

2 Affordances of the Digital Medium

Think of the computer not as a tool but as a medium.

—Brenda Laurel (1993)

The Four Affordances

All **digital** artifacts are made of a common substance: programmable **bits** that can be used for symbol manipulation. As a result, we can think of digital artifacts as part of a common medium, rather than as a diffuse collection of multiple “New Media.” Looking at the **computer** as a single new medium of representation, we can see its defining representational affordances: The computer is **encyclopedic**, **spatial**, **procedural**, and **participatory** (Murray 1997). These four properties constitute our design space, the context for all of our design choices. Individual projects will be located at various points in that design space, exploiting one property more than another. But thinking of the potential for any project to more fully exploit each of the four properties can help us to discover new directions that we may have previously overlooked. *In approaching interaction design as a cultural practice our aim is always to make an object that is satisfying in itself and that advances the digital medium by refining or creating the conventions that best exploit these four affordances.* This chapter provides a brief overview of these four properties as well as the most important contributory disciplines that designers can draw on to maximize the expressive potential of individual digital artifacts and of the larger digital medium.

The Computer Is a Procedural Medium

The biggest difference between the computer and earlier media of representation is its procedural property, its ability to *represent and execute conditional behaviors*. As I have argued elsewhere, one of the most striking early signs that the computer was a new expressive medium was Joe Weizenbaum’s Eliza, an interactive program that was widely mistaken for an actual person (box 2.1). Eliza was experienced as real even

Box 2.1

Dialog with Eliza

Men are all alike.
IN WHAT WAY
They are always bugging us about something or another.
CAN YOU BE MORE SPECIFIC
Well, my boyfriend made me come here.
YOUR BOYFRIEND MADE YOU COME HERE
He says I am depressed much of the time.
I AM SORRY TO HEAR YOU ARE DEPRESSED
It's true. I am unhappy.
DO YOU THINK COMING HERE WILL HELP YOU NOT TO BE UNHAPPY
I need some help, that much seems certain.
WHAT WOULD IT MEAN TO YOU IF YOU GOT SOME HELP
Perhaps I could learn to get along with my mother
TELL ME MORE ABOUT YOUR FAMILY
(From Weizenbaum 1966)

though there were no images or sound track to reinforce the illusion. The program played the role of a nondirective psychotherapist responding to free text input (unrestricted words typed in ordinary natural language) into a Teletype terminal. Eliza had no understanding of what was said to her, but looked for keywords like “depressed” or “mother” and applied cleverly formulated rules of response, including the use of simple grammatical inversions to echo back statements. Despite her many inappropriate replies, and to Weizenbaum’s dismay, many interactors believed she was an actual person. She was animated by the power of her procedural design, by the ingenuity of the rules of behavior that determined her reactions to novel input (Weizenbaum 1966, 1976; Murray 1997; Wardrip-Fruin 2009).

Draw on Computer Science Concepts and Conventions

The discipline most relevant to procedural design is of course computer science. *Even for those who will not be doing the coding, understanding how computer science describes objects and processes is crucial to making sound design decisions.* A key element of this process is **abstraction**. Computer science strives to master **complexity** by creating abstract representations that describe elements of systems in the most general terms that most accurately describe their most salient features. For example, to describe the items [apples, bananas, grapes] in a single abstraction we might choose fruit, food, or groceries depending on the context. Of these three terms, fruit would be the best term

for limiting the category and establishing common qualities such as sweetness or perishability. If those qualities are not relevant to the task we are performing and if we want to add other items such as [bread, meat, cereal] to the system later on, then food would be a good choice. And if we are making a shopping list and distinguishing between hardware store and supermarket items then groceries might be best. The computational frame of mind sees the world as composed of multiple alternate abstraction systems such as this.

Most important, programmers abstract *behaviors*. Before a programmer writes a line of code they conceptualize the processes the computer will be executing as an abstract set of instructions and rules known as an **algorithm**. A useful algorithm describes a process at its most generalized level but in such a way that every important variation and condition can also be accounted for and responded to appropriately. *Deciding how flexible a program should be, how many possibilities it should anticipate and accommodate is an important part of the design process.* The more possibilities it encompasses, the more powerful and widely useful it will be, but also the more challenging to design.

Legacy media practices like animation and film editing specify unconditional sequences that are always executed in the same order. They are often represented in a **timeline**, and some computer-based authoring systems also use timelines as an organizing framework (figure 2.1). The problem with this approach is that it reinforces our tendencies toward **linear** or **unisequential** design. Programmable bits can imitate **legacy media** and present unisequential **documents** and film clips, but they are particularly well suited to more complex **multisequential** objects that can be assembled and navigated in more than one order. Computational structures allow us to describe entities as **variables** that can have different values at different times, and to make **conditional statements** that have more than one possible outcome. Objects in digital form can have multiple **instantiations**, existing as identical copies or as variant examples of a common pattern. Computational systems change over time, exhibiting different **states**. When we make something with computer code we are creating not a single version of an object or event, but many possible versions with interesting variations. Computational artifacts exist not as fixed entities, like books or movies (even though we may think of them that way), but as a set of easily altered bits governed by conditional rules.

By harnessing the procedural power of the computer to represent objects and processes we can create **simulations**, working models of complex systems that can be run with controlled variations and that aspire to reproduce the complexity we recognize in natural and social systems like the human body or the global financial markets. Such systems are made up of multiple independently operating **objects** or **agents**, whose complex interactions produce results too numerous and multicausal to be predictable in advance. Although many beginning programmers see the core computational structures as a **branching tree**, in fact the inner workings of computer code

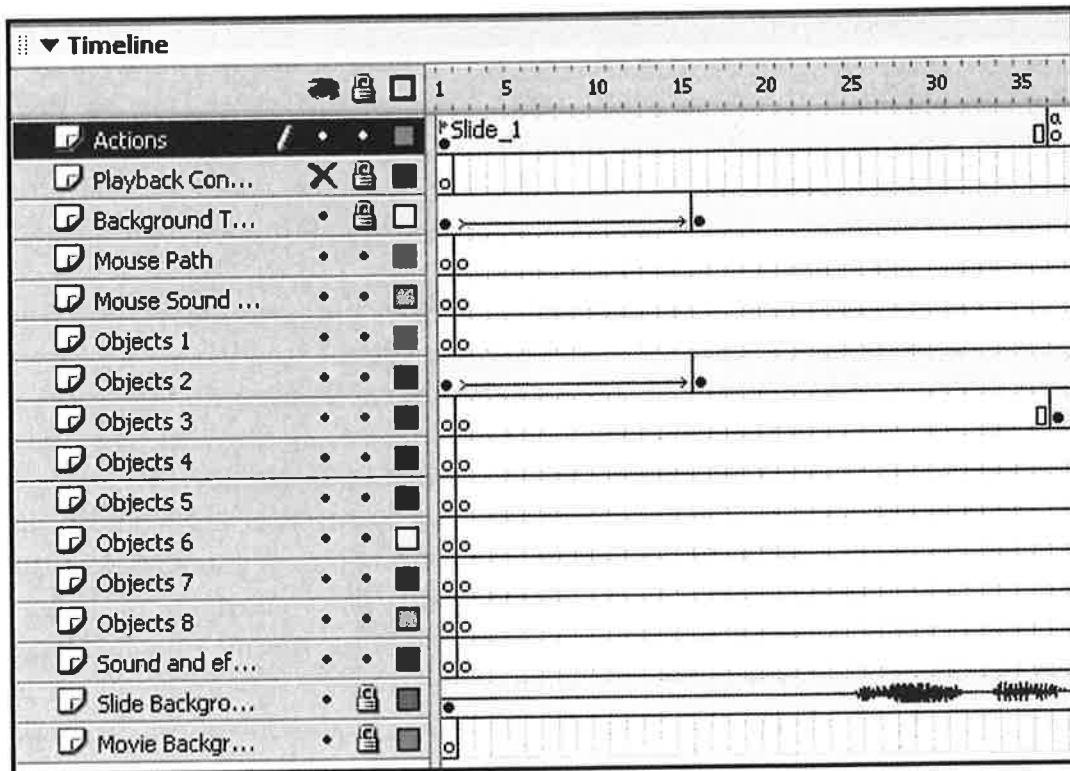


Figure 2.1

Timeline from Adobe Flash authoring environment, which is based on the **metaphor** of a movie. Flash also includes object-oriented scripting for more fully exploiting the procedural and participatory affordances of the medium.

have been growing more interactive: more like the ecology of a pond rather than the command structure of an army. Systems are now written that exhibit **emergent behavior**, behavior that is more than the sum of its parts, and therefore can be seen as similar to life forms.

Computer programs are judged by how efficiently and reliably they perform. Programmers aim for **robustness**, for not failing under a variety of error-inducing conditions, including coping with unpredictable user input and portability to multiple systems; and for **scalability**, for being able to accommodate more users, more data, more related procedures without having to be reengineered. The programmer values predictability and formula, aiming for the most generic, reusable solution that is also the most adaptable. The digital media designer, like the computer scientist and the programmer, should think of the process of representing meaning on the computer as a process of abstracting objects and behaviors as efficiently as possible, and should be aware of the possibility of representing complex systems as composed of multiple abstract actors in multivariable configurations.

In part III we will be looking more closely at the procedural affordance and identifying ways in which designers can make use of the power of computational abstraction for the purposes of describing objects and processes in the most coherent and expressive manner.

The Computer Is a Participatory Medium

It is surprising that the Eliza program fooled people when it first appeared, since it now seems so primitive in its ability to simulate a conversation that the tendency for people to assume that computer programs are more capable than they actually are is often referred to as “Eliza effect.” But Eliza also owed her success to an element of Weizenbaum’s design that is often overlooked: Weizenbaum did not merely script the machine; by framing his experiment in natural language processing as the highly conventionalized and familiar scenario of a therapy session he was also *scripting the interactor*. His character was successful because it exploited the participatory affordance of computer environments as much as it did the procedural affordance.

The relationship between the interactor and any digital artifact is reciprocal, active, and open to frustrating miscommunication. *The designer must therefore script both sides so that the actions of humans and machines are meaningful to one another.* Sometimes the script is quite rigid, as in a touchscreen ATM machine or a phone-based automated customer service system. Such rigid systems work best with screen-based multiple-choice answers, where there are a limited number of routine transactions, with each one taking few steps to complete. When used for more complex transactions, like the decision-tree diagnosis of a consumer problem, they can be maddeningly frustrating.

Sometimes the script is more flexible, offering the interactor an array of props with which to improvise their part of the exchange, like the iconized tool bars of word processors or role-playing adventure games. Some digital conventions are so familiar that they script us in a **transparent** way. For example, a blinking insertion point in a text box is a transparent cue to type in information; blue or underlined text on a website cues us to the presence of a link; arrows at the left edge or bottom-left corner of application windows cue us to the possibility of scrolling. *A large part of digital design is selecting the appropriate convention to communicate what actions are possible in ways that the human interactor can understand.*

Because the computer is a participatory medium, interactors have an expectation that they will be able to manipulate digital artifacts and make things happen in response to their actions. They will therefore become frustrated and impatient when they are not allowed to act. The responsiveness of digital media, whether accessed through a keyboard, mouse, joystick, touchscreen, scroll wheel, or gesture sensor,

excites our desire to do something, to see what will happen if we drag something around, click on an underlined word, or otherwise poke at the environment.

Participation in digital media increasingly means social participation. Previous mass-communication technologies provided either one-to-one (telephone) or one-to-many (books, television) **transmission** channels. The computer provides one-to-one (e.g., email) and one-to-many (e.g., DVDs) communication. It also provides new forms of many-to-many communication, most notably on the **World Wide Web**, which provides the same potential distribution to an album of baby pictures as it gives to CNN.com's coverage of a presidential inauguration, disrupting the media conventions for both. New participatory **genres** such as chat rooms, bulletin boards, discussion lists, blogs, wikis, instant messaging formats, virtual environments (Second Life), social networks (Facebook, Myspace, Twitter), and media-sharing technologies (Flickr, YouTube) convene communities of participants in discussions that are **synchronous** and **asynchronous**, spoken and written, individual and collective. They also quickly assimilate the functions of one another, posing new design challenges in providing coherence to a sustained collective conversation (see figures 2.2, 2.3).

Mass participation in digital environments has also raised questions of security and privacy. Designers have the power to capture input from users without their intentional action, and even without their knowledge and consent. As citizens of a digitally enabled world we are subject to many kinds of monitoring. Digital cameras record our images, digital networks record our purchases, and global positioning technology in our cars or cell phones pinpoints our personal location (figure 2.4). We are open to photographic and sound recording through portable devices that were once the domain of superspies and are now available at mass consumer prices. This capacity can be reassuring, as in devices that allow the elderly to call for help even if they are far from a telephone; or it can be menacing, as in spyware programs that surrepti-

 **News Feed****Top News** · Most Recent **21**

What's on your mind?

Figure 2.2

Status line on Facebook. The status line, cued here by a generic question, is a convention that has been widely accepted as a way of formatting participation in large common spaces. The box provides a constrained container (even when longer posts are allowed) that focuses contributions so they can be composed easily, distributed widely, and collected in lists that can be scanned quickly by the recipient. This simple convention—a limited size, rapidly posted container for messages that can be targeted at one or many recipients—has created new, globally accessible channels of communication for personal and organizational use.

Most Recent Customer Reviews

★★★★ Great product!

Used in a small room this product produces a noticeable difference in the air quality in a short period of time.

[Read more](#)

Published 5 days ago by

★★★ Air Purifier

It's a bit on the small side and I'm not sure how effective it is but, so far, I'm satisfied.

Published 6 days ago by

★★★★ honeywell air purifier table model

so far the air purifier has performed very well making a large contribution to my health and well being.

Published 6 days ago by

★★ Crappy little fan filter at a crappy little price

I generally makes good products (I own another all that I really like) but this model is very cheaply made and noisier than it should be for its size.

[Read more](#)

Published 8 days ago by

★★★★ Great purchase - want to get one for each room in the house!

For the price, this ~~is~~ ~~was~~ air purifier was a fantastic deal for me. I have terrible allergies and I often wake up with my eyes crusted over, my nose congested, and my throat... [Read more](#)

Published 8 days ago by

Figure 2.3

It has now become standard to provide a space for visitors to a website to talk back, like this customer review utility on Amazon.com. Other standard feedback conventions include email links, live chats, Frequently Asked Questions (FAQs), and polls. The refinement of these features and the invention of new ways to support participation will be a growing area of digital design for decades to come.

tiously record our keystrokes and page views in order to steal passwords or deliver unwanted advertising. The increasing mediation of our actions by digital technology has led to new kinds of counterfeiting, including the crime of identity theft in which possession of data about someone becomes the power to act as that person. For the purposes of many social transactions, we are identical with the sum of the digital information we have input. *Designers must therefore take care to limit access to information and to give interactors control over their own information and knowledge of how it is being collected and used.*

The participatory nature of the medium has also profoundly affected legacy media enterprises. Media now appears to us as something to be cut, pasted, reassembled, and




Figure 2.4

Google Map Street View of San Francisco includes a figure in a potentially embarrassing moment, an example of how once-private experiences are becoming publically—and even globally—visible.

distributed with ease. The music industry underwent a tremendous change starting in the late 1990s when users began trading digital music files over the Internet. As sales of CDs fell and threats of prosecution escalated, Apple Computer came up with a design solution that included a commercial solution: the introduction of the iTunes website and the iPod device gave users a legal method for collecting digital music files. As increasing bandwidth and memory make it easier to collect moving-image files, movies and television face similar challenges and the need for similar design solutions. *One of the key tasks for information designers is to satisfy increasing demand for access to media while preserving the property rights of those who create media artifacts.*

The production of entertainment is also becoming more participatory, with digital recording and editing equipment, and development platforms for interactive gaming, that were formerly the prerogative of studio professionals now available at consumer prices. One result is the exponential explosion of user-generated and self-published content, often indexed by other participants through popularity ratings (figure 2.5).

Participatory structures are increasingly being incorporated into legacy frameworks. In England, television viewers have access to a remote control with a special button for

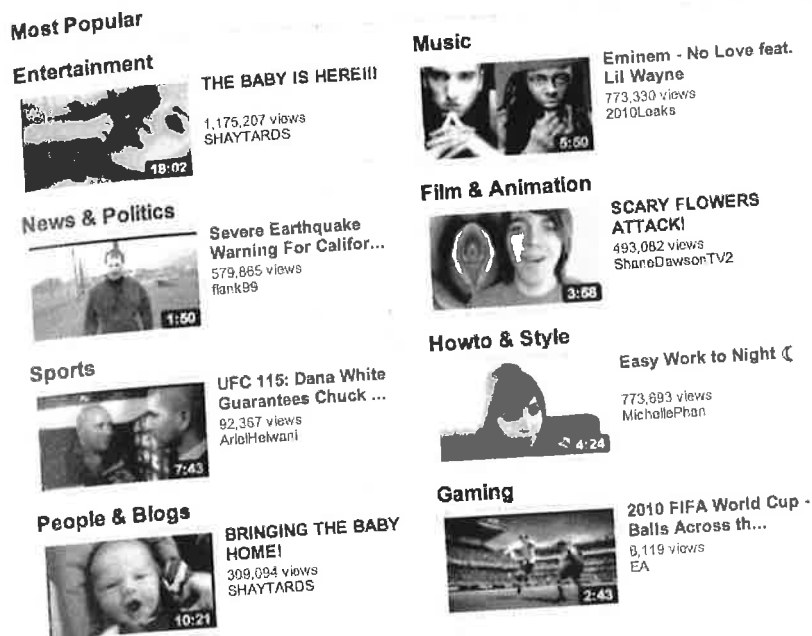


Figure 2.5
YouTube displays user-created videos based on popularity.

interactivity, and interactive enhancements are added to hundreds of television shows every year, giving viewers a choice of which tennis match or garden tour to watch, letting them vote for their favorite British person, inviting them to take an IQ test in sync with a live audience, or allowing them to participate in interactive stories (figure 2.6). As television merges with games and the Internet and more powerful computing is incorporated into media players of various kinds, the opportunity for inventing new participatory conventions for entertainment and for information resources will expand.

Draw on HCI Concepts and Conventions

The discipline that is explicitly concerned with the study of the relationship between human beings and computers is **human-computer interaction (HCI)**, which has its roots in industrial design. In fact, one of the most influential books for the interactive designers, Donald Norman's *The Design of Everyday Things* (1988), hardly mentions the computer. Norman, who began as a cognitive psychologist, called attention to the importance of the **mental model** formed by the user based on the appearance and behavior of the object. Mental models can derive from existing conventions and past experience. For example, we expect wall switches to turn on overhead lights. But they also derive more generally from what we perceive as the **affordances** of the artifact,



Figure 2.6

Red Button TV in Great Britain was among the first widespread interactive services to be delivered as part of the standard television package. Interactive features are accessed through a dedicated red button on the remote control, including on-demand content like news and traffic reports, local government public services, and multifeed coverage of sports events and concerts.

that is, the functions that the physical properties of an object make possible. Wood, for example, affords carving and burning; blackboards afford writing and erasing. Other affordances are more like cultural signals to the users, like the door handle that affords pulling, triggering the mental model that the door will move toward us in response to our action. Much of our confusion in coping with an unfamiliar environment like a hotel shower in a foreign country, or a rental car in an unfamiliar make and model, derive from poor signaling of what functions a hidden knob or mysterious gizmo affords. Mental models are particularly important for machines whose workings are usually hidden from us. We need to know not only the results of a command but

also the causal relationship: such as why a digital music player has chosen a particular song, a search engine has prioritized certain sites over others, or a computer operating system has suddenly caused all my working windows to disappear.

In digital design we must ask ourselves what mental models the interactor will bring to the object. For example, if we make a computer “notepad” how far should we follow the model and mimic the affordances of the paper notepad? Should we allow the user to turn its pages? To write on the “reverse side” of a page? To tear out a page and save it separately? With machines whose workings are complicated and hidden Norman advocates the key design principle of **visibility**. We should be able to see what the machine is doing—not by looking at the gears, but by understanding its behavior as determined by a logical system of cause and effect. In particular, we should have clear **feedback** from the machine on the effects of our own actions. For example, in an operating system, the mini-animations of a clock face or an hourglass are feedback mechanisms that acknowledge our commands and indicate the machine is working on executing them. In an electronic notepad, an animation may reinforce the model of turning a page as feedback for the interactor’s gesture and reinforcement of the legacy model.

In digital as in mechanical design, our goal should be to remove distraction, to allow the user to focus on the task at hand rather than the tool, by making the interface elements transparent, and allowing the interactor **direct manipulation**, as we do, for example, when we drag a file folder to the trash can. **Usability** is a key concept in HCI. The usability design process starts with the gathering of all the stakeholders of a system, with an understanding that there will be multiple kinds of users of the end product and multiple criteria on which that final product will be evaluated. Such a process identifies tasks and models alternate ways of performing them, relying on repeated testing and revision with small groups of users, including groups that work with paper **mock-ups** before anything is built. **User testing** often focuses on the performance of tasks, and takes quantitative data on how long it takes new or experienced users to complete them, and how many mistakes these users make in performing them. Users may also be asked to provide **think-aloud protocols** describing their assumptions, intentions, and reactions as they use an object, which can provide insight into the mental models they are forming. Such testing can be part of an *iterative and formative evaluation* in which the designers repeatedly test and redesign versions of the product. *Human-centered design* is an important usability concept, emphasizing the physical and cognitive capabilities and limitations of human beings. Keyboard design, for example, is the subject of ergonomic standards that take into account the size of the human hand and the positioning that creates the least strain (figure 2.7). (For more on usability guidelines, see Jakob Nielsen’s website <<http://useit.com>>; government guidelines at <<http://www.usability.gov>>; and the guidelines issued by the International Standards Organization at <<http://www.iso.org>>).



Figure 2.7

The Microsoft Natural Ergonomic Keyboard addresses the problem of repetitive motion injuries by providing wrist support and an angled split design to create a more “natural” rotated positioning for the hands. The buttons in the middle just above and below the spacebar are dedicated to zooming in and out on and moving between web pages, reducing the need for mousing.

Usability studies work particularly well within the model of the computer as a tool or machine, requiring a “user” whose main goal is efficiency in performing a task. But interactors can be understood in larger contexts than as tool users and task performers. They also make complex judgments about what they want to know, what they want to do, and where they want to go in the digital realm. They make these judgments based on social and cultural contexts that can resist quantification.

Social anthropology, for example, has introduced the useful concept of situated action, as a way of understanding human–computer interaction. **Situated action theory** sees users not as interchangeable processors of information, who all behave the same way according to pre-scripted, machine-friendly rational plans, but as uniquely positioned, complex actors whose frame of reference is shaped by the surrounding social and material world. Lucy Suchman, studying the interaction of users with copier machines in the late 1980s, demonstrated the many ways in which their model of where they were in the process of copying and binding a document at any given moment differed from the machine’s model, and from that of the Help system. Instead of asking the user to conform more fully to the **machine model** of a generic process, Suchman argues for an embodied and relational model, in which humans and machines are constantly reconstructing their shared understanding of the task at hand (Suchman 1987) (figures 2.8, 2.9).

Ethnographic research, grounded in cultural anthropology, is a fruitful way of looking at interactors in online social networks and virtual communities as cultural



Figure 2.8

A GPS navigational system must understand action as situated, adjusting to the current position of the car and the driver's real-time changes in position, rather than expecting drivers to conform to a single invariant route.



Figure 2.9

YELP provides location-sensitive information, including many user reviews of the same venue, and it does so for multiple platforms, illustrating three different ways in which the same interaction can be contextualized or "situated."




Figure 2.10

Celia Pearce has documented the ways in which players who were exiled from the short-lived Uru, a massively multiplayer online role-playing game (MMORG), have reconstructed elements of the game's culture in other online environments. Shown here is the Uru Fountain, an important gathering place, in (clockwise from upper left): Uru, There.com, Adobe Atmosphere, and Second Life.

beings, actively creating and negotiating the meaning of shared actions and objects. For example, Celia Pearce has found striking resemblances between the community-making behaviors of a gaming community displaced from one virtual world to another and those of similarly displaced émigré communities in the real world (Pearce 2009) (figure 2.10).

Another useful approach to capturing the many ways in which objects participate in systems of meaning is **participatory design**, which engages the potential users of a new system in every step of the design process, not merely as informants or testers but as collaborating members of the design team. Participatory design emphasizes the empowerment of marginalized stakeholders in proposed designs, such as poor people receiving a government service like a homeless shelter, workers rather than managers in a business introducing a new IT system, or people without technological expertise who are expected to be the users of a new device. Participatory

designers look to such groups as providing valuable knowledge of the relevant real-world contexts and of the core human needs underlying the design brief (Dearden, Rizvi 2008).

Individual experiences with interactive environments can vary widely from one another and from the designers' expectations. The emotional component of interaction, which is important to understanding players' engagement with game environments, is particularly hard to capture with experimental user testing. Players may suppress emotional information in nondirective think-aloud protocols, and directed questioning can limit responses to predetermined categories. One promising approach is to ask game players to provide feedback by manipulating evocative but nonspecific tactile figures that serve as symbols for whatever emotions the individual player associates with them (Isbister et al. 2007) (figure 2.11).




Figure 2.11

The Sensual Evaluation Instrument (SEI) used by Katherine Isbister and her collaborators to engage interactors in playful self-assessment of their emotional responses to computer systems. The system is composed of eight clearly differentiated, warm-to-the-touch, nonrepresentational objects created by a sculptor. The objects are meant to be appropriate to the kinds of feelings reported by interactors, such as confusion, surprise, and flow, without imposing any particular meanings. Subjects tend to associate more negative feelings with the spiky objects and more positive feelings with the smooth objects.

Participation in digital environments is always part of larger social and cultural contexts that the designer should bear in mind. In part V we will be examining several social and cultural models of interaction that designers can draw on to script the interactor.

The Computer Is an Encyclopedic Medium

The most capacious medium ever invented, the computer can contain and transmit more information in humanly accessible form than all previous media combined. Technically, this larger capacity is a function of engineering successes in electronic circuitry, signal processing, and a symbolic system of logical codes. Since the 1960s, advances in silicon technology have led to the creation of ever-larger repositories of bits in ever-smaller and cheaper machines. In 1965, Gordon Moore noticed that new computer chips seemed to be released every eighteen to twenty-four months and that each chip doubled the processing power of the preceding one. This rapidly accelerating growth became known as **Moore's Law** and has held true for almost half a century. The removable tape drives of the 1960s held two million characters each (two **megabytes**, or approximately two books worth of information), and were the size of a household refrigerator. By the end of the first decade of the twenty-first century we were measuring personal computer capacity in terabytes and carrying around on our key chains removable storage equivalent to dozens of those refrigerator-sized tape drives. The size of the available storage space and the ease with which it can be retrieved and displayed also has meant an expansion in the kinds of media formats that can be inscribed and transmitted as electronic bits. New technologies of signal processing have made it possible to represent audio, images, and moving images in digital form and to display them in a single environment. The phenomenon of media **convergence** challenges designers by creating a convergence of conflicting conventions with which to structure information. The encyclopedic capacity of the medium raises the expectations of the designer and the interactor, making it important to communicate the limits of any collection, and to make clear decisions on what lies inside and outside of the project's scope (figure 2.12).

I use the word "encyclopedic" to refer to both a technical and cultural phenomenon: to the unequalled storage potential of the new medium and to its promise of an infinite tablet, a library as big as the world. As a culturally encyclopedic medium the computer inherits age-old traditions of human culture, expressive of our core need to collect, preserve, and transmit knowledge across generations. Homer's epics, for example, include encyclopedic set pieces, such as a lengthy catalog of ships, similar to lists in other bardic works of preliterate **orature**. With the advent of a written alphabet the encyclopedic impulse appeared in scrolls of magical spells, detailed histories, Aristotle's knowledge compendia, and Pliny's thirty-seven-"book" survey of

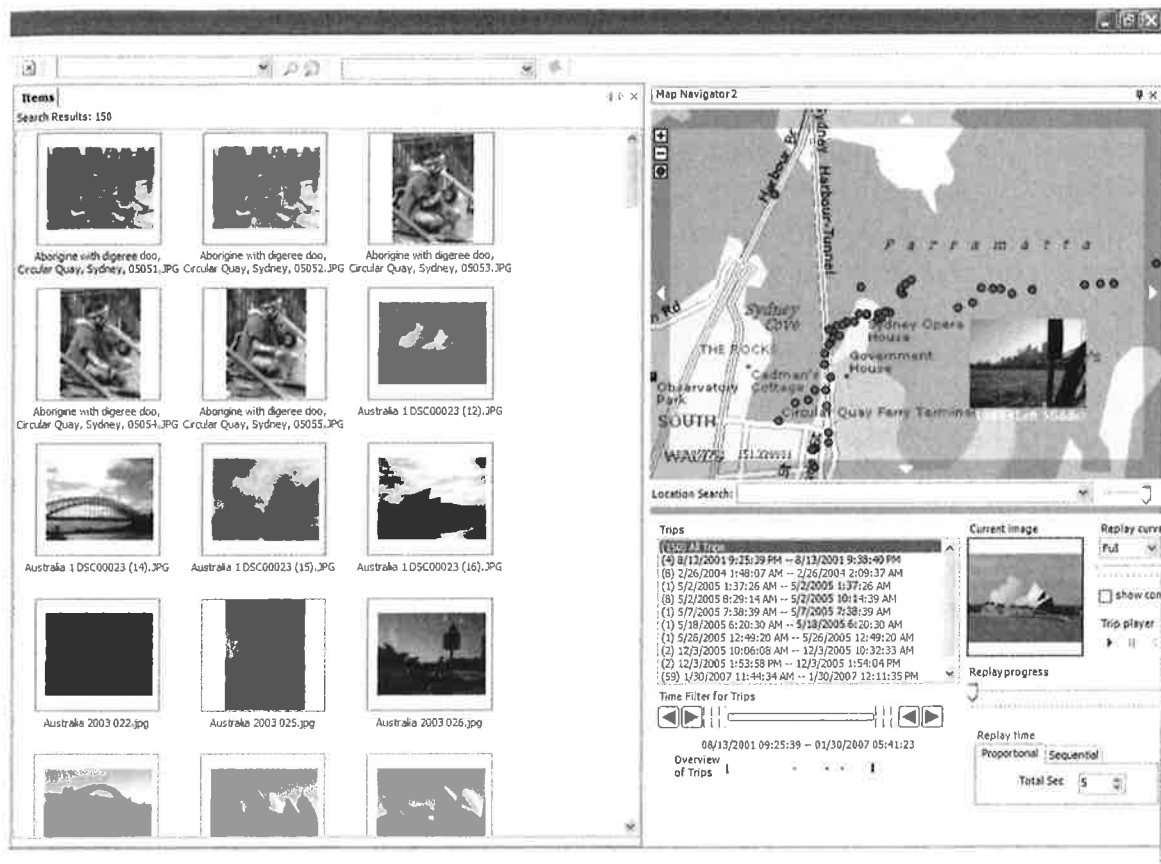


Figure 2.12

The “My Life Bits” Project is an attempt to record the life of Gordon Bell, a Microsoft researcher, as completely as possible, including photographs, video, and audio recordings of events as they happen. The challenge it poses is how to annotate such a wealth of data so that it can be retrieved. Merely recording everything would be similar to recording nothing, since human memory relies on forgetting and selecting. See <<http://research.microsoft.com/en-us/projects/mylifebits>>.

“Natural History” published in the first century A.D. With the invention of the printing press, the accumulation of knowledge took on new, standardized forms, from scientific treatises to picaresque novels. By the middle of the eighteenth century, the editors of *Encyclopedia Britannica* felt confident enough to attempt an alphabetical survey of all of human knowledge, a project that lasted for two hundred years, with increasing numbers of volumes, and served as a respected and useful reference work for shared authoritative knowledge until overtaken at the turn of the twenty-first century by Internet information resources. The human encyclopedic aspiration is also reflected in data collection, for example census taking; in narrative traditions, such as Shakespeare’s history plays and the novels of Charles Dickens, Honoré de Balzac, and George Eliot; and in the increasing development of film and television series that

follow multiple related characters and stories across time and space. All of these legacy genres are encyclopedic in scope, presenting sweeping panoramas and close-up detail of a particular complex society at a particular moment. *Their strategies of organization are part of the rich palette of media conventions that the designer can draw upon to create encyclopedic digital artifacts.*

The designer can also draw upon the conventions of emerging encyclopedic digital genres: vast databases such as world health surveys, stock market prices, soccer statistics, and recordings of electromagnetic impulses collected by spaceships exploring distant planets; massively multiplayer virtual worlds in which hundreds of thousands of people share a common landscape; a World Wide Web of “pages” describing everything in the world that can be captured in text, image, or sound; participatory information structures like Wikipedia that are created by multiple contributors acting independently within a common information structure; maps and satellite image systems that reproduce the whole earth in navigable form at multiple levels of **granularity**. The creation of encyclopedic networked resources will be a continuing challenge for digital design, posing the problem of how to shape our expanding capacity into coherent **genres** so that we can access information without being overwhelmed.

Draw on Information Science Concepts and Conventions

The exponential increase in information presents us with the challenge of finding more powerful means to organize it than the ones we have developed for legacy media. In doing so we can build on the methods of information science, which are themselves changing to accommodate the convergence of libraries, document archives, and museums. Information science is faced with the task of inventing the transmission technologies—the common standards for preserving, describing, cataloging, and accessing an exponentially expanding global knowledge base, the vast repository of human culture in all its artifactual manifestations.

The digital medium differs from legacy media in that it can store the **artifact** itself—the equivalent of the book, the audio tape, the film canister—as well as a description of the artifact, and it can **link** one to the other. The Google Book Search Library Project (<<http://books.google.com>>) is digitizing millions of volumes and making them directly searchable in the same way that web pages are searchable (figure 2.13). As scholarly journals, newspapers, magazines, and books move to digital delivery formats it will become increasingly possible to archive them in a manner that makes them open to free text searches for words and phrases. But free text search of large information resources (like typing words into a web search engine) produces too much information. The ability to find every mention of the word “freedom,” for example, in every book in the Library of Congress is a daunting prospect. *The challenge for designers*



Figure 2.13

Google Books has scanned over seven million books held in research libraries, and it provides an online search that returns titles, images, and full or partial access to the text, depending on copyrights. It can also search for words and phrases within a single book.

is determining how to organize information so that it is retrievable in a coherent form that minimizes confusion and maximizes understanding.

One of the key affordances of moving encyclopedic resources to digital form is the ability to index and retrieve information at multiple levels of granularity. We can search or browse by titles of books or by chapter headings or single words within multiple 100,000-word texts. We can retrieve whole books, or particular pages within a book. Furthermore, when we bring multiple media into a single format we have a chance to rethink the **segmentation** conventions with which we index, retrieve, and display them. Should films be accessed by title, chapter, shot, or frame? Or perhaps by line of dialog? When we make new resources for digital delivery we have the opportunity to segment information for the screen, and to avoid physical interruptions—like the end of a page—of sections that should be kept together by content. Design of encyclopedic resources involves active decision making on granularity and segmentation conventions, and careful assignment of appropriate **labels**.

An important goal of archive design is support for **intellectual access** to individual artifacts, that is, retrieval of the appropriate item not by its size or shape or media type or donor's name, but by the relevant conceptual terms. This usually involves an

act of **classification**, of assigning objects to standardized categories, often through the use of **metadata**. Classification involves establishing **controlled vocabularies**, words that everyone agrees to use in the same way to refer to a class of artifacts. The subject index of the Library of Congress is a controlled vocabulary, as is the Getty Art and Architecture Thesaurus, which provides a **taxonomy** for everything from arches to urns. The virtue of a controlled vocabulary is that it allows us to juxtapose separate items that relate to one another by virtue of the fact that they carry the same label. In a library environment the number of labels by which we can juxtapose individual units of information is limited. We can sort them onto separate shelves, and we can cross-reference the shelves by assigning multiple subject labels. But one book can only be in one place at a time, and so can one index card describing the book. If a book is to be listed by various subjects it must have separate cards for each one, which limits the number of entries per book. But a digital artifact can be called up in multiple contexts without having to be duplicated. It takes up no more room in the computer's memory, although it may appear on hundreds of virtual shelves. Every query to a search engine constructs a new shelf, customized to the individual searcher.

The values of information science are concerned with minimizing ambiguity and maximizing accessibility. A well-designed information structure is open-ended and can accept new elements without confusion about where they will be placed. It provides a means of finding items without memorizing the controlled vocabulary words, either by providing them as a navigation system, as in the stacks of a library, or by offering an intermediary—human or mechanical—who can find the right label when offered a synonym or related term. With the right labels attached to the right items, we can find any needle in the rapidly expanding global haystack. But our methods for labeling items of information are lagging behind our technologies for retrieving and juxtaposing them. Designers in the twenty-first century are faced with the need to bridge this gap, and must find new ways to maximize the power of the digital environment to retrieve information with precision and to juxtapose information from varying sources in coherent ways. We will explore these approaches further in part IV.

The Computer Is a Spatial Medium

Since space and time are the two foundational coordinates of human cognition, we experience everything spatially and we have many genres for representing it, such as paintings, sculpture, architectural plans, and film. But the computer constructs space in a different way from other media: it creates virtual spaces that are also navigable by the interactor. Digital space is created out of bits rather than bricks, and it rests upon the procedural and participatory affordances of computation: it is navigable because it responds to our navigational gestures in a consistent manner.



Figure 2.14

The Xerox Star (1981) was the first commercially produced computer with a GUI (graphical user interface). It included icons for files and windows for applications and it was mouse driven, which became standard features of the familiar desktop interface. There was no fixed desktop menu; the keyboard had several dedicated function keys. The mouse was used to select an icon, but a physical “Move” key was used to reposition it, in contrast to the later click-and-drag convention. The functionality of the interface is well documented in videos available at <http://www.digibarn.com> or <http://video.google.com>.

The **spatial affordance** of the computer does not rest on its ability to present us with images of real-world spaces (Murray 1997): for example, we perceive web pages as occupying “sites” that we “visit.” But the addition of spatializing images to computer displays has been an important part of the evolution of information organization. The familiar desktop **graphical user interface**, the world’s most successful GUI, was invented at Xerox PARC in the 1970s, further developed by Apple Computer as the commercially successful Macintosh in 1984, and fully established as the dominant world standard when it was deployed in a robust PC version in 1995. The objects and the space that make up the computer desktop—the file folders, windows, manipulable icons, and menus—were developed over time, not as literal reproductions of physical objects, but as **abstractions** that include only those spatial properties that reinforce their functions as **chunks** of information and programming code (figures 2.14–2.16). Computer file folders have customizable labels like real file folders but they are not depicted as vertical objects stacked in wooden drawers. Because there is no need to mimic the extraneous tasks of the physical world, like opening a drawer, we can focus our attention on the placement of the file within the **hierarchy** formed by the labels. The desktop GUI is now so well established that it is regularly lamented as too limited,

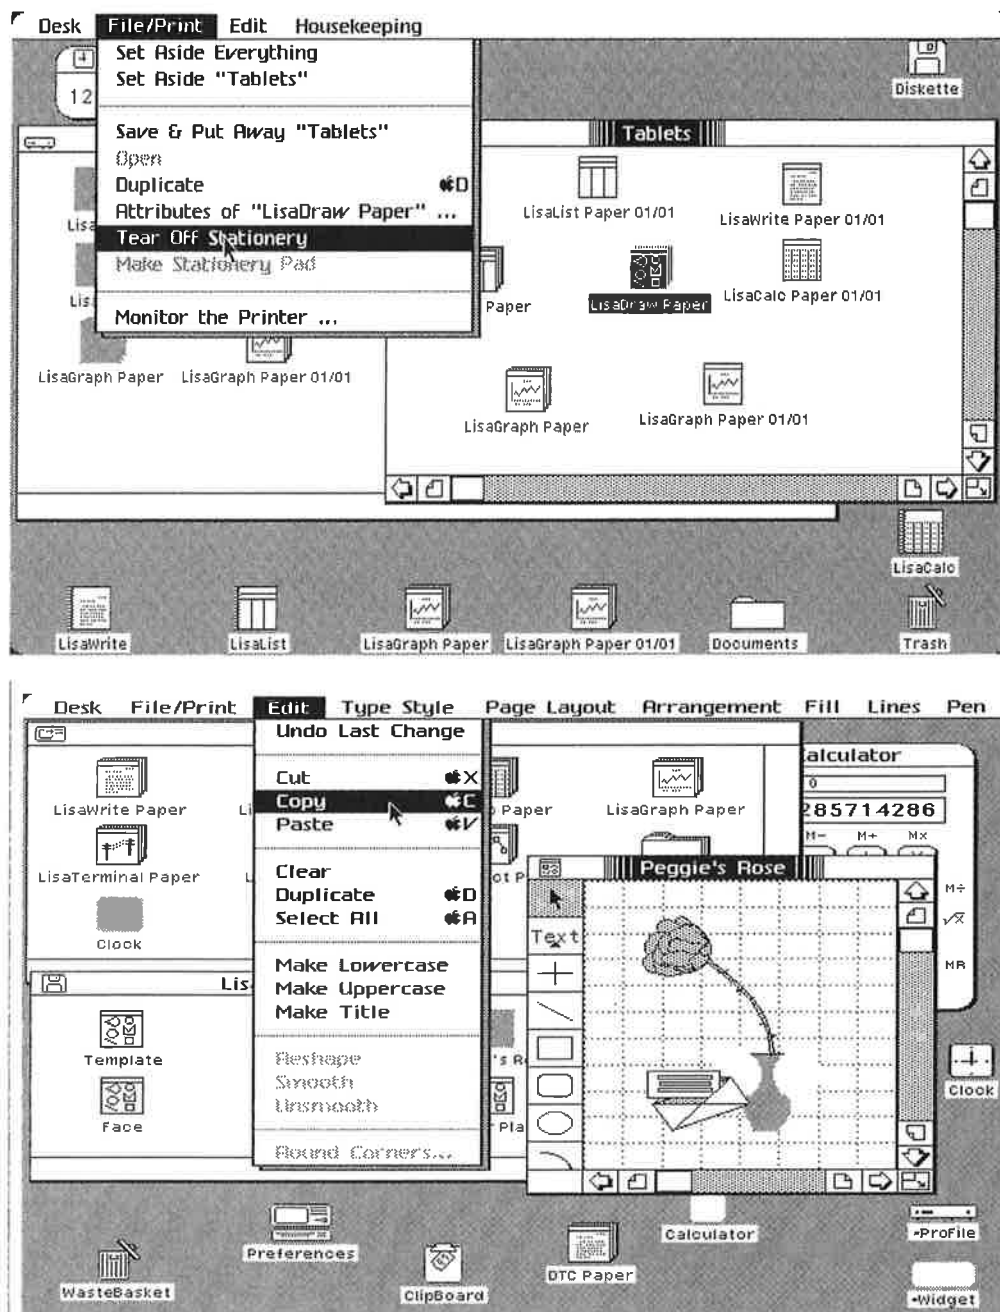


Figure 2.15

Apple Lisa GUI (1983) includes windows, files, icons (including the trash can), and a (rather verbose) top **menu bar**. The operating system, called Desktop Manager, was document driven like the Xerox Star, rather than application driven like the later Macintosh. To create a new document the user would “tear off” a new piece of virtual “Stationery.” Applications (or “tools” as they were called) like word processors would start up only when documents were opened and were otherwise hidden from the end-user within the operating system.

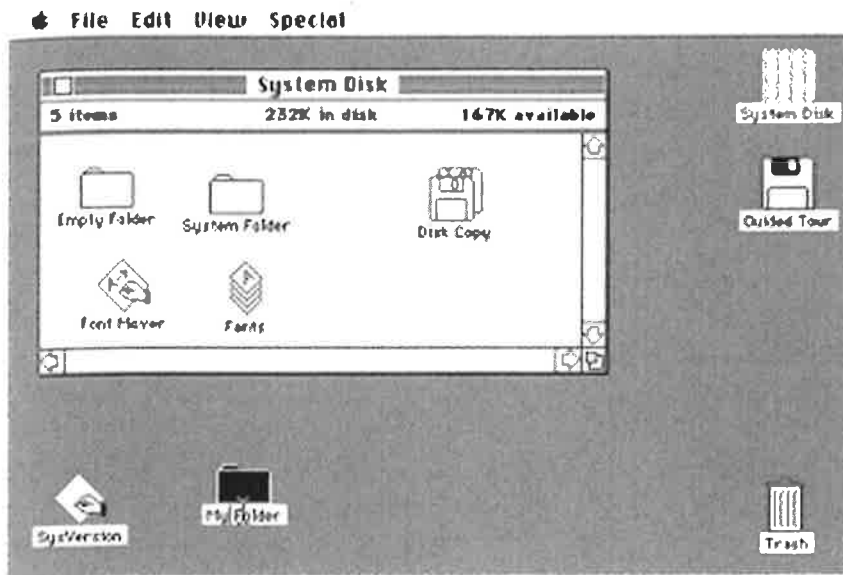


Figure 2.16

The Apple Macintosh GUI (1984) included windows, files, click-and-dragging of icons, and a menu bar at the top. The Macintosh, inspired by the Xerox Star, initiated the full WIMP interface we have come to think of as the standard for personal computers and workstations.

and even denigrated with the acronym **WIMP** (Windows, Icons, Menus, Pointing devices), but it has yet to be superseded, although web browsers, mobile devices, and entertainment systems offer partial alternatives.

By the 2000s proliferation of applications on the personal computer led to the addition of a “dock” at the bottom of the desktop screen, an abstraction of the function of docking a ship that allows users to park favorite applications in easy reach of top-level pointing. Desktops, windows, sites, tabs, and docks all occupy very different frameworks in the real world. We don’t find file folders floating with yachts, or windows arranged in tabs within our physical world. But as abstractions of spatial functions these conventions have proved to be learnable and are now transparent to users. However, the space of **applets** on mobile devices remains disorganized and will no doubt be an active focus of design in coming years.

Some of the earliest applications of computers, even before the wide availability of graphical displays, were spatial applications including *Spacewar!* (1962), the first action video game, and the *Aspen Movie Map* (1978), which preceded the Google Street View by two decades (figures 2.17, 2.18). Digital technologies have developed the two-dimensional screen into increasingly realized three-dimensional spaces, both for conventional movies and for interactive environments. For example we can fly through highly detailed models of real places like Disney World (figure 2.19), virtual



Figure 2.17

The Spacewar! game created for the PDP-1 computer (1962) was the first representational video game. It was created as a demonstration of the new graphics system by a group of researchers who were avid pinball players and science fiction fans (Graetz 1981).




Figure 2.18

Aspen Movie Map (1978) one of the first demonstration projects of the group led by Nicholas Negroponte and Andy Lippman that went on to found the MIT Media Lab. It made use of early optical videodisk inscription to create a computer-driven, interactive experience simulating a drive through parts of Aspen, Colorado, including the ability to make right-angle turns. In 2007, Google Street View introduced a similar interaction pattern.



Figure 2.19

Fly-through of Disney World 3D model in Google Earth.

communities like Second Life, and inaccessible abstracted places like a detailed model of the inside of a cell or a chemical molecule (figure 2.20). As with video games, we can often switch camera positions within these environments and zoom in and out from overview to close-up. Such environments can be both captivating and overwhelming. Designers face challenges in providing a consistent model of the space as a whole and in making individual areas memorable and findable (Nitsche 2008). These challenges are similar to the design problems of filmmakers, landscape architects, and urban planners, combining problems of layout with consideration of point of view.

Draw on Visual Design Concepts and Conventions

The visual organization of digital environments should reflect the core principles of good graphic design, including the design values of *contrast and regularity*, which allow us to focus the interactor's attention on the key informational elements and their relationship to one another. For example, we establish table-like **grids** for web pages and assign specific horizontal and vertical areas to specific purposes such as a site logo, a navigation menu, or a news box. Regularity of proximity, size, color, and font allows users to recognize similar items and to distinguish dissimilar items from one another (figures 2.21, 2.22).

Because we are attuned to the spatial characteristics of digital environments, *users assume that spatial positioning is meaningful and related to function*. If an item calls attention to itself with a different color or size we will try to touch it or click



Figure 2.20

Three-dimensional, navigable visualization from the Allosphere, a huge scientific visualization facility similar to a planetarium, designed by JoAnn Kuchera-Morin at the University of California in Santa Barbara. Image is from an informational video at <http://www.allosphere.ucsb.edu>. Such installations hold the promise of new insight into otherwise unobservable phenomena, and they pose new challenges for the navigation of three-dimensional space.

on it, assuming that it has some behavior or link associated with it. If two items are next to one another on a list we will assume that they are parallel and behave in similar ways. Designers must therefore be careful not to introduce purely decorative elements that resemble icons or linked text. Such items trigger expectations of interactivity and lead to user frustration when they do not respond to the user.

Good graphic design matches visual elements with the meaning they are meant to convey, avoiding distraction and maximizing meaning. Graphic presentation of information can help to explain relationships and serve as an aid to brainstorming and discussing structure at every stage of the design process. Even those who are not good at sketching can use the organizing power of graphic design to clarify and communicate complex ideas by making use of graphic design software, including the many free charting and visualization programs available on the web. Information designer Edward Tufte, an influential exponent of the explanatory power of graphics, refers to extraneous decorative detail as “chartjunk” and insists that all the aspects of an image serve to convey salient information (Tufte 1983), such as magnitude, time, and causality. By these standards, Tufte has identified Charles Joseph Minard’s 1869

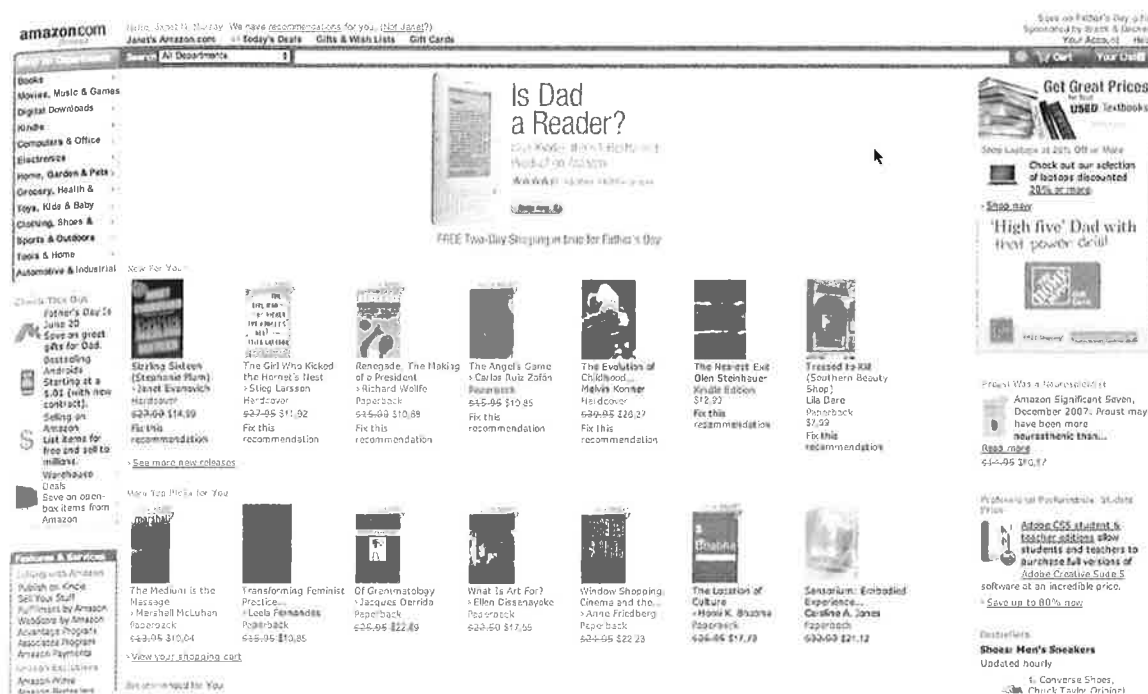


Figure 2.21

Amazon.com provides a great deal of coherent information on this page by using a grid that places navigational menus in the left-hand column, individual thumbnails of books in their own uniform grid within the middle column, and promotional displays in the right-hand column. The generous white space makes the page more readable and less aggressive. The banner at the top of the page includes the standard conventions of shopping cart and account access but is minimized to focus attention below.

map of Napoleon's campaign against Russia as "probably the best statistical graph ever made" because it conveys so many aspects of a complex event in a single unified presentation, including the geographical route of the army in advance and retreat, the size of the army at each point in the journey, and severe winter temperatures (figure 2.23).

Tufte's focus is on visual design, not as an end in itself, but as a means to communicate unambiguous information. Digital designers must be equally conscious of the weight of each design choice and equally critical of the merely decorative. But the culturally informed designer is also skeptical of claims to impartial factuality and to a simple one-to-one correspondence of code to meaning. As Tufte points out, the Minard map reflects antiwar values in its selection of details: it does not include the name "Napoleon" and is designed to dramatize the contrast between the large number of soldiers who began the campaign and the few who returned. Furthermore, the inclusion of so many data points does not mean that the presentation is

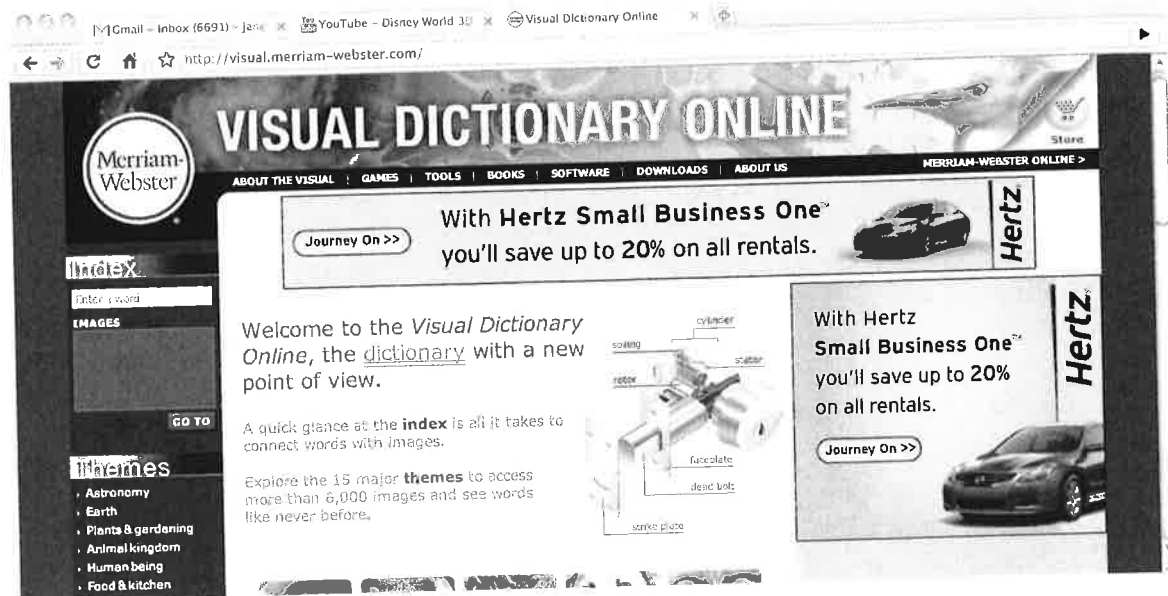


Figure 2.22

This site, like Amazon.com, is a three-column grid, but less well organized, with a mixture of digital elements and reproduced print elements. The navigation is confusing and the advertisements poorly separated from the editorial content. Stronger content is hidden at the bottom of the page, which visitors are unlikely to scroll to.

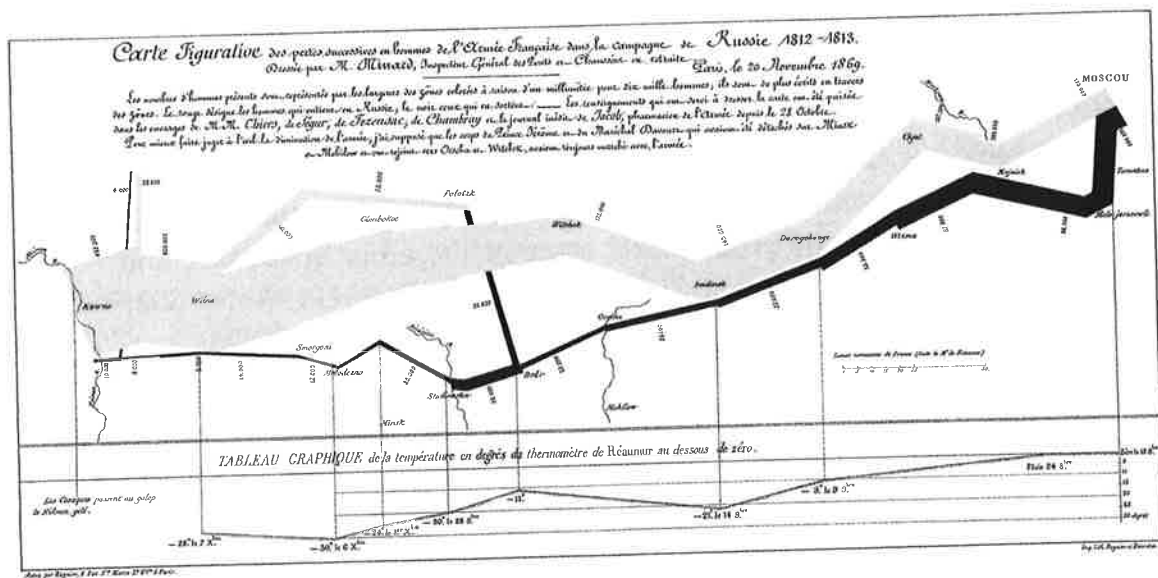


Figure 2.23

Charles Joseph Minard's 1869 map of Napoleon's invasion of and retreat from Russia, which Edward Tufte has called "probably the best statistical graph ever made." The brown lines were originally red, and represent the soldiers' progress from France to Russia; the black line represents the return. The thickness of the lines indicates the size of the force.

comprehensive. For example, Minard does not include statistics on the casualties among the Russians or the income level of Napoleon's soldiers, or the number of fatherless children they may have left behind.

The genre of statistical presentation is itself a reflection of cultural values characteristic of Minard's time, which saw a great growth of census taking and bureaucratic data gathering. Quantitative displays are important tools for understanding complex events, but they can prevent us from framing more open-ended questions. If we are to assess Minard's map as a comment upon the costs of war we might compare it to work in other genres, such as Tolstoy's *War and Peace* (1869) which takes a great deal more space to cover some of the same ground from the Russian perspective, or Picasso's *Guernica* (1937), an abstract image memorializing civilian victims of fascist bombing in the Spanish Civil War. Minard's graphic is a powerful representation of the event, but it is not necessarily more informative or objective than other media artifacts, though its use of numbers may give it the appearance of impartial and comprehensive authority.

Since the digital medium is more than a conveyer of quantitative information, aesthetic and stylistic strategies are an important part of the designer's craft. Strategies that would work well in a paper poster can result in static and self-referential design in a participatory environment, so designers must remember that no element can be beautiful if it interferes with **agency**. But many poster-based strategies, like the grid, can be adapted to digital formats. The key to good design is to integrate the aesthetic elements so that they serve the meaning of the artifact and support (rather than distract from) interaction (figure 2.24).

Graphic design in digital environments must always be in the service of interaction. Designers should therefore avoid purely decorative clutter, but they should also resist the seductions of extreme minimalism. The minimalist style adopted by graphic designers in the late twentieth century limited the number of colors and valued smaller fonts over larger ones. This aesthetic of understatement is an understandable reaction to the overly commercial, overly emphatic, visual shouting of the advertising and mass entertainment environments. The softer voice of the art poster or intellectual book jacket—a style soon co-opted by upscale advertising—draws attention to itself by its restraint. But minimalism can interfere with clear communication when the poster or printed page is so faint and the typeface is so small that the words cannot be read, and the problem is exacerbated in interactive environments where the user is looking for cues to how to behave or to what the computer is doing or how the information is organized. An interactive artifact may have great visual authority but will not be well made if the visual elements do not support the interactive function. A Macintosh may be delightful to look at but if the on/off switch is so elegantly hidden that infrequent users can't find it, then visual design values have been allowed to subvert acceptable interaction design.




Figure 2.24

Heart-shaped menu for the British Heart Foundation Annual Review site for 2008, a whimsical but readable alternative to a conventional menu or **tag cloud**. Mousing over any of the words causes related items to remain red while other choices fade, providing an organized overview of the contents of the site. Created by Sennep.

DESIGN EXPLORATIONS: EXPLORING THE AFFORDANCES OF THE DIGITAL MEDIUM

Designing Coherent Behaviors

- Game developers often paper-test a game before it is built to refine interaction and gameplay and to gain a sense of users' actions and potential frustrations. For this

exercise, you will form teams to do the reverse. Take an actual computer-based single player game (like Tetris) with which you are very familiar and create a paper version of it. Take turns being the player and the operator of the game. What kinds of actions and decisions does it take to run the game as the computer?

- Based on the previous exercise, reverse-engineer the algorithm for the game. What variables would you have to keep track of? What player actions would you have to anticipate? What conditions would you have to check for? Are there randomized processes involved? How should they be constrained? (E.g., Should it generate more of one kind of shape than another? Should it become less random if the player is closer to winning?)

Write out the rules of the game. For example, for Tic Tac Toe the rules might look like this:

```
TurnTaking is Player, Computer (randomized first)
First = X, Second = O
One mark per square
Three X or O horizontal, diagonal, vertical = win
```

After you write out the rules of your game, consider how you might change any of them to make the game harder or easier. For example, to make Tic Tac Toe easier, one could allow the player two turns in a row:

```
TurnTaking is (Player, Player), Computer (randomized first)
```

Use **pseudocode** to describe the step-by-step game play: Give names to all the variables and objects in the game and write a sequence of instructions that is in human-readable language similar to high-level computer code for the major game actions. Assume that there is one player per game and that the computer is running the game and playing the opponent.

The goal of writing pseudocode is to improvise a notation that lets you specify the logical steps without worrying about the syntax of any particular programming language. Use // to indicate that what follows is a comment rather than an instruction to the computer.

For example, pseudocode for Tic Tac Toe could be something like this:

```
Define BoardSquares ABC DEF GHI
Define Computer as O, layer as X
Random choice (0=Computer first, 1=Player first)
IF Computer first, place O in BoardSquare E // E=center square
ELSE Go to PlayerTurn
Loop until WIN:
PlayerTurn (X, SquareN)
  IF X (ABC) OR (DEF) or (GHI) or (AEI) or (CEG) or (ADG) or (BEH) or (CFI)
```

```

THEN WIN(Player)
ELSE MAKE_BEST_MOVE //blackboxing this for now
IF O (ABC) OR (DEF) or (GHI) or (AEI) or (CEG) or (ADG) or (BEH) or
(CFI)
THEN WIN (Computer)
ELSE LOOP AGAIN

```

I have simplified the system, assuming the computer would always play O and **black-boxing** the mathematical part of the game into an unspecified module called `MAKE_BEST_MOVE`.

Doing this exercise should make clear how tricky it is to specify procedures so that the computer—which can only do what we explicitly tell it to do—will not make mistakes that any six-year-old would avoid.

Scripting the Interactor

- Games often use expectations derived from other forms of entertainment or from highly gendered activities to script players' actions and expectations. Take a familiar game that is mostly played by members of one gender or by a particular subculture (such as teenage girls or boys) and identify the elements in the game that set up player expectations. How would these elements be different if they were aimed at encouraging participation by a different group of people?
- Create a digital scenario that scripts the interactor without explicit words. Put a visual element on a screen (or in some other interactive or sensor-driven artifact) that elicits an action from an interactor. Reward the action in a surprising way.
- Create a **storyboard** of the scenario (like a comic strip showing the sequence of events: before the interaction, the interaction, after the interaction) and then implement it if you are able to.

Example: Display a picture (or animation) of a group of birds and some breadcrumbs. The interactor will be “scripted” to pick up the breadcrumbs in order to try to feed the birds. But when the interactor mouses over the breadcrumbs, the birds will fly away.

Note: Extra credit for avoiding guns, swords, and other forms of overly familiar conventions of game violence.

- Create a digital scenario like the previous one, but this time, make three distinct items with three different effects. Create a reason for the interactor to try each of them. Remember you may not use words, so depict a scenario that will motivate interaction and items with strong dramatic or causal associations.

Example: A crying baby with a choice of a teddy bear, a bottle, and a rattle. The baby will only quiet down if the items are given in the appropriate order: first the

bottle to soothe hunger, then the rattle to tire himself out playing, then teddy bear to snuggle to sleep. Giving them in a different order should produce an escalation of unhappiness, possible destruction to the toys.

Without providing any explicit directions or offering any hints, script the interactor to behave appropriately to elicit meaningful behaviors from your scenario. Take turns playing with one another's scenarios while providing think-aloud commentary on your expectations, intentions, and reactions. Watch people playing with your environment without offering any comments or direction, and take notes on their self-reports. What worked best in your environment? What worked best for you in other people's environments?

Identifying Encyclopedic Design Elements and Issues

- Identify a digital encyclopedic resource and compare it to a similar resource in legacy media, such as in box 2.2.

Box 2.2

Encyclopedic Resources

Digital	Legacy
Search engine	Telephone book or card catalog
Online encyclopedia	Multivolume print encyclopedia
Online dictionary	Print dictionary
Electronic Programming guide	Newspaper TV listings
Electronic map	Paper map

What is the smallest unit of each? How is it standardized? How do you know what items are available? How are items chosen for inclusion? How reliable is the information? How is the resource updated? Which resource (digital or traditional) would you choose to use if both were available? (The answer may be different for different circumstances.) What use does the encyclopedic digital resource make of the procedural, participatory, and spatial affordances of the medium?

- The proliferation of encyclopedic resources has raised questions of freedom of access, criminal stalking, identify theft, and government and corporate surveillance of individuals. What countries are currently limiting their citizens' access to online information resources? How have large corporations responded to the demand to restrict access to information or to turn over private records (e.g., consider the behavior of involving Google and Yahoo in China)? How have librarians responded to demands for government access to Internet use? How have corporations sought to amass

encyclopedia data on individual purchasing through covert and overt means? How have information resources like YouTube and Twitter been used to oppose authority?

- Make a list of all the kinds of information that are produced by or about a particular individual over the course of their lifetime. Make a diagram by placing the image of a single person at the center of concentric circles. The circles represent information that is known about the person, with the innermost circles being the most secret (diary entries, love letters) and personal, and the outermost circles the most public (birth certificate, professional credentials, winning the lottery). Label each circle with who should gain access to the information, such as Best Friends, Immediate Family, Trusted Professionals, Accredited Institutions, General Public, and so on, and indicate what kinds of information should be known by/hidden from each group.

- How well does current information organization match your ideal categories? Where are the points of difference?

- Compare your diagram with those drawn by others. Are there common categories?
- Look at the kinds of information in the innermost circle of your diagram. Using the Internet—with the exception of illegal methods or sites—attempt to obtain this kind of information about yourself or other people. (Either use people whose permission you have obtained or use only voluntary self-revelations of public figures such as bloggers, celebrities, etc.) What conventions or standards make it harder or easier to obtain personal information? What are the trade-offs in designing a more secure environment?

Designing Coherent Spaces

- Select at least five pages of a large website, such as the site of an online department store, a university, or a government agency. Is there a common grid underlying the visual layout of all the pages? Can you draw it? How does this grid stay the same and how does it vary from the home page to pages farther down the hierarchy? How do these inner pages resemble or differ from one another in layout?

- Explore a virtual environment in a game or online community that provides a visual landscape that can be navigated. Choose two distinct places and describe all the ways that one can move between them. Is there a map of the domain? Can you navigate by clicking on the map? Can you fly over the landscape? Can you move rapidly on the ground? Is it pleasurable to move between places? Could you draw a picture of where the places are in relation to one another?

- Create a map for a confusing environment on or near campus that you know well, indicating landmarks as well as street names, and the best path from one familiar area to another, less-well-known area. For example, you might start at a central quad or public transit station and draw the route to a new coffee shop or specialized lab. Or you might start from your current classroom and find the nearest wheelchair-accessible bathroom or exit route.

Trade maps with another member of the class and try to follow the directions. Notice what helped or hindered your decoding of the map. How well did the diagram abstract the salient elements of the space? What did it leave out?

You can use a **think-aloud protocol** for this exercise in the following way:

- The person following the map describes what they are doing, what their guesses are about what the map means, what decisions they are making, and what their reasoning is about the route.
- The person who made the map should follow behind, remaining completely silent, giving no physical cues to the correct directions, and making no explanations about the written document. They should only record the behavior of the person they are following by taking notes or using recording equipment.
- Think-aloud protocols are powerful tools for designers because they make clear the gaps between what we think interactors will do and what they actually do. Take a confusing situation or information domain that you are already engaged with and are motivated to understand better. It could be a technical area, such as visualizing the human genome, or a recreational area, such as navigating television listings or choosing wines. Sketch a representation of the most important features of the subject you have chosen using three different visual frameworks such as:

- A table
- A **flow chart**
- An organizational chart for a hierarchy of people or other actors
- A taxonomy (also a hierarchical chart) of categories, such as the genus and species of living things
- A network
- A map of a geographical area
- A timeline
- A bubble diagram
- An arbitrarily chosen ordered arrangement, such as a subway system map from a particular city

Notice how the act of diagramming forces you to clarify qualities and relationships. Notice how the different approaches change your perspective on the domain. Explain the topic to others in your group, using your illustrative drawings, and noting which specific visual features enhance or detract from your explanation.

3 Maximizing the Four Affordances

All media as extensions of ourselves serve to provide new transforming vision and awareness.

—Marshall McLuhan (McLuhan [1964] 1995)

The Grid of Affordances

The digital designer has two responsibilities: to create the artifact that best serves the needs of the people who will interact with it, and to advance the digital medium as a whole. *One way to think about advancing the medium is to aim at maximizing all of its affordances.* As we have seen, anything made of computer code has the potential for encyclopedic capacity, for modeling navigable spaces, for describing and performing conditional behaviors, and for supporting participation by interactors. By exploiting these properties we increase the expressive power of the particular artifact, and we advance our ability to create other expressive artifacts.

The principle of maximizing the affordances of the medium does not mean that every artifact should include every possible application of digital technology. It does mean that the designer should imagine the many ways in which each of these affordances can be applied to the design task. Even though a particular artifact will only exploit some of these potential applications, it is still helpful to think of it as positioned within a more expansive design space in which all of these development paths are possible.

Let us imagine the approach of two designers, Charlie Competent and Victor Visionary. It is the mid-1990s, when the **World Wide Web** is just beginning, and Charlie and Victor are immersed in making websites for retail businesses moving toward the **Internet**. Charlie does not see beyond the immediate task. His goal is to **remediate** (translate from one medium to another) existing print artifacts like a sales catalog and print advertising for the web. He builds the site as a set of individual pages, cutting and pasting the content, and limiting the offering to selected items from the store. Given the era, this is received as an innovative project. Victor, however, sees the encyclopedic affordances of the web as offering the possibility of listing more items

than could be displayed in any print catalog. He understands that he can facilitate the entry, search, and display of inventory items by putting them in a database. He knows that existing 1990s technology can generate individual web pages off line and he anticipates that in time there will be increasingly powerful ways of generating web pages on the fly in response to interactors' queries. Victor's team creates a website that contains much more content than the one Charlie's team builds and that is better positioned for expansion as the web evolves. Eventually Charlie's clients have to redo their site, building up a new information structure from scratch.

Poor design is often limited to exploiting the novelty of superficial features like hyperlinks, blinking buttons, a presence on the web, social networking, or whatever is hot at the moment. For example, early producers of CD-ROMs relied solely on the capacity of the disk to bring value, flooding the market with products that came to be derided as "shovelware," in contempt for the lack of discrimination in their creation and the difficulty of navigating all the texts and images that they contained. Early **hypertext** fiction and educational applications included so much associational linking that interactors found them confusing and lost interest in engaging with them. But encyclopedic websites like the IMDB and Wikipedia have endured and grown, in part because they were built with an understanding that they would contain large quantities of information and therefore needed computationally structured back ends (like databases and wikis) and well-segmented and labeled site navigation tools in order to support massive participation and reach their full encyclopedic potential (figures 3.1, 3.2).

Successful applications are designed from the beginning with an awareness of the appropriate directions of growth, and an understanding of growth as part of the collective effort to maximize the four key affordances of the digital medium. We can think of the four affordances as a grid, offering **conventions** and **genres** that have become sufficiently standardized to be adopted across projects and **platforms** (figure 3.3). When we begin a new project it is useful to think of it as situated at the center of such a grid, and to ask ourselves which of the characteristic applications, such as those listed in the four quadrants displayed here, we could potentially apply to the design problem. We can accustom ourselves to thinking this way by analyzing existing applications in terms of how well they exploit each of these affordances.

The success of the original Google search engine is a good example of a sophisticated approach to the affordances of the medium. Google searches exceeded other first-generation web searches because they had a more processing-intensive **algorithm** that took on the challenge of assessing reliability and relevance. Other competing search engines used brute-force approaches, reading every word on a page and looking for matches to search terms. The Google engine analyzed the links between web pages, to capture information about who was pointing to a site and how they were describing it (in the words of the link). The Google team **reframed** the problem, thinking of the




Figure 3.1

Graph showing the growth in articles written for Wikipedia, between January 2001 and 2008. In its first year Wikipedia grew to 20,000 articles and by 2010 it had reached over 3 million articles in English and over 15 million articles in all languages.

web that they were searching not as a vast array of text strings but as a participatory social structure whose knowledge base could be mined by building social reasoning into the search procedures. As a result, Google is not just a more reliable search engine; it also has become a useful barometer of public opinion and a real-time index of public concerns.

If we were to place the original Google Search Engine on a grid representing the four affordances it would be strongest in the procedural quadrant (the algorithms of the search) and the encyclopedic quadrant (for its control over vast quantities of information), but also significantly present in the participatory quadrant (for its exploitation of user-created links and its later observation of which results users followed from similar search terms). But it clearly was not concerned with exploiting **spatial affordances**, keeping the home page free of clutter in an age when its competitor Yahoo was becoming an all-purpose portal. Google's only use of space in its original implementation was to organize the search results in a **list** with highest relevance at the top. By the end of the 2000s, however, it had expanded strongly into all four quadrants (figures 3.4, 3.5).

Of course no single artifact can exploit all the possibilities of the medium, and design can be quite successful in serving human needs by focusing on one or

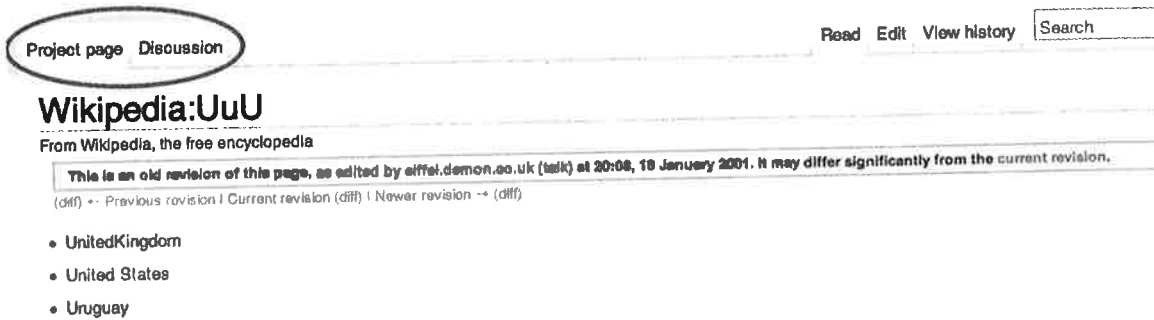


Figure 3.2

The earliest still-extant Wikipedia editing entry. Founded as an adjunct to a traditionally authored online encyclopedia project, Wikipedia is based on wiki software, which was designed explicitly for collaborative networked projects and includes a particularly useful architectural feature that separates discussion about a collectively edited page from the text of the page itself.

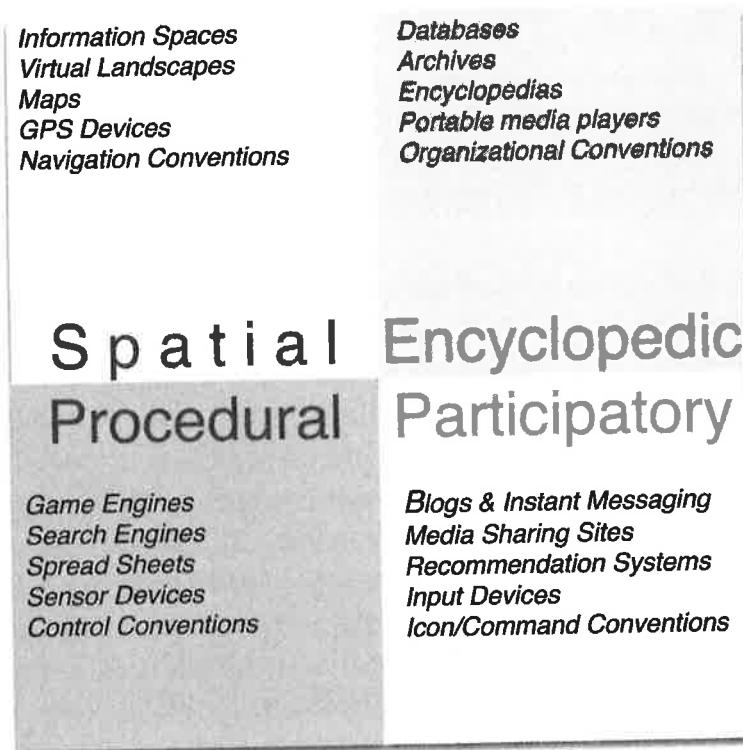


Figure 3.3

The four affordances of the digital medium with examples of standard formats and genres that exploit each of them.

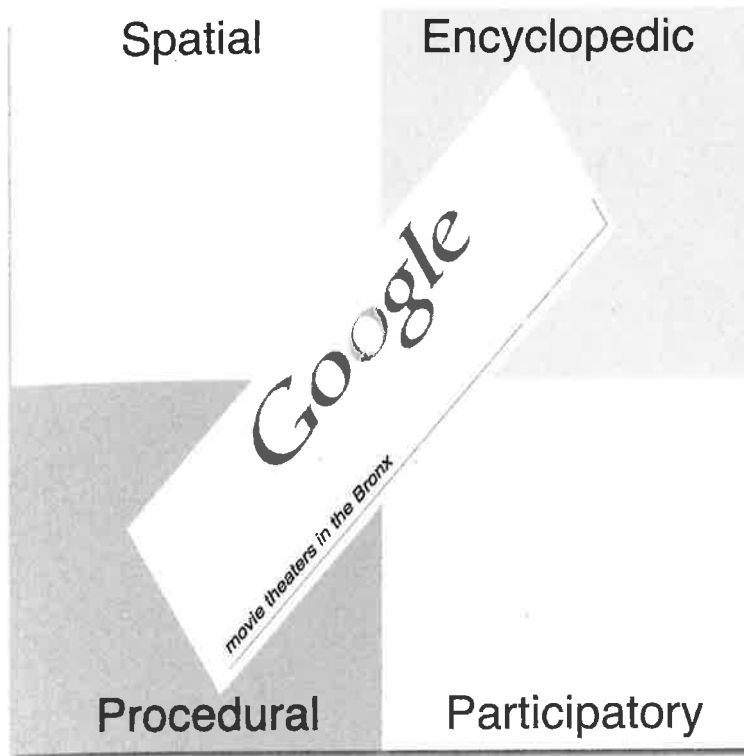


Figure 3.4

Google's original search engine mapped against the four affordances of the medium. It used little of the spatial affordance, presenting the results of searches as a rank-ordered list with the most relevant item at the top.

two affordances (see figures 3.6–3.8). But it is always useful to map an existing or proposed artifact against the larger design space in order to identify opportunities for growth and to predict the direction of media innovation. Affordance mapping can help to focus conversations about priorities. Perhaps the project is ignoring the possibility of spatial strategies. Perhaps it is overemphasizing participatory demands in the form of special features for individual users. Perhaps the major programming problems are coming from a particular procedural feature that is proving too hard to implement, but is not essential to the success of the project. The affordance grid can help to identify and clarify such design issues, and to organize thinking about paths not taken that may be worth revisiting on a future project. It can also suggest alternate ways of solving a problem, such as addressing a problem of user confusion (a participatory problem) with better information organization (an encyclopedic strategy) or a more orderly grid (a spatial approach). The affordance grid reminds us to contextualize design within the possibilities of the whole medium.

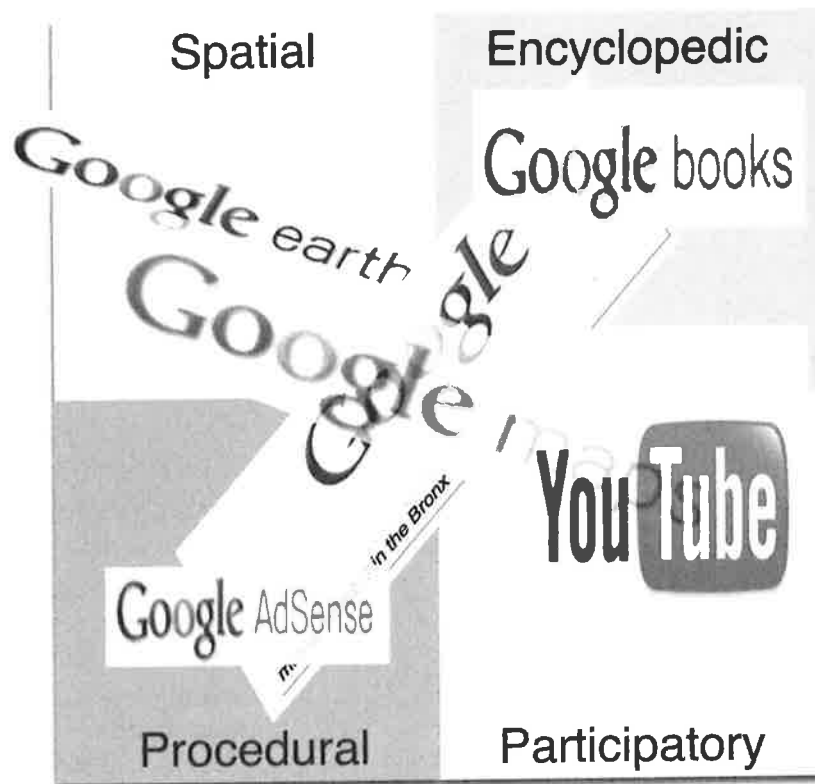


Figure 3.5

In the first decade of the twenty-first century, Google did not merely grow in size, it also built on its procedural strength to expand its use of the other three affordances of the medium. The foundational search engine expanded its coverage and its algorithms became more complex, including the participatory ad auctions system being built on advanced mathematical models; the Book Digitizing Project, Google News, and the acquisition of YouTube expanded its encyclopedic coverage. Gmail and YouTube expanded Google's exploitation of participatory structures; Google Maps and Google Earth established the standard for encyclopedic spatial exploration, from satellite image to road map to street view, while also providing for both procedural and participatory expansion with mash-ups.

One way to use the grid is to make it the focus of brainstorming at the beginning of a design process, placing an activity rather than an application at the center. For example, if we were making a new music player we might start with something like figure 3.9, thinking not of specific music-playing appliances but of the general activity of listening to music and asking how some of the existing digital formats and genres (like those listed in figure 3.3) might support that activity.

Alternately we can start not with the existing media structures but with the human activity of listening to music, asking ourselves in what ways it is spatial, encyclopedic, procedural, or participatory (figure 3.10).




Figure 3.6

The photo-sharing site Flickr.com was among the first successful networked media-sharing sites, relying on the affordance of vast amounts of storage, easily accessed and made available to others. It expanded its participatory affordances by partnering with social networking sites to syndicate content. Potential areas of growth would be to increase the procedural affordances, such as automated search and retrieval based on visual features or specialized tools for collages and photo editing; or to expand the encyclopedic affordances by including more kinds of media (such as moving images or sound files) or more structured annotation of images.

Or we could identify existing music applications and place them on the grid in order to get an overview of the design space (figure 3.11).

Finally, we could ask ourselves what we wish we could do in relation to this activity if we had a magic wand to make it possible. The magic wand concept is useful because it takes us away from the existing media conventions and existing behaviors, and puts us in the realm of the desired innovation. The ideas that surface in this part of the exercise can then be examined to see if they are implementable by exploiting the appropriate affordance.

Magic wand brainstorming can be based on questions like the following:

Procedural brainstorming What actions do I wish that [the proposed device or application] could magically do for me?




Figure 3.7

Wikipedia.org has established itself as a global information source, redefining the encyclopedia as a digital rather than print genre. It has created a coherent information space by segmenting knowledge into consistent categories and establishing standardized formats for entries, making it possible for thousands of contributors to coordinate their work and build upon one another's contributions.

Participatory brainstorming What actions do I want in my power that are not now in my power? What choices do I want to make explicitly? How do I want to perform the relevant task?

Encyclopedic brainstorming What is the extent of the information I want to command? What large resources do I want to control that I can't reach or can't organize with current technology? What kinds of media do I want to access, create, and share with this device?

Spatial brainstorming What kinds of mapping and visual organization of information would help me to use this device more fully? What kinds of information spaces do I want to enter or control? How do I think of physical space when using this device? How could the device help me to think more clearly about physical space? Should there be a geographical map and if so, at what granularity and with what information foregrounded?




Figure 3.8

Second Life, founded by Linden Lab in 2003, became the first successful shared virtual social space for a mass audience, with participatory opportunities for constructing, buying, and selling artifacts within the virtual world.

Designers should first ask these questions of themselves in fantasy thought experiments and then use them to design surveys and conduct interviews, focus groups, and market research with real consumers. It is important not to limit our own ideas or the brainstorming of research subjects to functions that already exist or that we already know how to create: that is why the concept of the magic wand is useful. It brings to the surface unmet needs and wished-for powers that can become the basis of innovative features and conventions. The goal in eliciting and analyzing this information is not to design a specific artifact that will have all of these magical qualities, but *to more fully understand the potential design space*. The exercise helps us to expand our expressive palette by reframing the task beyond the most familiar legacy and digital conventions that we might otherwise limit ourselves to. It gives us insight into the long-term direction of innovation, providing a context for prioritizing goals and making design choices throughout the development process.

Affordance mapping is also a useful way to synthesize the results of competitive analysis, which is used to survey the features of all the products in a design domain. In addition to creating the standard chart that lists which applications have which



Figure 3.9

Brainstorming for a new project using the affordance grid to ask how the core activity might be served by a wide range of existing digital formats and genres.

features, one could take all the features and map them against the affordance grid to get an overview of the whole field that will also make clear what areas remain uncovered in the design space.

Although the affordance grid can help us to anticipate next-generation features, designers must resist the temptation to extend functionality just because it is possible. Good design is shaped by **constraints** as well as affordances: constraints of time, resources, screen space, physical form, and especially human attention. Sometimes mature applications, like a popular word processor or simulation game, grow less usable with each new release because increased functionality gets ahead of interaction design. Even in a single development cycle design teams can easily fall into patterns of “feature creep,” adding functionality without first finishing a basic version of the artifact, delaying the project and confusing testers.

Shaping and Satisfying Interactors’ Expectations

Users do not think in terms of abstract affordances, but they expect new digital artifacts to have features similar to other digital artifacts, and even to have magical new

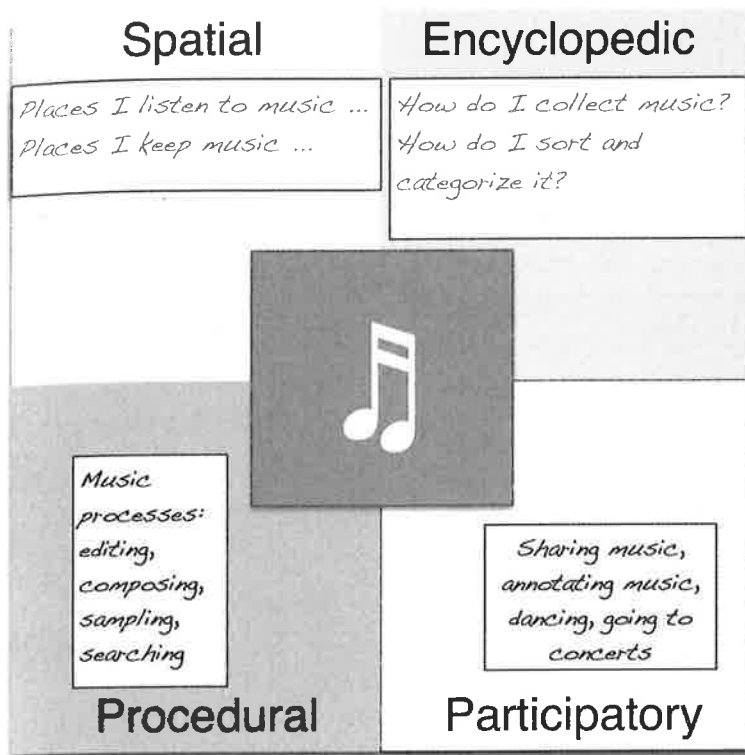


Figure 3.10

We can also use the affordance grid to brainstorm about the ways in which we perform the core cultural activity at the heart of a design process, independent of any particular device or media framework.

qualities based on the general kinds of things that they think computers are capable of. In order to avoid potential confusions, the designer should attempt to anticipate the interactor's expectations by imagining the interactor asking these four basic questions throughout the design process:

Procedural Question: What does it do?

Participatory Question: What can I do?

Spatial Question: Where am I in relation to the whole?

Encyclopedic Question: What are the boundaries to this domain?

The best way to answer such questions is to actively shape the expectations of the interactor by drawing on familiar genres and conventions that can focus the interactor's attention on productive actions. The more we can identify and create standardized conventions of interaction the less confused the interactor will be. Designers should be alert to their choice of conventions, avoiding the awkward and ad hoc solution, such as "quick links" on a poorly organized website or wordy instruction screens for an application that is not sufficiently self-explanatory.

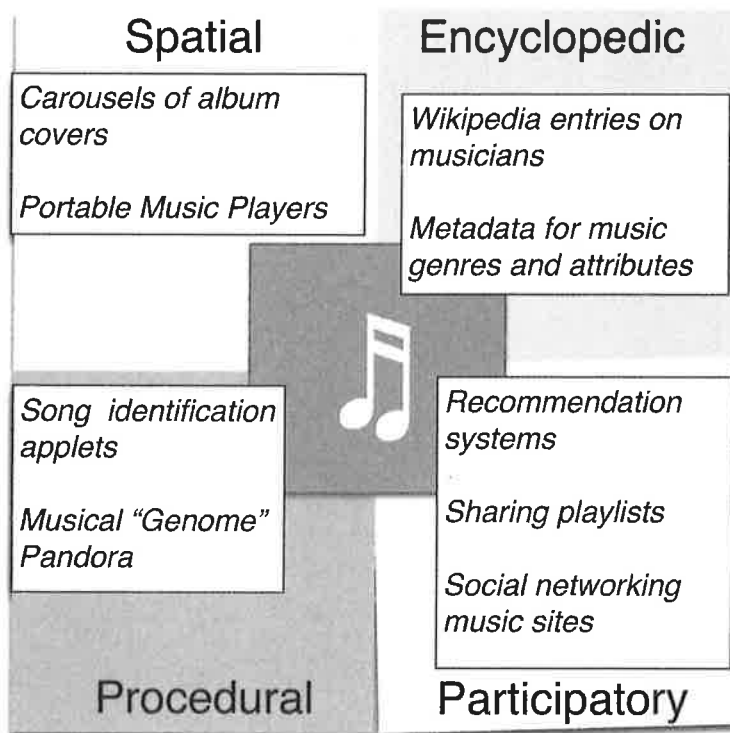


Figure 3.11

Using affordance mapping to generate an overview of existing applications.

The goal is to make the interaction easily learnable if it cannot be made immediately familiar.

Awkward strategies are sometimes unavoidable, but they should be recognized as symptomatic of an immature medium. As the medium matures, we collectively develop conventions to handle formerly awkward situations, often replacing indirect, multistep interaction patterns with **direct manipulation**. The examples in table 3.1 reflect the ongoing process of inventing and refining conventions that characterizes design in a new medium. Keeping in mind the four affordances can help us to focus that effort on specific areas within each project, looking for ways in which we can make the current application more coherent and thereby advance the expressiveness of the medium itself.

We can also tune up our awareness of the areas in which conventions are insufficient by noticing our own frustrations in digital environments and our own daydreams for magical functionalities. Whenever we use a digital resource or device that we have not used before, we can ask ourselves what our expectations are for each of the digital affordances and how well these expectations are met. For example:

Table 3.1

Answering the Four Questions with Immature and Mature Conventions

General User Question	Specific Question	Immature Solution	More Mature Convention	Why Better
What does it do?	What will happen if I click on this icon in a word processor?	Paper cheat sheet with icon labels and brief explanations.	Rollover text associated with the icon.	Information available in context and as needed.
What can I do?	How do I record a television program?	Consult newspaper listings to determine time and channel; enter time and channel using onscreen menu and remote, with help of manual.	Use remote to select program by name from onscreen menu.	Transmission codes (time and channel) eliminated, interactor's mental model simplified.
Where am I?	How do I find the movie listings in a newspaper site starting from the home page?	Provide "quick links" pull-down menu for frequently used resources buried in the site structure.	Put local listings with movie reviews under the Entertainment section of the site's hierarchical navigation system.	Logical use of space that reinforces the interactor's mental map of the categories of information covered by the site, as well as providing a memorable association for the location of the listings.
What are the boundaries?	How do I know which site I am on when I follow a link?	Links to pages inside and outside a website are displayed the same way and they all open in a single window.	Internal and external links are labeled and displayed differently, with external links opening in another window or tab.	Reinforces interactor's perception of a distinct, consistent site, part of but separate from the larger Web.

A new phone might

- not allow me to access all the features of an old address book (not encyclopedic enough);
- make it difficult to switch between dialing and looking at my calendar (failures of procedurality and participation);
- make me scroll through long lists or flip through multiple screens in order to find the right application (failure of spatializing information).

A new virtual community might

- provide poor tools for building assets (poor procedurality);
- make it hard to operate my avatar (poor participation);
- leave me confused about how to get from place to place (poor spatializing);
- provide too few choices of places to visit (not encyclopedic enough).

A new media player might not

- let me to open one application at a time (weaknesses in both procedurality and participation);
- play videos in a particular widely used format (not encyclopedic enough);
- make me scroll through screenfuls of applications to find the one I want (poor use of space).

Having the four categories in which to think about these common frustrations can help us to generalize our difficulties with particular devices so that we will recognize the same problem on other projects, and look for the media conventions that will provide solutions that can work across multiple applications.

Designing for Agency and Immersion

When we maximize the affordances of the medium in any one artifact we create satisfying experiences that it is useful to think of as **agency** and **immersion**.

As mentioned in the introduction, and as I proposed in chapter 3 of *Hamlet on the Holodeck* (1997), **interactivity** is rightly used to refer to the combination of procedural (computer processing) and participatory (human-directed) affordances of the digital medium. When these two affordances fit together well, they create the satisfying experience of agency. A highly interactive environment is one in which the machine puts a high degree of processing power into our hands and is highly responsive to our input. When the behavior of the computer is coherent and the application is designed so that a human interactor knows what to do and receives clear and immediate feedback on the result of their actions, the interactor experiences the pleasure of agency, of making something happen in a dynamically responsive world.




Figure 3.12

The Pac-Man arcade videogame (1980) offered the player multiple kinds of agency: moving Pac-Man, making the dots disappear, evading the ghosts, chasing the ghosts, navigating the maze, collecting fruit, completing the level.

Agency is therefore a more useful and specific design goal than interactivity. The designer creates agency by exploiting the interactive affordances of the medium, by shaping the rules by which the computer processes information and the conventions by which the interactor engages with the machine, so that human expectations and actions mesh well with computational expectations and actions. Agency is a concrete design goal, but it is not a property of an artifact. It is present or absent in the experience of the interactor (figure 3.12). Its opposite is powerlessness, confusion, or frustration (figure 3.13).

As with interactivity, there is a widespread sense that “new media” are immersive, but not much agreement on what **immersion** is or how to achieve it. Just as interactivity is often confused with a volume of activity, immersion is often confused with sensory stimulation. A high-resolution image on a very large screen or an enriched localized audio stream coming from a powerful speaker system may arrest our attention more fully and hold it longer than the same content delivered on a smaller screen and less powerful speakers, but we soon get used to increases in scale and intensity and learn to filter out stimulants that are merely very large or very loud. Immersion, like agency, cannot be a property of an artifact. Immersion is experiential and not a function of the size or volume of the stimulus, but of the engagement or distractibility of the interactor. An immersive environment is

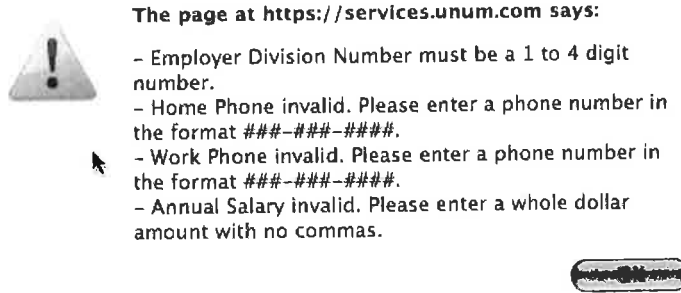


Figure 3.13

Frustration of agency is the mark of poor design. This error message is the result of an overly rigid data-entry form that rejects phone numbers separated by spaces instead of dashes. A better system would supply the dashes on input, automatically constrain input within three separate boxes, or accept multiple conventions for representing a ten-digit number.

one that captivates and holds our interest because it feels expansive, detailed, and complete.

In digital environments, the experience of immersion is reinforced by the encyclopedic and spatial affordances of the medium, by its enormous capacity and its openness to navigation in simulated landscapes or information spaces. We can also become immersed in old media, in compelling stories that hold our attention to the page or the image, in rhythmic experiences that focus us on music or movement, in puzzles or games that take over our thinking, causing us to lose awareness of the world around us. Immersion is related to the experience of being completely and pleurably absorbed in a challenging task, identified as **flow** (Csíkszentmihályi 1990), and it can be produced by a wide range of complex artifacts and challenging tasks. Digital environments can create all of these varieties of immersion, with the added sense of navigation of a landscape (like *Second Life*) or an information space (like Wikipedia) or a magically generative and renewing space of possibilities (like Tetris or the Spore creature editor) and a perception of the artifact as encyclopedic in its variety and extent.

Immersion is the experience produced by the pleasurable exploration of a limitless, consistent, familiar yet surprising environment. *As designers we create immersion by increasing scope, detail, and consistency while establishing clear boundaries and means of navigation.* These guidelines apply to information spaces such as Wikipedia and virtual worlds such as Spore or Second Life. Immersion is created by rewarding exploration with consistent, detailed variation, and undermined by poorly defined boundaries, gaps in coverage, and confusion in navigation.

Agency and immersion are mutually reinforcing. When we engage with an immersive world and it responds to us as we expect it to, revealing deeper levels of content, greater detail in its coverage, we become more deeply immersed. In the context of

interactive narrative I have called this effect the **active creation of belief**, to emphasize that digital immersion is created not merely from what Coleridge called “suspension of disbelief” but also from testing out the imaginary world and seeing how responsive it can be and how deeply it is described. Developing each of the four affordances increases the potential expressive power of the other three.

DESIGN EXPLORATIONS: EXPLORATIONS FOR MAXIMIZING DIGITAL AFFORDANCES

Using the Affordance Grid

- Make a grid of the four affordances of the medium and write down at least five specific examples (such as specific websites, databases, game levels, or devices) that make particularly good use of each of them. Justify your choices. Do your examples in the same quadrant have any common features or conventions? Are there some examples that are good illustrations of multiple affordances?
- For one of these artifacts you placed on the affordance grid, identify the most promising areas of unexploited potential outside its current boundaries (e.g., how can a strongly encyclopedic application better exploit the procedural power of the computer?).
- Find an example of a software application that is at least five years old and has gone through multiple releases (for example, Microsoft Word, Adobe Photoshop, Intuit TurboTax, Maxis's SimCity). Compare a screen shot (or video or running version if available) of an earlier version with the current version. What features have been added? Is it easier or harder to perform the original functions now that the program has been expanded? Is the application making appropriate use of the four affordances of the medium? Is it displaying feature creep (or perhaps feature stampede!)?
- Create a comparative analysis chart for competing digital products in the same market space, listing all the features of each in a table that makes clear what features are matched and which are competitive advantages. Now take the same products and map them on the affordance grid. Are there areas that none of them are exploiting? Can you suggest product extensions in these areas?
- For a device that you use regularly, do a magic wand exercise, brainstorming what other functions you can imagine it performing within the four quadrants of the affordance grid. Be as imaginative as you like, assuming anything you wish for will be created with no effort, out of magical technology. Think very concretely about how the magical functions would ideally work.

After detailed brainstorming, take a break and then return to your formulations. Are there existing or conceivable technologies that could be applied to implement some of the desired functions? Are there impediments of compatibility preventing the wish from becoming a reality? What kinds of expertise would have to be brought together to make the magic wand product a real product?