CCC Cyber-Social Learning Systems Workshop 1

Narrative Summary

Kevin Sullivan, Charles Friedman, Joshua Rubin, and Organizing Committee Members

Background

This two-day invitational Computing Community Consortium (CCC) scientific workshop convened in Seattle, WA on August 29-30, 2016, and was the first in a series of three workshops on cyber-social learning systems (CSLS). These workshops and the issues they will address evolve from multiple converging trends that together present deep and novel research challenges and opportunities. This first workshop, framed largely from a computer and information science perspective, focused on crosscutting basic science and engineering research challenges, grounded in the realities of sectors including health and healthcare, smart communities, and education.

Over the last decade, we have made enormous progress establishing scientific and engineering principles for cyber-physical systems (CPS). We are now on the threshold of a world of physical systems that are deeply computational, dynamic, learning, and connected at all scales, yielding radical improvements in physical systems properties. The next major frontier in science and engineering research and development is the integration of cyber-physical with human and social systems and phenomena across all major sectors and at all scales. Closing the loop from sensing to performance at all scales will give rise to CSLS. CSLS are socio-technical systems that support continuous and rapid learning – by machines, people, institutions, sectors, government, and society as a whole – enabling rapid adaptation to evolving conditions and ultimately radical improvements in system function and performance. CSLS hold promise to transform all major sectors of our society: healthcare delivery, education, community services, transportation, justice, civil rights, environment, homeland and international security, bio-security, business and commerce, to name some. Research and development in CSLS holds the potential to enable and drive radical improvements in the behavior and performance of human and social systems, and to do so in a manner consistent with the values of our open, modern, democratic society.

Challenges and opportunities in the CSLS space related to three emerging trends and their synergistic fusion: i.) progress in computer and information science as well as engineering, including (deep) machine learning at scale, CPS, data science, and artificial intelligence (AI); ii.) advances in the cognitive, behavioral, social, and systems sciences that has focused on processes of change across levels of scale including individual, group, organizational, and societal; and iii.) trends emerging from within crucial areas of human endeavor such as health and education and smart communities, from which has evolved the concept of a learning system that can continuously study and improve itself through iterative cyclical processes that generate data, convert data to knowledge, and apply that knowledge to change practice and improve performance.

The first such visioning workshop in this series – aiming to ultimately produce an agenda for developing the scientific, engineering, and design principles that will underlie the development, evaluation, operation, and evolution of future CSLS – had the following preliminary stated goals:

1. To expand the community interested in CSLS research and development, as established by earlier (2013 and 2015) National Science Foundation (NSF) funded Learning Health System workshops, to include the broader Computer and Information Science and Engineering

(CISE) and NSF community and participants from major sectors beyond health and healthcare;
2. To explore issues involved in, and develop an initial rationale (to be refined in subsequent workshops) for, CSLS research and development;
3. To develop an initial agenda for basic and applied research in CSLS (to be refined in subsequent workshops), integrating basic research in computing, engineering, and the social sciences, with applied research in multiple sectors, in part building on results from the 2013 and 2015 NSF funded workshops.

The Workshop

The workshop began by describing deep societal challenges spanning health, education, cities and communities, energy and the environment, defense, biosecurity, and beyond. We framed the challenge as working together to "integrate and anneal concepts, methods, and tools from across these disciplines to develop principles, methods, and tools for transforming critical systems into cyber-social learning systems exhibiting breakthroughs in function, performance, evolution, trustworthiness, and trust, across domains and scales, and in ways consistent with the values of an open, democratic society." Collectively, participants would lay the groundwork for upcoming workshops in November and January, ultimately aiming to advise the research community and policymakers vis-à-vis CSLS.

The discussion continued with a series of short talks providing participants with an opening perspective on CSLS. William Rouse of the Stevens Institute of Technology led off with a discussion of learning in complex systems. Dr. Rouse distinguished single-loop learning from double-loop learning. William Griswold of the University of California, San Diego, gave a talk titled, "It's (all about the) People!". Using health as an example, he described how even in a world of plentiful computing, people (including clinicians and patients) are not information processing machines. Elizabeth Churchill of Google spoke about full-stack infrastructure for CSLS. After describing social and political factors affecting the development of electrical power grids, she highlighted examples of socio-technical layers of infrastructure projects at Yahoo!, eBay, and Google. William Scherlis of the Carnegie Mellon University School of Computer Science discussed abstraction, assurance, and trust for CSLS.

A series of short talks on "The Edge of CSLS" followed. Tarek Abdelzaher of the University of Illinois at Urbana Champaign discussed social sensing. Dr. Abdelzaher posed a question: now that we are in an age where data is being collected at a large scale how do we integrate that data and use it to understand systems better? Dr. Abdelzaher looked at how humans and their connected devices detect signals in the world, and how in the age of social media, events in the world impact social media in ways that can in turn be detected by machines. Jim Spohrer of IBM spoke about AI. IBM is taking an approach to AI as a form of augmented intelligence, scaling the expertise of people. He spoke about AI serving the public good. He described AI in the context of tools, assistants, collaborators, coaches, and mediators. Lori Clarke of the University of Massachusetts Amherst College of Information and Computer Sciences discussed trustworthiness of CSLS. She raised the question as to whether, as continuous learning systems, CSLS can themselves learn to be more trustworthy over time? Dr. Clarke raised three key trust concerns: reliability, security, and continuous evaluation. Ben Shneiderman of the University of Maryland spoke about design tools for governance in CSLS. Calling for boldness, Dr. Shneiderman suggested the importance of incorporating ethnographic and qualitative methods into the work. John Seely Brown (JSB) of the University of Southern California spoke about learning and unlearning in exponential

times, with a focus on the importance of and challenges associated with tacit knowledge (especially as the rate of change increases).

The plenary discussion raised several important areas for further thought. One included protecting people from themselves and protecting a learning system from malevolent actors who will take advantage of these systems' processes. The concept of unlearning (e.g., of habits or unhealthy behaviors) and challenges associated with doing so garnered further discussion. Qualitative methods were further discussed, as was what was meant by "learning" in various contexts (e.g., human learning, machine learning, systemic learning).

There was a series of practical short talks about CSLS spanning three domains: health, education, and smart communities. Jonathan Silverstein of the Kanter Health Foundation discussed health and the LHS vision. He began by illuminating challenges healthcare and health are facing, and then pointed out ways informatics and predictive analytics are transforming other domains – asking what about biomedicine? Stephanie Teasley of the University of Michigan spoke about a learning system as one "with the capacity to continuously study and improve itself" and discussed realizing this vision in the context of higher education systems. With the creation of new data sources that are invisible to the learner, learning analytics is now possible. Jennifer Clark of the Georgia Institute of Technology Center for Urban Innovation spoke about smart cities and connected communities. She framed the challenge as, "design, development, and deployment of an emerging class of cross-platform, service-integrated, technology products to enhance performance and/or create a platform for economic development in cities and communities."

The participants then divided into six cross-disciplinary breakout discussion groups. The six groups were all given the same charge: to identify hard but potentially solvable problems that stand in the way of breakthroughs in function and performance across domains as envisioned in our discussion of CSLS. The groups were also asked to identify corresponding opportunities for basic and applied research to advance the science, engineering, and design of CSLS to address these problems and enable such breakthroughs, possibly including new research methods, communities, and testbeds. After meeting separately, the groups reconvened to share slide decks they prepared and to report on their discussions.

Group 1 identified key hard problems that included: developing a maturity model; defining success; scalability and composability; measuring progress, studying, "optimizing" a complex, dynamic system that changes; developing appropriate research methods; and driving a culture shift from blame to safety and continuous improvement.

Group 2 focused on moving from the edge of CSLS to the core. They stressed the importance of understanding architecture, including representing and developing a common syntax, modeling and creating different incentives across the diverse stakeholders, and testing the synthetic environment.

Group 3 identified challenging problems and research opportunities that included: catalyzing self-evolutionary systems that redefine input and outputs; developing language to define and automate CSLS; innovatively addressing rigorous challenges of eliciting, representing, and adapting to stakeholder needs at different levels of granularity and scale; defining and dealing with ethical, social, and legal issues (ELSI) in CSLS; and identifying and reducing bias in CSLS.

Group 4 envisioned human data interaction and data democratization as being key goals and challenges. They recognized a potential new discipline in data reasoning and data design.

Group 5 questioned, "what is the value and impact of data for the future of CSLS?" They noted that data is decontextualized, and cited the example of incredibly low completion rates for Massive Open Online Courses (MOOCs) to illustrate the point.

Group 6 identified the hard problems as, "design, prototyping, development, operation of CSLS (i.e., large-scale, long-running, open, and adaptive) systems; measurement and methodology; reasoning about emergent behavior; and validating system properties (including: functionality, safety, privacy, compliance, etc.)."

The second day of the workshop took on the tone of a participant-driven, self-organizing "un-conference". It began with a "game reset" plenary talk by Dr. Friedman and Dr. Sullivan. After synthesizing some key points from the previous day, they looked ahead to what shape deliverables and outcomes from the workshop series may take. The plenary group discussed what a paper might look like, including references, theoretical underpinnings being built upon, and questions addressed. The discussion also raised where metrics, tools, theories, and research methods spanning diverse disciplines fit in the new models envisioned.

The participants organized into three working groups. The first group was tasked with looking at the cyber-human ecosystem, including considering new forms of double-loop learning and designing for emergent behavior. The second group was tasked with considering trustworthiness of and trust in CSLS, including considering the role of responsibility and issues related to privacy/security/safety/ethics/ELSI. The third group was tasked with exploring issues related to CSLS reference architecture and infrastructure, including enterprise level issues and governance.

Each of the three working groups was charged with addressing the same set of questions for its respective topic area. These questions included: Why is this area important to CSLS (and why is this area unique to CSLS)? What would be the key research challenges? What are the key opportunities? What would a testbed look like? What would the relevant metrics look like? What do the research methods look like (and how do scientists actually do this research in the real world and realize impacts)? What are guiding examples worth looking at and building upon? What are related ethics/ELSI issues to consider?

The first group looked at the cyber-human ecosystem. They distinguished retrospective from prospective learning, and suggested social cognitive theory could underpin further research in this arena. The role of responsibility entered into the discussion, distinguishing ultimate responsibility by humans, with limited responsibility of machines in the system. They envisioned a new enterprise with novel needs for determination and verification, and new approaches to making informed decisions that can be far more remote.

The second group considered trustworthiness and trust in CSLS. They defined trustworthiness as related to attributes of the system (whether it should merit trust), versus trust as representing human attitudes toward it. People trusting an insufficiently tested autonomous system in a self-driving car might be an example of over-trusting. Some attributes of trustworthiness include: security, safety, privacy, reliability, resilience, and ethics. Trust attributes vary according to roles and backgrounds of participants and stakeholders and can be influenced by extrinsic events. After identifying key tensions in CSLS, research topics identified include: attributes of trustworthiness, measuring trust as it evolves over time, challenges in the certification of trustworthiness (and certification of processes of adaptation), and the concepts of "trust engineering" and "trust management".

The third group explored CSLS reference architecture and infrastructure. The discussion began with IBM's MAPE-K model as straw architecture. The group recognized various roles for experimentation at different levels of scale, and noted that mechanism design must include considerations of incentives that drive desired emergent phenomena.

Following presentations by the three groups, a plenary discussion of emergent topics took place. The group assessed commonalities and differences across the three CSLS domains (health, education, and communities) in light of what was learned from the groups – as well as forward looking next steps. A discussion of testing and deploying CSLS ensued. The group asked: how far can we get with modeling and simulation and how do we get from something a research group can do to something that's in the real world? They discussed a failure of imagination that comes with over-focusing on data, as opposed to driving knowledge generation and informed decisions that advance the social good and the public good. In the smart communities domain, there was discussion of the emergent component called urban science that is struggling to evolve a cross-disciplinary cadre of researchers with separate and collective domain expertise fit for a world where real time sensor data and other developments could drive more informed and empowered communities (and better outcomes). The concept of a CSLS testbed (empowering double-loop learning) was discussed. That led to a series of questions, noted as: "What disciplines do you need to bring together to succeed in getting good information coming out of a testbed environment? And how does this go to the question of research methods? Do social scientists really need to be a part of it? How is research evaluated based on how applied it appears to be?" How to define success and how to measure (continuous) improvement were two other key themes discussed.

A closing panel discussion followed. Mary Czerwinski of Microsoft Research presented first. She began by discussing honest signals. Honest signals are signals people emit without being consciously aware of them. Dr. Czerwinski stated that CSLS need to be polite and mimic people's conversation and interaction styles to maximize their effectiveness. She discussed affective computing and gave an example of precision psychology: using deep machine learning what aspects of particular therapies work for specific people, contexts, and problems. Dr. Czerwinski described challenges in scaling what is recognized to work well.

Susan Graham of University of California, Berkeley spoke about system level issues in CSLS. She recognized the dynamic properties of CSLS. Dr. Graham discussed the need to communicate better between disciplines. She recognized a practical challenge: in academia especially, these CSLS challenges often need to be addressed within domains, because that is where researchers expect to publish. Dr. Graham stressed the need to develop mental models of CSLS, and to address what we want CSLS to accomplish and why in more precise terms.

Starting with an image of drinking from a firehose, Rahul Basole of the School of Interactive Computing at Georgia Institute of Technology endeavored to integrate various pieces of the multi-faceted CSLS puzzle discussed. He stated that the component disciplines are largely disconnected; it is our challenge and our opportunity to close the gap. Some CSLS characteristics Dr. Basole observed include: complex ecosystems, multi-stakeholder, interdisciplinary, amorphous boundaries, and contextual. After discussing approaches to analyzing CSLS, Dr. Basole illuminated how CSLS must (and will) shape policy, and in turn, policy will shape CSLS.

Lise Getoor of the Computer Science Department of University of California, Santa Cruz spoke about deep machine learning in the CSLS context. Dr. Getoor recognized that data underpinning CSLS is, "multi-modal, multirelational, spatio-temporal, and multi-media." She described processes for

"flattening" data to render it more useful, coupled with the need for "machine learning method that that take into account relational context." After reemphasizing the value of double-loop learning, she closed by discussing three types of research opportunities in CSLS. The first involves tailoring machine learning models to handle causation and social interference, as well as evolving social science to integrate big data (machine learning) and small data (ethnographic studies) plus culture and social science theory. The second involves developing research approaches that include humans in the loop (and in the data and in the system), as well as communities and societies in the loop. Building upon the "three c's" of context, comparative modeling, and collective reasoning, the third set of research opportunities are underpinned by a recognition that, "There are common abstractions that go across a broad number of the most compelling (CSLS) domains including heath, education, smart cites, citizen engagement and more."

Following the discussion, Dr. Friedman and Dr. Sullivan led the "Conclusions and Next Steps" closing session. Looking toward future workshops, Dr. Friedman stressed the paramount importance of ethics and ELSI issues in CSLS.

Dr. Sullivan emphasized the incredible opportunity afforded by the broad perspective the CCC is taking. He acknowledged that the computer science community (and no single discipline or stakeholder group or sector) can tackle this problem alone. What is needed is to converge and integrate, centered around advances in computing, a broader set of disciplines. These diverse disciplines can in turn come together and collaboratively make advances around pressing societal challenges. Major advances in deep machine learning, AI, and other areas of computing truly need to be informed by ethics, the humanities, and social sciences, if they are to realize their positive disruptive potential. Dr. Sullivan believes the diverse scientists involved must move forward as an increasingly coherent and identifiable community. Looking forward, a key goal of the workshop series will be to compellingly describe CSLS and the science (and corresponding investments and policies) needed to realize its potential to the next Presidential administration; the final workshop will be convened in Washington, DC just prior to the next President taking office. Dr. Sullivan opined that CSLS is part of the new tidal wave of epistemological change of what we mean by research. CSLS is not reductionist in nature. CSLS will drive transformative interdisciplinary collaborations that will innovatively address real problems in the real world, will be contextually based, and will tackle pressing societal needs in ways that promise to realize profound, sustainable, and disruptively transformative impacts that touch society and people's lives.