Proceedings Summary

MIT-NSF Workshop: Smarter Service Systems through Innovation Partnerships and Transdisciplinary Research

November 20-21, 2014
Le Meridien Cambridge-MIT

Organizing Committee
Richard Larson, MIT
Kent Larson, MIT
William Rouse, Stevens Institute of Technology
James Spohrer, IBM

Editor
Nicolene Hengen, MIT

Event site: http://web.mit.edu/mitssrc/nsf/

Sponsored by
The National Science Foundation
Fujitsu
IBM

Table of Contents

Executive Summary

The service sector is the driving engine of the US economy, contributing 80% of GDP and employing the majority of high-skill workers. Health, education, transportation, telecommunication, and utilities are but a few examples. These are complex socio-technical systems comprising technologies, people, organizations, and information that provide services, creating value. Sustained economic growth requires research directed at fundamental understanding and improvement of these systems, leading to engineering of new transformative "smarter" service systems.

This two-day NSF-sponsored workshop brought together 100 individuals – academics, private sector leaders, government professionals and NSF leadership – to move us closer to a new exciting problem framing and research agenda setting for Smart Service Systems. The participants included eleven teams recently funded by NSF for their innovative Smart Service systems research programs.

A key goal was to move towards creation of service systems that make intelligent use of technology, impedance-matched to the needs and responses of humans, thereby maximizing their impact. The technologies - both hardware and software - encompass virtually all fields of engineering. They include information extraction and interpretation from sensors, actuators, smart phones, and people. They include robotics, machine learning and industrial engineering. Real-time technology adaptation, customization and personalization are increasingly important. As we go forward, research will range from small to large, from the impedance match of one-on-one bi-directional human-technology interactions to the agile adaptation of larger 'systems of systems'. Creating integrative technologies that match with people's needs and capabilities is key for productivity growth and higher quality of life.

The objective of "Engineering Smarter Service Systems" is to discover basic fundamental principles that govern cooperation and adaptation in human-technology interactions and thereby develop new foundational theories for service system design. Such research will lead to new methods, including those that support personalization, customization, and adaptation at multiple scales. Given the ubiquity of service systems, this field has the potential to address broad societal needs, from individual needs in health and education, to urban infrastructure including transportation, to national and global issues such as security.

Design and integration of technology in service systems is centered in engineering, but contributions from social and behavioral science, computer science, and design science will be essential. In the two-day workshop, speakers having diverse backgrounds in engineering and social science emphasized this engineering-social/science symbiosis.

The overarching goal is to foster a new discipline of service systems engineering that supports human needs via smarter adaptation, customization and personalization through fundamental crosscutting research. Creating innovative impedance-matched human-technology partnerships has the potential dramatically to transform society at all levels.

This Proceedings Summary provides an informal "notes-level" distillate of the presentations and discussions in the two-day workshop. More details, including videos of all presentations and speaker biographies, are available on the Workshop web site: http://web.mit.edu/mitssrc/nsf/index.html. In addition to this Proceedings Summary, a Final Report of the entire event will be available soon.

Richard C. Larson, Chair
MIT-NSF Workshop
Smarter Service Systems Through Innovation Partnerships and Transdisciplinary Research

Session 1: Welcome

Cynthia Barnhart: Service systems align well with her research and observed that evaluating solution approaches is much less difficult than translating the research into practice, especially as the number of stakeholders grows. From her own experience, for example, changing scheduling practices and resource allocations based on analysis was relatively easy when only industry was involved, but once the FAA, airports, airlines, and passengers were included it was much harder. She congratulated the 11 NSF grantee teams, noted the range and importance of the problems, and thanked the organizing committee.

Session 2: NSF Perspectives

Grace Wang: Over the last two years people at NSF have been working on defining smart service systems to build a tangible research agenda – the ultimate goal of this workshop. The definition is based on the work of James Spohrer and Paul Maglio. A service system is human centered; it requires interactions between people and the physical and virtual realities. By integrating enabling technologies with a goal of serving people better, it also incorporates value (co)creation. A smart service system is capable of learning and dynamic adaptation; is based on data the system receives, processes, and transmits; and is capable of decision making. It then is also able to respond in (near) real time with a goal of creating value for humans.

The challenge is to come up with a very solid research agenda. Applications include smart manufacturing, smart cities, personalized health care, etc. But there is a common set of design principles that need to be defined; the real challenge is determining the fundamental knowledge and principles we need to really enable smart service systems.

To understand the current and future human and social context, and to understand future market need it is essential that there be academic-industry collaboration, not just at the translational research stage, but starting with fundamental research. In order to design the system architecture, we need to understand who is being served and the human and social behavior aspects.

Day 1 is focused on academic industry partnership, along with examples of integrating technologies into a smart service system, while day 2 is focused on the fundamental knowledge and principles needed.

Academic-industry partnership goes far beyond tech transfer; it is about reciprocal knowledge transfer. Innovation is not linear, and you often go back and forth between stages and between academics, industry, and government, as noted in the 2012 PCAST report. NSF believes that the innovation ecosystem ideally should include the university, industry, customers, private investors, entrepreneurs, and startups working closely together. They should exchange ideas, knowledge, talent, and capital. NSF's role should be as catalyst: seeding good ideas, pulling communities together, building innovation capacity, expanding it, and making sure it's healthy.

Based on this concept, NSF has realigned, restructured, and initiated new translational research programs: e.g., Accelerating Innovation Research (AIR), as well as Innovation Corps that incorporated the SBIR and STTR programs with a focus on technology commercialization. We also restructured the Building Innovation Capacity (BIC) program to focus on smart service systems, both to build a community and as an experiment in thinking about the future and learning how to tackle many grand challenges and ensuring the research is relevant.

Session 3: Partnerships for Innovation: Building Innovation Capacity (PFI:BIC)

This workshop is the second round of its focus on smart service systems. PFI:BIC focuses on academic-industry partnerships led by an interdisciplinary research team with at least one US-based industry partner that has commercial revenues, to design and create/transform a smart service system. Additional information is in the solicitation released in September 2014.

Session 4: Ingredients for Successful Academic-Industry Partnerships

A. Jim Spohrer, IBM: Translational research requires industry partnership; academics can't do it alone. From IBM's perspectives, universities have four major missions:

- learning (teaching)
- discovery (research);
- engagement with the community
- convergence in a world that is increasingly complex and specialized.

Universities have a lot of knowledge in one place; they are also great labs for smarter service systems because they are basically small cities.

IBM has multiple criteria for rating universities: research relationship, recruiting (of students), revenue from research partnerships, responsibility (e.g., IBMers becoming adjunct faculty), IBM wants "T-shaped people" who have breadth across many disciplines, sectors, and/or regions (see tsummit.org).

Watson is an example of research collaboration; part of the collaboration is inspiring each other to go after such grand challenges. IBM takes the research and turns it into platforms that universities can use, e.g., smart cities platforms. The next big thing is cognitive platforms: builds on the resurgence in AI and cognitive science to understand how to augment and scale human expertise. Example: Baylor University biochemists identified 60,000 papers they'd like to have read; the IBM cognitive system ingested those but also had a model of the researchers' interests. With the system the team discovered 10 new molecules in three months; a 10x creativity increase.

IBM buys many companies (about 12/year at $100MM revenue each) and scales the revenue; 66% of those were started in a university ecosystem; e.g., SPSS. It takes 15 years on average from startup to IBM acquisition. Other parts of IBM are shrinking; often employees from those areas move to academia and use IBM platforms to create new startups. This constant flow of people is translational research.

Key question: what % of a company's products/services change every year? Answer 30 to 50%. What % of university courses change every year? Note the impedance mismatch. Translational research helps overcome that.

B. Joe Sussman, MIT: Many global issues require an interdisciplinary approach, e.g., climate change, energy, mobility. Universities contribute innovation and interdisciplinary research. MIT has worked for 25 years with the East Japan Railway Company (JR East). JR East endowed a chair when it feared it wasbecoming too insular. MIT has learned not to push too hard too fast when language, culture, and missions are barriers. Develop relationships with people at different levels, foster exchanges and partnerships.

C. Daniel Siewiorek, CMU: Five models for industrial relationships.

1. Industrial consortium, e.g., Data System Storage Center, includes all the disk drive companies, stable, get together a few times/year, students flow between industry and university (both ways). 1a. Ecosystem e.g., Quality of Life Technology center also includes end users. Core: people who provide services to others with disabilities. Also technology providers (e.g., Fitbits); thought leaders (e.g., AARP). Semiannual get togethers of about 70 people with amazing energy.
2. Startup companies, e.g., QoLT Foundry with entrepreneurs, interns, spinoffs
3. Industrial labs, e.g., Intel created a Science & Technology Center for looking at cloud computing over five years (joint across five universities)
4. State and academic partnerships e.g., Pittsburg Infrastructure Technology Alliance (PITA), addressing issues such as infrastructure, transportation, telecoms, energy, public health, water; involves development and outreach, and collaboration with a PA-based company providing direct or in-kind support with 2:1 leverage ($2 from industry per $ from PITA)
5. Academic industrial partnership e.g., PA Digital Greenhouse, starting in '97 or '98. Goal: spawn jobs in technology (systems on a chip). CMU, Penn State and Pitt produce over 1000 EE, CS, and CE grads/year. 300+ companies in the area with over 10,000 employees. There are multiple membership categories, members provide research funding and can direct projects; process of RFPs and a Tech Advisory Board that helps select the projects. Typical project: $200K/one year with shared IP. The center also provides courses for members and students.

Overall a broad range of sustainable models. They all take time but can be long-lasting and a source for new ideas and funding.

D. Patrick Harker, University of Delaware: The JP Morgan Chase-U of Delaware Partnership started with eight-month, full-time internships so students learn about working with massive information systems. Now JPMC has a joint PhD in financial service analytics with UD.

Key lessons:
1. Follow your alumni to build industry relationships.
2. It's a collaboration, not a transaction. Don't let the lawyers take over IP.
3. Stay true to your mission. They run a bank. We create and disseminate knowledge. But as a land grant university also do economic development.
4. Communicate. With different clock speeds and levels it takes time to learn to work together.
5. Keep it growing. And think broadly about different disciplines you can include and be a bridge for (e.g., with cybersecurity initiative with the banks and military and intelligence communities).
6. Trust each other. Trust has to be earned.

There are 4,700 institutions of higher ed; average student age is 36; costs are far too high. We are at an inflection point with an opportunity to reinvent higher ed with help from industry partners.

Session 4 Q&A

Q1: How do you resolve potential conflict in terms of publication? A1: most of IBM's collaborations are under our Open Collaborative Agreement, so it's not an issue. For the rest we have mutual review; it's important to establish that up front. With IP, you create it, you own it; if we create it, we own it; joint creation is joint ownership with royalty-free licenses for both parties to commercialize. That's the industry standard we work under. B1: Open source, unless it's a national security issue that has protocols for moving to rapid publication.

Q2: How do you interface with smaller businesses? C2: with smaller companies we often do projects in capstone courses e.g., HCI, software engineering or EE. Companies bring us a problem, students develop a proof of concept that companies can then work with. D2: We call it spin-in, same thing. Also, some companies need physical incubator space U Delaware can provide. We want to do this so companies stay in our state.

Q3: How do you navigate IP with a single partner vs. a consortium? C3: Engineering research centers use an IP-sharing model – the NSF model initially; subsequent we have pre-competition models where companies don't bring anything they consider proprietary. B3: At MIT many consortia are about creating a safe harbor where industry, government and large/small companies can come together off the record. A3: some of the most successful multivendor relationships IBM has had involve a professional association or other nonprofit organization with a share purpose for e.g., standards, skill development, etc.

Session 5: Industry and Government Perspectives—Health Care & Smart Service Systems

A: Hans Hofstraat, Philips Research: Smart services are enabled and facilitated by IT in health care. In health care information is highly distributed, and there are multiple users: providers, administrators, and patients/families/consumers who all need a different interface and capabilities. Philips has developed and is introducing the Health Suite Digital Platform, which is designed to bring together the information available across the continuum of health, from prevention to diagnosis and treatment and rehabilitation.

It's a cloud-based platform with our own API's and also ones we co-create with research partners in clinic and academia – open innovation with all stakeholders. We have a collaboration cloud for interaction among people, a device cloud (for imaging systems, vital sign monitoring systems, etc.), and a data cloud for analyzing data and information. Solutions are dedicated to each type of user, with different modes of interaction.

Extended example of work with Banner Health in Arizona, an ACO where they have jointly with Philips and Academic partners, developed a link between hospital and home care delivery. Key is implementation and measuring outcomes, continuously improving and learning from feedback.

Definition of service innovation: Research with stakeholders in the system to co-create methods and tools that identify local needs and are adapted to the local context and the local user. Development of service experience flows. Strategy of engagement happens via different types of interfaces, which is a big challenge, since not all stakeholders are equally motivated.

B: Heather Woodward-Hagg, VA (VERC, Veterans Engineering Resource Centers): Focus on challenges in moving forward in transition to a service delivery system (vs. silos/stovepipes that are traditionally the VA). Health systems engineers in the VA (five-year old program, 350 staff across four centers) integrate engineering into health care delivery.

Four approaches:

- Professional development of clinicians into systems thinkers
- Health systems point improvements done locally e.g., Lean and Six sigma

- Manage systems
- Design delivery and deployment of "big VA" enterprise health systems aligned with DOD

No commercial partnership but many academic ones. Work onsite (156+ locations) to build relationships with providers and patients to understand the problems at the site.

Lessons learned:

- Partnership with nontraditional systems engineering departments and younger faculty work better: greater willingness to do applied work vs. pure research.
- Most systems engineering approaches need heavy translation to be relevant in health care delivery because of the large social aspects. VERC does the translating, not providers. Health services researchers and implementation science researchers provide social contexts, so this relationship is key. They benefit from the robustness of our technical tools. We benefit from their perspective.
- Hire for resiliency, not skillset, because large government bureaucracies in particular are not ready for systems approaches, even if they think they are. OK to sacrifice some technical validity initially to get in the door and build relationships.
- Journey is far more important than the destination. It'll take a long time to get there.

VA future state service delivery model: to date VERC has focused on system solutions within silos, but with the changes Bob McDonald is charged with, goal is to integrate information and processes across VA: Veteran benefits, VHA, and National Cemetery Association. Policy has driven many of the disconnects and the organization supported them. So to change: MyVA is envisioned as single point of entry; requires integrating technical and service systems as well as large organizational restructuring.

Key challenges: limited in ability to design and deploy enterprise systems because of upstream and downstream effects, e.g., primary or specialty care capacity limits, unknown care standards, unknown ownership of processes, and tension between patent needs, what can be legally provided, and what staff is able to do. Also IT, contracting, HR and other shared services organizations and their processes and infrastructure are outside the control of VHA.

Example of work done: My HealthyVet portal. Not well integrated with EMR, low utilization (40% of patients use it) but has online benefits application process and mass customization to design appropriate delivery for primary care, and integrated clinical decision support within some specialties (e.g., stroke in ED, stroke rehab), but no integration across them. Many standalone patient-driven apps.

Key questions for this audience: How do we ensure that our new VA service delivery model, which is focused on connecting patients with the benefits they qualify for, also integrates our clinical delivery systems? That is, how do we make sure that what we do in VHA leverages the information you get from the larger VA service delivery system? How do we ensure that our methodologies are aligned so we can continuously improve and align across agencies so we don't end up in the same silos as before?

And a general challenge: how do we not repeat prior attempts and failures to transform health care, such as initial promises made around the electronic medical record? More specifically, can we define the burden of evidence for a robust health IT system, along with standards for a robust health IT system, so that we don't disrupt operations but more importantly don't inadvertently hurt patients?

C. Tetsuro Takahashi, Fujitsu: Fujitsu has many businesses including hardware and software and many devices. Recently Fujitsu has set out a vision of a human centric intelligent service society, with three dimensions: people, information and infrastructure. Such a system needs to empower humans and create intelligence by analyzing big data and information. All this requires connected infrastructure.

Here is an example: the KanTa electronic health record system that integrates across all of Finland. The system has three key features: eArchive, ePrescription, and eAccess. eArchive is a medical record that doctors can see. ePrescription allows doctors to write prescriptions, that pharmacies can access; while eAccess allows patients to access their own medical records from anywhere. The project is led by the Ministry of Social Affairs and Health, which helps because such a system has many stakeholders.

Session 5 Q&A

Q1: To what extent are your partners university researchers who are developing technologies for integrating into these smart health care delivery systems? B: Almost every technical product we have at VERC originates with our academic partners; we do a lot of translation and adaptation of these. We have a lot of flexibility, which is rather unique for a government agency. I'm happy to take your card. A: We have multiple collaborations, and partnerships with academia play a role across them, such as tools to speed up surgery or interactive intervention tools. We really like small and simple, high quality devices. We also need to develop solutions that are robust and tested in clinical practice. We also have academic collaborators in industrial design and behavioral science. C: Collaboration with academia is necessary for Fujitsu also, especially in basic research. Moderator: and many universities I visit have academic medical research centers so there's a lot of interaction, for example at Stanford.

Q2: First, is VISTA still being used at the VA? And you mentioned failure of the Electronic Medical Record (EMR) system effort. Do you think that is because of technology failure, policy failure, or procurement management failure? B: VISTA is our historical EMR; we're currently developing the next generation. And I wasn't speaking specifically about the failure of EMR in the VA. Ours is 20 years old and was the first; it was really significant. I was referring to other EMR implementations I've personally seen that were disastrous. They cost billions. They have not improved; have not been integrated into work flow. I believe, optimistically, that we as an industry underestimated the complexity of the social system in health care. If you want to be pessimistic there's a big medical industrial complex. It may be a bit of both. The issue with Vista is that while it's robust, we've allowed unlimited customization, so now we have 156 versions in 156 hospitals. So any development has to be adopted 156 times. So we have a really robust EMR. But in my opinion we have not optimized its use.

Q3: If we really want to make a transformation, who are the right people to be at the table? Who is the disruptor? B: I think you need to think about what drives disruption in industry: the customer. In my opinion we don't have an engaged customer in health care. Right now the health care system is designed to benefit a small subset of stakeholders, such as medical device firms, big pharma, physicians. We need an engaged customer base that understands the deficiencies, and currently they don't understand how significant the issues are. So how do we create that engaged and informed customer that demands something different? That's now happened in the VA. We're a fishbowl for healthcare, but everything I describe is going on everywhere else. But we've gotten customer demand and that has accelerated and facilitated transformation for us. It's interesting to see what happens. A: I think transforming the health care system is complicated because there are so many actors. So we need to collaborate with people who are motivated to change. Banner Health is one example.

Session 6: General Questions

Alex Medina-Borja: The first part of the workshop was planned to be a discussion about industry, academe and the translational space in which PFI:BIC operates. And we'll have the poster session of 11 awardees at lunch. After lunch we believe the posters will inspire afternoon discussion about whether there is a fundamental research agenda that is not covered by current NSF programs. We hope that this is really a new agenda on smart service systems that is about more than engineering. This afternoon's discussion will be about the future and fun things we can't imagine today.

Q1: In terms of expanding access to trying to solve specific problems we should perhaps look at other communities working on some of the specific issues, like mobile health and wireless sensor networks that people in my community (UC Merced) have been working on, with medical teams; one of my colleagues at University of Memphis is leading a large NIH funded center on that. Are there other communities and funders also working on that?

Another issue, that was mentioned earlier, is the need for a meta-agreement, cooperative research agreement, between an industrial partner and an academic partner in which broad issues are agreed to even prior to a specific proposal. Because once there is a proposal there's not much time. Two solutions: keep the technology transfer office out, and get the lawyers out of the meta- negotiating conditions. You could have a template in the call itself that says you should have agreed upon the main issues in a relatively non-binding way – e.g., with respect to royalties, with respect to sharing IP should already be done. Then if the application is successful and gets funded you don't need to renegotiate everything.

Dick Larson: MIT has a lot of experience with that, and has one of the largest industrial liaison programs, as well as experienced lawyers, some of whom are of the 'can do' type, so we are very experienced in dealing with IP issues.

Q2: Based on the last comment; I think our discussion group on LinkedIn is a good place to share what types of issues should be in such an agreement, or other checklists. There's also a lot of good material already there.

Session 7: The Role of Human-Centered Design in Smart Service Systems

Paul Maglio, UC Merced (moderator): *Service Science* is our flagship journal published by INFORMS. I'm the editor-in-chief. Please submit your work that crosses disciplinary boundaries. Yesterday we published Pat Harker's piece on higher education as a service system.

Second, to summarize our first workshop funded by NSF on developing a research agenda for smarter service systems, held in March in Washington, DC: we had 40+ people from academia, industry, government; including about 12 who are also at this workshop.

We defined human-centered service systems, as configurations of people, information, organizations, and technologies that operate together for mutual benefit. Modeling is key, but it's hard due to unclear boundaries and it's hard to model people. Because of the interactions and interdependent behaviors, human centered systems have emergent properties.

Applications include urban infrastructure systems, population health, manufacturing with service additions, information and communications technologies, smart services, platforms. Our research agenda has six themes (see report): theory/principles, data, modeling, assumptions, design, and measurement. Also education.

A. John Carroll, Penn State: Key principle: co-production as a principle of service systems science. Co-production is an important model for us to understand if we want to facilitate and coordinate human activity. It's moral, efficient, and innovative. It's the source of disruptive innovation to frame services.

Often in services the recipient is passive – e.g., VA patients. In co-production the recipient takes an active role. Definition: in co-production provider and recipient contribute reciprocally to the success of the service. Most of the services Paul Maglio alluded to involve this type of co-production.

Co-production is not the only way to theorize services, but it is the way to do so at the level of human activity if we want them to be successful. Origins: Jane Jacobs, Death and Life of Great American Cities. Key example: elderly people at home keeping an eye on neighborhoods. Elinor Orstrom (won a Nobel in 2009) coined the term when analyzing the service innovation of putting Chicago police in cars rather than walking the neighborhood (performance got worse). Example: Edgar Chan, No Throwaway People, on Timebanks, in which time is treated as currency to value an activity by how long it takes to do. Timebanks have been shown to help bring marginalized people more to the center of their community, including by British NHS who use time banks as a prescription for mood disorders, as well as enabling aging in place.

Any service involving learning, emotion, health and wellbeing, and social support cannot succeed if it is not co-production. The VA is learning this, and universities know this. Co-production emphasizes shrinking the top-down, central mediators. It is not pie-in-the-sky; people in HCI have been working on this for 35 years.

Lessons from HCI: Identify the active role of the recipient. It's more than receiving a service. In HCI we use scenarios and personas for that vs. service blueprinting (flow charts) that are used in service technology today. You need more than flow charts. You also need to focus on breakdown and workarounds from the recipient's point of view in service relationships. Think about services as appropriable by recipients, instead of customer journey maps. Probe, articulate and refine the recipient's experience of the service. In HCI we do role play and experience prototypes. Don't just enumerate touch points, that's not enough. In HCI we learned 'you can't add it later'. No one can add usability to a poorly designed piece of equipment. My guess is this will be true with services. We need to take on co-production from the start.

B. Kent Larson, MIT Media Lab: Cities are very important, as 90% of population growth will be in cities, as well as most innovation. But there's poor understanding of. The key research question is: what enables vibrant, entrepreneurial, high performance cities. We think there are three key parameters: density, proximity, and diversity. Density is number of residences, offices, cafes per unit area. Proximity is how many options do you have from e.g., your home to different work places, recreation, shopping, etc. Diversity includes demographics but also housing options, activity options. All these are managed through new technologies and services. Modeling integration, sensing, etc. to understand human behavior and thinking about how people move, eat, are housed. If you can do that, quality of life goes up, jobs go up, and resource consumption goes down.

Key components: people-centric sensing; enabling infrastructure. Started 10 years ago with a highly instrumented living lab in an apartment building. Sensors everywhere so we knew where people were, what they were doing, the systems they interacted with, state of the environment. We also could use sensors to make tiny changes, e.g., use blue light for some tasks, which saves energy. Typically young people love it; Boomers don't. Now we just started with a 'sensible floor' that recognizes activities; since people don't like having cameras everywhere.

A second area is "places on demand" – transformable homes and co-working; goes beyond internet of things. [Shows prototype of 200-square-foot apartment with a queen-sized bed that converts to an office with an eight-foot desk or a dining space for six people.] It's possible that housing will become a service model rather than being rental or ownership.

A third area is food for cities: producing food at scale near where it's consumed. Aeroponics multiplies productivity/square foot by six; but imagine skinning the Prudential tower with lightweight food production for a 1080x multiplier; along with 90% less water and 60% less fertilizer. Biggest challenge: behavior change and determining the service models that will work best.

Fourth: mobility on demand, e.g., service models for lightweight shared use electric vehicles of many modes. Key mode is walking. Also shared bikes, trams, on-demand; no fixed routes. We like working with industry on this. [Shows example of city car]. Now doing a three-wheel, one-person vehicle. Most city trips are 3-5 km; vehicle also converts from people mover to goods mover when demand is low. We don't know the service model but we think the goods mover could pay for it if done right. It's most powerful to combine autonomy with vehicle sharing and electrification, with sensor networks to know where people are; then you can dynamically meet demand [visual of Kendall Square simulation model] with shared vehicles eliminate parking, turn lines, and accidents. But it requires a compact high-density district.

We're looking at putting it all together, and with companies: urban co-creation. Our ongoing analysis of Kendall Square using real time twitter feeds shows where people are active. We can project a walk score under various scenarios. We are building 3D scenario tools to map data, construction cost, energy, occupancy and physical elements. We are now also focused on human behavior, e.g., seeing there is no interaction between people in two office buildings. But add cafes and shops and that changes significantly.

Key message: we tended to design cities for machines, e.g., optimizing traffic flow. Cities designed for people work better. These people-centric cities will lead to innovation, and then you get smart cities, eco-cities, low-carbon cities by default. Start with people and services and add technology when useful.

C. Brian Scassellati, Yale: I look at services that are 1:1 engagements. I work with kids. Our vision is to create robots that supplement educators, therapists, parents and clinicians, for example, teaching social skills to children with social and cognitive deficits to accomplish some long-term behavior change. [shows two videos, not shareable because HIPAA issues, that describe this vision with a 12-year-old with autism practicing vocal intonation, and English language learners learning when to use 'do' vs. 'make' for the Spanish verb hacer].

Key points: we need systems that can engage over long periods with kids, and also adapt to each child, as a peer, since this is how kids learn best. We have to do more to build intelligence, i.e., smart services, into that robot.

Session 7 Q&A

Q1: Please say more about Ostrom's notion of reciprocity, as well as the notion of reciprocity as part of a trust relationship over time. A: People don't miss the fact that both provider and recipient produce, but the miss the active role of the recipient. To enhance trust, try to make things reciprocal in the event, so there is no dependence on long-term interactions.

Q2: How will cities scale? Today you have central cities, suburbs, mega-cities. Do you see fundamental change in that structure? B: In Europe many cities like Paris already consist of neighborhoods. I think that is a model. Atlanta isn't like that but I met recently with the mayor and they're looking at changing zoning to create more neighborhoods or in effect micro-cities. So over time we will evolve to that model via policy, or else via the market.

Q3: In these talks it's about humans in these systems as a uniform mass, like a fluid. But in reality our systems are more hierarchical than that. How do we model that hierarchy, for example incorporate key individuals such as the policemen? B: You need good structure, and then think of it as a jazz lead sheet where top-down meets bottom-up. We want this balance. You can see it in good buildings where the architect has a clear coherent vision, but then people adapt the space to a particular use at a particular time. C: It's not just a heterogeneous group but we will shape and impact existing relationships; the technology is part of that social structure. For example, a colleague put a robot in a class of kids.

The robot responded more to the 'quiet kids' who then became significantly more popular over time. I think we have to be very careful as we think about this issue: we don't just use the network of social relationships, we also have an impact on it. A: I agree about the dynamics; to the extent that we see users as heterogeneous that means we don't know what we're doing yet; in HCI we have concepts of domain theories, requirements analysis; we use modeling techniques such as modeling people with key roles in scenarios; this helps to classify kinds of stakeholders, kinds of users and how they relate to the technology. So you're right and that indicates we have far to go to understand the domain well and design for it.

Q4: It's not just co-production, but also co-evolution, as you note. So an ethical question: will we like the change in society due to 200 square foot apartments? What if the change is something we don't like? A: isn't this a dilemma engineers have always faced? Q: But the evolution is faster now. B: So we need living lab experiments, testing with people earlier than traditionally done. A: Agreed. And living labs were first talked about by Robert Parker, an American sociologist, in 1925. It's a great idea.

Q5: Based on these talks, about co-production, personalization, and people flows and spaces, I observe that you have to build pretty significant models of 'the other' to have effective, moral, efficient co-production of personalization. I didn't see that much in Kent's work. A: Yes for co-production. B: There was no time for it but we've had a Smart Personalization Group in the past, and the Home Genome project. But we've moved away from personalization because transformation requires minimal personalization; better designed spaces require fewer personalization options. So with design comes more focus sometimes, and higher profits by the way.

Session 8: Framing the Smart Service Systems Field with Human Factors at the Core

Dick Larson: Services are usually defined by subtraction: anything that's not agriculture, mining or manufacturing. Technology has played a key role in service evolution through automation of tasks formerly done by people, such as telephone operators or elevator operators, or else by substituting customer labor for paid labor, facilitated by technology. Moore's law is very important in this (see Session 12).

Intellectual foundations for the science of services include queueing theory. A.K Erlang invented queueing for the Danish phone company to determine the required size of the central switch. But also an understanding of human behavior, as in the 'slow to arrive elevators' example, in which complaints fell dramatically once mirrors were installed in the waiting area, even though the elevator delays did not change. This can be generalized as the role of diversions in queueing situations.

Key questions:

- Does smart necessarily mean computers? [no].
- What is the appropriate balance across technology, human factors, systems analysis, and operations research?
- Is innovation by accretion a sufficient framing of the problem? [No; e.g., VA is now restructuring; adding a dean doesn't solve much]
- Is there any service system for which a single performance measure suffices? [Probably not – suggested axiom: there is no silver bullet].

Issue: we want both academic rigor and partnerships with industry, as well as a human touch. This is complex, especially for we silo dwellers. Academics tend to overly formalize. So most of our work is with the people who react, behave; humans first, all else second or third. Example: flu pandemic: lambda (contact frequency) is based almost entirely on human behavior; transmission rate is also affected by behavior. And R0 is the mean of a random variable, so not fixed.

How to look at this as a research issue: intersection of traditional engineering and social science; MIT ESD way, which is compatible with the NSF suggestion. We need an integrated analysis that will require an interdisciplinary team, working on a prairie, not people in canyons or silos. The results could be new fundamental research findings and/or a set of feasible imaginative policy alternatives that are implementable and satisfy the often conflicting goals of multiple stakeholders. But real problems don't self-organize by discipline.

Session 8 Questions

Q1: How does one tear down the walls? A: In universities: ESD at MIT is an example of that – large interdisciplinary research labs.

Q2: How do you protect assistant professors who don't have tenure? A: I think if you let them work on a prairie rather than in a canyon, but doctoral students are very risk averse. Q2: And part of the process protecting junior faculty requires senior faculty writing letters in tenure cases. So it's a chicken and egg issue. A: I think we're talking about a generational change. Universities change very slowly and the only way rapid change will happen is if it is forced by the outside. We should also discuss this with NSF.

Q3: Please comment on research conservatism, or focusing where the money is vs. where the excitement lies? A: The NSF team and this workshop is a wonderful counterexample; they're looking for out-of-the-box ideas. But being in a silo is part of a vicious cycle, and NSF now only funds 8c/dollar requested.

Q4: So NSF is a service system? A: Well it's not agriculture, mining, or manufacturing….

Session 9: The Story of Online Learning at MIT

Sanjay Sarma, MIT: MIT has a long history of educational innovation. 30 years of Project Athena, or technology-enhanced active learning and Open Courseware since the late 90s when MIT decided to give its curriculum away. But MIT got a little smug and did not actively pursue MOOCs until Udacity and Coursera out of Stanford were founded as for-profit companies around 2011. A MOOC is a video plus an assessment. What is new is that meaningful assessments, like differential equations, essays, circuits, can now be graded automatically, enabling thousands of students to take a course.

In late 2011 MIT's president and provost announced they would teach its first MOOC, 6.002X (circuits & systems) in the spring of 2012, even though at the time they had nothing prepared. They succeeded with the help of many professors, postdocs and students at CSAIL. They expected 10,000 students. In fact, 155,000 students started, from 196 countries; 8,000 finished (typically only 5-10% of students finish). Now more than 350 courses are offered.

Cognitive psychology and the neuroscience of learning: research shows the brain's capacity is about 1o minutes of content before you need to DO something with it to encode it into long-term memory. Therefore an apprenticeship model works well: small amount of instruction and learning by doing. A 90 minute lecture: not so well. In practice 10-minute, in-person lectures are impractical. But it can be done in online courses. It leads to 'flipped' learning or 'blended learning': lectures on your own and class time for going over problems. But it's not a fad; over 60 MIT courses are using the EdX platform on campus. We expect in 5 or 10 years on campus learning will fundamentally change: fewer lectures, more hands-on work.

MIT has a long tradition of *mens et manus*, doing things, having nontraditional students, working hands on. We are now pushing lectures back and reclaiming the lab work.

Session 10: NSF Vision for Smart Systems Research

Dick Larson: summary of day 1 and preview of day 2; we want a combination of ingredients: rigorous knowledge, peer-reviewed journal contributions and interdisciplinary or transdisciplinary work.

Pramad Khargonekar, NSF: My job at NSF is very simple: make wisest possible investment of tax dollars for societal benefits by investing in engineering research, education and innovation. This area might be great for that. Services are now 80% of GDP; manufacturing is now less than 11%. We will need high-quality, low-cost personalized solutions in many areas: education, health care, transportation, agriculture etc.

Economic growth is mainly a function of number of people and productivity. Productivity is the bigger part. Up to about 1970 labor productivity growth was high but it's now down to about 1%; it's negative in health care. Though reality is of course more complex. To maintain growth service productivity must increase much more than that.

Brynjolffson & McAfee's Second Machine Age shows that productivity no longer depends on mechanical or muscle power but on brain power, which will be transformed by things like machine learning, robotics, smart systems, etc. The combination of people and technology that makes smart service systems. NSF is doing a lot in advanced manufacturing, and there's a blurring of the lines between service and manufacturing that drives the internet of things, 3D printing, maker spaces, etc. We also see that retail is benefiting from engineering, e.g., Amazon has changed retail. Less clear but equally important is the challenge of the emerging cyber-physical-social future, in which our work, life, homes, etc. are infused with sensors, control systems, computation, networks, etc. We can build on that to add value to service systems.

So the research challenges include increased complexity, greater range of scales (global connectivity to personalized solutions) also faster time scales to (near) real time. The systems need to learn, from experience, from data; and they need to be human-centered.

So what fundamental knowledge is needed? What research communities need to be at the table to create that knowledge? Analogy: building machines in 1801 before the advent of mechanical engineering as a discipline. We don't know enough about service systems yet even as we're building them; but we're doing it because there's economic opportunity to do so. But without foundational knowledge we can't do it systematically. But what might we learn from experiences in development of new engineering fields? Two examples: materials science, and biological engineering, especially the latter.

So the key questions are, first, models. What might constitute models as we think about engineering of service systems? What math is appropriate? What is the role of social and behavioral sciences in these models? Second, what are the right analysis and design tools? How do we leverage advance in cyber-physical systems and machine learning? Cyber security is important. Third, and harder, how do we put people in the center of framing design questions? This goes towards economics, social sciences, organization security, regulatory and policy issues.

What can we learn from recent experience? In education teacher productivity has been static for many years. But knowledge in the science of learning and education is growing, in part due to NSF funding. We now have more tools also. What we need is learning engineers who can transform education delivery. Similarly in health, the 2014 PCAST report is in part a call for how to use systems engineering ideas to impact the health business – can one imagine health engineers. So the challenge would be a smart health care service system, designed in a way that can lead to a much faster learning cycle in the overall healthcare system, with obvious impacts. We fund a number of relevant programs for that. We want something that builds on those investments as a new way to create this foundational knowledge to create long-term impact.

Key issue: 47% of US employment is at risk for being automated away. (Brynjolffson & McAfee). We need to make sure we put humans at the center of the design question.

Key ideas:

- Is there a real opportunity to create learning engineering, health engineering, or others? Need to be creative but rigorous.
- Is the new foundational knowledge we create trusted and usable to create the value added, as per Charles Vest's idea of focusing on what is really needed and scholarship will follow?
- We value your ideas, research agendas, etc. We depend on the community to come up with the compelling ideas we can invest in.

Session 10 Q&A

Q1: How do we train students to do this kind of research? A: The evolving NSF vision of foundational knowledge is to meet the goal of capturing the knowledge in ways that can be taught. It's engineering AND social and behavioral sciences. So we need to build the bridges from engineers to people in social sciences, and also domain experts. Then capture the knowledge in a form that can be taught – models are a language that needs some translation. It's a hard question. We at NSF want to hear from you if there is potential there.

Q2: How do we ensure depth in this multi-disciplinary work? A: There is compelling need from society. But is there an intellectual case? Yes in my opinion, but we don't think enough about it. We need to think about the metaphor of knowledge as a network; control theory is a node, computer science is a node, etc. we get PhD's in nodes but need to pay more attention to the connections, build the fat pipes to connect the nodes. It's not bad between control theory and math, but others need to be built. That's the intellectual case: try to organize the knowledge that connects the domains in systematic and rigorous ways. We then would solve depth and rigor issues.

Q3: How can we better use gaming platforms in education? A: NSF currently funds that; we want to go further: understand how learning improves through gamin; systematize that knowledge about learning, gaming and human behavior.

Q4: We are trying to make design into an evidence-based science but it's hard to get people at NSF to take us seriously. Issue of silos; it's hard in today's funding world. A: Excellent point. NSF is us, this community. We pass judgment as panelists. So we need to work to articulate a vision, get it shared so panelists understand. NSF needs your help.

Session 11: Augmented Cognition in Service Systems

A: Dylan Schmorrow, SOAR Technologies: Developing a rigorous model of cognition and applying it to augment human cognition: "the next brain". The goal is to apply artificial intelligence (AI) technologies to make people more prepared, informed and capable. "Computational partners". Early example: making people part of the avionics platform, rather than just focusing on the human-machine interface. This may involve gauges that report if a person is paying attention, sense memory, executive function, sensory input. Cognitive architectures that go beyond Bayesian or neural networks. These can help in creating an interactive task learning type of system. [See video, which is still futuristic but gives an idea of the concept.]

Even though we may not understand the fundamental neuroscience mechanisms in the brain, we can do interesting things with the information we do have. AI systems will prove useful long before they become self-aware. Today's issue is how to deal with all of these overwhelming interfaces, technology, information. How do we tie it together? We have some of the keys.

Some involves looking inside the human brain. [See slides, which have more detailed definitions of new terms such as adaptive automation, brain-computer interface, and cognitive neuroscience.] A brain-computer interface is not for controlling something left to right, but help facilitate cognitive decision-making in individuals and teams.

The video highlights some emerging capabilities relevant to smart systems, including 'cognitive prosthetics' that can help mitigate cognitive bottlenecks – error detection, arousal, memory encoding, working memory, different verbal and spatial attention measures that can be fed into a support system. For example computers can tell whether you're switching between verbal and visual focus in less than two seconds, so can detect events such as mistakes even if you don't realize you've made a mistake. Using this support for DARPA in an in-flight missile retargeting task, which is mainly using working memory, there is a 642% improvement.

Key improvements were in task execution and mitigation strategies in real time with someone working in the field. We found the best strategy for mitigation was modality switching and scheduling.

Key point: yes you can detect cognitive state in humans, but the important thing is what you do about it. Our systems can anticipate and simulate how people are doing. Cognitive architectures integrate perception motor control, relational representation of situation and goals; key items that allow systems to learn. It is convergence of multiple domains focused on one outcome: natural human-technology interactions. It's not just big data, but this is only a means to an end. We need truly interactive symbiosis. What needs to be done now: understand the human element in social, cultural, technical systems; understand how to deal with unbounded uncertainties and optimize outcomes for the humans and the systems. Bottom line: how can smart systems help us reap the rewards of the digital age without overwhelming us with complex systems and impossible interfaces?

Session 11 Q&A

Q1: [Oli de Weck]: Optimizing the level of assistance provided to humans seems very important. Full automation might have prevented the Air France 447 accident; the other extreme is just people with paper and pencils. Should teenagers have spell checkers? What is the right mix? A: We really do AI, not full automation. Support humans to make them more prepared, capable, able to execute more appropriately. The underlying principle is to provide support. Initially we thought that would be for crises, but a lot of it is when you're bored because there is nothing to do: how does the system keep you engaged? In such cases the system might have you do a task it might otherwise do. Note however there is an ethics issue in systems knowing so much about me that could be used for good or ill.

Q2: [Jim Spohrer]: What is real today? And where is SOAR Technology? A: SOAR is a wonderful AI company, but there is a new nonprofit called Brain Hackers, with Chris Forsythe who is about to retire from Sandia Labs. 8^{th}- and 9^{th}-graders using these gauges for interesting projects. You can start hacking away today. Please contact me on LinkedIn or email. Slides are on Slideshare.

Q3: The military scenario in the video was interesting but may not scale. As a defense contractor we see size, weight, power, bandwidth constraints. How do you overcome bandwidth and cost constraints to viability? A: Two answers, DOD and commercial. For DOD there is a lot of acquisition already for medical sensing and computational stuff that we can leverage to make it more affordable, so we did that. Commercially we could leverage things like Android watches, health applications – some of this might be adequate; we have to be careful because there is a lot of noise. But a lot of work being done in health and fitness could apply.

Session 12: Exponential Forces of Technological Change: Moore's Law and More Laws

Jessica Trancik, MIT: Informed innovation: how do you accelerate technologies while conscious of factors like environmental impact, sustainability, and cost? Main example: energy systems.

My research group focuses on harnessing information in large data sets to try to understand the reasons behind technological improvement: the impacts of technology to inform design in the lab, and investment. We have learned important general lessons. One measure of technological change is prevalence: growth over time. In transistors it's a doubling of production every 1.2 years (log scale). Another is technological performance improvement, an example of Moore's law applied to the cost decline in transistors shows a halving time in cost of 1.4 years. This is quite fast. For IT overall improvement is 20-30%/year. The assumption was that only IT had exponential growth but we tested it in our paper for other technologies also. [Paper is online]. We collected data on 100 technologies, which was hard. We used hind casting to test six+ theories, including Moore's Law. Another eight theories performed very poorly. Hind casting provided us 40,000 error data points and used a statistical model to interpret it and pull out effects due to the data sets (treated as random) as well as the effects due to functional form of the equations. We found Moore's law does a good prediction job across technologies: exponential improvement is ubiquitous. We can use Moore's law for forecasting also.

We have looked at mechanistic models to explain why Moore's law is so prevalent. It matters for example in international climate negotiations: some of these technologies are improving fast so the situation may not be so bleak. See Duncan Watts: Everything is Obvious (once you know the right answer). We looked at design features in a design structure matrix that help us predict the rate of change. Other features also matter, like the degree that a technology is commodity-like, because more commodity like technologies can slow improvement. So IT's 20-30% improvement rate is quite high; other technologies like energy, are much slower, but still exponential. Photovoltaics is pretty fast, at 10%. Our large data sets help us predict both expected behavior and expected error, which is useful.

We live in a log world. Exponential improvement is all over the place. This has important implications for the structure of the economy, jobs, and nature of jobs. No single discipline can address this alone.

Key questions raised for me: can we plan for technological change? Can we design for system performance? Can we induce technological change to reduce environmental impact? While we're not social planners but it is important to inform these broader debates with rigorous and open-minded research.

Session 12 Q&A

Q1: How do you use your models to forecast or understand service productivity? A: The kind of model you're talking about is Wright's law, which is consistent with Moore's law if the number of improvement attempts is growing exponentially with time. For human systems that is tricky since the number of people isn't growing that quickly. Se we would see, and people have seen, slower rates of improvement. I think humans working with technology is the basis for exponential economic growth and productivity growth. But it is important to understand the difference in functional form between technologies and people.

Q2: First, what's on a log curve that would surprise us? And is there something NOT on a log curve, and if so would that change everything? A: on the first we're looking at things in terms of two common components: usage rates (technology costs) and input prices. This helps us explain many of the differences in rates of improvement, because usage rates are changing over time. Second, what would change the world? Storage technologies for energy. Cheap storage would change not just electricity, but transportation because both photovoltaics and wind are cheap enough to be competitive at small scales, except that they're intermittent. It might also reduce the need to carry batteries around.

Q3: You indicated exponential growth provides many opportunities. Does it also bring exponential problems and failures? As you miniaturize technologies, new forms of failure can occur. Devices capable of monitoring humans can fail. Are these new vulnerabilities? A: I'm sure there are, I don't work on that question. But our dependence on cell phones comes to mind; that is rather trivial. But aging may be a less trivial example. As we move to more advanced technologies some could be left behind. Some other issues: growth leads to increased carbon emissions, depending on your risk tolerance, many of us are concerned about climate risks, so that pushes us to trying to guide innovation to bring down performance metrics like carbon intensity of energy, emissions per unit energy.

Q4: [Dan Roos]: What about differences between new technology like Tesla cars vs. incremental change. A: so the price usage model shows that at a given cost, usage ratios determine performance as well. But I think that by starting with designing the entire system, there are opportunities to design such a system to improve quickly, and I think we're seeing that with Tesla. But Tesla has challenges with integrating into the existing system as well. Another concern is use of more exotic and expensive materials for say batteries. So how do we move to more common materials? I see it as a different kind of engineering, you're trying to do something you know how to do, but more cheaply. It's an interesting challenge.

Session 13: Grand Challenges for Social Innovation

Maja Mataric, University of Southern California: Socially assistive robots (SAR) can be used to take care of people in the health space. Changing behavior is traditionally very challenging and requires intrinsic motivation, appropriate environment, and social support. Robots could be used to provide:

- Monitoring
- Coaching/training
- Motivation
- Companionship / personalization

Example 1: With stroke recovery patients can robots be used to motivate people to do exercises to help recover limb function post stroke? Showed video of patient interacting with, and being motivated by, robot while doing stroke recovery exercises. Patient was quite engaged by robot and incentivized to do extra exercise to "please" robot. But questions still remain - will the effect last, and how do we do this ethically? Major knowledge gap: sufficient understanding of human behavior to design successful robots for such purposes.

Example 2: For autistic children can robots or computers be used to draw them out of their shell and stimulate social interaction? Showed video of autistic child who had previously been labeled as non verbal. They were able to elicit latent behaviors that had not previously been seen, including verbal interactions. Very compelling. With time, however, the child was less will to interact diminished to where the child was again essentially unresponsive.

The challenge is to identify methods to achieve persistent engagement of autistic child. Obviously, personality is key in human interactions. Therefore, one solution is to invest robots with back stories that will engage the children [she noted that there are lots of unemployed screenwriters in LA]. One important benefit of this research is that by studying interactions of humans with robots, we learn about how humans think, behave, and respond.

A: Kent Larson, MIT: MIT Media Lab has 22 groups in which group leaders can do anything they want. Ranges from lifelong kindergarten, biomechatronics, Opera of the Future, etc. Essential to have no fear of failure and be able to iterate very quickly. People come from a wide variety of backgrounds; not multi-disciplinary but trans-disciplinary (designer who can program, tech people who are sympathetic to good design, etc.) My job is to break down walls between disciplines that people come in with and show them that programming and data visualization involve design, thinking about people, and sociology. It's not pretty…

B: Don Norman, UCSD Design Lab: Focus on automation, but there's a problem when engineers automate what they can and leave the rest to humans, but this is bad because then the humans' skills deteriorate, and they are surprised when the unexpected problem arises (which it always does). Better: Think of a person as a coroutine: use collaboration and cooperation rather than automation. Take advantage of human capabilities AND machine capabilities, e.g., humans have a 5-10 second attention span; humans are always scanning and changing; alert. When you force people to act like machines they don't do it well. You get accidents, attributed to human error.

We shouldn't focus on disciplines but on problems and get people with the needed skills to work together and create solutions. We need to understand people since we build things for people. People should be creative, do strategic planning, change their mind, and be sensitive to environment. This is a key challenge for service systems: called distributed cognition at UCSD but a better name would be collaborative human automation systems. The big challenge for systems science is figuring out how people distributed in time and space in various organizations will work with machines of varying abilities in a collaborative way. Example: Deep Blue can beat the best human chess player, but not the best TEAM of chess players, which is a combination of humans and machines.

C: Calestous Juma, MIT: Application of science and technology to sustainable development. Example of Dr. Trancik's discussion earlier of exponential growth in technological capabilities: the first mobile phone (1956) weighed 42 kilos. Today we have 'mobile phone speciation': with modest technological modifications a phone is a banking device, and can start to create new industries such as money transfer in Africa; now $18B, expected to be $300B by 2018; growing at 82%/year. Money is transferred in the form of buying and transferring air time (minutes). Needed a few technical modifications to make it secure and transparent; everything else already existed. Question: are there other areas ripe for similar speciation? The challenge is social, not technical. One new application, coming out of the Ebola crisis is contact tracing.

In the process the government leaders in Lagos realized they were not well-informed of the issues or technology potential. They started by appointing a 'rumor manager' to manage the misinformation being spread but later realized they needed a Chief Scientific Advisor to help them become informed and make better decisions. The bigger issue is how to help governments leverage information and make better decisions; how do you design new advisor methods that work quickly; not spending 18 months designing a study on Ebola.

One way is to strengthen African science and technology academies to provide that kind of information to governments. It's not about large scale technology but about repurposing technology, and creating collaboration between those who know technology and those who understand the local social problems.

Q1: (to Calestous Juma): I worked in Africa for a few years and have experience with project failure due to rumor mongering. Please tell us more. A: I created a think tank in Kenya after finishing my PhD and a senior government official told me policy is what you make in the absence of knowledge. But things have changed and more African leaders have an analysis background; in 2012 six African countries elected engineers for president. And 15 African countries now have scientific academies that work closely with NAS. But most leaders don't yet have a listening ear, a Chief Scientific Advisor, so I'm working with them to get that. The academy I'm involved in creating is in agriculture; the request came from African presidents, not the scientists.

Session 14: Transdisciplinary Fundamental Research Problems: Rigor and Convergence of Cognitive/Behavioral Science, Computer Science and Engineering

A: Dimitris Bertsimas, MIT: Example of developing a personalized, comprehensive and dynamic system for diabetes and diet management (with Alison O'Hair who did her PhD in this area). Typically a pre-diabetic gets the non-actionable advice of eat less carbohydrates and exercise more. So we tried to develop a system that learns your food preferences, exercise preferences, metabolic situation – these vary from person to person. This required a combination of operations research expertise, nutrition science and endocrinology, computer science to implement it in code, and then we realized we needed to understand economic theory – specifically Kahneman and Tversky's work on how people decide. – This was the biggest challenge. We also worked with a cartoonist and psychologists to understand how to get people motivated to adhere to the suggestions. The system is called LiA – Lifestyle Analytics [demo here]. It asks for your personal characteristics, height weight, objectives for glucose, weight etc. It asks about your desired food variety, budget, how much time you have to prepare meals and lifestyle preferences. Also restrictions, allergies, etc. And we learn your preferences through A/B choices – usually in pictures. And we found it takes about 10 minutes to understand people's food preferences. Then we have a database with 700 recipes, that a culturally appropriate.

So you enter foods that you like and the system optimizes the rest. We found that we needed about 10-15 measurements over 2 months to make it very personalized. It takes about 2 minutes to update to the cloud but people felt 2 minutes was too long. So we added useful information for clients to see during that information and people enjoyed that and didn't mind the wait. So our model provides you with a daily plan and a prediction about your glucose levels – our model tries to keep you under 180 mg/dl all day. We have several hundred users now, but only 50% use it regularly – that is relatively high. We are now working on an iPhone and Android application and including GPS so we can find restaurants nearby and propose foods appropriate for a diabetic from the restaurants.

We also create grocery lists etc. All this is based on us observing people and finding appropriate disciplines to solve the problems. The cartoonist matters because experiments show that if a cartoon character gives advice it's more likely to be followed.

B: Nalini Vekatasubramanian, UC Irvine: We are looking at service-oriented architectures for smart cities, especially in emergency response. We work with FEMA and DHS looking at how IT can help in emergency response processes. It's a merger between technology and human processes. Cities are interesting because they have many stakeholders who require services (functional requirements) such as roads, energy, buildings, sewers. And nonfunctional requirements such as efficient, dependable, available, flexible. There's a transition from infrastructure-centric to information-centric (nonfunctional requirements). There's technology for that: Internet of Things, developments in sensing and control to support that. We need to ask can the city become sustainable? Resilient, scalable, secure, private? What are the systemic vulnerabilities? Are there new interactions? I'll focus on resilience, which is key to disaster recovery. It involves preparation and adaptation. Challenges include scale and density in cities; but these are also opportunities. Another challenge is aging infrastructures. So the notion is to recover forward, not recover to the prior state: can I bring in new capacity to improve things vs. before the disaster? It's key to deal with service continuity and recovery after disaster. New design principles are needed, not one-off solutions.

These principles include ways to structure the representation of the entities: observe, analyze, adapt, cycle.

Principles of computational reflection: methods to observe the evolution of a system represent the state of that evolution. How to determine how to adapt and improve the system its performance? These systems are by definition cross-layer but we need to develop lightweight rigorous methods to understand them. We have been testing this observe-analyze-loop at a campus level for fire awareness.

C: David Woods, Ohio State: Systems don't always do what its designer intended. For example, failure of IT systems in hospital pharmacy leads to medication errors for 36 hours. Or a software system in Dublin's air traffic control tower fails. Fortunately in both cases the people could shift back to previous processes and disasters were averted. This is a subset of complications, interactions and adaptations that are common. Failure is due to system brittleness. Why don't we have more failure? Because people adapt at the boundaries; they are unvalued sources of resilience. It's the fluency law: well adapted work hides the difficulties handled and the dilemmas are resolved so it's not seen by others (e.g., planners and developers). Those people could respond due to practice and they took initiative. This is a critical property of the Adaptive Universe, which doesn't work the way we think, the way we model it. Our systems are poorly calibrated and more brittle than we think since we don't model the adaptive humans.

Rigor comes via after-action reports, we study adaptive cycles, uncover the hidden sources of adaptation. We contrast case studies of success and failure; we have too much data on failures and kluges and workarounds and not enough success. Our unit of inquiry is the adaptive cycle. It's messy. Triggering events initiate a series of adaptations across different units, roles, levels and organizations. These changes often stall out and create too much extra workload – e.g., the 60-year history of failure in EMR is due to Grudin's law.

Grudin's law means you must think about who pays the workload costs and risks, and who receives the workload benefits? It means new computer systems are unlikely to be accepted or used as designed when those who pay the workload cost don't receive the benefits from the technology in their role. EMR is designed to enhance billing not clinical work; so clinicians don't see the usefulness and don't want to enter the extra data.

What we need to study is examples of fluorescence where change triggers adaptations that open new niches that other roles can take advantage of. What produces those? Graceful extensibility; similar to graceful degradation; it's a form of resilience. We use multiple observer teams who go out for months in the field not days in the lab.

D: Justine Cassell, CMU: We need to ensure there is a merger between technology and human processes and that we use an adaptive cycle. People often assume there is an intrinsic and unbreakable link between signals of emotion (smile, frown) and underlying psychological states. But there isn't. There are inferential links. We build models of the links between the looks on people's faces, their body movements and their putative underlying psychological states and we build 'virtual humans' and agents.

We know students learn better when they feel the teacher respects them. When people feel engaged they are more honest in census interviews, etc. It's a full loop but you can only get it by careful observation. For example in observing 14-year-old peers tutoring each other we found the number of insults they trade predicts their learning. But people have been building tutoring systems that are more polite. We find negative self-disclosure is necessary, so we build new kinds of architectures for dialog systems for virtual humans. We asses goodness of fit through a 'wizard of Oz' technology and then we build an implementation.

Q1: (To Prof. Woods): instead of assuming zero sum on who does the work and who gets the benefit (Grudin's law) why not look at total benefit and service-coproduction? [Prof Woods]: True, but we can go further because you can see adaptation grow across the network. For example we had a rural electronic ICU in which experienced nurses supported multiple less experienced people. Then we found others in the hospital wanted to take advantage of it as well. [Prof Venkatasubramanian]: the same phenomenon can be seen in computer aided dispatch centers and mobile crowd-sensing technologies: if that info could get back to the operator it doesn't help the provider except indirectly. [Prof Woods]: when you take the adaptive perspective there may be unintended impacts – fluorescence as well as conflicts because humans have multiple possibly conflicting goals. We'll see that as our design starts to trigger new layers of adaptation from the different stakeholders.

Q2: On insults… [Prof Cassell] ask your undergrads, you'll find insulting mitigates power differentials and that perhaps allows you to learn from each other.

Q3: In programming, we are finding that the program design paradigm has overly optimistic assumptions, e.g., data will never get lost, power is always up, makes the system brittle. It's a miracle we haven't had more disasters. So the paradigm is the fault, so modelers, programmers, paradigm creators have to think twice and make their assumptions explicit and then the rest becomes easier. [Prof Woods]: Yes, we should think bigger about the questions we are posing. Brittleness is one example; it's not very tractable so we don't focus on it and it's hard to fund. Databases are another example; schemas are hard to change but business changes all the time.

Q4: How have you been able to reconcile the differences in what is considered research rigor between social science and engineering? [Prof Cassell]: I think we're now in a liminal space between disciplinarity and no disciplines. I teach my students both how to make an argument for a social science audience and for an engineering/CS audience because I think that makes them stronger. It's like being bilingual. [Prof Venkatasubramanian]: Some of that comes from putting students in the field and helping them learn how to populate Cs/engineering tools with realistic meaningful assumptions that can be applied in practices – i.e. policies? We need to also start with assuming the base infrastructure is prone to change. We need to observe the change and design systems that can add new mechanisms: extensibility. [Prof Woods]: good work is good work and it starts with the phenomenon. And then innovation in method. Our problem is putting tractability too early in the exploration and discovery process, putting verification ahead of discovery. What we don't value we don't fund.

Team Presentations

The goal for this session was to explore the potential for fundamental interdisciplinary research in the realm of smart service systems. Nine teams chose a topic and then presented their ideas to the larger group. Topics and teams were as follows:

- resilience (teams 1 and 9)
- transportation (teams 2 and 4)
- aging (teams 5 and 7)
- energy (team 3)
- health (team 6)
- general approaches (team 8)

Team 1: Resilience

Questions:

- Can we collect real-time data (potentially from crowdsourcing) to develop a standard methodology to conduct postmortem data collection and analysis to result in agents of change in service system design?

How do we:

- halt propagation of the hazard impact or cascading failure?
- design routine infrastructure design, maintenance, and operations to be more resilient when a hazard occurs?
- promote or design a network of distributed micro service systems to increase the robustness and resiliency of a macro service system?
- categorize and systematize the complex interaction between multiple disciplines that impact resilient infrastructure?

- understand and plan to deal with the merging hazard response with sparse and potentially conflicting data?

Key issues:

- Understanding interdependencies and prioritization of human service infrastructure systems
- System-wide uncertainty in terms of performance
    - Level of use and human interaction
    - Ability to predict infrastructure system response
    - Knowledge and probabilistic assessment of system performance
    - System performance metrics
    - Minimum requirement of a system resilience and level of acceptable risk

Solutions/Goals:

- Minimize recovery time and cost
- Increase resiliency and readiness with respect to hazards
- Assess system-wide risk
- Quantify the productive role of crowd sourcing citizen science and hazard readiness and recovery
- Assess the value of incorporating product design methodologies to system infrastructure design

Team 2: Transportation

Questions:

- How will self-driving cars change urban transportation landscape?
- Will we need to buy these cars or will they be available? (Isn't this how Uber works now?)
- What does the future of smart urban transportation look like; how do we get there; and what do we need to know to get there?

Key issues:

- What would push us from "ownership" to "shared use"
    - Reliable transport from here to there?
    - How does this relate to other (long distance) transport systems?
    - What do people value when investing in cars?
- Urban area point of view
    - Given fewer cards, how can we design the city?
    - Policy decisions
    - Business decisions
- Stakeholders
    - Urban planner, highway engineer, car manufacturer/dealer, transit services,transportation providers, businesses, individuals
    - Information tracking, privacy, security

Problems:

- How do we combine multiple layers, multiple kinds of models into a single model?

Solutions:

- Multidisciplinary, team-based approaches
- Grand, unified modeling approach/formalisms

Team 3: Energy

Questions:

- How do we make solar energy economical?

Problems:

- Cost (rate of improvement for PV)
- Stability of the power grid
- Solutions:
- Source intermittence, storage, and load can be managed simultaneously to mitigate their impact on the power grid
- Short-term solar forecasting
- The post disciplinary device
- Behavioral change
- Value of co-production

Team 4: Transportation

Questions/Problems:

- What is a smart service? Do we think about it in short or long term scale?
- How do we make travel more efficient?
- How do we make traveling more comfortable/enjoyable?

Solutions/Goals:

- Systems across different time scales and looking at comfortable travel, minimizing travel time/congestion
- Develop a system that will reduce car ownership
- Use big data analytics to understand traffic patterns, travelers' behavior for short time decision and longer time scale planning
- Influence behavior using different policies, incentives (taxes, rewards)
- Looking at social and technical models, how do we integrate them and understand their interdependencies?

Team 5: Aging

Questions:

- What does mobility mean to the aging?
    - Psychological, physical, social meaning
    - Freedom of movement, movement without fear
    - Maintenance of social connections and of connection to city, public spaces, community
    - New explorations of the world during retirement
    - Diminish dependence

Solutions/Goals:

- Understand value of technology to needs of the aging, across the aging process
    - Don't try to solve all problems with one technology
    - Minimize risk
    - Allow scalability
    - Extend mobility in physical and social spaces

    - Design and iterate on design of technologies in urban, rural, and suburban communities
- How can technology be deployed to enable and extend mobility from healthy elders to disabled elderly?
    - How can the technology itself be mobile?

Team 6: Obesity

Questions:

- How do we counter obesity and associated diseases?
- How do we advance health informatics?
    - Prevention
    - Detection
    - Mitigation
    - Motivating (creating conditions)
    - Intervention

Key issues:

- Create a model of the service systems
- Identify interventions that can change behaviors
- Identify stakeholders and motivations to offer service
- Identify use cases and technologies that change behaviors
- Metrics that measure impact (polling/sampling)

Solutions:

- Design community for walking/exercise as a social activity
- Options extend to all (income levels)
- Incentives for healthy foods
    - Rebates for fresh and local foods
    - Food stamp incentives
- Incentives for exercise
    - Bus with treadmills
    - Multi-purpose spaces

Team 7: Aging

Questions:

- How do we help diversify elderly populations be safe and live independently (as desired)?
- Extend their working capability and allow them to contribute to the overall economy for longer
- Improve quality of life through interventions that allow them to function to the best of their abilities

Problems:

- Imbalance in the labor market
- User design (user experience studies needed)
- Emergency logistics driven by resource constraints
- Isolation (dispersion of family members who are potential caregivers)
- Handling preventable accidents and unnecessary hospitalizations

- Declining cognition

Solutions:

- Design of wearables
- New infrastructure systems/virtual reality
- Agent modeling and sensing
- Human factors
    - Prognostic → prevention technologies
    - Mobility needs for aging (transportation)

Team 8: General Approaches

Questions:

- Do we want to break the rules? How should we break the rules?
- Solicit diverse approaches focus on containing capability to adapt
- Try out promising candidate on two themes

Process:

- Value proposition (mission)
- Roles and relationship (coordination capability)
- Timing: pacing, tempo, synchronization
- Respect the trade-offs
- R&D finds:
    - Fundamental questions
    - Empirical scope
    - Modeling targets
    - Informs theory
    - Predicts economic impact

Team 9: Resilience

Questions:

- What hazards are there? (natural, terrorist, manmade)
- How do we define resilience? (build better, improve responsiveness)

Problems:

- Logistics (food, water, shelter, medical supplies to people in need)
- Time and spatial scalability
- Information extraction/data fusion from different sources and control
- Content design evolves with scale and interaction

Solutions/Goals:

- Need actionable information in as many hands as possible to increase resilience
- Need to get the right content to the right people at the right time in the right space
- Integration of knowledge STEM + human behavior