



When Everything Learns

Craig M. Kanarick and
Ben H. Davis, Editors

Contents

4	Foreword: “Learning A Few Things” Douglas Rushkoff
6	Introduction
14	“Thinking About Thinking” Craig Kanarick
19	Session 1
34	Images: Technology, Learning, Smart, Future, Life
36	Session 2
54	Session 3
66	Acknowledgements, Colophon, About Razorfish

razorfish

Razorfish, Inc.
©2001 All Rights reserved



Participants

Standing (left to right):

Ben Kleinman, B.J. Fogg, Stephen Turbek, Ben Davis, Ana Echeverry, Tucker Viemeister, Neil Crofts, John Weber, Peter Lunenfeld, Constance Adams, Adam Eeuwens, Joy Mountford

Seated (left to right):

Margaret MacLean, Joseph Busch, Sue Madden, Camille Habacker, Rebecca Odes, Craig Kanarick, Michael Ester, Eric Begleiter

Learning a few things

Douglas Rushkoff

Things don't learn. They can't. That's why we call them "things." Or at least that's how I've always felt most comfortable musing about the intelligence of inorganic matter. That is, until I fully considered the concepts being bandied about by the scientists in this book.

My interest in technology has always leaned toward the fantastical. The first time I posted to a FIDO-NET bulletin board or threw on a pair of primitive VR goggles was like being transported into a different world. It was an alternative universe, where the limitless possibilities of a vision quest or lucid dream were now suddenly on tap.

But somehow, because of my own naïve understanding of how these technologies actually worked, I was able to keep my experience of the technological relegated to the distant sphere of mystic journey or spiritual discovery. These tools might extend one's ability to participate the Gaian mind, but they certainly wouldn't achieve any lifelike dimensionality of their own. Why, that would challenge the notion that we thinking creatures are somehow special. The cyborg archetype always had a person inside.

But that was a different era.

See, back in the 1980's, when writers and artists like me got interested in the sudden explosion of new technologies, most of the real scientists were too busy actually figuring stuff out to sit around and muse on the possibilities of virtual reality, tele-robotics, cyberspace, or nano-technology. Instead, the responsibility for pondering the implications of these developments fell on us, the under-qualified poets and novelists. Romantic as we were, we always placed the human imagination above technological reality – and succeeded in vastly underestimating the true potential of the toys with which we were playing.

What a relief it is that Razorfish has decided to create opportunities for some real scientists do the talking for a change. Reading their uncensored, unrehearsed, and uninhibited conversations is at once refreshing yet terrifying, reassuring yet humbling. You hold in your hand a set of deeply informed reflections on technology's evolution towards something akin to human reason.

Scientists to the core, these people spend a good deal of time coming to consensus about just what it is they're discussing. And rather than simply ignoring the facts and plot points that might hurt their assertions or story lines, they mine these inconsistencies for new paths of inquiry.

Paradoxes are not treated as stumbling blocks, but as opportunities.

So let's dispense for a moment with our ingrained sense of God's good order, and suppose, along with these ladies and gentlemen, that science may one day soon actually develop a machine that can not only think, but learn.

Sure enough, it becomes easy to consider how almost everything learns in one way or another. Even a rubber band stretched around a book for too long "learns" by losing its elasticity, and "remembering" the shape of the object to which it was wed for so long.

But no, this isn't really learning, at least not according to the voices here – it's merely reacting. Real learning involves solving new problems using the experience of old ones. Today's most readily available examples might be digital video recorders like TiVo that study our viewing patterns and then record for us the shows we tend to like. Or cars that observe our driving habits, and then change their suspension and gear-shifting to optimum parameters for our individual styles.

When they talk about it that way, I don't worry so much about the ways in which things might learn. It's not as if our

machines are becoming alive – they’re just learning how to serve us better, without our spelling everything out for them.

But the romantic in me is still troubled. It’s easy to limit one’s consideration of technologies “learning” to the simple case of machines coming to understand what I, as a consumer or computer user, want. Which TV shows I want to watch, which books I want to buy, how I want to drive my car. Tivo might as well be a dog or a slave. It simply learns how to do its master’s bidding. So is this really learning?

I’m not so sure. Maybe that’s why other programmers are working on ways to make machines appear to be more sentient than they really are. New computer interfaces simulate the process of thinking by asking questions they already know the answers to, or pretending to make small mistakes. Again, though, these are pre-determined idiosyncrasies, and not real behaviors. But, then again, why do human beings make mistakes?

The men and women who assembled in Tarrytown traveled through just such difficult terrain, and then some. For example, they focused on how things that learn might be able to

impact our own human decision-making capabilities, in real time. One participant called this breed of intelligent device a “contextualizer.” Think of it like a global positioning satellite, but for information, ideas, or even ethics. Perhaps they could take the form of simple memory extensions, as Joy Mountford suggested: a pair of earrings that whispered to her the names of the people she was talking to. But how about extensions of our empathic or activist capabilities? What about a global-warming consciousness-raising implant, that reminded you of how much wasted energy is associated with each of the tasks you are considering? Or a thought monitor that measured the “neurosis quotient” of your internal monologue? Then, of course, comes the question of whether any of these enhancements are fair, appropriate, necessary, healthy, or even “moral.”

Don’t look for final answers on these pages. Instead, enjoy the opportunity to experience a bit of the kind of inquiry that engages the people who wrestle with these problems every day. They consider everything from how we anthropomorphize our machines to the unintended consequences of new technologies. Will machines allow us to build bigger conceptual taxonomies, or will they simply require us to learn

how to trust them to think for us? Will humans have any time to think at all when they’re so busy building, repairing, and investing in thinking machines? And what if our machines learn what we teach them, rather than what it is we’d really like them to know? Will they force us to quit smoking cigarettes and eating partially hydrogenated vegetable oil? Will they conclude that the greenhouse effect is a danger worth forcing us to face, or a terrific way to get rid of those troublesome biped mammals, once and for all?

Most of us are not used to thinking this way, except when reading Isaac Asimov or Fritjof Capra. And then it’s from within the safely sealed context of science fiction or quantum thought experiments. The men and women who conducted the following discussions in Tarrytown, 2000, are doing this work for real, 24/7.

Maybe that’s why the most valuable lesson for us to take from this series of conversations is that in the very act of pondering the ways things could learn to learn, these scientists give us a window into how we might learn a thing or two, ourselves.

Introduction

When Craig Kanarick, co-founder and Chief Strategic Officer for Razorfish came up with the idea of having a summit meeting that would address both the advent of new technologies and their social implications, the development of guest lists began in earnest. What he was most insistent about was having a diverse group of people that could grasp the issues, speak articulately, and develop new ideas at a run. And they had to be fun. Going to hundreds of “visionary” conferences on technology exposes you to most of the best minds in the world, so choosing people to be at your own meeting is problematic in an interesting way. You want to hear some new visions, ones that you haven’t encountered on the visionary circuit so many times that you can predict the future-speak after a few syllables. Not that those visions aren’t great and good, but a new wrinkle is worth the effort. Some of the people we invited of course had other things to do, other commitments, other events that conflicted with our targeted fall weekend just before Halloween a few miles from Sleepy Hollow in Tarrytown, New York. And a perfect weekend it was, fall

leaves raging, three days of bright autumn sunshine, crisp days, cool starry nights in an old estate by Washington Irving’s Hudson River.

But perfect is as perfect does. Who were these minds that instinctively grasped the significance of Craig’s notion of “When Everything Learns” being a singularity, a moment that would come very soon given the pace of technology and would have serious repercussions for everybody? The group was evenly divided between internal Razorfish people and ten of the most interesting thinkers that represented some of the most forward looking organizations in the world including Lockheed-Martin Space Operations, LightVisions and MicroVista, Interwoven, Inc., Luna Imaging, Inc., Stanford University Persuasive Technology Laboratory and Casio America, the Art Center College of Design, Cultural Heritage Management and Strategy, iBias, gURL.com, and the San Francisco Museum of Modern Art.


Easiest to rely on the alphabet for the telling of these lives...

Ben H. Davis
January, 2001



Constance Adams

is a Space Architect and Human Factors Engineer at Lockheed Martin Space Operations Company working with NASA on a variety of projects including the architecture and cabin integration for space crew return vehicles, second-generation reusable launch vehicles, crew outfitting and accommodation design, human centrifuge specifications, and solid-state lighting systems. She is also an adjunct Professor of Architecture at the University of Houston and Senior Designer/Professor at the Sasekawa Int'l Center for Space Architecture; and have directed design studios in collaboration with the Rhode Island School of Design, Yale School of Architecture, University of Texas (Austin), Lund University(Sweden), and the Technical University of Munich (Germany).



Eric Begleiter

is the patent holder of the “Prismagic” method for putting holographic illusions on unusual surfaces. He is the president of a research and licensing company DFC, in Boston, Mass., which carries out developments in the area of applied display optics. He is also the president and an owner of LightVision Confections, which manufactures edible holographic candies under the brand name Holopop TM. (holopop.com) with factories in Cincinnati, Ohio and the English midlands. He has been chewing over the concept of edible holograms since he was a fellow at the Center for Advanced Visual Studies at MIT and wondered what it would be like to eat a rainbow. Mr. Begleiter is also founder of a Boston based start-up called MicroVista, investigating microfabrication techniques for the biotechnology industry.



Joseph Busch

is a leading authority in the field of information science. His focus is on business strategy and market positioning for Interwoven’s content analytics solutions recently acquired from Metacode Technology. Joseph came to Interwoven, Inc. from the Getty Information Institute, where he was a Program Manager for ten years. He is widely published in the field of information science and maintains active participation in key professional organizations and standards committees. Joseph is the President-elect of the American Society for Information Science and Technology (ASIST). He has brought information software products to market and has extensive management experience. Prior to joining the Getty, Joseph was a Manager at Pricewaterhouse. He earned a BA from Portland State University and a Master of Library Science degree from the State University of New York, Albany.




Neil Crofts

came to CHBi, the UK forerunner of Razorfish, in the autumn of 1997, as Head of New Business. By the summer of 1999 he was building up a digital strategy capability in London. Neil became head of Strategy in Europe at the beginning of 2000 and led the small team that “wrote the book” that defines what digital strategy is and how it is done at Razorfish globally. Neil has applied his thinking to many clients including NatWest, Banca Mediolanum and C&A, and has influenced many more. Neil’s next moves include the development of a training course on “Transformational Digital Leadership” for clients and the writing of a book about how to re-orientate a company for success in the digital economy. Neil’s career began with racing cars, running his own team and racking up a few modest successes both in races and in sponsorship. He then started to create attractions for consumers first in magazine publishing and later in exhibitions, which led to running an exhibition company for four years. Following the sale of the company, to French mega-corp Havas, Neil wrote his first Web start-up business plan “software-house.com” towards the end of 1995. Finding little comprehension for the idea with venture capitalists, Neil returned to exhibitions for another year, running Multimedia 96 at the Business Design Centre. He still races but on mountain bikes now.



Ben Davis

was a Senior Scientist and a Strategic Development Director for Media and Entertainment with Razorfish, Inc. in Los Angeles. Davis came to Razorfish from the J. Paul Getty Trust where he was Manager of Electronic Publications and Manager of Communications for the Getty Information Institute. Davis came to the Getty Information Institute in 1995 from the Massachusetts Institute of Technology Center for Educational Computing Initiatives (CECI) where he was a Research Associate, Manager of the AthenaMuse Consortium, and Manager of the Project Athena Visual Computing Group (1987-1991). He was also an Instructor at the MIT Media Lab, a Fellow at the MIT Center for Advanced Visual Studies, and a lecturer in the MIT Visual Arts Program. He is the co-author of *Time and Bits: Managing Digital Continuity* with Margaret Maclean and has written for *Scientific American*, *Aperture*, *Visual Resources* and a number of other magazines. He is currently a private consultant for digital media asset management.



Ana Echeverry

led and developed innovation strategies for clients like 3Com, The Financial Times, and various start-up businesses. Internally at Razorfish, she works on diverse corporate initiatives involving the development of the global strategy and user intelligence practices. Ana concentrates on the development of strategies for clients dealing with convergent offerings that combine physical and digital components. Her skills combine business strategy, social science research, and user-centered design. Previous to joining Razorfish, Ana helped the Strategic Service Development department at Ernst & Young improve the speed and effectiveness of the firm’s innovation process. She has also done an innovation strategy leveraging digital technologies for a Chicago-based healthcare provider. Ana holds a master’s degree in design planning from the Institute of design, Illinois Institute of Technology, and a B.A. in industrial design from the Universidad Pontificia Bolivariana, Colombia.



Dr. Michael Ester

is currently President of Luna Imaging, Inc. Luna enables organizations to build and distribute high-quality visual collections in digital form, and provides sophisticated software to manage, access, and use rich media over the Internet. From 1985 to 1993 Dr. Ester was Director of the Getty Art History Information Program (AHIP), an operating program of the J. Paul Getty Trust. In collaboration with domestic and international institutions and organizations, AHIP worked at several levels of policy, standards and practice to help shape the direction of automation in the visual arts. Michael Ester was responsible for setting program direction and policy, and for managing its many projects based in the U.S. and in Europe. During his tenure at the Getty, he initiated basic research and technical development in the use of digital imaging as a reproduction medium for the visual arts. Prior to joining the Getty Trust, Dr. Ester was Information Systems Manager then Director at URS/Berger, a firm conducting remote sensing, environmental studies, and GIS services for the United States Government. He was also formerly General Manager at Technical Data Processing Associates, which provides CAD/CAM systems and services for architectural and engineering applications. Before entering the private sector, Dr. Ester was an Associate Professor at Rutgers University, where he taught courses in both computer applications and archaeology. Dr. Ester received his Bachelor's degree in mathematics and anthropology from George Washington University; he earned doctorates in the same disciplines from Brandeis University.




Adam Eeuwens

was a science writer/editor on Razorfish Reports. Originally from Amsterdam, the Netherlands, Eeuwens moved to Los Angeles four years ago, where he remains in a state of culture shock. After milking cows for six months in Iceland Eeuwens' career started in 1989 at Electric Word/Language Technology, as assistant to Louis Rossetto and Jane Metcalfe, the predecessor to Wired Magazine. He published his own magazine Flux in the early nineties and was editor of the night club event "TimeFuckers" in Paradiso. He played videogames professionally for two years meanwhile setting up the widely successful magazine Power Unlimited. He went on to become executive editor at WAVE magazine in Brussels, while also producing the new media event FaceTime at Paradiso, Amsterdam. After moving to California Eeuwens became the US correspondent for design and business publications in the Netherlands. He also was co-curator of the exhibition "Do Normal, recent Dutch design" and is currently working as a private design consultant.




Dr. BJ Fogg

(Ph.D., Stanford) directs research and design at Stanford's Persuasive Technology Lab (www.captology.org). An experimental psychologist, BJ also teaches courses in persuasive computing for Stanford's Computer Science Department, where he's been appointed to the consulting faculty. In addition to his Stanford work, BJ is the Director of Research & Innovation at Casio's U.S. R&D Center (www.casioresearch.com). In that role he leads Casio's efforts to create next-generation i-appliance products. BJ is also the Chairman of the Board at VentureNova, a seed fund for i-appliance ventures (www.venturenova.com). BJ holds several patents, and his work has been featured in Business 2.0, The New York Times, Wired News, and I.D. Magazine.



Camille Habacker

As experience lead for an approximately 90-member team, Camille Habacker is responsible for ensuring the delivery of high quality, measurable and innovative user experiences. She focuses mainly on developing new business, staffing projects, supporting team collaboration, and training project-level experience leads. As director of the content group, Camille is responsible for ensuring the quality of the content Razorfish develops for itself and its clients. In this role, she focuses mainly on recruiting, hiring, and training talented writers for digital media, as well as maintaining a consistent and dynamic sense of Razorfish style. Camille has been with Razorfish for three years, working as an information architect, writer and manager. Her skills have benefited several projects, including those for Charles Schwab, Armani Exchange, MoneyUnion, Derby Bicycles, the Cooper Hewitt, OmnicomGroup, Christie's and AT&T. Camille holds a B.A. in English from Franklin & Marshall College in Lancaster, PA, and a M.S. in technical communication from Rensselaer Polytechnic Institute in Troy, NY



Craig Kanarick

Craig Kanarick is the co-founder of Razorfish, Inc., where he served for over six years as Chief Scientist and Chief Strategic Officer, playing a critical role in the explosive impact of digital technologies on business practices. At the company, he was responsible for the creative and strategic direction of the company as well as for many of the projects created for its clients.

Craig Kanarick, 34, also founded Razorfish Studios together with Jeff Dachis in 1995, and has been responsible for establishing the creative and technical direction of the company. In addition to his role at Razorfish Studios, Mr. Kanarick serves as the EVP and Publisher of BUST, a women's lifestyle media company.

Moreover, Mr. Kanarick is the co-proprietor of The Slipper Room, a small theater/lounge in lower Manhattan. He is also on the Board of Directors of Rhizome.org, a digital arts non-profit organization, and Improv Technologies, a software company based in New York.

Mr. Kanarick holds a Master of Science degree in Visual Studies from the MIT Media Laboratory, as well as a B.A. in Philosophy and a B.A.S. in Computer Science from the University of Pennsylvania. Before starting Razorfish, Inc. and Razorfish Studios, he worked as a digital media consultant, designing network communications protocols and simulation software.



Ben Kleinman

works in the Infrastructure Group at Razorfish, helping the company know about, understand, and use its collective wisdom and experience. His focus is designing communities and tools that communities use to form, evolve, and thrive in a knowledge ecology. Previously, Ben worked for J. Walter Thompson in both New York and Chicago where his global branding accounts included Merck and Kraft. Ben is currently completing his Masters thesis in the John W. Draper Interdisciplinary Masters Program in Humanities and Social Thought at New York University, focusing on the current and future structure of communities in the digital space. Ben is a member of the Organization of American Historians.



Peter Lunenfeld,

Ph.D. is the Director of the Institute for Technology & Aesthetics (ITA), founder of mediawork: The Southern California New Media Networking Group, and a coordinator of the graduate program in Communication and New Media Design at Art Center College of Design in Los Angeles. He is the editor of *The Digital Dialectic: New Essays on New Media* (MIT Press, 1999) and the author of *Snap to Grid: A User's Guide to Digital Arts, Media, and Cultures* (MIT Press, 2000).




Susan Madden

As Managing Director, Susan Madden guides the overall operations of the Boston office. In this role, she is responsible for the leadership and oversight of client satisfaction, employee recruitment and retention, office culture, and revenue generation. Previously, Susan was the Vice President East Region and Practice Management within i-Cube's Client Services Group. She oversaw all i-Cube projects in the East region, and worked closely with Client Services management and project teams to develop and implement scalable processes and consistent high-quality, on-time project delivery. Susan's previous work experience consists of fourteen years in banking information systems at BankBoston. In this role, she provided leadership for major bank integrations and developed IT strategies for the Bank. Susan was a member of the Bank's IT due diligence team and participated in seven bank acquisitions. Her most recent experience was leading the Bank of Boston/BayBank integration.



Margaret Greenup Holmes MacLean

is an anthropologist and archaeologist whose work currently involves strategic planning and policy development for cultural organizations, designing and managing intensive seminars on a range of topics for several foundations, and conservation of heritage sites. Her interests include the history of technology, the history of the future, the built environment, and what culture chooses to value and protect. She spent a few years ambling around South America in her youth, six years on academic degrees in Northern California, four years managing the scientific research program for a granting organization in Boston, nearly a decade managing and directing international programs at the Getty Trust, and two surprisingly entertaining years as an independent consultant. Her Rhode Island accent is almost undetectable. BA in Anthropology, MA in Anthropology & Cultural Geography, Ph.D. in Anthropology & Archaeology [University of California at Berkeley].




S. Joy Mountford

has been designing and managing interface design efforts for over 20 years. Her experience encompasses a range of innovative and pioneering interface developments on different user systems from airplanes to PCs to consumer electronics. Most recently she was at Interval Research Corporation for over 5 years leading a series of musical development projects. Previously she was the creator and manager of the highly acclaimed Human Interface Group at Apple Computer for nearly eight years. Before joining Apple, Joy worked at MCC, an A.I. computer consortium and prior to that she designed advanced user interfaces for military avionics systems at Honeywell. Joy initiated Apple's International Interface Design Project and continued this at Interval. Both have been catalysts in setting up and funding various interdisciplinary courses at universities around the world. She frequently teaches classes at Art Center on various editorial boards across the design and interaction communities, and is an elected board member of International Design Conference in Aspen and contributor to the Discover awards. She conceived of and formed the team that wrote and edited *The Art of Human Computer Interface Design* by Addison Wesley, 1987.



Rebecca Odes

is the co-founder and creative director of gURL.com, a Web site for teenage girls and young women. gURL was created as an alternative to traditional girl-directed media. Founded in 1996, gURL.com has gained recognition for its content and design, and is the leading content/community site for teenage girls. In 1999, Odes co-authored and illustrated *Deal With It!*, A Whole New Approach to Your Body, Brain, and Life as a gURL, a national bestseller and winner in the 2000 I.D. magazine design competition. She received a Master's Degree from the Interactive Telecommunications Program at NYU and did additional graduate study in painting at the School of the Art Institute of Chicago. Her undergraduate degree is from Vassar College. Before entering the digital realm, she was a semi-punk recording artist.



Stephen Turbek


As an information architect, Stephen uses his skills in design, technology, and critical thinking to improve information accessibility and user experience in digital interfaces. By analyzing user goals and information structure, interfaces can be optimized to benefit the company and the user. Stephen Turbek has been an integral part of Razorfish's growth for the past five years, holding a variety of positions since joining the company as one of its first employees. He has designed and built projects for clients such as The American Museum of Natural History, AT&T, the online gaming site Bunko.com, Charles Schwab, Citibank, IBM, Toshiba, Verizon, and The Whitney Museum of American Art. In another capacity, he designed the Razorfish New York office environments, and created the Razorfish physical identity guidelines.

Stephen's work has appeared in ArtBytes, Art forum, How, The New York Times, Web Techniques, and Wired; in books including *Deconstructing Web Graphics* by Lynda Weinman, and *Web Navigation* by Jennifer Fleming and been awarded the Taka-Fuji International Design Competition.



Tucker Viemeister

Named after the car his father helped design, Tucker Viemeister, helped found Smart Design (they designed Good Grips and lots of other fun profitable products and environments). Opened frogdesign's New York integrated strategic design studio, before he brought a new dimension to Razorfish, helping to integrate all media to deliver a deeper user experience – which demands a new kind of designer. That's why Metropolis magazine called him "the last industrial designer." His work was selected for first Presidential Design Achievement Award (1984), Forma Finlandia (1987), annual ID Design Review (11 times), IDEA awards (7 times), Design of the Decade (1), and is in the Cooper-Hewitt and the Museum of Modern Art collections. He edited Product Design 6, is an Architectural League of New York Director, is President of the International Design Network Foundation, Cooper Hewitt's Professional Designers Advisory Committee Chair, and is a Fellow of the Industrial Designers Society of America. Tucker is currently a private design consultant in New York City.



John S. Weber

joined San Francisco Museum of Modern Art (SFMOMA) as curator of education and public programs in 1993. He oversees all SFMOMA education initiatives, including internationally recognized programs employing multimedia computers and other new digital technologies; public programs and events including lectures, classes, symposia and workshops; docent programs for school and adult audiences; programs for teachers, schools, youth and family audiences; community outreach; and other activities that augment the Museum's artistic program. In addition to supervising SFMOMA education efforts, Weber is active as an exhibition curator himself. His current curatorial projects include "010101: Art in the Technological Times," a collaboratively curated exhibition to be presented on the Web and in SFMOMA's galleries in early 2001. He was educated in studio art, art history, and criticism at Reed College and University of California, San Diego.

Editor's note

Our remote documentation facilitator for the event was Bernie DeKovan who's Technography, Inc. supplied us with a unique opportunity to have a technology assisted, on-line "smart meeting" experience.

As you read the transcripts of the meeting, keep in mind that our intention was to bring together a diverse group of thinkers associated with companies and organizations that valued them for their ability to not only "think out of the box" but to think "without the box" entirely. In many ways the meeting itself became a kind of "thing that learned" because it had no pre-determined end-game. If technology is changing the definition of thinking and learning then the context of the meeting should be able to shift to accommodate topics as diverse as city planning, lucid dreaming, digital democracy, the origin of sleep, environmental consciousness, and "Gilligan's Island." What you are about to read should surprise and delight you given the speed at which the world is changing. Everything will indeed have learned and changed since those golden autumn days in October of 2000 beside the Hudson River near Sleepy Hollow.

Ben H. Davis
January, 2001

Thinking about thinking

Craig Kanarick

I'd like to welcome everyone here, and to talk a little bit about why we're all here in two different contexts. One is the larger question of why are we here for this weekend getting together, talking about the subject that we're talking about and the other is to talk a little bit about the subject itself.

I've met many of you but for those I haven't, I'm Craig Kanarick and I will be the moderator for most of the discussion for the rest of the weekend. I'm one of the co-founders of Razorfish, founded about six years ago with an old friend of mine from nursery school, Jeff Dachis. We have come a long way in the last six years. It's been an interesting ride.

I don't know if everybody knows a lot about the company. About half the people here are actually from the company. The company now is about 2,000 people in fifteen cities and nine countries delivering what we call "transformational services"—helping companies adjust to what we believe is a new world, a new world that's heavily influenced by digital technologies in particular. We help create strategies for companies and then implement those strategies. We help companies continue to exist and hopefully take advantage of the great opportunities that are now available to them.

My job at the company is Chief Scientist. The idea behind that title is that what we do is a science—it's not magic, it's not luck.

It's a science that combines creativity, technology, service, and strategy. Being the Chief Technical Officer wouldn't cover everything that's relevant about what we do. So, the evolution of my title comes out of that situation, Chief Scientist.

What I do is run a small department in the company. It's about ten people and is simply called the Science Department. The Science Department is dedicated to what we call sustainable excellence. Making sure that the organization as a whole is sustainable and that the company is excellent at what it does. In other words, we think about what we do as the management and development of knowledge. We are, as a professional services firm—as a consulting company—hired because of our knowledge. Because of our expertise, not because of our hardware or our secret formula, or our manufacturing processes, or enhancements that we have. But because of our brains. And there's a theory behind the company that the mass of the 2,000 brains is better than the sum of the parts. The way that happens is by us sharing knowledge with each other. If someone does a project in Finland about filing insurance claims on cell phones, they can spread the knowledge from that solution to the rest of the company. People in the LA office who are doing a similar type of project can take advantage of that experience.

Our mission in the company is to set up programs and processes to share knowledge within the company as well as

to develop new knowledge, to develop ideas and perspectives, and to provide thought leadership for the company. This also helps ensure that Razorfish has a thought leadership position outside of the company. That's really the premise behind this weekend—to get together a group of very diverse, interesting people to help develop new philosophies and new knowledge. It's good for Razorfish because I get to trot out some of our really smart people in front of other thought leaders in the world. Theoretically, you know, you listen to the smart Razorfish people and go back to your community and say, "Hey those Razorfish people are really smart!" and that helps our company. What we get out of it is an opportunity to interact with you smart people and to get smarter as a result of the weekend. Hopefully, you get the same thing. Hopefully, that interacting with us will develop your knowledge and your perspective on the world and you will feel good about that.

This is the first time we've done a Science Alliance Summit. We're probably going to make mistakes. Hopefully, we'll also have the benefit of it being the first time and not having any baggage. I think we've collected a very interesting group of people. This is not a list of typical suspects for a weekend retreat to discuss digital technology. I think that's one of the things that we're really excited about. This is a incredible and interesting group. So we are excited about that.

We have basically all day tomorrow and a little bit of time on Sunday to talk about this subject, *When Everything Learns*. So what I'm going to do is try to give you some guidelines for how we're going to do it. We're going to work as one group. We're not going to break off into small discussion groups. You're going to get one thing, one exercise tomorrow afternoon for about ten minutes that require us splitting up into groups and we'll come right back together again. But in general, it's just going to be us as a cohesive unit.

What I want to make sure that we do is spend all our time talking about the subject at hand and not about semantics. So that's the only warning I'm going to make—to try and make sure that we don't get caught up in definitions and semantics. And that we can all agree that we're going to disagree about the specific definitions of words, but the intent of what we're trying to get across is what's important.

What I would like to try to do is take what I call a functionalist viewpoint of the world. Which is, something will fulfill a definition if it functions as though it does that thing. Let me try and be a little more precise. A thermostat is a thermostat if it controls the temperature in the room. It doesn't matter whether it's made of a metal coil and a switch, or if it's made of avocados or if it's a guy with a blowtorch that when he gets hot he turns it off and when he gets cold he turns it on. He's the thermostat for the room because he functions as a thermostat. So we won't get caught up in "well what do you mean by this..." hopefully. We can extend that to words like "think" and "learn." If an object needs to think in order to fulfill a function, we're just going to call it a thinking thing—just allow that to be the case. The same thing with "learning." And hopefully that will keep us on track and we won't have to spend a lot of time arguing about the definition of words like "is." That's about the only ground rule that I'm going to have. Other than that, we're hoping that everyone will participate. I will try and do my best to call on people if

they're incredibly quiet and redirect people who are dominating from taking over too much. But, that's about it. I hope that we'll have a nice lively discussion.

So a little bit about the subject at hand, *When Everything Learns*. Of course what I'll start out with is a definition although I just made fun of definitions. It's critical to talk about the word "learn" before anything else. We checked a lot of dictionaries and essentially what we came up with, what we discovered is that to learn is to add experience or skill or expertise. It's an increase in that sense. Typically learning happens one of two ways, it happens through instruction, you're told not to put your hand on the red hot stove because it's hot; or by experience, you put your hand on the stove and you realize it's bad and you don't do it again. Those are really the primary methods for learning. It's that type of concept that we want to talk about. Is this sort of increase in expertise experience or knowledge? To a certain extent, the reason that we chose this is because we think this is becoming incredibly important.

Each one of you got a copy of 'I, Robot' by Isaac Asimov in your packets. The reason I distributed that book is very selfish. I got a copy of the book when I was about ten or eleven years old and was completely fascinated and I absolutely credit that book with my desire to work in the field that I work in and my love of technology. I was fortunate enough to tell Asimov that when I was in college. I'm sure I wasn't the first person to talk about that with him. I got him to autograph a copy of *I, Robot*. I felt really bad because all I had was a paperback version. It was really beat up so he either thought "this is some guy who bought it for ten cents at a garage sale" or "he must really love this book."

The book to me is really fascinating for a lot of reasons. For those of you who haven't read it, it's a collection of nine short stories. The stories are told as a series by a Dr. Susan Calvin

who is the first robot-psychiatrist. She is a shrink for robots. The nine stories are about what happens when robots—who have very strict rules about their behavior—start to develop pathological behavior. How does she diagnose their pathology? The book is fascinating because it touches on all sorts of different subjects. It touches psychiatry, it touches decision-making—each story is her figuring out a puzzle, her figuring out what's wrong with the robot so it's a little bit of a mystery. Each one of them is a mystery novel, a whodunit novel. Each story is also a kind of math problem, a sort of logic problem because it's really about these rules of behavior and sometimes the rules don't make any sense no matter what they are.

But what was even more fascinating to me was how formerly inanimate objects, hunks of plastic and metal really became human through self-awareness and through certain types of behavior rules. These objects became so ingrained in people's lives. There would be a lot of conflict between the roles of the human in society and the threat of the robot taking over that role in society as the robots learn, as the robots get smarter. The robots are not programmed with knowledge and then they stay static. They're programmed with knowledge and the ability to learn and they get smarter through the course of time. The relationship that the humans have with them changes over time.

The book is very dated. I think the last chapter of the book takes place in 1992. They would have been off by probably 200 years in terms of us interacting with humanoid robots, but the challenges that they face in that book and the questions that they face basically does portray a view of the future. In order to get to that future, in order to think about those same concepts, what I want to do is take a quick trip to the past and explain what I think is an interesting logic chain that gets us to the questions that we want to discuss.

If we look back at human history, look back tens of millions of years ago, as far back as we can tell, humans have always tried to improve their lives. Always tried to increase their lot in life. First humans went out and found things to eat and then they killed animals with their hands and then they invented tools to kill, all those arrowheads and bows and arrows. They invented these mechanical objects to augment themselves, to make them stronger, better, faster, more reliable. They also began to use animals to replace themselves. They could carry things around and that worked pretty well. But then some smart guy decided that he could put these things on the back of a horse and he can carry more because horse is stronger, more reliable and isn't so cranky and now the guy doesn't have to do as much work.

So first we have animals replacing humans. Then, we augmented the animals. We put the carriage behind the horse to carry things. It's even better, even more efficient using this mechanical technology to improve the way people live their lives. Then we just replace the horse all together with an engine. In fact the legacy of the horse is still there, we talk about horsepower in the engine. How many horse's worth of pulling are in this mechanical object? We did it because, in a sense, machines tend to be more reliable, stronger, better animals than we are or the other biological animals that we use. Now of course there are still some challenges to this. Biological systems are self-repairing in some ways and there are different energy constraints. But, overall machines are viewed as better for mechanical activities. That has been a trend, it seems, for a long time.

But, we've also done the same thing with cognitive activities, some mental activities. Things like the abacus for math. It's really not a mechanical process, math, but really more of a mental process. Some devices help us think or make it easier for us to think. Charles Babbage thought it would be a really

great thing if we could just make a machine that was analytical, an analytical engine. We just make a machine to do analytical calculation and then we could have a lot more interesting mathematical problems. We have better thinking objects. We have more power at our disposal. The idea was to invent a mechanical device that would replicate not the physical things that we do, but the mental things that we do.

This has been something we have been striving for for a long time. In the last 40 or 50 years, we've made some huge advances because of digital technology. Have we gotten to the point where things actually think? Without even going into the semantics, I don't know. It's close. It's a toss up. Are some of these chess programs thinking when we play against them? Maybe, maybe not. Typically, what we're seeing is technology that is still being used to replace brute force. Even the chess playing games, even the things that we think are smart are primarily just a lot of math. And we don't have new algorithms and new methods and software for really replicating thinking. What we can do is just a lot of calculations really quickly. So when Amazon recommends a book for you, it doesn't really learn what people's preferences are, it just looks at a huge piece of statistics and analyzes it very quickly and recommends a book for us. But it doesn't know you and it doesn't really know what you buy. We're starting to see hints of thinking things. Just the fact that I can mention that example and we can pretend that it's smart and that it's thinking means that we're getting to that point. And it's obvious that people are trying to get there.

What I want to do is spend the next couple of days thinking about what the implications are of that. What's going to happen to us? I think there's a few interesting questions. One is, do we really want this? Do we really want things that learn and things that think? It's interesting to us now, it's new. It might be a novelty. It may be something that we want for a while but it's going to get boring after a while and we're going stop using it. Or,

more interestingly, it may become threatening. So like the people in I, Robot, if the machines really are better than us, if they really are more reliable, more dependable, don't need to sleep, don't need to eat and can think at the interactive level, will they make us to a certain extent obsolete? We're very comfortable with our position at the top of the food chain. That's going to be fairly threatened. It's not clear that our social structure is set up to take advantage of things that learn and things that think. It's also not clear that our legal structure is prepared. We have laws about people and we have laws about property and they're pretty clear. If I am sitting behind the wheel of a car that runs over somebody, I'm responsible for that device and I'm controlling it. Well when the car starts to think on its own, and it runs over somebody, who is responsible for it? Am I as the owner of the car? Or do we prosecute the car? How is our society going to deal with objects that are autonomous or are able to control themselves? And that's only things that think. When you start talking about things that learn, we really start to get into some interesting territory. I think that it's going to fundamentally change our relationship with objects.

The only object I can think of right now that really learns about me, in this abstract sense, are clothes, like jeans. Jeans adjust themselves to me the more I wear them and they become more comfortable the more I wear them. That's very different from almost every other object that I have, every other material object in my life. Pretty much my relationship to everything else is I buy it, I use it, it becomes obsolete and I stop using it. Now if I buy a television set and it starts to learn what I like to watch or it starts to learn that every time there's a commercial on I flip over to the weather channel my decision to replace that television set is very different from what's available today. Now I replace it because I want a bigger one and a faster one or it breaks. But if my choice is to buy a bigger television that doesn't know about me, a bigger, stupider television, I'm going to have to think about that process a lot differently.

I think that our relationships to material objects and the way in which we use them and what we expect from the things around us is going to change a lot. There are all sorts of other interesting questions about how these things are actually going to look. It tells you something the difference between instruction and experience. I'm trying to avoid referring to The Matrix as sort of the ultimate fear of computers taking over the world, but there was a great scene in The Matrix in when Neo learns Kung Fu. He learns Kung Fu in all of eight seconds because they download all of Kung Fu into his brain. That notion of learning as instantaneous learning. The actors studied Kung Fu for eight months, the entire time to make that movie. In the movie it just took eight seconds for him to learn all of Kung Fu and move like a master.

Machines might be able to actually do that and that changes our perspective on what learning is all about. Then what? Do we start giving these machines degrees? Do we give them Bachelors Degrees and Masters Degrees? Force them to have a Ph.D. thesis to prove that they've learned something and did some original research? Do these objects need psychiatrists when they start to learn? When they start to evolve? When they start to change?

It's those types of questions that we're going to be talking about over the next day and a half. That domain, that state that we're interested in. And...don't talk to anybody about it until tomorrow morning when we are all together! I think that there are a lot of things that we can learn about what happens when

things learn from each other. I'm looking forward to doing that tomorrow morning.

Thank you all for coming.

Craig Kanarick
October 20, 2000
Tarrytown, New York



Session 1

The purpose of the first discussion was to get a baseline understanding of “thinking” and “learning” and to reflect on the previous evening’s opening session. We join the conversation just after it started, on the subject of having a goal for the entire summit...

Craig Kanarick: For me the best way to measure success is to think about impact, what impact would we like to have? So, if we have a great conversation, what will we be able to do with that as we leave here? I don’t think we’re here to come up with any patent or a manifesto that we’re going to publish about the future or take up a full-page ad in the Wall Street Journal declaring our support for robots or things of that nature. But I would like to feel that we would walk out of here with ideas that inspire other people.

If we find that the result of this conversation is that more people want to come to the next Science Alliance Summit or that you all want to come back to this group or more people want to join this group, or that more people want to discuss this subject then we have had some impact. I don’t want to burden you with the responsibility to go back and share this information with people and make sure that it is written about. If we’re prescient here and we’re writing about that, then we’ll have some impact on what happens in the world as we approach a world where things do think. I want to feel that we have impact, that we’re here not only make ourselves smarter, which I think is a good thing. We should have a good time and learn things. But it would be nice for us to be able to get out into the world and take what we’ve learned and have some impact.

Michael Ester: The question I had about what the goals of the

meeting are was to give the interaction shape. One of the things that may give it shape might be some objectives that you have. The other thing that might give it shape is the scale of the problems that we deal with. Not to constrain the discussion, but to harness some very talented people in a particular direction. You might have also just taken a particular problem to solve and brought these people here to discuss it. Not that that problem is the most important, but it might help to surface ideas, concepts, and constructs that collectively we might engage.

Craig Kanarick: We tried to focus the afternoon discussions on subjects like “What two laws would we like to see?” Or “What two objects would we like to see?” I know that Tucker has been to a handful of events like this where people tried to invent something over the course of two days. And I don’t know how successful those are. Are they?

Tucker Viemiester: Yes, sometimes you can come up with something or at least spark some idea that later on has some effect.

Ben Davis: I think the motive behind the afternoon exercise was to actually see if we can come to some consensus about a thing all of us might actually like to see created or that might actually be useful—that would have some universal application.

Tucker Viemiester: The best part about that approach is that you focus on trying to come up with something real instead of just blabbing about ideas.

Eric Begleiter: An important way of breaking this up could be time. If we’re thinking five years in the future as opposed to thinking about 30 or 100 years, obviously you’re dealing with radically different issues.

Craig Kanarick: I think what we would probably talk about is the first generation of things, things that obviously go on for 100 years. I was at a conference about two years ago and ironically it was called “Two Years Out.” That was as far out as anybody dared to venture. Things were moving so quickly that any thought beyond two years was like saying 512k would be enough for any computer.

Eric Begleiter: You only see a lot of sociological problems if you start looking out 15, 20, 30 years. It can be kind of myopic to only look at five years. Just looking at the next little interesting hand-held gadget, you don’t have any real sense of its long-term implications.

Craig Kanarick: What is the time horizon for thinking or learning things? Do we have things now that actually learn? In the context of the definitions that I gave last night in “Thinking

About Thinking”—do we have things that actually think? I would suggest that we can pretend that chess games think about their moves. We certainly refer to them that way. We say, “Oh, it’s thinking.” Most people say that about the Internet when it’s trying to call up a Web page with a slow modem. Do we have things that think right now? Do we have things that learn?

Rebecca Odes: What about the Tivo for TVs?

Craig Kanarick: The Tivo is hardware that allows you to program television content. It’s essentially a hard drive for your TV. It’s a digital VCR that records while you watch, so it allows you to pause something. It’s always recording while you watch. So you can pause it, go to the bathroom, come back, hit pause again and it’ll pick up from where you left off. Actually it records two copies of the thing you’re watching simultaneously. So you can actually do an instant replay of something. You can go back on the disc and play something you just missed. It’s still recording. It just keeps going as a stream so you never miss anything. It also allows you to program shows ahead of time so it just like a programmable VCR.

Rebecca Odes: But it doesn’t figure out what you like from the week before and program other shows like those?

Craig Kanarick: The thinking part is the part that’s more interesting. The idea is that it keeps track of what you like to watch and records them automatically in the future. If you like watch McNeil-Lehrer all the time it records it everyday so you can watch the news whenever you want. The interesting idea would be that it starts to learn that you like happier news so that’s what it records for you. But it actually doesn’t do that. It only remembers if you like watching Gilligan’s Island and it records that every week. It doesn’t really do any deep thinking. But it could.

Neil Crofts: There are some cars that are supposed to learn the way that you want to drive, or the way you like to drive, and they adapt suspension settings and gearbox settings. I don’t know how well it worked. But, supposedly that is what they are capable of. I think BMW and Mercedes are doing it.

Joseph Busch: I think this raises an interesting question. To what extent are we talking about autonomous learning as opposed to some sort of feedback? This is an interaction with a human and a human is teaching the computer how to respond—and the computer generates rules. I think there are a lot of feedback systems being built into all sorts of appliances and applications—with or without computer technology. In the short term, this is a big area that I’d like to see a lot more of this sort of feedback being built in.

Craig Kanarick: What’s interesting about what we’re describing is being able to recognize patterns and repeat the patterns. Not necessarily being able to draw any inferences from those patterns, or any rules about why those patterns are happening. You said you like watching Gilligan’s Island, you must like watching Gilligan’s Island. So if Wednesday at seven that’s what you like to do—but without any extrapolation from that pattern at all—then all you will get is Gilligan’s Island again.

Constance Adams: We’re running into a problem with definitions already. I know you planned to avoid that. We should agree what we intend by “thinking.” There’s a lot of that sort of repeating or scanning through data with a predetermined set of things that a machine is looking for and identifying that and doing something with it. And that’s fine, that’s still another step forward. It’s certainly not metaphorical or analogical. So what do we mean then by ‘things that think’?

Craig Kanarick: Well it’s almost like a Turing test of the word ‘think’. If I say that it’s thinking, it’s thinking. If I can act like a

singer and convince you I’m a singer, then I’m a singer. Does one of those robot dogs think? Well it acts like it thinks, so I’m going to say it does, but I haven’t cut open its brain to see if it’s actually a computer.

Peter Lunenfeld: You know, it strikes me that if this were 1966 and we were at a robotic conference, we’d be talking about mimetic human action. I think robotics has really learned that is a sort of dead end. If we talk about thinking and learning machines, we mean thinking and learning as we do as opposed to what I think you were getting at, which is the difference between learning and reacting. Thinking and learning in the mimetic way is far off—like robots who come in and serve you breakfast. It doesn’t interest me at all. But working on various kinds of reactive systems and improving them beyond—if you like “Gilligan’s Island,” well then you must like... tick...tick...tick...Gilligan’s Island! Then you have something that’s a little bit more useful.

People are making billions of dollars creating that “tick, tick” noise. I think that if we should agree to not worry and it—I was a little bit hesitant that we were going to spend a lot of time talking about things that learn and then kill you. I don’t worry about that right now. Thinking things that help save time, and things that think about things you don’t necessarily want to think about—which is the equivalent of labor saving devices that help you do work you don’t want to—seems more useful.

Michael Ester: Things that augment human capacity.

Peter Lunenfeld: Yes. Going back to the sixties again, the difference between artificial intelligence, AI, and intelligence application, IA.

Joseph Busch: There are still a lot of issues. For example so-called “adaptive algorithms.” There are programs that do auto-

matic clustering of things that are similar but it's very difficult to "un-train" or retrain them.

Craig Kanarick: Or train in the first place.

Joseph Busch: A good example is on Amazon.com when you buy a book as a gift. They note it as if the book was something you are interested in and then give you a list of other books that you ought to be interested in. Now they've finally discovered that it's important for you to tell them if it's a gift so they won't suggest Dr. Seuss books for you anymore. There are a lot of issues that are very pragmatic like this. As it was defined it last night, learning is not outside the realm of what we might conceive of as software. If you're learning in chunks that are small enough and that can interact with each other, there's a lot of movement towards more apparently autonomous types of thinking. Much more sophisticated than the kind of thing we see now. I don't think it's hard to visualize things that would really begin to show both the pluses and minuses that we worry about and imagine in learning systems.

Craig Kanarick: In designing products for the consumer market, are learning or thinking objects considered a selling point?

BJ Fogg: Not that I know of. Part of our work at Casio involves generating new consumer products and it's not a direct improvement. I think most of the consumer electronic manufacturers are doing that. It's like the MP3 player with "X" more capability—let's get that out of the door. In terms of learning or thinking devices, perhaps in some research labs but I think in the U.S. we don't feel the need for the luxury of doing that—with Xerox Park closing and things like that. It's actually very sad. There's not a lot of really great visualization in prototyping the future.

Craig Kanarick: Is this stuff that was going on at Interval?

Joy Mountford: At some level, yes. I don't want to go off on a tangent, but I do actually think that it's quite critical that the time intervals that we're looking at in this country are so small and short. We don't actually address the really big problems which are sort of infrastructure at a boring technical level—let alone some philosophical and sociological impact of these things. This should be a country that has that sort of luxury to step back a little bit from its success and actually reflect on those things. I think that even this meeting is a microcosm of that. We need to be able to do that.

I think Marvin Minsky has been saying this for quite awhile—he's been writing a new book for at least ten years as far as I can remember. He's talked a lot about "personalities of the computer." I'd like to go even one step further. I don't know if people talk about computers that have humor, or that lie or worry. Because they're some qualities of what really make humans human. The serendipity of those factors coalescing with information that's factual is in fact I think the thing that computers are nowhere near touching. I don't know if it's important to actually think about it now, or whether you start experiencing little elements of this and see what the big collage gives you. Interval tried to do that, Xerox Park tried to do that and we've seen good examples of good hacks they did a long time ago. But no one knew quite what to do with that for a short-term financial success.

BJ Fogg: I think the best group right now is a group of students at RCA and Art Center. I think their work is pushing the frontiers more than anybody I know of. Of course, the MIT Media Lab is pretty high visibility. But the people at RCA are pushing the envelope of what technology's made to do and how they might fill needs. They're given different design briefs every year and they concentrate to create a user experience prototype in about two weeks, hands down. Then they present their solutions and they move onto a new project. At RCA there's a com-

puter related design program and there's one that's more prototyping oriented. I had four interns this summer from both of those areas and they were really cool to work with. They're not bound by corporate limitations so it was very exciting to work with them.

Going back to learning and thinking issues, I think perhaps the one way to look at a computer system that might be thinking or learning would be one that anticipates the needs of the user, and makes an attempt to fulfill those needs somehow. Tivo is an early example of that. I've been using Tivo for about six months and it doesn't anticipate my needs very well. It's pitched as that but it really doesn't. We have thought of computing systems that would anticipate our needs and be there like a really great butler or sort of our sidekick, saying "Oh, I can see what you're doing and I expect you'll be needing 'X.'" Here it is. It's ready for you."

Stephen Turbek: Even the perfect human butler, one who has worked with you for 20 years, cannot anticipate everything, so we should ground our expectations a bit. We accept mistakes, as butlers are only human. We are not so accommodating with our technological servants, but can we expect them to be better than an intelligent, conscious person?

Margaret MacLean: I'm not a technology expert like most of you—I'm an anthropologist. One of the things that I have been involved with recently is an effort to look at, with the Long Now Foundation, what might be good to have in a library 10,000 years



from now. We've been talking here about a horizon of two years out. Thinking of 10,000 years out requires you to step back a few paces so you see the impact of what we're doing differently. One of the things that comes out this discussion for me is that it's worth thinking about not designing machines that replicate the human brain or that try to anticipate it so much. But, rather to design machines that think in ways humans don't. Perhaps, complementing or filling in the gaps of human will.

For example, what should society be doing in terms of machinery or digital memory or saving things that we're not really paying attention to? We should be trying to keep this in mind, but we're not because we're interested in the marketplace and what's going to be the new thing six months from now. It would be great if the market-driven considerations could be balanced with some long-term thinking. What kinds of protocols could be set up in a machine that would remind us this is important to do now? Or that you need to build in some sort of facility that would take care of the things that we tend not to want to take care of right now.

Craig Kanarick: So the thinking thing, or learning thing, as a watch dog? Helping guide us morally? Or...

Margaret MacLean: A reminder. No one wants a school 'marm' watching over their shoulder, but something that could say "while you're designing this, don't forget about these things."

Craig Kanarick: So it's not just the robot that serves you breakfast, but that makes sure you eat it all?

Margaret MacLean: Well, maybe that there might be a sharing of breakfast with someone else.

Rebecca Odes: But who would determine the things that you

are reminded about? Would you tell it the things you care about that you often forget? Or would it be some kind of higher moral power that is set in the chip with every device that you buy, saying, "remember to be a good person and treat others as you would like to be treated?"

Margaret MacLean: There's probably room for both.

Constance Adams: You got your Big, Big Brother chip going on now.

Ana Echeverry: I think that while we'd really make a difference with these devices and things, devices that know "what is what," I'd really have to see, to understand why they behave the way they do. We need to go a step further. If we just stay on the "what are these things?" then we will keep just getting Gilligan.

Craig Kanarick: Well I think what I'm hearing is that one of the reasons we would want a thinking thing is to augment our capabilities in two different ways. One is to make our lives just a little bit more convenient, suggest books for us, record shows for us, anticipate some of our needs. The other is we would want something to augment our larger thinking—deeper "why am I here?" type of thinking – a moral guide or spiritual guide that we don't have access to.

Rebecca Odes: One thing I think Margaret was saying was about filling in the gaps of human intelligence. Not just the areas that we don't think about in terms of the content, but in terms of types of thinking—like taxonomies of things or creating hierarchies of things for yourself that you might not think of in your everyday work. Computers might be better at creating those kinds of organizations.

Tucker Viemester: I think that we're basically talking about

the kind of things that don't have to learn. You know, we can write a program that will do all those things and the program will send the stuff back to us when we want it. But, what I think is interesting is how the machine learns that stuff, or how we get them to go another step by themselves.

Neil Crofts: Yes. At the moment what we're talking about is basically just plugging more criteria in. The criteria that we plug into any device is only like two or three different criteria. But when we make judgments we're perhaps using a 1,000 criteria—way more criteria around what we're doing. I think if you could offer some sort of access point that fit with your beliefs or with your personal values then that would help you do your shopping in a way that fit in with those values or helped you make other decisions that fit in with those values.

Peter Lunenfeld: Stephen brought something up that I think relates to what Margaret was talking about. We have an expectation of success rates that is much higher for advice that comes from machines than from people. You're really annoyed with Amazon.com when it recommends Dr. Seuss to you. Yet, you notice, you're not nearly as angry when somebody says they really love a movie starring Bjork and when you go to see it and you hate it you just sort of assume "that's life." People will give you all sorts of different advice, but a machine-based recommendation seems to create the expectation, oddly enough, that they should be as reliable as a car ignition system.

I don't think I would ever be comfortable with a set of moral guidelines if it were from a machine. So I'm just not sure about bringing the machine into it. When you said a 1,000 criteria, it's not a 1,000 criteria, there's probably 10 million criteria—that's philosophy. So, once again I think that's pushing these various small little chips to have this mimetic relationship to the highest levels of pattern recognition that humans do so well.

John Weber: The things we've been talking about are relationships between individuals and machines. They can help us do this, they can help us do that. I was thinking about "how smart are our cities?" How smart are the ways in which they grow? How come they're not smart enough? We're all part of them obviously, but how come we're not smart enough to provide housing of a certain accessibility and quality? Why in San Francisco is it possible to wait for the bus and have no buses come for 20 minutes, and then have three buses come on the same line – and the last two are empty. Then if you drive or if you're on the bus, it's inevitable that the gas company and the phone company have decided to dig up the same street on the same day that they dug up last week and the week before that and the week before that. They're digging it up at the same time as the person who is dropping off the towels to the hotel is blocking traffic on the other side and the next morning you go through it again on another street. The same set of conditions will repeat themselves and so the whole thing grinds to a halt.

Machines may not even have to learn to help deal with the situation, but maybe it's possible to create some sort of a network in some sense. A structure that will make the city smarter as a thing because it could respond to its own fluctuations and know how to fix itself on the fly in some way that makes it a better living situation. This kind of condition could be extended in a lot of different ways too—how food moves around and other sorts of things. Thinking about an individual object used by an individual person is fine, but it seems like it does learn toward products and very short-term issues and not towards some of the things that would be more interesting to fix.

Ben Davis: I think another thing that we've learned from the Internet is that all of the objects will be global objects. They're not singularities. Smart objects are in the global network of other kinds of information. For instance, taking cell phones out of their country of origin. The cell phone doesn't know that it

won't work wherever you happen to be. This global dimension is startling in its impact. This water glass works anywhere right now, but if it was a "smart" water glass, it might not work in France. Protocols for that "water glass that learns" would have to shift as you shift geographically. And to further complicate things the smart learning object would have to know how things learn in France—which, of course, may be very different than how we learn in the U.S. As every "thing that learns" becomes global there are cultural implications as well as technical implications.

Constance Adams: I think you're starting to touch on something that's very interesting. Because of the history of technology and what we've developed—chopsticks are just an extension of fingers and that's the idea to augment, to extend—we're willing to accept the development of things that augment us or that advise us. That's an interesting thought right there. But, I don't think we're going to be willing to accept machines that tell us what to do. There's a fine line, but I think a really important one. A Global Positioning System can tell me to try a different street, but it better not tell me not to go down a street because by God I will go down that fucking street—I'm not going to take it from a machine. I see that, that interface right there, where is it helpful? And where is it troublesome?

BJ Fogg: Let me respond to that a little bit—it goes to the core of some of the work that we're doing at Stanford in the Persuasive Technology Lab, which is all about looking at computers designed to influence and motivate people. There are kinds of systems coming that will be self-adapting, or self-improving, self-helped, and so on. A system can help me navigate the city, whether it's persuasion or not. It seems like there is a category of technology out even now that's imposed on people. Like the company you work at, the government and so on. You don't have a choice whether you use it or not. And these surveillance systems—the simplest most blatant might

be the one that let's you see which workers are on the Internet while they are supposed to be working. Surveillance is a way of modifying behavior. You can either accept that or quit your job, right?

Craig Kanarick: It makes me think of, I think it's New Zealand or Australia, that requires every citizen to vote, which I think is unique. It's like paying your taxes. But the ultimate choice, the ultimate democracy is you have to make a choice. You could choose to abstain but you have to choose to abstain. You have the option to participate in this very free society that forces you to participate in this one activity. It's almost the danger of taking persuasion or force too far.

Ben Kleinman: You're not going to be motivated by the car that says "Go down that street." But a real learning machine will learn that's not what you personally respond to. Maybe in one case it's a more tactful statement such as "Why don't we try that one over there?" or "Hey, there's this great sight to see over there, so why don't you try that?" I think that's the one level more of distraction we need to potentially consider in the learning.

Rebecca Odes: Manipulative machines.

Tucker Viemeister: That's about how the information comes out back to you so...there's millions of ways to trick you into doing something. But, I think that our discussion is on the other end.

Peter Lunenfeld: TV does that already. TV is a box that tells you what you really need is a car that's better than the one you've got. That your skin would be better if you did this or that and we suggest these pads and these strips.

Ben Kleinman: There are people behind the scenes on TV

structuring this manipulation. The machine itself is not adaptive enough to learn that yet.

Peter Lunenfeld: Someone's going to have an economic interest in which street you will go on. There are going to be people who spend a lot more money convincing your machines to convince you to go down that street because that's where you will spend money. That's what the Zagat's guide is already. There are restaurants that aren't listed there and there are reasons for that. It's a somewhat naive assumption that machines are going to be neutral or that their recommendations will be neutral.

Craig Kanarick: We started talking about machines augmenting our behavior and what I heard Joseph talk about was the issue of autonomy. Will accept the machine (object or thing) that will augment human behavior? But we won't necessarily accept one that will replace it or be a counterpart to it. So a machine that improves my life is one thing. But a machine that just goes off and does its own thing, is completely autonomous, and decides that it's really important to clean up the environment and just goes out there and cleans up the environment on its own is something else. Isn't that, to a certain extent, a natural extension of things learning? They're going to make their own decisions? It's not that we tell the machine "just get really smart," but "check in with me before you go and do anything." That can put a limit on their learning. But what we might want is something that learns so much that it figures out it doesn't need to ask me what to do. You like to watch Gilligan's Island so you go to Gilligan's Island and extrapolate that to "it's a good thing to clean up the environment." The machine may go out there and clean up the ocean and just do it whether anybody asked it to or not.

Joseph Busch: Robots can kill the human race. You can have rogue algorithms—that's what viruses are—and you can have

rogue machines. Some of this can be good. You can invent bacteria to chew through some nasty stuff. We're going to let bacteria that chews nasty stuff out into the environment. On the other hand, we may decide that genetically engineered corn is not good so we're not going to allow it into the environment. But there are things that are being allowed out there as autonomous things. They can be either mechanical or biological. We should be thinking—is there a difference between a biological machine and a computer chip driven algorithm? There are differences, but there are similarities too.

We already have algorithms that we use to do things without our control and without our feedback. Surveillance is one of those things. This sort of a viral thing of deleting and cleaning out garbage on your computer is already being done. That's great. I don't want to think about cleaning the garbage out of my computer. There are these autonomous things, but they're not very smart right now. But I can imagine that they will begin to have some reasoning so they're more autonomous, so they are making choices and they begin to aggregate. Some of them are good and some of them are bad, some of them we want to think about and some of them we don't.

Stephen Turbek: It comes down to whether the technology is making choices or giving suggestions. You would not be pleased if your refrigerator starts ordering soy milk, because that's what you should really be having, instead of your favorite chocolate milk. How does it decide when your long-term goals and your short-term goals are in conflict?



Craig Kanarick: It sounds like what we have to do is control the learning. It's one thing to say, "I want you to learn what's good for me." Not, "I want you to learn what society thinks is good for me" or "what something else thinks is good for me." So we can still have things that learn if we just put a box around the primary motivation. To a certain extent, that's what happens in *I, Robot* by Isaac Asimov. The primary rule or robots is "I don't care what anybody else tells you, or what you think is important, you can't hurt human beings. You can't allow human beings to get hurt, you can't hurt human beings no matter what you do." I've read the sequel and I don't want to ruin it for people. A 100 books later, the robots have learned that there's actually a rule more important than the primary rule they are ingrained with which is "you can't allow humanity to get hurt." This was Asimov's learning over the course of 50 years of writing.

The primary rule became "don't allow a single being to get hurt." You can order a robot to punch through a wall, but you can't order a robot to hurt humans. The robots evolved, they learned. Asimov's philosophy evolved to humanity. I want to get back to the subject of "when everything learns." Say this thing is at the service of me and I'm defining what the service of me is which is to do things that I say or I want, not to learn these other things. Is that fair? Is it?

Peter Lunenfeld: I'm going to combine two laws: Moore's Law and the Law of Unintended Consequences. You get Moore's Law of Unintended Consequences. As computability increases—you know the whole Moore's thing, that eventually the law of unintended consequences moves to where all the rules that we've been talking about, all the biological elements that we've been letting out, all of this.

I love what Joy is saying about the trickster computer but what are the unintended consequences of engaging with the trickster? Maybe that's something we have to learn now as a

species. But we're talking in really a 19th century instrumentalist way about technology in that, "I tell you to do this. You do that." Which doesn't work with children—hence all the embarrassing moments—and it doesn't work with machines that aren't intelligent and learning in my opinion, but just reactive. I mean you put in the input and if you have only one reaction, that reaction can fail partially because the machine will eventually fail. There are no fail-safe machines. I think that's bad.

Ana Echeverry: I would just like to bring out another theory which is Darwin and evolution and why things survive or not. If there's no life there is no purpose. We all need to talk about computers and things that think with the Zagat Guide inside. The issue is the machine, it's in Zagat and people liking it. So it's really about what things stay. What things keep growing? If they serve my purpose they will stay with me and they'll know what they can do, what they want, and will have the choice to find another owner or go by themselves.

BJ Fogg: I think there's a bit of an assumption in our discussion that what individuals want will be the technology that will be developed. That's not true at all. Let me suggest maybe some categories for a slightly cleaner discussion. In the question "why do we want smart things" that "we," I think could be interpreted three or four ways. One is, "we" as individuals—and I think a lot of our discussion has been around "why do I want this?" The next "we" might be a social group like a family. And their reasons for wanting thinking things might be different than an individual's. Another "we" that's quite different is institutions. Business, for example, or non-profit institutions. Institutions are the drivers of what exists today. Their motives for smart things are very different than individual motives.

Yet another "we" is the community and larger level society. I think society's motives or reasons for having smart things are somewhat different than individual ones. So I guess if you were to boil

those four down there's really three, there's individual, social or group level, and institutional level. Ideally, human needs would drive the innovation process for the institutions. But there's a disconnect as we all know. I mean, if that were true, I'm sure Microsoft Windows would not be the monopoly that it is.

Peter Lunenfeld: Triumph of media!

BJ Fogg: Maybe 20 years from now, the evolutionary human needs will drive the innovation of today's technology.

John Weber: That's how you really extend that—you can conceptualize in many different ways from the object to the system. Some of those may be actually sideways systems that do require a certain sort of input from people. There are elements that are built into it that people simply can't do, whether it's got the scale or complexity or speed. But you're still building some system that requires both kinds of "thinking."

Craig Kanarick: I want to go back to this notion of unintended consequences because I think that when you talk about things that think, or things that are smart—there's an inherent difference between things that think and things that don't think, but also there's a difference between things that learn and things that think. I'm wondering if that is when the law of unintended consequences starts to appear? We can program something to think, it's controlled, it stays static and we can be dynamic around it. We can choose whether we want it or not through an evolutionary processes or just direct choice. But, what if we get to a point where we don't know what's going to happen when things learn? My guess is that learning has an element of unintended consequences—a sort of history for what's going to happen next—and what happens as it develops new knowledge and removes some of the predictability? Is that really the risk factor? Are thinking things OK? But learning things is where all the danger comes up?

Tucker Viemeister: Well, it's either a danger or that's where the safety comes from because hopefully if it learns, then it's not going to do the same stupid thing again. I mean the problem with unintended consequences is they (learning things) endeavor to figure it out. Something is going good and then what happens a 100 years from now and it turns out it's the worst thing that happened? I think the fact that the stuff is supposed to be learning is that it can take those things into consideration. It's not just like setting it off on a course—it ends up changing.

Neil Crofts: Well, my interpretation is that we're already living with enormous unintended consequences of everything that we've been doing for the last 10,000 years.

Eric Begleiter: Of course there is the example of the unintended consequence of cars, but still, I suppose it might have been possible for people to have looked much more carefully at the long-term consequences of the internal combustion engine from the start. A lot of information was there as to what effects they would have. So, to some degree, it's like saying "well, gee, all these (consequences) are sort of unintended and they all sort of just happen to us and so what can we do?" Sometimes inventions evolve in ways we could not anticipate, but many times we do not want to think about it, we do not want to make the connections.

Neil Crofts: That's true. That's where we don't consider the negative aspects because we don't want to because we're only focused on the positive aspects. Especially the fact that it's corporate type institutions who are selling them to us who are absolutely focused on the positive consequences.

Eric Begleiter: Right.

Neil Crofts: The second thing is a whole lot of even less obvious consequences that perhaps we even don't notice now

about cars. I don't know, levels of social division caused by people living in one place and working another place. This has reams of other consequences.

Eric Begleiter: I wanted to bring up two ideas. One is just the way in which we anthropomorphize objects and computers. I can think of two interesting examples of this. One is if you have a questionnaire on a computer. It's been found that you may get a better response if you ask the person to go to another terminal and fill out the questionnaire rather than the one they originally saw the questionnaire on.

There's this anthropomorphic projection that you don't want to upset the computer you were using. A lot of the sense of something being intelligent is our projection on to it. The second example in terms of anthropomorphic projection is that it's possible to create a very high impression of intelligence by having machines purposely making mistakes and ask you for advice—like “What did you mean by that? Or did you in fact really ask for me to do this, or do you want something else?” You can actually create programs in which the computer—even if it understood, will ask you this question to create the sense of its being more conscious.

The other idea I wanted to mention goes back a little bit more to a definition of consciousness itself, and how we think about it. If we imagine ourselves going back through different species, other types or levels of living things that are conscious, I think that to some degree, each one of these levels of consciousness is able to construct more or less complex models of what's outside in the world. We ourselves are on a continuum, beyond us is the ability to know more about the results of our actions, to cohere or model more space-time as it where, before us there is less understanding of how to do this. Intelligent things we create will be on that continuum too.

If you go back below that, below a two-dimensional structure to a one-dimensional consciousness, then everything would be sensation. If you have sensation, perception, and conception as coherence of the first, second, and third dimension, perhaps that's a continuum of consciousness. So the question is: What's the level of complexity of each of these systems? Modeling things in terms of whether they're one-dimensionally coherent, two-dimensionally coherent, or three-dimensionally coherent. Then you would have something like four-dimensional coherence, which would be organizing the results of the specific action or material. It is a sort of continuum of things as being aware, conscious through these levels.

Constance Adams: But is consciousness a necessary aspect of a learning thing? To me my easiest grasp on all of this in a larger, bigger sense is a very direct interest in creating a universe in which everything learns. Namely, we are going to send a crew to Mars and get them home? They can't be surrounded by the fallacy of “garbage-in, garbage-out.”

What a learning thing ideally does is get rid of that algorithm. Right? Garbage-in somehow gets corrected in time. I mean that's a big liberator. I've been covering a lot in the process of thinking about and working on the early stages for a concept on Mars exploration. You've got all the problems. I want fabrics that heal themselves when a micro-meteoroid penetrates the insulation layer on the outer shell. I want a storage box that automatically upgrades what my inventory is on items because six people are spending three years together. Everything has to be somewhere and you've got to be able to find it in a reasonable amount of time and know what the consumable level is. I need materials that remember what their latest experience of pressure and temperature have been so that they can keep the whole vehicle knowing what the story is because the vehicle has to be as close to a living organism as possible. It doesn't have to be conscious. It could have

had a lobotomy as long as all of the sensing data that the body has is there—the body has learned not to touch that hot burner a second time. I don't need consciousness. I need everything else.

Margaret MacLean: So then, are we back to just what Neil was saying? Is a learning thing one that is just able to process an enormous amount of data and develop patterns and make the right choices?

Constance Adams: And optimize that.

Neil Crofts: The other bit of that is specialization too. On one level you're talking about things that are learning in a macro sense, “everything that I want.” But at another level it's learning very specific things. “How do I fix myself after I've been hit by a meteorite? Or, it will know what drink I want to put in this glass?” So there's moving between those two levels as well.

Michael Ester: There's seems to be huge gulfs between community on the one hand and the repair of a fabric on the other. I'm just trying to think in terms of establishing value on this topic. What are the thinking things that I value most? It's my colleagues and my staff. That the excitement of managing any kind of project is being able to leverage more than you can possibly do—have them carry out tasks creatively to the goal that you have in your head. That's the excitement about development. So what are the thinking things that can help mediate that? Help improve that and extend it?

For me, this is the ultimate goal of the topic, in terms of my personal satisfaction. That can reach into the social community, communal factors—as well as the practical, almost operational things of what their machines do that enable them to work better. What specific things improve our communication with one another?

Craig Kanarick: Maybe the first thing we should talk about is people. People actually learn. We say “why don’t you learn now?” or “why don’t I learn now? Why do I keep making the same mistakes over and over again?” Maybe we shouldn’t worry about having a car that learns, or a computer that figures out what kind of email I’d want to read. What happens when every person learns from their mistakes or even just from their experience—and goes forward? It’s the same question about the city. Why would we want a city to learn, rather than just have a city? The difference between having a city that learns and having a city that’s just really smart is that cities change.

The reason that we need a city to learn is that we can’t set up a single set of rules that are simple. Rules that say here’s what you need to do, you need to do it this way. But a city is such a dynamic system that it must adapt to change and change those rules as it moves forward. I think that’s the interesting part about learning that cuts across whether it’s an object or a human or a society. Last night, I said learning was gaining knowledge or understanding. But it may not necessarily be “gaining,” it may just be a shifting of knowledge, a change.

Eric Begleiter: Learning could also be seen as a type of filtering. In many ways the definition of learning has changed perhaps more in the direction of over emphasizing information. A definition in which having more information is synonymous with more knowledge. That’s not really the case. Certainly you’d often rather have more information than less, but more importantly you want to have some type of filtering of that information too, in order to have some heightened understanding of it, to have some wisdom about it.

Tucker Viemeister: OK. We’re obviously never going to come the point where we know everything. Everybody’s going to have to keep on learning, otherwise we’re screwed. I think that

nobody’s thinking that this is all you need to know. I don’t need to know anything after this. You know, if you learn this, then you’re set.

Eric Begleiter: But still the ability to have some level of filtration, to be able to prioritize and organize information is necessary.

Tucker Viemeister: Well obviously there’s an infinite amount to learn, therefore the only way to deal with this is some kind of filter.

Joy Mountford: What’s bothering me a little bit is that I’m not sure that we’re able to think or problem-solve anymore in the large context. In addition, keeping a perspective so that you aware of any consequences. The way that we learn control is to think of things as tangible and smaller and smaller. Depending on the circumstance, the quantity of stuff is there so you then start to channel down to smaller and smaller piece. What I’m concerned about is companies actually focusing on the smallest elements of something. When you ask the very important question of “when everything thinks,” who are the people who are going to talk about the bits and pieces working together? I actually don’t think I know. Because the way I’ve learned to understand anything, which has changed in a rapid way, is to get small about everything and control it.

To get to the level where we really have to start thinking about it all being like a Matryuska doll is very difficult. I’m actually incapable of understanding it until I go back to what happens with these things and build up. Now there are obviously different philosophies on how you approach this. My question is can we bridge the gap between the technology approach—which is a bottom-up approach—and this sort of social worry of the city? Somehow, is there something between these which can be somewhat planned and thought about systematically? Because obviously, there’s a lot of stuff that will happen. Do we just give up and say, “well just let it happen then, because it’s

happened so fast we can’t control it?” Or write about it in a deep philosophical paper and panic. Which is the discontinuity that you fear, horribly.

Camille Habacker: I think that what really interests me is building something that keeps me aware of the larger implications and gives me practical, everyday choices that I can make to help either avoid those implications or mitigate those implications. I was really blown away the other day when I was at home in the afternoon watching global news on PBS and they were talking about how scientists in Great Britain and the United States were having meetings with their governments and saying, “remember all those models that we showed you about global warming years and years ago and we told you that these things were going to start happening and we talked to you about what we should do to avoid those things from happening? Well, guess what—they’re happening.”

They showed all these facts and basically their outcome, their recommendation was to cut carbon monoxide production by 60 to 70 percent which will never happen. So it’s totally freaked me out because you have this intense reaction—the world’s going to blow up unless carbon monoxide is cut by 60, 70 percent—and nobody’s going to do that! So what do you do? You’re either going to live in a world where it’s a 100 degrees hotter than it’s supposed to be, or try to figure out ways we could have been smarter about getting dumb, money-grubbing people to do smarter things to make life better in a more diplomatic, more persuasive, more “I-feel-your-pain-but-really-I’m-convinced-you-are-manipulating-me-to-do-something-else” way. That would be fascinating.

Stephen Turbek: If we’re going to make useful things that learn, they’re going to need an objective, consistent framework to make predictions. Currently, AI systems, such as Amazon recommendations, making dumb little tiny guesses based on

your most recent actions, which are not going to achieve your long-term goals. We often act in spite of our best interests. I'm not sure we can create a rational system to resolve all these irrational conflicting desires.

John Weber: But that's what we do everyday.

Neil Crofts: Isn't that the thing? That if we want to take account of all the consequences or actions, we simply wouldn't do anything. We couldn't do anything.

Rebecca Odes: Well that's how a lot of people live.

Camille Habacker: I'm not saying that computers or things that learn are going to make us solve this. It's human arrogance to think that. But even just a small amount of help from things that learn I think would be a big win.

Rebecca Odes: You can apply that same thinking to a personal level. I think that people do know that they want to accomplish whatever it is they want to accomplish. Whether it's not smoking or doing the things that are on their list of things to do. People use Palm Pilots for this to try to remind themselves of things to do. But maybe part of what these learning things can do in augmenting people's abilities is to somehow process and learn about priorities—learn about the things that they do and how they relate to their priorities in some way that the person themselves doesn't do.

Neil Crofts: Sort of extra will power. A tool that won't let you take a trivial journey.

Craig Kanarick: So every time you wake up after a night of heavy drinking, you say, "Why don't I learn? Why do I do this?" It may be that your little gadget says, "Just remember, you're gonna feel really bad in the morning." That becomes an inter-

esting question. Will that help? We may have these things that learn, or learn for us and then we say, "yeah, yeah, I know, I know. Thanks for reminding me about it." A sort of little mom, the little guy on your shoulder that says, "You know you said last time, 'why don't I learn?' I'm helping remind you." Some people have that inner voice. Maybe it needs to be louder. Things that learn can remind us of things we've already learned, already learned that we forget, or that we need reminders, or we say that we want reminders. In some ways that is an interesting conflict. People don't want a school 'marm' over their shoulder telling them what they should or shouldn't do, yet everything we're saying around the table says that we have to do all that. Is that the sort of inherent conflict that an anthropologist must have to deal with?

Margaret MacLean: I think it is. I think it's also a classic problem between what you want and what you need. I think both of these things are being talked about, some have more compelling energy behind them than others. This is exactly what I'm talking about.

Tucker Viemester: But I think that the difference between this virtual information we're talking about and real life is that in real life there's limits that our bodies know about. What I think is going to be interesting is when this information is bent back into all these things. You try and pick up this coffee cup and it knows that thousands of people are being killed in Argentina picking coffee beans, so it's very heavy. That kind of real feedback is going to be very interesting.

BJ Fogg: Its helpful to know that there's about 50 to 60 persuasion strategies out there. Not all of them are school 'marm' like reminders or nagging. There's about 50 others. For every strategy, there seems to be almost an infinite number of implementations computers can use to leverage that. Praise as a persuasion strategy can be implemented in many ways on a

computer. So if we do select technologies to remind us—it's a bad idea to drink until 3 in the morning—then there may be technologies that learn which strategies work for me as an individual, and it will adapt—it will choose one of those 50 that best work or move onto something else.

One of the strategies, at least for people that have high precognition, which is most of us in this room, is showing cause and effect relationships. Computers are great at showing that. Manipulate the causes, show what effects come out, whether it's global warming, or physiological behavior, or whatever. That seems to be a really promising area of computing systems starting to learn how to motivate people, but also then to display, to implement those strategies in some sort of way that may work. We're just in the very early stages of understanding how to do that, though that's the kind of landscape we're looking at right now.

Craig Kanarick: It seems like one of the things about learning is...making those things functional. If there are 50 persuasion strategies, and I need to program some box to make it do that, what I really want to do is have that thing learn which is the right one to give to me. Because if you can't program it and say "try one of every 50" or give it a rule specifically that says "if the person before used this strategy then..." unless there is some rule like "if you're a Capricorn, use this type of persuasion." And maybe there are those rules but they're undiscovered yet, but there's a good sized chance that if that thing learns what works and what doesn't work. If it has that system that then, it will make the right choices unless there is an error in the sense of some feedback loop for it to understand. So awareness may not be, or consciousness may not be a necessary condition, but it sounds like a feedback loop is absolutely a necessary conclusion for this thing to evolve. Because it needs to take it in and put it in adaptive behavior.

Peter Lunenfeld: The feedback loop concept is something that's been gnawing at me all morning. We haven't talked about maintenance. And I think that's one of the things that as more and more things are supposed to be able to talk to each other, or the things that we all know from our own experience is that there are generations of these interfaces and that getting even two things that are designed at exactly the same time to talk to just each other is an immense problem. The hours that have to be spent just getting a new printer and you're buying a computer at the same time, even now can take you forever. It's better now than it was. I'm thinking of about Constance's Mars shuttle now. I think that there's going to be a lot of emphasis put into sending people to Mars so that they're all going to be talking together and there's going to be a vast number of people working on a set of technologies all at once to try and integrate them. But then you're going to be dealing with exactly the same question. But when you're talking about a long-term issue, and you brought it up last night, your TV knows you, are you going to sell that TV?

Craig Kanarick: That was the setup. That was a setup because I saw at least one person in the room cringe thinking, and I know what they're thinking. Well, if we can invent a TV that learns about me, then your TV will just learn from the old TV what my favorites are, and that was why I set that up because we know that's not going to happen. If I could copy preferences about your desktop color from one to the other, you'd have to reset all of those things. There's no way a TV is going to know—you know, the new TV is going to say "hey, old TV, what did he like to watch?" The old TV is going to sit there and say nothing.

Peter Lunenfeld: How we can lessen that burden on the individual, to get back to that "we." The immense waste of time over a 30 year period of bringing I.T. into the work place. The endless number of hours that people spent not working, but

rather working on their machines. And only at a point now that we can start to build into the design of these objects, the design of these systems, understanding that if maintenance is built into it as opposed to an add-on. I'd rather have something that didn't learn as much, that didn't think as fast, but got up to speed quicker and maintained itself better. I'm willing to trade brute intellectual strength for the ability to fit into the network quickly on the assumption that I'm going to get rid of it in two years.

Craig Kanarick: We do want the machine to learn from the other machine, not just from the user of that machine.

Neil Crofts: Suppose we did all this and we freed ourselves up lots of floor space, what do we do with it? It seems most of what the past society is about is about distraction of living. So do we get to do anything profound?

Rebecca Odes: Make art or something, there's plenty of things that I could be considered profound.

Constance Adams: It's a nice idea; we've had this notion that a time-saving device is going to give us so much more time for how long, the past 100 years at this point? Don't we all experience how much more free time we have? We have all these devices and all this distraction, and less and less time for profound or meaningful activity.

Sue Madden: The law of unintended consequences.

Neil Crofts: That's the law of capitalism!

Peter Lunenfeld: There's no unintended consequences. Marx said what you'd do, you'd fish in the morning, and you'd write in the evening. It's not a particularly utopian moment right now, and I'm no Marxist, but it's still a good goal. Right? Fish in the

morning, write books in the evening. That's what I'd like to do. And in a way, machines freed us as a group from tending a factory right now, which is what my grandparents did. And so in essence, it's not a bad thing—if we're going to think about Robby the robot, at least let's think about what machines have been capable of doing. And you're right about distraction, but that's the moral and ethical implication. I think moral and ethical issues in our own culture are too pro-distraction and not enough towards contemplation. But they'll shift.

Camille Habacker: I think that's possibly a question of control over the distraction. Don't own a cell phone. I don't own a cell phone. What's the big deal? It's a choice that you make, right?

Joy Mountford: A thing that concerns me a lot now is, where are the opportunities to practice thinking? Because so much education is based on learning and always has been. I actually think the objects in our life should learn. I'm actually relieved that I don't want to learn. I'd rather think. There is too much obsession about I have to know this and this and this before I can have a meeting at Razorfish about things. I got the reading list and started to panic to read all these things so I can be intelligent. And then I realize, well, you're not going to do that. You're going to just think for a change. And then I realized what I hadn't done is thought for a long time because I tried to do problems like networking a printer.

And I actually believe that if we don't step back from it, allowing students to spend a bit more time thinking as opposed to learning. Students know now not to spend a lot of time learning tools. Tools will be wrong in two years, ten years, 50 years. I never learn anything about what I do in my job when I was in school. You learn to think. And many English philosophers have said it doesn't matter what you learn, you just have to learn how to think. And I really am concerned, I'm not thinking as much as that for some now 20 years as I used to when I was a student.

We all want to be reminded about the consequences of drinking. I know, because what I really want to have to do is every day get blindly drunk without knowing what happened at all because I should be able to enjoy that blasé feeling at the end of the day. And I'm concerned that we don't do more of that, hell-bent for leather kind of let's find out about the consequences as opposed to knowing that.

Craig Kanarick: We all know where we're going after dinner.

Joy Mountford: Our impression of children really has to do with that. They don't know the consequences yet. And a certain amount of that is exercising the mind—it's a muscle. And I'm actually concerned with just not using it very much and they're incapable of actually thinking at the Long Now because we're obsessed with the Short Today. And I think the only contradictory belief I have about the Long Now society is that they're too far away from me to understand it. Ten-thousand years is just too much for me. A 100 would try and be the better purpose for me.

Craig Kanarick: Joy, are you saying that there's an inherent conflict between thinking and learning? If we spend all of our time learning, we don't have time to think?

Joy Mountford: Absolutely. I try very hard not to read, not to use the Internet is an active pursuit of mine. I do things like needlepoint, which is ridiculous, right? You can't imagine how many men I've met who say "oh it's so refreshing to see a woman do this." In fact because I do needlepoint, I actually spend most of time thinking because needlepoint is a mindless activity. Can you actually look at the Internet and browse it mindlessly? It's a little different than needlepoint.

Peter Lunenfeld: But that's a nightmare. Needlepoint is a practice. And I think there are huge differences between a con-

stant state of entertainment and a practice. That there's an engagement that goes beyond the sort of glaze. I try and pick up things, not needlepoint, but other kinds of practices.

Joy Mountford: But there's a lot of interest in gardening now. We know why it's high tech, high touch, this is a counteracting force. And I actually believe that the things should learn and we should start learning as much. What do our friends do? I found someone that last night and said, "can you summarize 'I, Robot' because I need to know about it." And they did it. And that was a lot more useful than trying to read the book in three hours.

Craig Kanarick: It seems like we have a capacity issue, then, for us to learn and think at the same time. I mean, one might argue that while I'm learning, I'm thinking about what I'm learning and thinking about whether it is that relevant or not. So you're right, if I spend all of my time trying to learn how to speak French, I won't ever speak French, or replace whatever skill it is. We're living in a world where we need more of that. In order to use modern technology or to function among society, you have to learn a lot more than what someone had to learn a few 100 years ago. You got up, you picked your fruit, you ate it, you went to sleep. Or, you had your trade. You didn't have to learn how to operate the computer, how to operate a car, how to operate an ATM machine, how stop the clock blinking on your VCR. You didn't have to learn all these things. And I wonder is there a capacity issue for us? Can machines help get around that capacity issue? And learn and think at the same time?

Eric Begleiter: If you look at the time when books were the primary medium in terms of the information tools that we were using, it was so much more contemplative in terms of the methodology, maybe it was harder too. As we change those tools, the kind of underlying public discourse we have and everything associated to it changes. That's probably one of the most fundamental, unforeseen problems in using these types

of tools. It fundamentally changes what it means to know something and to think about something.

I think that it also pushes people towards a more, almost sort of dream-like state of perceiving, as opposed to going through more difficult structured forms, and you can see that in terms of U.S. politics and other things. They are being driven a lot more towards a sort of more superficial short-hand level of analysis. You can also clearly see that in terms of schools—the desire to have lots of computer literacy but not necessarily having the kind of analysis of information that tells us what the effect of these things will be. That's a fundamental problem related to the modeling of other things around the tools that we use. There are a lot of these other questions in regard to what's the ultimate effect on us will be.

Michael Ester: You talked about a subject that for me is really something I've been grappling with—not in the context of "when everything learns," but just in the sense of how do we learn? Education hasn't really changed for about 3,000 years. It's one smart guy standing in front of a room of 30 people all writing down about half of what that person says. That's the vast majority of education. That as you go up in education, that changes. You move into high school, college, graduate school, you start to do more experiments and more hands on education. But below those levels 50 percent of learning is pure verbal instruction in this sort of lecture mode, that sort of ancient Roman education.

John Weber: Yes, it seems like that's so true. Although there are other newer models that are being used a lot in the schools now. More what they call constructivist learning which is very hands on right now. But what I'm interested in is learning and thinking and getting back to what Joy was saying about wanting to have time to think, not willing to have to learn. I think our urge for learning is very structured.

One of the things that interests me while working in an art museum is what's going through people's minds when they're standing in front of all that stuff they're seeing. You know, whether pictures, or video installation, or what have you. People in museums now are thinking a lot about informal learning, which is almost impossible to figure out. If it's happening, or when it's happening, or how it's happening. We know it's happening, but it's very hard to pin down, it's very hard to test for. You can't figure out when it happened. You may learn something and not know you learned it until five years later and you look at a picture that suddenly means something to you that never meant anything to you all those other times you looked at it. But at a certain point, your brain picks up enough—I think a lot of it has to do with aesthetic patterns and having a sense of them and what not. It's where the line begins to blur between learning and thinking.

Craig Kanarick: It's funny because I have direct experience with that in museums. I used to complain a lot about museums—traditional art museums—and said I don't learn anything when I go there. They hung a bunch of pictures up in a straight line on a wall with one paragraph at the beginning and why should I go to a museum? I don't learn anything. But I realized you don't have learn something, you can just go there and think, or you can just go there and appreciate it or you can go there and—you know, for me, it was all about, "What am I supposed to learn?" The curator has this objective that I'm trying to teach you something about this artist or this movement and I've done it because I've crafted, I've put these in this linear row in exactly this way and put the paragraph at the beginning. And for me, well I can do that out of a book. It became very frustrating for me. I only let go of that six to nine months ago, actually going in and it doesn't matter to me. It really doesn't matter if you don't learn anything because that doesn't necessarily need to be your objective. Even further, it doesn't have to be that structured in that formal.

John Weber: If you had a good time, you've learned something. Even if you had no idea what it was.

Rebecca Odes: It's not about the knowledge, it's about having an experience with the object.

Joy Mountford: One of my favorite people actually is Frank Thomas. He's the original Disney animator of many, many scenes that captured our hearts and I believe will transcend various cultural and social statements that we're making here. And I often talked to him about computers and the value of them versus not. And his big concern is that a computer is too fast. And the reason that he doesn't think they're good for everything, he says, "There wonderful for scenery—you can do lots and lots of trees and fill-ins really quickly" and that was a really difficult thing for because he had to keep painting them when he really cared about who was in front. And he said the good news for him was he didn't have those tools (computers), he had to actually refine the level and the nuances of every character to the point where they really become very sophisticated, such that adults can enjoy them just as much as children, etc. And his belief is that the cultural characters that we built more recently, they pass really quickly. They come and they go.

The reason they come and go is they're quite thin and veneer-like. When he has to start thinking of a new character he said he likes to wake up, shave like Dopey, or try to drive like the snake in Jungle book to really understand the characters.

Craig Kanarick: I've seen him on the road. I've seen ad campaigns in New York about, I think it was about a computer or something and the tag line was "Write faster than you can think."

Joy Mountford: I love those who say there are three main ways

of looking at this. And I would like to be like that, instead of a "blurter." But I actually think that Craig has a very important point because, in a sense what you want the computer to be slow, to be an imparter of the things that you might want to know.

Ben Kleinman: I was thinking about the tacit knowledge aspect. If you're thinking faster, writing faster, your unconscious is moving and is ahead of you. Do you think a thing when you say it or do you think it first and then you say it? An example of this is the time Bruce Lee was asked, "If someone attacked you with a knife and you killed them are you guilty of murder?" He replied that, in a court, he would have to proclaim his innocence by saying "I didn't kill him. It killed him. I reacted before I could think." Is it going to become sort of instinctive learning? As people, we learn and it becomes instinctive and internalized. We no longer have to remember that we have learned that thing because it has become part of us. What's the level of complexity we want things to learn?

Peter Lunenfeld: Certain social rules can come up that try to do exactly what Joy is talking about. I think the slow food movement in Europe is one that's really interesting. It began in Italy and basically it was a group of Italian intellectuals, cooks, saying that they really hated everything about fast food. They hated its impact on young people. So they said, "let's start a new movement. Slow Food." And they sort of adapted various recipes so that slow food wasn't incredibly slow like 12-hour cooking food, but was like one hour cooking and then sitting down and having dinner. It's actually spread throughout Europe as a kind of cultural way to respond to what they see as a globalization that affects them in real personal ways, not just about how they're at work, but how they are at home. And I think that there's nothing wrong with thinking about that.

Thinking about how things that learn can also teach us how to take certain amounts of time. I had a student who, when I told

him I was coming here, said if you're going to have something that thinks, how about a coffee table that thinks about Nietzsche while you're off doing something else? So at least somebody in the house is thinking about it. Even if it's just a coffee table.

Sue Madden: How reassuring!

Neil Crofts: I want to come back to something we discussed a little earlier which is about having to learn to do lots of different things. If you could get these things to do some of that – thinking—then couldn't you just perform a few activities that you enjoy and the outcome would be all that other stuff that you don't enjoy?

Craig Kanarick: That's the augmentation value. The augmentation part is great when it's applied appropriately. If it does free me up to do everything, if it does persuade me to do the right thing, then it's a good thing. If it does all these other things and they're unintentional that makes me more stressed out, it makes me worry about learning more things, then it's a problem. And I we'll ask tomorrow morning are we optimistic or pessimistic about where this is going to go. Because obviously, we haven't decided yet. But I think what we're filling out are these two categories here. If I had the friends who can summarize all the books for me. In certain cases I might not need them, but in the times that I do want to slow down, and do the thinking, I do want to do that. I worry about work all the time. And I bought a lot of books, a lot of business books, a lot of books about design, a lot books about digital media, I don't read any of them. I read Harry Potter when I'm done at the end of the day. I don't want to think about work, I don't want to read those books. Harry Potter I'll stay up all night and read because it has nothing to do with what I'm doing. This is sort of a balance. So, I think that's an issue—if it works, it works—if it doesn't work, we're in trouble.

Tucker Viemester: For me, Harry Potter is all about work. It's how are we going to make this stuff that's going to be like that. I think that things that learn are all about this magic and how we're going to be able to control it. We might have some idea about what we want all these machines to do, but there's probably a whole bunch of people in India or something who have a completely different set of ideas.

Craig Kanarick: Well that goes back to a fundamental definition question of “why.” You're talking about embedding magic into an object, or embedding a behavior. Adding something to an object that doesn't already have it, adding the behavior to the broom that doesn't normally fly to make it walk. And to a certain extent, adding a behavior to a city to allow it to make the right choices or predict what it should be doing that it doesn't now do. It goes back to the idea either changing capabilities or adding certain new capabilities.

Eric Begleiter: We also project a lot into things. If you look at advertising in terms of magical or at least charged psychological qualities that are projected onto consumer products, they're conquering all the time. To some degree, in terms of intelligence and other aspects, it's difficult to separate out these factors. Even in terms of Harry Potter—looking at the interesting quality of the book itself, separated from the consumer media blitz and it's narrowing the range of things that children imagine about. And so there is this sort of quality of mental projection into objects of a specific theme, a specific element and its effects.

Craig Kanarick: But we are fascinated by it. I think that children's literature does, and what Tucker's saying is that we are at some point, and maybe this is childish romance, we are inspired by adding these functions to things. We are inspired at least at a young age about thinking that things can be magical. Not about people liking to be magical, but about Pokemon—

these monsters being able to fit in my pocket. Or, a broom that can fly. Or machines that can transform into robots or turtles that can transform themselves into Ninja masters.

Eric Begleiter: Well it's interesting also to think then about artificial intelligence not in terms of generating purposefulness, but generating playfulness as well. How would intelligent devices let us—not just free us up in terms of having more time, but let us ask the question: What's the relationship between purposeful activity and playful activity in intelligent machines?

Craig Kanarick: There's a new wireless service in New York, a sort of wireless community service. And the number one application is the celebrity spotter. You're out in public and you see a celebrity, you type in who they are and where they are and then people who are nearby can run over and catch a glimpse of them. So you talk about a whimsical application, it's not checking stock prices, it's not going shopping, it is...

Peter Lunenfeld: Stalking.

Craig Kanarick: There is also a cell phone application that is the “Insult of the Day.” Every day, the Insult of the Day. Every morning, this system emails to you some incredibly offensive insult—you are the equivalent of such and such—and people forward them onto their friends. It's the largest viral application ever, so this notion about practical versus whimsical is also something that I think is really important.

Images

The task: Groups of 3 people had an half hour to take two pictures of the following five words: technology, future, learning, smart, and life.


technology



future



learning




That's a close-up of a fly-trap that doesn't learn. No flies in it. They've all learned to stay out.



smart




Two people smiling at the photographer who's just figured out how to use the camera



life



The pattern shows intention, so someone must have been alive to make the pattern.



Session 2

The purpose of the second discussion was to explore brainstorming actual learning things we might find useful and interesting. It was also to explore who that “we” is—individuals or society and what the impact of those categories might be on things that learn.

Craig Kanarick: The idea behind the next session was to try and actually brainstorm some learning things; try to come up with some things that learn that we would all like to see. I think the interesting question there goes back to what B.J. (Fogg) was talking about which is who is the “we” that we’re talking about? Is the “we” an individual? Or is the “we” society at large? And does this group (Summit attendees), as a responsible body or group of very wise souls, feel more comfortable pointing in that direction? If “we,” society, could only invent one great thing that learns for the benefit of society, what would that be?

Joy Mountford: I attended a conference recently called Living Architectures which was kind of advance sensory media and what we realized was there was not case analysis being done at the moment in architecture, for conceptual thinking. That really changed people’s goals for the conference. In addition to such things as the Long Now Foundation, there needs to be something that designers galvanize, in an objective fashion, their thought processes around. Maybe there’s a post-case house study for design and the goal was to come up with something that we could collaborate on when we weren’t in the same rooms. That was what I was really posing back to some people here is what would that equivalent thing for us be? Because this is also a very quality gathering of people and it’s quite hard to

come together and get to know each other either way. The question really is, is there something mutual and beneficial that we actually feel we could support each other in being able to either conceptualize, write about or dream about that would help us all. I think what’s hardest is to think bigger than what we’re doing. It’s the, can-we-step-back-long-enough? And these are some opportunities to do that.

Michael Ester: Do you have any suggestions?

Camille Habacker: Well, going back to the global warming. I’ve been thinking about that since I heard about it. Outside of computer models, and whatever else environmental scientists showed governments, leaders of industry is it possible to create something that would help people bright enough to give up conveniences to save the environment? Would that be trying to solve too big a problem? I’m like getting depressed by this conversation, not inspired. But that’s just me personally. I’m feeling a little bit like we’re talking about the end of the world and I don’t really know the answers.

Eric Begleiter: Well, I have a hopeful invention that I’ve heard of I think would be interesting. As we’re looking at smart objects, one of the ways that probably makes sense to think about them, especially long-term, is an integration of

three basic areas of research over the next 100 years. The first area of research would be nanotechnology looking at the evolution of the construction of useful materials from the atom up. The second area, biotechnology; in terms of DNA control mechanisms and wet-wear. And then computing (software) or consciousness in terms of these control mechanisms.

So we can look at smart objects really as a fusion of these three areas, these three domains. One area that you might have heard of is research that’s going on now. It’s not smart exactly, in the sense that we’re defining it, but it’s smart in the sense that it’s an integration of these three things. It is the using recombinant DNA to identify in plankton the mechanism for the production of hydrogen. Plankton is capable of producing, as an alternate aerobic process, hydrogen. A few researchers now are looking at ways of identifying the gene that codes for the production of hydrogen so that, organically, we can grow hydrogen, again, as a nonpolluting renewable energy source.

As you might know, if you use hydrogen it’s non-polluting because it produces water when it combines with oxygen. I don’t know if it’s really smart in the sense we’ve been discussing—it’s not connected to a cell phone. It’s not smart in that sense. But it’s smart in the sense of it’s a functioning, liv-

ing thing that solves the major problem that we're looking at over the next 100 years. So that's an optimistic possible thing that could happen. That's a smart thing I want to see.

Craig Kanarick: Is it smart or are we smart for having for having figured it out?

Eric Begleiter: Well, it's smart in the sense that it's utilizing or altering living things in a purposeful way. It's on the edge of the definition we are now using. I mean, in the sense that...is an apple smart? How smart is an apple? Don't you think an apple is very smart? I do. To be an apple, well really that's a smart thing. To be able to do that, it probably is a smart thing in a way. There are qualities, I guess, of intelligence, in terms of its functioning that typify intelligence. But then that's a question of why is the manipulation of the DNA structure at that level attributing intelligence to a thing? Is DNA intelligent in that sense? Is it a form of intelligence? Or is it only a form of our brain? When DNA forms a brain, then does it becomes intelligent?

Peter Lunenfeld: This group isn't overstuffed with computer scientists. Definitely not bio-scientists working with plankton. I think this is a room filled with designers and I think we are talking about how design as a field is very thoughtful and can try to develop a practice that would allow us to make a contribution beyond designs like the origins of commercial art. The classic issue confronting our generation is to good by making good. And I think that's one of the things that we have to confront in an era after 1999, in capitalism, is what we do and what we are. It's very difficult to decide that. When thinking outside of the box is something that absolutely every middle manager demands his or her staff in the Midwest—thinking outside of the box goes right back into the box. So, the question for all of us is how do you develop a practice? The practice can not be developed in 16 hours

(length of a meeting like this). Practice is something that develops long-term and the company learns from that. But you have to go outside the of the boundaries. How do you develop things like the Long Now Foundation?

Constance Adams: At the risk of doing something kind of wacky and sticking my neck out it seems that one of the problems that people mentioned here, the problem with which they're struggling, is the absorption of their energies and focus by short-term market driven demands of producing an endless new stream of small things; of nifty little commercial products. I have the opposite problem, which is that I'm virtually alone with a small and disembodied cadre of people who are tasked to take seriously as designers some enormous questions. For example the question of how, if and when I have the money to do it, will I confront the holistic design of an environment that is responsive—the entire design of any spacecraft or space habitat is one big human interface. It's not a refrigerator that tells you you're out of milk. What we're talking about is not being driven by any specific economic force. When you sit and think about planetary exploration in support of humans in a non-terrestrial environment, you're thinking about things that are right over the edge. And if that is something that is a useful case study, something you can wrap your mind around—it certainly takes you out of the “two years from now timeframe,” gets you out the “we'll never make money on this project” problem.

Craig Kanarick: Unless there's gold on Mars?

Constance Adams: Those technologies that we foresee in solving the space exploration problem can come back and be developed for a market where they could be profitable. That might be an interesting discussion to bite into.

Ben Davis: Of course everything in the space habitat envi-

ronment is a tool of some kind. I never really thought of absolutely everything that we encounter being a tool. That is a strange condition we may find ourselves in—can you think of anything that isn't a tool in that kind of world? Any “thing,” any object with smarts will now be considered a tool—in fact it can tell you itself that it is a tool. We are now confronted as designers with everything in the environment being a tool. Adding intelligence to inanimate objects gives them the stature of tools. Shoes are not shoes anymore, they are tools for monitoring feet, or mobility, or health. Everything we design will have possible future functions—functions that the things themselves may point out. They may monitor or actuate or remind or project and they will all create data. That's why space design is so significant—it makes this notion of function so clearly. It's like Bauhaus with software.

Tucker Viemester: I think this gets back in a way to that capitalistic point of view—everything's for sale at the same time. I'm not sure how valuable it is to think that everything is a tool or up for sale. It fits into our culture now, but...then what?

Craig Kanarick: Well there's a theory that the rapid advance of technology will get rid of supply side economics and capitalism won't be as relevant. The new economy has been talked about—at least in San Francisco in the last couple of years, that will create infinite wealth because we have this new technology which we clearly saw lasted about three months.

Peter Lunenfeld: Exactly like the summer of love.

Tucker Viemester: We're going to have an endless orgasm!

Craig Kanarick: That is a very similar sort of philosophy. So. I think, you know, not wanting to say in the future we will live

like they did in Star Trek. But it is reasonable to say that we could erase that limitation from the discussion. If what we're really talking about is improving our life—we go back to the premise that I started the discussion with last night, which was that since the beginning of time we've been trying to make our individual lives better. That was why we invented tools. We're going to continue to do that. My guess is we won't be finished. I don't think anybody is going to say, "You know what? We're all OK. Life is great for me, and for everybody else. So now, I can read or think."

Stephen Turbek: That's true unless we change the problem. If the current problem is "making people happy," perhaps the solution is distributing 'soma'-like drugs.

Michael Ester: Right. Well that screws up the needs problem, doesn't it, you skip over food and shelter and go back to drugs. That's one way of turning it around.

John Weber: What's Peter's saying is not about trying to make people happy—he's saying its try to make money. The other things are a subset of that in some way. If I understand what you're posing.

Michael Ester: No, I'm trying to get back to what Joy was saying about a practice. I think that one of things about what I started talking is the kind of practice where it becomes exceedingly difficult to think in the way that El Lazitsky or another designer from the '20's or the '30's thought about themselves—in complete opposition to a kind of market culture. I think that very few people in this room would probably feel that way. It's really hard to find anybody with a lot of talent who thinks that way. The question then becomes how can we harness an event and a desire? A free-floating desire in this room to a concrete practice within a market culture?

Craig Kanarick: And if we try to brainstorm some of the great things that learn in this room, which one of us is going to run out and patent it first and try and make money off it?

Peter Lunenfeld: Right.

Craig Kanarick: That's sort of the inherent problem in this directed conversation.

Peter Lunenfeld: One assumes it would be the person who tried to organize the conference -that's sort of the dangers of these kinds of things. It's the natural reticence within the environment.

Michael Ester: Well maybe that raises the level of practicality. if we remove the of level of—can we build this in the next 20 years or 30 years or 50 years, then at least it makes it a little harder for me or anybody else to run out and do that.

Joy Mountford: But I think it's more to do with the idea of "let's think of something really big like going to Mars." The point of Mars is not to go Mars, in my opinion. It's the process that is liberating. You get to think about all sorts of other things. So I don't actually think that we need to say "it's a glass"—doesn't matter what you call it—I think we should start with something we understand—like trousers. Something really small. But my process of getting to smart trousers will hopefully be better than anybody else's. The point is not to patent smart trousers—if you want bigger ones we can make some. The point is that it's a marker—the important thing is to go towards it—it doesn't matter that we'll all go different directions because we will stumble onto things along the way. And those really will become pervasive pieces of technology. I personally don't think we'll make much use of Mars. You need a 'galvanizer'. And what is difficult for me is to know "what are we not talking about?" I

always ask myself that question. The problem I have is that everything is so big at the moment that I can't decide if anything I say is useful or un-useful. I think constraints really help you. When its really airey-fairey its terribly difficult. That's why I get depressed as well.

Michael Ester: Some of the characteristics of what Constance was saying is the long time horizon and the big idea that takes us out of our day to day focus on practical, let's say, directly commercial products. John (Weber) and I work in an area that certainly qualifies for that in the arts. Trying to bring the world of art to different casts of audiences, which include people that manage and design the presentation process in museums. To researchers, scholars, educators, and that also need these resources. The presentation to a general public, which has all the problems we've been talking about these past days. Can I anticipate what tools somebody will need? Can I draw in some way on my experience and on the experience of others that help refine the process? Can I find out things? One of my objectives in this is not just to provide another dull set of art images, but to delight and surprise people. How can I do that best, what can I learn from that person that would enable that? So it's certainly a domain that is also in that same turf. Keeps me getting up every day.

BJ Fogg: I think all of us in this room are lucky in a lot of ways. Some of the responsibilities of being a lucky person is to transcend the lowest level of needs and to create something that benefits society in general even if it does mean we are being a little bit arrogant that we know what good for everybody. Rather than throw up your hands you have to stake a claim and go in one direction—you may be right, you may be wrong but that's what needs to happen. Probably the reason why one of the issues you brought up this morning about global warming, or the issues of sustainability in gen-

eral to me has been, for awhile, a pressing issue. The Institute for the Future says that we are facing “inflection point” in the year 2020 where our population grows and our resource consumption grows to a point where we are going to come to a crisis. I feel that humans, as a species, will survive that, we will get through it but like may be dramatically changed. I was at a conference where we tried to understand what that meant and what we could do in the meantime. There are things we can do in the meantime. We probably can’t advert that crisis that is coming. But we can soften the inflection point. So one point we could start from in this discussion is to start looking at that. It might help to abstract environmental and sustainability issues.

Neil Crofts: One of the thoughts that is around is that if you can create some kind of hyper-reality then we don’t need to do so much damage to reality. We can make use of hyper-reality for all the trivial stuff. And therefore we’ll do less damage.

John Weber: What sort of damage do we do that that would alleviate?

Ben Davis: Travel.

Neil Crofts: Yes, instead of me traveling here, I could just experience exactly this, but from home.

Eric Begleiter: One interesting aspect of tele-presence which should be mentioned in the context of artificial intelligence is how doctors at the end of their work day send out their transcripts to India, having somebody type it up overnight and send it back. Certainly, one day you probably



will see third world countries becoming the artificial intelligence that we’re talking about here today. So some person in India or China is basically running my (smart) coffeepot. In terms of tele-presence, you’d probably need to think not just of replacing all of these things but having an almost virtual servant class that are highly removed in some very distant place that are taking on the functioning of things. Because humans are very good at doing that. And once you have the ability to have tele-presence, probably we’ll develop that as an intermediate step in the introduction of full artificial intelligence.

Camille Habacker: That beats selling your kidneys!

Eric Begleiter: The new economy?

Craig Kanarick: Perhaps what we should do is try to take on this initial idea of inventing a learning thing and through the process discover if there can be learning things. If we can even come up with one at all? One of the questions I asked last night was do we need them at all? Or is it just a novelty? Or are we better off out-sourcing the learning to the people who need some wages and can use technology to do the learning for us. Should we create objects that learn? And if we can come up with one that actually provides social value as well as personal value then maybe the answer is “yes.” If we spend an hour and we can’t come up with anything then, perhaps the answer is “no.” Or perhaps we just have to spend more time thinking about it.

Eric Begleiter: Well it’d be interesting to see something like a “contextualizer.” Let’s say I have a book and I can go onto the Internet and the contextualizer can look at all the text, it can create a hyperlink for every word, or concept, then make associations based on my interests. It could also have GPS (Global Positioning System) guidance, so I can say all right,

I’m sitting here in this building tell me everything about where I am. Or if I’m walking down the street, it could basically give you a tour. That’s sort of a simple device that probably is not more than five years off. I’m thinking of something that can contextualize either what I’m thinking about or where I am at any specific moment to give me a greater historical sense or more and more information. So it’s allowing me to learn. It’s sensing where I am, sensing what I’m doing and it’s giving me access to the interconnections between all things in a really profound way.

Craig Kanarick: It’s just a general smarts, general knowledge?

Michael Ester: It sounds like the book in Neil Stevenson’s “The Diamond Age.”

Eric Begleiter: Well certainly, I know that in cars there’s a map to tell you where you’re going, but it doesn’t really give you a tour. It’s more the concept of, instead of separating everything out into parts, it’s trying to find what the connections are between things to allow me to think in a more holistic way about everything, about every object.

Craig Kanarick: If I say “the biggest bank in Japan failed,” you’d say “so what?” If I say “the biggest bank in Japan failed and it happens to be the size of all of the banks in the United States put together,” that relationship, that context, provides a different set of values to that piece of information. So I wouldn’t disagree with context or relationship or more depth getting facts, is relevant. But, that object, that’s just a larger database of information, it doesn’t necessarily change—it doesn’t change, it allows you to learn about your environment, but it doesn’t learn itself, right?

Tucker Viemester: That another conversation about “things that teach.”

Ben Davis: Are they the same?

Eric Begleiter: It would depend on how complex the patterns that it's replicating are, because it might know things, again, only in terms of limited information. It might just be sensing the environment, or it might be knowing, in terms of say a book, if it's made hypertext of the entire document for me, the types of connections that it's making, if they can be related to me certainly would make it smarter and so I would learn from it.

Stephen Turbek: I was reading a comparison of the interface design philosophies between the Airbus A320 and the Boeing 757. The Airbus' design philosophy is "coddle the pilot" by putting a thick computer interface between the pilot and the physical machine such that they can't crash the plane. The interface actually prevents them from pushing a lever outside of the safe area. There was a recent crash in Italy where they needed to do something "wrong" in order to not hit the ground and the plane would not allow them to do it. In a Boeing plane, it is up to the individual to know the limits of the plane.

Tucker Viemester: So let me get this straight, if you're driving in the Airbus, you can crash because you can't control the airplane and if you're driving Boeing airplane, you could crash because you can drive the airplane?

Craig Kanarick: Right. The question is, can airbuses learn? And maybe the answer is when, everything learns, the ultimate thing that should learn is...capitalists or organizations should learn.

Ben Kleinman: Can you have a learning object without a learning organization behind it?

Craig Kanarick: You know the TV that learns what I like to watch, or the email system that learns what types of things I like to read and which type of things I don't like to read and just trashes them automatically. Learning media has been talked about a lot. Like the newspaper that knows what I like to read and it automatically rearranges the stories in the best way for me. It knows that I never read the Bridge column, so it just removes the bridge column from the newspaper over the course of a long.

Something that I've been talking about at Razorfish is interfaces that learn. Right now, interfaces on Websites are very static. I will never use certain functions in Amazon. I will never purchase certain objects from them, I will never call on some of those functions ever, ever, ever. Yet, if I visited a trillion times, that tab will still be the same size as all the rest of the tabs, it'll still be on the top of the page. It will be better for me, personally, if over time it learned not to waste my time with that. A great host at an organization, a great host at a party will only introduce you to the people that they know you want to meet. They won't waste your time introducing you to people that you don't want to meet. That happens only over time. That only happens if that host, that interface worked, not if it's just smart. There's no way to, at least I can't conceive of a way to have that information automatically programmed in an instant into the interface.

Michael Ester: But I don't want something that's that figured out. What if they did they that on a social level and pretty soon there are no black people around anymore because you'd never talk to one at the party. What are the things you're going to miss because somebody programmed them out of your interface?

Eric Begleiter: There is an important thing that will need to evolve and that is the ability to identify those things which

are intelligent enough to be responsive to us and so deceptive. I'll give you an example. Let's say I'm interested in a presidential candidate. I go to their Web site and I want to see what is important to them, and unbeknownst to me, something about my interest in the environment is known to them and so all I get from the site is information on their pro-environment stance. The system's pretty smart, but if there's no little sticker on that thing that says "we're making this just for you," then I'm going to have a very hard time separating what they're about from what their presentation to me is about. Without this ability to distinguish those things that are intelligent enough to be deceptive in some sense there is a problem. We don't have a mechanism to be able to distinguish those things clearly yet.

Craig Kanarick: But isn't that just a short-term issue?

Eric Begleiter: It is.

Craig Kanarick: Just thinking about it right now we expect certain things to be a certain way and after visiting 500 websites and reading tons of articles about the fact that they do know who you are and are modifying their behavior you'll change your relationship to them. We want technology to act like an appliance. We're going to drop that at some point because we're going to eventually culturally just learn it doesn't work that way.

Michael Ester: I think that we were describing market manipulations. Devices would tell you go down a certain street or we give you a search engine and you be shown only the stuff that's been paid for. I think that there will be a quality of society that will just demand to be told what is happening. I agree, it's a short term goal.

Craig Kanarick: I always laugh when I see, "this is a paid

advertisement,” disclaimers in magazines or for an half an hour of television. You know, if you can’t figure that out by now, you haven’t been watching a lot of TV.

Eric Begleiter: Well. No. Actually, there are mechanisms in newspapers that are clearer. I think that if the lines becomes so blurred that you can’t tell what is an ad that’s a problem. Eventually everything may be “ad-like.”

Peter Lunenfeld: I’m thinking again about maintenance and the idea of programming your interface to put away something that you don’t use very much. I think that’s not a bad goal. It moves away from a notion of interactive configurability which was such a big deal and still remains that way. When I look at myself, I never configure anything, really. I could spend half my life configuring things perfectly, so I just take it out of a box, use it, and then live with it in a half-assed, screwed-up manner, but at least it works. I would love to have an ability to have these things shift over time. I that may be reactive, rather than intelligent. But, then you get a simulacrum for it. Which would be fine. That’s all I asking a machine to do.

Rebecca Odes: Isn’t just gaining knowledge, learning in some way? That kind of reflection of your behavior and being able to accrue that kind of information and take that and use it as a sieve for all further information that comes into the object. Isn’t that kind of learning?

Craig Kanarick: That’s my point. That would be enough for me. I would have a lot more demands, but it certainly would make me happy if I came back to Amazon and it wouldn’t recommend a Miles Davis book to me because I bought a Miles Davis record. But more importantly, notice that I don’t ever buy cigars and just get rid of anything about cigars—just streamline my experience. Its also dealing a little bit with

information overload and living in a society where there are a lot of choices, and there are lots of configurations. I can’t find my way through the Microsoft Word menus, but if it only has the four things that I want it to do—which it did five years ago—I would be fine with that. I don’t want to get trapped into technology interfaces. I mean, that’s my personal bias, that’s why I was trying to make a metaphor of like a party or a good host.

Apparently, Christie’s has the best doorman on the face of the planet. He pays a lot of attention to learn each individual customer that comes into the door. Who they are, what department they visit, where they go, learn about their families, ask them questions, end up having a relationship with them when they come in. That relationship develops over time. So when I think about personal benefit to me, at least in my day to day life, I would love for the distractions, the things that are essentially distractions or obstacles, to go away.

Tucker Viemester: But I think what’s interesting about this is if you have this such personalized software that everything is honed to what you want, I think it’s going to be much more important for people like that to get out and connect with other people that have their own personalized thing. And I think this is going to foster a lot more social connection so that you’re going to be able to see it like one of your friends tunes into Amazon their thing looks totally different than yours and that’s how you’re going have a reality check.

Craig Kanarick: You mentioned Autonomy software this morning. Autonomy is essentially a search engine. And it actually tells you who else is looking for the same types of things you are. It learns not only the connections between the documents, but it also learns about the audience and says, “Hey, you are obsessed with this subject or this person

and did you know these other people are too? Maybe the two of you should get together and collaborate.”

Tucker Viemester: But see I’m not worried about trying to find all these other people who are like me. I’m trying to make sure I have some kind of object point of view about what the world is like. But like you, I don’t want to have all that other crap on the computer either. I like the combination of that customized thing and the randomness of real life.

Michael Ester: Is there a sense of urgency in this? Is there an inverse relationship between the explosion of information and an increasing need for the editorial hand? In other words, finding ways of grouping information, pre-selecting it. Even in our area (the arts), as we put up more and more collections, how are you going to know what what’s in those collections? There needs to be someone that helps aggregate interpret that. What’s in them, what you’re looking for. So part of this is the ability to filter, aggregate, and enhance something that is looking at it from a more of a macro view than you are. And it’s filtering that out for you.

Stephen Turbek: What if the user doesn’t think they are qualified to be rating the news stories? You authorize other people, like editors, to make decisions for you because you recognize their skill in that field.

Craig Kanarick: But you do rate the value of them. Merely by the about the amount of attention you pay to them and how far into them you read. If you read three sentences, that story isn’t valuable to you anymore. It’s a matter of instruction or, if the paper just watches you and sees what you read, then you don’t have to take any of the effort, or actually rate the stories consciously—“this gets a four, this gets a ten—it just automatically does that.

Stephen Turbek: Just because I don't read a story about the Middle East 10 times in a row doesn't mean it isn't important to me.

Craig Kanarick: No, but at the time, it's not important because you must have gotten your information elsewhere. And the theory is that that thing really learns effectively. Over time it will learn to understand that.

Michael Ester: There's an irony because you stopped learning—in other words, it starts shedding more and more things that you might learn as a casual observer. It's like reading the newspaper. If it starts only focusing on the things that you do by habit, it's learning and you're not.

Michael Ester: Right.

Michael Ester: It drops out the Bridge column. Now you have to skip over some other column. So now you don't read the sports section and that's gone. And pretty soon you're reading the one headline and that's it.

Craig Kanarick: That's imposing a value that says that I should be required to be very diverse. You are making a judgment that I'm missing out on something because I'm only reading the sports section. God damn it, I should get the bridge column—that's bad because it was removed. It's bad that you're only focused on the one thing. I don't necessarily believe that I'm required to have every media interface provide me with some serendipitous new content that I'm not interested in. Just in order to make me a more diverse person. I may not need my newspapers to provide me with that extra little edge, I may get that when I get outside and screw the newspaper. It's not the newspaper's job to do that, the newspaper's job is to just give me the facts. My guess is that if that thing really does work, is really smart, it'll figure that

out as well. And it won't delete the Bridge column a 100 percent of the time, it will delete it 99 percent of the time and the rest of the time, it'll pop up.

Margaret MacLean: I was wondering if we are waiting for the for technology to make these choices for us? What makes me so angry at Microsoft is when they anticipate how I should be designing a document. I didn't ask them to do that. I can do that all by myself. And maybe I'll want to ask it to help me with that at some point. I'd like to have that ability. But I don't it to be crammed down my throat. And maybe that's one of the problems here. We're assuming perhaps that the technology should be out in front of us, guessing what it is that we want. So it's pushed on us. What would be more useful would be something that waits for us to say what it we want then perform that activity.

Craig Kanarick: So the thing I heard this morning was anticipating needs. That's one of the values of smart things, anticipating our needs.

Margaret MacLean: That may change. What if all of a sudden I decide I want to play Bridge.

Ben Davis: Let me bring up two other areas because I think what we're talking about here is information overload and context. What about personal safety? What about surprise and delight?

John Weber: I think that what would actually kind of interesting if people individually came up with two different things. One which might be a very small appliance thingie another might be a very big, almost impossible thing, but two things that learn.

Ben Davis: We're talking about personal things and prob-

lems we'd like to see learning things solve like information overload and context. Context brings up issues of personal safety—how to survive in a particular environment. Its really two scales—a close-up and personal thing that may translate to a big thing—maybe personal safety can be extrapolated to violence in general. Perhaps if the small personal thing is universally useful it makes an impact on society in general.

Eric Begleiter: Or the other way—What were the other ways in terms of individual personal environments like sound cancellation to get out the noise and have pure air and wear your own smart things.

Margaret MacLean: Right, everybody has a smart recycling thing in their house that spits out something usable like a chip that becomes sort of a secondary form of money to use for something. It might make people want to recycle more.

Craig Kanarick: What I'd like to do for the next half hour is think about all these things that learn—specifically things that learn, not things that are smart, or things that will reduce carbon gasses or things that will make me stronger. But, real things that adapt over time or things that learn either about me, or about the world, that will make a difference. If you could snap your fingers and someone would walk in with that invention right now, what would that be? For me, it was if Amazon.com got smarter the more I used it. That to me was something that is within the realm of possibility and I don't even see that. I think we can think of much more fantastical into things that are much more difficult.

Adam Eeuwens: Based on my own experience something because of how hard it has been to move a lot in the last four years, I would like to see something that would help me get re-settled by disconnecting utilities and reconnecting utilities

in the new city I am moving to. Helping me understand the new neighborhood < where is the grocery store, the gas station? etc. I wish I could just tell my phone I'm moving to Chicago and it would set me up there. This device would choose my energy based on what type of energy consumption I had before. I wouldn't have to deal with all this stuff .

Craig Kanarick: Right now I know I use the wrong long-distance calling plan. And I, theoretically could call AT&T and say, based upon my calling habits they've learned how to give me the best deal and still make money so it do that for me. Your appliance is watching your habits and as a result chooses appropriately, even as you adapt to a new environment.

Adam Euwens: Exactly.

Craig Kanarick: Based upon how you use them.

Adam Euwens: And then if I moved today it would remember all that stuff and it would go with me so I don't have to set up that again. I don't have to think about it anymore.

Eric Begleiter: A small aside in terms of that. I did see a cradle for a mobile phone that when you get home, you put it down, it basically becomes your local service, and so it rings the other phones in your house. So you never—basically, you don't have a local provider, you just have a cell phone and it becomes your regular phone.

Craig Kanarick: But that thing doesn't learn everything about you. It doesn't know, doesn't watch your habits.

Eric Begleiter: It takes away the need, to some degree, to change your number, because you just have the same number wherever you go.

John Weber: What's the advantage to just having a cell phone and not having another phone?

Eric Begleiter: Because it basically allows you to have all your regular phones in your house run through it, and it acts like a transmitter.

Craig Kanarick: Multiple hand sets on one mobile number.

BJ Fogg: This is an idea that everybody's familiar with. It would have been in the learning category, as a sort of six degrees of separation device. Maybe it's enabled to a cell phone. When I walk into this meeting, for example, I—over lunch, Margaret and I realized we'd been in Peru at the same time in the early eighties, and it would have been great—maybe we had a mutual friend there, it'd be great to walk into a context and be able to speed up social relationships. To figure out what are the degrees of separation and over time know who I've met and who other people have met, and when you come into proximity, it should know the outcome of the degrees of separation, the little links there.

Eric Begleiter: And then you could have the opposite, which would erase all knowledge of wherever you've been.

Margaret MacLean: I guess you could use it to say "I don't want any exposure to this kind of social circle."

Eric Begleiter: It would actually go back and just erase all your tracks.

Craig Kanarick: So over time it learns all the people that you've met and all the people that they know, or whatever circle it is or whatever you can discover about that

BJ Fogg: I'd like to be matched up with which kinds of

degrees of separation are interesting to me or not. Maybe it's more a I want to meet more of the academic types, or the industry types, for example.

Craig Kanarick: Just as a note—that was attempted on the Internet.

BJ Fogg: Yeah, that's why I said it's not a new idea, but in terms of a mobile device.

Tucker Viemeister: It could be a phone service.

Craig Kanarick: It is becoming interesting to me is that what we want is these things to learn without our instruction. That seems to be one of the best things about learning. When I talked about construction versus sort of experience. It sounds like we don't have time to instruct our devices, we don't have time to teach, you know, to configure the device, to teach the device, to type in information. That was the thing about SixDegrees.com. I never wanted to type in all of the people that I knew which is going to take a long time in order to get the perceived benefit. But if it just noticed who I sent email to and did it on that, it might work. So, it seems like on the criteria side, one of them is learning through observation is much better.

BJ Fogg: That's undoubtedly true. It may be possible with the location sensing technology that's coming to cell phones within the next year.

Michael Ester: I'm having a little trouble with the "it" part that's learning. In these examples, these are sensors that are picking information. Where does the "it" part reside? Who owns it? Who manages it? It's not the devices that are so smart here, it's the system above the devices. I'm very interested in the last examples where that system resides.

Neil Crofts: If your cell phone could check the address and everyone else's cell phone in the room and refer those to six degrees, then it could come back to you and tell you about them.

Craig Kanarick: Is there some magic software program that we're talking about, and does it learn?

Michael Ester: The difference is that you can have smart devices that solve the problem locally and we can have other kinds of smart things that look smart, behave smart because the information might come through your phone, but it's not where the learning and the intelligence is taking place. And there's probably social and other consequences to the fact that this information doesn't reside in a little device, but resides somewhere out here that somebody's managing.

Neil Crofts: And economic consequences.

Craig Kanarick: Which I think is discussion of the legal issues which we'll get to as we develop these concepts.

Joseph Busch: I think it's important to point out that these are not necessarily new ideas, but they're difficult to implement because they require integration and inter-operation. We sometimes say that we're swimming in a sea of information or a fire hose of information, but the interesting thing is that with the technology today, we can't harness and localize all that information to do things that we know we should be able to do. We could specify how to do it, but there are all sorts of barriers to accomplishing it. They are caused by the way in which this technology evolved as well as economic barriers.

There are good reasons—and great difficulties in—why your phone doesn't talk to your Palm Pilot. And there are lots of

people who don't want them to talk to each other. And there are lots of dreamers who say if we could only talk, it would be great. These are mechanical as well as social-political issues. If there was a passion and a compelling reason to do this it might happen. If we go back to the sustainability context—globally and locally—that might become a compelling reason that could mobilize people to overcome these barriers. This a how, instead of a what. But I think it's important.

Craig Kanarick: That's like the space analogy. We put a man on the moon. It wasn't why. It was just because it's a good idea. Then all of those barriers of economic or technology were just a challenge.

Stephen Turbek: Part of the motivation for putting the man on the moon was just to beat the Russians. Different people get motivated in different ways. In regards to Neil's earlier point, environmental consciousness could be a good application of information technology. No one wants to be constantly minded, but it's nice to have guidance when you need it. For example, if there was a system that could give you context and consequences of your decisions, you might be better able to be socially conscious. If you could see how your decisions effected people from another country, such as the people working in the factory where your purchase was made. Summing all these up could create a network for you to decide the relative consequences of all actions.

Tucker Viemester: Like an eco-tamaguchi or something.

Neil Crofts: I'm thinking of some sort of aspirational assistant that you could put your aspirations into and then it would guide you to eat food that has been grown organically and without damage or whatever or grown sustainably. Or to only buy products that have been created sustainably. And your aspirational assistant would collect those aspira-

tions and help you to—or learn those aspirations—and then help you to fulfill them.

Craig Kanarick: Sort of both sides of one coin. One side is the consequence device that tells you the consequences of you're going to do and the other one's the aspirational loop. I mean, they're both the same thing, but they're just a different way of conveying that. It's the "what should I do?" device and then provides the right sort of way to get there.

Neil Crofts: Oh yes, but one gives you the right aspirations.

Craig Kanarick: Right.

Neil Crofts: And then helps you fulfill them.

Rebecca Odes: Well I have a related item which is similar to that in some ways. It's also a sort of multi-disciplinary filing system that knows what stuff you have going on and knows what to do with it. I work in music. I write creatively and I have tons of list management, writing stuff to do, plus I'm a visual. My notebooks are a horrible mess and they're very interesting to look at for the outsider, but I can never classify my information properly, so if there's some way to create something that knows what your categories are and to be able to take information as you input it, both in a musical instrument device and put that with input of words in a certain case—that are not in list format, knows those are lyrics and put those with, with that information. So you're able to more easily have your personal taxonomist.

Craig Kanarick: It's your digital left brain.

Rebecca Odes: Basically. That's what I mean. Clearly, it didn't work out in the biological domain.

Craig Kanarick: Or vice-versa for the people who are left-brained, they need the digital right-brain to help them. Maybe that's the aspirational line, help them become more creative in what they do, remind them there are other ways of doing things or there are less organized ways of doing them.

Margaret MacLean: It just puts the two wrong things together.

Craig Kanarick: Right.

Joseph Busch: When you talk about taxonomies, these should allow you to put things in more than one place at the same time. There's sometimes a tendency in these systems to put something in a buffer, but that's not the way taxonomies work best. The smarter way, which a lot of software doesn't do, is to let you put the same things in lots in of places.

Peter Lunenfeld: It seems to me we're falling into something that is incredibly common, which is this notion that what we're always driving for is the ultimate goal of efficiency. That's the money shot of business form, right? Efficiency. That's it. I'm starting to think maybe we need smart machines and we also need smart-ass machines. Things that can remind us that efficiency is not always the goal. A friend was using a Palm Pilot the other day and we started talk to about the fact that it's a new addiction, like email is an addiction. Now he sees a blank spot in his agenda and he fills it. It becomes his obsession to have full pages.

BJ Fogg: It's like a tidy desk.

Peter Lunenfeld: Right, it's like a tidy desk. Everyone's been talking about efficiency this, efficiency that. How hard is it to pass the Bridge page in the newspaper? It's one of the eas-

iest things I do when interface with the newspaper. Bridge? Turn the page!

Tucker Viemeister: You don't even think about it.

Peter Lunenfeld: Devoting teams of programmers to making sure that Bridge never pops up, who cares? I think we need to address some of the ways that we can, you know, have a little bit more flexibility?

Craig Kanarick: While I agree with what you're saying, my guess is that you're addressing that to this audience, but there's a whole load of people out there who need to be more efficient. I mean, who are not as self-motivated or as hardworking as the people in this room. So when talking about extremely productive people, it's easy to say, "hey, why don't we all relax?"

Ben Davis: I bought this a while ago. It's a DigiWalker for \$ 29.00. It's like a pedometer. You program your stride into it. The theory is that if you walk 10,000 of your strides a day you've had a complete aerobic workout. You have to walk till the digital display says 10,000. The bad part, of course, is when you get home and it says only 300. It is a good example of a body-monitor that can reinforce your behavior in terms of exercise. It can encourage you to be more efficient about your exercise. I'd love it if it also told me how many good ideas I had in a day and where I was when I got most inspired. Sort of hook it to mood ring somehow, so it told me something else.

Craig Kanarick: Which sounds again like efficiency. You want to have more good ideas, so therefore, you'd want it to learn where your good ideas were so that you could go back there and do it again.

Ben Davis: To me, it's an example of the body monitor that could do something else—could be combined with a geographic location system and could learn where the "sweet spots" are for getting new ideas.

Joy Mountford: Something's troubling me here about the notion of efficiency. We've begun discussion at Stanford recently on how people write in different countries. Americans write totally efficiently and English people on average tend to write a bit more abstractly, verbosely, at least that's the tradition. What's an interesting observation is that many of the early observations of scientists in the late 1800's were story-based, for example Charles Darwin. They are wonderful stories and you learned a lot from these stories because you became embroiled in the presence of the story. Now, we have these wonderful little abstracts that are summarized to the point of dryness and brittleness. I can't read those things anymore, and I get nothing from them. So there is a point on a continuum where I think things break down horribly. I'm sure there's probably a service that could summarize USA TODAY. I mean, how summarized does summarize really have to be?

We choose our sources of information to accommodate our differences in style. Some people buy USA TODAY, some people read the NEW YORK TIMES. . I think it's very important that it's not just always efficiency. Because dragging yourself through the process the same way as dragging yourself through the movie "Schindler's List," we don't want those people to die, but it's really useful for us to go through that much of an experience because it is in itself something valuable, that it doesn't just mean efficiency as a goal.

I'm a bit troubled that it's one or the other. I think there are ranges and grades that are very valuable to us for learning . Reading long treatises in the 1800's actually helped us

understand because we took longer to read and write. The time is a valuable construct when you come to really understand information as well. So when we come back to really mundane things, like “would it be useful to have the Bridge page go away?” we can consider things like fifteen years ago, when there was some early work done at the University Maryland by Ben Schneiderman where menus configure themselves based upon your usage of them. In fact, it drove people absolutely bananas because things were always moving around. So the thing that you used the most was always at the top. But the really important part of that study was really that your musculature actually learns to go here and there—we have all these little built-in maneuvers. You actually sort of learn everything physically and by changing these things, even with the Bridge page going away, they’ll put some other shit up there. It’s all going to become a giant and big so the word “burp” is always there. I mean, one has to stop and think about what you’re designing for.

So to get back to what one learning thing I’d like to see created < I’d like to have a pair of earrings that whispered the name of the person who I’m talking to.

Rebecca Odes: Excellent idea.

Joy Mountford: I feel as if that’s the sort of thing that I’m never gonna make money out of , I don’t know how to solve the problem. But I’d love us to think about it.

Neil Crofts: That could be connected to B.J.’s idea.

Constance Adams: It could be connected; but let’s resist that temptation, the temptation to take a good, straightforward idea and add a thousand features onto it when you only want one. My goal in designing is to optimize everything for its specific task, so that it is typologically determinate and

not schizophrenic. I think if those earrings started communicating with other earrings or sharing too much information I would rip them off my head. On the other hand, when Joy first stated the simple idea, almost everybody went “yeah.” You know you’re onto a good idea when you can get visceral reactions from people. Now that’s the sort of energy that I think people need to build towards a greater future.

BJ Fogg: And what I love software companies that make software that could do that and stay alive without having to add features.

John Weber: Absolutely.

Peter Lunenfeld: And it would work in 10 years. The same software would actually work.

Ben Davis: Would you settle for earrings that just whisper?

Craig Kanarick: Remove some features.

Ben Kleinman: We’re focusing a lot on devices that help us to learn, not the devices themselves that are doing the learning. It’s really just basic pattern recognition. I mean, at the most basic level and they’re helping us to sort of do that. I don’t think we’ve really come up with a learning object or a learning system.

Craig Kanarick: The interface learns, it doesn’t help me learn, it just helps me shop. I guess my thing is not necessarily about efficiency, or about being able to do more. It’s more about pleasure. For me, it’s about removing the distractions or the things that clutter my life. It’s not necessarily about just being more efficient, it’s really about having more pleasure. I mean, if there’s less negative things around—I’m not saying I want a news broadcast that is all

good news—but to a certain extent, things do get in my way. It’s not that hard for me to flip the page of the Bridge column, but it is hard for me to wade through all of the stuff, all of the obstacles that get in the way when I do want to be efficient. So it’s not that I need overall efficiency, it’s just that there are times when I do want efficiency and at those times I like to be as efficient as possible.



Ben Kleinman: It should learn to be efficient when you want to be efficient.

Tucker Viemiester: How are they going to learn that? Now there’s the trick—that would be fun. If it can know that you’re cranky that day and you don’t want to see the Bridge page at all. Or, on other days, when you’re like just cruising it would leave the Bridge page alone. I think that would be very interesting to have a machine that could learn—

Rebecca Odes: It’s like a mood ring.

Craig Kanarick: There was a great conference at that MIT media lab about the emotive computing. When everything has emotions, but more importantly, when everything understands your emotions. A user-interface that knows you’re in a bad mood and gives an appropriate error message versus one that knows you’re in a good mood. You know, great people do that. People adjust their discussion style, if they understand you, they pick up on your cues. So if I know my Dad’s in a bad mood, I’m just going to say something quick

and get out of there. If I know he's in a good mood, I might talk to him for a long time. Computers don't do that, right? Same error message pops up every time. It doesn't say, "you know I don't really want to mention this but I know you're in a bad mood but it seems to me you're making a lot of mistakes." I would also like a computer, to go back to another thing that I use every day, that learns what I do in different emotional states and responds appropriately, so that's where the learning comes in. Not that the smarts are responding to the states, the learning is. I don't want to teach at that when I hit the keys really hard, I'm in a bad mood. Because I don't think I could delineate those things. So what I would like it to sort of figure out what my emotional state is based upon my behavior.

Neil Crofts: Do you want it also to respond to that emotional state with appropriate stimulus to bring you back?

Craig Kanarick: Potentially.

Neil Crofts: Play soft music?

Craig Kanarick: To take it to the next level, right. What it could learn then is what I would want it to do in that situation.

Ben Davis: You seem very excited, Dave.(reference to the Hal computer in 2001: A Space Odyssey)

Craig Kanarick: That's right. Yes.

John Weber: I'd like a car that knows where to park. And it might even tell me, "you'd better call a cab tonight."

Craig Kanarick: And how does that car learn? How does it increase its awareness?

John Weber: Not my problem. I just want it to tell me where to park. How it learns I don't care. I'll tell it where I want it to go—you know, "I'm going to go here now." It'd be nice just to be able to say that out loud. And then it can say, "OK, here are your options"—and if there's going to be some problem I might have to adjust it. I might suggest this, that, or the other. But, something that just deals with that situation.

Craig Kanarick: Anne?

Anne Young: I'd like a mirror that would tell me when I need to moisturize every day. Take a look at my body, keep track of what I'm going to take every day like vitamin E, and have that just grow with me and help me as I get older.

Constance Adams: I'd like a power grid that uses the most ecological sources available at any given point in time along the city's power use curve.

Craig Kanarick: And it would predict—it would probably learn the behavior of the power consumption and know that at half-time at the Super Bowl, the water supply goes down.

Constance Adams: Yes, and it can also store data and deliver reports on what new facilities need to be built in order to improve itself.

Craig Kanarick: Sort of the smart city...the concept of learning from the history of how this very complex organism behaves.

Constance Adams: And its metabolic structures and how that effects the ecological balance.

Ben Kleinman: One of the things I came up is related to these two issues, which was a learning highway system. I

want the road to know how I drive and how everyone else prefers to drive. So I like to drive exceedingly fast and other people don't, so it helps guide them out of my way. Maybe it talks to the car and then the cars learn.

Rebecca Odes: Road rager is now behind you.

Ben Kleinman: The fast cars go at 7AM or 9AM or whatever and it starts to adjust the whole spread of things. It potentially ties into mass transit etc., but the system really learns rather than my individual car or that individual block of concrete that I'm on as part of the whole.

Tucker Viemiester: Drives right to the police station.

Constance Adams: See, that's the flip side of that noise, right there.

Eric Begleiter: I realize because of the ambiguity of what we're looking at now—in terms of what's sensing and what's learning < that over the next 5 or 10 years every product that senses anything is going to be called "intelligent." We won't be able to stand the amount of "intelligent" things around us!

Craig Kanarick: That was the crux of this discussion about why things learn, versus why things think. I think, we're getting to the point where we are saying that things are thinking or at least acting like they're thinking. But, when they start to learn, it does become a different animal because we have a different relationship with them because they change over time and I think that's the thing about Ben Schneiderman's work that was so interesting about menu interfaces. At any individual moment were they arranged appropriately? Yes. But it sucked that every time it was different.

We're not used to that much change in objects. We are used

to change in people, but we're not even used to change that much in people. People are incredibly shocked at a limited amount of change. Wow, you changed your hair color! You got new glasses! You look great! That's a shock that people have modified themselves, even a small thing about appearance. I guess we expect moods to change quite a bit, but we're definitely programmed. We haven't learned yet to accept that objects will change at a huge rate. I'm wondering if we will be able to?

Joy Mountford: I think this is an interesting problem also with it being wireless. You know, I'm sick of people telling me that interfaces have to be transparent. That was a few decades ago and that we tried to do that. Now I'm actually very concerned that they are transparent. This is sort of born out by the fact it's wireless. So now you have these things that are so-called potentially smart and how do I know if this is a dumb company car or a smart one, and how do I know if this is a dumb company car communicating with your smart one? And yours is picks something up—you can go on, and yes that might sound absurd. You're talking about a very similar thing like at the Web where you put your credit card numbers going elsewhere. But now I imagine it, when there are physical things around you, so you don't have to carry a Powerbook you can use TV monitors in hotel room and do useful things with them—I mean, there are good positive qualities about this as well.

But the interesting thing is "How will I know what I let talk to what? In what situation? Many of these things will be under my control, in addition to the situation, the environment's control." So in some way they now have move from being transparent, to being explicit. So of most of these devices we'll say: "This is always embedded, this will be great." Well—what is embedded? How do I know it is embedded and how will I query it? So I find myself thinking of a totally

new device, which sounds a bit bizarre, but it's a device that I could see check out the room possibilities and find out what options are available for me given my suite of peripherals that I'm carrying today? Maybe I can learn a lot. I can configure and use the energy valuably, the communication protocols valuably, blah, blah, blah. That is a very different interface problem than we've ever seen. A I'm very upset that the Media Lab, when they did their "things that think" stuff, never really addressed that sort of problem.

Craig Kanarick: I don't want to get too buried in this sort of "dangers of monitoring and privacy and all of those things," but it does sound like anything that does learn, we want to know that it's learning. Just like anything that communicates, we want to know that's going to communicating. We want to be aware—WE want to be aware of what's going on at all times.

Eric Begleiter: In a way if we're lucky, we will want to. The other side is that you basically generate a culture that says, fine, give me a T-shirt with a logo and I don't care.

Craig Kanarick: The classic Media Lab example of the thing that learned was the coffee machine that learned your general preference for coffee. Now of course what it didn't know, again, was a simple pattern matching thing. It's that every time you come and you get two sugars, so it gives you two sugars every time. It didn't know that on the mornings that I'm really hung over, I don't want any sugar. It didn't create rules, any learning system, it just created a sort of instant pattern. That may have been its down fall or maybe that's just sort of another level of learning. The first step is to recognize one thing to do, but there's this combination of case-based learning and rule-based decision making. There probably needs to be some balance of the two for anything to really learn about.

Stephen Turbek: Will we accept things that do more than just suggest things to us? For example, we'd all probably like a coffee machine to offer us our usual, but I don't imagine we'd accept things that make decisions for us.

Eric Begleiter: I think that we could find ourselves in a situation where as a culture; we look at these issues of privacy and become very upset at what we are losing. And yet, at the same time, I can easily see that it's also not like a "1984" kind of situation, but it's kind of like the case where people will just care less and less about those issues. If this is so, it is even much more of a problem...

Craig Kanarick: Perhaps. That's why I don't know if we want to end up having a discussion about this sort of implication about privacy. We'll spend the next hour complaining that we don't want it to happen. And yet, that's all we're going to end up doing. Just like, we don't want global warming to ruin the planet and we can all agree and we can for the next hour talk about why we don't want to. But, it will be six 'o clock and we'll agree that we don't want it which is where we started. So.

Eric Begleiter: What I would be concerned about is the sort of "brave new world" model in which we would all give these things up and just be absolutely thinking this is the greatest thing—to give up those things.

Craig Kanarick: Well that's sort of the value of the capitalist model. For every person who is afraid of these things, there's an opportunistic person who is going to invent a way to circumvent it or prevent those things from happening. So that's my ultimate faith in the capitalist model. But, it doesn't get us back to things that we want.

Tucker Viemeister: It seems to me that all these things that we've been talking about that we want are pretty much hyper

normal. They're like you want to know what somebody's name is, you want to look cool, oh, you don't want to look like an idiot, so you have this thing that helps you do that. But, what about all that other stuff that's not normal? My sort of vision for the smart world with all this stuff embedded is it's like sitting around in a cave with a fire and eating some meat with your pals but its all connected and smart. In other words, the extension of all this smart stuff is back to like what a normal life was. So what about all this other technology that we can do that is outside of that?

BJ Fogg: There's only a fixed set of human needs that humans really care about. One of them is entertainment, others are for security, others and so on. I think that's in some ways, I think all of our ideas have addressed basic human needs that we care about. Then sitting around with pals, it's sort of an entertainment experience, right?

Tucker Viemeister: I meant like your gizmo that you wanted so that you could walk into a room and find people that you wanted to talk to. It's like living in a small town where you know everyone. Where's the high tech, you know?

BJ Fogg: I guess that's sort of what my point is. It's tapping into a common need. That has been consistent for thousands of years. That is the interesting thing about psychology and technology. The technology changes incredibly, but human psychology isn't changing.

Neil Crofts: It's an evolution -compensation thing, isn't it? We've not evolved as fast as our capability to move around. We actually still want to live in a small town with all the benefits of living in a small town, so that we know everyone, but we actually don't have a small town, so we want something to compensate for our lack of ability to evolve at that speed.

Craig Kanarick: It is just this sort of augmentation that we invent devices to make us incrementally better at a particular thing. Whether that be remembering people's names or, you know, putting a spear through an animal. We just create these individual little pieces. Are we hitting all of these five categories? We've talked about we want an earring, we want a six- degrees-of-separation device, I want a digital left-brain or digital right-brain. And a motivator, or reminder—this sort of digital conscience. Then, a couple things about larger state –an organization—a society that learns what the patterns for that society are. I don't want to say efficient because I don't think the power grid is about being efficient, it's more about achieving some other goal which was, I think, a little bit different—optimal resource use.

Michael Ester: I think of efficiency as doing what I'm focused on better. What I heard a lot of people talking about was kind of cutting away the distractions—things I don't want to know about. Whether it's taking things to the laundry or going through the Bridge page so that I can do other things that might be a very creative—kind of cutting away the distractions so I can do creative things. Or focus my mind in other ways.

BJ Fogg: I think one of my issues is that I think there's so many barriers to getting in the flow of experience, and a lot of that has to do with technology. I've come to realize, every little email I get, there's no way I can get into the flow and answer email. Answering email is not a flow for me at all because it's just a series of little pieces and having to skip over the Bridge page, that breaks up the flow, so maybe part of it is a learning technology that helps keep us in flow somehow. Knows what experiences are like, and put us into the optimum flow.

Michael Ester: So maybe it's making us more effective, rather than efficient?

BJ Fogg: Optimal experience. Bring us to optimal experience. Whether that's getting things done quickly, or getting us into more of that creative mode.

Craig Kanarick: Or getting them done better. You know, having them be more rewarding in some way.

Michael Ester: When I looked at the materials that I was sent—the idea of all these learning devices being embedded in things—and there certainly are such things—but every single example is a sensor or a responder that is local. But the information coming from those devices is being aggregated and interpreted external to those local devices as well as the value coming from the integration of many of these interpreters and processors. There's implications there. There's important implication there.

Craig Kanarick: Like what?

Michael Ester: It means that I've got to figure out how information gets aggregated. I've got to figure out how that information gets shared. I've got to figure out does it aggregate with me personally? Something that is personal to me, from a lot of different sources, that's in my area? Or does it aggregate to some central area? I can imagine walking around where other things have to query my personal resources to find out how to respond to me.

Craig Kanarick: But are those implementation details?

Michael Ester: No, no...

Craig Kanarick: I guess what I'm wondering is, we don't need to figure out how these devices will work. We need to figure out what the impact of those devices will be. So what I'm trying to understand is—there's a lot of questions about how we actually build those things. But, do the answers have any impact?

Michael Ester: Yes. It means, first of all, they're not autonomous. And that may be a good thing, from the point of maintenance. But, it means that I'm depending on being, in fact, wired generally, not wired locally. There's a sense of less control. There's an implication of central control. Not of distributed control. I don't think it's just operational.


Joseph Busch: This is all about ubiquitous computing. There are lots of strategies. There's a period of turmoil when you need to distribute stuff. When you have an environment in which all this information is being accumulated in lots of different systems, those problems are going to be solved. The technical issues are important, but to me the question is to what end? Whether we like it or not this accumulation of personal information is happening. Cell phones are amazing data collecting units and lots of people are collecting data on us. The Web is an amazing data collection facility and there's a lot of data that's being collected. If there's all this stuff being collected, what would we like to see done with it? What do we do with this wonderful, information technology?

Craig Kanarick: To me, being digital implies a universal language to communicate with each other. It doesn't mean being able to use that information to draw inferences, and improve upon the performance of the object. Do I think that there are issues about privacy and data collection control? Absolutely. They all boil down to we all want privacy so we should be able to have control over where information goes. I've heard a few things that say we want things that learn.

But it doesn't sound to me like there's really a lot of them. A lot of them are about wanting things that think. Or things that will tell us where we are or what we're doing. But not necessarily things that change over time or things that modify themselves all the time. I don't know whether there are another set of issues that get raised when that occurs.

Eric Begleiter: It also depends on the time span we are talking about. If we look out 30 to 60 years from now, when the neural nets and synthetic intelligences we are talking about will have amazing capabilities, there will be a strong temptation to use them and that will cause untold changes in our lives. Imagine starting a company—what could do with an inexpensive supply of the power, the power of 100 people in a network box on your desk? It can do a lot, voice, intelligence, its almost super human let's say; it can run the office, customer service, computer design, and manufacturing. So now we see why we need to be starting our own company, so few other companies need us. But then what is going to happen if not enough people can be entrepreneurs or if every consumer just wants just one hot brand, maybe two. Then conformity will concentrate so much power. I think that kind of circumstance would lead to other forms of redistribution of wealth. We might move to a situation in which each person is a shared resource company vying for a one billionth share of a global market. But that might not happen.

Constance Adams: Well if 29 out of 30 fewer people were employed, there would be that many fewer people to whom



to sell whatever Goddamn product we thought was brilliant in making. The question I find really pressing at this point, I've listened and participated for a few hours now, is can we identify anything that learns that is a necessity and not a luxury? All I keep hearing—and I agree about a lot of, I would love to have a pair of earrings that told me everything, you know, but I think it's an interesting question.

Neil Crofts: I think there's a rider to that which is...one thing that learns which is a necessity, perhaps, is children. But the outcome is completely unpredictable.

Rebecca Odes: Which is part of its charm.

Neil Crofts: Exactly. And so our only experience of something that learns is children, dogs, perhaps? But actually, it's pretty limited. And the outcome is unpredictable. So to try to put ourselves into conceptual space where we're thinking about something that learns predictably and desirably and essentially is tough.

Rebecca Odes: Or relationships, I mean, part of what people are talking about, you know, when you get tired of your television, do you get rid of it? But then, the new television won't know what you like to watch. I mean that's what people do when they change relationships all the time. You sort of throw out the old information with hopes that the new information will be, or the new situation will be better. But, you have a learning curve and you're creating intimacy with something and then you let go of that thing. You need to create intimacy with something else in a way that you would with someone else. It's all human relationships I think that learn.

Craig Kanarick: I would venture to say that the answer to the thing that does need to learn is us. I mean, the one thing

that needs to learn and change. Learning is about modifying, right? It's about changing or increasing knowledge in order to promptly react to a changing environment. The environment doesn't change. The thing doesn't learn, necessarily, unless it's about getting efficient. The environment would stay static and it could learn to become more efficient, which would be good cause it would create efficiency, but if it's sort of doing what it needs to do, it doesn't let me change. Since the world changes, we need to learn. But we can't apply the same strategies that we used 400 years ago to anything we have now, because we just don't have the same environment. We don't have the same situation. So I think we need to learn what we're going to do when the ice caps melt because living in New York isn't going to work anymore because we'll be underwater. So, society is probably one thing that definitely needs to learn.

Stephen Turbek: There are many groups of people our age that are addressing this issue in ways that you might not expect. I think you guys are familiar with the work of Dolce International [the organization that ran the conference center where the conference was held]? I was reading their sign that says that, they their approach is that they are global leaders in learning through meetings and leisure.

Craig Kanarick: So are they learning or are they're just helping other people learn?

Stephen Turbek: They're leaders in learning.

Tucker Viemiester: But it seems to me that there is a survival issue about this learning, you know. That, we're going to have to learn how to deal with the melting of ice caps. Either we learn or we don't. Or we die, basically.

Craig Kanarick: So evolve or die is true?

Tucker Viemiester: I believe it is. I was upset when Gore in that debate said the environment was his number one issue and then a little while later he said education was his number one issue. He didn't relate the two. It seems to me that would have been so easy. You know, it's like, we're not going to be able to deal with this environmental issue unless people know about it. And it means every person on the street, we're not going to be able to ignore the people who don't want to know.

Eric Begleiter: The intelligent tractor might be an example of a non-trivial application. A tractor which uses radar satellite information to figure out what area to plant and when .

Craig Kanarick: I was thinking about a food product that would learn and would adapt to the changing environment. So as it retains its nutritional value, but it learns how to grow based upon the way in which the planet changes. And if we can't teach it what is going to happen, it just has to learn how to adapt and if learning is adapting—that is, it's taking information in and changing as a result. Then maybe we need to invent a food source. This is the magic bean that right now it grows a certain type of product, but in 50 years, when the climate is different, in that same location, it has modified itself. It's learned how to adapt to its environment and continues to fulfill the same things. And maybe even learns where the environment is going so that it can make that change ahead of time.

BJ Fogg: It learns how to grow the ozone layer somehow.

Craig Kanarick: And then when there's too much of that gas, it learns to go back the other way.

John Weber: If we look at things that are necessities they are so huge, it's hard to get a grip on them. So maybe one of

the things we need to learn is how to break down survival into small enough components that can be designed for by companies that can sustain their economic existence by designing them. And that the notion of a "survivability factor," eco-logically, socially, could be legislated..

Peter Lunenfeld: God forbid that word.

John Weber: But we're already legislating all sorts of stuff—tax credits and what have you. So we're doing this all the time. So let's not freak out about it too much. But there could be a way of understanding the pressure points on the eco system and sort of the pressure points on the different social systems then you could design specific things. And you would have a smart car that would be, maybe only 10 percent more ecologically efficient, but in fact that would be good enough if there was also a smart house that was also 30 percent more efficient, then maybe from your clothes you're getting an additional percentage of lowering the necessity for heating systems that are creating pollution—accumulating enough of these incremental savings will in fact result in a situation that is a much smarter economy.

Constance Adams: It's also true that for example you can break the smart car and smart house down into smaller components too. You use an LED based light fixture for example and when you turn down the rheostat, you have a direct reduction in the power. Whereas, you rheostat an incandescent light, it's the same amount of power going into support the ballast. It's just that it's reducing the amount of light that you get. So you're still burning a 100 watts, even though you just got, you know, 20 watts worth of foot candles.

Neil Crofts: But even if you do that, or even if you make a car that won't drive above certain pollution levels, just turns itself off, the education actually still has to be with the people

because people still have to want that.

Craig Kanarick: I had a car that wouldn't start if everyone in the car wasn't wearing their seatbelts, we learned to put the seatbelt in the little thing so that when you sat down, that sensor went off. And as a result, nobody wore their seatbelt because they were already in the little thing behind them. So that's the other problem—we learn just as fast as these things do.

John Weber: It wasn't a very smart car.

Craig Kanarick: Right, it should have then learned how to detect that ridiculous situation.

Constance Adams: It was controlling, but not smart.

Craig Kanarick: Well it didn't learn. We did as a result. It was smart. We learned how to be smarter..

Michael Ester: No one's raised the issues of traditional crime and danger. Either the smart bullet, or smart gun—or

there's the protective suit that inflates around you. I don't know. Clearly, if you look at the issues that people are concerned about in elections, crime, danger, all that is very high on the list and we haven't mentioned that at all.

Craig Kanarick: I guess what's interesting to me is that smart seems very expensive. It seems very expensive to make things that are actually smart because you have to figure out all these possibilities. Like the car. Like the seatbelt in the car. It had to of time be told that whenever someone sits down, right, if the seatbelt's already in before the pressure is on the seat, then they obviously figured out some way to circumvent it, so turn it off. And then they would have come up with some other method. And I'm wondering is it cheaper to make something dumb that learns? Then try to make it smart to begin with.

Neil Crofts: Try to predict the permutations. It needs to learn the permutations in practice.

Ben Davis: So that's the thought I'm going to leave with you as we wrap this up this evening.



Session 3

This final session on Sunday morning was to see if anything from the previous day's discussions had created any new ideas, dreams, or insights. Participants were asked to sum up what "When Everything Learns" meant to them.

Craig Kanarick: Nobody dreamt about learning? Nobody dreamt they were in school about to miss a big exam?

BJ Fogg: I dreamed I was planning some kind of park space—sitting in a room with space designers.

Craig Kanarick: We had a little discussion at breakfast this morning about the value of dreams. One reason for having this on the agenda is as another "get to know you" idea. But it appears that "things that think," us being the best example, do dream and there's a debate about what value or what impact that dreaming has. Would we need things that learn or think, to dream as well. I guess it depends on your interpretation of why people dream or why people sleep. Tucker had a good explanation for why people sleep.

Tucker Viemiester: That's because when mammals or animals are young, the parents need to have time to get food for them and so while the kid's asleep they stay in one place, sort of stay out of trouble, so you as the parent can go away. That was the usefulness of it, but it's just evolved. When you grow up you still have to sleep.

Eric Begleiter: Well actually there are separate species of nocturnal animals and daytime animals—so you have two complete shifts of animals. It's more efficient. So that's another

reason why you can go off for awhile and not be eaten.

Craig Kanarick: Not being eaten is very important.

Eric Begleiter: It is. And then in terms of dreaming it's maybe so you can get up and run away fast. So your brain doesn't take that much time to get started and react.

Tucker Viemiester: It didn't shut down completely.

Eric Begleiter: It didn't shut down completely. It's still doing a little work.

Tucker Viemiester: It's like a screen saver.

Craig Kanarick: With a nice geometric pattern.

Constance Adams: That's a hell of a screen saver.

Craig Kanarick: Anybody else have anything else to say about dreams? Before we head back into learning?

Eric Begleiter: Well, I actually had something left to say. I have been involved with experiments dealing with human computer interfaces for dreaming enhancement. What you do is use the computer to determine when you're in REM

sleep, and at first it will wake you up so you can just remember what your dreams were. You do this for a while, a few weeks. Then you train yourself, when you see a specific stimulus, to ask yourself if you're dreaming and you have to prove to yourself that you are, in fact, not dreaming. The stimulus that I was using was a flashing light. So whenever you would see a flashing light during the day, you'd ask yourself, okay, is this a dream? You have to prove to yourself "where was I five minutes ago?" Does it make any sense, the fact that I'm here? If I read something and I look away, does it change? Does gravity feel different? You go through this and you really prove to yourself. OK, this is a dream or it's not a dream.

Then what happens is, while you're sleeping, because you've built this up as a response and you're connected to the computer, it's looking for REM state, it generates a stimulus, but it's not enough to wake you up this time so in your dream you see a flashing light and you've created enough of a repeated behavior that when you see a flashing light you then ask yourself in the dream, all right, is this a dream? And you go through the little process of saying, OK, I read something, I look away. How did I get here? But what's interesting is eventually in the dream you come to the conclusion that this must be a dream. And then you go off in your sort of aware/dreaming state.

Constance Adams: That's a pretty high state of consciousness to come to the conclusion that you are, in fact, dreaming. Another way of putting it is to think of dreams as a kind of shadow on the system. While I was working in East Berlin, the ISDN lines hadn't gone in yet so were still working with Weimar Republic telephone lines laid underground that were in a bad state of repair. Whenever I was online just trying to compose mail and it rained, bits of electronic garbage would slowly start washing down the screen as the rain would pick up. Because it was raindrops actually affecting the phone line connection—it was this shadowy artifact of nature hitting the system and I think that that's the dream.

Craig Kanarick: Eric, you say that the computer detects the REM sleep and then would flash a light in the room that you are in?

Eric Begleiter: Yes, at the second stage. The computer waits about five minutes so you store enough memory into your mind's short-term memory, then it'll wake you up. Because you need to develop the ability to remember your dreams in order to have this work. And the second stage it is able to bring stimulus into the dream. But you need to have developed this automatic response to trigger the lucid stage.

John Weber: Is that a learned response?

Eric Begleiter: Yes, a conditioned response. Every time you see a flashing light you ask yourself the same question so eventually it becomes sort of ingrained.

Craig Kanarick: Learn how to behave in dreams? To a certain extent.

Eric Begleiter: Because then once you've developed this alternative state of consciousness, in terms of being able to

be aware in the dream, then it's possible to begin manipulating the dream based on your anticipation of what's going to happen. It normally only happens at the end of the night because you need to have a good night's sleep first. You need to be really well rested in order to awake enough to dream in this way. It doesn't create a kind of confusion between sleeping or dreaming and waking state. It actually makes you much more aware of the distinction between them because it really takes a very sharp mind to question the basis of reality at that level. To ask yourself whether or not this could be a dream is hard. Because normally your assumption is, and this happened to me actually last night, very often the assumption is, you come up with the conclusion that magic is real or that something else has altered, but not that it's a dream. Not that the underlying basis of reality is you. That's like a really bizarre result. So usually you'll come up with anything other than that.

Craig Kanarick: I guess, the opposite of what some of the robots were in *I Robot* because for them it was pretty much the other way around. Sort of very Descartes influenced.

Eric Begleiter: If you look at a large scale three dimensional binocular display, a large projected image, before you go to sleep—something that's not really changing or moving very much—let's say an image of the Earth, or yourself approaching a space station—for 10 or 20 min then that will much more get into your dream in terms of a virtual reality experience than some really strong moving, running kind of thing. Also something interesting is that you'll think you've woken up—it's like a false-awakening in this lucid state, but in fact you're still dreaming. One interesting way to actually wake up is to have yourself go to sleep in the dream. Because it lowers the stimulus enough. So there's perhaps some deeper meaning to that.

Ben Davis: I have had those experiences though where you're in such a deep sleep that you question—that you have that false wake up. You hear the alarm clock in the dream and you get up and get dressed and you're still lying in bed.

Tucker Viemiester: Could you flash a light just to make sure I'm not dreaming? (LAUGHTER)

Craig Kanarick: Yes, read it again to see if it changes.

Ben Davis: Oh boy.

Craig Kanarick: At the beginning of the weekend at the dinner on Friday night I went through a little bit of a logic chain. It said that since the beginning of time, we have tried to help ourselves as individuals, we have tried to make our own lives easier. We've done that through the invention of various devices. Those devices have either augmented us or replaced us. Embedding a sense of intelligence or at least a false sense of intelligence or some representation of intelligence made those objects even better and even stronger. The idea was that thinking things were better than non-thinking things. If learning is an increase in that ability to think or a shift in that ability to think, hopefully for the positive, that that made learning things better than non-learning things. So we have, dumb things, thinking things, and learning things in the scale of good to better, if you will.

Then I tried to challenge that notion and see if that was actually true and see if in fact we did want thinking things or, in fact, we did want learning things at all or if they were purely a novelty, purely something of interest, purely our desire to want things that were actually magical and not helpful and just that we were fascinated by what those objects could have as capabilities. Whether our social structure was pre-

pared to allow thinking things or learning things into it. Whether our legal structure was allowed or prepared to allow thinking and learning things into it.

I'm wondering now if we're any closer to answering those questions. We asked each of you to prepared some sort of refection or statement on that subject or on any other related subject. What is the epilogue? What is the take-away reflection of the last day and a half? Are we any closer to either defining "when everything learns" or any of the specific sub-questions that were asked during the day?

Peter Lunenfeld: B.J Fogg talked a lot about basic human needs and I think design has been looking at those things that are not really part of those basic human needs. I'm still not sure, if after this weekend, that "things that think" or the "things that learn" are addressing basic human needs as opposed to engaging with a kind of aesthetics of the light. This notion that novelties please less than they impress—which is a line from Don Juan. I think Byron's really onto something there. I'm not sure, after this weekend, that I'm any closer to thinking that any of this is any more than just "make a really, really, really neat watch." I like really, really, really neat watches but it doesn't get to Mars and it doesn't get to global warming. I'm wondering if there's a way to bring that to a certain kind of resolution. Just trying to figure out might be enough—and it's fine if that's where learning things stop.

Craig Kanarick: Right—and in the interest of not wanting to be overly efficient.

Peter Lunenfeld: Right.

Craig Kanarick: I don't think we need to come to "The Conclusion." So I think that's great, thanks. Should we just go around the room?

Margaret MacLean: Well, I can't speak for anyone else, but I've reorganized a little bit since I arrived. Thanks in part to something Joy mentioned yesterday. I've reversed my notions on thinking and learning. Thinking is the sort of quintessential human action and learning is something that is more like reacting and reassembling information. You can be a machine or an animal that can be taught to do something. It can learn to do something simple, linear, but the thinking part is the more creative and less linear. And less accidental. So when we say, "what happens when everything learns?" that's a much easier thing to consider than "what happens when everything thinks?" I think the big problem is when every thing starts thinking and talking.

Craig Kanarick: Conferences we don't want to attend.

Margaret MacLean: When all the objects in the room are talking at you. So that's my only revelation. Small and personal one.

Craig Kanarick: Great. Thanks.

Constance Adams: Well, I would certainly concur with the fact that learning is much easier than thinking. And much simpler. It struck me in thinking about the various utility of what we were talking about that there are different levels at which these things are useful and interesting—and of value. It seems that we all seem to be people here who learn and who think and who agree that what we really need is for humans to learn more and think more. And the machines doing that is not necessarily a great virtue on a small scale, though on the big scale it could be.

To me, the revelation about a learning thing, a cybernetic learning thing, that seemed most fundamental in terms of building one is that the virtue of it learning is that it will overcome "garbage in garbage-out." We can make machines

that don't continue to endlessly transmit our own fallacies in building them. That by itself is a good thing. That's a useful thing, rather than transmitting human failures to even greater extents. But I think once we get beyond that, an important rule—an important aphorism to substitute for "garbage in garbage out" would be "first, do no harm." Or at least, "first do less harm." That would be a thing to focus on. That if our machines, all of them, at the simplest and most fundamental and basic level, don't need to be luxury items, and can learn to compensate for failures in their original education and structure and then can help us learn to do less harm—then they're of value. Otherwise, I don't know that they are.

Tucker Viemiester: Hardly anybody starts out to do something intentionally to harm someone. So I think the big problem is to deal sort of with that harm after it's been done.

Constance Adams: You know, my analogy might have gone too far. I was borrowing an aphorism from medicine: "First do no harm." Doctors don't set out to do harm either, but the reason medicine is taught with aphorisms is that aphorisms contain important reminders about priorities for people who are dealing with balancing enormous amounts of critical information. It's easy for a doctor to become fascinated with the disease and to forget that they're hurting the patient in dealing with the disease. So what I mean, to address is that forgetfulness...and that hopefully we can build machines that will remind us of those priorities.

Michael Ester: As learning and thinking creatures we get more information, we get more bad information, I don't know that even the ability of things to learn necessarily overcomes the basic fundamental issue of garbage in, garbage-out. Those are interesting issues. What's the tolerance for increments of learning and how does learning improve that tolerance for error and the volume of information?

Joseph Busch: We know from information theory that we can tolerate a lot of noise in the system and still be able to get the message. There's always going to be garbage-in. There's always going to be garbage out. The question is how do you mitigate it so that you still get the message through. And whether those tolerances can be tweaked. What I've been thinking about is the theme of the connections between things. It doesn't matter whether it's garbage or not, the interesting thing is that there are lots of connections between things. We haven't talked about the notion of the global brain. I think our discussion of ecology is important for us to think about. Think about the growth and development of the brain, the ability to generate connections and to regenerate different connections over time. A physical, physiological model of learning. We haven't talked too much about this. I think there are analogs in language and in the development of information systems that try to replicate these processes. And it's related to natural systems as well.

Tucker Viemeister: I really like that concept. It just gels in my head that the reason why I like digital technology so much is because I always assumed that it was going to a way of saving the world and I realize now that it's sort of like the evolution of the world, like we have a single cell organism and now we're sort of building a neural network for it that's going to make it more intelligent.

Eric Begleiter: I was also struck by the relationship between learning and thinking. I think that in some ways we've talked about thinking a little differently than being contemplative, I think we've also used the term "thinking" more in terms of sensing, the distinction between sensing the environment and predicting and observing the result of those predictions and then learning. And also how we learn, for good or bad, from the things we make that learn. I did think that that was a kind of interesting aspect of the idea of the filter, of both

increasing and decreasing the flow of information as one of the things that we would want to do.

In terms of being more efficient or effective the other thing that I thought was particularly interesting were the ideas about how dangerous a total model efficiency might be in relationship to the playful or to the dreaming mind. One thing that occurs to me in the work that I was doing with the human computer interface for dreaming is that it is something that allows you to become more efficient, and yet the very quality of being able to do it is completely frustrated if try to do it in a work-like efficient way. It's almost as though it's like a bubble that you're trying to catch. If you reach out too fast, the bubble will break. There's some aspect I found interesting about creating an efficient model, but then not using that completely, allowing this sort of playful, creative mind to also manifest itself.

Ben Kleinman: I thought about fear a lot. If a city runs efficiently because all the parts know how it's supposed to work—sort of an organic thinking—well this device will do this and this device will do that, that element of fear sort of came in. It started to come along a continuum for me. It's when everything learns—whatever that boundary is, we hit it. At the one end of the continuum is to influence and manipulate. The people behind it are teaching it so that when I react this way, it will do this to make me react another way. Very manipulative. Then it enters the OK range and then it starts to become too uncontrolled, when it starts to go from learning into thinking. When things start thinking, I think, is when we become very fearful of it. I think that's when it approaches what is human. Are they going to take us over? When everything thinks really depends on where it falls along this spectrum. How are we going to react to it? How are we going to treat it socially?

Constance Adams: Don't you expect, therefore, if people really feel that way—which I believe we do—that it will never happen? You think it's inevitable that we will develop thinking machines of which we are afraid?

Ben Kleinman: I don't know if it's inevitable, but I think there's a strong possibility we might get there. If you take a short term view of things and not a long term view of things, a lot of incremental steps go unnoticed, the big jumps, you know. So if we go 200 years into the future, with what we know now and all of a sudden we'll say, "Oh my God, I can't believe what happened." But, for the people who are actually living through it, it's just part of the daily life.

Tucker Viemeister: It's the old frog in the boiling water....

Ben Davis: Nuclear energy is a good example of that.

Rebecca Odes: Yeah, plenty of things that people were terrified of have happened.

Neil Crofts: I have a concern coming out of this which I think picks up all the themes that we've have around here, which is that there's a choice really between thinking machines, or, having things that think and things that learn, which allow us to abdicate from thinking—allow us to simply focus on entertainment or distraction < or having things which think, which actually augment our ability to think and to consider. It seems to me that the way that our society has developed, would pitch it very much at the level of helping with distraction and not with helping us to deal with the profound, which is kind of ironic. I don't know what we can do, or what anyone can do to influence that. We also talked quite a bit about the effect of capitalism. It seems to me that capitalism is the engine that would push us towards creating thinking machines that allow us to just entertain ourselves. And that

may be, if we're to avoid that, we have to move to whatever goes beyond capitalism.

Stephen Turbek: I tend to agree with you, particularly about the abdication of responsibility for thinking. I've been suspecting for a while that consciousness is a kind of burden for thinking people. It seems we've been making things that make us not think, for example alcohol. Every society has always created some way to stop intelligent thinking. And uses it quite often. It sort of makes you wonder if thinking is what we want.

Peter Lunenfeld: I want to stand up for alcohol.
(LAUGHTER) And drugs.

Tucker Viemiester: Drugs make you think more.

Peter Lunenfeld: And so does booze. Used the right way. They're wonderful ancient technologies of thought amplification and thought manipulation. (LAUGHTER) And inhibition reduction, all those sorts of things.

Neil Crofts: I wasn't referring to those. I was referring to TV and stuff like that.

Peter Lunenfeld: I'm just trying to defend alcohol. I should become a lobbyist. (LAUGHTER)

Joseph Busch: Why make a distinction between TV and drugs, I mean, we can just say they're the same thing.

Ben Kleinman: I actually agree with you. With both you're not thinking in a sort of rational. I've never used any drugs at all, so I can't speak to that. It's a sort of hallucinatory effect, sort of like an entertainment. It's internal.

Stephen Turbek: But television doesn't affect the actual

mechanism of the mind as directly as alcohol does. You're not changing the process of thinking when you're watching television.

Peter Lunenfeld: Drugs can become a practice. A mental practice.

Ben Kleinman: Next conference's in Amsterdam.

Craig Kanarick: "When Everything Drinks."

Adam Eeuwens: I guess what I want to take away from this meeting is more than just a group of people sitting here being very familiar with technology to some degree, and excited about technology and coming up with ideas, but that we are concerned with a sense of responsibility and effectiveness. I had this feeling that technology is actually way more developed than we are developed to accept it. Thinking again about this whole issue of the individual, the family, the society, community, and institutions and seeing how people do need community and need each other. How can the "thinking thing" help? I do not have the answer, but I think that it would be very interesting to see how those relationships develop. What is it one expects from people to help you, to trigger new thoughts in you, to entertain you? So I was just thinking about my expectations of people, like the hotel staff here. I expect my room to be clean, but I don't necessarily want the staff to ask me how I feel today and ask about my family and give me a recommendation on where to go.

BJ Fogg: I think I'll start out here by listing a few favorites from the event. I think the favorite big idea is the concept that maybe we don't want everything that learns. I think that was a bit of a surprise to me. It's a bit of a reality check. It's an interesting idea. We may look back 20 years from now and say, "no, of course we want everything to learn but...."

Favorite phrase I think I'll take with me describes a concept that's very interesting is "aspirational assistant." The most comforting idea that came out of this discussion was that a lot of you share my concerns about where we're headed with the environment as well as where we're headed with capitalism if you want to call it that, but just corporate institutions and how that is shaping what we do with the minutes of our day and where we channel our creative energy. So that was actually nice to hear from you all.

I do want to stand up a little bit in defense of learning. This conference has helped me put down more concretely my definitions of learning. I know we purposely didn't go there, but it was so hard not to do that on my own as I was listening. It seems like there's at least three levels of learning, I don't think these are new, education people have already sketched these things out. The lowest level is probably knowledge acquisition, simply learning facts. I think often in our discussion we were conflating learning with that. I think that's just one level of learning. Another one was skill acquisition. Learning how to do something. And the highest level, which I think is true learning, is behavior change. Actually doing something differently. We talk about learning machines, and learning computers, I think this is the kind of learning we're talking about. Machines that adapt autonomously, without our necessarily asking them to learn. And I guess, to me that's circular, they adapt on their own, trying to fulfill our needs in some ways. And often they're wrong. Often it surprises us, which brings up some of those ethical and usability issues.

The last set of ideas I think I want to share in summarizing is the "what and the how" of those things from our discussion yesterday. First of all, what can learning things do? What is the "what" of what they do? One is, they change their functionality so they offer a different value or they add something

new. So the “what,” the functionality can change. The other thing that can change that we didn’t talk about was the style of the interaction, or their approach, whether computers are dominant or submissive in their interaction. I think that is a type of learning. We do that with people all the time. The content may be the same, but our style, our approach may differ. I think that might be a kind of learning that we would welcome in our technologies.

Then, moving onto the “how,” or how to design things that learn? I think that one interesting idea that came up yesterday was the pace, or the speed of something that changed, changes too dramatically, I think we sort of agreed that was an issue, if things changed all of a sudden that we would have a problem with it. So I think we have a lot of learning to do about what’s the optimal pace for each individual?

The next thing, and Joy brought this one up was the transparency of the change, knowing that something has learned or changed and being able to find that out and making that apparent for a couple of reasons. Knowing what we’re missing, I think somebody talked about that, so it’s changed, so what am I missing, am I missing the Bridge column? And also, for giving humans a sense that they’re in control of their devices, which, at least right now in our current mental states is important. It may not be a generation from now. And I think the final thing about how to design learning things that I’m taking away with me is the idea of how to undo something. So it’s trying to learn and trying to adapt and fit my needs. How do you undo that easily and quickly? Because we talk about how they’re going to make mistakes and all that kind of thing. How do you undo that? Whether it’s as simple as “show me the rest of the damn menu,” you know, “you changed the menu and I didn’t want that.” Or, “you keep ordering plain yogurt for me and delivering it to my fridge and my visiting cousin from New York is no longer here

to stop that kind of thing.” So maybe there’s a rhetoric or maybe over time we’ll develop a rhetoric for these kinds of things.

Craig Kanarick: What if we went through life with everybody thinking there was an ‘undo’ button. To a certain extent that would be liberating to feel we don’t have to be as guarded as we normally are. On the other hand I wonder how much responsibility that removes from each of us.

Tucker Viemeister: I think people have always wanted to be able to go back and fix things. But there is something very different about a machine that can instantly ‘undo’ a mistake.

Joseph Busch: In the digital context it is an essential property of the digital condition. You can really undo something. In life if you hurt yourself you can’t undo it. If a doctor makes a mistake they can’t undo that. But in the digital environment you do have the option. When we talk about these virtual realities becoming more powerful, more pervasive, more invasive then ‘undo’ become a very big issue. Its become an icon for us now that is synonymous with “I can screw-up with no penalty.”

Craig Kanarick: Is the analogy in our context here ‘undo’ is “unlearn?” We can’t un-learn something without a tremendous amount of de-programming. The process for unlearning may be as hard as learning for us. Would it take the ‘smart city’ as long to un-learn something as it does for us? Would smart digital things be able to instantly “unlearn.”

Peter Lunenfeld: In sports it takes a long time to unlearn some bad habit. It takes the body a long time to do that.

Tucker Viemeister: I thought the most interesting idea about all these machines that learn was that they would learn with-

out us teaching them. That was going to be a labor saving aspect, they would just learn by observing us. The other idea was the spectrum of everything being a tool to the endless orgasm. I think you have to step back and ask “Is this just my point of view that this all a good thing, or is it going to be good for other people?”

Craig Kanarick: The intent of this event was not to take the scientific view of these technologies but to ask, as we do at Razorfish, not just “Everything that can be digital will be” but “Should it be?” The critical debate is therefore “if it will be—its our responsibility to discuss whether it should be or not” if we are indeed helping that condition to happen. And if it is inevitable, what can we do to make it as comfortable as possible. If the analogy is “Everything that can learn will learn” it doesn’t mean it should. That was a primary motivation for this discussion. If it is like global warming we have to ask what we can do to lower the inflection point, make it more valuable and less dangerous. I’m glad those social and moral implications keep popping up.

Eric Begleiter: The session we did before on what we want could be modified to discuss what we need, what we don’t need, and what we don’t need, and we don’t want.

Rebecca Odes: When I was thinking about this summary I was first struck with thinking about fear. It seemed that there was a real anxiety about maintaining control over the information we receive. The thing that struck me about that was to assume we have control over that. I don’t feel that we do have any control over the information we receive. I feel misunderstood by most of the media that is supposed to be targeting me. I also thought about the implications of what would happen if this thing you wanted to learn failed and why that is more disappointing if you expect success. Even with a person you get upset if that connection fails. I was

also thinking about priorities, a feedback loop of priorities. If something was going to learn that it would be learning our priorities and it would not give us more information than that. Which brings up a larger issue < if everything is going to learn then we need to make sure the things that we're teaching are good. Even if people are afraid of things learning there is a responsibility of making sure the content is good and that you must give the thing feedback that you are happy with it.

Craig Kanarick: When we talked about this event Tucker kept saying that we should talk about who the teachers are and what their responsibilities are. The difference also about learning through observation vs. learning through instruction.

Rebecca Odes: I think it is the issue of the machine reflecting back to you the smart things you have put into it—unless you want the big brother situation—which I don't.

John Weber: I've been struck by several things. One is the flux between looking at things as a micro- gadget level, and looking at things that are at a very big systemic level. It's totally logical that we would do that. Our society is really good at making gadgets, but this can also be very big and very interconnected. That productive capacity seemed very directed towards increasing the capacity of gadgets and figuring out ways to sell and market them and then the big step toward systems just sort of happens in some way. The systems, the governments, the cities, the networks sort of happen. Another interesting idea is the augmented learning and that the notion that gadgets can learn. Then there's the notion that several people have mentioned about how they're all interconnected to the people. One of the questions that I take away is "what do we want to know that we don't know now?"; and if we think it's not knowable—is there a way that something could be designed that would help us learn what that is and then help us to know it.

I'm struck by the fact that a lot of the comments have to do with the fact that the world's too full, that there's too much stuff. There are too many messages. Digital technology makes it easier to replicate a lot of stuff and the world is filled up very quickly. We've gotten very good at building gadgets. Their capacities have increased to the point that we're now afraid of—where those capacities might take us. We're good at teaching our machines and our technologies to increase their capacities and their skills. Although a tremendous amount of money and effort has been invested in our school systems, our learning systems it, I don't know that it's actually that much better than it was 20 years or 40 years ago. It's interesting that it's not. That, the computers are wasted on it and I'm not sure that's a smart learning place.

And then one post-script: People who work in schools and who work with teaching a lot are very aware today that one of the best ways to learn is by doing things. That learning is a very sort of active—hands-on, minds-on sort of thing. It's hard to learn more by doing less.

Michael Ester: I think one thing that's come out of this is that things that learn is actually an important topic. And if nothing else, taking that perspective is an interesting way of looking at things. One kind of measure of that is some of the areas that it could affect in terms of some of the benefits that we talked about, contextualizing information, making us more effective to do things that we as humans are good at. Maybe there's some things that learn and they let us think. Also, the aggregating of information as it gets increasingly dispersed.

The other thing that came out of it was the idea that it is systemic and that came up in everything from looking at the environment and cities to the idea that all these things are interconnected. That was a thought that still provokes me.

You ask why is it important? I don't think that the sensing part of it is learning. Or at least I would take that as a position in that if it is dis-articulated from these larger systems of interconnectivity, it's not interesting—there's no action that results. Maybe there's a new adage that if a tree falls in the woods and its sensor is disconnected, do we know that the tree fell? But, also, there's an autonomy. At least if there's a sensor I can go step on it. Once all this is interconnected, we do get into those issues that I thought Peter raised, which are things like unintended consequences and the idea of maintenance. So the complexity of these systems, I think, multiplies. Finally, the whole issue of pace of change, and how fast we can take advantage of new ways of learning. We have software where our users already are resisting change from a release a year ago. Even with obvious benefits, more power, more ability, they've already become wedded to an interface they only started using roughly a year ago. So, how fast do we take advantage of new learning environments?

By the way, I do have one more point which is a metaphor. Neil and I had a kind of side conversation and I just kind of leave it as a last question. We assume the way things learn is similar, somehow analogous to the way humans learn. What if there are better ways of learning? Or better ways that a computer can help us learn. We were talking about a game that would have all the characteristics of entertainment, but might solve a real world business problem. Maybe we ought to be looking at new kinds of learning that we haven't explored yet.

Ben Davis: I want to try to update the old story of the pious man and the flood. Anyone know this story? The pious man, maybe today is a very brainy theologian, understands computer systems, understands "the force" and is up to his waste in the water of a flood. A row boat comes over and says "come on, get in the boat." And he says, "No, you go

on, help someone else. I'm really connected to a higher force here, I'm going to get out of this." And later a canoe comes by with someone else and the water is up to his shoulders, and he says, "no, go on, go on, I've got a direct line here to a higher force, I don't need help." And pretty soon, the water's up to his chin and the helicopter comes. And he thinks it over and goes, "No, go on, help someone else. I'll get out of this because I'm really connected here." And pretty soon the water's over his head and then he drowns. And then he's in heaven and he gets his five minutes with "the force" and he says, "I don't understand what I'm doing here. I really had it together here and what's the deal?" And "the force" looks at him and says, "Well I don't know what the deal is either, I sent the helicopter for you!"

And I think a little update on this to me is that these things we're talking about can be really great gifts and we should really take responsibility. One of the things I'm always struck with at these kinds of meetings is the startling revelation, "Oh, maybe the people in this room are really the ones that make a difference." That it's not another group of smarter people somewhere else—it actually might be us.

And finally, some of things that John was saying about moving from the local to the global, moving to the personal to the societal, the notions of contextualizing. I think this notion of being overloaded, that we're drowning in the flood of information and we really do need the helicopter to get us out of it to some degree. Personal security, personal safety, violence in the world—big problems that a thing that learns can help us with somehow. This notion of surprise and delight. I think that's where the playfulness of the technology comes in. Make something really nice and efficient—then let's play with it. I think this is a really nice notion for things that learn.

Joy Mountford: I guess the first comment is that I think Craig and the Razorfish team have been very courageous in trying to initiate this conversation and I would like to thank them for that. Particularly, struggling with the notion of not having a fixed goal, which I think is a very easy track to get people stuck in. So that's one of, thank you for letting us be here.

My comments fall into three main areas. One is that we've been struggling a lot with the matter of scale. There are various points on the scale we can gravitate towards. For example, size specific scale. It's very easy for me to think of things just being the things in my hands or the world that I walk within. Some of the bigger thinkers here think more about cities and world issues. I think that there's a future continuum there and it's easy to sort of flit from one end or the other.

I also think there's another scale—the size of a group. Is it just me or is it, you know, a couple of my friends or is it a whole societal group?

Another scale is utilitarianism. And again, because we're coming from this wonderful, capitalistic view, efficiency is good, more menus are good—the new Microsoft version model. I think on the other end of that continuum is a whole notion of desire or a play. There seems to be very big discontinuities in that. So we either say that things are nice and playful or they're work. I often wonder if we can spread ourselves a bit more over that. What constrains me a little bit is that learning seems to be something that only occurs over time. So earlier we were taking photographs and kept struggling with how do we show something in the photo learning because it's an active passage of time?

The airport, which of course is my favorite place to shop nowadays, I found a new magazine which you probably should be aware of, I don't know how many of you have seen

it: "Tomorrow's Technology Today"(T-3). This should be changed to say "Thinking Today About Tomorrow" and they should be our answer to the T-3. So that was one clump of things, the scale thing I think is really important when you talk about this to clarify our questions.

I think we're talking about a desperate need for serendipity. The reason I love traveling often is weird things happen. I came on the plane with a man who collected badges. It was just so intriguing to me that a man could be going on trips around the world to meet other men who collect badges. He got his little suitcase of things out and we looked at badges. I never would have been interested in badges, except the way he talked about them was wonderful. The detail that he saw in these little badges was really interesting to me. I think serendipity is very important and that's the really only interesting motivating thing in my life, actually. Maybe I've had a boring life. Serendipity is very important to add to our experiences. The other one I think is querying. I guess I am not clear that I've thought at this event about thinking. I have thought a bit about learning and I think that I believe that all living objects learn and all manufactured objects don't. Maybe I'm wrong about it, but what puzzles me is in the future when I see manufactured objects, how will I query it about whether it's learned anything or is close to thinking.

We talk too much about windows that we look through as opposed to mirrors that reflect parts of us back. Another way of saying it is to know more mirrors, that would maybe reflect us.

I'd like to see if a group of people like this can challenge themselves to create something for longevity's sake. We were joking last night about the facade downstairs, that the pillars that looked like marble were made of wood. How I had learned in America to notice how distressed they actually were. Were they genuine stone or the degrees of corrosion

that had occurred. I often think about this in terms of furniture. Are we actually creating things that we'd like our great, great grandchildren to have as a heirloom? I can't really understand why any of these (digital) things would be passed down to generations of families. I'd like to see us as a group of responsible people continue this effort in some way and it's frightfully difficult to do these sorts of things. But can we get started on something that we think has some redeeming qualities, not just desiring qualities that would go on for generations and generations. Aesthetics take a lot of time. I just don't believe that we're thinking about these things in a big enough way to believe that there's anything really there beyond immediate gratification. I'd like us to try to and add some value to Razorfish and challenging ourselves somehow to do some more things like this. I think groups like this are actually better at thinking about bigger problems than individuals.

So thinking is very, very rare. I'm more and more aware of it than ever. And it's terribly expensive, especially today, because I feel that we have to do some learning and I'd much rather just think.

Craig Kanarick: This magazine is actually pretty fascinating. It's actually been out in the U.K. for about a year and half and they just brought it to the United States. You made a comment about the woman on the cover and there's always a woman on the cover who then does not appear anywhere else on the magazine. They had a great letter to the editor from an Anglican priest who said he loves the magazine, but he was embarrassed to buy the magazine and bring it home. Why did they insist on putting a woman on the cover? And the answer was, because the people in the newsstands are stupid and this way, they'll put this magazine with things like GQ and Maxim and FHN, and People rather than putting it with Computer Shopper, and Windows World, and PC World.

They wanted that audience rather than in with the technology magazines So it's purely a device for the stupidity of the people who run magazine shops to put it in the right place.

Sue Madden: I think that a couple of the important things I'm taking away is taking the time to think in our busy lives, and just stepping back taking that time. It's valuable time, taking the time to think is important enough for us to remember. I was interested that we got to a place where we could say "everything doesn't need to learn and maybe everything doesn't need to think." But we do. And I think that it will take a little longer to get to the point where we can figure out what exactly does need to learn besides us. I was interested, Craig, when we talked yesterday a lot about removing distractions and removing the Bridge page from what you see, and I think that there is some risk there in not having any unintentional consequences because some of those can be delightful. And making sure that things don't take away that possibility of delight.

I think there's value in our diverse backgrounds—having diverse backgrounds come together and talk about these things. And then my final thought is from breakfast this morning when, I don't know that we coined a new phrase from last night's movie is that we don't want someone to go "Yul" (Yul Brenner as the mad robot in "West World.")

Craig Kanarick: It's a neat update on "going postal."

Camille Habacker: Last night BJ and Neil and I and Constance and others were having a really interesting conversation about designer's ability to revolutionize with a conscience in a capitalist society which is difficult. I mean, "T-3's" (magazines) are great examples. You put the woman on the cover because it gets your message across, but it has far-reaching implications and so how do you negotiate those two

ends of the spectrum? And, you know, I was looking Craig's watch with the pictures on it (wrist camera). We all sort of leaped over there to look. The big question is, why? I'm not a gadgetphile, personally. I said yesterday I don't have a cell phone and I'm always interested in talking about that with people because nobody wants to talk to me about it.

Tucker Viemester: Maybe they just don't want to talk to you.

Camille Habacker: You know, I'm going to have to find out what it is. Because I have such an opinion on it that it annoys people. Asking that same question about gadgets like "What's the purpose? Why?" puts me sort in the realm of looking for "objects that learn" to be in the realm of digital conscience. It starts to make me think about what we could possibly be ourselves or what we could influence to get created as a thing. I have babies on the brain personally because I'm thinking of having one. But, I think that we could really do something interesting for children that acts like a cool parent.

It's sort of a strange idea I just came up with this morning, but there's a couple of books that you can get that I've seen being sold in the Museum stores. I can't remember what the titles are, but they're books that teach kids about manners, but it's done in a way that's not overly didactic. It's actually a really interesting read and I think really smart in its tone and style. So, it'd be kind of interesting to get at these issues of the greater good and respecting mankind and respecting the environment and humanity at an early age. And it goes to what Joy was just talking about in terms of heirlooms. Wouldn't it be fascinating if it (an object) started learning as you were younger and got passed down from generation to generation and all of a sudden, you have this digital thing that is a representation of family and community over time. I come from a small town in eastern Pennsylvania that started because of coal-mining and is very ethnic. It's all about com-

munity there. I think that had a huge influence on making me the kind of person I am today. I knew so many people growing up in college, that didn't have a community and I was often times very depressed at major gaps that I saw in their happiness. I always linked it back to the fact that they didn't really have a home town, lots of people move around a lot, two parents working, these are sort of the issues I grapple with on a daily basis considering that I'm planning to give birth and I work.

The other interesting thing that I just wanted to end with was if we are giving human qualities to these objects that we're thinking about be learning things. This goes in contrast with there being different ways to learn, but if we're thinking about learning objects learning like humans learn, I think figuring out other aspects of human relationships, referring back to issues of compromise and negotiation are really interesting. Having a learning object that knows when to shut up, or that knows when to be one way toward you versus another way, compromise and negotiation in my opinion are crucial to making human relationships work and get somewhere. So how do we deal with that in the technology itself.

Craig Kanarick: I think it's interesting to take the first story in the book, *I, Robot*. It is about a robot nurse-maid. It's a story about feeding the fear of trusting a thinking device with our children. It's one thing to have machines that are out in the dangerous coal mines digging up things, it's another thing to say that we can leave a child alone with a robot. That becomes the crucial debate, that the father thinks it's OK and the mother doesn't. And they struggle with that through the story. When you talk about heirlooms and things, "Bicentennial Man," the movie made from the Asimov book is essentially about that. It's about a robot that is given to two girls as their nurse maid when they're young and tracks them through four generations of family. He works for the family and stays as part of that fam-

ily through each of them becoming more human and learning about himself, he learns about emotions, but also learning about the family and about the people. My dad gave me his old stereo but that doesn't quite count. It's nice, it's quaint, it's beautiful, but finding the tubes for it are kind of a hassle. Things that learn, the things that think that maybe have more value than just that individual function might be able to outlast some of the component issues.

Peter Lunenfeld: We've done that, this notion of handing our children over to a machine. We did that with television. Television was sold to the American public as something that was going to be for children in large part. There's this wonderful ad from the mid-fifties for TV of "Be a Columbus in Your Own Home." It's a family sitting there and they're all Columbuses and little Columbuses. And Columbusinas. When you look at the rhetoric of that, then you look at the rhetoric of the computer—first generation PCs and then the Internet—there's always a huge corollary: that whatever gets out of the office and into the home is always about how your children will fall behind if they don't have it. That thing will open this huge opportunity for your children. As a marketing strategy it's got 50 years, incredibly powerful legs. So I don't think that's going to be a problem. The question is can we 'undo' some of the horrendous damage we've done with television with these new learning objects?

Craig Kanarick: I wonder if, to go back to Tucker's original idea about the purpose of sleep is occupying the kids so you can go out and get food. If in the future, we won't need to sleep, we'll just watch TV.

Tucker Viemiestter: But you know your brain waves will do the same thing while you're watching TV or sleeping.

Craig Kanarick: Put the infant in front of a PC, or a TV and

go out and gather some food. That's why I feel so rested after watching Gilligan's Island.

Stephen Turbek: If only kids slept that much.

Tucker Viemiestter: It's enough to run around and grab some carrots and bring them back to the hutch.

Stephen Turbek: Two themes came up several times in the discussions. The first theme was 'control'. Whenever you make a tool, it's to exert some control on the environment. Making things that learn is an attempt to exert more control. There is a flip side to this as we've perhaps had too much control, or too much of an effect on the environment, to its degradation.

If we ever get around to creating machines that learn, people are ready to fear a loss of control over them. Frankly, I don't know if we deserve to fear the loss of control because, regarding the environment we've not done a very good job of exerting responsible control. Perhaps we should lose control.

Or perhaps we can be helped to exert responsible control. The ecological context adviser device I mentioned before could help us understand the consequences of our decisions. Right now we make these decisions in a vacuum, which is like a global game of 'the prisoner's dilemma'. This is a situation where two criminals are questioned separately. They're told if they rat out the other, they'll serve less time, but if they both keep to their alibi, they'll go free. With the environment, one person can cheat, as long as the others don't. Unfortunately, we're always cheating. If we were better informed, there might be less cheating.

The second, more subtle theme that has been running through discussion is respect for one's fellow human. Many cultures do have helping things that learn, they're called ser-

vants. Because of our culture, we've accepted the idea that servile labor should be replaced with machines. Otherwise, this discussion is easily solved: just have a servant do it. I find that quite nice. The fact that intelligent machines are years away doesn't make us go back to using servants. I think you all should be congratulated for wanting to have people fulfill their higher potential.

Craig Kanarick: I had a little internal struggle with doing this because it seemed that everybody was very struck by the notion of wanting to have time to think and not to learn. I had a slight twinge that maybe this session was—well, what have we learned? Then people said, “You know I didn't learn anything, I just had a really good time thinking.” And probably one of the things that I was struck by most is how complex that relationship is. I know that I value learning a lot and may be progress oriented, I need to learn something, I need to get to the next level, I need to move forward. And not so excited about not getting somewhere. So for me, I think that learning was more important, not ever realizing that all that learning is hopefully contributing to my ability to think, not just my ability to do. If I spend all my time learning, and none of my time thinking, then it's a little bit for naught.

So I'm happy that I have a better understanding of the relationship between the two. One of the things that was really dramatic for me—and also something that clearly was for other people because they talked about it a lot—was scale—that the we can be an individual or group or anywhere in between. That a thing can be an object, or a city. Early this week the Science Department did the same exercise of drawing their favorite thing just as you did on Saturday morning. Most of the people in the department drew pictures of people as their favorite thing (partially because the person administering the exercise used “wife or family” as an example). We didn't mention it this time on purpose to see if the

results would be different. In fact, they were different. Nobody put down a person. I don't think anybody put down a person as their favorite thing. Did we?

Rebecca Odes: I had my eyes.

Craig Kanarick: Your eyes. So you put a part of your person. It was very exciting to think about things on a different scale. Think about the planet as a thing, and the earrings as things. I do think that it came up a couple of times—children need to learn and they're a thing. We, as society, needs to learn—and society's a thing. I thought that was an exciting way to discuss things.

The other scale thing that I was interested in was the thing that came up at the very beginning and filtered through all of the rest of the conversations but was never explicit again and that was the idea of the time horizon and our scale of looking at things that were important. Now, versus important in a 1,000 years or important in 10,000 years. We weren't outside talking about the time horizon, but we were clearly all focused on the short time horizon, and at the same time very concerned about the longer one. So we had this passive-aggressive relationship with the time horizon and the future. There was sort of a good balance of “I'm excited about the future because we can do all these things, but it might also cause these problems—but we can hit undo, but we really can't hit undo.” So it seems like we're all still intrigued by it. None of us have lost hope about where we're going. I guess, at a minimum, that's exciting. But I think that this notion of perspective, whether it be in scale or whether it be a now versus later or a small thing versus, the entire universe needs to learn was very helpful for me in something that I don't think I do a lot of the time.

I also realized that we still strive to improve ourselves and

that's something that I think is what makes us humans. It's something that I hope never goes away. Will we continue to talk about things that we want in terms of improvements for the future? That to me is something that was helpful and probably not surprising. I didn't expect anyone to come in and say, “Ah, everything's fine. I don't need to worry about it, it's really not that important because I am in a blissful state.” Until SOMA is distributed to everybody and drugs are accepted as the appropriate way of replacing television, things that think and everything else in the end, we probably will continue to strive for that. So we do have different ways of getting there. Some people need to reduce clutter. Some people need to augment themselves, but nonetheless, everyone has their own small incremental change that they want to make, it seems, to get better. And to make not just themselves better, but everything around them better.

The other thing that was interesting to me was also our passive-aggressive relationship with efficiency. And I thought that was a very interesting discussion, especially for people like us here who I know are all very busy. And is, I guess, a good segue into thanking everybody for actually taking the time to come here. Because, it wasn't clear that this would be the most efficient use of people's time in terms of where they need to go in this world, or what they needed to do. It wasn't as though we were all going to come out of here with a great invention or something that we could take back to our lives. Yet, it was about spending time together with a group of smart people and talking and seeing what would happen. So I'm glad that we were conscious of the relationship with efficiency and sort of catching ourselves, saying, “Yeah, we do want to be more efficient, we do want things to be more effective, maybe not efficient.” But, efficiency in and of itself isn't necessarily a great thing. So, whether or not I've learned anything over the weekend, for me, isn't really that important any more. But I am glad, really glad that I've had a

chance to think for the last day and a half. And to do that with all of you.

So I want to thank you for putting so much energy into this and for participating. The other thing I want to do is thank Anne Young for putting together this entire flawlessly executed weekend. Anne does that for me every single day so I'm lucky that I have somebody who is so fantastic to help

me out every day. And Ben Davis who did a lot of work on this as well. So thank you, Ben, for all that.

Ben Davis: My pleasure.

Craig Kanarick: And our faithful Dutch note-taker, Adam Eeuwens. After that whole discussion on on "Gilligan's Island," he turned around and said, "I have no idea what this is."

Acknowledgements

Thanks to Anne Young, Ben Davis, and the staff at Tarrytown House for organization, thanks to Adam Eeuwens for note taking, thanks to Bernie Dekovan and Technography, Inc. for facilitation, thanks to Joan Wai for tape transcriptions, thanks to Ben Davis for text editing and image processing, thanks to Gail Abernathy and Jackie Goldberg for design, and thanks to Stephen Turbek for seeing the book through to completion. A special thanks to the participants for being so patient for these results.

Colophon

This electronic book was set in Helvetica 55 Roman, with titles set in Helvetica 65 Bold, both distributed by Adobe font Library.

About Razorfish

Founded in 1995, Razorfish is a global digital solutions provider that helps leading organizations generate competitive value by leveraging the power of technology.