COMPUTER GRAPHICS

PROCEEDINGS

Annual Conference Series 1997

SIGGRAPH 97
Conference Proceedings
August 3–8, 1997
Papers Chair: Turner Whitted
Panels Chair: Barbara Mones-Hattal

A Publication of ACM SIGGRAPH

Sponsored by the ACM's Special Interest Group on Computer Graphics

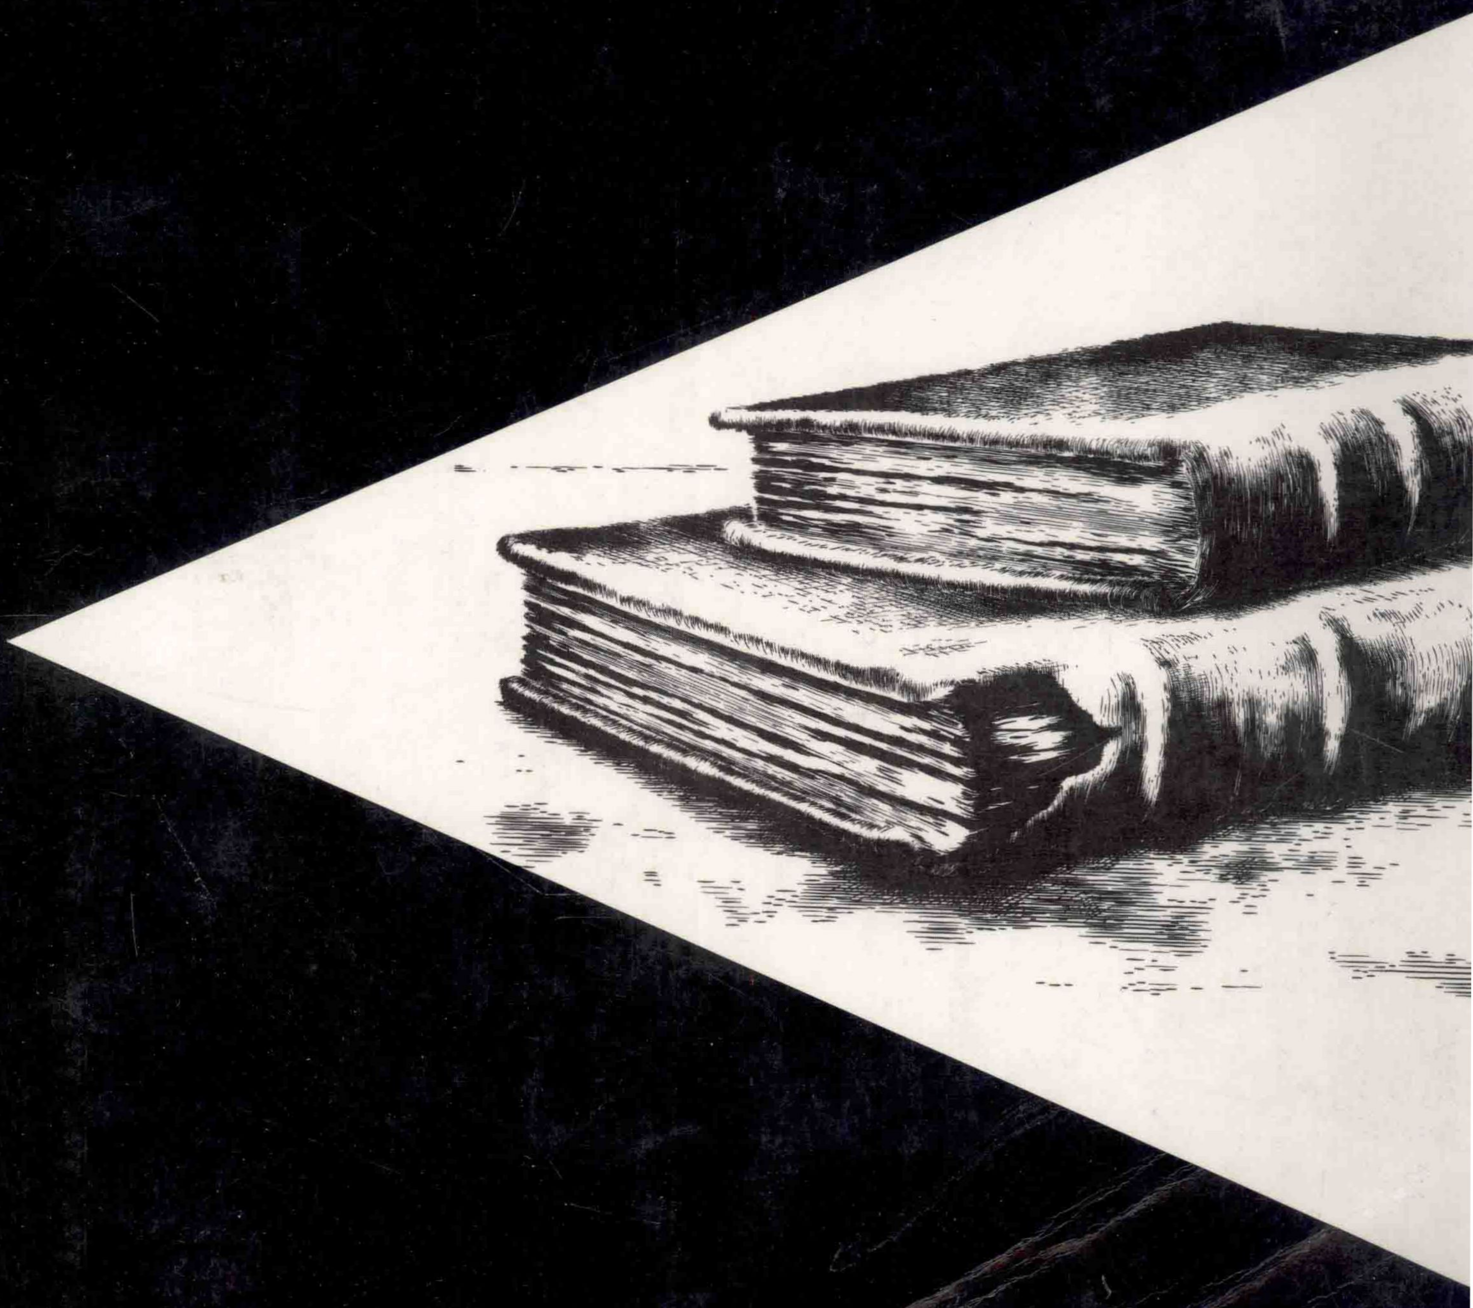

COMPUTER GRAPHICS

PROCEEDINGS

Annual Conference Series 1997

SIGGRAPH 97
Conference Proceedings
August 3–8, 1997
Papers Chair: Turner Whitted
Panels Chair: Barbara Mones-Hattal

A Publication of ACM SIGGRAPH

Sponsored by the ACM's Special Interest Group on Computer Graphics

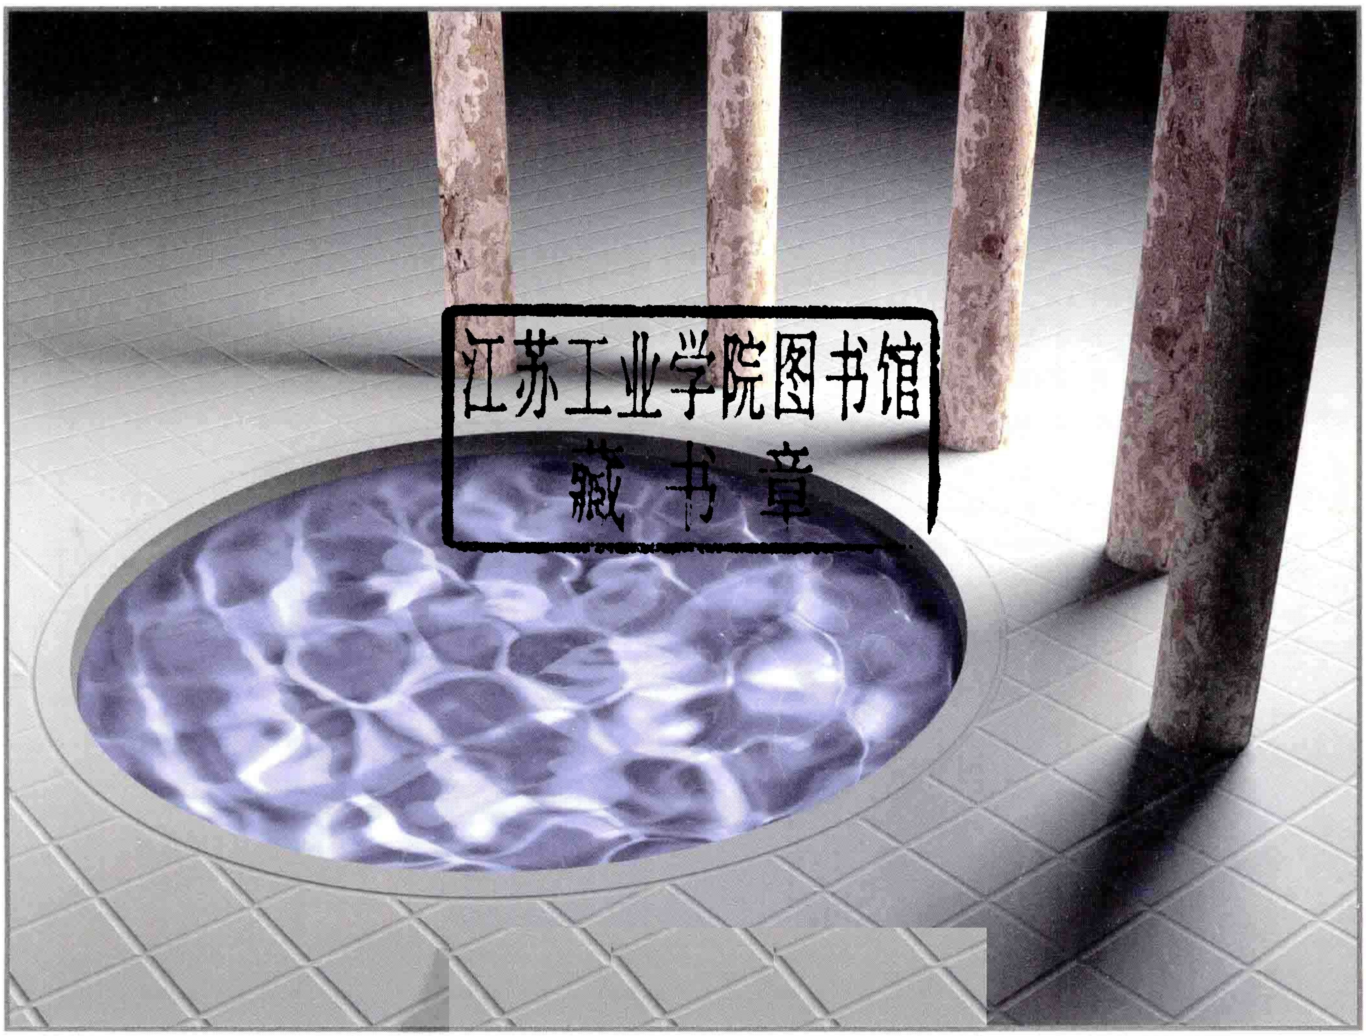

The Association for Computing Machinery, Inc.
1515 Broadway
New York, NY 10036

Sample Citation Information:
...Proceedings of SIGGRAPH 97 (Los Angeles, California, August 3–8, 1997). In *Computer Graphics* Proceedings, Annual Conference Series, 1997, ACM SIGGRAPH, pp. xx – yy.

Orders from nonmembers of ACM:

U.S.A. and Canada:
ACM Order Department
P.O. Box 12114
Church Street Station
New York, NY 10257
Telephone: +1-800-342-6626
Telephone: +1-212-626-0500
Fax: +1-212-944-1318
E-mail: `orders@acm.org`
URL: `http://www.acm.org/`

All other countries:
ACM European Service Center
108 Cowley Road
Oxford OX4 1JF
United Kingdom
Telephone: +44-1-865-382338
Fax: +44-1-865-381338
E-mail: `acm_europe@acm.org`

ACM Order Number: 428970
ACM ISBN: 0-89791-896-7
ACM ISSN: 1069-529X

Orders from nonmembers of ACM:

Addison-Wesley will pay postage and handling on orders accompanied by check. Credit card orders may be placed by mail or by calling the Addison-Wesley Order Department at the number above. Followup inquiries should be directed at the same number. Please include the Addison-Wesley ISBN with your order:
A-W Softcover Proceedings and
CD-ROM Package ISBN:
0-201-32230-7

United States:
Addison-Wesley Publishing Company
Order Department
Jacob Way
Reading, MA 01867
Telephone: +1-800-447-2226

Europe/Middle East:
Addison-Wesley Publishing Group
Concertgebouwplein 25
1071 LM Amsterdam
The Netherlands
Telephone: +31-20-6717296
Fax: +31-20-6645334

Germany/Austria/Switzerland:
Addison-Wesley Verlag Deutschland GmbH
Hildachstraße 15d
Wachsbleiche 7-12
53111 Bonn
Germany
Telephone: +49-228-98-515-0
Fax: +49-228-98-515-99

United Kingdom/Africa:
Addison-Wesley Publishers Ltd.
Finchampstead Road
Wokingham, Berkshire RG11 2NZ
United Kingdom
Telephone: +44-734-794000
Fax: +44-734-794035

Asia:
Addison-Wesley Singapore Pte. Ltd.
15 Beach Road
#05-02/09/10 Beach Centre
Singapore 0718
Telephone: +65-339-7503
Fax: +65-339-9709

Japan:
Addison-Wesley Publishers Japan Ltd.
Nichibo Building
1-2-2 Sarugakucho
Chiyoda-ku, Tokyo 101
Japan
Telephone: +81-33-2914581
Fax: +81-33-2914592

Australia/New Zealand:
Addison-Wesley Publishers Pty. Ltd.
6 Byfield Street
North Ryde, N.S.W. 2113
Australia
Telephone: +61-2-878-5411
Fax: +61-2-878-5830

Latin America:
Addison-Wesley Iberoamericana S.A.
Boulevard de las Cataratas #3
Colonia Jardines del Pedregal
Delegacion Alvaro Obregon
01900 Mexico D.F.
Telephone: +52-5-660-2695
Fax: +52-5-660-4930

Canada:
Addison-Wesley Publishing (Canada) Ltd.
26 Prince Andrew Place
Don Mills, Ontario M3C 2T8 Canada
Telephone: +1-416-447-5101
Fax: +1-416-443-0948

Contents

Papers Sessions, Thursday, 7 August

Papers Sessions, Friday, 8 August

Panels Sessions, Friday, 8 August

Special Session, Thursday, 7 August

Special Session, Friday, 8 August

Papers Preface

Whether this issue of the SIGGRAPH Conference Proceedings serves as your introduction to computer graphics or as another half inch along your already full bookshelf, you can be assured that it contains the best examples of research from our discipline. It is worth reminding each of you who attend the conference's Technical Program that the quality of this result comes not only from the effort and ingenuity of the authors and the experience and diligence of the reviewers, but also from the enthusiasm of the attendees. To the authors presenting papers at SIGGRAPH 97, your acceptance of their work is powerful motivation and a rich reward.

Submitting six copies of each of 265 papers along with their associated videotapes is not only a chore for those who create them, but a logistical challenge to those on the receiving end. This year, the Computer Science Department of the University of North Carolina at Chapel Hill generously set aside an office purely for the purpose of receiving and sorting papers. My administrative assistant, Nereida Segura-Rico along with Cathy and Chris Whitted spent days unpacking and sorting papers. In the meantime, Michael Cohen, SIGGRAPH 98 Papers Chair, traveled to Chapel Hill so that he and I could assign papers to senior reviewers. Illness did not slow Michael down, and he completed the job in two days.

Papers review has evolved into a streamlined and extremely fair process. Each submission is forwarded to two senior reviewers (members of the Papers Committee), one of whom solicits reviews from at least three other reviewers. Each senior reviewer coordinates the review of 10-12 papers, as well as reviewing 10-12 papers he or she is asked to review. Finally, each senior reviewer collects the reviews from the outside reviewers, summarizes the reviews, and consults with the other senior reviewer to whom each paper is assigned.

The 24 members of the Papers Committee met in Atlanta during the weekend of March 7, to confer, occasionally argue, and to ultimately select 48 outstanding papers from the 265 submissions. It is impossible to adequately praise these men and women for their dedication and hard work.

Accepted papers were forwarded to Stephen Spencer, SIGGRAPH Director for Publications. Similarly, videotapes accompanying papers were forwarded to Jim Rose. Both have done an outstanding job of assembling the printed and video versions of the Conference Proceedings and deserve sincere thanks.

In a recent survey attendees were asked to name their most memorable SIGGRAPH conference. The most common answer was "my first." If this is your first SIGGRAPH conference, welcome, and I hope that you find it memorable. If this is your 20th time, we all hope that you find the Technical Program as exciting and valuable as your first.

Turner Whitted
SIGGRAPH 97 Papers Chair

Panels Preface

The SIGGRAPH 97 Panel Committee vision was to support the broadly based and growing SIGGRAPH community, with outstanding, stimulating,and provocative panel topics and presentations. Concurrently, we made it part of our mission to invite the many new communities that have emerged as a result of evolving technical developments to participate in the dialog. The definition of our vision led to the design of three new methodologies in order to encourage the full participation of these diverse, international and intercultural audiences.

We will extend the reach of the panels this year by introducing the use of simultaneous translation for our Japanese speaking participants in order to provide for them a more comfortable environment for information exchange. In order to expand the horizons of our geographical connections, and improve the quality and quantity of the interactive experience, we have offered, for the first time, four pre-conference Online Panels that commenced in early May. They were the following:

- Putting a Human Face on Cyberspace: Designing Avatars and the Virtual Worlds They Live In
- Sounding Off on Audio: The Future of Internet Sound
- Motion Capture and CG Character Animation
- Medical Illustration & Visualization: Why Do We Use CG and Does It Really Make a Difference in Creating Meaningful Images?

Those panels will culminate in onsite panel presentations during the conference. Fortunately Janet McAndless, the SIGGRAPH 97 Online Chair, worked tirelessly to "make this happen" for us.

In addition, several panels will present complementary technology displays and interactive experiences before, during, or after their panels in the Creative Applications Laboratory, thus providing an innovation in the overall conference experience. The opportunity to enhance the information presented, with a "hands on" experience, or simply to take a closer look at the material presented during the panel, will encourage the integration and extension of the whole technical program experience.

The expanded desciptions of the SIGGRAPH 97 Panels are documented here in the Conference Proceedings. And this year, panel participants had the unique opportunity to include color images with their individual panel descriptions. We proudly present to you the fruit of their efforts. Credit for the excellent text and images captured here goes to the Panels Committee, panel organizers and speakers, and the SIGGRAPH Director for Publications, Stephen Spencer.

Just a few words about the process used to design the Panels Program. Our panels, by intent, highlight emerging technologies, provide a forum for the debate of technical and creative controversies, allow for the potential of diverse opinions and present the effects of these technologies on the graphics and animation communities.

Our Committee continued the practice of reviewing early proposals in order to provide constructive and useful feedback for the final proposal process. In December, we received and reviewed twenty early proposals. In January we received thirty-seven final proposals, many of which were revised versions of reviewed early proposals.

In early March, the Committee met to evaluate the final proposals and recommend final selections. The Panels Committee worked carefully, with enthusiasm, and expertise to review the final proposals and mentor the sixteen panels and one special session that were selected for the conference. The goal in selecting these panels was to select a combination of panels that would cover important and developing topics, represent new insights, and illuminate different viewpoints. It was especially important to present controversial approaches and opinions. After the selection process, our Committee members provided extensive ongoing help to the selected panels by acting as individual mentors to the Panel Organizers. The mentors and Panel Organizers enhanced communication among and between the individual panel members and the rest of the conference community.

This year, the technical program chairs defined several content tracks to make it easier for attendees to investigate and pursue their interests that pass through more than one part of the technical program. Those participating in courses, papers, panels and educators program may find it useful for scheduling their time spent in the technical sessions.

High praise goes to the members of the Panels Committee who defined and nurtured this program to its completion at the conference: Leo Hourvitz, Alyce Kaprow, Mike Mcgrath, Celia Pearce, Theresa Marie Rhyne, Carl Rosendahl, Alan Turransky, and Mary Whitton.

Our Panels Administrator Dawn Truelsen provided great assistance, support, and focus.

We also worked closely with the other technical chairs to ensure the highest quality content for all of the venues. We trust that this will result in a well organized and coordinated experience for the conference participants.

I would like to personally thank all of the SIGGRAPH 97 technical chairs, and especially Turner Whitted, the Papers Chair, whose wealth of experience, knowledge of the graphics community and collaborative spirit was an invaluable resource to me and to the whole Panels Committee. In addition, I thank Scott Owen, this year's Conference Chair, for his support in helping us reach for the vision. We hope that you will enjoy the results of our labor and find the Panels Program a worthwhile learning experience!

Barbara Mones-Hattal
SIGGRAPH 97 Panels Chair

1997 ACM SIGGRAPH Awards

Steven A. Coons Award for Outstanding Creative Contributions to Computer Graphics

James Foley

The 1997 Steven A. Coons Award for Outstanding Creative Contributions is presented to Dr. James Foley for his strong and sustained leadership in computer graphics education and research, and for his dedication to the profession through books and his work with ACM/SIGGRAPH and ACM publications.

Foley received the BS in electrical engineering from Lehigh University in 1964 and completed graduate studies at the University of Michigan in 1969. His interest in computer graphics began when he took an early course from Bert Herzog: he was instantly "hooked" and, recalling a lecture from the course, chose distributed graphics computing as his Ph.D. topic. He began his professional career at the University of North Carolina. After a stint at the Bureau of the Census he joined the faculty at the George Washington University (GWU) in 1977. In 1991 he moved to Georgia Tech and in 1996 became Executive Vice President, Mitsubishi Electric Information Technology Center, MERL.

Dr. Foley is the lead author of *Fundamentals of Computer Graphics* and of *Computer Graphics: Principles and Practice*. He is recognized as the organizer and motivator whose vision made possible these complex multi author texts and their subsequent updated editions. More than 300,000 copies have introduced an entire generation of graphics students to computer graphics not only in English but, via translations, in Chinese, German, Japanese, French, Polish, Russian and Spanish.

Foley has been a driving force in the graphics community by recognizing that the power of computer graphics can be achieved only through a carefully crafted user interface. At the University of North Carolina he co-authored the important 1974 paper "The Art of Natural Graphic Man-Machine Conversation." This inspired a career-long emphasis on research and teaching dedicated to the integration of computer technology with human-centered concerns. He launched an on going series of research projects focused on user interface management systems and development environments. In the late 1970's, at the George Washington University, he introduced one of the first courses on user interface design, a short form of which many of us took at SIGGRAPH, SIGCHI, or NCGA conferences.

Having recognized the importance of user interface studies within computer graphics, Foley combined his research talent with his people skills to build and nurture growing organizations. His group at GWU and, more recently, the world-class Graphics, Visualization and Usability Center (GVU) he established and led at Georgia Tech demonstrate this outstanding leadership. At Georgia Tech he set an interdisciplinary standard for graphics and user interface research. He successfully integrated computer science, human factors, cognitive science, graphics and multimedia design, and engineering disciplines. In a remarkably short period, Foley created an environment in which over thirty faculty and a hundred graduate students worked together. Foley not only built the GVU Center but also taught students and led research projects. He won the College of Computing Graduate Student Award, "Most likely to make students want to grow up to be professors," and the Sigma XI sustained research award.

He has devoted time and talent to fostering today's flourishing SIGGRAPH organization and conferences. He and Paul Oliver organized the first short courses at the 1974 conference in Boulder, precursor of the SIGGRAPH conferences. As Vice Chair of SIGGRAPH (1974-76), he established the annual SIGGRAPH conferences beginning with Bowling Green in 1975 and Philadelphia in 1976. His GWU colleague John Sibert and he proposed that SIGGRAPH support student and faculty attendance through volunteer positions and scholarships. He was an early influence on computer graphics standards. With Ira Cotton, he organized the 1974 NBS (now NIST) workshop on graphics standards. He later co-chaired, with Dan Bergeron, the team that specified the 1977 SIGGRAPH Core Graphics Standard. He was section editor of the *Communications of the ACM* for graphics and image processing from 1975 to 1982, and editor-in-chief of *ACM Transactions on Graphics* from 1991 to 1995.

Through his books, courses, papers, organizational, and professional contributions, Foley has made a broad and lasting impact on our field. He was an early and vigorous champion of the science, technology, and art of computer graphics, and remains a leader in his efforts to support and strength the computer graphics community. In recognition of these accomplishments and contributions to Computer Graphics, SIGGRAPH is pleased to present Dr. James Foley the Steven Anson Coons Award.

Selected References: Books

J. Foley and A. van Dam, Fundamentals of Interactive Computer Graphics, Addison- Wesley (IBM Systems Programming Series), Reading, MA, 664 pp., 1982. Translated into Chinese, Japanese, and Russian.

J. Foley, A. van Dam, S. Feiner, J. Hughes, Interactive Computer Graphics: Principles and Practice, Addison-Wesley, Reading, MA, 1174 pp., 1990, C Edition,1995.

J. Foley, A. van Dam, S. Feiner, J. Hughes, and R. Phillips, Introduction to Computer Graphics, Addison-Wesley, Reading, MA, 559 pp., 1993. Translated into German, French, and Spanish.

Selected References: Papers

Sukaviriya, P., J. Foley and T. Griffith, A Second Generation User Interface Design Environment: The Model And The Runtime Architecture, CHI '93, ACM, New York, pp. 375-382.

Foley, J., W. Kim, S. Kovacevic and K. Murray, The User Interface Design Environment - A Computer Aided Software Engineering Tool for the User-Computer Interface, IEEE Software, Special Issue on User Interface Software, 6(1), January 1989, pp. 25-32.

Foley, J., Interfaces for Advanced Computing, Scientific American, Special Issue on Advanced Computing, 257(4), October 1987, pp. 126-135.

Foley, J., V.L. Wallace and P. Chan, The Human Factors of Computer Graphics Interaction Techniques, IEEE Computer Graphics and Applications, 4(11), Nov. 1984, pp. 13-48.

Garrett, M. and J. Foley, Graphics Programming Using a Database System with Dependency Declarations, ACM Transactions on Graphics, 1(2), April 1982, pp.109-128.

Foley, J. and V. Wallace, The Art of Natural Graphic Man-Machine Conversation (Invited Paper), Proceedings of the IEEE , 62(4) (Special Issue on Computer Graphics), April 1974, pp. 462-470.

Previous Award Recipients

1995	Jose Luis Encarnação
1993	Edwin E. Catmull
1991	Andries van Dam
1989	David C. Evans
1987	Donald P. Greenberg
1985	Pierre Bézier
1983	Ivan E. Sutherland

1997 ACM SIGGRAPH Awards

Computer Graphics Achievement Award

Przemyslaw Prusinkiewicz

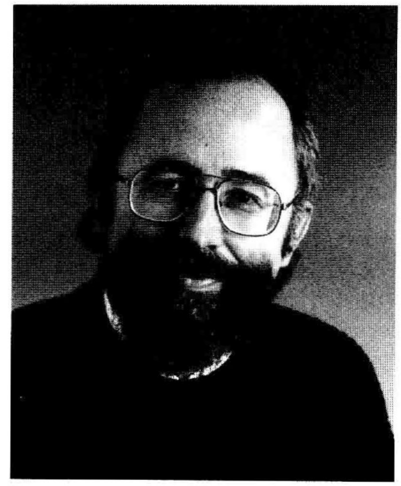

The 1997 SIGGRAPH Achievement Award is presented to Przemyslaw Prusinkiewicz for his work pertaining to modeling and visualizing of biological structures.

Dr. Prusinkiewicz's interest in computer graphics began in the late 1970s. By 1986 he originated a method for visualizing the structure and growth of plants based on L-systems, a mathematical theory of development of multicellular organisms introduced by the late Professor Aristid Lindenmayer. Professor Prusinkiewicz, his students, and collaborators transformed L-systems into a powerful programming language for expressing plant models, and extended the range of phenomena that can be simulated. Specifically, parametric L-systems facilitate the construction of models by assigning attributes to their components. Differential L-systems make it possible to simulate plant growth in continuous time, which is essential to the animation of developmental processes. Environmentally-sensitive and open L-systems provide a framework for simulating the interactions between plants and their environment. The power of these concepts is demonstrated by the wide range of biological structures already modeled, from algae to wild flowers to gardens and stands of trees competing for light.

In addition to the important extensions of L-systems, Professor Prusinkiewicz's research also includes studies of fundamental problems of morphogenesis - emergence of patterns and three dimensional forms in nature. This includes the modeling of spiral phyllotactic patterns in plants, and developmental patterns and forms of seashells.

Professor Prusinkiewicz received his M.S. (1974) and Ph.D. (1978) degrees in Computer Science from the Technical University of Warsaw. His initial research interests were in digital design, fault-tolerant computing, computer arithmetic and computer music. He held Assistant Professorships at the Technical University of Warsaw (1974-1979) and at the University of Science and Technology of Algiers (1979-1982). He joined the University of Regina in 1982 and was appointed to his current position as Professor of Computer Science at the University of Calgary in 1991. He has also held Visiting Professorships at Yale University and l'Ecole Polytechnique Federale de Lausanne, and was a visiting researcher at the University of Bremen and the Center for Tropical Pest Management in Brisbane.

As a result of his research, plants can be modeled with unprecedented visual and behavioral fidelity to nature. The book, "The Algorithmic Beauty of Plants," his contributed chapters to other books, and many papers demonstrate that plant models can be combined artistically into stunning and inspiring images. Growth of realistic and artificial life forms now can be included in computer graphics animation. His modeling methods have been incorporated into commercial products and reproduced in public-domain programs.

Dr. Prusinkiewicz's work stands out for its scholarly approach and for his collaboration with biologists, agronomists, horticulturists, theoretical computer scientists, and mathematicians. Biologists, inspired by these thoroughly researched models, have initiated international research programs including a study of the impact of microclimates on the growth of crop plants, the modeling of interactions between plants and insects for crop pest control, and a study of the relationships between plant genetics and the development of plant architecture.

These achievements produced a large impact by making complex natural environments a visible part of computer graphics. The impact can only increase as these environments become richer and even more realistic. In recognition of these contributions SIGGRAPH is pleased to present the SIGGRAPH Computer Graphics Achievement Award to Przemyslaw Prusinkiewicz.

Selected References

Prusinkiewicz, P., Lindenmayer, A., Hanan, J. S., Fracchia, F. D., Fowler, D. R., de Boer, M. J. M., and Mercer, L.. The Algorithmic Beauty of Plants, 1990 and 1996.

Mech, R. and Prusinkiewicz, P.. Visual Models of Plants Interacting with Their Environment, in Holly Rushmeier, editor, *SIGGRAPH 96 Conference Proceedings,* Annual Conference Series. ACM SIGGRAPH, Addison Wesley, August 1996.

Prusinkiewicz, P., James, M., and Mech, R.. Synthetic Topiary, *SIGGRAPH 94 Conference Proceedings,* Annual Conference Series. ACM SIGGRAPH, Addison Wesley, August 1994.

Prusinkiewicz, P., Hammel, M., and Mjolsness, E.. Animation of Plant Development, in *SIGGRAPH 93 Conference Proceedings,* Annual Conference Series. ACM SIGGRAPH, Addison Wesley, August 1993.

Fowler, D., Meinhardt, H., and Prusinkiewicz, P.. Modeling Seashells. In *Computer Graphics (SIGGRAPH 92 Conference Proceedings),* volume 26. Addison Wesley, July 1992.

Fowler, D. and Prusinkiewicz, P.. A Collision-based Model of Spiral Phyllotaxis. In *Computer Graphics (SIGGRAPH 92 Conference Proceedings),* volume 26. Addison Wesley, July 1992.

Prusinkiewicz, P., Lindenmeyer, A., and Hanan, J.. Developmental Models Of Herbaceous Plants For Computer Imagery Purposes. In *Computer Graphics (SIGGRAPH 88 Conference Proceedings),* volume 22. Addison Wesley, August 1988.

Prusinkiewicz, P.. Graphical Applications of L-systems. In Graphics Interface 1986 Conference Proceedings.

Previous Award Recipients

1996	Marc Levoy
1995	Kurt Akeley
1994	Kenneth E. Torrance
1993	Pat Hanrahan
1992	Henry Fuchs
1991	James T. Kajiya
1990	Richard Shoup and Alvy Ray Smith
1989	John Warnock
1988	Alan H. Barr
1987	Robert Cook
1986	Turner Whitted
1985	Loren Carpenter
1984	James H. Clark
1983	James F. Blinn

Quantifying Immersion in Virtual Reality

Randy Pausch[1], Dennis Proffitt, George Williams[2]
University of Virginia
pausch@cmu.edu, drp@virginia.edu, gcw@best.com

ABSTRACT

Virtual Reality (VR) has generated much excitement but little formal proof that it is useful. Because VR interfaces are difficult and expensive to build, the computer graphics community needs to be able to predict which applications will benefit from VR. In this paper, we show that users with a VR interface complete a search task faster than users with a stationary monitor and a hand-based input device. We placed users in the center of the virtual room shown in Figure 1 and told them to look for camouflaged targets. VR users did not do significantly better than desktop users. However, when asked to search the room and conclude *if* a target existed, VR users were substantially better at determining when they had searched the entire room. Desktop users took 41% more time, re-examining areas they had already searched. We also found a positive transfer of training from VR to stationary displays and a negative transfer of training from stationary displays to VR.

INTRODUCTION

In 1968, Ivan Sutherland implemented the first virtual reality system. Using wire-frame graphics and a head-mounted display (HMD), it allowed users to occupy the same space as virtual objects [Sutherland]. In the 1980's, VR captured the imagination of the popular press and government funding agencies [Blanchard, Fisher]. Potential VR applications include architectural walk-through [Brooks], simulation [Bryson], training [Loftin], and entertainment [Pausch 1996]. For the purpose of this paper, we define "virtual reality" to mean any system that allows the user to look in all directions and updates the user's viewpoint by passively tracking head motion. Existing VR technologies include HMDs and CAVEs^{tm} [Cruz-Neira].

[1]Currently at Carnegie Mellon University.

[2]Currently at Fakespace, Inc. 241 Polaris Ave. Mountain View, CA 94043

Figure 1: *Users Stood in the Center of This Room and Looked For Target Letters.*

The National Academy of Sciences report on VR [NAS] recommends an agenda to determine when VR systems are better than desktop displays, and states that without scientific grounding many millions of dollars could be wasted. Ultimately, we would like a predictive model of what tasks and applications merit the expense and difficulty of VR interfaces. In this paper, we take a step towards quantifying immersion, or the sense of "being there." We asked users, half using an HMD and half using a stationary monitor, to search for a target in heavily camouflaged scenes. In any given search, there was a 50/50 chance that the target was somewhere in the scene. The user's job was to either find the target or claim no target was present. Our major results are:

1) VR users *did not* find targets in camouflaged scenes faster than traditional users.

2) VR users were substantially faster when no target was present. Traditional users needed to re-search portions of the scene to be confident there was no target.

From these two findings, we infer that the VR users built a better mental frame-of-reference for the space. Our second two conclusions are based on search tasks where the users needed to determine that no target existed in the scene:

3) Users who practiced first in VR *positively transferred* that experience and improved their performance when using the traditional display.

4) Users who practiced first with the traditional display *negatively transferred* that experience and performed *worse* when using VR. This negative transfer may be relevant in applications that use desktop 3D graphics to train users for real-world tasks.

In a practical sense, the only way to demonstrate that VR is worthwhile is to build real applications that have VR interfaces, and show that users do better on real application tasks. That can be expensive, and new technologies take time to mature. But the computer graphics community has not even achieved a lower standard: showing, *even for a simple task*, that VR can improve performance. We show improvement in a search task and discuss *why* a VR interface improved user performance.

RELATED WORK

Several researchers have attempted to qualitatively define immersion with taxonomies [Robinett, Zeltzer] or subjective ratings by users [Heeter]. Others have measured "fish tank VR" head-tracked performance [Authur, McKenna, Ware 1993, Ware 1996], or compared variables such as resolution and frame rate in virtual environments [Smets]. We know of no work that formally measures that VR is better than a desktop interface for any search task; the closest is Chung, who compared VR against hand-based manipulation of an object, rather than the viewpoint [Chung].

COMPARING VR AND DESKTOP INTERFACES

To see if VR is useful, one could pick a representative task, such as finding an object in a scene, and compare performance with the best possible VR and desktop interfaces. That introduces many variables, as shown in Table 1. We do not wish to ask if *current* VR interfaces are useful, but rather if VR will *ever* be useful. Simply put, do users perform measurably better when controlling the viewpoint with their head instead of with their hand?

	Desktop	HMD
resolution	1280x1024	240x120
horizontal FOV	40 degrees	93 degrees
vertical FOV	30 degrees	61 degrees
input device	mouse or joystick	6 DOF tracker

Table 1: *Typical Values for Displays*

To hold the variables constant we used the same HMD as both the head-tracked display and the stationary monitor. In both cases, the scenes were rendered in stereo (The use of stereo was probably not significant, as all objects in the scene were at least two meters from the user). Figure 2 shows the stationary condition, where we bolted

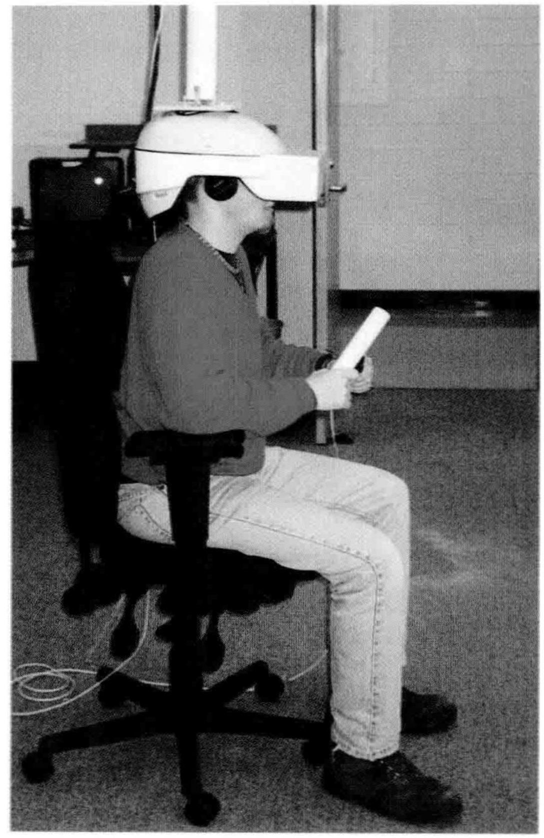

Figure 2: *Using the HMD as a Stationary Monitor.*

the HMD onto a ceiling-mounted post, thus turning the HMD into a stationary monitor. This provided the same resolution, field of view, and image quality in both VR and desktop interfaces. Table 1 gives the values for this particular HMD, the Virtual Research Flight Helmet^{tm} [HMD]. We chose a task where the display resolution was unimportant because the targets were large and easily visible. Using a mouse or joystick as the desktop input device would have introduced variables in lag and sampling rate. Therefore, we used the same magnetic 6DOF electromagnetic tracker [Tracker] from the HMD as our hand input device. All we did to create the desktop interface was to seat the user in a chair and take the 6DOF tracker off the user's head and place it in a comfortable device held in user's hands. By holding all other variables constant, we can claim our results are dependent on head-input versus hand-input. For the remainder of this paper, we refer to these groups as "the VR users" and "the desktop users." While we acknowledge that our desktop users are hardly using a conventional configuration, we claim their setup contains the essential components: a stationary monitor and a hand-input device.

We attempted a pilot experiment [Pausch 1993] where 28 users searched for easy-to-find, uncamouflaged, targets at random locations in a virtual room. VR users found the targets 42% faster than desktop users. We feared we had measured how fast users could move the camera, rather than how immersed they were. For example, finding the red 'Y' in Figure 3 is a *pre-attentive* task, where the human visual system can find the target without having to consider the camouflage. In a room surrounding the user, the time to find a red Y might be limited only by how fast one could move the camera. But the time to find an object like the black 'Y' in Figure 3 is limited by one's ability to serially examine the items. Searching for a black 'K' in Figure 3 is another mentally limited task; there is no 'K', and the only way to be certain of that is to systematically search the entire scene. We claim that VR users are much better at systematic searches because they can better remember where they have already looked in the scene that surrounds them.

We placed users in the center of a simple virtual room, 4 meters on each side. The room contained a door and two windows which served as orientation cues. During each search task, the room contained 170 letters arranged on the walls, ceiling and floor. Figure 1 shows a third-person view of the scene, with one wall removed. Letters measured 0.6 meters in length and were easily visible through the display. Users needed to apply some degree of concentration and focused attention to locate the target letter among the similar looking "camouflage" letters. In any given task, we chose target and camouflage letters from either the set "AKMNVWXYZ" (whose primarily visual features are slanted lines), or "EFHILT" (whose primary features are horizontal and vertical lines). We began each search by displaying the target letter in a fixed location over the door, and waiting for the user to say the target letter in order to begin the search. On the user's cue, we rendered the 170 camouflage letters, placing the target letter at a random location. When they found the target letter, users said "there it is," which we confirmed by watching an external monitor which displayed what the users were seeing.

48 users participated in the experiment, 24 using VR and 24 using the desktop configuration created by bolting the HMD into a fixed position. Desktop users controlled their viewpoint with the hand-held "camera" controller shown in Figure 4, which contained the same 6DOF tracker used to track the VR users' heads. We did a large number of informal experiments to design a reasonable hand-held camera controller. Based on that experience, we also removed the roll component of tracking for the hand input device. The end-to-end system latency in all cases was roughly 100 milliseconds, measured by the technique described by Liang [Liang], and we rendered a constant 60 frames per second on an SGI Onyx Reality Engine2.

AVMNVWXZWXZWXZAVMNVWXZWXZAVMNVWXZWXZAVMNV
AMNVNYVAVMNVWXZMNVAVMNVWXZWXZAVMNVAVMNVWX
ZWXZANVWAVMNVWAVMNVWXZXZXVMAVMNVMWXZNVZAV
MNVWMNZAVMNVAVMNVWXZWXZWXZWXZAVMNVWMAVYNV
AMNVNVAVMNVWXZWXZAVMNVAVMNVWXZAVMNVWXZMXZ

Figure 3: *Find the red Y, the Black Y, and the Black K.*

Figure 4: *The Hand Input Device, Containing the 6DOF Tracker.*

RESULTS

Graph 1 shows the average time users needed to locate a target.

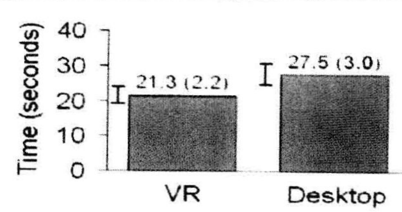

Graph 1: *VR versus Desktop Performance. The difference is not statistically significant.*

Each user performed five searches which we averaged together to form a single data point for that user. The bars in Graph 1 are the average of the 24 VR users and the 24 desktop users. The error bars show the standard error for each data set. The VR and desktop times are very similar, and their difference is not statistically significant. We constructed a cognition-limited task, so it is reasonable that the VR and desktop times are similar. We informally observed that users never physically turned the camera as fast as they could have. The cognitive portion of the search task slowed the users down.

Practice did not appear to be a factor. We asked users to do practice searches until they were comfortable; we required two practice, searches, and some users did three. We did not count practice searches in the results. Users took roughly 15 minutes to perform the searches, and none appeared fatigued. To measure practice and fatigue, we ran separate control groups with eight users each, who ran double the number of trials on both the VR and desktop interfaces. These users showed no statistically significant differences between their earlier and later trials. Users made essentially no errors. All users were between 18 and 25 years old, mostly undergraduates with no VR experience. Both groups were evenly balanced by gender. All users said they could easily see the targets. In addition to the 48 users we report, 3 other users began but did not complete this study. All 3 felt slightly nauseous, and they all reported that they were generally prone to motion sickness.

We now consider searching for a target which is *not* in the scene. Of course, if the user *knows* the target is not there, then the task is pointless. Therefore, we had users perform a sequence of searches, each of which had a 50% likelihood of containing a target. Users were instructed to either locate the target, or claim no target existed. In this way, we measured the time users needed to *confidently* search the entire scene.

If the targets are dense, and the users are efficient in their searching, we can predict how long this will take. Working backwards, consider an efficient user who takes 40 seconds to completely search a scene, with no wasted effort. On average, when a target *is* present, that user should find it in 20 seconds. Random placement may make the letter appear earlier or later in the search process, but on average the user will find the target halfway through the search. We know how long it takes users to find targets when they are present. If the users searched perfectly, it should take twice that long to search the entire room and confidently conclude the target is not there. Any time over that would imply that the users were re-examining portions of the room that they had already searched. This prediction is shown in Graph 2.

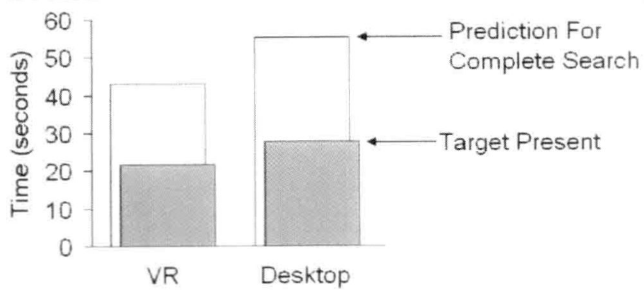

Graph 2: *Predicted Times for A Complete Search.*

Graph 1 showed the results of users who each performed five searches for targets that were in the room. In fact, these users each performed a sequence of ten searches, where on any given search, the target might or might not have existed. For each of the ten searches, the user was told to either find the target, or announce that it was not there. Users did not know beforehand whether a target would be present in any given search. Graph 3 shows the average time users required to locate a target that was in the room (Graph 1 results), the predicted time to search the entire room (Graph 2 results), and the observed time to search the entire room and conclude that no target existed.

The VR user data is only 1.4% above the prediction for efficient search. This concurs with our personal observations of VR users, who appeared to search the entire room without rescanning. However, desktop users typically examined portions of the room a second time. As shown in Graph 3, the desktop users spent 41% above the time that a perfect search would take.

IMPLICATIONS

The VR community claims that a head-tracked, egocentric camera control provides a stronger sense of immersion, or "being there," than does a desktop display. Our results indicate that VR can help users remember where they have and have not looked. The ratio shown in the "desktop" performance in Graph 3 implies that back-

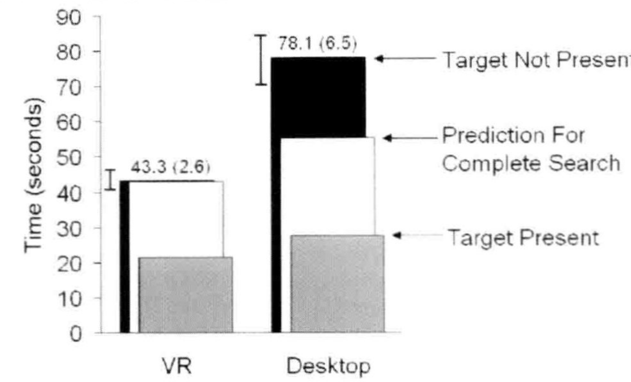

Graph 3: *Observed Times To Search the Entire Room and Determine that No Letter is Present.*

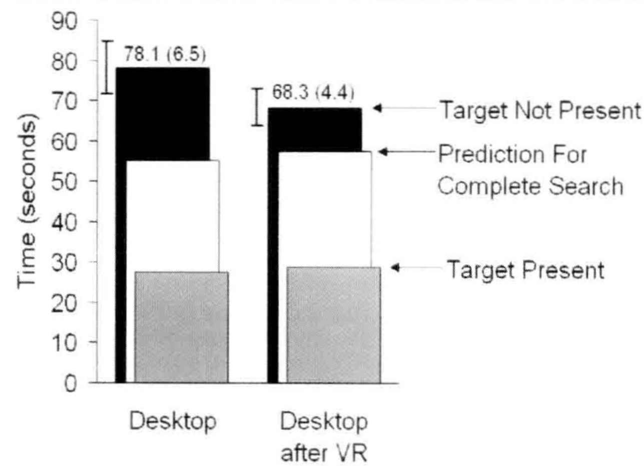

Graph 4: *Positive Transfer: Users Who Practice in VR Improve Their Performance on the Desktop*

tracking is occurring [Braddick]. If the desktop users were slower for some biomechanical reason, such as our choice of input device, we assume it would have also slowed them when the target was present.

TRANSFER EFFECTS

We wondered how users would perform the desktop search tasks if they first practiced in VR. If VR allows the user to develop a good frame-of-reference for a space, perhaps that memory would carry over to a desktop interface. We had each of the VR users perform their ten searches, rest for five to ten minutes, and then perform ten more searches using the desktop interface. In this way, we could see if the experience with VR affected later use of the desktop interface. The ten desktop searches, just like the first ten VR searches, contained five with a present target and five without. Graph 4 shows a *positive transfer effect*, where practicing in VR improves performance of the same task when using a desktop interface. This result is statistically significant ($p < 0.0096$). We also performed the reverse experiment — we had the desktop users rest and then perform ten more searches using the VR interface. Graph 5 presents the results, which show a *negative transfer of training*. Practicing on the desktop decreases performance of the same task when using a VR interface. This result is statistically sig-