

EventWeb: Developing a Human-Centered Computing System

Ramesh Jain, University of California, Irvine

Dealing with emerging applications of computing requires taking a fresh look at our tools.

EventWeb is a human-centered computing system that will give users a compelling experience by combining quality content, carefully planned data organization and access mechanisms, and powerful presentation approaches.

In a sense, all computing is human centered. Once humans successfully developed machines to augment their mechanical strength, they focused on developing machinery to augment their analytical abilities. The first analytical ability they addressed was simple mathematical calculations. This resulted in different types of calculators, leading finally to electronic computers in the middle of the past century.

Augmentation of human analytical facilities remains the primary goal of computers today. Computing has evolved rapidly, with advances in processing, storage, communication, sensing, and related areas. Another driving force is the emergence of increasingly challenging sensory-data problems, including video and audio, that computers could solve.

Current interest in human-centered computing suggests new winds blowing in the computing community. HCC combines many powerful and independent approaches in different aspects of computing, ranging from human-computer interfaces (HCIs), computer vision, speech recognition, and pervasive computing to virtual reality systems. Most computing systems allow creation of powerful interfaces using audiovisual techniques. However, designing HCC systems that give users a compelling experience requires quality content, carefully planned data organization and access mechanisms, and powerful presentation approaches.

Content is ultimately what's interesting to users, so its quality is important. Content quality includes its credibility, depth, and timeliness. The challenges that HCC researchers face aren't limited to HCI or gestural

approaches, but go deeper, into correctly organizing multimodal data from disparate sources, finding the best combination of multimedia sources to communicate the message or experience, presenting and distributing these sources for the best subjective quality of experience, and helping advance human knowledge and build stronger communities using these approaches.

Most current approaches in computer science and applications evolved from alphanumerical data—the dominant data in the early days of computing. Researchers have tried to extend these approaches to increasingly multimodal dynamic data. To deal with emerging applications of computing, ranging from biology to entertainment, and from security to business, we need to take a fresh look at our tools.

A combination of technological advances, a reduction in barriers to interactions among different parts of the world, and the quest for solving increasingly difficult problems has created a situation that's unique in its potential to impact the course of human civilization.

DISRUPTIVE INNOVATIONS IN COMPUTING

Computing has already gone through two major disruptive evolutionary stages, and is now on the verge of a third. Although the first two were revolutionary, the third stage will effect the most long-term fundamental changes in how computing influences human civilization.

Table 1 shows the basic features and characteristics of the three stages of this evolution. In addition to the nature of input and output, the nature of applications and user expectations has changed. Also, each phase

Table 1. Evolutionary stages in computing.

Feature	Stage 1	Stage 2	Stage 3
Input	Data	Documents	Multimodal data
Output	Results	Information	Experiences
Processing	Computations	Information extraction	Experience processing
Devices	Mainframe, workstations	Personal computers, Internet	Mobile phones
Applications	Computing	Information and communication	Insights and entertainment
User level	Trained	General population in the developed world	All humans
Interaction environment	Command languages	Windows	Experimental

builds on the previous phase's advances, so it subsumes and even enhances the previous phase's functionality.

Data and computation

The invention of electronic computers marked the first major event in computing. These computers processed data at an unimaginable speed and could perform calculations millions of times faster than humans. The ability to program computers opened avenues for many previously unimaginable applications.

This initial evolutionary phase focused on scientific and engineering computations. Soon, the business community and other organizations realized the potential of computing with alphanumeric data. This extended computing to businesses and large organizations. Early mainframe computers and workstations represented this style of computation. The terms "computer" and "computing" are the legacy of this phase.

Information and communication

The second major evolutionary stage brought personal computers, including laptops, and the Internet. In this phase, the emphasis shifted from data to information and communication technology (ICT). The rise of PCs and what-you-see-is-what-you-get (WYSIWYG) word-processing and spreadsheet programs brought computing from trained computer operators to the general population in the developed world.

The Internet and World Wide Web accelerated the ICT revolution. By networking computers to form a global pervasive network, users could connect alphanumeric data sources, including documents, and communicate data and information. The ICT revolution started unimaginable applications and affected human life in most developed and developing parts of the world.

Insights and experiences

The easy availability of multimodal sensory data and devices that can capture, play, store, and process this data is propelling the third major evolutionary stage. This phase will bring insights and experience to the forefront in the same way that the second phase focused on information and communication.

Insights refer to the perception of the true nature of things made possible when a human understands the relationships among a thing's different components. Insights help in deeper analysis and problem solving. We need insights to make decisions, and information to implement these decisions. Insights are closely correlated to experiences. Dictionary.com defines "experience" as the "active participation in events or activities leading to the accumulation of knowledge or skill" or as "knowledge or wisdom gained from what one has observed, encountered, or undergone."

People experience events and activities using their sensory tools: sight, sound, touch, smell, and taste. These observations constitute experience of the event resulting in insights related to the event and objects in the event. In the digital world, approaches that use sensing technology convert sight, sound, and touch to electronic form and then, after processing, convert it back to sensory data for human perception. Smell and taste remain relatively difficult to capture, convert to digital form, and then convert back to the original senses for presenting to humans.

It's common, however, to experience the sights and sounds of remote events, and touch is finding increasing use. Fortunately, people experience the world around them and create knowledge about their environment mainly through sight and sound. Thus, we can naturally extend digital experience to these dominant human senses.

This third stage is about experiencing events, saving experiences, gaining insights and knowledge from these experiences, and sharing these insights and experiences with others. It's no wonder that the past few years have seen the emergence of companies like Flickr, YouTube, and Facebook because these companies provide environments in which users can share experiences through photos, videos, and multimedia. Progress made in this stage could lead to revolutionary approaches for expanding access to the more than 80 percent of the Earth's population that computing has yet to reach.

APPLICATION ENVIRONMENTS

Given the emerging nature of data and computing and projecting emerging applications, several applica-

tion areas seem to be the natural extension of current popular computing application paradigms.

Immersive telepresence

Virtual reality systems use computing as a storytelling mechanism, with the user an integral part of the story. In these systems, the user is immersed in an environment. Developers are making immersive environments increasingly realistic by combining computer graphics technology with image processing, audio, and even tactile information. Videogames, for example, routinely use tactile processing in the form of force feedback and similar mechanisms.

In an immersive system,¹⁻³ a rendering engine creates a detailed model of the environment in response to user actions. On detecting a user action, the engine renders an appropriate segment of the model to keep the user immersed in the environment. The system's realism depends on its ability to render the model and adapt the environment to the user's actions.

All these actions occur in an interactive setting. The system must have less latency than a person's perceptual limits. These systems usually combine a predefined environment model with known user behaviors to generate the situation model synthetically.

Suppose we modify such a system. First, we assume that the model is from a real, rather than synthetic, environment. Let's assume that we're in the US and the environment model is a football game taking place in India. Let's also assume that we've placed as many cameras and microphones as required for creating a realistic model of the sights and sounds of this event—including what's happening on the sidelines and what spectators are doing. And suppose this dynamic model is properly indexed using events. We'll call this a situation model to differentiate it from the static environment model. (A simpler version of this scenario is a live telecast of the game.)

But now let's assume that we can request what we want to see and from where on the field, and the system will render exactly that view using the model that exists at each moment. Effectively, we could be at any location on or off the field at any point in the game. Of course, in this scenario, we can't act, except to change our position. But everybody else—potentially millions of people—can enjoy the game from their own perspective and, if they desire, share their view with others.

Figure 1 shows this architecture. On the left side, different sensors and data sources feed the system through arrow A. Using the environment model, the system assimilates and indexes this information. It creates a situation model to represent the physical world of interest at every time moment. It also saves and indexes all sensory data, perhaps for later use. Arrow C is the inter-

action with different remote users. Each user requests to experience a different segment of the situation model, and the system generates that model using information from different sensor and data sources.

This is an immersive telepresence experience, based on reality that is “real” because the system's environment model comes from the real, rather than virtual, world.⁴ Moreover, the situation model is prepared at every instant to represent the real world as captured using relevant sensors. This makes it different from the systems that use these models from artificial worlds. This model could prove equally effective in entertainment, videoconferencing, telemedicine, personal communications, scientific explorations, and education.

Information assimilation. Most current multimedia information systems deal with archived video, audio, and images. Continuous queries—that is, persistent queries that are issued once and then logically run continuously over live and unbounded streams, have recently become a major research area in data management. A media stream comes from a sensor device such as a video, audio, or motion sensor and produces a continuous or discrete signal, but typically a data-stream processor can't directly use it.

To evaluate queries on media streams, a system needs to extract features continuously and assimilate them to form domain events.

When we place multiple sensors in a physical environment, their placement and the models of the events they're supposed to capture play important roles in information extraction and assimilation. My research focuses on knowledge about the location of sensors in a physical space and the role of spatially and temporally correlated information obtained from disparate sensors.^{2,5,6} An environment model captures the physical placements and constraints on the information obtained from these sensors. My research in multimedia stream queries addresses issues related to live multimedia data.

In an immersive telepresence system, multiple sensors placed in the physical environment capture an event. In most cases, a significant amount of metadata associated with the environment, sensors, and event is available. The system should continuously process and assimilate this data to form a unified situation model. The unified model provides all information about the events and objects in the environment. In assimilating the data from sensors and other sources, the system considers each source as an observation source that contributes to the complete model.

Semantic indexing. Databases and search engines have traditionally used indexing to efficiently store and retrieve data. In addition to dealing with indexing at that level, an immersive telepresence system must deal

In an immersive telepresence system, multiple sensors placed in the physical environment capture an event.

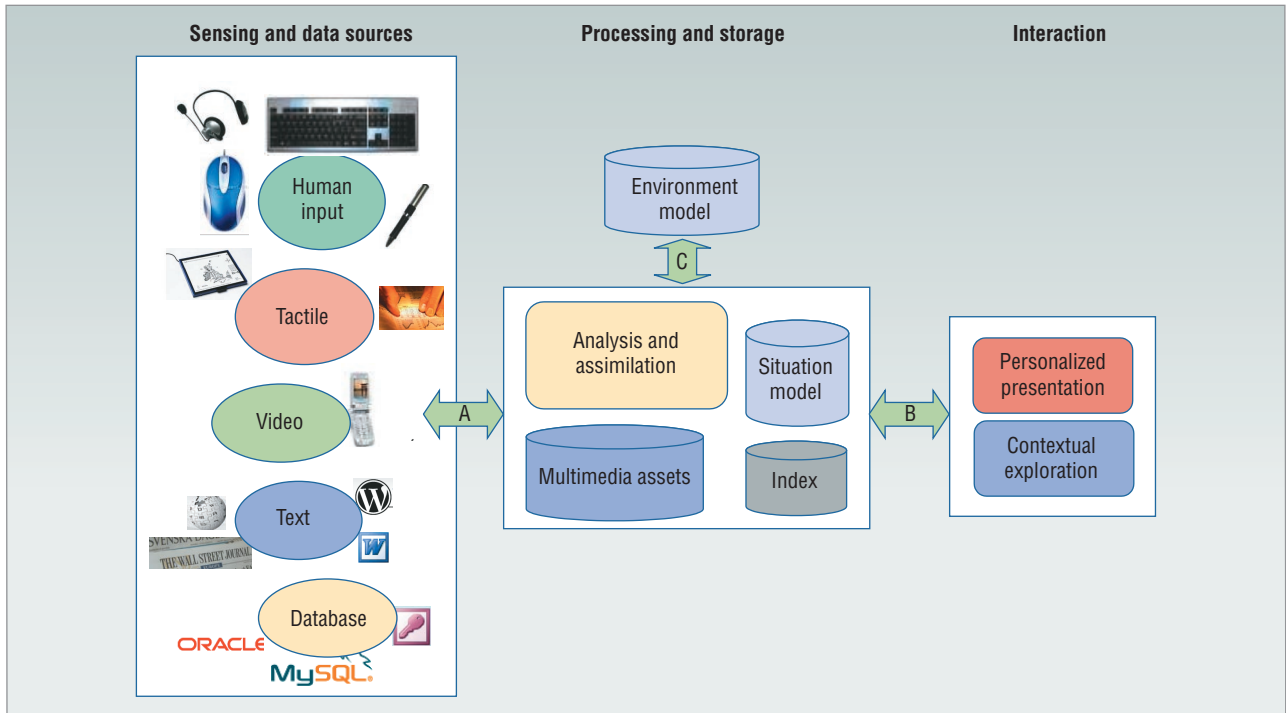


Figure 1. Immersive telepresence system architecture. On the left side, different sensors and data sources feed the system through arrow A. Arrow B is the interaction with different remote users.

with indexing at the semantic level. Current indexing techniques for different data types depend on metadata for that particular type. Metadata plays a key role in introducing semantics and determining how to use data. Schemas provide semantics in relational tables. XML is increasingly used to introduce semantics in strongly human-mediated environments.

Current information systems create data silos. They define and introduce the metadata for a particular data type, then index it and stash it neatly in a silo. Breaking down these silos can unify information. I proposed a unifying indexing system that introduces a layer on top of each data silo's metadata layer, or disparate data source.^{7,8} The layer uses an event-based domain model and metadata to construct a new index that's independent of data type. This system can model an application domain in terms of events and objects.

An event ontology parses the data as it comes from the sensors and data sources and assimilates it to build a situation model that reflects knowledge about the event on the basis of information collected so far. An event index is essentially a list of spatiotemporal events as they occur. An event base stores the event's name, type, and all other relevant information.

We might not have access to the relevant information when the event is created. If this is the case, when it becomes available, the system attaches the information to the event. Thus, the event base is an organic database that grows as a result of many different processes running, in contrast to the current database form. The event

base also stores links to original data sources, so the system can present the appropriate media in the context of a particular event.

Personalized distribution mechanisms. Different perspectives of the event might interest system users. The environment model contains all information used to generate different perspectives. In most cases, the system renders these perspectives by sending each user a different combination of sensor and other data streams. Moreover, because many different camera streams exist, the system can switch the streams at different times, depending on the user's request. In a way, the system performs dynamic semantic remixing for each user.

Current streaming systems are designed for fixed data streams. Some systems prepare the data stream at the receiver by combining multiple streams. Even in this case, the receiver renders a single stream. In dynamic switching systems, however, a user's actions or requests cause the system to switch the data stream—for example, from Camera 1's video to Camera 28's video. Because each user's requests are unique, the switching combination and timing differ for each user.

This dynamic switching ability will require new media-streaming techniques. This difficult real-time resource-management problem should consider quality-of-experience issues for different users.

EventWeb

The Web has revolutionized many aspects of human society in just over a decade. However, in the current

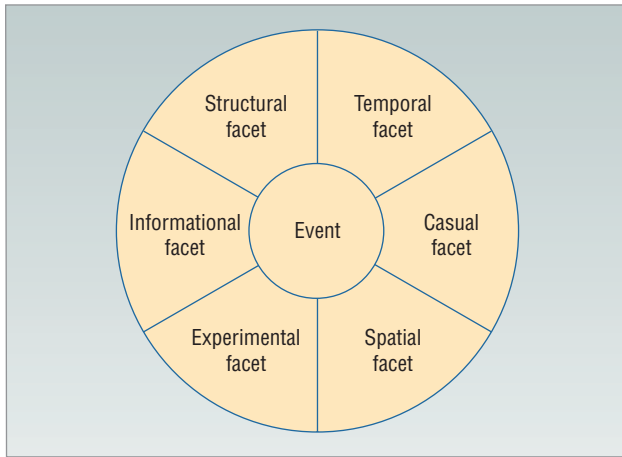


Figure 2. Different facets of an event. In EventWeb, each node (event) is represented by its basic properties—time, location, and type—as well as informational and experimental attributes.

Web each node is a document connected to other nodes using manually created referential links. Because of emerging digital-media devices and technology, we're now in a position to develop the EventWeb. Creating EventWeb will require developing technology to produce events using heterogeneous media elements, represent each event as a node, and then create explicit links between events and between events and information on the current Web.

EventWeb organizes data in terms of events and experiences and allows natural access from users' perspectives. For each event, EventWeb collects and organizes audio, visual, tactile, textual, and other data to provide people with an environment for experiencing the event from their perspective. EventWeb also easily reorganizes events to satisfy different viewpoints and naturally incorporates new data types—dynamic, temporal, and live. The current Web is document-centric hypertext. Unlike events, hypertext has no notion of time, space, or semantic structures other than often ad hoc hyperlinks.

Applications. EventWeb has significant applications, ranging from education to healthcare and from games to government. Events on this Web can offer multiple perspectives of important events, permitting remote personalized participation in live events such as meetings, lectures, concerts, and sports. Because users could conceivably take part in a meeting and change its course, they would find remote participation in a meeting different from remote participation in a concert or sports event.

Users would participate in a sporting or entertainment event only through observation from a particular perspective. In addition, they could archive these events to experience later, maybe from a different perspective each time. This could provide insights leading to valuable knowledge creation.

In EventWeb, each node (event) is represented by

- its basic properties—time, location, and type;
- informational attributes—participants and characteristics of the event; and
- experiential attributes—text or audio reports, photos, video, and other sensory information.

As Figure 2 shows, the structural and causal relationships among the events create the EventWeb. Event creators could specify these relationships explicitly, or different processes or people could discover them. Events captured in this Web could be as diverse as an application demands and at the desired granularity.

Dynamic structure. Anyone could produce an event node on EventWeb, whether it's the World Cup finals, a tsunami, a wedding, an accident, a bus arriving at a bus stop, a sale at a local store, my daughter's first music lesson, or me typing these words. You could use text descriptions, photos, audio, video, and haptic, smell, infrared, and other suitable sensor data to capture the important information and experiences related to the event.

EventWeb is fundamentally a dynamic Web structure that's linked to physical locations and uses familiar natural sensory characteristics. It uses text when needed. The EventWeb will link to the current Web, as content (for example, blogs and news sites) frequently describes events and experiences. Thus, EventWeb and the Web will work synergistically.

Realizing EventWeb will require technological advances in many areas. Some key areas are the same as for immersive telepresence. However, EventWeb also requires novel innovative concepts and tools from media processing, databases, Internet technologies, media creation and presentation, Web crawling, social networking, computer architecture, arts and architecture, and media search.

Archiving and indexing. People can use an Event Markup Language to post their events and related information and experiential data in the form of photos, audio, videos, and textual data. An EML will also provide an environment for expressing and creating relationships among events. Combining this language with event capture and a media-processing tool will help users identify events of interest in the EventWeb.

Events have an interesting life cycle. They're planned, they take place, then people store their experiences of the event in the form of experiential data and relate them to past and future events. Interestingly, past events play a more important role in our lives than current and future events.

For example, all sciences rely on the analysis of past events. Most paintings, novels, movies, and news reports are related to past experiences. An environment for capturing, archiving, and indexing events is therefore essential, as is the facility to continuously add new information and experiences and links to these events.

When billions of events (past, current, and future) exist, how do we find events of interest? We should somehow aggregate and develop an environment that lets users discover events of interest. We also need powerful and efficient indexing approaches for accessing all of these events. Current search-engine and relational-database techniques might not allow for event indexing. Instead, event indexing might require a combination of multidimensional and inverted file approaches.

Exploration and presentation environments. The keyword box on the current Web is practically unusable in EventWeb. We'll need novel approaches that combine navigation and search environments for finding appropriate events. Moreover, because photos and video are better than text for capturing and representing events, we'll need a novel presentation environment. Such a presentation environment will require a unique combination of ideas from visual arts and HCI to present event experiences to users.

Realizing immersive telepresence and the EventWeb will require advances in several research areas. Two areas in particular are essential to the development of such emerging applications and to HCC's advancement. These areas require a new perspective on current approaches and present some challenging issues. There's some research in these areas, but at best it's in its early infancy.

EVENT MODEL

Current information tools deal well with entities, objects, and keywords. However, information management in dynamic multimedia environments requires new concepts and techniques. Clearly, current concepts and tools are good for text-oriented and structured information systems that deal mostly with static information. But these tools aren't good for dealing with images, video, audio, and other sensory information. Consider, for example, major search engines' poor results for images and video. These search engines try to apply text-oriented search tools to the text associated with images and video, without processing images and video to extract meaningful indexing information from them, with surprisingly bad results.

Experiential systems

Current information tools evolved before the wave of mobile phones, digital cameras, and broadband systems changed the information system landscape. With all these advances, experiences are becoming an integral part of information systems. The recent flurry of activity in community-oriented systems such as YouTube, MySpace, and Facebook is a good example of experiential data's increasing popularity.

This data's popularity will likely increase even more rapidly in the developing world because of the abundance of mobile phones and the number of non-English-speaking people. Experiential systems deal with sensory data in sensory space without linguistic abstraction to bypass language and keyboard issues.

Event concept

The concept of "event" can serve as the fundamental organizational principle for multimedia systems. Strong conceptual, engineering, computational, and human-centered design principles support the use of event as a primary structure for organizing and accessing dynamic multimedia systems.

The definition of event depends on context and granularity. Dictionary.com defines an event as "something that occurs in a certain place during a particular interval of time." The term also refers to a significant occurrence or happening, or a social gathering or activity.

Events indeed have different contexts. A wedding is an event, as are the wedding reception and the cake-cutting. The bride and groom's first meeting is as much an event as the bride's birth, her parents' wedding, and so on. And, yes, the World Cup soccer final between Italy and France was an event, and so is my grandson's first soccer kick in his backyard. Theoretically, even moving my finger to a specific key is an event. So, events depend on context and occur at different granularities or resolutions.

The concept of "event" can serve as the fundamental organizational principle for multimedia systems.

Capturing and combining events

As Figure 3 shows, you can combine events in many ways to define other (compound) events. And you can again combine these combinations of events with other events to define yet another set of events. So this process of defining events continues. An application clearly determines these definitions. On the other side, an event is the result of one or more past events, which were in turn results of other events, and so on. Similarly, an event might result—maybe in combination with other events—in multiple events, which in turn might result in many other events. So this process of event creation has been ongoing and will continue into the future.

But, if all of these things are events, how can we capture them in our computing systems—or can we? At first, this situation appears confusing, but objects are equally confusing. Objects could be physical or conceptual. Objects can also exist at many resolutions. So, I'm an object, and so are the shirt I'm wearing and the buttons on the shirt.

Defining event aspects

Object-oriented programming and object-oriented design concepts have dominated computer science for

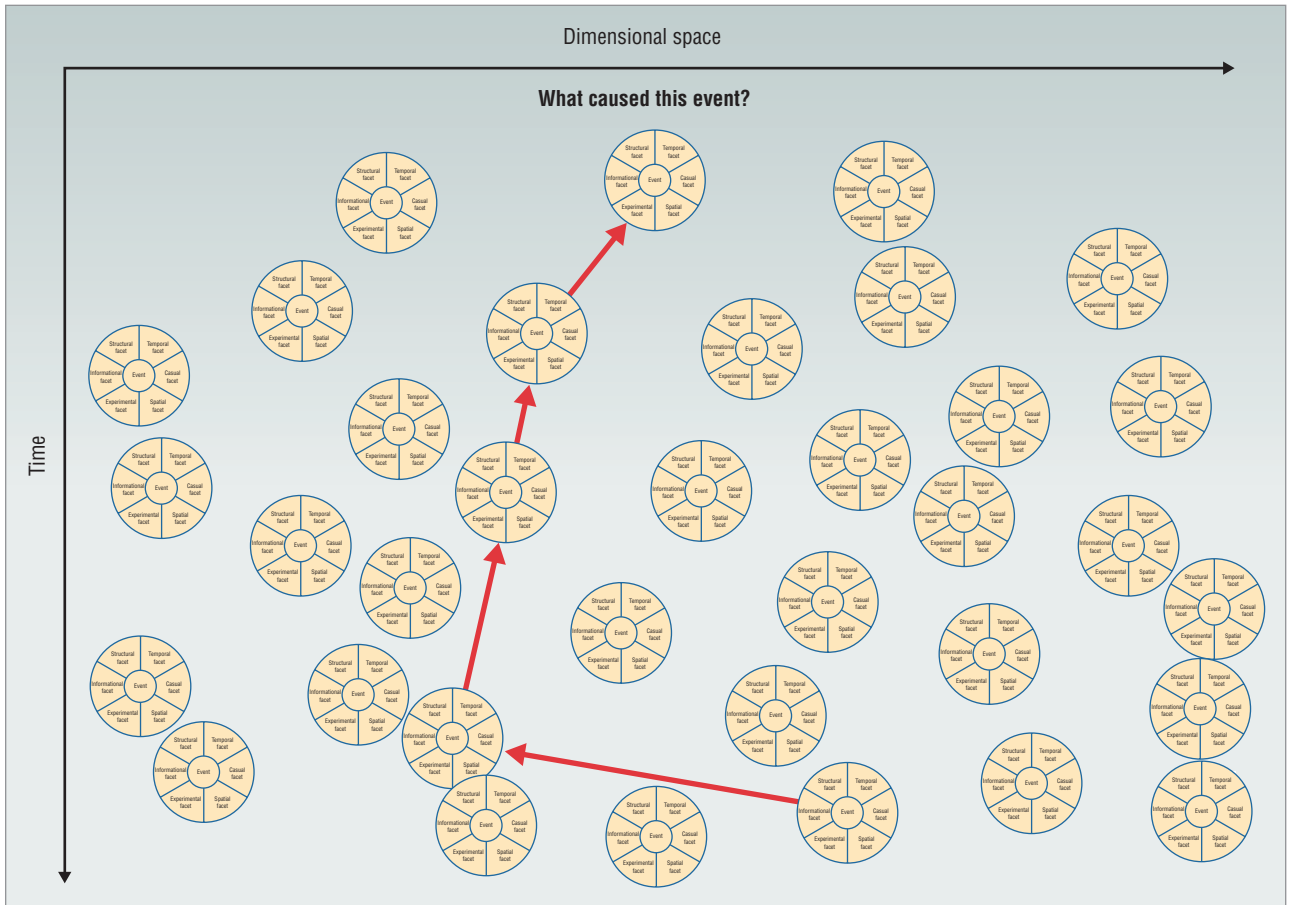


Figure 3. Time and space. We could consider a similar event occurring at a different time and space a different event.

more than a decade. This has been a powerful paradigm. In the context of computer science, “objects are a language mechanism for binding data with methods that operate on that data” (<http://en.wikipedia.org/wiki/Object>). This avoids any reference to physical, conceptual, or other types of objects that we use in our regular language, and it also provides an elegant new functional definition. Objects become a mechanism for binding data with methods that operate on that data.

Similarly, an event should explicitly define three important aspects:

- information about the event,
- experiences related to the event, and
- the event’s structural and causal relationships with other events.

An event in computational form should represent data associated with these aspects as well as the processes necessary to acquire and present them. By providing flexible and expressive mechanisms to define these three components and associated methods, we could effectively define events. The event environment should provide tools for defining any event of interest from many disparate application domains. We could, for example,

define event classes, and each event in the system could serve as an instance of a class.

An event’s basic characteristics are its identification, time, and location, with the latter two becoming the event’s fundamental defining characteristics. We could consider a similar event occurring at a different time and space a different event. In this sense, an event is defined in spatiotemporal space. Point events are just points in the spatiotemporal space, while interval events are regions in spatiotemporal space.

My team at the University of California, Irvine, is developing multiple applications using event models to validate our hypothesis that events can effectively capture multimedia semantics and help us build efficient systems to deal with multimedia information.^{5,6}

EXPERIENTIAL ENVIRONMENTS

Assume we create a system containing events and all their related information and experiences. We’ll then need an appropriate environment for interacting with such systems. The interaction environment should be human-centric and should work synergistically with human strengths and limitations.

Current information environments actually work against the human-machine synergy. Humans are effi-

cient in conceptual and perceptual analysis but relatively weak in mathematical and logical analysis, while computers are the opposite. In an experiential environment, users apply their senses to observe data and information of interest related to an event, and they interact naturally with the data based on their particular set of interests in the context of that event.

Support direct communication

People like to work in their natural environments using sensors. Unfamiliar metaphors and commands create confusion and make systems difficult to use. Similarly, a keyboard is an effective tool in some text-oriented situations, but most people would rather talk and gesture than type on a keyboard. An experiential environment should present a user with data that the human senses can easily and rapidly interpret. We should make user interactions with the data set as natural as possible. Gesture interfaces might play an important role here.

Same query and presentation spaces

WYSIWYG word processing and spreadsheets facilitated the personal computing revolution. WYSIWYG environments can provide quick feedback because they use the same space for queries and results. In fact, the notion of interface becomes just an environment for getting things done.

Different query and presentation spaces make most current information systems difficult to use. These spaces make users feel they're in a structured, rather than natural, environment. Popular search engines provide a box for entering keywords, and the system responds with a list of thousands of entries spanning hundreds of pages.

Most users never go beyond the first page. Contrast this to a spreadsheet, where users' actions result in a new sheet, showing new relationships. A WYSIWYG environment that merges query and presentation spaces would allow for easier interaction with spatiotemporal data.

User state and context

People feel comfortable in situations with static or gradually changing contexts and states. Computing systems should know the user's state and context and present information that's relevant to that state and context. Current information systems, including databases, were designed to provide scalability and efficiency, which are better achieved in stateless environments. This design was justified in the early days of computing, when computers were very expensive compared to human time. The situation has changed dramatically, however, and computing time is cheap compared to human time. So, we'll need to design systems that maintain context and state to maximize the efficiency and quality of users' experiences.

Perceptual analysis

Humans can use their perceptual facilities more rapidly and efficiently than machines, hence the popularity of visualization. Minimizing text-oriented displays and presenting information using appropriate sensory modalities can significantly improve performance and experience. Video's increasing popularity over any other medium is due to its ability to combine multiple modalities. Videogames and many simulation systems are so engaging because they provide a powerful visual environment, sound, and, in some cases, tactile inputs.

As speech recognition and computer vision become more sophisticated, we'll likely see more multimodal interfaces. Experiential environments must organize information as well as maintain user context and state. WYSIWYG applications are the first step in the direction of experiential environments.

It's no surprise that videogames are so successful. They're easy to learn and natural to use, and they provide a compelling experience by engaging our senses. General-computing-environment designers can learn a lot from the interaction environments in videogames.

How can we bring computing to the more than 5 billion of the world's 6 billion people it hasn't yet reached? Although many think selling a computer for less than US \$100 would help achieve this, sound HCC practices would definitely help.⁹

Mobile phones—which are easy to use and outnumber computers by two to one—can bring the Internet and computing to the masses, even in remote parts of the world. Providing audiovisual-tactile interfaces in phones can help people create and access content. The iPhone and advanced phones like Nokia's N95 are more powerful than PCs were less than a decade ago, and they have built-in cameras, microphones, and other sensors. These devices could make content creation and access less dependent on language and education level than in current systems.

It's time to reduce our dependence on text by using all other data modalities and organizing and presenting information in ways that are more perceptual and cognitively meaningful to humans. By providing an environment for experiencing events to gain insights and share event experiences, technology will accelerate knowledge growth at an unprecedented rate. Moreover, the ability to effortlessly share event experiences with fellow humans will help identify those with similar interests and subsequently build communities across the globe. ■

References


1. P. Rander, "A Multi-Camera Method for 3D Digitization of Dynamic, Real-World Events," tech. report CMU-RI-TR-98-

12. Carnegie Mellon Univ., 1998; www.ri.cmu.edu/pub_files/pub1/lander_peter_1998_1/lander_peter_1998_1.pdf.
2. R. Jain, "Experiential Computing," *Comm. ACM*, July 2003, pp. 48-55.
3. S. Vedula et al., "Three-Dimensional Scene Flow," *IEEE Trans. Pattern Analysis and Machine Intelligence*, Mar. 2005, pp. 475-480.
4. H. Rheingold, *Virtual Reality*, Simon & Schuster, 1992.
5. R. Jain, "Out of the Box Data Engineering: Events in Heterogeneous Data," *Proc. Int'l Conf. Data Eng. (ICDE)*, CS Press, Mar. 2003, p. 8.
6. G.U. Westermann and R. Jain, "Events in Multimedia Electronic Chronicles (E-Chronicles)," *Int'l J. Semantic Web and Information Systems*, vol. 2, no. 2, 2006, pp. 1-27.
7. G. Utz Westermann and R. Jain, "Toward a Common Event Model for Multimedia Applications," *IEEE MultiMedia*, Jan. 2007, pp. 19-29.

8. A. Katkere et al., "Towards Video-Based Immersive Environments," special issue on multimedia and multisensory virtual worlds, *Multimedia Systems*, Mar. 1997, pp. 69-85.
9. R. Jain, "Folk Computing," *Comm. ACM*, Apr. 2003, pp. 27-29.

Ramesh Jain is the Donald Bren Professor in Information and Computer Sciences in the Department of Computer Science at the University of California, Irvine. His research interests are multimedia information systems, image databases, and machine vision. Jain received a PhD in computer science from the Indian Institute of Technology. He is a Fellow of the IEEE and the ACM. Contact him at jain@ics.uci.edu.

stay on the Cutting Edge of Artificial Intelligence



IEEE Intelligent Systems provides peer-reviewed, cutting-edge articles on the theory and applications of systems that perceive, reason, learn, and act intelligently.

The #1 AI Magazine 

www.computer.org/intelligent