THE HORIZON REPORT

2004 EDITION

NMC: The New Media Consortium

TABLE OF CONTENTS

EXECUTIVE SUMMARY

This first edition of the NMC's annual *Horizon Report* details the recent findings of the NMC's Horizon Project, a research-oriented effort that seeks to identify and describe emerging technologies likely to have a large impact on teaching, learning, or creative expression within higher education. The *2004 Horizon Report,* produced with support from the Corel Corporation, highlights six technologies that the research suggests will become very important to higher education within three adoption horizons over the next one to five years.

The project draws on an ongoing series of interviews with scientists, engineers, technologists, and other knowledgeable individuals in business, industry, and education to identify potentially promising technologies, as well as published resources such as the Gartner Group's Hype Cycles and other examinations of emerging technologies. The Horizon Project's Advisory Board provided another perspective on the relative utility of each of the technologies identified.

The technologies chosen for the *2004 Horizon Report* are highlighted over three adoption horizons: The first assumes a high likelihood of broad adoption within the next year; the second, adoption within two to three years; and the third, adoption within four to five years. The choices in the first category, scalable vector graphics and learning objects, are already being applied on dozens of leading campuses across North America. Applications for both can be expected to grow substantially within that timeframe.

The four other technologies are spread over the two remaining timelines. They are described in the following paragraphs and also show considerable promise for higher education. All are already the subjects of substantial development in the private sector — but their educational applications are still being explored. As the adoption horizon moves out, the applications and examples provided are more conceptual and prototypical, but they do provide a sense of how each highlighted technology might be applied, and why it has been included here. Within

two to five years, all four are expected to be in broad use by colleges and universities.

Technologies Detailed in the 2004 Horizon Report

- **Learning Objects**
  Learning objects are assemblies of audio, graphic, animation and other digital files and materials that are intended to be reusable in a variety of ways, and easily combined into higher-level instructional components such as lessons and modules. The primary purpose behind the development and use of learning objects is to increase access to quality content, and to avoid wasteful replications of effort by making that content usable in a variety of contexts.

- **Scalable Vector Graphics (SVG)**
  SVG uses XML for describing two-dimensional graphics, holding the information needed to draw an image in a text file. Scaling is smoothly achieved without jagged edges. Graphical objects can be styled, transformed, grouped, or placed into previously rendered objects. Text is searchable and selectable. SVG is an especially powerful tool for instructional developers on college and university campuses, with potential applications in virtually any discipline, but especially the sciences and engineering.

- **Rapid Prototyping**
  Rapid prototyping refers to what amounts to 3-D printing, e.g., building three-dimensional physical objects from digital data files. These files may be created in a variety of ways, such as computer-aided design (CAD), computer-aided tomography (CAT), or even X-ray crystallography, then output to a rapid prototyping machine that creates a physical model of the object. This technology already is widely used for a variety of manufacturing, design, and engineering applications, but as cost decreases, is finding new applications in the arts and the classroom.

■ Multimodal Interfaces

Multimodal interfaces provide ways for humans to interact with computers beyond the traditional mouse and keyboard, using inputs and outputs that target not only each of the five senses, but also take advantage of nonverbal cues common in human conversation. Considerable development is taking place in simulations that use multimodal techniques (haptics or force feedback, for example) to great effect.

■ Context-Aware Computing

Context-aware computing refers to computing devices that can interpret contextual information and use it to aid decision-making and influence interactions. Contextual cues may include what the user is attending to, the user's location and orientation, the date and time of day, lighting conditions, other objects and people in the environment, accessible infrastructure in the immediate vicinity, and so forth. Context-aware applications can make decisions based on such information without the need for user input.

■ Knowledge Webs

Knowledge web is a term that describes a dynamic concept of individual and group knowledge generation and sharing, with technology used to make connections between knowledge elements clear, to distribute knowledge over multiple pathways, and to represent knowledge in ways that facilitate its use. Work in knowledge webs overlaps considerably with that going on around communities of practice, and holds the potential to help such communities share, create, analyze, validate, and distribute existing and emerging knowledge.

Imbedded in the description of each technology is a discussion of its relevance for teaching, learning, and creative expression. Links to several example applications are provided in each section, as well as an annotated list of additional readings.

Several important themes emerge from the choices listed in the *2004 Horizon Report.* Technology is providing learners with an ever greater access to learning materials, and because of the rich variety of interaction pathways being developed, these materials are becoming more "real" and more responsive all the time. As computing becomes more ubiquitous and embedded in more and more everyday objects, the potential applications that are emerging will draw more on contextual cues. These cues will feel at once more invisible and more pervasive, with the result that users will be able to generate, use, and share knowledge in ways we can only imagine today.

LEARNING OBJECTS

Time-to-Adoption Horizon: One Year or Less

This technology is in use in a rapidly growing number of universities today, as evidenced by the tremendous growth of both general and discipline-based learning object repositories, and will likely find its way into broad use within the next 12 months.

Overview

Learning objects have garnered considerable attention within higher education in recent months, and the concept is viewed as an important strategy in making substantive use of the growing collections of digital materials. At the same time, confusion continues to exist in practice about what a learning object is, and how it might differ from content files such as a photo, a video clip, or a research report. A practitioner might ask at what point do typical digital assets like these become learning objects? What distinguishes a learning object from any other sort of learning material?

The most common view is that a learning object is a collection of digital materials — pictures, documents, simulations — coupled with a clear and measurable learning objective or designed to support a learning process. This view distinguishes a learning object from an "information object" (akin to a simple fact) — which might have an illustration or other materials attached to it — or from "a content object" such as a video or audio clip, picture, animation, or text document. The key distinguishing feature between these kinds of objects and a learning object is the clear connection to a learning process. This definition is built on the assumption that by combining learning objects in different ways, higher-level learning goals can be met, and ultimately, entire courses constructed (Johnson, 2003).

The National Learning Infrastructure Initiative (NLII) has adopted a broad definition for learning objects, drawn from the work of David Wiley. Wiley defines a learning object as any digital resource that can be reused to support learning. NLII adds that learning objects are "digital resources, modular in nature, that are used to support learning. They include, but are not limited to, simulations, electronic calculators, animations, tutorials, text entries, Web sites, bibliographies, audio and video clips, quizzes, photographs, illustrations, diagrams, graphs, maps, charts, and assessments. They vary in size, scope, and level of granularity ranging from a small chunk of instruction to a series of resources combined to provide a more complex learning experience" (NLII, 2003).

Reusability is a key attribute of a learning object, and means the content can be accessed and used in a variety of settings and for a variety of purposes. The concept of standard deviation, for example, is important to a wide array of disciplines. A well-designed learning object that illustrates the concept would be useful as part of an economics curriculum, a psychology course, a statistics course, a course in epidemiology, or any one of a number of other courses that require understanding of the concept. A library of such objects would make the job of assembling a web-delivered course much simpler, and students could use the resource to explore topics or self-remediate.

Relevance for Teaching, Learning, or Creative Expression

The work of the NLII, mentioned earlier, included an exploration of a variety of potential learning object applications, including the use of learning objects in "learner-centered, collaborative, cost-effective, high-quality, and accessible" learning environments. What they found was that learning objects, if designed to support deeper learning, offered considerable potential to increase student ownership of the learning process and could be well suited for learning activities that are social, active, engaging, and contextual.

One of the benefits of object-based authoring is that it is very compatible with learning designs that support such learner-centered approaches. Among the many promising uses driving development of learning objects is the potential to individualize learning in ways that have not been possible before. By linking learning objects with adaptive learning assessment approaches, it should be possible to highly customize a learning experience, and deliver it "on the fly," in real time.

Building learning materials in ways that make it easy for them to be reused in a variety of situations represents a shift in thinking about the design, development, and deployment of instructional content. Considerable work is taking place in the design of ontologies, topologies, and metadata approaches to make learning objects easy to catalog and find. A number of important learning object repositories have begun to emerge, and one of the most useful applications of these repositories is that they make it much easier for faculty and instructional designers to provide prerequisite or background knowledge to students. Similarly, these sorts of repositories facilitate the inclusion of material supporting more advanced explorations of a topic.

Although the "whys" for using learning objects have become clearer to practitioners in recent months, the "hows" are still being determined. Standards that permit interoperability are still being developed, and many organizations do not yet have the infrastructure to deploy learning objects. Deeper learning requires that learning be infused with the context in which the learning is used in the real world. Aggregating or combining the components into learning objects and then into instructional elements adds context but optimal design of instructional elements using learning objects is still being considered and debated. Designing object-oriented learning requires new ways of thinking, methodologies, and technologies, and these changes to the status quo present potential barriers to adoption.

Despite these challenges, learning objects have emerged as an approach with considerable potential to impact the design and construction of learning materials. The list of learning object repositories grows daily. The content in these collections is also growing rapidly, and already spans the range of disciplines in higher education.

Learning Object Examples

The following links provide examples of learning objects and learning object repositories.

MERLOT Multimedia Educational Repository
http://www.merlot.org

> MERLOT is a free and open resource designed primarily for faculty and students of higher education. Links to online learning materials are collected here along with annotations such as peer reviews and assignments. Community members help MERLOT grow by contributing materials and adding assignments and comments.

CAREO: Campus Alberta Repository of Educational Objects
http://careo.ucalgary.ca/

> Adopted by many institutions as a way to manage digital resources for use in teaching and learning. CAREO currently has almost 4,000 learning objects available to its 700 users, and adoption is increasing every month.

Maricopa's Learning eXchange (MLX)
http://www.mcli.dist.maricopa.edu/mlx/

> The Maricopa Learning eXchange is an electronic warehouse of ideas, examples, and resources that support student learning at the Maricopa Community Colleges.

EdNA
www.edna.edu.au

> EdNA Online provides access to 18,000 Australian quality-assured education and training resources for school, early childhood, vocational education, adult continuing education, and higher education.

NLII List of Learning Object Repositories
http://www.educause.edu/nlii/keythemes/ learningObjects.asp

> The National Learning Infrastructure Initiative (NLII) provides an extensive list of sources for learning objects

For Further Reading

The following articles and resources are recommended for those who wish to learn more about learning objects and related topics.

Making Sense of Learning Specifications & Standards
http://www.masie.com/standards/S3_Guide.pdf

> This white paper, facilitated by the S3 Working Group of the e-Learning Consortium, was created "to help the average person understand the rationale, development, and implication of learning standards and to accelerate their adoption." The first part of the paper serves as a primer for those who have little to no knowledge of learning standards.

Learning Objects and Learning Standards: Everything You Ever Wanted to Know But Were Afraid to Ask.
http://www.learnativity.com/standards.html

> In this paper, Hodgins first describes the need for standards as they are related to maximizing learning technologies, then he raises several key questions driving standardization projects

The Instructional Use of Learning Objects
http://www.reusability.org/read/

> An online version of David Wiley's edited book that explains what learning objects are and how they can be used for teaching and learning.

NLII Reading List
http://www.educause.edu/nlii/keythemes/LearningObjects.asp#resources

> Learning objects are an emerging NLII theme because their use has the potential "to provide learning customized for each specific learner at a specific time, taking into account, their learning styles, experience, knowledge and learning goals" (Schatz, 2000). The project reading list spans a broad range of topics related to learning objects.

SCALABLE VECTOR GRAPHICS (SVG)

Time-to-Adoption Horizon: One Year or Less

This technology is already in use in many universities today, and is especially prevalent in cartographic and geographic applications. Because of its many powerful features, within 1-2 years SVG will also find significant applications in other fields. As colleges continue to embrace open-source standards and approaches, SVG implementations can be expected to grow. Within the next 3 to 5 years, it is highly likely that SVG will be in mainstream use and applied in a great variety of learning contexts and applications.

Overview

Scalable Vector Graphics (SVG) is a language created in 1999 and developed under the banner of the World Wide Web Consortium (W3C), the non-profit, open-standards group that created HTML and XML. SVG is designed as an open standard for the description of two-dimensional graphics using eXtensible Markup Language (XML), and has broad industry support from companies such as Corel Corporation, Macromedia, Adobe Systems, AOL/Netscape, Apple Computer, Microsoft, Sun Microsystems, and many others.

SVG incorporates three types of graphic objects: vector graphic shapes, images, and text, as well as the basic graphical forms — rectangles, circles, ellipses, and polygons. With these components, virtually any image can be created. Graphical objects can be styled, transformed, grouped, or placed into previously rendered objects. Text is searchable and selectable.

SVG files hold the graphical information necessary to render an image in a text file that contains a sequence of commands that is used to draw the image. Because the image is actually redrawn when zoomed, scaling is smoothly achieved without jagged edges. SVG produces clear, crisp, high-resolution output on any display, and offers unlimited color and font choices, as well as linking and scripting control for animation and interactive events.

Because of these capabilities, SVG has many advantages over other image formats, and particularly over the most common web graphic formats, JPEG and GIF, which are composed of bitmaps or individual pixels. The following points highlight the ways in which SVG graphics differ from these bit-mapped formats.

- **Plain-Text File Format** — Because SVG files are simple text files, they can be read and modified by a range of tools. Additionally, the files are usually much smaller and more compressible than comparable JPEG or GIF images.

- **Scalable Output** — SVG is a vector format, which means SVG images render with high quality at any resolution, without the degradation or "jaggies" you see when scaling bitmapped images.

- **Resolution-Independent Zooming** — SVG images allow zooming on any part of the image without degradation in quality.

- **Searchable and Selectable Text** — Text in SVG images is text, and thus can be manipulated and searched easily. In one commonly cited example, a street or other location could be found within a large SVG map file with a simple text search for the street name.

- **Bi-Directional Data Exchange** — Using scripting, SVG files can be manipulated in real time to create sophisticated applications using dynamic and interactive graphics.

- **True XML** — Because SVG is an XML grammar, SVG graphics can easily be generated on web servers "on the fly," using standard XML tools. For example, a web server could generate a high-quality, low-bandwidth graph using SVG from real-time sensor data.

The SVG feature set includes nested transformations, clipping paths, alpha masks, filter effects, template objects and extensibility, and gives developers the ability to control and implement sophisticated

animation techniques. Filter effects give developers the ability to add effects directly to shapes and text in a manner similar to those achieved using professional imaging tools. SVG provides the same features typically associated with bitmap images (GIF, JPEG, PNG, etc.) including blending, tiling, transforming or rotating elements, morphing, offset, merge, and special lighting effects. Event handlers such as *onmouseover* and *onclick* can be assigned to any SVG object or element.

Within an SVG image, every element and every attribute of an element can even be animated, and any portion of an otherwise static graphic can be externally processed and modified. Because the Document Object Model (DOM) for SVG includes the full XML DOM, this information can be manipulated through a variety of methods (scripting, for example) to display application-specific graphic images that can be rendered based on user input or other data in real time.

Relevance for Teaching, Learning, or Creative Expression

The combination of sophisticated graphics and animation effects, high-resolution output, and bi-directional data interchange make SVG an especially powerful tool for instructional developers on college and university campuses, with potential applications in virtually any discipline, but especially the sciences and engineering. In fields like cartography and geography, which take unique advantage of its high resolution output, layers, and degradation-free transformations, SVG has already achieved a critical mass of adopters.

As the examples at the end of this section illustrate, SVG is being used to develop high-quality simulations and learning objects for a variety of disciplines. Since SVG is an XML grammar, the language is compatible with the specialized XML-based markup languages being developed for scientific, technical, medical, and engineering applications. For educational applications, SVG allows developers to easily add engaging new components to instructional web sites, such as:

- Dynamic data links to and from controls, monitoring devices, or real-time sensors;
- Rich simulations of processes or experiments;
- High-resolution, zoomable maps and illustrations with user-selectable content; and
- Tables and diagrams responsive to user input or displaying live data.

SVG also shows considerable promise as a tool for the digital arts and creative expression, where it is already being used to create art forms that can be influenced by user inputs or dynamic data streams.

A growing number of campuses are building repositories of reusable learning objects, and interactivity is a key feature that faculty and developers wish to include in learning materials. With its strong XML heritage, SVG offers considerable promise as a platform for the creation of specialized learning objects. Applications across a sampling of disciplines could include:

- **Sciences.** Interactive simulations and experiments with the results presented graphically on the web;
- **Engineering.** Scalable, zoomable plans and drawings, simulations, and interactive animations;
- **Medicine.** Live web monitoring of patient data, high-resolution illustrations, patient simulations;
- **Business.** Data-driven dynamic charts and graphs of stock and other business data; and
- **Mathematics and Statistics.** Graphical display of complex data and models.

With new tools such as the Corel Smart Graphics Studio streamlining the process of creating SVG files and making the technology more accessible, the number of faculty, instructional designers, and content developers adopting SVG as a platform for creating dynamic, data-driven simulations with interactivity and rich graphic output is expected to increase significantly.

SVG Examples

The following examples demonstrate a variety of ways in which SVG can be applied to learning contexts.

Vienna Social Patterns & Structures
http://www.karto.ethz.ch/neumann/ cartography/vienna/

In this site, an SVG-driven map interface is used to present data on social patterns and structures in Vienna.

Interactive Mortality Charts and Graphs
http://www3.cancer.gov/atlasplus/charts.html

Cross-tab style charts and graphs are created interactively by selecting from lists of variables.

Interactive Maps: DBx Geomatics
http://www.dbxgeomatics.com/SVGMapMaker Samples.asp?Language

This Canadian interactive web mapping company provides a wide range of examples of interactive maps showing population, weather, fishing areas, etc, based on GIS information.

Real Time Data Monitoring
http://www.wpsenergy.com/JayNick/CTG/ DNRStats.html

This site illustrates how SVG can be used to create an interface for monitoring real-time data.

Animation: Signal Read-Out
http://www.usbyte.com/SVG/CDROM Optical Pickup.html

An SVG Animation of the signal read-out from the optical pick-up of a CD Player.

For Further Reading

The following articles and resources are recommended to those who wish to learn more about Scalable Vector Graphics.

W3C SVG Site
http://www.w3.org/Graphics/SVG

The official SVG site provides complete documentation and information on the SVG recommendation.

Scalable Vector Graphics (SVG): An Executive Summary
http://wwws.sun.com/software/xml/developers/ svg/

This article by Sun provides a concise overview of SVG.

SWF is not Flash
http://www.oreillynet.com/pub/a/javascript/ 2002/05/24/swf_not_flash.html

In this article on the O'Reilly site, the SVG and SWF formats are compared and contrasted.

IDC White Paper on SVG
http://www.nmc.org/projects/horizon/Corel White Paper.pdf

IDC details the advantages of SVG, and describes the Corel Smart Graphics Studio SVG authoring platform.

Smart Graphics Technology Overview
http://www.nmc.org/projects/horizon/Corel Graphics Solution Paper.pdf

Corel's white paper on SVG and the technology underlying the Smart Graphics Studio.

Adobe SVG Zone
http://www.adobe.com/svg/main.html

Adobe's SVG site contains information on SVG and links to example content.

RAPID PROTOTYPING

Time-to-Adoption Horizon: Two to Three Years

This technology is currently in wide use in industrial engineering, but is quickly beginning to be adopted in other educational settings. 3-D printing and prototyping technologies will likely find their way into broad use across university campuses within the next two to three years.

Overview

The term *rapid prototyping* refers to technologies that construct physical objects from three-dimensional (3-D) digital content such as computer aided design (CAD), computed aided tomography (CAT), and X-ray crystallography. A 3-D printer (rapid prototyper) builds a tangible model or prototype from the file, one layer at a time, using an inkjet-like process to spray a bonding agent onto a very thin layer of fixable powder. The bonding agent can be applied very accurately to build an object from the bottom up, one very thin layer at a time. The process even accommodates moving parts within the object. Using different powders and bonding agents, color can be applied, and prototype parts can be rendered in plastic, resin, or metal.

This technology is commonly used in manufacturing to build prototypes of almost any object (scaled to fit the printer, of course) — models, plastic and metal parts, or any object that can be described in three dimensions.

An aerospace engineer might use rapid prototyping to output a scale model for flight worthiness testing in a wind tunnel. A cell phone designer might build a prototype of a new phone model and ask potential users whether it feels right. A surgeon could construct a model of a child's spine to plan an upcoming surgery, and then use the model to explain the process to the child's parents. In all three cases, access to a physical model has great advantages over a picture. Models and prototypes make excellent visual aids for communicating and testing ideas.

Designers have always utilized prototypes and models but rapid prototyping makes this process faster and less expensive than traditional methods. Rapid prototyping has literally revolutionized industrial engineering by reducing the time and cost of product development and is now beginning to find applications in the arts, education, and many other fields where rapid visualization is advantageous. In the past two years, rapid prototyping has moved into other areas of design as a common tool, and has found considerable application in the arts as well.

The cost of rapid prototyping printers has declined dramatically in the last ten years, and relatively inexpensive 3-D printers (under $1,000) are on the near-term horizon.

Relevance for Teaching, Learning, or Creative Expression

The combination of relatively low cost, rapid turn-around, and a variety of media formats (plastic, resin, and metal, with hybrids in development) make rapid prototyping an especially useful tool for anyone needing to render three-dimensional models. For faculty and students on college and university campuses, potential applications can be found in virtually any discipline. In fields like engineering and medicine, where three-dimensional data sets are common, rapid prototyping has already achieved a critical mass of adopters.

As the examples at the end of this section illustrate, extensive applications of rapid prototyping exist in the health sciences, and the technology is

increasingly being used to render complex data sets from chemistry and biology. The resulting models enable students to experience the data not only cognitively, but also visually and tactilely. Rapid prototyping also shows considerable promise as a tool for the digital arts and creative expression, where it is already being used to redefine sculpture and 3-D design.

A growing number of campuses are using rapid prototyping not only for fast and cost-efficient creation of models and parts, but are also expanding the science behind rapid prototyping to enable the creation of functional devices with moving parts and even circuitry.

Applications across a sampling of disciplines include the following.

- **Architecture.** Some new approaches to architecture and engineering do not lend themselves to traditional model building, which is time consuming and expensive. Rapid prototyping allows models to be easily rendered at any stage of design.

- **Engineering.** Prototypes of forms and parts can be created at any stage to permit rapid evaluation and refinement.

- **Arts & Design.** Sculpture of forms not feasible in other media are being created, or that incorporate moving parts in ways not possible using other approaches.

- **Mathematics.** New approaches to modeling visualizations of data or geometric forms are emerging that allow tactile as well as visual interpretation.

- **Science.** Molecular models can be built directly from crystallographic and other data, providing additional insights into structure and function.

Allowing students to design and create physical models and prototypes should encourage them to experiment and learn in new ways. The implications for design, sculpture, and other disciplines that work in three dimensions when ideas can be cheaply and quickly rendered are many.

Rapid Prototyping Examples

The following examples show how rapid prototyping has been applied in manufacturing, art, medicine, and education.

Digital Sculpture
http://www.bathsheba.com/

> Bathsheba Grossman is an artist exploring the region between art and mathematics, using rapid prototyping to produce artwork of unusual and distinctive design. Grossman uses a direct-metal printing technique, which produces a composite steel-bronze metal with a rich texture.

Models of Proteins & Other Complex Molecules
http://www.rpc.msoe.edu/3dmd/

> 3D Molecular Designs produces accurate, three-dimensional physical models of proteins and other molecular structures using rapid prototyping technology. Models are based on x,y,z coordinates obtained from protein Data Bank files of known structures.

Anatomical Models for Visualization, Surgical Planning, and Patient Education
http://www.medicalmodeling.com/flashsite/ surgeons.html

> This company uses rapid prototyping to produce highly accurate medical models from medical imaging modalities such as CT and MRI.

Biological Models and Implants
http://www.anatomics.net

> This company employs a wide variety of rapid prototyping technologies to produce highly accurate medical models based on medical imaging data for surgical, teaching, and analysis.

For Further Reading

The following articles and resources are recommended for those who want to learn more about rapid prototyping and related topics.

The Learning Factory Homepage
http://www.me.psu.edu/lamancusa/rapidpro/index.htm

> Produced by the Learning Factory in the Department of Mechanical Engineering at Penn State, this site includes an excellent primer on rapid prototyping.

The Rapid Prototyping Homepage
http://www.cc.utah.edu/~asn8200/rapid.html

> A frequently updated resource page maintained by the Department of Mechanical Engineering in the Rapid Prototyping Research Group at the University of Utah.

3-D Printing's Great Leap Forward
http://www.wired.com/news/technology/0,1282,59648,00.html

> Wired Magazine explores 3-D printing devices.

More about Rapid Prototyping Printers

To get a better understanding for what these printers look like and how they operate, it may be helpful to visit websites of some of the primary manufacturers of these products. Z Corporation has a Flash animation of one of their printers in action that is especially helpful for understanding how 3-D printers work (see http://www.zcorp.com/products/demo.asp?ID=1).

Z Corporation *http://www.zcorp.com/*

3-D Systems *http://www.3dsystems.com/*

Solidimension *http://www.solidimension.com/*

Objet *http://www.objet.co.il/home.asp*

MULTIMODAL INTERFACES

Time-to-Adoption Horizon: Two to Three Years

These technologies are in wide use in simulation, gaming, and other environments today, and are beginning to appear in disciplines like medicine and related health professions. Because of their compelling applicability to simulations, they will likely find their way into broader use across university campuses within the next two to three years. ADA-related applications are already commonplace, and will become more so over that timeframe.

Overview

In most cases, interaction between humans is multimodal, which means we use a variety of senses during the give and take of conversation. We understand others' intentions, for instance, not just through words but also through gestures and expressions, and even the measured pace of breathing. Communication occurs on many levels and via many pathways.

When interacting with computers, however, we primarily communicate in a very rudimentary fashion, using approaches that have been largely unchanged for almost two decades. The most common way of interacting with computers today (labeled WIMP by some researchers — windows, icons, menus, and pointers) is an approach that depends highly on the visual sense. The human directs the computer by manipulating or reacting to changes on the computer screen, and through his or her input via the mouse and keyboard.

This limitation is changing. An increasing number of researchers are developing and evaluating technologies that allow humans and computers to interact using multiple sensory pathways. The research promises to greatly enrich the way we communicate with and use these tools.

Called multimodal interfaces, these new approaches provide avenues for interacting with computing devices that span (and even extend) the range of human sensory ability. Visual, auditory, haptic, olfactory, and flavor pathways are all being explored, and that work is developing in two related directions.

In the first of these, the aim is to free users from having to use the traditional interface and thus make interactions with computers easier, more natural, and more intuitive. The second, which taps very similar technologies, aims to enrich the computer experience and make whatever is happening — communication, interaction, or a simulated activity — seem more real to the user.

Multimodal interfaces are already being applied to a variety of situations where multi-sensory information is especially valuable. Modern medical training, for example, makes use of a wide variety of simulators that provide detailed haptic, auditory, and olfactory feedback to the user.

In one example in common use today, the student practices an angioplasty using a real catheter in a simulated patient. As the catheter is inserted, the student "feels" the differences in various fascia and tissue, as well as key resistance points. A simulated fluoroscopic display illustrates progress from the femoral insertion point toward the blocked coronary artery. As the catheter enters the heart, the "pull" of the heartbeat affects its progress. As in real life, the student must avoid puncturing an arterial wall using only the fluoroscopic display and "feel."

In an example pulled from the headlines, security officers at an airport under high alert for a terrorism threat use facial-recognition systems to pick suspected terrorists out of large crowds. The system creates a facial print of each person in the crowd and cross-references it against a database of faces of

suspected terrorists. The system greatly augments typical video surveillance, and allows officers to be deployed more precisely. In another aspect of the same example, a computerized "sniffer" checks baggage for chemical cues that might indicate the presence of explosives.

Multimodal interfaces are seen as particularly important for mobile devices, which are limited by small screens and awkward keypads. An example is the "hands busy, eyes busy" scenario. In such a case, two-way speech communication is especially helpful.

Two multimodal technologies, speech-to-text and speech recognition, have seen considerable advances in the past few years, to the point that many companies use them in conjunction with automated attendants in telephone systems, and as part of the "intelligent agent" technologies being applied to customer service functions. Sprint Corporation, for example, uses speech technologies to drive their customer support agent, known as *Clair.*

Gesture-recognition systems are under development by a number of companies, and are being applied to situations where communication often depends on gestures. A person providing directions will often point to ensure that the directions are understood. Visual interaction systems build on gesture recognition, in which users can communicate by turning their heads or pointing with their arms, and aim to enable humans to communicate with computers implicitly, using facial expressions and gestures. Facial expressions are the focus of a related body of research, which is being applied to avatars, such as Kurzweil's *Ramona,* that are able to replicate human facial expressions and gestures more and more accurately.

ADA-related applications for multimodal interfaces are increasingly commonplace, and specialized multimodal interfaces such as eye-tracking, breath-controlled mice, and other technologies have provided computer access and control to people that may be unable to effectively use computers because of the complex series of steps required

or the motor skills needed to make use of current interfaces and input methods.

Relevance for Teaching, Learning, or Creative Expression

The potential of multimodal devices to greatly enhance the realism of simulations make these tools especially appropriate for designers of such experiences both in business and on college and university campuses. The potential applications could apply to virtually any discipline, but especially in disciplines in which realistic simulation environments can enable a suitable level of experience to be gained in a safe and reinforcing manner. In fields like medicine, which take unique advantage of haptic feedback in training medical professionals, multimodal interfaces have already achieved a critical mass of adopters.

As the examples at the end of this section illustrate, multimodal approaches are being used to develop not only very high-quality simulations, but also engaging conversational agents that may find considerable application in learning.

Haptic devices in particular show considerable promise as tools for the digital arts and creative expression, where they are already being used to create art forms in virtual space. Museums are increasingly interested in the technology as a means of providing visitors a "hands-on" experience that poses no risks for valuable artifacts.

A growing number of campuses are building applications for mobile devices that provide students access to campus services and information, and that use speech recognition and text-to-speech approaches to make these services easier to access.

Applications across a sampling of disciplines include the following.

- **Languages.** Researchers are studying the use of conversational avatars in ESL or foreign language learning. Computers that understand speech and intonation provide feedback to

learners on both meaning and pronunciation.

- **Art & Design.** Haptic devices allow sculptors, designers, and others to "feel" the virtual artworks created. Disabled artists may soon be able to create by simply describing the art they envision so the computer can create it.

- **Medicine & Health Sciences.** Catheter simulators simulate the feel of fascia and tissues involved in angioplasty. Patient simulators mimic the outward symptoms of a heart attack such as shortness of breath and clammy skin.

- **Engineering.** Force feedback allows engineers to experience the "feel" of elastic distortion and stresses in structural components.

- **Adaptive Support.** Screen readers read web pages to users with visual disabilities. Speech-to-text software allows students with motor difficulties to write research papers. Gesture recognition allows students to communicate with the teacher and with other class members. Eye tracking allows the mobility-impaired user to advance text from one screen to the next.

Technologies that make computer interactions easier, more natural, and more realistic are already finding their way onto campuses, and applications that make use of them can be expected to grow substantially in the next two to three years. Over that same time period, simulation technologies will expand from their current fairly common applications in fields requiring skill and dexterity into a much wider range of disciplines.

Multimodal Interface Examples

The following examples show multimodal interfaces in use or development.

Haptics and Surgical Simulation
http://cs.millersville.edu/~webster/haptics/

A description of a research project that is developing software for use in simulating a suite of surgical procedures that provides several examples of haptics-based simulations. Virtual reality surgical simulators provide sensitive touch feedback along with realistic 3D imagery.

SensAble Technologies Virtual Sculpture and Design Tools
http://www.sensable.com/

SensAble's FreeForm Concept system includes the PHANTOM® Omni™ haptic device that allows designers to use their sense of touch to transform the images in their mind's eye into 3D designs in ways that aim to be as easy and natural as sketching and painting. A number of demonstrations are available on their site.

Haptics Products from the Immersion Corporation
http://www.immersion.com/products/

This company is producing haptic interfaces and devices for the medical and automotive industries, and also provides haptic add-ons for many software products. A number of video demonstrations are provided on the site, including several on catheterization and endoscopy.

Ramona
http://www.kurzweilai.net/index.html?flash=2

This conversational avatar is Ray Kurzweil's alter ego. Click on the Ramona link, and install the Life FX Player to interact with Ramona. Ramona provides a good demonstration of several multimodal techniques, and illustrates how such approaches can heighten engagement.

Mobile Multimodal Access to City Help Demonstration
http://www.research.att.com/projects/ MultimodalAccessToCityHelp/

AT&T's Multimodal Access To City Help project is a working city guide and navigation system that enables users to access information, about places of interest, restaurants, and the subway for New York City using speech recognition technology, text-to-speech, handwriting recognition, and gesture recognition.

For Further Reading

The following articles and resources are recommended for those who want to learn more about multimodal interfaces and related topics.

MIT's Touch lab
http://touchlab.mit.edu/

> The goals of the Touch lab are to understand human haptics, develop machine haptics, and enhance human-machine interactions.

Visual Interaction
http://oxygen.lcs.mit.edu/Vision.html

> MIT's Project Oxygen focuses on technologies enabling pervasive computing. This article describes how visual interactions can be used to facilitate computer control.

Can Computers Communicate Like People Do?
http://www.nsf.gov/od/lpa/news/press/
pr9714.htm

> National Science Foundation description of a $10 million Speech, Text, Image and Multimedia Advanced Technology Effort (STIMULATE) to fund university researchers investigating improvement of interactions with computers.

Presence Research's List of Interface Technologies
http://www.presence-research.org

> This not-for-profit study group maintains a current listing of technologies and vendors.

Institute for Defense Analysis Review of Virtual Environment Interface Technology
http://www.hitl.washington.edu/scivw/
scivw-ftp/publications/IDA-pdf/TOC.PDF

> This report, while dated, provides an excellent overview of the range of interface technologies that have been in development for a number of years, including many that are now in production. Among those discussed in this in-depth report are visual, tracking, auditory, tactile, kinesthetic, full-body motion, and olfactory interface technologies.

Natural Language Software Registry
http://registry.dfki.de/

> The Registry provides a summary of sources and capabilities of academic, commercial, and proprietary natural language processing software.

CONTEXT-AWARE COMPUTING

Time-to-Adoption Horizon: Four to Five Years

These technologies are the subject of considerable research focus today, not only because they represent a logical next step in the evolution of computers, but because of their potential for security and related applications. Convergence of work in multimodal computing, global positioning, pattern recognition, and other areas is also fueling increased development in context-aware computing. While much of this work remains in the laboratory today, the underlying technologies are well-defined. Context-aware computing will likely find its way into broader use across university campuses within the next four to five years.

Overview

As computers become more ubiquitous, an obvious next step is for these devices to become more invisible and responsive to human needs. When humans talk to each other, context naturally plays a large role in our interactions. We react to facial expressions, body language, intonation, age, gender, and many other subtle and less subtle indicators. We know where we are, and have considerable information about our surroundings. We use all of this information all of the time, even in the most mundane situations. Computers, by and large, are not programmed to react to such contextual indicators, but that is changing.

Context-aware computing refers to computing devices that can interpret such contextual information and use it to aid decision-making and influence interactions. Contextual cues may include what the user is attending to, the user's location and orientation, the date and time of day, lighting conditions, other objects or people in the environment, accessible infrastructure in the immediate vicinity, and so forth. Context-aware devices and applications can make decisions based on such information without the need for user input.

Computing is increasingly embedded in everyday objects in ways that add this kind of functionality, and just as often, that computing power is not obvious. Nonetheless, many of the devices we use every day — cars, televisions, telephones, video cameras, even refrigerators — can already communicate with us and with each other on a rudimentary level. Computers embedded in motion sensors, for instance, trigger an alert when a delivery person is at the front door, simultaneously communicating this information

between the sensor and a control station. Cars that lock the doors when the car is underway provide a straightforward example that has been commonplace for a number of years. Lighting systems that turn on when motion is detected are another rudimentary case.

The emergence of inexpensive networking and next-generation sensors will move pervasive and context-aware computing forward rapidly over the next few years, and greatly expand its utility. Commercial buildings have been built for years with imbedded infrastructures that allow them to monitor structural and other components, and those technologies are now becoming part of "intelligent" homes. For example, in a wirelessly networked home, it is an easy matter to view pictures or video stored on a media server on any television or computer screen in the house. Likewise, it is a fairly simple matter to monitor home functions remotely via the web, viewing live video streams from security cameras, and even operating appliances and lighting while away from home.

Smart badge technology provides a means of locating individuals within a building by determining the location of a small device worn by personnel that transmits a unique infra-red signal that can be picked up by sensors placed throughout the workplace. Thousands of badges and sensors are deployed throughout several European universities including Cambridge University, the University of Kent, Imperial College, Lancaster University, and the University of Twente. In the US, Xerox PARC, DEc Research Laboratories, Bellcore and MIT Media Lab are all using similar systems. Such systems can be used

not only to locate individuals when needed, but also as a trigger to automatically direct relevant resources to a room in which a particular work team is located, and to set up a collaborative workspace for shared documents. Similarly, in a security conscious area, resources may be withheld if a person not having the required clearances is present.

"Smart dust" is a new wireless networking technology that will speed the development of context-aware computing tremendously over the next few years. Smart dust allows extremely simple collection of information about objects and conditions in the physical world and the simultaneous distribution of this information over data networks. Tiny sensors, called "motes," are about the size of a grain of rice, and contain wireless communication capacity that enables automated awareness and control of the surrounding environment. Because of their small size and generalizability, smart dust sensors will be easily incorporated into all kinds of every day objects. The possibilities are dazzling.

For example, alarm systems in residential-care facilities could use smart dust sensors to monitor the health of the residents in a building, tracking heart rates, blood oxygen levels and body temperatures, and relaying electronic alerts when the measurements rise above or fall below a pre-specified range. A smart-dust-enabled case of wine will "know" its temperature and humidity history, its handling history, and its owner's name and address, and all that information will be accessible via a wireless network. Medical equipment tagged with a smart dust locator could be dynamically tracked and located in times of need.

It is expected that technologies like smart dust will increasingly be built into mobile computing devices and cell phones to allow them to take advantage of contextual cues. A simple example is a cell phone which shuts down its ringer when a student enters a classroom or a theater, but it is very likely that such devices will become increasingly useful tools used to monitor all manner of processes and environments.

Relevance for Teaching, Learning, or Creative Expression

The technologies underlying context-aware computing are developing very rapidly, but applications specifically directed at teaching learning or creative expression are still on the horizon. Nonetheless, given the range of work going on, it is not hard to imagine what some of these applications may be.

- **Awareness of Social Settings.** Devices could adapt to the number of people in the room, changing resolution or font size if more than one user was present in front of the monitor. Music or video streams could be set to automatically pause if a voice or a conversation is detected.

- **Collaboration.** An application could be automatically shared if team members were perceived to be in the room or the immediate vicinity.

- **Smart Monitoring.** Smart dust could enable curators and preservationist to monitor the condition of artifacts or paintings, along with related ambient information such as ambient humidity, UV levels, and temperature.

- **Location Sensing.** Embedded wall sensors that monitor where a user is and what direction he or she may be facing could be used to provide way finding information in conjunction with campus maps.

- **Just-in-Time Information.** Sensor-reading devices could obtain detailed information about a piece of equipment, its current operating state, and operating instructions. A related application for museums could provide background about an artist, or a particular painting or sculpture.

- **Highly Personalized Instruction.** Instructional materials that can adapt to a user's emotional state or level of understanding might change in colors or tone if user appears frustrated, or additional help on the topic might be offered if appropriate.

- **Adaptive Use.** Devices that automatically adapt to user needs or conditions of use could make computing more "friendly." Examples are computer monitors that change the contrast of the display to match lighting conditions, or pattern recognition systems that recognize the outward indications of motor difficulties and trigger a switch to a verbal interface.

As with multimodal interfaces, the goal of context-aware computing efforts is to make computing at once less obtrusive and more responsive to the wide range of human needs. Although these technologies are largely in experimental stages at this time, development is taking place on so many levels and on so many fronts that it is very likely that universities will be using context-aware devices pervasively within the next four to five years.

Examples of Context Aware Computing

The following links will showcase some context-aware computing projects.

Smart Dust: Applications for Reliable Low Power Wireless Networks
http://www.dust-inc.com/applications/main.html

Dust, Inc. is a company formed to develop applications for low power wireless networks using tiny sensors the size of a grain of rice — "smart dust."

Context-Aware Computing MIT Media Lab
http://cac.media.mit.edu:8080/contextweb/jsp/projects.jsp

Descriptions of more than 20 common devices in which a level of computing ability or intelligence has been incorporated, including tables, cups, floors, beds, and other household items.

Input and Output, and Augmented Environments
http://depts.washington.edu/archi/0.context-aware/0.default.html

A variety of context-aware devices that are being designed and studied by the Department of Architecture, University of Washington, Seattle.

The Active Badge System
http://www.uk.research.att.com/ab.html

Developed by Olivetti & Oracle Research Laboratory in Cambridge, UK, the system uses computerized badges and networked sensors to locate people within a building.

For Further Reading

The following articles and resources are recommended for those who wish to learn more about context-aware computing and related topics.

Special Issue on Context-Aware Computing
http://hci-journal.com/editorial/si-context-aware-intro.pdf

Introduction to a series of articles published in Human-Computer Interaction in 2001.

Seeking a Foundation for Context-Aware Computing
http://www.dourish.com/embodied/essay.pdf

This article presents a conceptual framework for various forms of embodied interaction inherent in ubiquitous and context-aware computing.

Out of Context: Computer Systems That Adapt To, and Learn From, Context
http://lieber.www.media.mit.edu/people/lieber/Teaching/Context/Out-of-Context-Paper/Out-of-Context.html

MIT Media Laboratory paper on how context can be addressed in computing.

Understanding and Using Context
http://www.cc.gatech.edu/fce/ctk/pubs/PeTe5-1.pdf

Future Computing Environments Group at The Georgia Institute of Technology describes how context can be used in computing and describe methods by which this may be accomplished.

Context-Aware Computing Applications
http://seattleweb.intel-research.net/people/schilit/wmc-94-schilit.pdf

Paper presentation at an IEEE Workshop on Mobile Computing Systems and Applications describes how software can respond to context and mediate users' interactions with hardware, software, and other people.

KNOWLEDGE WEBS

Time-to-Adoption Horizon: Four to Five Years

The concepts and definitions behind knowledge webs continue to evolve and much of the work in this area is still experimental. Nonetheless, the first forays into developing prototype knowledge webs are already taking place, as the concept offers considerable promise as a way to make better use of greatly expanding knowledge, particularly as a discipline-specific resource. Knowledge webs will likely require four to five more years of refinement before finding their way into broader use across university campuses, but are definitely an area to watch.

Overview

Knowledge Web is a term that has seen a number of variations over recent years, but it is beginning to be applied to a group of convergent technologies and ideas based on a dynamic concept of individual and group knowledge generation and sharing, with technology used to make connections between knowledge elements clear, to distribute knowledge over multiple pathways, and to represent knowledge in ways that facilitate its use.

The emergence of a term to describe the concept is recognition of the convergence of work being done in digital libraries, mind mapping, communities of practice, and emerging technologies such as RDF, RSS, and data mining. Knowledge web development overlaps considerably with that going on around communities of practice, and holds considerable potential to help such communities share, create, analyze, validate, and distribute existing and emerging knowledge.

As knowledge webs are still evolving, the following scenario may help illustrate the concept:

A comprehensive set of excellent resources supports a community of scholars, researchers, and students, who cannot only easily find and use what they need, but contribute new materials and comment on existing ones. The relevance and validity of content is easily judged as it relates to the user's needs, and relationships to other resources and content are indicated in ways that facilitate serendipitous discovery. The resource is equally valuable to practitioners and learners, and facilitates dialog on a number of levels. Navigation is easy and intuitive, and supports the needs of the professional researcher, university students, and informal learners.

Knowledge webs are increasingly seen as a strategy to bring this sort of vision to everyday reality, and include processes and tools for gathering, validating, organizing, representing, navigating, and sharing knowledge. Knowledge webs are seen as particularly useful for disciplines and other communities of practice that need to manage not only existing knowledge, but streams of newly created knowledge.

The Chemical Abstract Service (CAS), for example, helps chemists and other scientists manage a vast collection of scientific knowledge that grows substantially every day — and about which they must stay abreast. CAS provides multiple pathways and tools to these professionals for accessing published research in the world's journal and patent literature for chemistry, life sciences, and a range of other scientific disciplines. Among the several notable features of this service are STN, which provides a federated search of more than 200 quality databases of evaluated information in science and technology, and the CAS Science Spotlight, which among other useful services, provides a real-time, real-usage ranking of its most requested documents, articles and patents derived from millions of requests from scientists around the world.

In such a collection, the amounts of information are vast and the content complex. Efficiently searching and navigating such a body of knowledge is the focus of a great many researchers. Knowledge representation is a developing body of work that aims to make navigating such a storehouse easier and more natural. Knowledge representations are visual depictions unique to specific knowledge domains

that are meant to show structure and relationships within the domain. Prototypical knowledge webs have relied on concept maps or 3-dimensional navigational tools as a way to make these mental models clear. These approaches provide easy and intuitive navigation and access to the knowledge elements that comprise these models.

Advances in knowledge representation and navigation are converging rapidly with other aspects of knowledge sharing and collaboration. Resource description frameworks (RDF), for example, provide a XML-based methodology for describing digital resources that is very flexible and machine-readable. RDF-based approaches to metadata allow information to be shared more easily, and important context to be transparently conveyed with it. Related to RDF is RSS (RDF site summary; also called "really simple syndication"), which provides a way for information posted on one web site to be incorporated into another automatically.

RSS provides a compelling simple way to share knowledge, especially new knowledge. The technology is an important component of weblogs (blogs), which use RSS to support communities of shared interests. A typical blog is the work of a single author, with comments from visitors, but as the technology evolves, community blogs have begun to emerge in forms that offer considerable promise for knowledge sharing within a community of practice.

The multi-author shared discussions that characterize community blogs are a natural fit with a more established technology — collaborative working environments. These tools are designed to help people work, sharing documents and resources, and developing ideas collaboratively. Each of these categories of tools offers promise in supporting the needs of a dynamic community.

Academic disciplines are increasingly reframing themselves in ways that mirror the development of communities of practice, and the convergence of the technologies described above make them an increasingly good fit with advances in this body of thought.

Relevance for Teaching, Learning, or Creative Expression

The vision of a space where knowledge sharing tools, knowledge generation tools, dialog and discussion tools, and quality assurance approaches come together is a compelling one for educators of all levels, but especially for higher academe. The tools are developing rapidly, and in fields like the sciences, which must manage a torrent of new knowledge every day, considerable progress has been made in a number of areas. Many of the professional societies are already using advanced knowledge management tools, and supporting dynamic communities around that knowledge.

Knowledge webs are increasingly seen as a natural way for scientific and other disciplines to evolve the ways they apply technology to their traditional roles, and to help them both manage the flow of emerging knowledge, and make new findings and ideas easily and intuitively available to faculty and students. The design of these webs is likely to be very much tied to the way the disciplines already organize themselves.

As the examples at the end of this section illustrate, knowledge web prototypes have assumed a number of forms and approaches, and such uniqueness is likely to remain an important attribute of such systems. Many scientific and engineering societies have already begun to organize themselves into prototypical knowledge webs, and as collections of digital resources are built, more disciplines will begin to move in that direction.

Knowledge Web Examples

As knowledge webs are a collection of ideas, tools, and concepts that is still in the development process, a considerable amount of variation exists in how the ideas are finding expression. The following links demonstrate some of the ways knowledge webs are evolving.

CAS Science Spotlight
http://www.cas.org/spotlight/
> CAS provides a real-time, real-usage ranking of its most requested documents, articles and

patents derived from millions of requests from scientists around the world. Also provided is a list of the "most intriguing" developments in chemistry. CAS scientists analyze over 200,000 documents per quarter, including articles from over 9,000 journals, and patents from more than 45 patent-issuing authorities from around the world. A document is characterized as "intriguing" if it contains new, novel or trendsetting scientific research that is likely to be of growing interest over time.

STN — Chemical Abstract Service's Scientific Information Service
http://stnweb.cas.org/

STN International is an online scientific and technical information service dedicated to meeting the information needs of scientists, engineers, and information professionals throughout the world. STN provides quick, direct access to more than 200 quality databases of evaluated information in science and technology.

James Burke's Knowledge Web Project
http://www.k-web.org/index.html

This video clip (requires Real Player) provides an in-depth overview of the K-Web project. Additional information is available at http://www.k-web.org.

World Music in Contemporary Life Knowledge Webs
http://trumpet.sdsu.edu/M345/knowledge_webs/knowledge_webs.html

Concept maps of complex relationships related to music.

Stephen Downes' list of RSS feeds on Educational Technology
http://www.downes.ca/cgi-bin/xml/topics.cgi

This service provides subject-specific RSS feeds culled from a collection of some 130 educational technology weblogs.

Science Blog
http://www.scienceblog.com/community/index.php

The Science Blog provides a range of current science news with an RSS feed, along with simple mechanisms for rating the articles presented, commenting, and forwarding the article via email. Links are provided to the source materials.

For Further Reading

The following articles and resources are recommended for those who wish to learn more about knowledge webs and related topics.

The Educational Promise of Knowledge Webs and Virtual Communities
http://www.tie.smartlibrary.org/NewInterface/segment.cfm?segment=2244

This article explores how knowledge webs help learners sift through the glut of information resources and make connections.

Seven Principles for Cultivating Communities of Practice
http://www.askmecorp.com/pdf/7Principles_CoP.pdf

This article by Etienne Wenger and Associates provides an excellent overview of current thinking about communities of practice.

National Science Digital Library
http://nsdl.org/

The National Science Foundation's digital resource center for science, technology, engineering, and mathematics.

THE HORIZON PROJECT

The NMC's annual *Horizon Report* details the current findings from the NMC's ongoing Horizon Project, a research-oriented effort that seeks to identify and describe emerging technologies likely to have a large impact on teaching, learning, or creative expression. The *2004 Horizon Report,* produced with support from the Corel Corporation, highlights six of these technologies that the research suggests will become very important to higher education within the next one to five years.

The project draws on an ongoing series of interviews with scientists, engineers, technologists, and other knowledgeable individuals in business, industry, and education to identify potentially promising technologies for the project's Advisory Board to consider.

Since the project launch in March 2002, the NMC has held an ongoing series of conversations with more than 400 technology professionals in the Silicon Valley, senior IT reps and faculty leaders from colleges and universities, and representatives of leading corporations. The "horizon" in the project's name refers to the time horizon expected before a technology is adopted for use by a significant number of colleges and universities.

The project considers three time horizons in its research: a very near-term horizon of 12 months or less before broad adoption; a mid-range horizon of 2-3 years; and a long-term horizon of 4-5 years. The project operates as a recurring cycle of five components: discovery, knowledge-gathering, vetting, exploration, and knowledge sharing.

Discovery. Potentially promising technologies are "discovered" (identified) through a series of interviews and conversations with technology professionals both within and outside academe. Qualitative research tools and techniques are used to generate an initial list of technologies that these groups and individuals feel may have promise for advancing teaching, learning, and creative expression in higher education. At this point, little is known, or even can be known, about the appropriateness or efficacy of these technologies for these purposes, as the Horizon Project expressly focuses on technologies not currently in widespread use in academe. In a typical year, 15-20 technologies may be identified in this phase.

Knowledge-Gathering. Enough information is gathered to allow the Advisory Board members to form an understanding of how each of the discovered technologies may be being used in settings outside of academe, of the potential the technology may have for higher education settings, and to envision applications of the technology for teaching, learning, and creative expression.

Vetting. (Assessment level one). The technologies are discussed and ranked by the Advisory Board for their potential relevance to teaching, learning, and creative expression. Only a handful of technologies which the Advisory Board ranks highly in this phase are carried forward.

Exploration. (Assessment level two). The potential applications of these important technologies are explored in this phase by higher education practitioners who are interested in new uses or new ways of thinking about how they might be used. Of particular interest is finding educational applications that may not be intuitive or obvious.

Knowledge Sharing. The experiences and understandings gleaned from the Horizon Project are shared in several ways, but the centerpiece of this phase is the annual *Horizon Report,* a document detailing recommendations for key emerging technologies to watch.

HORIZON PROJECT ADVISORY BOARD