Meeting report

Expert workshop report organised by SATW and SAMW on "Artificial Intelligence in medicine - can it be trusted?"

Thursday, 25.10.2018, 13.15 – 16.45

SATW, St. Annagasse 18, 8001 Zurich

Background

The constantly increasing computing power and the exponential increase in available digital data have led to a renewal of the hopes to develop useful artificial intelligence (AI). Substantial progress has been achieved in fields such as image classification and natural language processing. The combination of these algorithms with hardware connected to the internet of things, with social media, with voice interfaces, and more, is producing useful tools in many domains, including health and healthcare. Image recognition can be used to identify potentially cancerous skin lesions with performances on par with medical experts. Natural language processing helps unlocking and summarizing mountains of clinical information and connect them with scientific knowledge to recommend personalized therapeutic approaches.

There is a growing expectation that this new generation of AI will outperform and even replace humans in many professional domains, also in healthcare. While this looks far-fetched, it is already clear that AI will fuel important shifts in the production of healthcare, redefining roles and tasks of professionals, patients, and their families. Concurrently, due to many demographic and socio-economic factors, healthcare systems in most countries - rich and poor - are in danger. Innovative ways to reduce waste and errors, improve collaboration, and increase safety, quality, and efficiency are required. AI is top on the list of new tools which could provide these much-needed solutions.

Given the hopes, the hype, and the significant societal implications of the development of AI in medicine - all still poorly understood - there is a clear need for education and debate involving not only professionals but also a wider, general-public audience.

Digitization - main topic of the Swiss Academies

In the context of the reflections led by the Swiss Academies on the “impact of the digitization on economy and society”, the Swiss Academy of Engineering Sciences SATW, in collaboration with the Swiss Academy of Medical Sciences SAMW, organised workshops and public events to address issues related to the development and application of AI in medicine.

Themes of wide interest - including societal, political as well as economic issues – were discussed in the present expert workshop as well as two public events - one in Geneva on Sunday 7 October 2018 as part of the salon planète santé and an evening event in the Careum in Zurich on Thursday 25 October 2018 – the TecToday on “Dr. KI – Arzt Ihres Vertrauens”.

About this workshop

Many AI systems are data-hungry applications, launching a hunt for data with challenges to societal values such as privacy, ownership and sovereignty. When it comes to their reliability and trustworthiness, many of these tools are poorly understandable. Their underlying algorithms are unable to explain their behaviour in human-comprehensible terms. Furthermore, being as value-neutral as any technology, AI's potential for malicious use is obvious. This is particularly worrisome when dealing with life-sustaining devices or life-threatening decisions.

In this workshop selected participants with proven expertise in AI and its application in medicine will discuss current developments and their need for robustness of applications, control mechanisms and regulations, ethical guidelines and norms. Keynotes introduce certain topics to initiate the discussion. Driven by current and future challenges in the health sector with a need for action, examples of AI systems in medicine are presented and solutions required in the future will be discussed.

Additionally, general and solutions-specific challenges and risks to be addressed when deploying AI in medical applications are discussed. These include trust in AI systems and the need for explainability of their recommendations. Finally, the current regulatory requirements for software as medical products is presented. Liability for recommendations and potential risks regarding bias and cyber security are further aspects to be addressed.

The goals of this workshop are to create an overview on the current state-of-the-art of different aspects in context with AI applications in medicine, identify and discuss concrete fields of actions to address specific risks, to use the promised potentials and to initiate recommendations for politicians, regulators, practitioners and researchers.

Agenda

13:00 - 13:15	**Registration and welcome coffee**
13:15 - 13:30	**Welcome and introduction** Dr. Rolf Hügli, Managing Director SATW Prof. Thomas Szucs and Prof. Urs A. Meyer, University of Basel
13:30 - 13:45	**Memetic pattern based algorithm to diagnose or exclude coronary artery disease** Keynote by Prof. Michael Zellweger, University Hospital of Basel
13:45 - 14:30	**What is Missing in Switzerland for the AI Wave in Medicine** Keynote by Prof. Philippe Cattin, University of Basel Q&A
14:30 - 15:00	Coffee break
15:00 - 15:15	**ETAPH - Establishing Trust in Artificial Intelligence Powered Health Systems** Keynote by Dr. Verena Pfeiffer, ISPM - UNIBE and Dr. Aitana Lebrand, SIB
15:15 - 16:00	**Software as Medical Device: The Regulatory Requirements** Keynote by Prof. Christian Johner, Johner Institut GmbH Q&A
16:00 - 16:45	**Discussion and final synthesis**
From 16:45	**Apéro**

Welcome and introduction

Rolf Hügli welcomes the workshop participants and thanks them for their attendance. For the activities of the academies it is important to get inputs from experts by means of such workshops.

Urs A. Meyer introduces the topic with a few slides. Demographic changes require new solutions. He states that AI in medicine is now and there are already numerous applications in use. Some of them perform as good as doctors in specific tasks. In this workshop the potential as well as the limits of AI in medicine should be discussed, too.

Memetic pattern based algorithm to diagnose or exclude coronary artery disease

Prof. Michael Zellweger, University Hospital of Basel

Michael Zellweger presented a memetic pattern-based algorithm (MPA) to support the diagnosis of coronary artery disease (CAD). Only 16% of the cases where patients complain about chest pain are related to cardio vascular problems. The algorithm aims at determining the probability of CAD based on easily available patient data focussing on a very low rate of false negatives. The algorithm is constituted of a variety of different AI methods.

The comparison of Framingham risk scores of the MPA with the gold standard based on male patients in the age of 65 +/- 10 years is promising. The algorithm was also validated for other patient populations with good results.

The application provides results with high diagnostic accuracy and is already in daily practice. However, the data size is still relatively low and patient screening is not yet feasible. Nevertheless, it has the potential to cover a broad risk-spectrum of patients who need CAD evaluation.

What is Missing in Switzerland for the AI Wave in Medicine

Prof. Philippe Cattin, University of Basel

While in its early days deep learning was claimed to be for people who lack engineering skills to solve a problem properly, today the number of publications is increasing rapidly and deep learning features in more than 50% of medical papers. The number of start-ups transforming healthcare with AI is continuously increasing, too. Furthermore, these days the first medical analysis system to support radiologists was approved by the FDA.

In recent years a paradigm-shift has been ongoing: AI methods are readily available but what is missing is the right kind of data. AI applications will be introduced first in tedious tasks which include a lot of manual work of medical specialists, e.g. segmentation. For such applications the scores between a trained algorithm and field experts already are very similar. A “Hamster race” is ongoing for such challenges to increase the score by a few percent’s e.g. by tweaking some parameters, specially prior to a conference on the topic.

Philippe Cattin’s research focuses on weakly supervised machine learning. In this field algorithms learn e.g. how to cure a brain with a disease. The approach is a step towards explainable AI since it displays areas which are involved with a disease. For such kind of studies, however, data of both

healthy and unhealthy persons are required. This is one of the reasons why he thinks that Switzerland needs a national healthy cohort study with a collection of healthy patient data. The US, the UK, Germany and other nations already have ongoing projects in that field.

Discussion part 1

Medical data and national healthy cohort study:

? What steps are involved in setting up a national healthy cohort study?
? What kind of data should be tracked?
? What are potential funding sources?
? How could the academies support such an endeavour?
- The healthy citizen cohort study should be compatible and complementary to existing projects, thereby supporting **international collaboration**.
- Do we already ask (enough) for the **collection of the public health data** for the benefit of the society? We should probably put more effort in sensitizing the public and provide good arguments for people to share their health data. We probably should maybe do it more aggressively. Up to now the government does not do enough in this respect. The current consent procedures might need to be revised, too.
- Insurance data up to now was mostly ignored. It mainly involves claims data. Why are insurers not involved in these discussions? Specific reasons were not discussed in detail but probably involve conflicts of interests. One common denominator, however, was mentioned to be the decrease of health care costs.
- What **further types of medical data** are available and should be linked? The medical data base of the Swiss Army covers basically the whole male population of Switzerland. Another source is the data from school doctors. Private hospitals host huge data sets, too.
- One main question is **how to use the available data in the best way** and **how to supplement it with cohorts**. The SPHN should be doing that – collect the inputs from different databases and combine them.

Campaigning AI in medicine

- Who would be the best addressee for AI systems?
- Doctors and practitioners could benefit most and their work would become easier and more reliable with the support of such systems.

ETAPH - Establishing Trust in Artificial Intelligence Powered Health Systems

Dr. Verena Pfeiffer, ISPM - UNIBE and Dr. Aitana Lebrand, SIB

Verena Pfeiffer and Aitana Lebrand presented their ETAPH project which was submitted to SPHN. AI has tremendous potential for realising personalized medicine and health. AI applications in the health care sector can roughly be separated into three different areas: clinical setting where they are used e.g. as decision support tools; clinical research and health/wellness apps. Various challenges exist for all of these areas which ultimately all base on trust. This project focuses on enabling trust in AI applications in the clinical setting to realize their full potential in personalized medicine. Its goal is to define a framework for best practice clinical use of AI applications, to

ensure clinical-grade quality, transparent methodologies with validated performance on different patient groups (fairness between e.g. women vs. men; child vs adult), compliance with ethical and legal principles, as well as proper training of end-users with the ultimate goal to define, if needed, the structure and processes of a Swiss AI governance body for healthcare.

These challenges will be tackled by a Swiss-wide working group, bringing together experts in AI, clinical bioinformatics, bioethics, clinical research, legal policy and various medical specialities, as well as physicians and patients representing bodies.

This working group will tackle open questions and provide - amongst others – the following:

- List wishes, needs and constraints in context with AI applications in medicine
- Provide guidelines for the development, testing and use of AI applications
- Set up an ethical code for AI use in clinical settings
- Describe legal implications of AI applications in the clinical setting and their possible solutions
- Answer the question whether a new Swiss governance body on AI is needed

The project aims at both, a political as well as an executive approach.

Further open questions that could be taken into account are:

- Who will pay for the use of AI applications in the clinical setting (the health insurance?). The economic difference between an algorithm and a drug is not necessarily clear to everybody. How is the insurance question to be answered and who pays – specially in case of failure?
- What is the intellectual property on AI and who owns it? The legal link between an algorithm and the data which was used to train the algorithm is not clear to all.

Software as Medical Device: The Regulatory Requirements

Prof. Christian Johner, Johner Institut GmbH

The European medical device regulation (MDR) describes in which context software is classified as a medical device. Generally speaking it depends on the intended use. Is the software applied in the market or is it just a tool? Does it affect the decision of the practitioner or not?

Life cycle development of software is one important requirement. Furthermore, it must be demonstrated that the benefits of the application outbalance the risks. Thus, a thorough risk analysis is demanded which requires sufficient clinical data.

The MDR covers new aspects like mobile platforms and requires instructions for their use. One major request is the documentation of software lifecycle processes. Standards - which may not be available i.e. explicitly defined yet - need to be fulfilled for different requirements. The Johner Institute is working on guidelines for secure applications. This is a mandatory documentation and different standards need to be fulfilled. However, no specific AI requirements are available.

The current classification of software as medical devices is problematic. In the MDR, it is only based on the severity of potential failure but not on the risks. Failure probably is not included in the categorization process. The continuous reduction of the number of notified bodies adds to the

problem. The remaining ones often have an overload of cases resulting in long evaluation times for companies.

Boldly put the MDR does the opposite to the FDA – it kills the innovation process rather than fostering it. In this context the academies and every involved party should campaign for a more liberal approach towards software in the medical field. A post-market evaluation of products is required. There is no lack of rules but a lack of enforcement.

Discussion part 2 and final synthesis

Data

- Available sources should be linked and combined by bringing together the different parties. Collaboration with SPHN is advisable.

Campaigning and recommendations to politics

- Talk to politicians and militia, inform and give recommendations to the government (see ETAPH project)
- Call for a working group as suggested by the project ETAPH to make recommendations for politicians and regulating bodies to foster better approaches balancing the risks and benefits of certain applications. Collaborations with existing certification bodies.
- Updating guidelines on a frequent basis, e.g. quarter-yearly.
- Historically medicine is evidence-based. We should also foster the acceptance of data-based knowledge in the medical domain.

Education and product approval

- A transparent reimbursement model regarding AI applications is required and there is a lack of understanding the economic aspect of AI related to IP questions.
- The black box issues should be tackled involving the problem that work based on DNNs is not peer-reviewable.
- Physicians lack the statistical background to critically assess the performance and output of AI applications. One suggestion was to increase the importance of statistics in the medical education. Another approach is to make the approval of products more reliable such that physicians can rely on the results provided certain boundary conditions.
- Performance before explainability – the latter is not defined in humans, either. Either way, there is an increasing number of papers and publications on explainability of AI.
- Focus on clinical physician tools