

**COMMUNITY
REFERENCE MEETING:
HEALTHCARE**

REPORT

May 2021

Introduction

From data analytics to policy, from patient empowerment to behavior support systems, artificial intelligence (AI) is significantly changing the healthcare field for all those involved. AI-driven technologies are changing the way clinical providers make decisions and how patients find and engage with their own health. AI plays a key role in health and clinical decision support as it delivers data to professionals and patients to aid in diagnosing, treatment planning, patient self-support, distance medicine and population health management. AI has already been successfully applied to improve the speed and accuracy in the use of diagnostics, give practitioners faster and easier access to more knowledge, and enable remote monitoring and patient empowerment through self-care. This will all require bringing new activities and skills into the sector, and it will change healthcare practice and education. This requires a multi-disciplinary approach to research, development, implementation and evaluation of the impact of AI and autonomous systems in healthcare and medicine.

In this first Community Reference Meeting, WASP-HS (the Wallenberg Artificial Intelligence, Autonomous Systems and Software Program – Humanities and Society) brings together Swedish researchers, practitioners and other interested parties, to discuss how the rapid evolution of the technology and public awareness requires impacts research and practice. Following an introductory note on human-centric care in the age of digital by Dr Maja Fjaestad, State Secretary to the Minister for Health and Social Affairs of Sweden, participants discussed how to shape the direction of deployment and use of AI in healthcare in four different application areas. Participants in the meeting included researchers from most Swedish universities, industry, and national and regional governments, as well as general public and international organizations.

The discussions highlighted the importance of social and democratic principles – at the core of Swedish society and political tradition – for the development and use of AI in healthcare and medicine, including responsibility, participation, inclusion and diversity, grounded on a fundamental respect for human agency and self-determination. A research roadmap that ensures alignment with these principles should include efforts in

- Infrastructures, including structured information models in our healthcare systems, and laws that enable 1) multidisciplinary research providing new evidence-based knowledge based on shared data across organizations, and 2) co-creation of new AI systems across academia, healthcare organizations, industry, public authorities, citizens and NGOs.
- Methodology and instruments that allow citizens and other stakeholders to engage in an AI system design process.
- Multidisciplinary research on individuals' identity in an AI future, their relationship to their data and health trajectories, to healthcare and society, to AI systems, privacy, and legislation about the use of AI systems.
- Data governance, including theory and practice on data aggregation, translation and harmonization, considering the role and aims of different actors
- A sound legal and ethical scaffolding framework to ensure trust and deal with issues of liability and responsibility.
- The need and instruments to question and contest the idea that prediction is only possible if enough data is available.

WASP-HS Community Reference Meetings (CRMs)

CRMs are aimed at helping public and private organizations in Sweden with challenges and questions regarding their interests, as well as developments within WASP-HS. This is done to identify opportunities for collaboration between different sectors.

Harmonizing Diverse Patient Data for Person-Centered Care

Main Challenges

- Data harmonizing awakes questions about what is potentially lost when subjective, embodied, data are translated into computational format;
- Lessons for AI can be learned from existing research methods applied to diverse data which have been harmonized in the theoretical and clinical practices of medicine;
- Care should be taken to address ethical and ontological issues that arise when data collected for one purpose are used for another;
- Multidisciplinary approaches are important in relation to decision-making, accountability and responsibility in data translations and decisions made with the help of AI.

Combining and translating different data sets – of different forms (written, visual, embodied), and in different locations (private companies vs. medical institutions, different municipalities) – is a challenge which is often addressed through the concept of harmonizing. This practice raises questions about what is potentially lost when subjective, embodied, data are translated into computational form, and when data travel, or are shared in-between locations and knowledge communities, and what can actually be derived from the data. Take, for example, the use of wearable watches for noticing epileptic seizures. These watches can acquire knowledge of biomarkers and vital parameters that occur just before an epileptic seizure and can provide patterns of such indicators. However, it can be difficult to know if these indicators actually point to a pathology (epilepsy) or if they are caused by other (albeit perhaps related) reasons, such as anxiety due to anticipating that the seizure that is about to occur. This leads to difficulties regarding how to use the data in medical practices.

While some practices of standardizing and translating data can be seen as problematic, for some of the reasons mentioned above, it should be emphasized that the medical community has professional experience in combining different datasets, and that enabling AI to do similar things is a matter of adjusting already existing translation and standardization methods. Diverse data have been harmonized for years in the theoretical and clinical practices of medicine.

There are also more ontological questions about what happens to the data itself in processes of translation and automation. The practices of processing data – for example, documentation and standardization – are part of the onto-epistemological production of data, and this raises particular questions in relation to the increased use of machine learning practices in processing data and about how to think and talk about data when such processing is automated.

A final interesting point that the harmonization of data articulates is the relationship between the aim of gathering data and how data ends up being used, which is particularly relevant when data gathered for commercial reasons are used in a medical context, in relation to a particular pathology. In addition to the ethical issues this raises, it is also important to ask if something happens (and what that something is) when the use of data shifts from what the original goal of collection was.

Initiatives such as WASP-HS are important because voices from the humanities and social sciences often come with questions about aspects of the ethics of data-sharing otherwise missed: relations to surveillance, privacy, and so on. Domain experts, for example, medical practitioners might not always feel comfortable speaking to such ethical issues, and cross-disciplinary conversations and collaborations can be helpful in aiding such discussions.

A multidisciplinary approach can also be important in the implementation of new technologies in healthcare, as this can require long and in-depth engagement with the field of research. For example, fieldwork was brought up as a potentially fruitful method for humanities and social science scholars to access the daily details of implementation practices. This is a way of working that medical experts also recognize as important. It speaks to the need to “go around and talk to people” in order to determine what sort of technological development is actually needed.

Multi-disciplinarity is, likewise, important in relation to decision-making, accountability, and responsibility. The question of who is responsible for both data translations, and decisions made with the help of AI, needs ethical and legal attention by scholars from wide-ranging fields of research.

Social Robotics and Trustworthy Human-Robot Interaction in Healthcare

Main Challenges

- Human oversight and input from domain experts is a key requirement in human-robot interaction, especially in socially assistive settings, making sure that the robot does not undermine human autonomy.
- Expectation management: A challenge of having robots with human-like traits is that we set expectations on what the robot can deliver.
- Multi-disciplinary perspectives are crucial to understanding users and their vulnerabilities, which are necessary to build truly human-centric social robot interaction in health care.
- Multidisciplinarity helps to reveal characteristics in the interaction that may be unseen otherwise.
- Involving users in the design process is an important aspect in user-centric social robotics.
- The trust issue depends on if we can deliver something that is useful to the user

AI is transforming healthcare by reducing, and preventing, invasive procedures and bringing healthcare into patients' homes, and further closer to their communities. As such, this requires us to go beyond individual disciplines to think much more closely about interdisciplinary research. This is especially important for robots in socially assistive roles, where they provide both physical and social support to humans. From robot-assisted diagnosis of women's depression around childbirth to social robots that positively influence human behavior, and robotic platforms such as telepresence robotics for social interaction between people to robot pets in elderly patients with dementia, robots (representing applied artificial intelligence) will become an important part of people's everyday lives. However, in all these scenarios where we have social interaction including robots and users, multidisciplinary perspectives are key for revealing characteristics in the interaction that may be unseen otherwise. For example, from a cognitive science perspective, social interaction is not only reacting to what we see but also inferring what goes on in people's minds, such that we are acting in line with the expectations of our interaction partner. Therefore, (especially) when we utilize human-like design cues in social robots, we create expectations regarding the skills of robots, which can negatively influence the dynamics of trust when the robot inevitably fails to 'live' up to those expectations.

Questions such as how to involve patients and clinicians in the design of robots, to what extent the robot can act autonomously, how to measure the quality of social interactions between people and how robots can act as mediators in human-human interactions, what the overall purpose of the robot and how can it help the user, ought to be explored from a multidisciplinary lens. Furthermore, multidisciplinary approaches go hand in hand with human-centric care visions for improving individual and societal well-being. In other words, the design, development and evaluation of AI research and technologies should follow

the principles of safety, accountability, diversity, fairness, and societal well-being, among others, to be fully human-centric and trustworthy. Two key requirements are (1) human oversight and domain expert input, making sure that the robot does not undermine human autonomy, and (2) explainability, making sure that the robot's behavior is understandable and predictable by humans. An important component in this is to manage expectations of what the robot can deliver.

The COVID-19 pandemic has boosted new multidisciplinary insights, particularly in the way that virtual worlds could be integrated into user experiences, but also in unveiling new research questions and challenges. For example, how were robots used before and how they will be used after the pandemic? And, how are robots used now when we do not have as many social interactions with humans?

As we move forward to the far-reaching socio-technological transition, new technologies are developed in leaps and bounds. A case in point is how robots are projected to detect cognitive decline from human behavioral signals without performing intrusive samples of spinal fluids, or how robotic arms enable independent eating for elderly people with limited mobility. Nevertheless, there is one clear thing: technologies have and will continue to have major effects on our society. As such, multidisciplinary approaches are crucial for understanding the challenges and impacts of AI in healthcare and think about design principles for responsible AI.

Citizen Perspectives on Human-Centered AI for Illness Prevention and Health Promotion

Main Challenges

- AI systems have the potential to enhance participation in society for people with disabilities, who are now to large part excluded, as well as to empower individuals in preventing illness and managing their health.
- Co-design and participatory design is key in AI systems development, in which patients, citizens, healthcare professionals, medical experts, researchers, and industry are engaged.
- Embedding AI in healthcare systems can push Swedish healthcare to be more proactive and predictive. However, this requires research to be tied to clinical practice.
- Multidisciplinary research is very important for addressing issues such as the identity of individuals in an AI future, the relationship to their data and health trajectories, the relationship to healthcare and society, and privacy and legislation regarding the use of AI systems.

AI is increasingly applied in systems aimed to support individuals before one seeks contact with healthcare, and also in the first line of care. Moreover, AI is increasingly applied in technology-based interventions and aids for compensating for decreased function and ability. However, there are multiple challenges related to achieving empowerment, enhanced capability, and human-AI collaboration for achieving health-related goals in an everyday context. While there is a substantial profit to be made in the health sector driving the development of commercial applications, Swedish healthcare is primarily reactive and research is lagging behind.

We took, as a starting point in our dialogue, the vision of human-centric AI systems collaborating with humans, enhancing human capabilities, and empowering humans to achieve their health-related goals [1].

However, a major challenge is how to empower the individual in an AI-induced society. AI systems may enhance participation in society for people with disabilities, who are now to a large part excluded. These systems, however, may also hinder participation. The individual can be engaged in designing AI systems and contribute to innovation, provided there are accessible technology-based tools for this. Individuals may also become empowered to take control over their health through person-adapted systems that can be used in everyday life and outside of the healthcare system, but these individuals may not want to share data or information. How can individuals, who are not yet patients, use predictive systems for preventing disease and injuries? Can an AI system keep a secret? What will privacy mean, and what will it mean to be human, in a future AI society?

The main challenge identified was how to connect patients and citizens, healthcare, academia, and industry in co-designing and developing AI-based systems, which could push Swedish healthcare to become more proactive and predictive, in order to prevent illness and intervene earlier in the progression of diseases. This was seen as instrumental in addressing the related challenges: to advance evidence-based knowledge, and to apply AI technology. To apply AI technology requires that developers understand what the problems in healthcare and the needs of the patient and citizens are and that healthcare professionals, and patients and citizens, understand and engage in the question of how AI technology can be applied -- and its limitations.

A related challenge is adapting legislation to new ways of accessing, sharing, and using data. This is important for innovation as well as research purposes, to advance AI technology and evidence-based medical knowledge, and for the industry to develop new AI systems. Moreover, in order to develop new knowledge and AI systems for managing rare diseases, it is necessary to generate sufficient amounts of evidence across healthcare-providing regions, on a national or Scandinavian and international level, rather than a regional level.

[1] A. Nowak, P. Lukowicz and P. Horodecki, "Assessing Artificial Intelligence for Humanity: Will AI be the Our Biggest Ever Advance? or the Biggest Threat [Opinion]," in *IEEE Technology and Society Magazine*, vol. 37, no. 4, pp. 26-34, Dec. 2018, doi:10.1109/MTS.2018.2876105. <https://ieeexplore.ieee.org/document/8558761>

Data Work in Biomedical AI: the Hidden Challenges of Data, Pre-Training, and Ground Truths

Main Challenges

- AI systems in biomedical research promise a knowledge revolution, but we need to move forward with caution and not have overconfidence in the abilities of the systems. Healthcare and health research cannot be built on AI systems alone but requires a deep understanding of the phenomena and data in question.
- The use of AI in biomedicine thus requires interdisciplinary collaboration, which demands good communication and understanding of the challenges of biomedicine, computer science, and statistics.
- AI systems are always dependent on data of different sorts, and that data can in itself have a multitude of challenges. In biomedicine the challenges of the datasets are multiple and complex, ranging from the confusion of subjective data with objective data, to differences in data contexts, and the use of historical data to predict current trends.
- The travel of pretrained structures and data sets between domains, areas, and contexts leads to a number of challenges that need to be understood in more depth.

AI systems are providing new possibilities for biomedical research by processing and analyzing a diverse set of data. The use of these new systems gives the possibility to analyze and work with multimodal data. For example, by drawing together data stemming from questionnaires, medical diagnoses, genetics, or different types of medical imaging.

These systems create possibilities for biomedical research that depend on interdisciplinary collaboration between domain experts in different fields of biomedicine, but these interdisciplinary collaborations are also challenging as they depend on people from very different knowledge domains that speak different analytical languages.

These types of AI-supported analyses create an analytical revolution in biomedical science, but they also lead to risks of overconfidence in the systems. In multiple applications, the use of AI systems have been shown to work well on training data, but when new data are introduced they work less well, due to overfitting the AI model to the training data. Another challenge is overconfidence in the predictive power of the systems. There is furthermore a big risk that handling of different groups becomes biased, due to the mathematical impossibility of treating different groups fairly [2].

These systems lead to particular challenges when it comes to handling data in a responsible and ethical way. For example, when we train AI systems on biased data we get biased pre-trained structures that are potentially reused in very different places, contexts, and times. This can be seen when systems travel between contexts, such as, for instance, using AI systems for self-driving cars that are trained in a US context in a Korean context—with a different language, textual representations, and traffic patterns. A related challenge is that AI systems are always working on historical data. The historical data set might contain patterns that are not valid in the current situation, and can thus the system can make predictions and inferences that are related to a situation that seems to be the same at first glance—but in reality, is quite different.

The use and travel of pre-trained structures between different contexts and applications also pose challenges for interpretation and validity in treating very diverse phenomena as the same thing. For example, when systems designed for predicting earthquakes are used to predict crimes, are there unintended consequences in how the phenomenon is understood and treated?

The multimodality of data also leads to challenges in that data are used without deeper knowledge about the characteristics of the data. For instance, subjective data—that depend on human interpretation and judgment—but that looks quantitative—such as a doctor’s diagnosis that are coded according to the ICD standard—can sometimes become treated as objective data by the creators of the AI model. In such cases, subjective data can easily become treated as objective data, with overconfidence in the results from the AI system. This poses a big risk in the multimodal analyses of subjective data.

The current data hunger in the use of AI systems might also lead to the use of suboptimal and bad quality data. Sometimes data are simulated to fill in gaps in the dataset, but these simulations run the risk of pushing the system further from the real context.

In sum, AI systems show great promise for creating new knowledge in biomedical science, but there are many risks with the diverse and multimodal datasets that are used in biomedicine that can lead to overconfidence, incorrect analyses, and in the end risks for patients. It seems of particular importance that we do not over-automate biomedical research so that we end up with a system that automates “crashes”—like the automated trading systems in finance.

[2] D. Sumpter, “Impossibly Unbiased,” in *Outnumbered: From Facebook and Google to Fake News and Filter-Bubbles – the Algorithms That Control Our Lives* (London: Bloomsbury Sigma, 2018), 59–70.

WASP—HS

The vision of the Wallenberg Artificial Intelligence, Autonomous Systems and Software Program – Humanities and Society (WASP-HS) is to realize excellent research and develop competence on the opportunities and challenges of artificial intelligence and autonomous systems with a strong investment in research in humanities and social science.

The WASP-HS program is planned to run 2019 – 2028 and will form an independent and parallel program to WASP, The Wallenberg Artificial Intelligence, Autonomous Systems and Software Program.

Request or more details: contact@wasp-hs.org

Authors

Maria Arnelid, Department of Thematic Studies (TEMA), Linköping University
Magnus Boman, Division of Software and Computer systems, Stockholm University
Andreas Brännström, Department of Computing Science, Umeå University
Erik Campano, Department of Informatics, Umeå University
Ginevra Castellano, Department of Information Technology, Uppsala University
Virginia Dignum, Department of Computing Science, Umeå University
Katherine Harrison, Department of Thematic Studies (TEMA), Linköping University
Erica Johnson, Department of Thematic Studies (TEMA), Linköping University
Francis Lee, Division of Science, Technology, and Society, Chalmers University of Technology
Helena Lindgren, Department of Computing Science, Umeå University
Ingeborg Nilsson, Department of Community Medicine and Rehabilitation, Umeå University
Alicja Ostrowska, Division of Science, Technology, and Society, Chalmers University of Technology
Isabel García Velázquez, Department of Thematic Studies (TEMA), Linköping University
Katie Winkle, Division of Robotics, Perception and Learning, KTH Royal Institute of Technology

How to cite this report

WASP-HS. Community Reference Meeting: Healthcare. Report. May 2021.

Design and typesetting

Tatyana Sarayeva