

Principles of New Media

By Lev Manovich

The identity of media has changed even more dramatically than that of the computer. Below I summarize some of the key differences between old and new media. In compiling this list of differences, I tried to arrange them in a logical order. That is, the last three principles are dependent on the first two. This is not dissimilar to axiomatic logic, in which certain axioms are taken as starting points and further theorems are proved on their basis.

Not every new media object obeys these principles. They should be considered not as absolute laws but rather as general tendencies of a culture undergoing-computerization. As computerization affects deeper and deeper layers of culture, these tendencies will increasingly manifest themselves.

1. Numerical Representation

All new media objects, whether created from scratch on computers or converted from analog media sources, are composed of digital code; they are numerical representations. This fact has two key consequences:

1. A new media object can be described formally (mathematically). For instance, an image or a shape can be described using a mathematical function.
2. A new media object is subject to algorithmic manipulation. For instance, by applying appropriate algorithms, we can automatically remove "noise" from a photograph, improve its contrast, locate the edges of the shapes, or change its proportions. In short, media becomes programmable.

When new media objects are created on computers, they originate in numerical form. But many new media objects are converted from various forms of old media. Although most readers understand the difference between analog and digital media, a few notes should be added on the terminology and the conversion process itself. This process assumes that data is originally continuous, that is, "the axis or dimension that is measured has no apparent indivisible unit from which it is composed." (7) Converting continuous data into a numerical representation is called digitization. Digitization consists of two steps: sampling and quantization. First, data is sampled, most often at regular intervals, such as the grid of pixels used to represent a digital image. The frequency of sampling is referred to as resolution. Sampling turns continuous data into discrete data, that is, data occurring in distinct units: people, the pages of a book, pixels. Second, each sample is quantified, that is, it is assigned a numerical value drawn from a defined range (such as 0-255 in the case of an 8-bit greyscale image)."

While some old media such as photography and sculpture are truly continuous, most involve the combination of continuous and discrete coding. One example is motion picture film: each frame is a continuous photograph, but time is broken into a number of samples (frames). Video goes one step further by sampling the frame along the vertical dimension (scan lines). Similarly, a photograph printed using a halftone process combines discrete and continuous representations. Such a photograph consists of a number of orderly dots (i.e., samples), although the diameters and areas of dots vary continuously.

As the last example demonstrates, while modern media contain levels of discrete representation, the samples are never quantified. This quantification of samples is the crucial step accomplished by digitization. But why, we may ask, are modern media technologies often in part discrete? The key assumption of modern semiotics is that communication requires discrete units. Without discrete units, there is no language. As Roland Barthes put it, "Language is, as it were, that which divides reality (for instance, the continuous spectrum of the colors is verbally reduced to a series of discontinuous terms) :'⁹ In assuming that any form of communication requires a discrete representation, semioticians took human language as the prototypical example of a communication system. A human language is discrete on most scales: We speak in sentences; a

sentence is made from words; a word consists of morphemes, and so on. If we follow this assumption, we may expect that media used in cultural communication will have discrete levels. At first this theory seems to work. Indeed, a film samples the continuous time of human existence into discrete frames; a drawing samples visible reality into discrete lines; and a printed photograph samples it into discrete dots. This assumption does not universally work, however: Photographs, for instance, do not have any apparent units. (Indeed, in the 1970s semiotics was criticized for its linguistic bias, and most semioticians came to recognize that a language-based model of distinct units of meaning cannot be applied to many kinds of cultural communication.) More important, the discrete units of modern media are usually not units of meanings in the way morphemes are. Neither film frames nor halftone dots have any relation to how a film or photograph affects the viewer (except in modern, art and avant-garde film--think of paintings by Roy Lichtenstein and films of Paul Sharits--which often make the "material" units of media into units of meaning).

The most likely reason modern media has discrete levels is because it emerged during the Industrial Revolution. In the nineteenth century, a new organization of production known as the factory system gradually replaced artisan labor. It reached its classical form when Henry Ford installed the first assembly line in his factory in 1913. The assembly line relied on two principles. The first was standardization of parts, already employed in the production of military uniforms in the nineteenth century. The second, newer principle was the separation of the production process into a set of simple, repetitive, and sequential activities that could be executed by workers who did not have to master the entire process and could be easily replaced.

Not surprisingly, modern media follows the logic of the factory, not only in terms of division of labor as witnessed in Hollywood film studios, animation studios, and television production, but also on the level of material organization. The invention of typesetting machines in the 1880s industrialized publishing while leading to a standardization of both type design and fonts (number and types). In the 1890s cinema combined automatically produced images (via photography) with a mechanical projector. This required standardization of both image dimensions (size, frame ratio, contrast) and temporal sampling rate. Even earlier, in the 1880s, the first television systems already involved standardization of sampling both in time and space. These modern media systems also followed factory logic in that, once a new "model"

(a film, a photograph, an audio recording) was introduced, numerous identical media copies would be produced from this master. As I will show, new media follows, or actually runs ahead of, a quite different logic of post-industrial society-that of individual customization, rather than mass standardization.

2. Modularity

This principle can be called the "fractal structure of new media." Just as a fractal has the same structure on different scales, a new media object has the same modular structure throughout. Media elements, be they images, sounds, shapes, or behaviors, are represented as collections of discrete samples (pixels, polygons, voxels, characters, scripts). These elements are assembled into larger-scale objects but continue to maintain their separate identities.

The objects themselves can be combined into even larger objects-again, without losing their independence. For example, a multimedia "movie" authored in popular Macromedia Director software may consist of hundreds of still images, QuickTime movies, and sounds that are stored separately and loaded at run time. Because all elements are stored independently, they can be modified at any time without having to change the Director "movie" itself. These "movies" can be assembled into a larger "movie," and so on. Another example of modularity is the concept of "object" used in Microsoft Office applications. When an "object" is inserted into a document (for instance, a media clip inserted into a Word document), it continues to maintain its independence and can always be edited with the program originally used to create it. Yet another example of modularity is the structure of an HTML document: With the exemption of text, it consists of a number of separate objects--GIF and JPEG images, media clips, Virtual Reality Modeling Language (VRML) scenes, Shockwave and Flash movies which are all stored independently, locally, and/or on a network. In short, a new media object consists of independent parts, each of which consists of smaller independent parts, and so on, down to the level of the smallest "atoms"-pixels, 3-D points, or text characters.

The World Wide Web as a whole is also completely modular. It consists of numerous Web pages, each in its turn consisting of separate media elements.

Every element can always be accessed on its own. Normally we think of elements as belonging to their corresponding Web sites, but this is just a convention, reinforced by commercial Web browsers. The Netomat browser by artist Maciej Wisniewski, which extracts elements of a particular media type from different Web pages (for instance, images only) and displays them together without identifying the Web sites from which they are drawn, highlights for us this fundamentally discrete and nonhierarchical organization of the Web.

In addition to using the metaphor of a fractal, we can also make an analogy between the modularity of new media and structured computer programming. Structural computer programming, which became standard in the 1970s, involves writing small and self-sufficient modules (called in different computer languages subroutines, functions, procedures, scripts), which are then assembled into larger programs. Many new media objects are in fact computer programs that follow structural programming style. For example, most interactive multimedia applications are written in Macromedia Director's Lingo. A Lingo program defines scripts that control various repeated actions, such as clicking on a button; these scripts are assembled into larger scripts. In the case of new media objects that are not computer programs, an analogy with structural programming still can be made because their parts can be accessed, modified, or substituted without affecting the overall structure of an object. This analogy, however, has its limits. If a particular module of a computer program is deleted, the program will not run. In contrast, as with traditional media, deleting parts of a new media object does not render it meaningless. In fact, the modular structure of new media makes such deletion and substitution of parts particularly easy. For example, since an HTML document consists of a number of separate objects each represented by a line of HTML code, it is very easy to delete, substitute, or add new objects. Similarly, since in Photoshop the parts of a digital image usually kept placed on separate layers, these parts can be deleted and substituted with a click of a button.

3. Automation

The numerical coding of media (principle 1) and the modular structure of a media object (principle 2) allow for the automation of many operations involved in media creation, manipulation, and access. Thus human intentionality can be removed from the creative process, at least in part. I"

Following are some examples of what can be called "low-level" automation of media creation, in which the computer user modifies or creates from scratch a media object using templates or simple algorithms. These techniques are robust enough so that they are included in most commercial software for image editing, 3-D graphics, word processing, graphics layout, and so forth.

Imageediting programs such as Photoshop can automatically correct scanned images, improving contrast range and removing noise. They also come with filters that can automatically modify an image, from creating simple variations of color to changing the whole image as though it were painted by Van Gogh, Seurat, or another brand-name artist. Other computer programs can automatically generate 3-D objects such as trees, landscapes, and human figures as well as detailed ready-to-use animations of complex natural phenomena such as fire and waterfalls. In Hollywood films, flocks of birds, ant colonies, and crowds of people are automatically created by AL (artificial life) software. Word processing, page layout, presentation, and Web creation programs come with "agents" that can automatically create the layout of a document. Writing software helps the user to create literary narratives using highly formalized genre conventions. Finally, in what may be the most familiar experience of automated media generation, many Web sites automatically generate Web pages on the fly when the user reaches the site. They assemble the information from databases and format it using generic templates and scripts.

Researchers are also working on what can be called "high-level" automation of media creation, which requires a computer to understand, to a certain degree, the meanings embedded in the objects being generated, that is, their

semantics. This research can be seen as part of a larger project of artificial intelligence (AI). As is well known, the AI project has achieved only limited

success since its beginnings in the 1950s. Correspondingly, work on media generation that requires an understanding of semantics is also in the research stage and is rarely included in commercial software. Beginning in the 1970s, computers were often used to generate poetry and fiction. In the 1990s, frequenters of Internet chat rooms became familiar with "bots"--computer programs that simulate human conversation. Researchers at New York University designed a "virtual theater" composed of a few "virtual actors" who adjusted their behavior in real-time in response to a user's actions. (11) The MIT Media Lab developed a number of different projects devoted to "high-level" automation of media creation and use: a "smart camera" that, when given a script, automatically follows the action and frames the shots; 12 ALIVE, a virtual environment where the user interacts with animated characters; (13) and a new kind of human-computer interface where the computer presents itself to a user as an animated talking character. The character, generated by a computer in real-time, communicates with the through user natural language; it also tries to guess the user's emotional state and to adjust the style of interaction accordingly."

The area of new media where the average computer user encountered AI in the 1990s was not, however, the human-computer interface, but computer games. Almost every commercial game included a component called an "AI engine," which stands for the part of the game's computer code that controls its characters-car drivers in a car race simulation, enemy forces in a strategy game such as Command and Conquer, single attackers in first-person shooters such as Quake. AI engines use a variety of approaches to simulate human intelligence, from rule-based systems to neural networks. Like AI expert systems, the characters in computer games have expertise in some well-defined but narrow area such as attacking the user. But because computer games are highly codified and rule-based, these characters function very effectively; that is, they effectively respond to the few things the user is allowed to ask them to do: run forward, shoot, pick up an object. They cannot do anything else, but then the game does not provide the opportunity for the user to test this. For instance, in a martial arts fighting game, I can't ask questions of my opponent, nor do I expect him or her to start a conversation with me. All I can do is "attack" my opponent by pressing a few buttons, and within this highly codified situation the computer can "fight" me back very effectively. In short, computer characters can display intelligence and skills only because programs place severe limits on our possible interactions with them. Put differently, computers can pretend to be intelligent only by tricking us into using a very small part of

who we are when we communicate with them. At the 1997 SIGGRAPH (Special Interest Group on Computer Graphics of the Association for Computing Machinery) convention, for example, I played against both human and computer-controlled characters in a VR simulation of a nonexistent sports game. All my opponents appeared as simple blobs covering a few pixels of my VR display; at this resolution, it made absolutely no difference who was human and who was not.

Along with "low-level" and "high-level" automation of media creation, another area of media use subjected to increasing automation is media access. The switch to computers as a means of storing and accessing enormous amounts of media material, exemplified by the "media assets" stored in the databases of stock agencies and global entertainment conglomerates, as well as public "media assets" distributed across numerous Web sites, created the need to find more efficient ways to classify and search media objects. Word processors and other text-management software has long provided the capacity to search for specific strings of text and automatically index documents. The UNIX operating system also included powerful commands to search and filter text files. In the 1990s software designers started to provide media users with similar abilities. Virage introduced Virage VIR Image Engine, which allows one to search for visually similar image content among millions of images as well as a set of video search tools to allow indexing and searching video files. (15) By the end of the 1990s, the key Web search engines already included the option to search the Internet by specific media such as images, video, and audio.

The Internet, which can be thought of as one huge distributed media database, also crystallized the basic condition of the new information society: overabundance of information of all kinds. One response was the popular idea of software "agents" designed to automate searching for relevant information. Some agents act as filters that deliver small amounts of information given the user's criteria. Others allow users to tap into the expertise of other users, following their selections and choices. For example, the MIT Software Agents Group developed such agents as BUZZwatch, which "distills and tracks trends, themes, and topics within collections of texts across time" such as Internet discussions and Web pages; Letizia, "a user interface agent that assists a user browsing the World Wide Web by . . . scouting ahead from the user's current

position to find Web pages of possible interest"; and Footprints, which "uses information left by other people to help you find your way around." (16)

By the end of the twentieth century, the problem was no longer how to create a new media object such as an image; the new problem was how to find an object that already exists somewhere. If you want a particular image, chances are it already exists-but it may be easier to create one from scratch than to find an existing one. Beginning in the nineteenth century, modern society developed technologies that automated media creation-the photo camera, film camera, tape recorder, videorecorder, etc. These technologies allowed us, over the course of 150 years, to accumulate an unprecedented amount of media materials-photo archives, film libraries, audio archives. This led to the next stage in media evolution-the need for new technologies to store, organize, and efficiently access these materials. The new technologies are all computer-based-media databases; hypermedia and other ways of organizing media material such as the hierarchical file system itself; text management software; programs for content-based search and retrieval. Thus automation of media access became the next logical stage of the process that had been put into motion when the first photograph was taken. The emergence of new media coincides with this second stage of a media society, now concerned as much with accessing and reusing existing media objects as with creating new ones.

4. Variability

A new media object is not something fixed once and for all, but something that can exist in different, potentially infinite versions. This is another consequence of the numerical coding of media (principle 1) and the modular structure of a media object (principle 2).

Old media involved a human creator who manually assembled textual, visual, and/or audio elements into a particular composition or sequence. This sequence was stored in some material, its order determined once and for all. Numerous copies could be run off from the master, and, in perfect correspondence with the logic of an industrial society, they were all identical. New media, in contrast, is characterized by variability. (Other terms that are often used in relation to new media and that might serve as appropriate synonyms of variable

are mutable and liquid.) Instead of identical copies, a new media object typically gives rise to many different versions. And rather than being created completely by a human author, these versions are often in part automatically assembled by a computer. (The example of Web pages automatically generated from databases using templates created by Web designers can be invoked here as well.) Thus the principle of variability is closely connected to automation.

Variability would also not be possible without modularity. Stored digitally, rather than in a fixed medium, media elements maintain their separate identities and can be assembled into numerous sequences under program control. In addition, because the elements themselves are broken into discrete samples (for instance, an image is represented as an array of pixels), they can be created and customized on the fly.

The logic of new media thus corresponds to the postindustrial logic of "production on demand" and "just in time" delivery logics that were themselves made possible by the use of computers and computer networks at all stages of manufacturing and distribution. Here, the "culture industry" (a term coined by Theodor Adorno in the 1930s) is actually ahead of most other industries. The idea that a customer might determine the exact features of her desired car at the showroom, transmit the specs to the factory, and hours later receive the car, remains a dream, but in the case of computer media, such immediacy is reality. Because the same machine is used as both showroom and factory, that is, the same computer generates and displays media-and because the media exists not as a material object but as data that can be sent through wires at the speed of light, the customized version created in response to the user's input is delivered almost immediately. Thus, to continue with the same example, when you access a Web site, the server immediately assembles a customized Web page.

Here are some particular cases of the variability principle (most of them will be discussed in more detail in later chapters):

1. Media elements are stored in a media database; a variety of end-user objects, which vary in resolution and in form and content, can be generated, either beforehand or on demand, from this database. At first, we might think that this is simply a particular technological implementation of the variability principle,

but, as I will show in the "Database" section, in a computer age the database comes to function as a cultural form in its own right. It offers a particular model of the world and of the human experience. It also affects how the user conceives the data it contains.

2. It becomes possible to separate the levels of "content" (data) and interface. A number of different interfaces can be created from the same data. A new media object can be defined as one or more interfaces to a multimedia database. I.,

3. Information about the user can be used by a computer program to customize automatically the media composition as well as to create elements themselves. Examples:

Web sites use information about the type of hardware and browser or user's network address to customize automatically the site the user will see; interactive computer installations use information about the user's body movements to generate sounds, shapes, and images, or to control the behavior of artificial creatures.

4. A particular case of this customization is branching-type interactivity (sometimes also called "menu-based interactivity"). The term refers to programs in which all the possible objects the user can visit form a branching tree structure. When the user reaches a particular object, the program presents her with choices and allows her to choose among them. Depending on the value chosen, the user advances along a particular branch of the tree. In this case the information used by a program is the output of the user's cognitive process, rather than the network address or body position.

5. Hypermedia is another popular new media structure, which is conceptually close to branching-type interactivity (because quite often the elements are connected using a branch tree structure). In hypermedia, the multimedia elements making a document are connected through hyperlinks. Thus the elements and the structure are independent of each other-rather than hard-wired together, as in traditional media. The World Wide Web is a particular

implementation of hypermedia in which the elements are distributed throughout the network. Hypertext is a particular case of hypermedia that uses only one media type-text. How does the principle of variability work in this case? We can think of all possible paths through a hypermedia document as being different versions of it. By following the links, the user retrieves a particular version of a document.

6. Another way in which different versions of the same media objects are commonly generated in computer culture is through periodic updates. For instance, modern software applications can periodically check for updates on the Internet and then download and install these updates, sometimes without any action on the part of the user. Most Web sites are also periodically updated either manually or automatically, when the data in the databases that drive the sites changes. A particularly interesting case of this "updateability" feature is those sites that continuously update information such as stock prices or weather.

7. One of the most basic cases of the variability principle is scalability, in which different versions of the same media object can be generated at various sizes or levels of detail. The metaphor of a map is useful in thinking about the scalability principle. If we equate a new media object with a physical territory, different versions of this object are like maps of this territory generated at different scales. Depending on the scale chosen, a map provides more or less detail about the territory. Indeed, different versions of a new media object may vary strictly quantitatively, that is, in the amount of de

tail present: For instance, a full-size image and its icon, automatically generated by Photoshop; a full text and its shorter version, generated by the "Autosummarize" command in Microsoft Word; or the different versions that can be created using the "Outline" command in Word. Beginning with version 3(1997), Apple's QuickTime format made it possible to embed a number of different versions that differ in size within a single QuickTime movie; when a Web user accesses the movie, a version is automatically selected depending on connection speed. A conceptually similar technique called "distancing" or "level of detail" is used in interactive virtual worlds such as VRML scenes. A designer creates a number of models of the same object, each with progressively less detail. When the virtual camera is close to the object, a highly detailed model is used; if the object is far away, a less detailed version is

automatically substituted by a program to save unnecessary computation of detail that cannot be seen anyway.

New media also allow us to create versions of the same object that differ from each other in more substantial ways. Here the comparison with maps of different scales no longer works. Examples of commands in commonly used software packages that allow the creation of such qualitatively different versions are "Variations" and "Adjustment layers" in Photoshop 5 and the "writing style" =option in Word's "Spelling and Grammar" command. More examples can be found on the Internet where, beginning in the mid-1990s, it became common to create a few different versions of a Web site. The user with a fast connection can choose a rich multimedia version, whereas the user with a slow connection can choose a more bare-bones version that loads faster.

Among new media artworks, David Blair's Wax Web, a Web site that is an "adaptation" of an hour-long video narrative, offers a more radical implementation of the scalability principle. While interacting with the narrative, the user can change the scale of representation at any point, going from an image-based outline of the movie to a complete script or a particular shot, or a VRML scene based on this shot, and so on. (19) Another example of how use of the scalability principle can create a dramatically new experience of an old media object is Stephen Mamber's database-driven representation of Hitchcock's *The Birds*. Mamber's software generates a still for every shot of the film; it then automatically combines all the stills into a rectangular matrix one shot per cell. As a result, time is spatialized, similar to the process in Edison's early Kinescope cylinders. Spatializing the film allows us to study its different temporal structures, which would be hard to observe otherwise. As in WaxWeb, the user can at any point change the scale of representation, going from a complete film to a particular shot.

As can be seen, the principle of variability is useful in allowing us to connect many important characteristics of new media that on first sight may appear unrelated. In particular, such popular new media structures as branching (or menu) interactivity and hypermedia can be seen as particular instances of the variability principle. In the case of branching interactivity, the user plays an active role in determining the order in which already generated elements are

accessed. This is the simplest kind of interactivity; more complex kinds are also possible in which both the elements and the structure of the whole object are either modified or generated on the fly in response to the user's interaction with a program. We can refer to such implementations as open interactivity to distinguish them from the closed interactivity that uses fixed elements arranged in a fixed branching structure. Open interactivity can be implemented using a variety of approaches, including procedural and object-oriented computer programming, AI, AL, and neural networks.

As long as there exists some kernel, some structure, some prototype that remains unchanged throughout the interaction, open interactivity can be thought of as a subset of the variability principle. Here a useful analogy can be made with Wittgenstein's theory of family resemblance, later developed into the theory of prototypes by cognitive psychologists. In a family, a number of relatives will share some features, although no single family member may possess all of the features. Similarly, according to the theory of prototypes, the meanings of many words in a natural language derive not through logical definition but through proximity to a certain prototype.

Hypermedia, the other popular structure of new media, can also be seen as a particular case of the more general principle of variability. According to the definition by Halasz and Schwartz, hypermedia systems "provide their users with the ability to create, manipulate and/or examine a network of information containing nodes interconnected by relational links." (20) Because in new media individual media elements (images, pages of text, etc.) always retain their individual identity (the principle of modularity), they can be "wired" together into more than one object. Hyperlinking is a particular way of achieving this wiring. A hyperlink creates a connection between two elements, for example, between two words in two different pages or a sentence on one page and an image in another, or two different places within the same page. Elements connected through hyperlinks can exist on the same computer or on different computers connected on a network, as in the case of the World Wide Web.

If in old media elements are "hardwired" into a unique structure and no longer maintain their separate identity, in hypermedia elements and structure are

separate from each other. The structure of hyperlinks--typically a branching tree--can be specified independently from the contents of a document. To make an analogy with the grammar of a natural language as described in Noam Chomsky's early linguistic theory, (21) we can compare a hypermedia structure that specifies connections between nodes with the deep structure of a sentence; a particular hypermedia text can then be compared with a particular sentence in a natural language. Another useful analogy is computer programming. In programming, there is clear separation between algorithms and data. An algorithm specifies the sequence of steps to be performed on any data, just as a hypermedia structure specifies a set of navigation paths (i.e., connections between nodes) that potentially can be applied to any set of media objects.

The principle of variability exemplifies how, historically, changes in media technologies are correlated with social change. If the logic of old media corresponded to the logic of industrial mass society, the logic of new media fits the logic of the postindustrial society, which values individuality over conformity. In industrial mass society everyone was supposed to enjoy the same goods--and to share the same beliefs. This was also the logic of media technology. A media object was assembled in a media factory (such as a Hollywood studio). Millions of identical copies were produced from a master and distributed to all the citizens. Broadcasting, cinema, and print media all followed this logic.

In a postindustrial society, every citizen can construct her own custom lifestyle and "select" her ideology from a large (but not infinite) number of choices. Rather than pushing the same objects/information to a mass audience, marketing now tries to target each individual separately. The logic of new media technology reflects this new social logic. Every visitor to a Web site automatically gets her own custom version of the site created on the fly from a database. The language of the text, the contents, the ads displayed--all these can be customized. According to a report in USA Today (9 November 1999), "Unlike ads in magazines or other real-world publications, 'banner' ads on Web pages change with every page view. And most of the companies that place the ads on the Web site track your movements across the Net, 'remembering' which ads you've seen, exactly when you saw them, whether you clicked on them, where you were at the time, and the site you have visited just before."

Every hypertext reader gets her own version of the complete text by selecting a particular path through it. Similarly, every user of an interactive installation gets her own version of the work. And so on. In this way new media technology acts as the most perfect realization of the utopia of an ideal society composed of unique individuals. New media objects assure users that their choices--and therefore, their underlying thoughts and desires--are unique, rather than preprogrammed and shared with others. As though trying to compensate for their earlier role in making us all the same, descendants of the jacquard loom, the Hollerith tabulator, and Zuse's cinemacomputer are now working to convince us that we are all unique.

The principle of variability as presented here has some parallels to the concept of "variable media," developed by the artist and curator Jon Ippolito. (23) I believe that we differ in two key respects. First, Ippolito uses variability to describe a characteristic shared by recent conceptual and some digital art, whereas I see variability as a basic condition of all new media, not only art. Second, Ippolito follows the tradition of conceptual art in which an artist can vary any dimension of the artwork, even its content; my use of the term aims to reflect the logic of mainstream culture in that versions of the object share some well-defined "data." This "data," which can be a well known narrative (Psycho), an icon (Coca-Cola sign), a character (Mickey Mouse), or a famous star (Madonna), is referred to in the media industry as "property." Thus all cultural projects produced by Madonna will be automatically united by her name. Using the theory of prototypes, we can say that the property acts as a prototype, and different versions are derived from this prototype. Moreover, when a number of versions are being commercially released based on some "property," usually one of these versions is treated as the source of the "data," with others positioned as being derived from this source. Typically, the version that is in the same media as the original "property" is treated as the source. For instance, when a movie studio releases a new film, along with a computer game based on it, product tie-ins, music written for the movie, etc., the film is usually presented as the "base" object from which other objects are derived. So when George Lucas releases a new Star Wars movie, the original property--the original Star Wars trilogy--is referenced. The new movie becomes the "base" object, and all other media objects released along with it refer to this object. Conversely, when computer games such as Tomb Raider are remade into movies, the original computer game is presented as the "base" object.

Although I deduce the principle of variability from more basic principles of new media-numerical representation and modularity of information--the principle can also be seen as a consequence of the computer's way of representing data-and modeling the world itself-as variables rather than constants. As new media theorist and architect Marcos Novak notes, a computer--and computer culture in its wake--substitutes every constant with a variable. (24) In designing all functions and data structures, a computer programmer tries always to use variables rather than constants. On the level of the human-computer interface, this principle means that the user is given many options to modify the performance of a program or a media object, be it a computer game, Web site, Web browser, or the operating system itself. The user can change the profile of a game character, modify how folders appear on the desktop, how files are displayed, what icons are used, and so forth. If we apply this principle to culture at large, it would mean that every choice responsible for giving a cultural object a unique identity can potentially remain always open. Size, degree of detail, format, color, shape, interactive trajectory, trajectory through space, duration, rhythm, point of view, the presence or absence of particular characters, the development of plot-to name just a few dimensions of cultural objects in different media-can all be defined as variables, to be freely modified by a user.

Do we want, or need, such freedom? As the pioneer of interactive filmmaking Grahame Weinbren argues, in relation to interactive media, making a choice involves a moral responsibility. (25) By passing on these choices to the user, the author also passes on the responsibility to represent the world and the human condition in it. (A parallel is the use of phone or Web-based automated menu systems by big companies to handle their customers; while companies have turned to such systems in the name of "choice" and "freedom," one of the effects of this type of automation is that labor is passed from the company's employees to the customer. If before a customer would get the information or buy the product by interacting with a company employee, now she has to spend her own time and energy navigating through numerous menus to accomplish the same result.) The moral anxiety that accompanies the shift from constants (to variables, from traditions to choices in all areas of life in a contemporary society, and the corresponding anxiety of a writer who has to portray it, is well rendered in the closing passage of a short story by the contemporary American writer Rick Moody (the story is about the death of his sister) (26)

I should fictionalize it more, I should conceal myself. I should consider the responsibilities of characterization, I should conflate her two children into one, or reverse their genders, or otherwise alter them, I should make her boyfriend a husband, I should explicate all the tributaries of my extended family (its remarriages, its internecine politics), I should novelize the whole thing, I should make it multigenerational, I should work in my forefathers (stonemasons and newspapermen), I should let artifice create an elegant surface, I should make the events orderly, I should wait and write about it later, I should wait until I'm not angry, I shouldn't clutter a narrative with fragments, with mere recollections of good times, or with regrets, I should make Meredith's death shapely and persuasive, not blunt and disjunctive, I shouldn't have to think the unthinkable, I shouldn't have to suffer, I should address her here directly (these are the ways I miss you), I should write only of affection, I should make our travels in this earthly landscape safe and secure, I should have a better ending, I shouldn't say her life was short and often sad, I shouldn't say she had demons, as I do too.

5. Transcoding

Beginning with the basic, "material" principles of new media-numeric coding and modular organization-we moved to more "deep" and far reaching ones-automation and variability. The fifth and last principle of cultural transcoding aims to describe what in my view is the most substantial consequence of the computerization of media. As I have suggested, computerization turns media into computer data. While from one point of view, computerized media still displays structural organization that makes sense to its human users-images feature recognizable objects; text files consist of grammatical sentences; virtual spaces are defined along the familiar Cartesian coordinate system; and so on-from another point of view, its structure now follows the established conventions of the computer's organization of data. Examples of these conventions are different data structures such as lists, records, and arrays; the already-mentioned substitution of all constants by variables; the separation between algorithms and data structures; and modularity.

The structure of a computer image is a case in point. On the level of representation, it belongs on the side of human culture, automatically entering in dialog with other images, other cultural "semes" and "mythemes." But on another level, it is a computer file that consists of a machine-readable header, followed by numbers representing color values of its pixels. On this level it

enters into a dialog with other computer files. The dimensions of this dialog are not the image's content, meanings, or formal qualities, but rather file size, file type, type of compression used, file format, and so on. In short, these dimensions belong to the computer's own cosmogony rather than to human culture.

Similarly, new media in general can be thought of as consisting of two distinct layers-the "cultural layer" and the "computer layer." Examples of categories belonging to the cultural layer are the encyclopedia and the short story; story and plot; composition and point of view; mimesis and catharsis, comedy and tragedy. Examples of categories in the computer layer are process and packet (as in data packets transmitted through the network); sorting and matching; function and variable; computer language and data structure.

Because new media is created on computers, distributed via computers, and stored and archived on computers, the logic of a computer can be expected to significantly influence the traditional cultural logic of media; that is, we may expect that the computer layer will affect the cultural layer. The ways in which the computer models the world, represents data, and allows us to operate on it; the key operations behind all computer programs (such as search, match, sort, and filter); the conventions of HCI--in short, what can be called the computer's ontology, epistemology, and pragmatics--influence the cultural layer of new media, its organization, its emerging genres, its contents.

Of course, what I call "the computer layer" is not itself fixed but rather changes over time. As hardware and software keep evolving and as the computer is used for new tasks and in new ways, this layer undergoes continuous transformation. The new use of the computer as a media machine is a case in point. This use is having an effect on the computer's hardware and software, especially on the level of the human-computer interface, which increasingly resembles the interfaces of older media machines and cultural technologies-VCR, tape player, photo camera.

In summary, the computer layer and the culture layer influence each other. To use another concept from new media, we can say that they are being composited together. The result of this composite is a new computer culture—a blend of human and computer meanings, of traditional ways in which human culture modeled the world and the computer's own means of representing it.

Throughout the book, we will encounter many examples of the principle of transcoding at work. For instance, in "The Language of Cultural Inter

faces," we will look at how conventions of the printed page, cinema, and traditional HCI interact in the interfaces of Web sites, CD-ROMs, virtual spaces, and computer games. The "Database" section will discuss how a database, originally a computer technology to organize and access data, is becoming a new cultural form in its own right. But we can also reinterpret some of the principles of new media already discussed as consequences of the transcoding principle. For instance, hypermedia can be understood as one cultural effect of the separation between an algorithm and a data structure, essential to computer programming. Just as in programming, where algorithms and data structures exist independently of each other, in hypermedia data is separated from the navigation structure. Similarly, the modular structure of new media can be seen as an effect of the modularity in structural computer programming. Just as a structural computer program consists of smaller modules that in turn consist of even smaller modules, a new media object has a modular structure.

In new media lingo, to "transcode" something is to translate it into another format. The computerization of culture gradually accomplishes similar transcoding in relation to all cultural categories and concepts. That is, cultural categories and concepts are substituted, on the level of meaning and/or language, by new ones that derive from the computer's ontology, epistemology; and pragmatics. New media thus acts as a forerunner of this more general process of cultural reconceptualization.

Given the process of "conceptual transfer" from the computer world to culture at large, and given the new status of media as computer data, what theoretical framework can we use to understand it? On one level new media is old media

that has been digitized, so it seems appropriate to look at new media using the perspective of media studies. We may compare new media and old media such as print, photography, or television. We may also ask about the conditions of distribution and reception and patterns of use. We may also ask about similarities and differences in the material properties of each medium and how these affect their aesthetic possibilities.

This perspective is important and I am using it frequently in this book, but it is not sufficient. It cannot address the most fundamental quality of new media that has no historical precedent-programmability. Comparing new media to print, photography, or television will never tell us the whole story. For although from one point of view new media is indeed another type of media, from another it is simply a particular type of computer data, something stored in files and databases, retrieved and sorted, run through algorithms and written to the output device. That the data represent pixels and that this device happens to be an output screen is beside the point. The computer may perform perfectly the role of the jacquard loom, but underneath it is fundamentally Babbage's Analytical Engine-after all, this was its identity for 150 years. New media may look like media, but this is only the surface.

New media calls for a new stage in media theory whose beginnings can be traced back to the revolutionary works of Harold Innis in the 1950s and Marshall McLuhan in the 1960s. To understand the logic of new media, we need to turn to computer science. It is there that we may expect to find the new terms, categories, and operations that characterize media that became programmable. From media studies, we move to something that can be called "software studies" from media theory to software theory. The principle of transcoding is one way to start thinking about software theory. Another way, which this book experiments with, is to use concepts from computer science as categories of new media theory. Examples here are "interface" and "database." And last but not least, along with analyzing "material" and logical principles of computer hardware and software, we can also look at the human-computer interface and the interfaces of software applications used to author and access new media objects. The two chapters that follow are devoted to these topics.