CCC Workshop Summary:
Research Opportunities and Visions for Smart and Pervasive Health

Elizabeth Mynatt mynatt@cc.gatech.edu Georgia Tech	Gregory D. Hager hager@cs.jhu.edu Johns Hopkins University	Santosh Kumar skumar4@memphis.edu University of Memphis
Ming Lin lin@cs.unc.edu University of North Carolina Chapel Hill	Shwetak Patel shwetak@cs.washington.edu University of Washington	Jack Stankovic stankovic@cs.virginia.edu University of Virginia

Helen Wright
hwright@cra.org
Computing Community Consortium

Version 1: March 12, 2017[1]

The past decade has witnessed tremendous improvements in an expanding number of computing capabilities—sensors, advanced analytics, networks, data infrastructure, advanced imaging, and cyber-physical systems—that have, and will continue to, revolutionize health research, and healthcare, with resulting improvements in health and wellness. As the cost of healthcare, the prevalence of chronic disease, and the scope of healthcare disparities increase, so too do emerging paradigms in smart and pervasive health that rely on the growing availability and robustness of sensors, mobile health capabilities and real-time analytics. These paradigms offer holistic approaches for health and healthcare delivery to enable precision medicine and personalized healthcare that can improve prevention and control the growing expenses associated with chronic disease management and persistent healthcare disparities.

The Computing Community Consortium (CCC) sponsored a workshop on Discovery and Innovation in Smart and Pervasive Health on December 5-6th, 2016 in Washington, DC. This workshop brought together leading researchers and policymakers to discuss the successes of Smart and Pervasive Health research activities, the evolution of relevant computing capabilities and the application of these technical innovations to health and wellness goals. This report highlights these paradigms and concludes with specific recommendations for the successful future implementation of an impactful smart and pervasive health research agenda.

[1] Contact: Ann Drobnis, Director, Computing Community Consortium (202-266-2936, adrobnis@cra.org). For the most recent version of this workshop report, as well as related reports visit: http://cra.org/ccc/resources/workshop-reports/

1.0 24/7 Monitoring for Physiological, Psychological, Environmental and Social Information

Always available health sensing and behavioral monitoring offers the transformational possibility of health and wellness that is bound by neither space nor time. It enables observations of dynamic changes in each individual's health state as well as key physical, biological, behavioral, social, and environmental factors that contribute to health and disease risk, anytime and anywhere that have not been possible in the past.

Although there has been a significant progress in the use of health sensors for remote care for specific diseases (e.g. heart rate monitors), **future strides will come from "holistic" sensing approaches** that integrate wearable, environmental, behavioral, and social network sensors and create new computational models that for human states can, for example, dynamically recognize the influence of environmental factors on human physiology and behavior. With the explosive growth of social media data, it is now possible to obtain temporally varying measurements of interpersonal exchanges, social network structures, and social capital, which can be combined with complementary physiological and psychological measurements of an individual collected by sensors. These dynamic measurements can also reveal network and community effects in health outcomes — quantification of which has been challenging so far due to the paucity of adequate social and community-level data. These approaches also aim to capture the *context* surrounding human behavior to effectively guide healthcare delivery. For example, the noise and bustle of an ICU should affect the content of ICU information displays just as the availability of healthy food and measures of social influence should affect diabetic care and nutritional coaching.

There are several challenging issues in analyzing data, designing decision-support systems, and developing strategic mechanisms for reaching informed responses from holistic sensing approaches:

- How can we measure key environmental variables, such as exposure to cues and triggers for adverse health-related behaviors, and infer a comprehensive characterization of individual behavior?
- How can we model the effect of environmental and social factors on behavior regulation to support behavior change and behavior maintenance?
- How do we combine large-scale population "big data" with sparse individual data, controlling for biases in each?
- How do we capture and effectively model the context of and around an individual to guide the "how" of healthcare delivery?
- How can we support decision-making where prioritization of key factors is critical to effective clinical care?
- How can we improve robustness and reliability of these systems operating "in the wild" outside of the traditional confines of healthcare environments?

Promising approaches include the development of explainable models that support inferences from digital biomarkers that facilitate interpretability in decision making; modeling advances that allow inference making in the presence of temporally imprecise labels; time series pattern mining methods to extract events of interest from multi-modality stream of data with varying temporal granularity; and provenance systems that capture both metadata as well as annotations of the entire data processing stage to facilitate both interpretability and comparative analysis.

Applying cyber physical systems technology to healthcare is also promising. This research trajectory includes solutions for real-time processing, plug-n-play components, human in the loop feedback control strategies, formal methods for safety specification and analysis, and dependability and robustness techniques for uncertain environments. We contend that combining big data analytics, machine learning, data mining, and natural language processing with cyber physical systems merits increased attention. There is also great promise in the revolution in sensing stemming from nano-devices and “smart skin” as these technologies become more inexpensive and pervasive.

2.0 Human-Centric System Approaches for Smart and Pervasive Health

Holistic data approaches enable a spectrum of system capabilities that act to integrate data insights and human action. **Fundamental progress will come from systems that actively incorporate input from patients and caregivers and emphasize human engagement** as a whole. These systems must grapple with a host of challenges including balancing human direction and autonomy, integrating real time input and control, creating scalable approaches that manage overlapping layers of input and action, and creating mechanisms for learning.

Mixed Initiative Systems and Closed Loop Systems: Mixed initiative systems balance automation and independent human action including patients, physicians and caregivers. For example, consider a wearable system that helps monitor activities and glucose for patients with diabetes. This system might report every week to the clinician a summary of the patient’s daily activity and note unusual spikes in glucose, it might offer supportive hints and reminders throughout the day to the patient, and send requests for help or feedback to family members as necessary. Mixed initiative systems must also create a flexible and time-varying division of responsibilities. Decisions that a patient might be capable of making under typical conditions might not be possible under periods of stress, or during episodes of depression or days when dementia symptoms are acute. While traditional mixed-autonomy systems rely on a single, engineered structure for division of task and responsibility, next-generation systems must have the capability to safely support different models of responsibility and risk management for individualized care communities and across changes in the capabilities of individuals. Mixed-initiative systems must have some underlying model of possible human action. Anticipating the range of “random” human action is a challenge, especially when the goal is to create a real-time, closed-loop system.

Closing the loop is ultimately the objective of any smart health system – using data, analytics, and interaction with patients and providers to adapt to the situation. However, adding real-time elements introduces the additional complexity of trading automation/autonomy against oversight, filtering, and feedback through a human. Furthermore, creating closed-loops systems typically involves choosing an objective to optimize, but in smart health, which objectives do we chose, and how do we measure them? For example, rehabilitation may be accelerated by “pushing the system” to the limit, but at the risk of patient comfort; coaching may lose effectiveness if it does not respect patient education, experience, bias, etc., things that are hard to measure and predict; decision-support may not be effective if it cannot take provider context into account.

Many of these systems are distributed – e.g. an ICU is really a complex cyber-physical system-of-systems with multiple controllers (decision makers) that either decide independently on different aspects of the overall system or collaboratively for the entire system, which raises the overall complexity level.

Multi-Tiered Sensing, Modeling and Control Systems: Multi-tier SMC system approaches create the foundation for mixed initiative systems that flexibly balance automation and human action and closed loop systems that integrate real-time data to optimize for specific health outcomes**.** Healthcare data, computed inferences and associated models have a natural multi-tier structure, with varying levels of rate, time-scales, abstraction and specificity. Therefore, it is natural to design "smart health" systems that mirror this multi-tier structure. Multi-tier SMC framework opens various opportunities for:

- Using different models at different time scales and abstraction to flexibly integrate different sources of data, from real time sensors to public health data.
- Creating techniques that help transition from cold models (population) to hot models (personalized).
- Dynamically assigning different roles to humans, e.g. across a coordinated care team or between patients and informal caregivers.
- Managing tradeoffs in sensing, model prediction, data sharing, and privacy.
- Engineering systems to optimize resources across multiple data gathering, modeling, and decision-making processes.

JITAI Systems: JITAI systems, just-in-time adaptive interventions, integrate mobile health applications with dynamic information about an individual's emotional, social, physical and contextual state to facilitate the adoption and maintenance of healthy behaviors and to discourage and prevent negative health behaviors and outcomes. They are designed to provide the right type and amount of support at the right time in person's natural environment by adapting to the person's changing internal state and context. For example, most (93%) unaided smoking cessation attempts fail in first week and 95% of lapses lead to a relapse to smoking. Stress, craving, alcohol, and proximity to tobacco outlets are some major precipitants. The challenge is how to monitor them in real-time to trigger timely interventions. Again, data analytic challenges arise as each JITAI intervention is optimized for individual behavior and his/her environmental context.

Overarching concerns for these human-centric systems include:

- How are inferences and predictions created, how is uncertainty represented, and how can humans understand the rationale and limits of these systems? For example, a neural net may produce outstanding predictions, but without clear justification or "error bounds"; an explicit model may be more explainable and, although it has reduce prediction performance, may be more "convincing" to a provider or patient.
- Ideally these systems should learn over time, but there is not a single loop in most cases, but rather multiple loops at different levels of reaction time and generality. "Inner loops" typically being more traditional real-time control loops (e.g., technology assessing and adjusting prosthetics in real-time based on use), "mid-tier loops" being more human-controlled loops (e.g.,

adjustment of prosthetics settings by physician) and also much slower or less frequent, and "outer loops" producing generalized knowledge (e.g. likely outcomes for different types of prosthetics)

3.0 Supporting New Health and Healthcare Paradigms

Advances in smart and pervasive health research create opportunities to support new paradigms in healthcare delivery and to address persistent healthcare disparities in the US. Specifically, these new holistic and human-centric approaches can provide the foundation for accountable and patient-centric care and for addressing systemic healthcare disparities based in inequities in environments, resources, education and access to healthcare.

Value-Based Treatment Plans: Currently, much of medicine is moving toward "outcomes-based" approaches to enhancing value – get the "best" outcome for the resources expended. However, this premise assumes there is a global and uniform definition of "value" – but value can be interpreted at many levels – patient value, family value, provider value, health system value, and economic value. For example, end of life is a focus point for value – a time of high expenditure with a perception of high value related to treatments (or lack thereof). What are lacking are tools that model "quality of life" alongside predictive models for healthcare treatments. Value-based treatment plans could also foreground the use of incentives in chronic disease management. For example, systems that coach patients with diabetes to sustain life goals such as outdoor activity and interaction with family members could prove much more effective than the more abstract task of managing glucose levels.

Accountable and Patient-Centered Care

Multi-tiered SMC systems can provide the foundation for accountable and patient-centric care by integrating complementary models for data collection, decision-making and collaborative, coordinated care. Likewise, systems that integrate more holistic information about patients' lives and information about networks of informal caregivers have the potential to shift current models of patient-centered from supporting physician collaboration to literally centering care around a patient's behavioral, social and environmental context.

Decreasing Healthcare Disparities

Different groups have been subject to and suffer from inequities in environments, resources, education, access to health care amongst other challenges, which can all adversely impact health outcomes such as quality of life and life expectancy[2]. These groups include, but are not limited to different racial, ethnic and cultural groups, low socioeconomic status, low resource, women, LGBTQI, people who have physical and cognitive and sensory impairments, older adults, veterans, rural populations, etc. Smart and pervasive healthcare systems have the potential to reduce disparities by creating systems that are situated in specific socio-economic contexts, for example providing nutritional and exercise guidance in the context of local resources and economic constraints. Moreover, pervasive rates of smartphone use create a more equitable playing field for access to mobile sensing and healthcare expertise.

---

[2] "Health and Social Conditions of the Poorest Versus Wealthiest ... - NCBI." 17 Nov. 2016, https://www.ncbi.nlm.nih.gov/pubmed/27854531. Accessed 6 Dec. 2016.

4.0 Cross Cutting Issues and Barriers to Success

Despite the promise and potential of smart and pervasive health capabilities, there are nonetheless several critical issues to address as possible barriers to success for the research community and for the long-term goal of providing meaningful health outcomes.

4.1 Infrastructure for Health

Sensing and actuation form the basis for smart and connected health research. More and more wellness and disease specific studies are being conducted in homes and assisted living facilities. However, many of these research efforts build one-off platforms. This approach is very inefficient and subject to many errors. What is required is common infrastructure that is easily modifiable for different health studies. This common infrastructure should include, at least, sensing and actuation libraries with associated software processing, networking and communication capabilities covering both local communication and communications to the cloud, cloud processing support that includes databases and analytics, user feedback mechanisms, and runtime monitoring to notify researchers as soon as any problems with the system or associated data recording experiences any problems. Having such a platform will also improve the robustness of the system as it is used in many different studies.

Another issue is that studies for similar issues, e.g., those studying depression, cannot easily share and compare results. Standards for data descriptions, reporting, privacy, and availability are necessary so that researchers can utilize results from prior studies and improve the results in subsequent studies.

4.2 Safety, Regulation, Security and Privacy

Digital health technologies and computation are now permeating through all of healthcare – wireless nursing stations, insulin pumps, implantable medical devices, home blood pressure monitoring systems, health records, etc. To inform the design of these system and to inform regulation, a comprehensive for approach to cybersecurity in health needs to be considered. Many of the security and privacy vulnerabilities found in traditional computing technologies both apply in the health setting, but more worrisome, they are amplified through the significant use of legacy systems, out of date operating systems, and many entry points for malicious activity. Security and privacy health is also about risk management and research is needed in how we balance the risk and benefits of these digital systems. Best practices in security and privacy also need to be re-considered for providers, patients, and developers given the workflows in healthcare.

4.3 Collaboration Challenges: data, expertise, personnel, and interventions

While there is a growing and improving collaboration between technologists and healthcare professionals, paradigm mismatches still exist. Some of the barriers to successful collaboration include:

- Disciplinary Silos
- Training of personnel
- Accessing expertise
- Availability of data
- Other issues in getting started

Enormous amounts of medical data have been collected in the past. Access to that data is often very limited, partly due to how the studies specified policies for their Institutional Review Boards (IRBs). Data formats and use of English sentences for annotations also complicate the ability to reuse the data. With new data mining, machine learning, and natural language processing capabilities, having access to this past data would be invaluable. Policy and data formatting issues need to be resolved to accelerate the pace of discovery and innovation in this field.

In addition, increased training of multidisciplinary personnel and access to appropriate expertise are critical to break down existing disciplinary silos to better facilitate closer interaction between researchers and practitioners of different backgrounds, education, and experiences. Finally, identification of mechanisms to enable collaboration, such as face-to-face interaction, special theme-based workshops, scheduled meetings, etc. is critical to facilitate inter-disciplinary research.

4.4 Healthcare Disparities

There are acknowledged challenges and tradeoffs in addressing healthcare disparities in smart and pervasive health research programs. For example, research focused on new technology innovation may have limited access and resources to simultaneously engage with diverse population groups. However, it is important for any research to appropriately integrate new information and long-standing knowledge of target populations. The assessment of smart and pervasive health research against intended uses and patient target populations is part of the intellectual merit of the work. Simultaneously, any work that does focus on specific communities or population engagement should adopt and publish appropriate methodologies (e.g., user and population-centered design principles; community-based participatory design; action-based research[3]) on how they effectively engage these groups in the research process to identify their values, assess the potential of technical approaches, and address their concerns in any technology development or other research initiative. Due to the complex and nonlinear nature of working with human populations, robust study designs should also acknowledge how the research process may change if and when different groups and outcomes are included (or discluded) in the process. Additionally, researchers must acknowledge what the limitations are of known missing data regarding target populations.

Research programs should have a commitment to engage with target populations with health disparities from conception of ideas through design, deployment, and evaluation. We encourage funding solicitations to require appropriate methods for design, deployment, and continued engagement with intended target populations as part of the *Intellectual Merit* or *Specific Aims*, being acutely aware of groups who have been and are affected by disparities. It must be kept in mind that disparities in economics, education and other areas can all result in health disparities. We further note that including populations with health disparities contributes to the science itself, because including these communities impacts the generalizability, usability, and appropriation of these systems.

[3] (2011, July 1). The relationship of action research to human ... - ACM Digital Library. Retrieved December 6, 2016, from http://dl.acm.org/citation.cfm?id=1993065

Proposed wording on solicitations to address and implore this topic: All research proposals should include in their Intellectual Merit/Specific Aims appropriate methods for design, deployment, and continued engagement with populations with health disparities.

Proposals can specifically address disparities in relation to identifying:

- Access to resources
- Inequities within, among, and between communities
- Needs, designing, building, deploying, evaluating, or meta analysis of various data sets from sociotechnical interventions
- Methods to engage people with health disparities (e.g., community-based participatory design[4], health activism[5]) throughout the research process
- Technical approaches to navigate different resources (e.g., rural communities with limited internet connectivity)

Advocating for including populations at risk for healthcare disparities is necessary but not sufficient. The research community needs (1) resources, (2) expertise, and (3) methods to ensure the contributions generated through this broadening research initiative can improve science through increasing generalizability and thus the impact to science. In terms of resources, the community needs mechanisms and partnership models to sustain engagement throughout the research process and between grant funding mechanisms. Regarding expertise, the computing community needs mechanisms - for researchers at in all career levels - to facilitate connecting with researchers who have valuable knowledge about populations associated with healthcare disparities. Finally, the computing research community needs to combine the needs of resources and expertise to adopt, appropriate, and build methods to design, build, and study sociotechnical systems to address health disparities.

5.0 Recommendations

In addition to pursuing the research and development paths described in this paper, we want to call out these specific recommendations for the successful implementation of an impactful smart and pervasive health agenda that results in improved health outcomes, cost-effective healthcare delivery, and improved patient satisfaction, engagement and utilization:

1. New health data platforms are needed to support holistic and unconventional data streams. Current industry platforms under development are either proprietary or implement a limited range of physiological data fields, severely limiting the platforms' utility beyond their immediate application goals. New platforms need to accept data produced from both internal and external sources (including both wearables and traditional health monitors) with varying levels of noise, fidelity, and sampling rates. These platforms need to support data triage and inspection capabilities to help ensure that data for an individual is being representing appropriately in statistical models. Data quality can easily vary from person to person and device to device

[4] (2015, July 30). Integrating community-based participatory research and informatics .... Retrieved December 6, 2016, from https://www.ncbi.nlm.nih.gov/pubmed/26228766
[5] (2012, May 5). Health Promotion as Activism: Building Community Capacity to Effect .... Retrieved December 6, 2016, from https://pdfs.semanticscholar.org/8617/5302ec9f77df7eedd55c62033dced4efea17.pdf

2. As mobile 24/7 technology is brought to bear on personalized human centric care, new efforts are required that fundamentally address human-in-the-loop participation based on principled psychological, behavioral, and physiological models which are exposed as explicit components of feedback loops. New multidisciplinary research efforts among control theorists, interaction designers, and health professionals are recommended.
3. Rapid progress in infrastructure to support human centric healthcare can only occur if a standardized, composable architecture for integrating devices and data is developed. Further, instantiations of this architecture must support privacy, security, dependability, safety, reusability, and be useable by non-experts. Significant new research efforts in achieving these goals are recommended.
4. Mobile technology offers unprecedented opportunities to improve prevention of chronic diseases. Realizing this promise requires development of new models and theories for just-in-time intervention. In particular, the timing, content, and delivery modalities need to be optimized so as to effect sustained behavior change in the long-term.
5. With the proliferation of healthcare apps and in-home technology, individuals are becoming more involved with their own health decisions, especially in rural communities. Due to many confounding factors, possible simultaneous treatments, and lack of medical expertise of patients there is a great potential for unsafe actions due to conflicts in treatments, including drug-drug interactions. New population-based solutions must be developed to create sound, evidence-based methodologies and to ensure safety in human centric systems.
6. Substantial and sustained collaboration is needed between healthcare practitioners and smart and pervasive health researchers to fully understand the potential of these approaches for new paradigms of healthcare delivery and to identify and address barriers to successful implementation.
7. We encourage the funding agencies to consider health disparities as an integral consideration in their scientific research portfolios and agendas, and to encourage appropriate methods for design, deployment, and continued engagement for the intended target populations, being acutely aware of groups who have been and are affected by disparities.

This material is based upon work supported by the National Science Foundation under Grant No. 1136993. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Smart and Pervasive Health Workshop Participants

Mooi Choo	Chuah	Lehigh
Rumi	Chunara	NYU
Kay	Connelly	Indiana
Richard	Conroy	NIBIB/NIH
Sandra	Corbett	CRA
Munmun	De Choudhury	Georgia Tech
Anind	Dey	CMU
Khari	Douglas	CCC
Ann	Drobnis	CCC
Noemie	Elhaddad	Columbia
Emre	Ertin	Ohio State
James	Fogarty	U. Washington
Carl	Gunter	UIUC
Greg	Hager	JHU
Brad	Hesse	NIH/NCI
Zach	Ives	UPenn
Sabrina	Jacob	CRA
Roozbeh	Jafari	Texas AM
Emil	Jovanov	Alabama
Santosh	Kumar	Memphis
John	Lach	UVA
Insup	Lee	Penn
Scott	Levin	JHU
Ming	Lin	UNC

Beth	Linas	NSF
Ben	Lok	Univ. of Florida
Lena	Mamykina	Columbia
Beth	Mynatt	Georgia Tech
Wendy	Nilsen	NSF
Shwetak	Patel	UW
Christian	Pollebauer	Notre Dame
James	Rehg	Georgia Tech
Laurel	Riek	UCSD
Brian	Scassellati	Yale
Julie	Shah	MIT
Lea	Shanley	South BD Hub
Ben	Shneiderman	UMD
Katie	Siek	Indiana
Ida	Sim	UCSF
Marjorie	Skubic	Missouri
Sylvia	Spengler	NSF
Bonnie	Spring	Northwestern
Donna	Spruijt-Metz	NSF-NIH on development
Miho	Tanaka	VA
Greg	Welch	UCF
Dana	Wolff-Hughes	NIH
Helen	Wright	CCC