THE IMPERMANENCE AGENT
Project and Context

Noah Wardrip-Fruin & Brion Moss
(with a.c. chapman & Duane Whitehurst)

Introduction

The Impermanence Agent explores a model for literature/art that is specific to the Internet-connected, office or home personal computer and a peripheral model of attention. It is not like a book, which we pick up and read. It is not like a painting, which we visit and inspect. It is not like a hypertext artwork, on which we click. And, similarly, it is not like a story on a Web page, a picture on a Web page, a Web-based hypertext, or any other stand-alone artwork that is intended for our full attention. Instead, *The Impermanence Agent* is an artwork that operates as functions of the user's Web browser. The browser is approached as a daily computer tool, in which the artwork becomes a peripheral part of the daily browsing experience.

When *The Agent* is engaged, user browsing causes a story to be told. This story is experienced in a corner of the PC screen, over a period measured in days rather than minutes. And while it is presented in a small Web browser window, *The Agent's* story does not act as other Web content. It will only move forward as the user clicks on *other* Web sites (those not associated with *The Agent*), and there is no way in which to "click on" or navigate *The Agent's* content directly. Simultaneously, *The Agent* monitors the user's Web traffic, and here takes on the semi-autonomous character from which parallels with more traditional software agents (discussed below) can be drawn. *The Impermanence Agent* continually alters its story using material from the user's browsing (and occasionally intervenes in

browsing itself). Over the time *The Agent's* story is told, the story's contents are altered until they are nearly entirely determined by browsing actions of the individual reader.

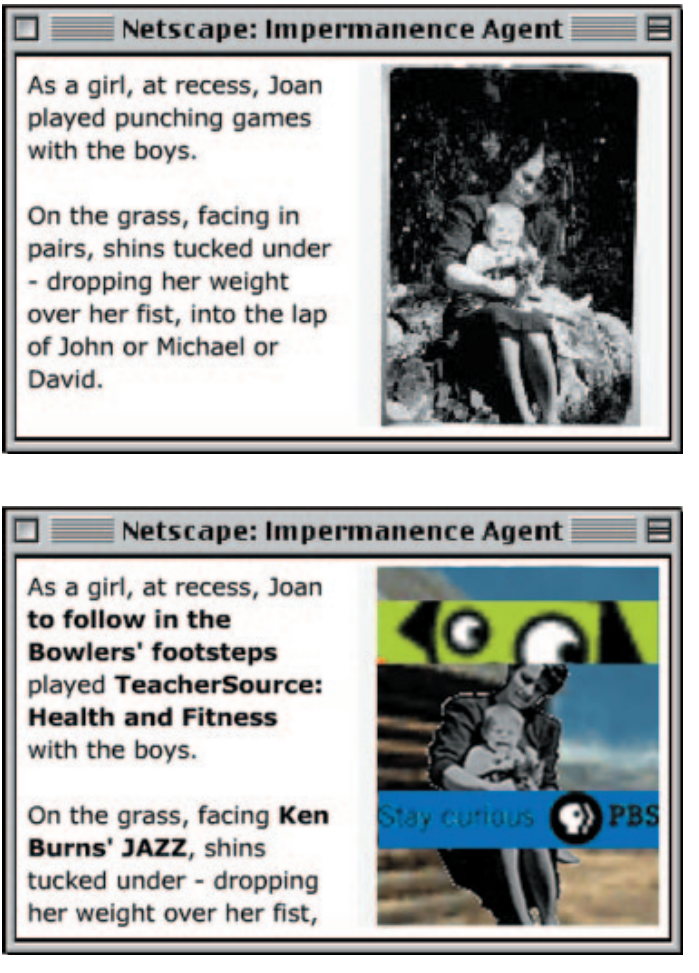

Figures 1 & 2. *The Impermanence Agent* window, and a sample of story contents after customization.

Part I: A Storytelling Agent

This section describes our collaborative development of *The Impermanence Agent*, a Web-based storytelling agent that customizes its story for each user (figures 1 & 2). *The Agent* is both a critique and employment of a number of Web technologies and mythologies. It begins with a story of remembrance and loss, a collage-story merging family history and personal narrative into relatively traditional fiction. Information from the user's Web browsing is used to customize this story, and this customization continues until almost none of *The Agent's* story remains. *The Agent* project bears some relation to Phil Agre's formulation of "critical technical practices" – and also to games the authors have played for decades.

Playing

We began playing together before we were even school age – when Brion's mom became Noah's babysitter. We like to make things, and we have a habit of finding materials in inappropriate places. One day we opened up Brion's front lawn to get two sprinkler assemblies for a robot we were building. They were brass, and the perforated heads rose up – clack – when water passed through. No one was pleased with us.

Brion's father worked for Ford Aerospace, and Noah's wrote for *Infoworld*. When we opened up software to make our own versions, to make parodies of them, no one at home minded. Brion was almost always the one at the keyboard, and we talked continually as he typed. We mostly rewrote the text portions of games, creating new narratives from the repurposed parts of their "interaction mechanics" – new meanings using the same programmatic structures. And we wrote the Klingon warships pink.

Recently we have been playing again, making a project called *The Impermanence Agent* (collaborating with a.c. chapman and Duane Whitehurst). We're playing using the tools of the ecommerce Web, within the technocultural narrative of "agents", and making an artifact that operates as a Web-based customizing storyteller, driven by the daily act of browsing. We have designed this agent to tell a story of impermanence, and its operations and technologies are part of its story.

An Agent

The Impermanence Agent project began in January 1998, when Stephen Pusey and Yu Yeon Kim invited Noah to contribute to *Omnizone – Perspectives in Mapping Digital Culture*. *Omnizone*, to be presented as a curated website on *Plexus* (www.plexus.org), sought mappings of digital culture from artists, curators, critics, and other cultural practitioners.

Over Vietnamese dinner in San Francisco, we started talking about *Omnizone*. We talked about ways that businesses were offering to provide maps of the Web – to provide experience guides that steered users away from the Web's "problems". These supposed problems seemed mostly rooted in the Web's overwhelming openness, resulting in "too much" information, no "quality control", and the impermanent roil of sites going up, coming down, and changing. We talked about websites like *Hotwired's* "Newsbot", *The Mining Company*, and *Ask Jeeves* that used "agent" metaphors to describe their services. We talked about agents as anthropomorphized maps, mediating our interactions with digital space, and giving us a means of circumventing these problems of digital culture. We thought we'd make a Web agent for *Omnizone* that provided an alternative, perhaps inverted, mapping. We thought we'd provide an experience emphasizing, rather than hiding, the supposed problems of the Web – the broken links, the fictitious data, and the things that don't interest us (or make us uncomfortable).

We did not embark on creating what, by most technically-oriented definitions, would be considered a "real" agent. For example, *The Impermanence Agent* completely fails to conform to Alan Kay's oft-cited interpretation of the term's founding meaning: "a system that, when given a goal, could carry out the details of the appropriate computer operations and could ask for and receive advice, offered in human terms, when it was stuck" (Kay 1984). What we sought to engage was not this type of construction but rather an agent that is as much a product of narrative as technology – narratives found in discourses ranging from Internet marketing to cultural studies to computer science.

Phoebe Sengers, whose work to bridge computer science and cultural critique will be discussed later, provides us with a useful formulation for discussing agents in narrative terms:

> I will take agenthood broadly to be a sometimes-useful way to frame inquiry into the technology we create. Specifically, agenthood is a metaphor we apply to computational entities we build when we wish to think of them in ways similar to the ways we understand living creatures. Calling a program an agent means the program's designer or the people who use it find it helpful or important or attractive to funders to think of the program as an independent and semi-intelligent coherent being... When a program is presented to its user as an agent, we are encouraging the user to think of it not as a complex human-created mechanism but as a user-friendly, intelligent creature. If 'actually' some kind of tool, the creature is portrayed as fulfilling its tool-y functions by being willing to do the user's bidding (Sengers 1999).

In the case of *The Impermanence Agent*, we set out to construct a Web artifact that we would describe as an agent, that would function as an agent engaging in a discursive practice of our design, and that would simultaneously serve to question the grounds and goals of other Web artifacts that are presented as agents. Sengers's formulation is useful because it describes our project, and those that fall under Kay's definition, as well as those projects we work to critique.

As mentioned above, projects such as *Ask Jeeves* (www.ask.com) are as much products of narrative (of portrayal) as technology – and it was aspects of narrative that we primarily sought to critique, such as implied labor relations in some agent marketing language. Of course, in much agent marketing these imagined relations are more explicit than implied, and do not require unearthing prior to critique. They need only be brought to our attention. For example, take this agent description from the current incarnation of the *Ask Jeeves* website, "As the world's first Internet Butler, I'm always at your service. I've made it my mission to humanize the online experience by making it easy to find the most relevant information, products and services". The kinship of agent marketing blurbage with fantasies of human servitude[1] becomes even more clear if we read the following text, recently removed from the *Ask Jeeves* site:

> [Y]ou probably view the Internet as being somewhat shaggy, wild, even downright uncivilized. (Jeeves thinks it no mere coincidence that we call it the "Net", implying that we can get tangled up in this "Web" of information. Think about it.)

> Well Jeeves is here to tame the Net for you. He loathes disorder and lives to serve. He wishes only to make your experience of the Net more civilized. . . .
>
> In short, Jeeves is delighted to be your personal guide through this expansive – sometimes too-expansive – resource. A butler for as far back as he cares to remember, Jeeves now lives to serve you and you alone. He's friendly, smart, periodically witty, and at your service 24 hours a day, 7 days a week. He never requests time off for holidays, and has yet to call in sick. Type a question for Jeeves just the way you'd phrase it if you were asking a particularly knowledgeable friend. Dutiful butler that he is, Jeeves will immediately and intrepidly begin navigating through hundreds of millions of web sites to find just the answers you need, and deliver the results quickly and thoroughly to your screen.
>
> Jeeves does not accept gratuities.

In seeking to construct a Web artifact that both functions as an agent and operates to critique phenomena such as *Ask Jeeves*, we would seem to closely echo Phil Agre's formulation of a "critical technical practice". He writes, "A critical technical practice will, at least for the foreseeable future, require a split identity – one foot planted in the craft work of design and the other foot planted in the reflexive work of critique" (Agre 1997a). We agree that finding a way of talking about – and a home for – projects with a foot in each of these types of work is essential.[2] However, we soon found that we needed a third foot.

Noah and Nana

Noah: Nana was my grandmother. After she died, in 1993, I began working on an essay about media, particularly hypermedia, and dreams of permanence. It was in her house that I'd first read Ted Nelson's 1974 *Computer Lib / Dream Machines*, while on a school holiday (Nelson 1986). In that book, and later more fully in his 1981 *Literary Machines*, Nelson lays out what has become a common vision, decades later, of the future Internet/Web (Nelson 1981). In this vision, in a not-so-distant future, we will read and write (view and draw, hear and compose) almost everything from and to a world-spanning computer network. Everyone will have the ability to produce their own documents and connect them to any other public docu-

ments. The author may constantly create new versions of her or his own document, and individuals may create their own versions of any public document; public connections made between one version of one document and another version of another will usually automatically place themselves in all the extant versions. Historical backtrack and degradation-proof storage will allow us to visit any version, any moment in the network's history. To have the ultimate archive, and yet have each element of this archive constantly in process. Dynamism without loss. Impermanence enfolded within permanence.

At Nana's house, in the summer of 1993, we were left with her letters, photographs, collections of news clippings. Already, for many pictures, no one knew the people in sepia tones and odd-shaped haircuts. The letters were those that belonged to her, not the ones she wrote. A lousy way to try to know someone, the paper trail, the box of letters and photos. But I was drawn to it, and not alone in that. It seemed she must be in there somewhere, in the possessions, in the records, in the writing.

As it turned out, it was much more at the memorial that I saw her, in the ways we treated each other with her caring. And yet, while I began to see the paper trail as no dodge for impermanence, I still felt compelled to write stories of Nana, to make a deposit in her name in the collective box of letters.

A Story

In early 1998 we contacted Duane Whitehurst, who had the previous year worked on a project called *Raku Writing*, organized by Noah for the *Port* exhibition. *Port*, curated by *artnetweb*, was presented online and at the MIT List Center for the Arts as the first exhibition dedicated to time-based performance on the Web (musEleanor 1997).

Duane was skeptical that a purely reactive architecture, giving continual commentary on user web browsing, would serve the purposes of our agent project. He felt strongly that *The Agent* had to provide, if not a "plan", a *context* for its actions – and that the most appropriate context, considering Noah's background and inclinations, would be a story. We agreed.

In time, the beginnings of Noah's 1993 essay from Nana's house became the first sections of an essay about *The Agent*, and retrospectively seemed an early source of the project concept. Further, through this now-explicit tie to Noah's grandmother, the project began to take on the shape

of another Web phenomenon – the Web shrine (Bernstein 1997). The first story used for the project was abandoned, and Noah began to write a collage story, combining fiction with stories of his grandmother and found texts. We also, in a group that now included a.c. chapman, agreed to create a "framing voice" for *The Agent* that would exist simultaneously (for more on this, see section 2 of the next part of this essay). As these elements took shape (and as we worked to define *The Agent's* visual form – see sidebar) the project began to come together in our eyes. We made our first public presentation, using a mockup created by a.c., as part of Thundergulch's "Twilight @ the Wall" series in the spring of 1998.

The Agent we presented at that time was firmly a storyteller. Duane's insight had led us to understand that *The Agent* project needed more than a "foot planted in the craft work of design" and a "foot planted in the reflexive work of critique". It also needed a foot planted in the work of personal narrative and storytelling, in the work of art. Such a formulation brings to mind Simon Penny's discussion from a panel at *Digital Arts and Culture 2000*. At this panel he, Sengers, Noah, and others spoke about critical technical practices:

> I've maintained for some years that well informed and productive digital media art practice demands a tripod structure, three legged things being inherently stable. These parts are: cultural studies (particularly science studies) and critical theory; technical knowledge and skills in both hardware and software; and artistic and creative skills, by which I mean not simply how to wield a paintbrush, but the basic skills of idea development, project design and development, aesthetic reasoning and self criticism. (Penny 2000)

In part III of this essay Penny's position will return.

Customization

The Agent's first actions within the context of its story were an outgrowth of Duane and Noah's work for *Raku Writing.*[3] The Raku project involved real-time writing by authors at distributed locations, which was posted, paragraph by paragraph, to a Web page. Audience reaction would then determine whether programmatic "stress" would be placed on particular sentences in each paragraph – and the resulting altered paragraphs would

be posted to the Web page. In the final performances, a second live writer rewrote each stressed paragraph, without looking at the initially-written paragraphs, resulting in a further set of effects (figure 3).

The Raku stresses involved literal translations of many common rules for "good writing". For example, using the "hyponym" feature of *WordNet* (a lexical database freely available from Princeton's Cognitive Science department) writing was automatically made "more specific". (Fellbaum 1998) Similarly, one of the most successful stresses employed by the Raku project was a form of lightweight sentence recombination – placing portions of previously-written sentences into the current one – to cause the automatic return of earlier word constructions and imagery. This latter technique became the basis of *The Impermanence Agent's* customization of textual content for each user – a technique that might in the art world be described as complex collage, or in the world of agents as an "extremely lightweight" intelligence model. We have also developed a complimentary method of customizing images, which is even more lightweight. The next part of this essay is dedicated to a discussion of *The Agent's* techniques, which function both as design and critique, and which result in customization that continues until none of *The Agent's* original content remains.

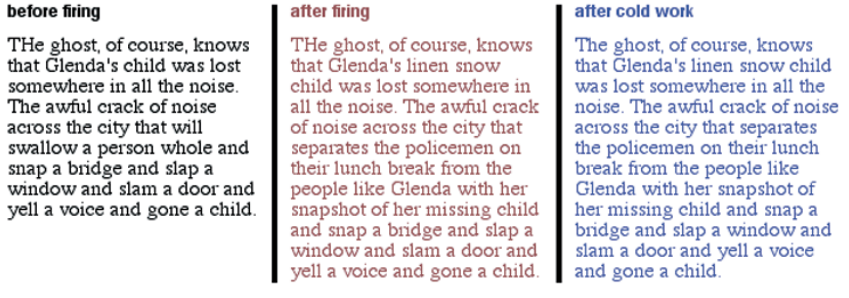

Figure 3. An example passage from *Raku Writing* in three versions: as written under time constraint, as "stressed" by audience reactions and programmatic devices, and as re-written by a second time-constrained writer working with knowledge of the second text but not the first.

For here, let us discuss in a little more depth the idea of customization, which by some is seen as the great virtue of computer-provided information over traditional distribution mechanisms. On the Web we have seen customization in several forms, especially the plethora of "My<X>" sites such as *MyYahoo!* which give the user a customizable space of largely corporate information with which they can be greeted. On *MyYahoo!* this currently includes such tidbits as "Barnes and Noble Showcase" and "Movie Releases" for those anxious to learn of the next mega-media offering. Another important Web customization type has been that of the "recommendations" that form the experience of those who do not prevent commerce sites such as *Amazon.com* from tracking their movements and purchases. These sites compare an individual's data against that of other customers, presumably consider what the site is in need of selling at the moment, and then suggest to users what they might express themselves by purchasing next. But the pinnacle of customization dreams seems to be, for one reason or another, customized news. Perhaps this has to do with the MIT Media Lab's "News in the Future" project, which has been highly publicized for more than 15 years. Walter Bender, the current Executive Director of the Media Lab, writes of the project on "Walter's Page":

> Modern telecommunications is leading us inevitably to the smallest news product imaginable: the personalized newspaper, or *Daily Me*, whose content has been tailored to meet an individual's needs and interest. Computerized "butlers" or "agents" are acting on the reader's behalf, culling articles of interest from traditional and non-traditional news sources, before sending them down the wire to the reader's home. Luddites see the *Daily Me* engendering a fragmented world populated by self-interested myopes. They argue that editors should continue to publish articles that establish the point-of-view for the community. They want news pushed upon them. The *Daily Me* proponents want to pull news in (Bender 2001).

Reading this, one might wonder: How many US readers would, after the "butler" learned their preferences, get a *Daily Me* with no news at all about Africa? With no news about human rights abuses? With no news of the world that might make them uncomfortable in any way? And then one might wonder how different this would be from our current news. US destruction of the Sudan's primary source of medication, in a mistaken

1998 missile attack for which appropriate reparations have not been made, is news about Africa involving human rights abuses that might make readers uncomfortable – and the ongoing death toll has received almost no coverage in the US media. So what does the media report? Well, despite its name, *Newsweek* gave its front cover to the movie *Men In Black* the week that Hong Kong returned to China. Barbara Ehrenreich, who writes for both *Newsweek* and *Time*, reports talking with the top editor at "a major national magazine" – pitching him again on an essay about the feminization of poverty. Finally, he said, "Okay, do your thing on poverty. Only make it upscale" (Ehrenreich 2000). Perhaps the *Daily Me* would be only different, for most readers, in that it would be a slight exaggeration of these tendencies.

But this slight news content change is not what excites the proponents of customization – it is the larger change in the means of getting news. It is that we will have the new freedom of pulling in the news we want, rather than having it thrust upon us. It is that we won't have to do any of this pulling ourselves, but rather turn such manual work over to a servant. And that this servant will, when facing us, be a mindreader – and when facing the news, an expert editor (as Bender writes later on the page, "such systems can personalize articles for individuals and communities of readers"). Our information servant, our agent, will learn what we like and present us the world in that image through judicious customization.[4] It is this line of imagining that *the Impermanence Agent* takes as the starting point for its customization processes, as outlined below.

Part II: Design and Presentation

This section describes the artifact called *The Impermanence Agent.* This is an artwork specific to the Web-connected PC, integrated with the user's Web browsing, and meant to be experienced peripherally. There is no place to click on *The Agent's* content, and no way to simply "visit" the piece. *The Agent's* content is displayed when the user actively browses other sites, and this browsing across the Web is the only means of interacting with *The Agent. The Agent* both adds content to the pages users visit and draws on the material from visited pages to create the unique content it displays to each user in its browser window. *The Agent's* content cycles

and repeats, slowly altering, continuously displayed in a corner of the user's desktop whenever browsing takes place. After about a week *The Agent's* story has been told.

Scrapbooks

A Web user is always viewing a scrapbook. After a Web server responds to a series of HTTP requests, the images and texts and other media it has sent (which together make up the pages requested) are stored on the browser's computer in a cache – a scrapbook – and collaged together for viewing in the browser window.[5] The cache has a maximum size, and in order to make room for more pages it is regularly purged of older material.

When a user browses with *The Impermanence Agent* a second level of scrapbook is created. This is a scrapbook kept for the user by *The Agent*. The user no longer makes HTTP requests directly to Web servers. Instead, all requests are made through *The Agent*. *The Agent* gathers the media requested by the user, pastes it into a scrapbook kept for that user, and then passes this down to the user's cache and browser for display.

The technology used by *The Agent* for this activity is known as a proxy server. Proxy software was first put into use as a building block for "firewalls" – virtual barricades between users and the (apparently) hostile and frightening Internet. When it is too risky to open one's browser directly to the Internet, one instead uses an agent – the proxy server – to perform the work of moving across the Web and fetching pages. The user/browser must ask the proxy to fetch the pages for it, pages that the browser cannot fetch for itself because of security policies and network design.

The pages requested may not come, or may not come as requested, because proxy servers have become more active agents over time. They are now also used to filter out "bad" content. This bad content ranges from potential viruses (or any program a user might install on their machine without authorization) to pornography (or perhaps any site that uses words such as "breast" – even in a health context) to sites that criticize corporate policy. Some proxies, rather than simply blocking access to pages, rewrite them to remove elements that are considered dangerous. Proxies are another part of the attempt to protect us from a Web that is too wild, too dangerous, too (as Jeeves would have it) "shaggy" to experience directly.

The Impermanence Agent's proxy is not a protective one. Instead it monitors the requests that the user makes and records them in its scrapbooks. *The Agent* keeps these scrapbooks (one for each user) for two reasons. The first is to employ each page of the scrapbook as a communication space from *The Agent* to the user – a space where occasional elements are added based on *The Agent's* opinions (rather than removed based on corporate policies). The second reason *The Agent* keeps scrapbooks is to learn of the user's browsing – that is, as a means of communication from the user to *The Agent*. *The Agent* employs this information to tell stories that draw on the materials the scrapbooks contain. *The Agent*, then, is an opinionated archivist, adding comments in the moment, and extracting meaning from the collection.

The alterations and insertions are simple, further expressions of what was characterized above as an "extremely lightweight intelligence model". For example, image elements may become decayed, or page titles altered. *The Agent* prefers to decay the most transient elements of most pages – advertisements. This preference is implemented through a straightforward analysis of file URLs (e.g., many advertisements are stored on servers with "ads" as part of their name, and this is reflected in the file URL).

One type of scrapbook insertion is the placement of images and text drawn from Web pages concerned with *The Agent's* themes – impermanence, hypermedia, preservation, agency – onto other pages browsed by the user. These insertions contain links back to the source pages. Another type of insertion is of sentences from *The Agent*'s story (which will be discussed in greater detail later). These insertions are a way that *The Agent* "writes back" to the user in the context of the web browsing experience. They are *The Agent*'s annotations, *The Agent*'s mark, in the scrapbook of the user's experience. They are also a reminder to the user that they are working through *The Agent*, even when *The Agent*'s own window is not open.

The Impermanence Agent does not simply record, does not simply create memory pages, or simply create and alter these pages. As discussed above, *The Agent* is a storyteller. By pressing the "start" button on *The Agent*'s website, the user launches a small browser window that *The Agent* uses for storytelling. This window contains two equally-sized frames – one on the left and one on the right.

Each of these two frames is continually displaying either a scrolling text file or a fading image. These pieces of media are transmitted as traditional Web pages. *Text pages* scroll upward, and when a piece of text has finished scrolling a new page is requested to take its place. The scrolling is accomplished using a piece of JavaScript code that continually increments the vertical offset of the page until the end of the page is reached, and then requests another page. It is instructive regarding the current state of the Web, however, to note that JavaScript is unevenly implemented across browser and operating system platforms – so these text pages must also include a simple "time-out" for cases in which the offset measurement is not properly implemented in the browser version for the user's operating system, and in testing we limit our target audience to users who have a copy of *Netscape Navigator 4.* (The initial JavaScript code was written by Joshua Vura-Weis for the project, later one of the pervasive "scrolling news" implementations was adopted instead, and now a new script must be adopted if *The Agent* is to move to standards-compliant browsers such as *Mozilla*.)

The Agent's *image pages* each contain an "animated gif" – a type of image file that contains a number of still

A Well-Dressed Agent: The Aesthetics of Impermanence

by a.c. chapman

When Noah first spoke to me about *The Agent* I was working as a web and print designer in Seattle, and Noah knew me as both an artist and writer. We spoke via the wonder of the telephone – and tried to talk through some of the concepts of *The Agent*, and some of the visual aspects, in the image-free environment of interstate telecommunication. We discussed the playfulness and beauty of an agent of loss – I felt the tug, again, of that paradoxical experience of realizing the importance of someone or something only after it's gone. And I was hooked.

Over the years since I have lead our efforts – Noah, Brion, Duane, and I – to determine how *The Agent* will look and feel, how it will exist as an artifact. This work began in earnest when Noah and I met in New York in 1998 and, over coffee and cheesecake I sketched a number of possible representations – a window with three frames (one for images, one

images that are displayed in succession. In *The Agent*'s window, these successions of still images each show an animation of a single picture fading in and out. When the animation is completed a new page is requested. Because animated gifs are a technology commonly used for eye-catching advertising, and (perhaps for this reason) were introduced relatively early in the development of the Web, they are considered a basic Web technology and function on nearly all Web browsers for all operating systems.

The content of *The Agent* window, after it is launched by the user, begins with an introductory sequence and then transitions into storytelling. The introductory sequence is a visual and textual introduction to the workings of *The Agent.* During this sequence the left frame contains the textual introduction, which includes basic directions for using *The Agent* as well as elements of the story of Noah and Nana. The right frame initially contains images from *The Agent*'s story. A series of these is shown, each one fading in and out. This series repeats. But after the second or third repetition, the images begin to change (as does their fading behavior). The images have begun to customize, to become collaged with images drawn from the user's browsing (in a manner that will be

for fictive story, one for quotes and reactive comments), four frames, two frames. My primary concerns were that it be unobtrusive and clean. If *The Agent* window was something we wanted people to leave open on their desktop for long periods of time, it had to be something that could sit in the corner without being too distracting. In the end, we agreed that two frames and as small as possible would be the most elegant solution (as described in section "*The Agent* Window").

Of course, we were endeavoring to make a work of art, so while *The Agent* to be unobtrusive, it also had to be beautiful. And its beauty, its aesthetic engagement, was to be of a sort I would later hear the artist Camille Utterback refer to as *ambient interactivity*. The idea of ambient interactivity is simple: the work functions and is aesthetically pleasing on its own without the need for any direct interaction, however the User is rewarded with an experience which becomes richer in direct relation to their level of interaction with the work. Thus, the work had to be aesthetically engaging in and of itself.

The appearance would be a subdued color palette, text on the left, images on the right. My cemetery photos would be black and white, images of Noah's grandmother would be sepia-toned as to nominal-

described in greater detail below), and they are now fading in and out in a pixilated (rather than smooth) manner. By the time the introductory text has completed, each image will have been altered at least once.

During storytelling, both frames alternate between text and image. The left frame contains *The Agent*'s framing voice. This is a collage of observations, occasionally punctuated by *The Agent*'s reactions to the user's web browsing. The observations include texts from multiple authors (especially Dogen) and photos of memorial imagery from multiple cultures (many of which are drawn from cemetery photos of a.c.'s). The reactions include helping the user through the Kubler-Ross stages of grief as 404 "not found" errors are encountered during Web browsing (e.g., "It must just be a typo. <pagename> can't really be gone." with <pagename> replaced by the name of the page from the URL that produced the 404 error).

The right frame is *The Agent*'s story. This story's text begins as a collage of family history and fictions more recently created. The images are a collection of family photos and memorial imagery. The story, like the images that were in the right frame during the introductory sequence, repeats. The first time it is told it is likely to be in a ly set them apart. The text would automatically scroll, the images would slowly fade up from and out to white. All of which was in stark contrast to the splashy, flashy aesthetic of the web – an anti-web as *The Agent* was an anti-agent.

In addition, there were two versions of every photo: the first being the photo in its normal state, the second being a version where either context or primary content of the focus of the images was dropped out – replaced by an image file grabbed from the User's browsing. The process for this is described further in section "An Extremely Lightweight Intelligence Model" (Image Customization). The toned down appearance of the original images combined with the calculated image loss to emphasize the alterations and collaged additions that took place as *The Agent* ran. The monotone images of memorials and memories were pasted over and recontextualized by the bright colors and bold type of banner ads and roll-over buttons.

The result is a work that looks and reads beautifully on its own, but provides the User with a richer interaction if they become more involved with it. *The Agent* can either be a collected history of the User's browsing, or it can be an actively interactive collage with themes the User determines by browsing different genres of sites.

linear sequence, but over time there is an increasing probability that some circling back and looping will take place. As with the introductory images, each time either an image or text that is part of the story repeats there is a chance that it will be altered using information drawn from the user's scrapbook (on this process, more below). As discussed above, this alteration can be seen as an example of the "customization of content" that many commercial Web technology vendors and even academic writers see as the future of information – as well as a critique of it. The right-hand text also in some ways enframes the material in the left frame, calling its outside status into question. However, the materials in the left frame are untouched by the alterations that take place on the right.

An Extremely Lightweight Intelligence Model

When we say that *The Impermanence Agent* project employs "an extremely lightweight intelligence model," that's jargon. It is jargon that serves to indicate two things. The first is that *The Agent* reads the scraps collected from the user's browsing like a box of letters – searching for meaning, or at least pattern, in a pile of missives meant for someone else. The second is that we admit that all *The Agent* is capable of as it does this, all that is possible, is collage, juxtaposition, layering, pattern recognition, and reconstruction by form. There are those in the artificial intelligence community, and in some cases in the Web products community, who do this in a much more complex fashion than *The Agent* does, and we indicate this with the first two words of the phrase. Some examples of customization strategies, using an extremely lightweight intelligence model, follow.

a) Interest Enhancement. Part of the scrapbook kept for each user by *The Agent* is dedicated to a list of words. This list is drawn from the titles and "meta tags" of pages the user browses.[6] *The Agent* periodically analyzes this list of words for use in a particular one of its types of text customization. The first step in this type of analysis is the disposal of certain words from this list that are too common in English to be particularly meaningful in *The Agent*'s context. Next, each word is given a weight corresponding to how often it has been encountered while browsing. Third, each word that is most commonly used as a noun is given an additional weight.

Finally, each word is given an additional weight if its frequency of usage in English is low. These last two operations are performed using *WordNet*.

The six words with the highest weights are then placed on a stack. We find that the method outlined for choosing these words is, for its simplicity, surprisingly effective. The six words chosen are almost always quite evocative of the user's browsing experience.

This stack of words is then used in a type of text alteration that, as customization, we refer to as "interest enhancement". When *The Agent* chooses to perform interest enhancement on a passage of text in its story, an "interesting word" is popped off the stack mentioned above. The passage of story text to be customized is then opened, and *The Agent* finds a word in the passage that is the same part of speech as the interesting word (or, at least, that is most commonly used as the same part of speech). Every instance of this word in the passage is found. The interesting word is then put in all the places previously held by the word in the passage.

A subject of the user's browsing then becomes a subject of the story. By the logic of agent and Web-customization hype, this should make the story more interesting to the reader.

b) Sentence Customization. Another type of text customization performed on *The Agent*'s story operates at the level of the sentence section, rather than the word. This is probably the most complicated piece of *The Agent*'s customization code, and it is the direct descendant of techniques designed by Noah and Duane for *Raku Writing*. For *The Agent* project, the Raku code was reworked and integrated with Brion's code for generating, storing, and serving the content for each Agent user.

When sentence customization is chosen for a passage of story text, one of the sentences within the passage is selected, and a very lightweight analysis of the sentence structure is performed. This analysis uses either lists of grouped keywords or a combination of punctuation and part-of-speech analysis. An example list of keywords is: he, she, I, they. An example part of speech and punctuation sequence might be: a comma, preceded by a noun, and followed by an adjective. There are two modes for this analysis – more and less precise. More precise analysis uses only the keyword lists considered most accurate. Less precise uses an expanded number of lists as well as the punctuation and part of speech method. (Again, part of speech analysis is performed simply by consulting *WordNet* for the most common usage of a word.) Like much software presented on the

Web using an agent narrative, these methods were developed and refined with the goal of evocative results, rather than accuracy – though here what is evoked is perhaps more loosely structured, linguistically.

Once this analysis has been performed, a piece of text is chosen from the user's scrapbook. The sentences of this text are then analyzed to see if they contain any phrases with characteristics similar to those in the sentence from the story. For example, if the sentence in the story contains a word from list A, and later a word from list C, a sentence in the scrapbook that contains words from the same lists (though not necessarily the same words) in the same order would be considered a match.

The matching phrase from the scrapbook sentence is then inserted in the place of the phrase from the story sentence, thus customizing that story sentence. Given this process, over time the story text converges toward words and phrases in which the user has expressed interest (by browsing pages in which they are contained). Each time this process takes place, the scrapbook text from which the sentence phrase was drawn is marked for disposal (however, it may receive several uses before disposal takes place). This avoids endless repetition, and makes *The Agent*'s scrapbook one that is continually eroded by its own activity, in a manner quite similar to the user's browser cache.

c) Image Customization. In addition to customizing the text portions of its story, *The Agent* also customizes its images for each user. As with sentence customization, these customizations involve the collaging of story material with material drawn from the user's Web browsing.

Each time an image from *The Agent*'s story is shown to the user it becomes a candidate for image customization. The process begins with *The Agent* choosing an image from the user's scrapbook. First *The Agent* checks the image to see if it is large enough to be useful (images that are too few pixels across are unlikely to lead to interesting customizations). Next, if the image is an animated gif, one image is pulled from the sequence and the rest are disposed. The image is then cropped to remove any solid-color border, after which it is checked to ensure that the width and height are still greater than 10 pixels. All of these operations are performed using *ImageMagick*, a set of freely-available image-manipulation libraries.[7]

If this is the first time the story image has been customized, a special procedure is used. When *The Agent*'s images were prepared by a.c., two

versions of each image were made. The first of these is the version initially displayed by *The Agent*, before any alteration takes place. The second is the "mask" version of the image, used in the first customization. The mask version retains many of the essential elements of the image, but the remainder of the frame is a solid color intended to "drop out" in the customization process. The customization consists of replacing this solid color with a portion of an image from the user's browsing, so that the user's image becomes the background for details retained from the original image. The background image is created by scaling a wide image so that its height is the same as that of *The Agent* window and then taking a cross-section along its length, or scaling a tall image so that its width is the same as one of *The Agent* window's frames and taking a cross-section along its height.

After the first customization, a different process is used for each subsequent customization. In this process a scrapbook image is again selected and screened in the manner outlined above. The image is resized and a horizontal section is taken from wide images, while a vertical section is taken from tall images. Alpha channel information (an alpha channel is akin to the single-color area meant to drop out of the mask images described above) is retained from gifs that contain it, so these sections often appear irregularly shaped. (Alpha channel information is common on the Web, to a significant degree, because it decreases the download time for additional frames of animated gifs used in advertising.) These sections are then collaged over the current image.

The image produced by either of these processes is then made into an animated gif (fading in and fading out) for display by *The Agent*. The image drawn from the scrapbook, regardless of whether it was suitable for use in customization, is disposed. As with the story's text, the story's images in this way move toward containing the maximum number of elements from pages in which the user has expressed interest.

The users of *The Impermanence Agent* intend to be watched. The users of Netscape Smart Browsing, office proxy connections, ecommerce cookies, and other common Internet technologies may not realize they are under surveillance. As outlined above, *The Agent* uses these common

technologies – cookies, a proxy, and IP address tracking – for its watching.

We are currently at work on a new presentation of *The Agent*, which is predicated on our ability to watch the products of each user's interaction with *The Agent.* The first versions of this were recently shown in two places. One presentation was at the theatre of the Guggenheim Museum, New York, in a program titled "The Brave New Word," organized around an issue of *American Letters and Commentary*, as part of the Works and Process series. The other presentation was at the Z Media Lounge of the New Museum of Contemporary Art, New York, in a show organized by Harvestworks. In both of these cases, where the long-term individual experience for which *The Agent* was designed was impossible, we chose to provides a viewing opportunity for watching the products of Agent's watching of others. These could be seen as another of the now-common (yet important) explorations of the panoptic potential of the Internet. But rather than Bentham's (or Foucault's, or Echelon's[8]) panopticon which asserts control through knowledge of the identities of those surveyed, these presentations dis-

The Agent as Artifact

by a.c. chapman

The Impermanence Agent is actually two related, but separate, works of art.

The first work is *The Impermanence Agent* itself – that is, the browser window with two frames which opens on a personal computer desktop, monitoring and reacting to the User's browsing. This is a process-oriented artwork, created for extended experiences of ambient interactivity.

The second work is *The Agent* as Artifact. As discussed in section "Surveillance, Voyeurism, Storytelling," when we first began showing *The Agent* in galleries, we quickly realized that it was so absolutely an Internet art work that it was very difficult to exhibit in any other sort of environment. In a gallery or museum space where an artist can expect an audience member to spend an average of 40 seconds on their piece, a work such as *The Agent* – which we ask the User to spend approximately a week "using" – gamered befuddled, bored, and frustrated reactions. To compound the problem of the duration of audience engage-

played *anonymous* projections of content customized for other viewers.

Anecdotally, we found these presentations of *The Agent* to be significantly better received than those in traditional art spaces where visitors have been invited to experience *The Agent* directly. At venues such as the galleries at *SIGGRAPH 2000* and *New York Digital Salon 1999* we had made *The Agent* available for use, and watched visitors seem confused by the piece itself. In these settings, the explanation of the piece seemed much easier for the audience to engage than the piece itself. That is to say, we found we had not only succeeded in creating a piece which *could* be experienced peripherally, over time – but in creating a piece that *could not* be experienced in the traditional art manner of short, focused attention.

By contrast, even without explanation of *The Agent*, audiences seem to be drawn in by the results of *The Agent*'s impermanence – by its customized sentences and images – when they are presented "center screen". This could be because these texts and pictures are themselves engaging (when not placed in a small window at the edge of a screen), and it may also be that they are attractive because of the anonymous voyeurism they represent.[9] Whatever the reasons for their appeal, many venues (even those specifically exhibiting Web art) lacked Internet or unimpeded Internet connections. Naturally, lacking a connection to the Internet, it is difficult to show an intrinsically net-based art work.

The solution was to create an autonomous artifact of *The Agent*.

It's normal to document one's work, and to make that documentation as compelling as possible – to receive grants, be selected for juried shows, etc. However, in this case, the documentation actually became the work itself; that is, the documentation is the artifact exhibited. In this way the exhibited Agent is moving closer to conceptual art than to process, performance, or artifact-based artwork.

peal, the piece we are now creating uses this material, and is tentatively titled *The Impermanence Agent's Story*. Through it we are exploring the possibilities for presenting *The Agent* project as a piece for short, focused attention while still telling a story related to what it tells when it is experienced as originally designed. As we create it, we are also finding ways in which this new piece can be evocative in ways the first Agent project could not.

The Agent's Story piece will, like the individually-experienced Agent, be in a state of constant change, with past configurations impossible to recover. It will be presented as a space of image and text frames, organized in the most likely first sequence for the material of *The Agent*'s story when it is experienced by an individual. Each image frame will present a selection of recently-customized versions of a particular image drawn from the files of individual Agent users. Each text frame will present the sentences of a section of the story serially. Each sentence will be shown in several recently-customized versions before the next sentence is shown (in several customized versions, and so on). This approach will, we hope, be less invasive of privacy than showing a complete customized story from any reader, while still allowing some sense of time passage, impermanence, and voyeuristic appeal to be retained. Also, while a snapshot of this type of presentation could be distributed in a fixed medium, we hope that some of the sense of continual change experienced in the small, peripheral, individual version can be brought to a form more appropriate for the type of attention given by those experiencing new media artworks in galleries and other public settings.

At the same time, we may find that this project, when completed, seems orthogonal to the goals that drew us to work on *The Agent*. It provides a way to watch the story through the customizations created for others. But is this the "watching" that is important to *The Agent* being *The Agent*? Is the experience we worked to create that of being watched by *The Agent* (which the Story project captures in material drawn from many Agent watchings), or is it watching oneself through *The Agent*? We suspect it may be the latter. Again, we find we are not quite in the panoptic situation. In our experience, *The Agent* is more a mechanism for watching our own daily browsing, defamiliarized so that it may be seen again. We find ourselves more likely to alter our behavior based on what we wish to see reflected than on the potential for being viewed within *The Agent*'s imagined surveillance booth.[10] When *The Agent*'s story reflects too much of

the content that would be in a corporate-sponsored *Daily Me* – for example, endless IT product announcement "news" – we, sitting at our desks, feel encouraged to browse elsewhere, to shape what we are showing *The Agent* in a different direction.

Part III: Critical Technical Practices

In this section, work in Critical Technical Practices (CTP), of which *The Impermanence Agent* may be seen as a part, is considered. First, what CTP might be is discussed, along with the possibility that current discussion of CTP might be anachronistic – anticipated by practices in 20th Century art or Participatory Design. Next, the arrived-at definition of CTP is used to consider work not currently discussed as CTP. Finally, *The Agent* project's three-legged practice is considered as a type of CTP, and the prospects for this tripod are briefly explored.

Critical Technical Practice?

Now we return to a question first raised toward the beginning of this essay: Is *The Impermanence Agent* project an example of critical technical practice (CTP)? In order to answer this, and determine the potential import of the answer, it may be useful to consider the term at greater length. As noted earlier, the CTP formulation comes from the work of Phil Agre. Agre rose to prominence from an Artificial Intelligence background – he is trained as a computer scientist, and worked at the MIT AI Lab in the 1980s. In his writings he positions CTP as a way of rethinking technical domains, and particularly of working through technical impasses that result from naturalized, mistaken assumptions about human nature. (Agre 1997b) Given that AI and Human-Computer Interaction are often the only areas of computer science that acknowledge their work to be deeply connected to (ideas of) human nature, CTP might seem to have restricted potential for application. However, recently the term has been given an increasingly broad range of use. For example, Sengers has recently discussed, under the umbrella of CTP, using insights from disciplines such as cultural studies and history to inform the development of "information appliances" – beginning with the types of living one might wish to sup-

port, determining what types of information might be useful for this, and then designing the appropriate artifact to communicate this information. Sengers, in describing this work, used the example of exploring designs for helping those concerned with conservation monitor their electricity usage. (Sengers 2001) This contrasts markedly with the more common engineering approach of building supposedly value-neutral "enabling technologies". It points to one primary way in which CTP differs from other models of creating technological projects: a wider view of what should be considered part of the intellectual ground for such work. What it leaves out, however, is that for Agre the *critique* which provides the "C" in CTP is not satisfied by cultural critique in general. Rather, the C in CTP requires a critical reflection on the technical practices themselves, which is in turn integrated with this technical work. As Agre writes, "A critical technical practice rethinks its own premises, revalues its own methods, and reconsiders its own concepts as a routine part of its daily work". (Agre 1997b, 24)

However, it is precisely this second facet of CTP that lead Penny, at our DAC panel, to question whether there was much new in the CTP perspective:

> When I first read Philip Agre's *Computation and Human Experience* in 1997, I felt it was the most important critique of AI since Hubert Dreyfus's (an opinion I still hold). Yet as an artist, I felt a hint of anachronism, specifically around the issue of Critical Technical Practices. Because while Agre argued that this approach was groundbreaking in Computer Science, it was in fact the ground upon which I had learnt to make art, in the late 1970s. That is – one of the very few continuities in 20th century visual art is a constant critical if not revolutionary relation with its past. Avantgardism is premised on the subversion, inversion or total trashing of accepted conventions of the immediate past, over and over again.

Penny's point is well taken, but it seems to run us into that old chestnut of "the two cultures". Are art and science two fundamentally incommensurable constructions of the world, or are they in fact close enough that the *modus operandi* of 20th century artists could make Agre's CTP anachronistic?

Before attempting to answer this, let us consider a second comparison which has led some to consider CTP anachronistic: Participatory Design (PD). Now that CTP is used by figures such as Sengers to describe a larg-

er universe of computer science projects based on question formulation that is informed by humanist disciplines, it begins to sound remarkably like Scandinavian formulations of PD (Ehn 1988). This is particularly true when the questions posed by a critical technical practice are of the types of living and work we wish to support, or of how systems can be made more successful by considering the user in new terms. These are the central concerns of PD (which has been developing as a practice for decades) and in general of long-term Scandinavian movements for technological democracy. If we can be forgiven for a lengthy quotation, these passages by Langdon Winner seem particularly appropriate:

> In the context of theoretical debates about determinism and social shaping of technology, the Scandinavian projects offer a number of very promising departures. One distinguishing feature in some Scandinavian approaches, for example, is to take seriously the design of technological devices for the qualities of social life they sustain and the everyday political habits they nurture.... Another point of departure in some of the Scandinavian approaches is that they affirm as both social policy and research method that impending technological developments should not be regarded as something external to the lives of those who will eventually be affected. There is an affirmation that having a say in the design and application of new instruments is a basic right that derives from citizenship, not just property ownership.... Yet another point of departure is an affirmation that ordinary people are capable of being directly involved in shaping new technologies. They already know a great deal that is useful and, beyond that, can educate themselves further in areas of technical knowledge usually supposed to be the sheltered domains of experts.... Another promising feature in much of this work is the recognition that technological development can fruitfully draw upon a much richer array of human fundamentals than the mechanistic technical and economic models that have prevailed until now. One can draw upon models in philosophy and anthropology, and sociology to ensure that systems are not spawned and nurtured by a one-dimensional rationality. This means that beyond the critique of instrumental rationality lies a body of understanding and fruitful practice that one can begin to teach the next generation of technical professionals and ordinary citizens... The overall promise is that [we] will see the rise of an orientation toward planning and design that can produce qualitatively superior

> systems, ones that are fully respectable in an economic and technical sense, but which incorporate a much wider spectrum of democratically relevant features in their shape and performance. Hence, democracy can be manifest in the process, in the evolving creation of technical knowledge and practice. Perhaps it will even be tangibly apparent within the lasting forms of the technological devices and systems in widespread use. (Winner 1994)

The first and last portions of this quotation outline areas of Scandinavian work which are clearly mirrored in projects such as Sengers' information appliances: informed by social goals and humanist disciplines, and with an aim of building better systems. However, the middle section offers an interesting departure from what we have discussed so far of CTP. It emphasizes the direct participation of users in the definition of technologies – not simply through "user testing" or focus groups, but through co-creation that may have the larger public interest as a goal.

Does this make critical technical practices simply a less-radical version of Scandinavian work – and late to the party, at that? We don't think so. Neither do we think that the 20th century art practices described by Penny make CTP anachronistic.

This is not to say that these pre-existing practices did not anticipate important elements of CTP. Rather, it is to argue that anachronism is not the most fruitful formulation of their relationship with CTP.

Toolness and Artness

Let us return to the question of the two cultures. We've had some interesting conversations in this vein recently. One friend said that she'd heard Rich Gold say that artists and scientists are alike because they're both elites that believe themselves oppressed. Another friend argued that art and science might be more alike than either is like engineering. He held this to be the case because both art and science make a request of their practitioners – that if they begin a journey in one direction, but find something more important along the way, they change course to follow their discovery. Engineers, on the other hand, are more often charged to continue "on course" toward the pre-established goal of their project. Clearly this distinction doesn't entirely hold up (do scientists become engineers

when they have National Science Foundation funding that includes reporting and project-completion requirements?) but it does usefully draw our attention to the *products* of our work.

Scandinavian-style participatory design is generally viewed as producing tools (and specifications for tools, and new work structures around the introduction of tools). These tools are, specifically, computer tools that continue the tradition of skilled craft tools. The well-known UTOPIA project, for example, undertook this task for graphic workers employed at newspapers. Such projects drive forward technology and its applications, but from a specific, culturally-engaged position.

20th century art, on the other hand, is rarely seen as primarily engaged in the creation of tools. Instead, it is engaged in the creation of art, which may range from tools to concepts to performances to wall hangings. It is probably safe to say that artists are more often involved in the use of tools, and the critique of naturalized tools, than the creation of tools for others. Similarly, artists in general are involved in culturally-engaged creation, but only a small subset create specifically with computer technology.

Given how divergent are the products of participatory design and art, it's not surprising that they haven't often been compared. To our knowledge, no one has spoken of participatory design as anachronistic given that democracy of production was a prior theme of 20th century art.

So where do CTP's technical practices lie that both PD and art have claims to precedence? Let us consider a few more examples of current work that is called CTP before we answer this question.

The Industrial Graveyard

The Industrial Graveyard is another product of Phoebe Sengers's critical technical practice (Sengers 1998). It is a virtual environment in which a small lamp tries to survive while being overseen in a hostile manner by a larger lamp (figure 4), which in turn is overseen by the user. The large lamp is controlled by a reactive agent architecture of a type commonly used in the AI community. It is made up of a set of relatively independent behaviors, and selects its current behavior through a stimulus-response model. The small lamp is controlled by The Expressivator, a system de-

Figure 4

signed by Sengers that builds upon, and departs from, the standard reactive agent architecture.

While reactive agents have allowed AI to spring back from a low point (after the hype of expert systems turned sour), Sengers points out that even those who currently create reactive agents are often unsatisfied by them.[11] Somehow, they seem to be mechanistic, to lack the "juice" of life (a word used to describe this lack by Rodney Brooks, famous practitioner of reactive robotics at MIT). Sengers believes that this missing juice can be found in narrative psychology. That is to say, she believes the juice missing from reactive agents – the lack that makes them seem mechanistic in comparison to living things – is a lack of narrative. She argues that we understand living things in narrative terms, and within a background of narrative, but that agents are not currently designed to be understandable in narrative terms. In fact, the very atomization of their behaviors into independent components seems opposed to narrative. As she points out, in humans we call this sense that we are made up of largely independent, alternately-controlling machines by a particular name: schizophrenia. Current reactive agents express their schizophrenia by making jumps from one behavior to another in ways and at intervals that seem mechanistic to viewers.

This, then, is the perfect ground for a critical technical practice. An impasse has been encountered – reactive architectures create schizophrenic-seeming agents. The impasse is arguably the result of a misformulation of how we understand human beings – the traditional domain of the human

sciences. Sengers takes this opportunity to create The Expressivator, which is an architecture for creating agent systems that are explicitly designed to be understood in narrative terms. For this architecture she has reformulated the basic building blocks of agent architectures. Rather than manage her agents as a collection of behaviors, she manages them in terms of signifiers and transitions (transitions connect signifiers into narrative structures). Sengers calls her specific critical technical practice, of which The Expressivator is an example, "socially situated AI".

Her writing on the subject traces the evolution of this work out of the critiques of AI's practices – beginning with the familiar critique of classic AI by alternative AI practitioners that an agent must be engaged in a *context* for its intelligence to be meaningful, and then proceeding to critique alternative AI's assumption that this context is purely *physical,* and arguing that it must also be *socio-cultural*. Sengers traces the similarities between a reactive, physically-bounded view of agents and the view taken of workers as objects in Tayloristic industrial capitalism (in an interesting counterpoint to Pelle Ehn's writing on the subject – see endnote 1). Socially situated AI frames its questions in terms of communicating subjects – and in projects such as *The Industrial Graveyard* it emerges as an alternative that operates culturally and technologically, while informed by practices of the narrative arts.

Terminal Time

Terminal Time is a CTP project involving Michael Mateas and two collaborators, Steffi Domike and Paul Vanouse (Mateas & Domike & Vanouse 1999; see also Mateas 1999). *Terminal Time* is presented in "performances" made up of screenings and discussion. At each screening the group presents a TV-style documentary outlining the history of the last millennium (the years 1001 through 2000 of the Common Era). However, the documentary is different each time – produced anew in response to audience choices during ideologically-loaded surveys (Figure 5). Using classical AI techniques[12], an "ideological model" of the audience is created, and then exaggerated. A documentary is then constructed which makes a strident argument that the events of the past millennium support this exaggerated model of audience ideology. *Terminal Time* is sometimes presented straightfaced – as though nothing could be better:

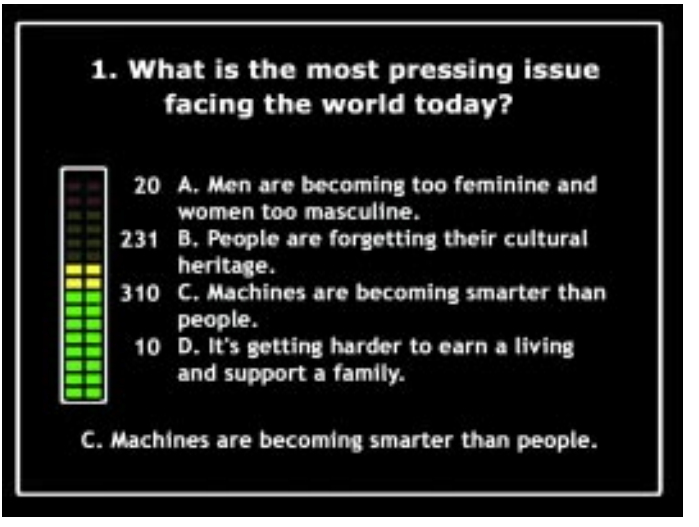

Figure 5. A sample survey question from *Terminal Time*.

> *Terminal Time* is a history "engine": a machine which combines historical events, ideological rhetoric, familiar forms of TV documentary, consumer polls and artificial intelligence algorithms to create hybrid cinematic experiences for mass audiences that are different every single time. History as it was meant to be told![13]

It's obvious that, given its goal of creating TV-style documentaries, *Terminal Time* is at least as concerned with narrative construction as is *The Industrial Graveyard*. Both of these narratives are also explicitly *performed* for the audience/user. But what may not be as clear to those first encountering these projects is what Mateas states in his writings on "expressive AI" – his name for his practice:

> In expressive AI the focus turns to authorship. The AI system becomes an artifact built by authors in order to communicate a constellation of ideas and experiences to an audience... The concern is not with building something that is intelligent independent of any observer and their cultural context. Rather, the concern is with building an artifact that seems intelligent, that participates in a specific cultural context in a manner that is perceived as intelligent.

> Expressive AI views a system as a performance. Within a performative space the system expresses the author's ideas. The system is both a messenger for and a message from the author. (Mateas 2000)

By formulating "expressive AI", Mateas might be seen as making an end-run around the turf war between GOFAI (Good Old Fashioned AI) and alternative AI. But this is more than a clever linguistic/political maneuver. Again, it is a re-seeing of the field of inquiry. While AI had primarily seen itself as creating tools for intelligence, Mateas and his collaborators created a project that worked with AI elements but saw itself as creating tools for expression. And, like Sengers's *Industrial Graveyard*, it achieved a technical success that had eluded those who saw the field in the previous terms.[14] Sengers's success was in creating agents that did not seem schizophrenic to observers. Mateas's was in creating a dynamic narrative system with both more flexibility and coherence than those that had relied on techniques that were not formulated in terms of expression. For example, *Terminal Time's* narrative construction process didn't use off-the-shelf language generation techniques, because the deep natural language generators previously created did not provide the ability to represent the complex pragmatic (e.g., syntax, semantics, pragmatics) information necessary to describe *Terminal Time's* particular writing style. On the other hand, simply choosing between pre-written texts, however fragmentary, would not have provided the flexibility necessary for *Terminal Time's* undertaking. Instead, *Terminal Time* took the first steps toward a new class of interactive text systems – providing support for authors writing many carefully-crafted pieces which the generator works to recombine (using a "heavyweight" intelligence model) while preserving style.

But what must also be noted about Mateas's project is that another type of critique is at work within it. This critique is similar to that at work in Sengers's work on information appliances (described earlier), and to the thematics of the *Industrial Graveyard*, and to the work of Warren Sack described below.

Conversation Map

Warren Sack is another of the emerging group of "second wave" critical technical practitioners. His *Conversation Map* is not an agent-centered

system. This differentiates it from the work of Agre, Sengers, and Mateas – our prior CTP examples. Like the work of Sengers and Mateas it has a focus on communication – but in this case the communication is not between the system and the user, but between very large numbers of users.

The Conversation Map brings together computational and conceptual tools developed in sociology, linguistics, and information visualization to create a tool for interacting with very large-scale conversations (VLSCs). Examples of these VLSCs are messages sent to Internet newsgroups and mailing lists. Sack's tool works to support navigation of these conversations, while providing an alternative visualization of "the public," and while embodying what he calls a "perhaps ridiculous" optimism (figure 6).

The panes of the interface shown here display a number of visualizations of a VLSC from the Usenet newsgroup soc.culture.albanian, a group devoted to the discussion of Albanian culture in general, but at this period in time (16 April 1999 – 4 May 1999) especially the war in Kosovo. The upper-right pane displays which participants are responding to others. Those who respond to each other more often are connected by short lines. The fact that much of the activity is mapped in the middle of this pane shows a very cohesive conversation. The next two frames (top middle and

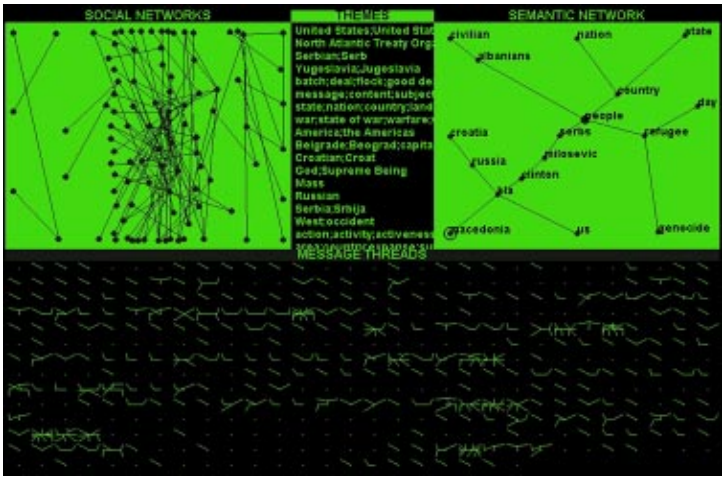

Figure 6. A sample image from *Conversation Map*.

right) use the *WordNet* tool that *The Agent* uses, but in a very different manner. The top middle frame attempts to draw out themes from the conversation, while the top right frame – perhaps the most optimistic portion of the project – attempts to create a mapping of terms that are used similarly. This "semantic network" is, in a sense, attempting to produce metaphors through analysis (about which more soon). Finally, the lower portion of the interface is given over to visualizing the discussion threads. Each thread begins with an initial message. A response to that message, and any followups along the line of a particular response, creates an "arm". Those threads with many responses begin to network outward like a spiderweb.

Now let us return to the semantic network pane. Sack writes of that portion of the example pictured here:

> This conversation illustrates a non-dialectical exchange in which, potentially, no common ground might ever be accomplished. In the words of the philosopher Jean-Francois Lyotard, this conversation may very well illustrate a differend, a difference so vast between participants that it can never be bridged (Lyotard, 1988). But, the machinery of the *Conversation Map* (i.e., those functions which automatically compile a rough-draft thesaurus for a set of messages) works in a strictly mechanical manner that sums and then averages together the language of the group. *The Conversation Map* is doggedly dialectical. Because of the way it is built, it *cannot not* find a common ground. Consequently, even for an argument so vicious or incoherent that a skilled, human negotiator might find no place to start building common ground, the *Conversation Map* will diagram – through its mechanical operations – a potential synthesis. (Sack, forthcoming)

We read the *Conversation Map* to be ridiculous and sublime. It is a rethinking of the conversation interface that attempts a very different type of presence than the advertisements of Web interfaces to VLSCs, or of the tight platform integration of the Microsoft offering. It attempts a presence that is well-established in the tradition of mapping, from medieval monstrosities at the map's corners to the still-prevalent maps that exaggerate the sizes of the Northern Hemisphere countries relative their Southern Hemisphere colonies. This is the tradition of mapping as ideology. And here, Sack has chosen to map the "flame wars" that often characterize VLSCs as spaces of potential critical reflection and hope. In this way he

invites a new participation in these conversations, and a new view of the public that is engaged in them.

Sack's project, of course, involves a critical re-seeing of technical domains, just as our previous examples do. But it also places before our eyes a differentiation from Agre's formulation that we haven't had to face directly since our discussion of Sengers's information appliances. Agre's work clearly proceeded from a technical motivation, and was caught in a technical impasse, which was worked through using the insights of CTP. But Sack's work, like Sengers's on information appliances, did not begin with an impasse recognized by the technical community. Rather, it began with one explicitly not recognized by the technical community – one to which the technical means of proceeding were blind, but which could be identified via social and critical engagement, and then (like Agre's technical problem) only effectively grappled with using the insights of CTP.

Another way of putting this is that – in "second wave" CTP – using critical tools to re-see technical problems and solving technical problems via these insights are still primary activities. But critique may be operating in another important fashion. The CTP may have been entered into because of motivation from, and the CTP project may itself be infused with, critical engagement with the culture. Or, instead, as in Agre's formulation, the projects may have been initiated in an attempt to address a technical impasse. But whether the philosophical/critical issues come into play upon running into an impasse in the technology, or the technical issues come into play upon running into an interesting critical/philosophical problem, if the motivating problem is addressed through a creation of technological artifacts and a cultural engagement that are pursued *as one activity* – informing each other, and calling any set of assumptions (however recent, however "technical") into question, then the practice can be called CTP.[15]

Converging the 'Two Cultures'

The presence of these levels of critique – combined with using critical tools to re-see technical problems and solving technical problems via these insights – not only ties the above projects together. It also ties second-wave CTP to Scandinavian PD and certain technologically-engaged art.

Given that we now have a better view of this connection, we see we may be well served to conceptualize CTP's relationship with PD, and oth-

er practices, in terms other than anachronism. One amusing alternative might be "anticipatory plagiarism" – a phrase used by the Oulipo. The Oulipo explore methods of writing using mathematically-inspired rules for constraint and generation. Obviously, others had explored portions of this territory long before the Oulipo began (in 1960), but none had articulated this exploration specifically as their area of focus. The Oulipo has used "anticipatory plagiarism" – tongue firmly in cheek – to refer to these prior explorations.

If we decide to take the emerging area of second-wave critical technical practices as "using critical tools to re-see technical problems and solving technical problems via these insights, a combination which may be motivated either by reaching a technical impasse or by a critical/social engagement (or both)" we can find many examples of anticipatory plagiarism, while still finding the articulation and pursuit of CTP as such to be a useful and new activity. Beyond the examples already stated we could take no less a luminary than Norbert Wiener. Wiener, who coined the term "cybernetics" from which we get cyber-prefixed phrases and concepts such as "cyberspace," became dedicated after the Second World War to defining a new type of scientist. This scientist would be actively concerned with the results of science, with technology and the way it shapes and is shaped by culture. Wiener became a cultural critic, writing and lecturing widely, and arguing for the sort of work he also attempted to move forward: cybernetic technologies that would help support a more just society.

In addition to historical figures, we may also recruit to CTP the work of those working now who do not yet use the label themselves. For example, Natalie Jeremijenko, who has pursued technological development with critical cultural/technological engagement in areas ranging from biotechnology to wide-area video networks. Jeremijenko's *OneTree* project, for which she created a series of genetically-identical walnut trees, has been one of her most public undertakings. These trees are being planted around the San Francisco Bay Area, where they will develop differently depending on conditions: environmental, socioeconomic, and so on. The result is both the development of a new type of "instrument" and a critique of the concept behind genetically-determined development. The result will be a visualization both of the readings of the instrument (a visualization readable even by those without scientific training) and of the fact that genetic identity is only a part of life's picture[16] (Jeremijenko 2001).

A final example is the work of Andruid Kerne, whose *CollageMachine* was an influence on our thinking about *The Agent*. *CollageMachine*, rather than being integrated with the user's daily experience of a Web browser, creates an alternative browser, and with it imagines a different relationship to the Web. *CollageMachine* creates a continual, user-guided collage of elements drawn from websites. It may begin with search terms or the URLs of particular sites, and from there it forages. What it finds is then removed from the presentation mode of Web "pages" and becomes part of a temporal, rythmic, interactive collage presentation. (Kerne has an extensive background in composition, and has worked both in traditional African musical contexts and with Cage-inspired avant-garde composers.) The user may at any time rearrange elements of the collage, express a desire for more or less content of a particular type now on the screen, or open the original page from which an element comes. *CollageMachine* exists as both a tool and an artwork, and is refined and imagined in both terms – Kerne has undertaken traditional user testing, as well as engaged in artistic dialogue and friendly critique, in order to create an agent of Web recombination that presents a new mode of relation to the Web's media (Kerne 2000).

The Third Foot and Dynamic Text

Noah: Gordon was about to become a father when he killed himself. I was visiting. He had an idea of the paper trail continuing after. He left a note, he left a diary. I feel, toward the end, he must have written in his diary to be read – and in it he wrote of reading a diary he was not meant to.

It is because we live in stories, in language, in the paper trail, in impermanence, that we created *The Agent*. As we wrote above, this is the "third foot" of our practice. And what is clear, from examining the work of Sengers, Mateas, and Sack, is that the emerging definition of second-wave CTP coming from work such as theirs has room for narrative, for performance, for play, for the exploration of human experience. And when we take the elements of these three "feet" and translate them into capabilities, we find ourselves with a tripod structure quite similar to that outlined by Penny. By the logic of anticipatory plagiarism, let us consider this work a form of CTP as well. And let us place *The Agent* within this CTP field.

The Agent's CTP can be seen as of the sort in which "critique" operates in two manners. There is the culturally-engaged critique of agent servitude and dreams of permanence, proceeding from sources ranging from the humanities to Zen, and which is described in the first part of this essay. There is also the critique that was employed in re-seeing the technical domain of storytelling systems, which was motivated by strands of storytelling practice – of reading and writing practice – different from those encapsulated in "how to" books advising writers on the structures of the ongoing traditions of the neo-Victorian novel and neo-Aristotelian drama.[17]

In the practice from which *The Agent* project proceeds, previous interactive story systems can be divided into two sorts: fixed text systems and dynamic text systems. Most interactive story systems, from videogames to hypertexts, are fixed text systems (if they use text to a significant degree). A few systems designers have explored the potentially interesting area of dynamic text systems, which create a new text for each user. Unfortunately, dynamic text system designers have generally viewed the task as: creating a program that "understands story," which is used for dynamic generation of the story (perhaps based on (user) input), followed by some sort of natural-language generation (NLG) to produce the actual fiction (e.g., Lang 1994). This hasn't worked for creating what we would consider compelling fiction, and we would contend that no such system is likely to succeed at this task.

A reason for this is that natural-language generation technology is simply not aimed toward creating unexpected sentences, and these are a necessity for nearly any compelling storytelling. This also brings us toward a deeper reason, from the point of view of *The Agent*'s practice: many of the most interesting writers and artists of the 20th century haven't viewed their work as beginning with story and working up to text. They have begun with language, exploring permutations of language, evocative juxtapositions of language, and ways that language play can get at the world in ways that story play cannot. It is this, in fact, that has lead us to formulate the problem (as in the paragraph above) with divisions based on text, on language, rather than story (as has usually been done in the technical practice and in "how to" writing books alike). *The Impermanence Agent* begins with language, and locates its dynamism at the level of language. This is not only more appropriate to meeting our artistic goals, it is also a much more computationally-tractable problem than locating the dynamism at the level of story and hoping that NLG can take care of the rest.[18]

Yet the formulations above require some clarification. First, it is important to note that the goal of *The Agent*'s system is not to exactly reproduce the pleasures of recent language-centered fiction practice, even though this work made possible the re-seeing of the technical domain. Rather, the goal is to explore an in-between space toward which some of this work points. For example, while Burroughs produced cutups for others to read, he also presented his processes and encouraged others to produce their own cutups. A goal of *The Agent* and our ongoing work is to explore what can be specific to dynamic text work that explores the space in-between Burroughs's cutups and his instructions: something that includes pre-written fiction, pre-written language, but also gives a significant role in language formulation to what happens in the user's time and place. Which leads us to another clarification. The distinction between dynamic and fixed text systems should not be read simply as one of granularity – of how large the pieces are that the system recombines. A focus of *The Agent* project (and, again, of ongoing work) is on the introduction of *outside* text to the system at the time of user interaction. With *The Agent* this takes place via Web browsing. But it could also take place via voice recognition, or tapping into a newsfeed, or the monitoring of email, or even the direct typing of text by the user at a *Zork*-style prompt.[19]

The structures of the story, of how it is told, such as the details of how text is collected and what is done with it, are also *part of The Agent*'s story – connected to its themes of impermanence and documentation, agency and attention. This is another departure from the technical field of story systems, which tend to operate with a conception of story as that delivered by the system, rather than of the system and its structure being story elements. Our conception, of course, brings with it the conclusion that each dynamic text system of this sort will likely be one-off. The form of the text manipulation will be specific to the goals of each piece. What happens with language will be as much what the fiction is about as anything else is – so just as new text will have to be written, so will a new manipulation system. Certain tools, of course, may be reused – just as some were reused from *Raku Writing* in *The Impermanence Agent.* But the tools reused between these two projects were more akin to *WordNet* than to an overarching story system, and we expect that this will continue to be the case for projects of this nature.

Why tell this kind of story? This kind question has been asked and answered many times. Rather than survey the generic responses, let us

answer from the specifics of this project, from the specifics of personal experience. We do this because to us Nana's story and Gordon's story and stories like theirs and ours are not epistolary novels – they cannot be got at by only the words left on the pages left behind.

Postscript

In 1995 and -96 Adrianne Wortzel and Noah offered a course in "Writing and New Media" at NYU's Graduate Film and Television program. Wortzel taught that the importance of new media work could lie not in imagining that it represents the true structure of thought, or in attempting to give greater freedom to the reader, or in whizbang effects – but in that it offers a model of continual process, a sidestep from the weight of the masterpiece, an escape from the illusion that we write for posterity. We heard her saying: writing in a different mode, a different engagement with narrative and the paper trail, impermanence. It is perhaps, finally, in these terms that we see this work.

This is a revised and expanded version of an essay first published in "Intelligent Stages: Digital art and Performance," PAJ 70, Vol. XXIV: 1 (Winter 2002), from MIT Press.

NOTES

1. Pelle Ehn, one of the founding figures of the Participatory Design movement, has written forcefully on the issue of 'intelligent tools' and human oppression in his *Work-Oriented Design of Computer Artifacts* (pages 401–402). He underscores that the newfound control promised by these dicourses is actually in conflict with the idea of tool intelligence itself:

   > Neither the problem of 'tool intelligence' nor the use of 'intelligent tools' are new phenomena, slavery being the most inhuman historic example of human use of other humans as 'intelligent tools'. Slaves seen and used as 'intelligent tools' by their oppressors could to some extent be controlled by violence and ideology. However, this widely practiced non-human use of human beings as 'intelligent tools' later continued in the emerging capitalistic society. Here it took the (by then socially accepted) form of divided wage labor. However, this human use of human beings represents a fundamental control problem for the oppressing 'owner' of the 'intelligent tools': He might get the commanded job done, better than he could ever have done himself, but he lacks control or understanding of *how* it is done. The human 'intelligent tools' are by no means transparent. Taylorism and 'scientific management' were an attempt to regain this control and transparency – on behalf of management.

2. This notion of having "two feet" should not be viewed as a full definition of Critical Technical Practices, but rather as an evocation of one of its central characteristics. More exploration of how CTP might be defined appears in part III of this essay.

3. This project was a collaboration also involving Chris Spain and Nathan Wardrip-Fruin, as well as a number of guest writers. The project took its name from the raku pottery technique. While many definitions of raku could be given, the most

essential element is probably that found in the last sentence from this paragraph of the *Encyclopædia Britannica* online's definition:

> The most significant fact about raku pottery is the technique: instead of warming and maturing the pottery in a cold kiln, glazed ware is placed in a hot kiln for only about one hour, then removed and forced to cool rapidly at air temperature. The process subjects the pottery to extreme stress and creates unique effects throughout the glaze and, sometimes, in the pottery itself. Reduction firing, in which the hot pottery is placed in a flammable substance to deprive the surface of oxygen, increases the chance aspects and dramatic surface variation of the glaze. Chance and process are the key elements in the raku aesthetic.

www.britannica.com/bcom/eb/article/8/0,5716,64128+1+62544,00.html

4. Not all the *News in the Future* projects share this tone. Others suggest more meaningful community building and filtering possibilities. These, however, have attracted fewer attempts at public implementation. One may presume this is because they do not contain a clear path to corporate profit in their conception.
5. Web browsers make requests to Web servers (where "pages" are placed for network access) using HyperText Transfer Protocol (HTTP). Each page (as well as each image, sound file, or other element not included in the page's main text) has a unique address that is used when making the HTTP request for it. This address is a URL (Uniform Resource Locator).
6. Meta tags are a type of text found in Web pages meant to give information about the page and that is not displayed when the page is viewed in a browser. Common types of meta tag information include copyright information, keywords that describe the page, and a short summary of the page's contents.
7. See http://www.wizards.dupont.com/cristy/ImageMagick.html/
8. "Echelon" is the popular name used for a massive communications-surveillance system developed by the U.S. National Security Agency and partners around the world. Its existence is still formally denied, but with a formality somewhat

akin to how the NSA's existence was itself denied for a period after its presence was well known. As the American Civil Liberties Union's *Echelon Watch* reports, this system monitors "as many as 3 billion communications everyday, including phone calls, e-mail messages, Internet downloads, satellite transmissions" in the United States, the United Kingdom, Canada, Australia and New Zealand. The collected data is examined for unusual flows, pre-defined keywords, or emergent patterns. The information found is used not only against targets chosen in advance (which apparently range from suspected terrorists to charitable organizations to foreign companies involved in bidding against domestic ones), but also to *identify new suspects*. Previously, of course, we were lead to believe that the government would only violate the privacy of those already suspected of crimes. If the evidence for this suspicion wasn't good enough, the warrant for a search or authorization for a wire tap might be denied. Now we find ourselves entering a total surveillance situation of truly dystopic proportions.

9. Anonymous voyeurism is itself a common Internet occurrence, and one apparently seen as having significant commercial potential. To return to our previous example, Ask Jeeves, within its agent discourse, invites users to – instead of asking their own question – look into the question-asking interactions that "Jeeves" is having with other users. And webcams, of course, have largely proven popular only when they allow – or seem to allow – illicit forms of watching.
10. Self-surveillance is also a major theme of *Office Plant #1*, a project created by Michael Mateas and collaborators. This piece is a robotic "plant" situated in the user's office which responds visibly to the user's email activities. Mateas is a major figure in what we will characterize, in the next section, as second-wave critical technical practices.
11. This cycle of AI hype and collapse has, of course, taken place several times – and is far from unique to AI.
12. Including formal knowledge representation via higher-order predicate calculus, constrained search, a logic engine providing inferential capabilities, and a task-net-like planner.
13. See http://www.terminaltime.com/project.htm

14. In fact, Sengers's socially-situated AI and Mateas's expressive AI are in many ways moving along similar lines. Perhaps the most important difference is Mateas's concept of artistic "affordances" for authorial and interpretive purposes, as described in Mateas 2000.

15. In this discussion we are indebted to personal communication from Sengers, Mateas, and Sack in response to the previous version of this essay. Sengers has also playfully suggested that what we are calling "second wave critical technical practices" might instead be called "technical critical practice" (given that its definition departs from Agre's). We have also drawn on her discussion of the distinction between the two trains of thought.

16. However, it is possible to imagine a CTP "purist" (if such were to come into existence) arguing that this is not the most felicitous example from Jeremijenko's work. The reason being that walnut trees, even genetically-identical ones, are arguably not technical artifacts, even if their use is both scientifically and culturally innovative.

17. We are indebted to Phil Agre for his personal communication in response to the earlier version of this essay, which has pushed us to define more clearly the relation of *The Agent* to outside concepts of story and previous story systems, as well as pointed out the need to deal more explicitly with the concept of technical impasse in the sections above.

18. It is also possible to scale up our approach to a relatively large number of Web-based users without depending on special client-side processing, while most other techniques are too processing-intensive. However, in Web system design, a lack of client-side processing is certainly a convenience for which the price is sometimes too dear.

19. In much of this thinking and working on dynamic text we are walking in the footsteps of poets such as John Cayley. In fact, in the *Cybertext Yearbook for 2000*, Cayley mentions a dynamic text implementation of his classic *Speaking Clock* that draws on a newsfeed (p. 99). Also, returning to the question of granularity, the range of Cayley's morphing text projects presents an interesting limit case for the fixed/dynamic text distinction – as, of course, does work such as *Terminal Time.*

REFERENCES

Agre, Phil E. (1997a) “Toward a critical technical practice: Lessons learned in trying to reform AI”, in Bowker, Geoffrey C. & Star, Susan Leigh & Turner, William & Gasser, Les (eds.), *Social Science, Technical Systems and Cooperative Work: Beyond the Great Divide*. Hillsdale, NJ: Lawrence Erlbaum.

Agre, Philip E. (1997b) *Computation and Human Experience*. Cambridge, UK: Cambridge University Press.

Bender, Walter (2001) “Walter’s Page”, http://nif.media.mit.edu/people/walter/#personalization/ [still on page as of July 2001]

Bernstein, Mark (1997) ”HypertextNow: Web Shrine“, http://www.eastgate.com/HypertextNow/archives/Shrines.html/

Ehn, Pelle (1988) *Work-Oriented Design of Computer Artifacts*. Stockholm: Arbetslivscentrum. (Distributed by Lawrence Erlbaum, Hillsdale, NJ.)

Ehrenreich, Barbara (2000) “The Vision-Impared Rich”, *The Progressive*, http://www.progressive.org/ehrn0500.htm/

Fellbaum, Christiane (1998) *WordNet: An Electronic Lexical Database*. Cambridge: MIT. See also http://www.cogsci.princeton.edu/~wn/

Jeremijenko, Natalie (2001) “OneTree”, http://www.cat.nyu.edu/natalie/OneTree/OneTreeDescription.html/ [still on page as of July 2001]

Kay, Alan (1984) “Computer Software“, *Scientific American*, Vol. 251:3 (September).

Kerne, Andruid (2000) “CollageMachine: An Interactive Agent of Web Recombination”, *Leonardo* Vol. 33: 5, 347–350.

Lang, R. Raymond (1999) “A declarative model for simple narratives,” in *Narrative Intelligence: Papers from the 1999 AAAI Fall Symposium*, Technical Report FS–99–01. Menlo Park, CA: AAAI Press.

Mateas, Michael (1999) “Not Your Grandmother’s Game: AI-Based Art and Entertainment”, in *Proceedings of the 1999 AAAI Spring Symposium, Artificial Intelligence and Computer Games*, SS–99–02, 64–68. Menlo Park, CA: AAAI Press.

Mateas, Michael (2000) “Expressive AI”, *ACM SIGGRAPH 2000 Art and Culture Papers.*

Mateas, Michael & Domike, S. & Vanouse, P. (1999) “*Terminal Time*: An Ideologically-biased History Machine”, in *Proceedings of the AISB’99 Symposium on Creative Language: Humor and Stories*, 69–75. The Society for the Study of AI and Simulation of Behavior.

musEleanor (1997) “Venues of Process: The Road Unraveled”, *Intelligent Agent*, Vol. 1: 11 (March). Http://www.intelligentagent.com/march_road.html/

Nelson, Ted (1981) *Literary Machines*. Sausalito, CA: Mindful Press.

Nelson, Ted (1986) *Computer Lib / Dream Machines.* [Originally self-published 1974] Redmond Wash: Microsoft Press.

Penny, Simon (2000) Comments at “Critical Technical Practices” plenary session, *Digital Arts and Culture 2000*, Bergen, Norway. Confirmed through personal communication.

Sack, Warren, “What Does a Very Large-Scale Conversation Look Like?”, in Wardrip-Fruin, Noah & Harrigan, Pat (eds.), *First Person: new media as story, performance, and game*. Cambridge, MA: MIT Press, forthcoming.

Sengers, Phoebe (1998) *Anti-Boxology: Agent Design in Cultural Context.* Ph.D. Thesis, School of Computer Science, Carnegie Mellon University. *Technical Report CMU-CS-98-151.*

Sengers, Phoebe (1999) “Practices for Machine Culture: A Case Study of Integrating Artificial Intelligence and Cultural Theory”, *Surfaces*, Vol. VIII.

Sengers, Phoebe (2001) “Intelligent Expressive Systems”, special departmental colloquium, Department of Computer Science, Courant Institute, New York University, January 2001.

Winner, Langdon (1994) “Political artifacts in Scandinavia: An American perspective“, *Scandinavian Journal of Information Systems* 6: 2, 85–94.