Curriculum Vitæ

James D. Hollan

Contact

Mail	Department of Cognitive Science University of California, San Diego La Jolla, California 92093-0515
Email	hollan@cogsci.ucsd.edu
Web	http://hci.ucsd.edu
Phone	Office: 858.534.8156 Fax: 858.822.2476

Education

Stanford University

PostDoc (1973-1974) Artificial Intelligence Laboratory and Institute for Mathematical Studies in the Social Sciences

University of Florida

B.A. (1969), M.S. (1971), and Ph.D. (1973)

Employment

8/05 to 12/05	Visiting Professor, Computer Science Department, Stanford University
7/04 to Present	Adjunct Professor, Department of Computer Science and Engineering, University of California, San Diego
7/97 to Present	Professor, Department of Cognitive Science, University of California, San Diego
10/93 to 7/97	Professor and Chair, Computer Science Department, University of New Mexico
12/89 to 10/93	Director, Computer Graphics and Interactive Media Research Group, Bellcore
6/89 to 12/89	Visiting Scientist, Department of Cognitive Science, University of California, San Diego
2/87 to 6/89	Director, Human Interface Laboratory, Microelectronics and Computer Technology Corporation
9/77 to 2/87	Lecturer to Associate Research Cognitive Scientist and Director Intelligent Systems Group, Institute for Cognitive Science, University of California San Diego Director, Future Technologies, NPRDC, San Diego
9/74 to 9/77	Assistant Professor, Department of Psychology, Clarkson University, Potsdam, New York

Research Overview

My research goal is to understand the cognitive, computational, and social ecology of computationally-based media. It is motivated by two beliefs: (1) that cognitive science needs to move away from the notion of cognition as a property solely of isolated individuals to a distributed view of cognition as a property of larger social and technical systems and (2) that we are at the beginning of a paradigm shift in thinking about representational media, one that is starting to appreciate the importance of representations that are not only dynamic and interactive but that also adapt to the structure of tasks, the context of activities, and even our relationships with others.

My research interests span across cognitive ethnography, distributed and embodied cognition, human-computer interaction, multiscale information visualization, multimodal interaction, and software tools for visualization and interaction. My current work involves four intertwined activities: developing theory and methods, designing representations, implementing and evaluating prototypes, and understanding the broader space in which these activities are situated.

Summary of Past Research

In the first part of my career I explored dynamic graphical representations to support simulation-based training. This work resulted in one of the first object-oriented graphics editors and a series of seminal training systems. The science that accompanied the development efforts made significant contributions to the understanding of direct manipulation interfaces and played an influential role in initiating mental models research.

The next phase of my research focused on multimodal interfaces to high-functionality systems. I led the Human Interface Lab at MCC, the largest HCI research lab in the world, in designing and building a multimodal interface prototyping environment. We were among the first to demonstrate integration of gestures, graphics, and natural language within a common interface development framework. One significant contribution was a hybrid software architecture that combined neural networks with symbolic representations using an integrated knowledge base. Other work begun at MCC on history-enriched digital objects and collaborative filtering continued when I moved to Bellcore. This work resulted in a patent and a series of early demonstrations of the effectiveness of collaborative filtering. Collaborative filtering is one of the technologies that has made Google search so extremely effective.

At Bellcore I started a computer graphics and interactive media research group to explore information visualization. Among other efforts, I initiated and led the first large scale project to explore multiscale information visualization. When I moved to the University of New Mexico and subsequently returned to UCSD, this became an expanded multi-institutional (Bellcore, University of New Mexico, New York University, University of Maryland, University of Michigan, and UCSD) effort that enabled the first exploration of zoomable multiscale interfaces. The resulting system, Pad++, has been widely used by the research community and was licensed nonexclusively to Sony for $500,000.

My work on multiscale interfaces and visualization has continued, focusing primarily on information navigation of complex web-based domains, personal collections of scientific documents, and tools to assist analysis of video and other time-based activity data. Supported most recently by an NSF grant and funding from Intel, we implemented Dynapad, the third generation of our multiscale visualization software. The approach views interface design as the creation of a physics for information that is specifically designed to exploit our perceptual abilities, reduce cognitive costs by restructuring tasks, and increase the efficacy and pleasure of interaction.

Summary of Current Research

Ed Hutchins and I co-direct the Distributed Cognition and Human-Computer Interaction Laboratory (`http://hci.ucsd.edu`). We are united in the belief that distributed cognition promises to be a particularly fertile framework for understanding cognitive, social, and technical systems. A central image for us is environments in which people pursue their activities in collaboration with the elements of of the social and material world. Our core research efforts are directed at understanding such environments: what we really do in them, how we coordinated our activity in them, and what role technology should play in them.

Work in our lab contributes to the current shift in cognitive science toward a view of cognition as a property of systems that are larger than isolated individuals. This extends the reach of cognition to encompass interactions between people as well as interactions with resources in the environment. We are dedicated to developing the theoretical and methodological foundations engendered by this broader view of cognition and interaction.

My current research includes the following projects:

Analyzing Activity Dynamics

Recently in collaboration with Ed Hutchins and Javier Movellan I obtained NSF funding to explore a multiscale framework for capturing and analyzing activity dynamics.

The big question that motivated the proposal is: *What conditions can facilitate rapid advances and breakthroughs in behavioral science to rival those seen in the biological and physical sciences in the past century?*

We argued that the emergence of cognitive science and the converging view across multiple disciplines that human behavior is a complex dynamic interaction among biological, cognitive, linguistic, social and cultural processes are important first steps. In addition, we argued that while empirical and theoretical work is rapidly advancing at the biological end of this continuum, understanding such a complex system also necessitates data that capture the richness of real-world human activity and analytic frameworks that can exploit that richness.

In the history of science, changes in technologies for capturing data, as well as those for creating and manipulating representations, have often led to significant advances. The human genome project, for example, would have been impossibly complex without automatic DNA sequencing. Recent advances in digital technology present unprecedented opportunities for the capture, storage, analysis, and sharing of human activity data. Researchers from many disciplines are taking advantage of increasingly inexpensive digital video and storage facilities to assemble extensive data collections of human activity captured in real-world settings. The ability to record and share such data has created a critical moment in the practice and scope of behavioral research. The main obstacles to fully capitalizing on this opportunity are the huge time investment required for analysis using current methods and understanding how to coordinate analyses focused at different scales so as to profit fully from the theoretical perspectives of multiple disciplines.

Our plan is to integrate video and multiscale visualization facilities with computer vision techniques to create a flexible open framework to radically advance analysis of time-based records of human activity. We are combining automatic annotation with multiscale visual representations to allow events from multiple data streams to be juxtaposed on the same timeline so that co-occurrence, precedence, and other previously invisible patterns can be observed as analysts explore data relationships at multiple temporal and spatial scales. Dynamic lenses and annotation tools will provide interactive visualizations and flexible organizations of data.

The goals of our project are to (1) accelerate analysis by employing vision-based pattern recognition capabilities to pre-segment and tag data records, (2) increase analysis power by visualizing multimodal activity and macro-micro relationships, and coordinating analysis and annotation across multiple scales, and (3) facilitate shared use of our developing framework with collaborators.

This new work builds on our long term commitment to understanding cognition *in the wild*, developing multiscale visualizations, and recent experience automatically annotating video of freeway driving. We are extending the theory and methods developed in our earlier work and integrating them with new web-based analysis tools to enable more effective analysis of human activity. As initial test domains we are focusing on understanding activity in high-fidelity flight simulators and the activity histories of workstation usage and the process of writing. We also plan to evaluate a novel technique to assist in reinstating the context of earlier activities. Our long range objective is to better understand the dynamics of human activity as a scientific foundation for design.

Collaborations with TDLC, LIFE, and Calit2

I recently received an augment to the current NSF grant to build on an existing collaboration with Temporal Dynamics of Learning Center (TDLC) and with Falko Kuester, who is leading HIPerSpace display-wall visualization developments at Calit2, to reestablish an earlier informal collaboration on Diver video annotation facilities with Roy Pea of Stanford. Stanford is a partner along with U. Washington and SRI International in the Science of Learning Center called LIFE (Learning in Informal and Formal Environments). We thus have an interesting opportunity to explore a linkage between these two NSF Science of Learning Centers.

Paper Augmented Digital Documents

For several decades, experts have predicted that the advent of more powerful and compact computers would result in creation of paperless offices. Yet, as Sellen and Harper point out in "The Myth of the Paperless Office", consumption of paper is on the rise, and with few exceptions, office work still relies heavily on paper.

Ethnographic work makes clear that the root of this apparent paradox is the tension between the affordances provided by printed and digital documents: (1) *Paper documents* are easy to annotate and provide large, inexpensive, high-resolution display surfaces either by using large sheets of paper or by arranging several documents together. They are light, flexible, and convenient to carry and manipulate. (2) *Digital documents* are easy to edit, index, search, and transform using increasingly powerful computers. They also can provide instant access to large repositories of information on the web. Their intangibility make them inexpensive to store, duplicate, and distribute, thus facilitating sharing.

For the last several years I have been collaborating with Francois Guimbretière, from the University of Maryland, to explore Paper Augmented Digital Documents (PADD). In this approach the digital and paper worlds are on equal footing: paper and computers are simply two different ways to interact with documents during their life cycle. When paper affordances are needed, a document is retrieved from the database and printed. The printer acts as a normal printer but adds a pen-readable pattern to each document. Using an experimental digital pen, the document can now be marked like a normal paper document. The strokes collected by the pen are combined with the digital version of the document using a software system we are developing. The resulting augmented document can then be edited, shared, archived, or participate in further cycles between the paper and digital worlds. The PADD prototype system is currently running in Guimbretière's lab at Maryland and in my lab at UCSD.

We recently developed PapierCraft, a gesture-based command system that allows users to manipulate digital documents using paper printouts as proxies. Using an Anoto digital pen, users can draw command gestures on paper to tag a paragraph, email a selected area, copy selections to a notepad, or create links to related documents. Upon pen synchronization, PapierCraft executes the commands and presents the results in a digital document viewer. Users can then search the tagged information and navigate the web of annotated digital documents resulting from interactions with the paper proxies. PapierCraft also supports real time interactions across mix-media, for example, letting users copy information from paper to a Tablet PC screen. In a recent journal article we described the design and implementation of the PapierCraft system and user feedback from initial use.

We are in the process of extending this research to some new technology that allows simultaneous capture of pen strokes and audio. This enables one to touch paper notes with a digital pen and hear what was being said at the time those notes were taken. With the addition of PapierCraft's gesture-based command system we expect to be able to enable a wide collection of interesting new applications. I expect this to become a major component of the Digital Ethnographer's Workbench we are building as part of the new NSF project mentioned above.

UCSD: Ubiquitous Computing and Social Dynamics

The miniaturization of commodity computing devices is giving rise to a demand for wide-scale ubiquitous applications that can enrich our world with computation. This presents enormous new cognitive, social, and technical challenges. Not only is the monolithic "computer" being unbundled into fragmentary, appliance-like components, but experience in constructing the first generation of applications reveals that we have much to learn about designing such systems such that they mesh with and improve real-world activities, and much to teach our students that we don't yet know how to teach.

In collaboration with Bill Griswold in Computer Science and Adriene Jenik in Visual Arts I helped to create the UCSD ActiveCampus project. With funding from Microsoft and HP we explored the problem and opportunity of sustaining community through mobile wireless technology. A major thrust was the development of infrastructure for community-oriented ubiquitous computing. ActiveClass supports classroom activities such as anonymous asking of questions, polling, and student feedback. ActiveCampus Explorer supports several location-aware applications, including location-aware instant messaging and maps of the user's location annotated with dynamic hyperlinks of nearby buddies, digital graffiti, and other information.

This collaboration has evolved into the Ubiquitous Computing and Social Dynamics Research Group and now also includes Barry Brown from Communications and Louise Barkhuus from Computer Science. We have received funding from Microsoft, Intel, and HP to explore applying ideas from ActiveCampus to a new project on the Campus of the Future. (See `http://activecampus2.ucsd.edu/drupal` for details and this and other projects.)

Multimodal Interaction

I have recently initiated a series of projects to explore multimodal interfaces. This includes a project with Kevin Li, a computer science graduate student I advise jointly with Bill Griswold, and with Patrick Baudisch from Microsoft Research, on tactile feedback to allow devices to communicate with users when visual and auditory feedback are inappropriate.

Since current vibrotactile feedback is abstract and not related to the content of the message, this often clashes with the nature of the message. We are extending the repertoire of haptic notifications to including tapping and rubbing. Unlike traditional vibrotactile feedback, tapping and rubbing convey a distinct emotional message, similar to those induced by human-human touch. To enable these techniques we built a novel device we call soundTouch. It translates audio wave files into lateral motion using a voice coil motor found in computer hard drives. SoundTouch can produce motion from below 1Hz to above 10kHz with high precision and fidelity.

In collaboration with my graduate student Anne Marie Piper we have implemented a Shared Speech Interface (SSI), an application for an interactive multitouch tabletop display, for an experimental DiamdondTouch table, designed to facilitate medical conversations between a deaf patient and a hearing, nonsigning physician. In the development we involved members of the deaf community as well as medical and communication experts. Our initial evaluation compares conversation when facilitated by: (1) a digital table, (2) a human sign language interpreter, and (3) both a digital table and an interpreter. Our research reveals that tabletop displays have valuable properties for facilitating discussion between deaf and hearing individuals as well as enhancing privacy and independence.

Jan Borchers, a Professor at Aachen University and head of the Media Computing Group, spent a sabbatical with me this past year. We initiated a collaboration on FTIR multitouch interfaces. We have constructed an FTIR table in the DCOG-HCI lab and I am collaborating with Falko Kuester to build a similar table to control the huge HIPerSpace display-wall at Calit2.

In addition, I am collaborating with Jan, and students from each of our labs, on constructing a novel interface technology for multitouch tables we call SLAP (for Silicone Illuminated Active Peripherals). While multitouch tabletop technology has spread rapidly in the research community, these tables do not provide standard UI controls with tactile feedback. For example, users cannot feel an on-screen button before pressing it. In particular, typing on on-screen keyboards is difficult without looking, interrupting visual activities and increasing typing mistakes. SLAP peripherals, resolve this issue by synchronizing tangible and intangible representations. They are simple flexible physical widgets cast from translucent silicone and other materials that can be placed on a tabletop. Rear projection allows us to dynamically alter the appearance and function of a keyboard, slider, dial, lens, or other component. Input is sensed using FTIR and a camera below the surface of the table.

Project GreenLight

In July 2008 we received a large NSF equipment grant ($2.6 million) for development of instrumentation for Project GreenLight. Tom DeFanti of Calit2 is the PI. The information technology industry consumes as much energy and has roughly the same *carbon footprint* as the airline industry. Tom has assembled a collection of UCSD researchers and engineers to build an instrument to test the energy efficiency of computing systems under real-world conditions – with the ultimate goal of getting computer designers and users in the scientific community to re-think the way they do their jobs. My role in the project is to examine computer users' mental models of energy consumption and explore how visualization might influence their behavior.

Here is a brief list of other recent reserach projects:

- Asynchronous Negotiated Access
  A continuing research focus is on the myriad problems arising from people being constantly available for interaction because of increased connectivity via the internet and associated wireless technologies (e.g. pagers, cell phones, email, and instant messaging). There is a crucial need for practical methods to negotiate access and avoid unwanted interruptions. Scott Stornetta and I have devised a very promising method of asynchronous negotiated access. It mitigates problems in arranging access to people and coordinating information sharing. It supports asynchronous interaction and provides flexible boundaries between less urgent and more urgent access. In addition, it allows control over timing of access and permits tailoring of access level for specific individuals and groups. We have applied for a patent and the university has agreed to allow assignment to a nonprofit organization.

- Multiscale Information Visualization with Dynapad

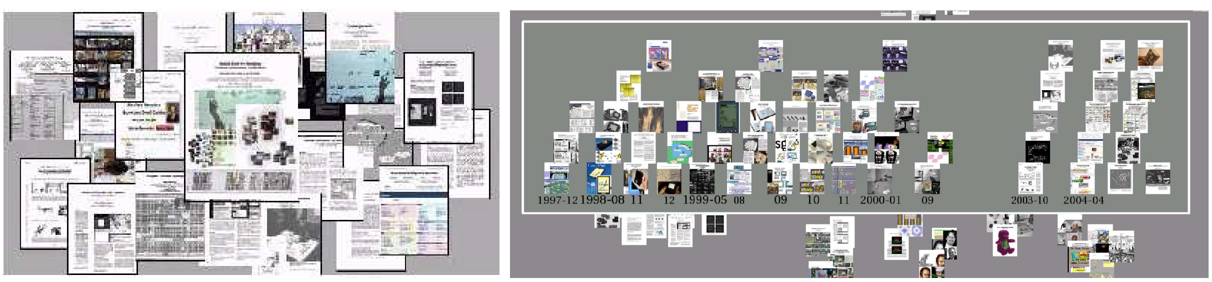

Figure 1: *On the left are sample Dynapad portraits (augmented thumbnails in which the first page from a paper is merged with a montage of images from it) of a paper and its references. On the right a timeline lens is positioned over document piles in a Dynapad workspace. The lens provides a temporary chronological ordering (by either PDF-creation date or date of entry into an individuals paper collection) of the documents.*

Dynapad is the third generation of our multiscale interface and visualization software. The goal of the effort is to understand the cognitive strategies people use in managing information collections in visual workspaces and how to design multiscale representations and a versatile infrastructure of tools to support those strategies. Dynapad currently supports multiscale organization and management of collections of digital photos and iconic representations of articles. It is the basis for our NSF supported work on image-based access and organization of information and of our Intel supported work exploring an informational physics approach to human-computer interaction. This approach views interface design as creation of a physics of appearance and behavior for informational entities. The long term vision is of dynamic entities operating according to multiple informational physics specifically designed to exploit our perceptual abilities, reduce cognitive costs by restructuring tasks, and increase the efficacy and pleasure of interaction.

- RUFAE: An Augmented Environments Research Network
  I was one of the founders of RUFAE (`http://www.rufae.net`). RUFAE is an international network of research laboratories that are designing and studying interactive spaces. RUFAE's industrial and academic research labs, from France, Germany, Sweden, Russia, and the U.S., cover architecture, cognitive science, computer science, and psychology. The collaboration currently manifests itself in three main aspects: (1) Spaces: Each member lab is creating and using a prototypical augmented space. (2) Ideas: In monthly phone/video/application sharing conferences between our augmented spaces, we present our research, discuss common projects, and share technical expertise. This helps to improve our spaces, their interconnectedness and compatibility, and to establish good practice through real use. (3) People: We have begun to exchange students and researchers between member labs, to transfer technology, and to work on common research projects.

Research Funding

Current Funding

NSF *A Multiscale Framework for Analyzing Activity Dynamics($875K, 9/2007- 8/2010)*

I am the PI on this project. Along with Co-PI Ed Hutchins, we are integrating video and multiscale representation facilities to create a flexible open framework to radically advance analysis of time-based records of human activity. We combine automatic annotation of activity data and improved support for research practices with the development of multiscale visual representations to better understand the cognitive ecology of human activity. These representations and linkages allow events from multiple data streams to be juxtaposed on the same time line so that co-occurrence, precedence, and other relationships can be observed. Previously invisible patterns can emerge as the analyst explores data relationships at multiple temporal and spatial scales. Dynamic lenses provide interactive transformations of data representations. Automated vision-based pattern recognition capabilities pre-segment and tag searchable data records, saving time by accelerating analyses.

Our long-term goal is to better understand the dynamics of human activity as a scientific foundation for design. .The work builds on our long term focus on understanding cognition Òin the wildÓ, developing multiscale visualizations, and on recent work on automatic annotation of freeway driving data collected in an automobile instrumented with ten video cameras. We are extending the theory and methods developed in our earlier work and intergrating them with exisiting tools to enable more effective analysis of human activity. Specifically our focus is on understanding activity in flight simulators and the activity history of normal workstation usage. In addition, we will explore a novel technique to assist people in reinstating the context of their earlier activities.

Recent Research Funding

Nissan Research Center and UC Digital Media Innovation Program *Human-Centered Intelligent Driver Support Systems: A Novel Multimodal "Driving Ecology" for Enhanced Safety ($1.3M, 05/10/02 – 05/09/05)*

I am one of the PIs for a multidisciplinary research project sponsored by Nissan Research Center, Japan and the University of California Digital Media Innovation Program. The project involves a collaboration between the Department of Cognitive Science, the Department of Electrical and Computing Engineering, and the Department of Psychology. The scientific goal is to understand the cognitive ecology of driving and design instrumentation and controls to improve driver safety.

We have instrumented a car with 10 video cameras as well as access the state of the car's controls and instrumentation. This provides us with a unique rich data set of digital recordings of driving behavior. This includes a first-person video from a small camera mounted on glasses or in a headband worn by the driver (3rd Eye). A computer in the car's trunk is equipped to digitize and record multiple video and audio streams on a array of disks. The video and audio data streams are synchronized. In addition, timestamped data from car controls and instruments provides details of steering, breaking, acceleration, as well as other parameters of the car's operation. The data is synchronized and linked to GPS data and notes made during driving sessions.

My specific focus is on ethnographic studies to explore sense making of drivers and how cognitive factors influence driving behavior. We are particularly interested in identifying the meaningful distinctions drivers make of their driving experience. To investigate this we are collecting data as drivers follow prescribed routes designed to allow us to see how they behave in a variety of driving conditions. During these drives we collect think-aloud protocols and notate interesting events relevant to our interests or to those of our collaborators (e.g., situations that may be challenging for automatic identification from sensor data).

I am particularly encouraged by our success in using computer-vision algorithms to automatically annotate video to assist analysis. For example, we automatically compute the lateral angular velocity of the head from the 3rd Eye camera video. This allows identification of small head position adjustments, glances to the rear view mirror, glances to the left or right side mirror, and very clearly shows large over-the-shoulder head movements. Such annotations accelerate and ease analysis. Being able to access the data in terms of meaningful characteristics rather than low-level video timecode is invaluable. We also threshold the amplitude of recorded audio to automatically index times when someone is speaking in the car. Foot motion and lateral foot position are extracted from a "Foot-Cam" video using another simple detection algorithm. In combination with recordings of brake pedal pressure this enables easily determining, for example, when drivers move their foot to the brake pedal in preparation for braking. We have also developed code to determine where the hands are positioned on the steering wheel and to automatically compute lateral position of the car as a basis for detecting lane changes.

NSF Grant *Image-Based Access and Organization of Information (NSF ITR $325K, 8/15/01 – 7/31/05)*

There is an astonishing amount of information on the web and it is constantly increasing. To avoid being overwhelmed by the volume of information available and confused by its uneven quality, people need assistance in efficiently finding task-relevant information and in effectively managing complex dynamic information collections. Current interfaces primarily employ textual representations for accessing and organizing information collections. Access is either via taxonomies or queries to search engines and results are typically organized as lists or hierarchies of web page titles. Given the ability of images to assist memory and the common exploitation of space in everyday problem solving to simplify choice, perception, and mental computation, it is surprising that so little use is made of images and spatial organization to aid information access and organization.

In this project we are examining whether spatial and temporal organization of images can serve as effective interface components and attempting to determine what advantages they offer over textual lists of titles and URLs. We are developing and evaluating a series of applications: visual browsing session summaries, visual augmentation of search engine results, hierarchical cache visualizations, and activity visualizations. We are also exploring a variety of image-based techniques: image selection to provide access, methods for choosing images that effectively represent web page content, and algorithms for scaling images. Our primary focus has been on the development and exploration of flexible lightweight multiscale piles for information management and the use lenses to provide multiple views of the structure of a document workspace.

NSF Grant *A Distributed Cognition Approach to Designing Digital Work Materials For Collaborative Workspaces ($1.6M 07/01/99 – 06/30/02)*

Ed Hutchins and David Kirsh are co-PIs with me on this project. Our work is motivated by the belief that effective design of digital workplaces and work materials requires a new theoretical perspective that views work materials as more than stimuli for disembodied cognitive systems. We are exploring a distributed cognition perspective in which work materials become elements of the cognitive system itself, and cognition becomes an emergent property of the interactions among people and work materials. This work led to the development of Dynapad.

Our primary research interest is in understanding annotation in domains ranging from collaborative scientific research to commercial aviation. In each domain, we are conducting ethnographic and experimental studies as a basis for understanding how to better design new digital work materials. My focus in the project is on developing a series of prototypes of history-enriched work materials and multiscale information visualization facilities to support collection and selective sharing of personnel and group activity histories and annotations.

Intel Gift *Active Multiscale Information ($100K 07/01/98 – present)*

A wonderfully flexible Intel gift funding continues to help support my work on multiscale software that was originally developed under DARPA funding. In this project we are pursuing an informational physics approach to human-computer interaction. This approach views interface design as creation of a physics of appearance and behavior for informational entities. It is part of a larger research program that led us to develop techniques to make task-relevant characteristics of information perceptually available and create active work materials that record and advertise salient aspects of the context of their use. The long term vision is of dynamic entities operating according to multiple informational physics specifically designed to exploit our perceptual abilities, reduce cognitive costs by restructuring tasks, and increase the efficacy and pleasure of our interaction. Based on this informational physics approach we are using Dynapad to develop active multiscale interfaces to aid information navigation of complex web-based domains, personal information collections, and activity data.

Prior Research Funding

DARPA Grant *Beyond Imitation: A Strategy for Building a New Generation of HCI Design Environments ($3.7M)*

I was the Principle Investigator on this research contract awarded when I was at the University of New Mexico. The contract moved with me when I came to UCSD. As part of the effort there were subcontracts to Ken Perlin at New York University, to George Furnas at Bellcore and later at the University of Michigan, and to Ben Bederson at the University of Maryland.

This was a very successful research project. We developed Pad++, a software substrate to allow exploration of zoomable multiscale interfaces and visualizations. Pad++ provides a radical alternative to traditional interfaces and is the first serious exploration of zoomable multiscale interfaces. The project produced over 30 publications from our research group. Pad++ has been used in a wide range of application domains and continues to be available and maintained. It has been downloaded from our web site by over 4500 people around the world. Example applications include: map-based information systems for command and control, our PadPrints system that provides multiscale graphical histories of web traversals, the MIT intelligent room prototype, the Xerox Parc fluid documents prototype, and multiscale visualizations of the human genome. The Pad++ software was licensed to Sony for $500,000.

NSF Grant *Effective Information Access: Computer Science Research Fundamental to Creation of an National Information Infrastructure ($1.25M)*

I was the Principle Investigator on this CISE research infrastructure award at the University of New Mexico. It provided support for visualization facilities, servers to support distributed simulation, and a high-speed cluster network. As part of this effort I established two new laboratories in the department. The grant also supported research collaborations with the Santa Fe Institute and the National Laboratories at Sandia and Los Alamos.

We made significant progress in several areas: (1) immunological modeling and construction of a computer immune system, (2) development of a multiscale web browser, (3) refinement of ccr, an evolvable metasystem that allows the group members themselves to control many aspects of the computation and communications system that hitherto have been fixed by a priori software design, (4) development of AR-NET, a network of collaborating automated reasoning processes, (5) porting the communication and local resource management features of the Puma operating system, then running on the Intel Paragon supercomputer at Sandia, to a collection of workstations running Linux and connected by high speed networks (Myrinet and ATM).

Bellcore *Computer Graphics and Interactive Media Research Group*

I established the Computer Graphics and Interactive Media Research Group at Bellcore. Research focused on information visualization and construction of 3D visualization and interface prototyping environments. Projects included unified graphical interfaces to heterogeneous databases, visualization of network and switching activity, visualization of software systems and programmer activities, information filtering, prototyping and exploration of interactive animations, and empirical studies of history-enriched digital objects (Bellcore Video Recommender, one of the earliest demonstrations of the effectiveness of collaborative filtering). Additional efforts were concerned with theories of telecommunications and exploration of alternatives to imitating face-to-face interactions for supporting informal communication.

MCC *Human Interface Laboratory*

As Director of the Human Interface Laboratory (annual budget $5M) at MCC, I coordinated the efforts of approximately 40 researchers. Areas of research were graphics, knowledge editing, natural language, neural networks, computer supported cooperative work, and new metaphors for interaction design. Our goal was to develop the foundations for principled and efficient construction of collaborative interfaces to high-functionality systems. Research within the laboratory was coordinated around the construction of an integrated interface prototyping environment and its application to challenging interface problems. The vision was to evolve a set of human interface tools (HITS) into a general user interface design environment (GUIDE). HITS and GUIDE were experimental vehicles for grounding, motivating, and coordinating the lab's scientific and technological efforts. They served as prototypes supporting the rapid implementation, exploration, and demonstration of new human interface concepts.

ONR *Intelligent Systems Group and Future Technologies Group*

In my earlier work at UCSD, in collaboration with Ed Hutchins and Don Norman, I served as Director of the Intelligent Systems Group. Our research group was concerned with application of artificial intelligence and cognitive science to the design of human computer interfaces and development of graphical simulation-based training systems. At NPRDC I was head of the Future Technologies Group and in collaboration with Ed Hutchins and Michael Williams led efforts to build advanced training systems (Moboard, Semnet, and Steamer). I was PI on a number of research projects: Theory of Graphic Representation, Declarative and Procedural Representation, Steamer: An Advanced Intelligent Computer-Assisted Instruction System (in collaboration with Larry Stead, Bruce Roberts, and Al Stevens at BBN), Qualitative Interfaces to Quantitative Process Models, AI-Based Tools for Building Simulations, and Computation via Direct Manipulation.

Software Systems

A major portion of my intellectual activity is devoted to the design and implementation of software systems. Such systems are fundamental to my research. I find creating software and sharing it with the students and the wider research community to frequently have a more significant impact than traditional forms of academic publication. Software is an artifact that can mediate very productive interactions and collaborations.

Pad++ *Zoomable Multiscale Visualization Software*

I led the effort to develop Pad++. This software has made possible the first serious exploration of multiscale interfaces. It consists of 164,714 lines of code and was recently licensed to Sony for $500,000. My main collaborator has been Ben Bederson. The development of the software benefitted from interactions with Ken Perlin and Jon Myer at New York University, George Furnas at the University of Michigan, and feedback from a surprising number of the 4500 people who downloaded the code.

STkPad *Scheme-Based Zoomable Multiscale Visualization Software*

STkPad is a version of Pad++ developed at UCSD. It consists of 97,737 lines of code. It's main features are the replacement of the Tcl scripting language used in Pad++ with Scheme and integration of MySQL, a relational database. The database serves to store STkPad content and to provide a mechanism to support collaborative applications. The research focus for STkPad is to explore shared activity histories and image-based information navigation. I developed it in collaboration with two recent postdocs, David Fox and Ron Hightower.

Dynapad *Scheme-Based Zoomable Multiscale Visualization Software*

Dynapad is the third generation of our multiscale interface and visualization software. The name Dynapad was chosen to reflect the software's heritage from our earlier Pad++ and STkPad software as well as ideas from Dynabook and Sketchpad. It makes scale a first-class parameter of objects, supports navigation in multiscale workspaces, and provides especially effective mechanisms to maintain interactivity while rendering large numbers of Tabletop Displays for Small Group Study: Affordances of Paper versus Digital Materials Anne Marie Piper, University of California, San Diego James D. Hollan, University of California, San Diego graphical objects. Dynapad employs Scheme to provide a high-level programming interface to the multiscale graphical and interaction facilities in the C++ rendering substrate.

Dynapad implements multiscale graphical objects that are interactive (e.g., they can be scaled or moved via user interaction) and dynamic (e.g., they can have behaviors that result from the running of attached code). Behaviors can be associated with an object, a set of objects, or a region of the multiscale workspace and are triggered by user actions, the behavior of other objects, various events, or timer interrupts. It is build using PLT Scheme (`http://www.plt-scheme.org`) and is the basis for our exploration of informational physics and our current NSF supported work on piles and lenses to support image-based access and organization of information.

Publications

Anne Marie Piper and James D. Hollan. *Tabletop Displays for Small Group Study: Affordances of Paper versus Digital Materials*, CHI2009, in press.

Malte Weiss, Julie Wagner, Roger Jennings, Yvonne Jansen, Ramsin Khoshabeh, James D. Hollan, Jan Borchers, *SLAP Widgets: Bridging the Gap Between Virtual and Physical Controls on Tabletops*, CHI2009, in press.

Malte Weiss, Roger Jennings, Julie Wagner, Ramsin Khoshabeh, Jan Borchers, and James D. Hollan. *SLAP: Silicone Illuminated Active Peripherals*, IEEE Tabletops and Interactive Surfaces, October 1-3, 2008, Amsterdam, the Netherlands

Chunyuan Liao, Francois Guimbretiere, Ken Hinckley, and James D. Hollan. *Papiercraft: A Gesture-Based Command System for Interactive Paper*, ACM Transactions on Computer-Human Interaction, 14:4, 1-27, 2008.

Timothy Sohn, Kevin A. Li, William G. Griswold, and James D. Hollan. *A Diary Study of Mobile Information Needs*, CHI2008: Proceedings of ACM Conference on Human Factors in Computing Systems, 433-442, 2008.

Kevin A. Li, Patrick Baudisch, William G. Griswold, and James D. Hollan. *Tapping and Rubbing: Exploring New Dimensions of Tactile Feedback with Voice Coil Motors*, UIST 2008: Proceedings of ACM Symposium on User Interface Software and Technology, to be held October 19-22, 2008. IN PRESS

Anne Marie Piper and James D. Hollan. *Supporting Medical Conversations Between Deaf and Hearing Individuals with Tabletop Displays*, CSCW 2008: Proceedings of ACM Computer Supported Cooperative Work, to be held November 1-10, 2008. IN PRESS

James D. Hollan and Edwin L. Hutchins. *Opportunities and Challenges for Augmented Environments: A Distributed Cognition Perspective*, In Designing User Friendly Augmented Environments: From Meeting Rooms to Digital Collaborative Spaces, Saadi Lahlou (Ed,), Springer, 2009.

Amaya Becvar, Jim Hollan, and Ed Hutchins, *Representational Gestures for Developing Theory in a Scientific Laboratory*, book chapter, In Artifacts in Workiplace Practice, Kluwer Academic Publishers , in press.

Saeko Nomura, Hiroshi Tamura, and Jim Hollan. *Information Management Centers in Everyday Home Life*, Proceedings of HCI International, 11th International Conference on Human-Computer Interaction, 2005.

Amaya Becvar, Jim Hollan, and Ed Hutchins, *Hands as Molecules: Representational Gestures for Developing Theory in Scientific Activity*, Semiotica, 156, 89-112, 2005.

Dan Bower, Pierre Fastrez, and Jim Hollan. *Spatial Tools for Managing Information Collections*, Proceedings of Hawaii International Conference on System Science, 2005.

Erwin R. Boer, Carrie A. Joyce, Deborah Forster, Monal Chokshi, Tayopa Mogilner, and James D. Hollan. *Bridging Ethnography and Engineering through the Graphical Language of Petri Nets*, 5th International Conference on Methods and Techniques in Behaviroal Research, The Netherlands, 2005.

Amaya Becvar and Jim Hollan, *Envisioning a Paper-Augmented Digital Notebook: Exploiting Digital Pen Technology for Fieldwork*, Poster, 9th European Conference on Computer-Supported Cooperative Work, poster, 2005.

Erwin R. Boer, Carrie A. Joyce, Deborah Forster, Jean-Baptiste Haue, Monal Chokshi, Tayopa Mogilar, Elaine Garvey, and James D. Hollan. *Mining for Meaning in Drivers' Behavior: A Tool for Situated Hypothesis Generation and Verification*, 5th International Conference on Methods and Techniques in Behaviroal Research, The Netherlands, 2005.

Ryan Y. Sit, James D. Hollan, and William G. Griswold, *Digital Photos as Conversational Acnchors*, Proceedings of Hawaii International Conference on System Science, 2005

Joel C. McCall, Ofer Achler, Mohan M. Trivedi, Jean-Baptiste HauÃl', Pierre Fastrez, Deborah Forster, and James D. Hollan, *A Collaborative Approach for Human-Centered Driver Assistance Systems*, Proceedings of IEEE Conference on Intelligent Transportation Systems, October, 2004.

Daniel Bauer, Pierre Fastrez, and Jim Hollan, *Computationally-Enriched "Piles" for Managing Digital Photo Collections*, Proceedings of Visual Languages and Human-Centric Computing, Rome, Italy, September 26-29, 2004.

Daniel Bauer and James D. Hollan. IRYS: A Visualization Tool for Temporal Analysis of Multimodal Interactions, *Proceedings of the ACM 5th International Conference on Multimodal Interfaces*, Vancouver, British Columbia, Canada. 285-288, 2003.

Wakefield Scott Stornetta and James D. Hollan, *Device and Method for Mediating Access*, US Patent Application 20020144136, 2002.

James D. Hollan, Edwin L. Hutchins, and David Kirsh. Distributed cognition: Toward A new theoretical foundation for human-computer interaction research. In J. M. Carroll, editor, *Human-Computer Interaction in the New Millennium*, 75–94. Addison-Wesley, 2001. (This is a revised version of our TOCHI article)

Jon I. Helfman and James D. Hollan. Image representations for accessing and organizing web information, in *Proceedings of the SPIE International Society for Optical Engineering Internet Imaging II Conference*, 91–101, 2001.

James Hollan and Scott Stornetta. Asynchronous negotiated access. In S. McDonald, Y. Waern, G. Cockton, editors, *Human Computer Interaction 2000*, 17–26, 2000.

James D. Hollan, Edwin L. Hutchins, and David Kirsh. Distributed cognition: A new theoretical foundation for human-computer interaction research. *ACM Transactions on Human-Computer Interaction*, 7, 174–196, 2000.

James D. Hollan. Human-computer interaction. In R. A. Wilson and F. C. Keil, editors, *Encyclopedia of the Cognitive Sciences*. 379-381, MIT Press, 1999.

Ron R. Hightower, Laura Ring, Jon Helfman, Benjamin B. Bederson, and James D. Hollan. Graphical multiscale web histories: A study of PadPrints. In *Proceedings of ACM Conference on Hypertext (Hypertext 98)*, 58–65. ACM Press, 1998.

Benjamin B. Bederson, James D. Hollan, Jason Stewart, David Rogers, David Vick, Laura Ring, E. Grose, and C. Forsythe. A zooming web browser. In C. Forsythe, J. Ratner, and E. Grose, editors, *Human Factors in Web Development*, 255–266. Lawrence Erlbaum, 1998.

James D. Hollan, Benjamin B. Bederson, and Jon Helfman. Information visualization. In M. G. Helenader, T. K. Landauer, and P. Prabhu, editors, *The Handbook of Human Computer Interaction*, 33–48. Elsevier Science, The Netherlands, 1997.

Allison Druin, Jason Stewart, David Proft, Benjamin Bederson, and James Hollan. KidPad: A design collaboration between children, technologists, and educators. In *Proceedings of ACM CHI 97 Conference on Human Factors in Computing Systems*, volume 1 of *Design Briefings: Understanding Users*, 463–470, 1997.

N. Wardrip-Fruin, Jon Meyer, Ken Perlin, Benjamin B. Bederson, and James D. Hollan. A multiscale narrative: Grey matters. In *ACM Visual Proceedings of Computer Graphics (Siggraph 97)*, 141. ACM Press, 1997.

Benjamin B. Bederson, James D. Hollan, Ken Perlin, Jon Meyer, David Bacon, and George Furnas. Pad++: A zoomable graphical sketchpad for exploring alternate interface physic. *Journal of Visual Languages and Computing*, 7:3–31, 1996.

Brad Myers, James Hollan,and Isabel Cruz. Strategic directions in human-computer interaction. *ACM Computing Surveys*, 28(4):794–809, 1996.

Benjamin B. Bederson, James D. Hollan, Jason Stewart, David Rogers, Allison Druin, and David Vick. A zooming Web browser. In *Proceedings of the SPIE — The International Society for Optical Engineering*, 260–271, 1996.

Benjamin B. Bederson, James D. Hollan, Allison Druin, Jason Stewart, David Rogers, and David Proft. Local tools: An alternative to tool palettes. In *Proceedings of the Ninth Annual Symposium on User Interface Software and Technology*, 169–170, New York: ACM Press, 1996. Press.

Benjamin B. Bederson, James D. Hollan, Jason Stewart, David Rogers, Allison Druin, David Proft, and Mohammad Ijadi. Pad++: A zooming web browser. In *Proceedings of ACM User Interface Software and Technology Conference (UIST 96).* ACM Press, 1996.

Benjamin B. Bederson and James D. Hollan. Pad++: A zooming graphical interface system. In *Proceedings of ACM CHI'95 Conference on Human Factors in Computing Systems*, volume 2 of *Demonstrations: Information Navigation/Usability*, 23–24, 1995.

William C. Hill and James D. Hollan. History-enriched digital objects: Prototypes and policy issues. *The Information Society*, 10:139–145, 1994.

Benjamin B. Bederson, Larry Stead, and James D. Hollan. Pad++: Advances in multiscale interfaces. In *Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systems*, volume 2 of *Short Papers: Virtual and Visual Environments*, 315–316, 1994.

Benjamin B. Bederson and James D. Hollan. Pad++: A zooming graphical interface for exploring alternate interface physics. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, 17–26, 1994.

William C. Hill and James D. Hollan. History-enriched digital objects. In *Proceedings of the ACM Conference on Computers, Freedom and Privacy*, 9.16–9.20. ACM Press, 1993.

Laurence Brothers, James Hollan, Jakob Nielsen, Scott Stornetta, Steve Abney, George Furnas, and Michael Littman. Supporting informal communication via ephemeral interest groups. In *Proceedings of ACM CSCW'92 Conference on Computer-Supported Cooperative Work*, 84–90, 1992.

James D. Hollan and Scott Stornetta. Beyond being there. In *Proceedings of ACM CHI'92 Conference on Human Factors in Computing Systems*, Perspectives on the Design of Collaborative Systems, 119–125, 1992.

William C. Hill, James D. Hollan, Dave Wroblewski, and Tim McCandless. Edit wear and read wear. In *Proceedings of ACM CHI'92 Conference on Human Factors in Computing Systems*, 3–9, 1992.

William C. Hill and James D. Hollan. Pointing and visualization. In *Proceedings of ACM CHI'92 Conference on Human Factors in Computing Systems*, 665–666, 1992.

James D. Hollan. Visualizing information. *Bellcore Exchange*, 1992.

Bill Curtis, Roy Kuntz, James Hollan, S. Joy Mountford, and George Collier. Retrospective on the MCC human interface laboratory. In *Proceedings of ACM CHI'91 Conference on Human Factors in Computing Systems*, Panels, 373–376, 1991.

James D. Hollan and Michael D. Williams. The craft of exploiting AI techniques in human interface design. In *American Association for Artificial Intelligence Tutorial Notes*, 1–90, 1991.

William C. Hill and James D. Hollan. Deixis and the future of visualization excellence. In *Visualization '91*, 314–320, 1991.

James D. Hollan, E. Rich, William C. Hill, D. Wroblewski, W. Wilner, K. Wittenburg, J. Grudin, and the MCC Human Interface Lab Members. An introduction to HITS: Human interface tool suite. In J. Sullivan, editor, *Intelligent User Interfaces*, pages 293–338. Addison Wesley, 1991.

James D. Hollan. User models and user interfaces: A case for domain models, task models, and tailorability. In *Proceedings of AAAI-90*, 1990.

James D. Hollan, Edwin L. Hutchins, T. P. McCandless, M. Rosenstein, and L. Weitzman. Graphical interfaces for simulation. In B. Rouse, editor, *Advances in Man-Machine Systems Research*, 129–163. 1987.

Edwin L. Hutchins, James D. Hollan, and Donald A. Norman. Direct manipulation interfaces. In Norman, D.A. and S. W. Draper, editors, *User Centered System Design: New Perspectives on Human-Computer Interaction*, 87–124. Lawrence Erlbaum Associates, 1986.

Henry. M. Halff, James D. Hollan, and Edwin L. Hutchins. Cognitive science and military training. *American Psychologist*, 1131–1140, 1986.

James D. Hollan. Review of artificial intelligence: The very idea by John Haugeland. *Contemporary Psychology*, 31:971–972, 1986.

James D. Hollan, L. Weitzman, and M. Rosenstein. Interface design for simulations. In *Proceedings of the First Annual Artificial Intelligence and Advanced Computer Technology Conference*, 357–368, 1985.

Edwin L. Hutchins, James D. Hollan, and Donald A. Norman. Direct Manipulation Interfaces. *Human-Computer Interaction*, 1:311–338, 1985.

James D. Hollan, Edwin L. Hutchins, and Louis Weitzman. Steamer: An interactive inspectable simulation-based training system. *AI Magazine*, 5(2):15–27, 1984.

James D. Hollan. Intelligent object-based graphical interfaces. In G. Salvendy, editor, *Human Computer Interaction*, pages 293–297, 1984.

James D. Hollan, Edwin L. Hutchins, M. Rosenstein, and L. Weitzman. Tools for graphical interface design. In G. Kohl and S. Nassau, editors, *Combining Human and Artificial Intelligence: A New Frontier in Human Factors*, 24–40. 1984.

Michael D. Williams, James D. Hollan, and Al Stevens. Human reasoning about a simple physical system. In D. Gentner and Al Stevens, editors, *Mental Models*, pages 131–153. Hillsdale, NJ: Lawrence Erlbaum Associates, 1982.

James D. Hollan, Michael D. Williams, and Al Stevens. An overview of Steamer. *Behavior Research Methods and Instrumentation*, 13:85–90, 1981.

Michael D. Williams and James D. Hollan. The process of retrieval from very long term memory. *Cognitive Science*, 5:87–119, 1981.

James D. Hollan A. Stevens and Michael D. Williams. Steamer: An advanced computer-assisted instruction system for propulsion engineering. In *Proceedings of Summer Computer Simulation Conference*, pages 400–404. Society for Computer Simulation, 1980.

James D. Hollan and Mark R. Wallen. The LNR Laboratory. *Behavior research methods and instrumentation*, 11:151–155, 1979.

James D. Hollan and Ken L. Bowles. An introduction to UCSD PASCAL: A (nearly) machine independent software system for micro and mini computers. *Behavior Research Methods and Instrumentation*, 10:531–534, 1978.

James D. Hollan and Ken L. Bowles. Microcomputer-based mass education. In *Proceedings of Conference on Micro-computers in Education and Training.*, 1978.

James D. Hollan. On the use of a cross-assembler, compiler (PL-11), and simulator with the LSI-11. In *Proceedings of the Fall Decus Computer Symposium.*, 1977.

James D. Hollan. Features and semantic memory: Set-theoretic or network model? *Psychological Review*, 82:154–155, 1975.

James D. Hollan and Robert E. Osteen. ProjectX: A program for processing acyclic digraphs. *Behavior Research Methods and Instrumentation*, 105–106, 1975.

James D. Hollan. A graph-theoretical analysis of Kelley's repertory grid. *Behavior Research Methods and Instrumentation*, 7:382, 1975.

James D. Hollan. Clarification of some effects of hierarchical organization upon free recall. *Psychological Reports*, 34:1227–1231, 1974.

James D. Hollan. Some effects of epistemological structure on memory. *Memory and Cognition*, 2:670–676, 1974.

James D. Hollan. A model of epistemological structure in mathematics. In W. L. Vogt and M. H. Mickle, editors, *Modeling and Simulation.* Pittsburgh School of Engineering, University of Pittsburgh, 1974.

Presentations and Invited Addresses

I have given hundreds of talks and invited addresses over the years. A list is available on request.

Teaching

My primary teaching goal is to excite students about the topics we are covering and provide them with the strong foundation needed to do good work. In addition, I always attempt to expose students to the intellectual history of key ideas and to what leads researchers to care about particular framings of questions and to approach them the way they do. Ideas, like people, have histories and, like people, can best be understood in the context of their histories. My experience is that when students appreciate that science and the process of research are very human activities, they are better prepared to move beyond just reading about science to doing science.

I recently led the effort to redesign the undergraduate curriculum for our department and created a new major specialization in human-computer interaction. I currently teach a graduate course on information visualization and undergraduate courses on the cognitive consequences of technology, cognitive design, human-computer interaction, and programming techniques for interactive systems. I have also assisted in teaching the graduate software engineering course in the Computer Science and Engineering Department.

Management Experience

Much of my effort over the years has gone into creating and managing research groups. My most recent management experience was as Chair of the Computer Science Department at the University of New Mexico. At UNM I assisted the department in increasing faculty size by 40 percent and research funding from less than $2M per year to over $5M per year. I was PI on a $1.25M NSF five-year research infrastructure grant for the department. I established the Computer Graphics and Interactive Media Research Group at Bellcore and served as Director. As Director of the Human Interface Laboratory at MCC, I coordinated the efforts of approximately 40 researchers and managed a $5M annual budget.