SYSTEMS ENGINEERING RESEARCH CENTER

SERC TALKS

WELCOME

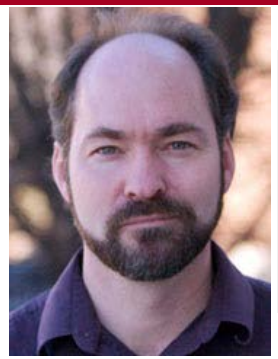

What Are Cyber-Social Learning Systems And How Will We Form Them?

Dr. Kevin Sullivan, University of Virginia

June 7, 2017 | 1:00 pm ET

- ❑ Today's session will be recorded.
- ❑ An archive of today's talk will be available at: www.sercuarc.org/serc-talks/
- ❑ Use the Q&A box to queue questions, reserving the chat box for comments, and questions will be answered during the last 5-10 minutes of the session.
- ❑ If you are connected via the dial-in information only, please email questions or comments to Ms. Mimi Marcus at mmarcus@stevens.edu.
- ❑ Any issues? Use the chat feature for any technical difficulties or other comments, or email Ms. Mimi Marcus at mmarcus@stevens.edu.

The Systems Engineering Research Center (SERC) is a federally funded University Affiliated Research Center managed by Stevens Institute of Technology.

Any views, opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Department of Defense, ASD(R&E), nor the SERC.

This material has been approved for public release and unlimited distribution.

Credit

- Credit for *many* of the ideas in this talk goes to the participants in the workshops on cyber-social learning systems that I co-organized with Chuck Friedman of the University of Michigan
- Sponsored by the Computing Community Consortium (advisory to and funded by the National Science Foundation to work with the research community to develop audacious new research ideas)
- We had nearly 100 participants in the series of three workshops: from computer science, systems engineering, the social sciences, and from healthcare, education, and smart cities/communities
- This talk is based on a preliminary, *personal* draft of a report. An official consensus workshop report is forthcoming. Blame for any shortcomings in this version rest with me alone.

The future of systems

- We envision a class of radically new socio-technical systems
- We are calling them cyber-social learning systems (CSLS)
  - —Human-intensive, socially complex, often ultra-large-scale systems
  - —Providing the most critical services across many sectors of our society
  - —Deeply infused with computing technology, information processes
    - Sensing, big data, machine learning/AI, distributed, cloud, mobile, decentralized
  - —Reflecting, projecting, predicting, to *learn, adapt,* and achieve high fitness
  - —Spanning many scales: individual, dyad, group, enterprise, sector, society
- In defense, healthcare, education/learning, government, city and community services, policing, criminal justice, transportation, etc.
- *The future of systems is one of cyber-social learning systems*

Example: Learning *Health* Systems

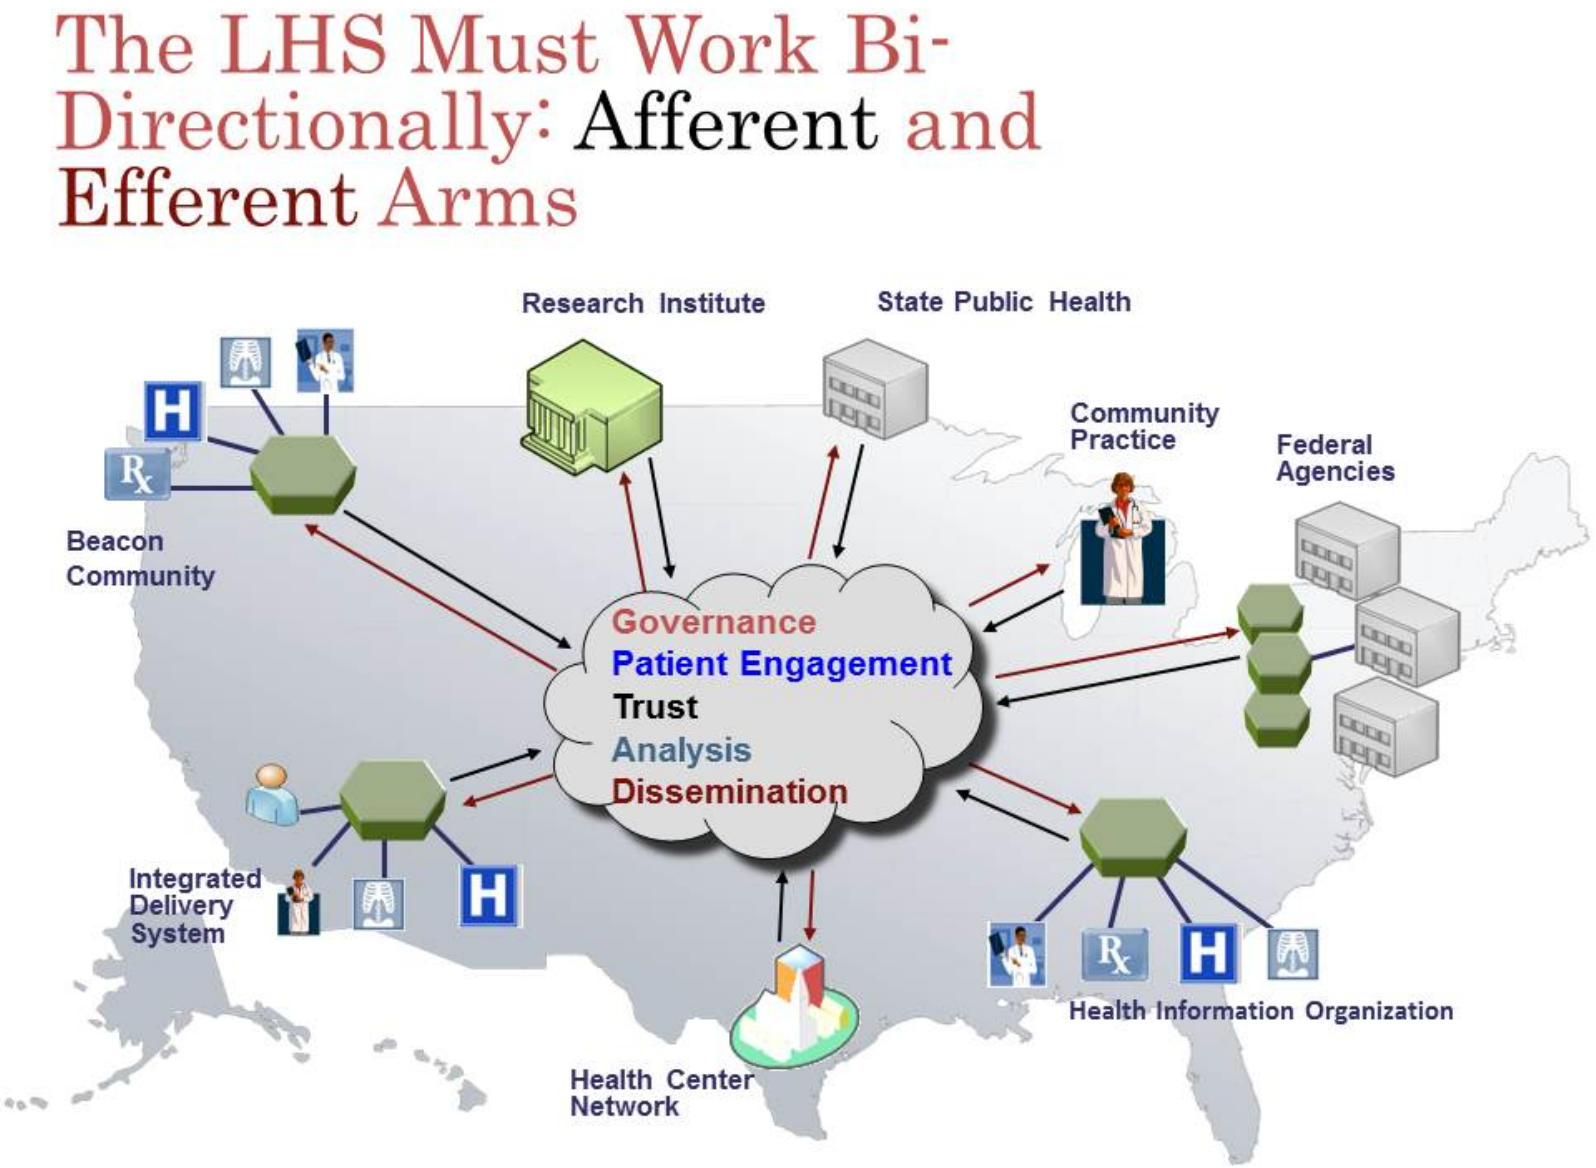

SYSTEMS ENGINEERING RESEARCH CENTER

The Reality Today

- Notwithstanding promising islands of excellence …
- We're not on a path to building such systems today
- We're not even well configured yet to get ourselves on this path
- It's vitally important that we figure out how to get on that path
- Because the problems we face can no longer wait for solutions

SYSTEMS ENGINEERING RESEARCH CENTER

The Nature of the Problem

- Consider two vastly different cultures
- The new-world culture of computing system
- The old-world culture of service systems

Two Cultures

- Culture of New World Computing Systems
  - Long-term exponential improvement in compute, store, sense, connect
  - A culture of plenty: speed, customizability, personalization, connectedness
  - Revolutionized information-intensive systems (bits: finance, music, etc)
  - Revolutionizing physical systems (atoms: CPS, robotics, IOT)
  - Amazing new functionality: assistants, collaborators, coaches, mediators
  - Is now running headlong into human and social systems and phenomena

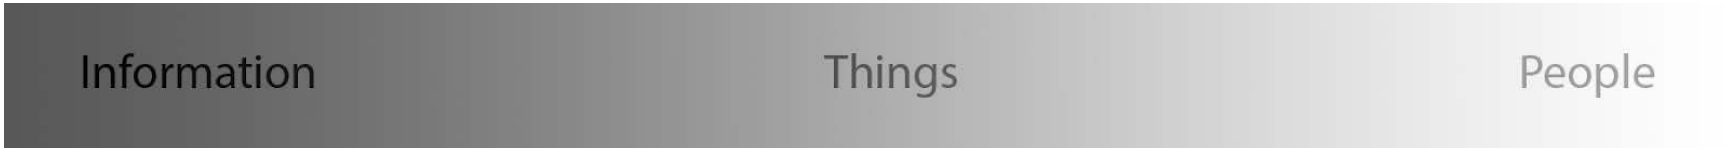

Two Cultures

- Culture of Old World Service Systems
  - —Healthcare, education, government (incl. procurement), etc.
  - —Substantially lacking meaningful, integrated, and deeply analyzed data
  - —Substantially lacking in effective, modern data & computing infrastructure
  - —Often unscientific in the use of what limited data are available
  - —Not organized to learn from experience, and learning far too slowly, if at all
  - —Not only woefully underperforming but in many cases getting even worse
  - —While computing has gotten billions of times better since early 1960s, …
    - o Obesity grew from 10% to 50% of US while costs grew from 5% to 18% of GDP
    - o Death by avoidable medical errors: perhaps 400,000/yr; serious injury @ 5-20X
    - o We've also spent $30B+ on interoperable EHRs but got close to no interoperation
    - o In 2015 only 25% HS seniors proficient in math, 37% in reading (down vs 2013)
    - o Broad familiarity with dysfunction in major government (procurement) systems
  - —These systems are really "eating us alive" economically, politically, socially

A Computing Systems Solution!

- The culture of computing wants its qualities for human systems
  - —Continuous sensing of physiology, psychic/sleep status, social status, etc.
  - —Improved situation awareness: radical transparency to critical phenomena
  - —Induce healthy individual & organizational behaviors across the board
  - —People, teams, firms, societies that analyze, envision, explore, learn, improve
- Make old world systems perform like new-world systems
- One idea: just apply CS and systems thinking to the problem!
  - —Leverage advances in sensors, mobile/cloud computing, AI, alerting, etc.
  - —For example, take blood pressure *every day (hour, minute),* c.f., Nyquist
  - —Inform subject in real time on need to adjust diet, exercise, sleep, etc.
  - —Gather data at scale (millions/billions) using mobile and cloud computing
  - —Learn from data, monetize, grow, and achieve better health outcomes!

Health Apps

- Conceptual architecture of typical health app

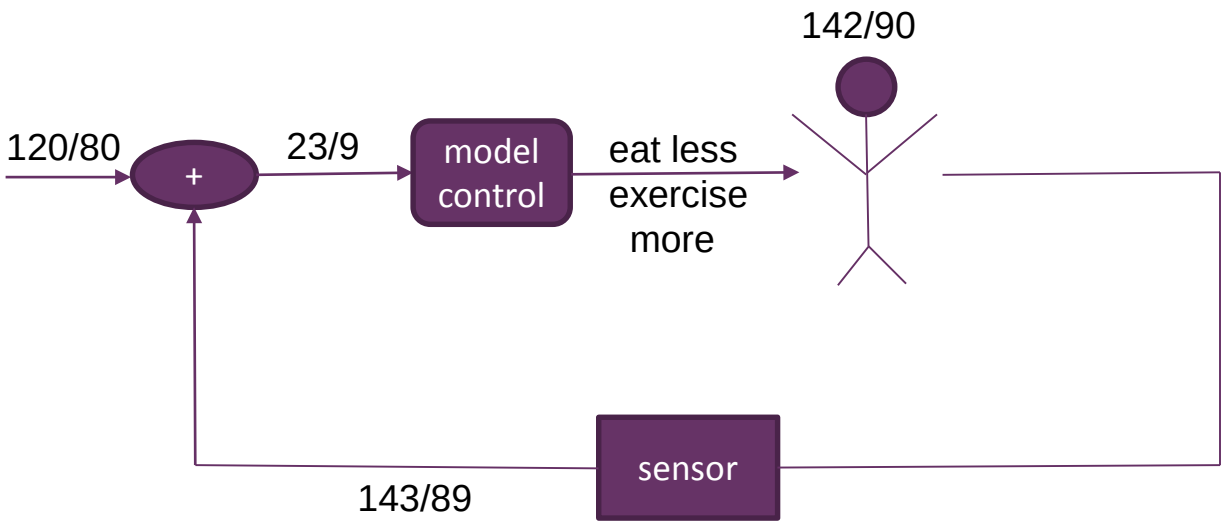

Health Apps

- Conceptual architecture of typical health app

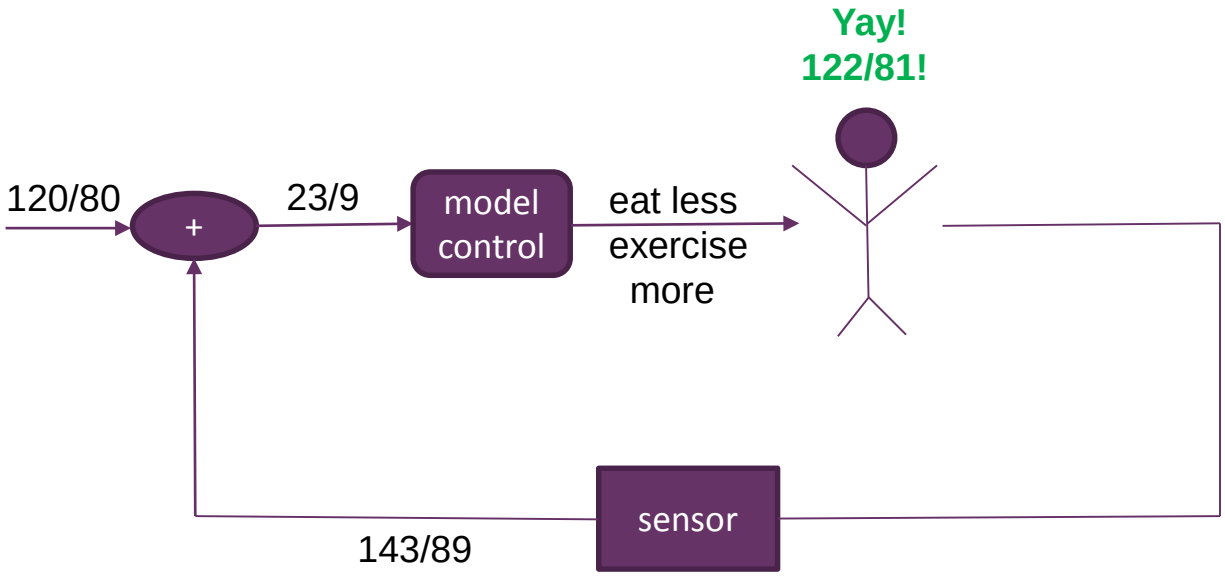

Well, It Basically Doesn't Work

- Unfortunately, this approach mostly doesn't work out either
- Tens of thousands of health apps in app stores, mostly useless
- One big problem: assumes *control theory* paradigm for people
- But people are not machines; treating them as such generally fails
  - —When the "plant" is human, organizational, social, you have to deal with …
    - cravings, values, competition, corruption, economics, incentives, objectives, power, psychological fear and reward systems, irrationality, bias, cognitive dissonance, intentionality, agency, ideology, norms, culture, laws, policies, responsibility, social networks (real ones), treachery, creativity, learning, fatigue, illness, forgetful, etc.
  - —People profoundly affected by environment (e.g., what's on their shelves?)
- The principles underpinning CS/SE/CPS, e.g., control & dynamics, are fundamentally inadequate for human/social domain

A Single Cyber-Social Systems Feedback Loop

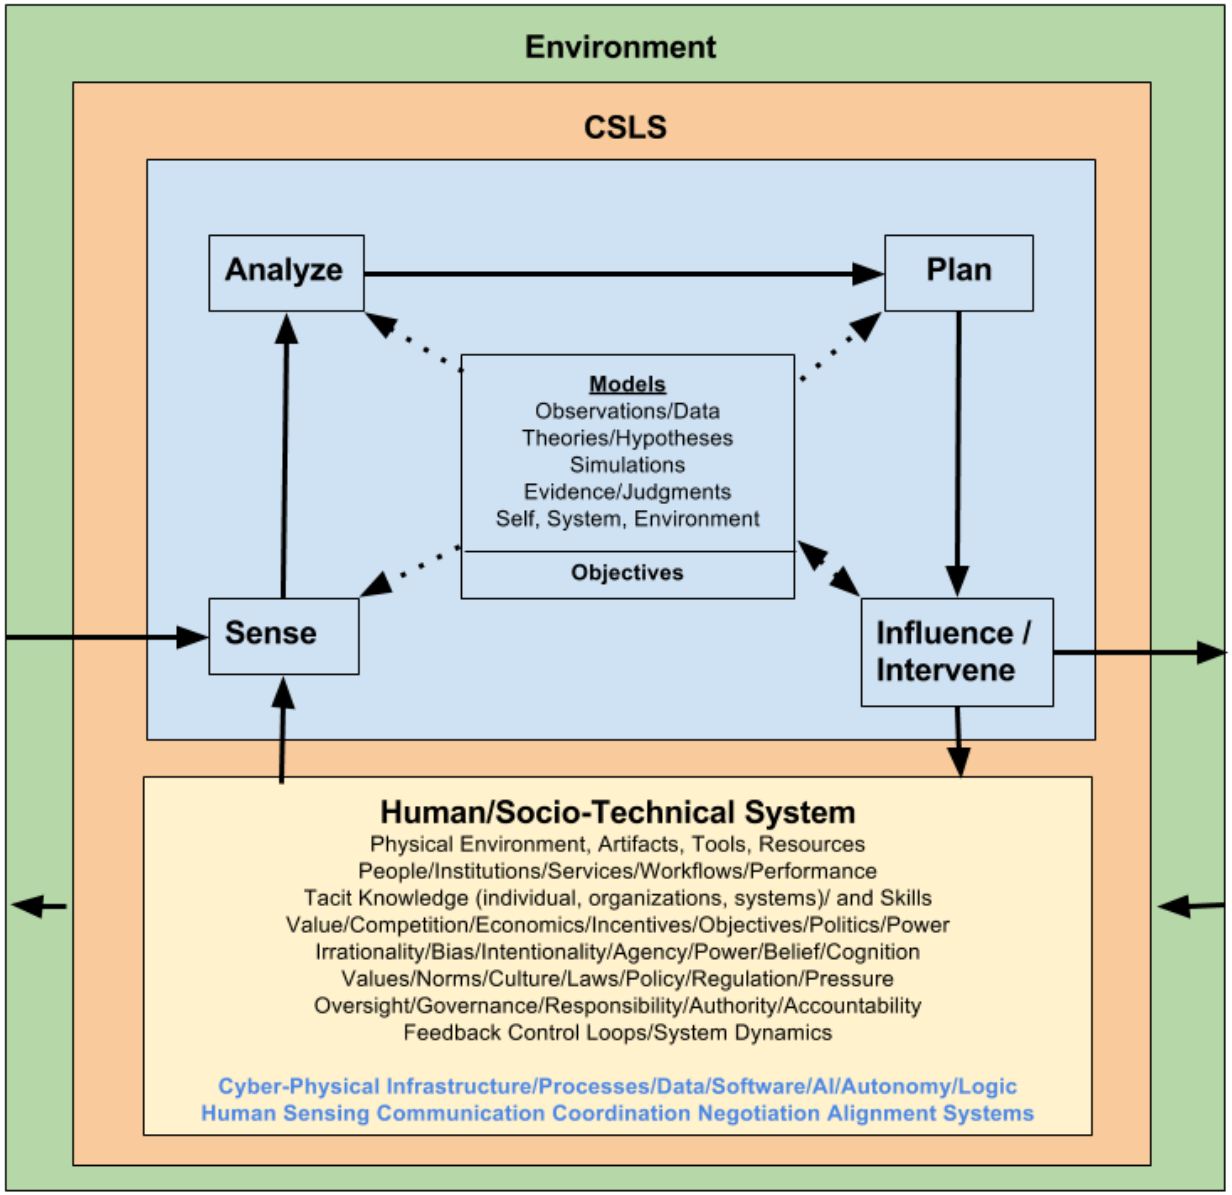

SYSTEMS ENGINEERING RESEARCH CENTER

Shortcomings of Computing/Systems Paradigm

- Often lack physics-like *system models* for humans & social things
- Interventions are far from mere adjustments to *actuators*
- Interventions often don't have deterministic *effects*
- Can have unexpected, undesirable *side effects*
- The "plant" is *smart* and constantly *learning and changing*
- And subject to diverse and complex (e.g., social) *disturbances*
- Effectiveness of feedback can be sensitive to *local conditions*
- The concept of a set point or *objective* is fraught (who sets it?)
- CS/SE/CPS thinking vitally important but also woefully inadequate

Goals

- Short-term: Bring together broad diversity of disciplines to identify audacious aims and outline a research agenda to establish conceptual foundations and principles for analyzing, developing, operating, integrating, and evolving cyber-social learning systems across scales and sectors

- Note: Critical is integration of perspectives and knowledge from other fields: basic, such as the social, behavioral, and economic sciences, the humanities (e.g., ethics), architecture, policy, law & business; applied, such as implementation and service science; and domain-specific, such as healthcare, education, defense, etc.

- Long-term: Catalyze a self-reinforcing transformation of critical sectors into dramatically better performing cyber-social learning systems

The Rest of This Talk

- Meeting these goals will require multiple paradigm shifts
- CSLS will be based on both traditional and innovative elements

CSLS Require Multi-Dimensional Paradigm Shifts

- From *machines* to *people* (individually and in social arrangements)
- From *control* to *influence* (over people, organizations, societies)
- From *knowledge* to *knowhow* (from propositions to tacit skills)
- From *management* to *politics* (variant of control -> influence)
- From *centralization* to *decentralization* (even crypto-anarchism)
- From proprietary knowledge *stocks* to open knowledge *flows*
- From *racing against* to *racing with* artificially intelligent machines
- From *optimization* to *adaptation, hedging,* and *acceptance*
- From *disciplinary* research to *convergence* (CS, SE, SBE, Hum, Art)

Paradigm Shift: From *Machines* to *Human/Social*

- We've always known it's about both (socio-technical systems)
- But most disciplines based on principles better suited to machines
  - —E.g., CPS rooted in integration of real-time systems, sensors, and control
  - —E.g., software engineering rooted in logic, formal representations, algorithms
  - —E.g., systems engineering emphasizes optimization, control, formal models,
  - —E.g., ICSE 2017 Research Track topics: documentation, analysis, concurrency, refactoring, dependability, search-based SE, process, testing, mobile app security (a separate track did publish papers on software in society)
- And when we do model people, it's often through a technical lens
  - —E.g., modeling people as expected value optimizers

SYSTEMS ENGINEERING RESEARCH CENTER

Paradigm Shift: From *Control* to *Influence*

- E.g., Fogg's behavior triggering in motivation-difficulty plane
- Behavioral economics applies psychological, cognitive, social, emotional insights to help explain and influence real behavior
- Cyber-social systems requires integration of CS&E with models & theories from the social, behavioral, economic, humanistic fields
- Examples
  - Trans-theoretical model of behavior change, Prochaska et al. 83, 92
  - *Influence: The Psychology of Persuasion*, Caldini 84
  - *Connected: The Surprising Power of our Social Networks and How they Shape our Lives*, Christakis and Fowler 09
  - *Mechanism Design: Analysis and Synthesis*, Erdman and Sandor, 1997

SYSTEMS ENGINEERING RESEARCH CENTER

Paradigm Shift: From *Knowledge* to *Knowhow*

- Question: How do you turn left on a bicycle?
- You know *how* to do this; but do you know *what* to do?
- Tacit knowledge can foster bad mental models
- Answer: To turn left, you must first turn the wheels *to the right*
- Critical to learning are *unlearning*, frame breaking, new lenses
- CSLS learning is ultimately about *performance* not *propositions*
- CS has largely focused on explicit propositional knowledge (that)
- Still a key; but people, groups, enterprises, sectors, societies must also unlearn old & learn *how to perform* in new and better ways

SYSTEMS ENGINEERING RESEARCH CENTER

Paradigm Shift: From *Management* to *Politics*

- Computing enables high impact, distributed, autonomous action
- Such actions do not always aggregate to positive social outcomes
- Lack of consensus values can drive divisions, conflict, and strife
  - —E.g., individual freedom vs. public safety around guns
  - —E.g., strong encryption, onion routing, and dark web
- Management challenges are transformed into political challenges
- Scaled solutions must positively aggregate autonomous initiatives
- Incentive structure, mechanism designs as *architectural* elements
- Governance is a major, potentially the largest, challenge in CSLS

Shift: *Centralized Control* to *Healthy Emergence*

- Traditional engineering uses centralized design management
  - Architect builders
  - Architect integrators
  - Prime contractors
- Emerging platforms foster disintermediation and emergence
  - Peer-to-peer communication platforms, e.g., Signal, Napster
  - Blockchain: peer-to-peer, authenticated, decentralized ledgers, e.g., Bitcoin
  - Envisioned decentralization of cloud computing, e.g., Tlon's Urbit system
  - Open production, design, and distribution networks in emerging countries
- How do we design the transformational platforms of the future to incentivize participation, effective governance, and emergence of healthy CSLS at scale?

SYSTEMS ENGINEERING RESEARCH CENTER

From Knowledge *Stocks* to Knowledge *Flows*

- From exclusive knowledge *stocks*
  - —Creating, protecting, exploiting proprietary, authoritative knowledge assets
  - —Scalable *efficiency*
  - —Yet return on assets for public companies in US collapsed by 75% since 1965
- To open knowledge *flows*
  - —Participating in more and more diverse knowledge flows at scale
  - —Main centers of *institutional innovation* are now in India and China (!!!)
  - —Scalable *learning,* open production/distribution systems, design networks
  - —Creating new knowledge, with a strong tacit *knowhow* component
  - —Scalable platforms enabling *pull* from 10s even 100s of thousands of peers

From *Racing Against* to *Racing With* AI Machines

- New symbiotic relationship between us and computation
  - —Race with not against "smart" machines
    - Freestyle chess vs Kasparov/Deep Blue
  - —Now "human + machine" teams play against each other
    - Humans still fully in control of ultimate decision making
    - Increase level of play to unprecedented heights
    - Achieve both tactical and strategic excellence
    - Help observers gain insights into excellence in human and machine play
  - —Gen. S. Crystal changed Iraq JSOC from running 10 ops/mo to 10 ops/night
    - People responsible for context and decision making
    - Machines for helping with intel, analytics, ISR (intel, surveillance, recon)
  - —*On threshold of new kind of human-machine symbiosis, freeing people to focus on context, imagination, decision-making*

Optimization to *adaptation, hedging, acceptance*

- Traditional engineered systems
  - —Objectives, dynamics, constraints measurable & tractable => optimize
- Much of contemporary software development
  - —System response time < environment change rate => sense & adapt
- Complex, long-lasting systems
  - —Multiple, viable alternative futures plausible => hedge / real options
- Else accept

Paradigm Shift: From *Disciplinarity* to *Convergence*

- Problems of a kind not susceptible to intra-disciplinary solutions
- Structure of disciplines countervailing against multi-disciplinarity
  - Emphasis on "proprietary" (intellectually siloed) knowledge stores
  - Positive and negative implications of the peer review system
  - Interdisciplinarity is always vulnerable to cutbacks
- Cross-cutting funding, integrative workshops important, insufficient
- We haven't explored how to use new technology and research methods to integrate across disciplines for convergent research
- *We must build community among researchers across all fields critical to the realization of cyber-social learning systems?*

SYSTEMS ENGINEERING RESEARCH CENTER

Some Key Aspects of Cyber-Social Learning Systems

- Multi-disciplinary pull (CS, SE, SBE, A&H)
- Design / architecture / emergence
- Platforms / infrastructure
- Human / social feedback loops
- Governance models and mechanisms
- Assurance of critical CSLS properties
- New research communities, aims, methods, results

Multidisciplinarity

- Problem: Fragmentation of research societies
  - —Need: Cross-community transactive memory systems
    - Respecting our own and other expertise
    - Knowing the boundaries of what we know
    - Knowing where we should defer to other expertise
    - Creating incentives & mechanisms for voluntary participation at scale
- Bridge building between CS, SE, Social/Behavioral/Econ Sciences
  - —What can we build ← → what is worth building
  - —*How do we build community and communication among researchers in the Learning Sciences and related fields?*
  - —Cross-cutting funding and integrative workshops important but insufficient
- *Apply CSLS principles to construction of CSLS research community!*

Multidisciplinarity

- Attempted solutions and enduring problems
  - —The structure of disciplines itself serves to perpetuate the problems
    - Positive and negative implications of the peer review system
    - Problem is that interdisciplinarity is always vulnerable to cutbacks
  - —Cross cutting funding goes a huge way towards solutions
    - Science of learning centers is a great example
    - Research on some large, interdisciplinary programs shows this doesn't always work
    - Workshops like the CCC workshop series are extremely productive
      - — Huge progress in big data for education
      - — But we all go back to our departments and labs
  - —*We have not explored how we can use technology to fundamentally change how disciplines are structured*

Multidisciplinarity

- Need new foundations for acquiring, relating, using models drawn from knowledge housed in multiple, still largely siloed disciplines
  - —System models are at heart of any feedback-driven and/or learning system
  - —Learning captured to a significant degree in evolving computable models of what's known (but see later comments on criticality of tacit knowledge)
  - —Must augment physics-like models with models of human/social phenomena
  - —Many such models will (have to) be learned from data, and will be contingent
  - —How to develop, validate, integrate, evolve, and use CSLS system models?

Architecture: Some Elements

- Data & computation; human/social intervention mechanisms
- Computable models of human/social phenomena
- Communication, coordination, governance, oversight subsystems
- Sensing, double-loop learning & adaptation: stores, flows, waves
- Ecosystem governors and stabilizers
- Human/social forces/glue (e.g., induce voluntary participation)
- Catalysts of healthy emergence (e.g., network effects)
- Continuous assurance of critical properties of learning systems

Architecture: Composition of Social Components

- Key
  - **Components,** *connecting relationships,* dynamics, metrics
- Example
  - **Governments** *regulate*, and …
  - **Journalists** *review* and *report* on ….
  - **Professionals, managers, researchers, designers, analysts**, …
  - *Threatened by* **malicious actors** …
  - While *providing service* to and *obtaining feedback* from ...
  - **Patients, nurses, administrators, passengers, students, soldiers**, …
  - All of whom are *enabled* and *required* or *incented* to ...
  - *Participate* in **technology-mediated value co-production processes**
  - Evolving to track high states of fitness relative to changing **objectives**

Cyber-Social Platforms / Infrastructure

- E.g., Large-scale, distributed "learning / doing" fabrics that capture contexts and actions for reflection, learning, archiving
  - —Every piece of data contextualized, citable, machine-checkably interpretable
  - —Data flows instantly and seamless through evolving analysis pipelines to uses
  - —Automated support for discovery, analysis, important relationship finding
  - —Platforms integrating devices, computing, human actions and interactions
  - —Model discovery as creating and updating of contextualized digital assets
- Example: Discovery environment for relational information and versioned assets (Deriva: http://bd2k.ini.usc.edu/tools/deriva/)

> "...the key question is not whether a problem is a "social" problem or a "technical" one. That is putting it the wrong way around. The question is whether we choose a primarily social or a technical solution, or some combination. It is the distribution of solutions that support infrastructure in different ways at different moments." Bowker et al, 2010

Interaction and Governance

- Communication & Coordination Infrastructure
  - Deliberative system design tools for thousands (millions, +) of participants
  - Large-scale teamwork monitoring and improvements systems
  - Collaboration toolkits & remote teamwork systems
  - Online conflict forecasting and dispute resolution tools
- Governance systems
  - Objectives, tradeoff decisions – who gets to set and make them?
  - Responsibility – duty and ability to take action
  - Accountability – responsibility to report on events
  - Consequences – penalties and reward for desired and undesired behaviors
  - Controls on large-scale system behaviors – regs, compliance, reporting, etc
  - Processes, communities, norms/values, etc. E.g., IETF/internet/WWW

Research : Evidence-Based CSLS Assurance

- Process
  - Assurance considerations addressed at outset and continually thereafter
  - Co-production of design artifacts, propositions, and evidentiary structures
  - Governance (a community cyber-social process) to manage evolution of common framework and API models, data models, process invariants, etc
- Evidence
  - Diverse kinds of evidentiary data, informal and formal
  - Dependency models and argumentation structures to link evidence
  - Analytic models for hazards, safety, security/threats, privacy, compliance
- AI components: Integration of three technical approaches
  - Algorithms and explanations
  - Models and reasoning
  - Safety and systems engineering

Research Example: Learning

- Greatly elaborated analysis of learning to underpin theory of CSLS
  - Much more than education, and much more than machine learning
  - Similarities and differences between human and machine learning
  - Detailed analysis of ML, e.g., adaptive (autonomous) vs. adaptable (guided)
  - Design methods for complex learning systems, e.g., meta-design, libertarian paternalism, end-user development
  - Models for academic-industry collaboration around systemic problems

Research Example: Learning

- Start with theories of human learning and posit whether and how distinctions operate in a health care system, smart city, etc.
  - —Differences in how we can instrument, guide, and assess learnings
  - —Wide chasm between understanding micro processes like neural firings and making a good life decision. What are the analogs for systems that learn?
  - —Memory: Retention in wetware vs retention in the cloud must have unique metadata requirements
  - —Metacognition: When an individual fails to learn something or learns it poorly, the consequences may be severe. When a system fails to learn something or learns it poorly, there can be consequences. What is the systems analog of meta-cognition (awareness of own though processes)?
  - —How do we assess the summative and integrative effects of hundreds of individual learning processes to understand the effectiveness of an entire smart community, health care system, defense establishment, etc?

SYSTEMS ENGINEERING RESEARCH CENTER

Research Example: (Machine) Learning

- Data from social phenomena are multi-modal, multi-relational, spatio-temporal, multi-media, etc
- Most ML techniques assume tabular and IID data
- Need ML methods that account for relational context of data

SYSTEMS ENGINEERING RESEARCH CENTER

CSLS: The Future of Systems

- Smart cities: Design, development, deployment of an emerging class of cross-platform, service-integrated technology products to enhance performance and create a platform for economic development in cities and other communities
- Education: "If the ladder of educational opportunity rises high at the doors of some youth and scarcely rises at the doors of others, while at the same time formal education is made a prerequisite to occupational and social advance, then education may become the means, not of eliminating race and class distinctions, but of deepening and solidifying them." --President Truman, 1947
- Healthcare: An audacious goal: reduce mortality from avoidable medical errors by 10X within 20 years, and morbidity similarly, saving economy $1T+/year
- Defense: Learn how to learn faster than every major adversary at all scales: in space, organization, time, and in doctrine, strategy, tactics, and technology (we are at risk of losing the learning arms race before we even see that it's started)
- Policing and criminal justice: Assure equitable and just treatment of all people

SYSTEMS ENGINEERING RESEARCH CENTER

We're Not in Kansas Anymore

- These are all problems in the field of cyber-social learning systems
- A report from a speakers in one workshop's closing session

> "… there was a growing sense of enthusiasm and discussion within the room that it was starting to feel like we're not in Kansas anymore … once you start putting all of these pieces into the same design constellation we were looking at something that to us felt pretty different."

- I hope to have conveyed a sense of why some people felt that we are now looking at a *real phase transition* in the world of systems
- From separate old and new world, service and computing, systems to a future of very high-performing cyber-social learning systems

Thank You!

For listening, and to the many workshop participants whose ideas and insights, some published elsewhere, are presented in this talk

Sullivan@Virginia.EDU

SERC 2017 Annual Events

REGISTER NOW

For more information or any questions regarding this event, please contact:
Ms. Monica Brito or Ms. Megan Clifford

SYSTEMS ENGINEERING RESEARCH CENTER

SERC TALKS

UPCOMING TOPICS:

Cybersecurity Series

How Do We Prepare the People Who Will Need to Manage the Real-time Responses to Cyber Attacks on Physical Systems?

Dr. Barry Horowitz, , University of Virginia

August 2, 2017 | 1:00 pm ET

What are the Top Ten Software Security Flaws?

Dr. Gary McGraw, Vice President Security Technology, Synopsys

October 4, 2017 | 1:00 pm ET

Thank you for joining us!

Please check back on the SERC website for today's recording and future SERC Talks information!