.

Can divide behavior into powers of ten seconds
0.1 seconds ("psychological moment" is magic number); if two sounds get to ear within a psychological moment, they are perceived as simultaneous actions; need to be at least this fast (10 fps) for animation
1 second (700ms) - reaction time; and so interacted with display, you can't modify instantly (think clicking a link or a node on a tree) because it doesn't give feedback as to what actually happened, but if you animate the tree, it gives that feedback; you want to animate it within a second, because that's the amount of time it takes to react to something - you couldn't do usable work in that time even if you wanted to, is the point (trying to cram as much information as possible into as many portals as possible).
10 seconds - unit task; nominal amount of time it takes to do like a text-editing operation
100 seconds - search, information foraging, trying to solve a problem (heuristic search, which goes way wrong if you take the wrong path, because you add expontentially to the time). Can get a qualitative change here.

Can play tricks with perspective representation - things closer more relevant, etc.; possible to cram more information in in that way?

Web Forager - one of the most notable things is the rack where the books are

Limits the user's mobility in the virtual world (in Information Visualizer) so that they can't get into trouble (problems of occlusion, etc.) or enter a disadvantageous space

Time Warp (Mac app - 3D UI to their OS)

Big question - what is your interest in it? (he asks me)
Big question i'm asking - when are 3D interfaces more appropriate than alternatives? In virtualizing something, there is semantic essence we're losing; we're gaining efficiency of represenation/storage/search/retrieval, but there is some essence we're losing - what is it that we're losing, what is it that is beneficial about the computerized representation?
"Well you generally win when you talk about something that is already intrinsically 3D." Take Google Body as an example.

On 3D interface to a copier (1996)
Navigation, low training time
Six key concepts
1. Overview
2. Physical Metaphor
3. Revealed ___ Articulation and Zooming
4. Manipulation Augmented Property Sheets
5. Animated Transitions
6. Combinatoric
Used Information Visualizer to prototype...
(on 50k SGI machine with less graphics capability than iPhone)

Alan Kay - "did realistic representations and then added magic to it"

3D Book - developed by Professor Card; generates true 3D representation from a PDF; 4 levels of representation so that animations are fast enough; when turning a page go to a higher (less detailed) level, and when viewing the page normally go to the lower (most detailed) view.

If 3D UIs weren't successful before their time, then why aren't they successful now?
"You can be way too early"
Such that it ends up hurting the field in the long run?
Well, two things are needed:
1) Technological readiness
2) Dominance of design (GUI interface been more or less the same in the last twenty years)

Take the typewriter (Remington 2000 was the groundbreaking trend-setter) - many early competing models were very very different (different key configurations, different designs, etc.) - after that the typewriter had a more or less common design. A lot of people liked the Devorac keyboard (better arrangment of keys for like 10% faster typing, apparently) competed with the QWERTY. And so for such a minor incremental improvement (10%), it will never displace the dominant design. Going against the accepted standard. To overcome that resistance in the marketplace, you have to show a real advantage over the traditional GUI desktop.
Professor Card still believes it can; he likes the 3D desktop surface, possibly with physics so you can do some interesting effects.

Time Machine on Mac

Value in providing 2D/3D toggle? What he told George Robertson - three things you need to get into the OS:
1) application redirection (his solution to that was to provide 3D components to a 2D interface, and to then colonize the interface with these gadgets, such as the 3D Book, for example)
Multitouch provides a cool interface - the non-dominant hand provides an easel/foil to support the dominant hand's actions (pallete of a painter, versus the paintbrush)

Value in a gesture based interface?
Yes, but you get tired!

Engelbartian principles - need four hands. Note that two-handed interfaces enable a 3D environment.

The Alan Kay formula: metaphor + magic

"The wrong way to design 3D interfaces is just off of metaphor. Now, there is a role for metaphor (metaphor + magic, the Alan Kay formula), but what you want is to have some of the analysis of the mechanics that underline it."
Want good value for (locality, cost of knowledge characteristic).
The current GUI moves from recall to recognition (CLI need to remember the commands, versus in a GUI you can open a menu and recognize the right command).
3) input on output; asymmetric - can take in a lot of information (pixels on-screen are changing and "washing over you with information) but can't output as much
Whirlwind in 1950s, became SAGE North American air defense (zillion dollar) computer

Book - thickness of the pages tells you information; 3d representation can tell you information, intrinsically.

One way of making something people like - go for familiar metaphors.

Me: value of a user interface - 1) usability, 2) efficient (speed + power), 3) cool/fun/x-factor
Professor adds 4) ease of use; "It depends on the group that you're going for. If you take Englebart - Englebart could care a less how easy to use it was. He was doing an I/F that would allow the elite members of society to get performances that were never before possible to solve society's problems. And Englebart said this to me a lot - "You don't not have violins just because ordinary people can't play violins! Out of a violin or a harpsichord you can get experiences which are possible no other way." And so he was building a machine to do that, and the chord set..." Faster, and took about ten hours to learn.
Christiansen defined disruptive technology as "one that can be done by a person of less skill in a place of greater convenience." Applies to a GUI I/F. Market is a pyramid, go down each level there are ten times more people.

Important to find what users find appealing, but then you also need objective measures to find out what is best. The market doesn't necessarily choose the best thing in the long run but it avoids the worst thing.

"Games are one way in which new interface technology could come in, because there's sort of a critical mass for it." Gone into the military for that reason. "Someone has put in the money to develop the engines, which makes the adoption of it in this other market cheaper, maybe cheap enough to have a shot at it." "3D interfaces work well for the games, because they have a functional role in the games, but you've got to show that you can do something with them that you can't otherwise."

Gestures as abbreviating actions - 'flicking' things into place

"I was a pop and flash guy - never missed a chance for gratuitous animation."

Consulting Professors for Minority Report - could be MIT Media Lab, or Michelle and Wendy (visiting Profs here).

Augmented reality - where you project pieces of paper on your desk, and you can do hand gestures, and you can input documents into the computer, and extract them from the computer.

Apple Knowledge Navigator Tape - radical interfaces for the future. Apple did this tape - supposedly I/Fs of the future, and they tagged it "I/F of the past"; futuristic stuff was voice interface stuff "that you were never going to do", and virtual desk as augmented reality I/F. Pierre Welmer's augmented reality I/F virtual desk.

Re. voice recognition - more relevant (as Prof Heer said) when your sensory modalities are spoken for; voice recognition great while you're driving; important part of I/Fs that you can do searches quickly (shout "pizza" into your iPhone, and you'll be presented with all pizza restaurants in your local area).
Requip or Requill - note taking program for iOS(?).

"In one conversational term. There's kind of a threshold - voice is the way in which you can get the largest branching factor for the smallest amount of time. So if I say something and stuff comes back and I say stuff I can get that loop going faster and faster; if I can get that loop going fast enough, it's almost as if that knowledge was in my head."
Mike Williams thesis showed that even for stuff that is in your head (long term memory), you have to go through a process like this - we're getting to time constant that are sort of similar. "If you can do this, you have sort of a strap-on cortex of knowledge that you don't have but you do, if you can get the time constant down fast enough - and voice is a key thing for that."

Summary - when are 3D I/Fs more valuable than 2D alternatives
1) When stuff is already intrinsically 3D (take Google Body, for instance - harder to do with a 2D representation) vs. using 3D for compaction (2D directory to 3D cone tree, for instance)
2) Stuff where you can get 3D effect that does some sort of separation (Time Warp on Mac)
3) Place where you can use perspective (like in Web Forager - can keep information constant & ready on bookshelf, also Task Gallery that George (other Prof?) is doing)
4) Places where you can use physics, like tossing around windows and such (coolest one he's seen is where desktop is sort of slanted (Bumptop?))

Re. real vs. virtual metaphors...
Familiar metaphors are useful, sure. Helps in organizing.

"The art of interface design is the art of finessing the stuff you can't do."

Appendix C – Meeting w/ Professor Heer

June 10th, 2011

Largely settled (past) controversy - on the relative merits or lack thereof of 3D representations...
Some issues:
1. occlusion - things blocking other things
   If there are good 2D representations, they might be more beneficial
2. take more time to work with - rotating/navigating 3D space is physically and cognitively more complicated
3. perspective distortion for certain kinds of data - distance judgments; foreshortening, how things align in the axis grid (can be misleading) ... ("there are a number of perceptual and interactive issues that often lead people away from 3D representations");

"where 3D visualizations provide more value - where the data is inherently 3D - you're not creating or designing an artificial mapping from an abstract data set to space; rather the spatial coordinates are given, and you need to reason about them."
Different research literature - scientific research literature, statistical data graphics. "Sure you can take that abstract data and map it into a 3D display, but in most cases studies have shown that the 2D displays - if you can do them well, will be more effective."
One of the takeaways - "what are the interesting and valuable and really important types of data that do have that natural 3D element to them ... which would be a more effective candidate for the technologies you're describing."
- Architectural planning, landscape design - domains where you have to reason about 3D properties - have "more traction and a clearer value proposition."

Real entities vs. their virtualized corrollaries
Illusion of 3D as projected on 2D monitor
Monitor in 3D space
Real vs. virtualized - "how does that object come to BE in someone's perception?" Made of atoms, holographic projection, "more interesting perspective is how it is experienced by the person, regardless of how that perception came to be."
What does it mean for objects to be physical, and what it means for an object to have haptic feedback.
Spatial memory - the way we organize objects in space is actually part of our thinking process as well. "How we structure our environment is part of what makes us smart."

Why were Data Visualizers not successful?
Were successful in research way. "By design, it was ahead of it's time."
Paid for cutting edge, custom systems that people just couldn't afford; "commodity computing hardware wasn't nearly caught up with what was required to support those environments." For most of what they did, "3D was an unnecessary component". More 2.5D, like Perspective Wall. Tree visualizations - suffered from occlusion (lots of elements you couldn't read). Simplicity in UI is key - if I'm solving a challenging problem, I want the tool (the UI) to get out of the way; hopefully my reasoning about the problem is dictating the pace, not some impedance imposed by the UI.

And so given that hardware and game engines are readily available now, why are they still not available?

Cool factor - think about it more deeply. If you start to peel back the layers of the onion - "are we talking about people's emotional response - affective pathways of the brain; what does that mean?"
Also important to think about is the context of the task. Depends on context.
Different models - people trying to make sense of data
1) The scientist - you have questions you're trying to reason about or hypotheses you're trying to prove, and the data is going to help you get to the bottom of it (scientific endeavor)
2) Business person - trying to optimize my operations, or prevent disaster. (operations research)
3) Casual user - personal life - music you've been listening to, contacting on social media; engagement with data, but it's also entertainment.

Bloom.io - data visualization in a consumer and game-like setting; using game engines to build visualizations; take music you're listening to and turn it into the orbit of planets around the sun; more about personal reflection and social interactivity and less about accurrate communications about patterns in the data. Context of exploration/play/reminiscence.

Minority Report - really cool, but "worst idea I've ever seen" because it's not ergonomic - people's arms would be falling off by the end of the day. Ergonomic nightmare.

On Tony Stark's computer - appealing for the future.
"All depends on what you're trying to do."
"What are things I want to be able to accomplish."

"Speech is really interesting in places where sensory/output modalities are spoken for."
- Driving
- Cooking
(or distance really) - per the demands of that particular environment

Natural successor to a WIMP interface?
"I think we already have it" - successor though? "I don't see WIMP going away - I just see other models of interaction being simultaneously realized."
Desktop apps and web apps are designed differently - how we navigate (single/double click), do we expect to be monitored/logged, do we expect to communicate with other people? We form these perceptions just by virtue of the chrome of the application.
iPhones/Androids/Tables - gesture based interaction. Interesting to note the ways in which touch will work it's way back into the desktop. Will future monitors be touch? Maybe, if so, will be relatively limited use. But if you think of other things - what's the next version of the drafting table?
"Again, it's task-based - what is the type of work I'm trying to do? And so how I design my environment and my tools to support those activities?"
"Mobile and gestures are big... in a sense that's the next thing, but in a sense they're already here, so what's the real next thing?"

Would these be termed NUIs? "Presumptuous to call it NUI" ... "designed, symbiotic form of communication" that was not natural. NUI as a term was really "a marketing vehicle for Microsoft."

"Impedance match - using a mouse doesn't slow us down". A mouse is the most natural extension of our arm - using a joystick or a trackball is much more

clumsy; some of Professor Card's most famous early work was in suggesting this phenomena. Excellent example.

From an entrepreneur's standpoint - cool vs. better not really an enabler; even "X %" better or more efficient, however you quantify that, is not really compelling. What is compelling is something that is an enabler. "Here's an area where you couldn't do it before, and now you can. You've got yourself a business case."

On Helsinki Project
"As these things filter more into day to day life you're going to see more diversity." In every day life I'm not necessarily trying to be more efficient, I'm trying to do things that are beautiful/compelling/playful.

"Surface computing" really the logical evolution of a WIMP interface. Works really well for consumption (reading/watching/listening)

Paper that Professors Kolton and Klemmer - "Beyond Being There" - thinking not about how computers replicate things you've done before, but how do you think through that to enable experiences that maybe weren't possible before? Pushes us to be more innovative, and think about the unique role that computers can play, without falling back on the 'it's more natural' - no, it's not copying life, it's different - but how do we take that difference and use it to do something that we simply couldn't do before.