Re-humanising automated decision-making

Workshop report from the ADM: Nordic Perspectives research network

Minna Ruckenstein, Stine Lomborg & Sne Scott Hansen

November 2020

UNIVERSITY OF COPENHAGEN

UNIVERSITY OF HELSINKI

LUND UNIVERSITY

DANMARKS FRIE FORSKNINGSFOND
INDEPENDENT RESEARCH FUND DENMARK

Content

Introduction

As the uses of automated decision-making (ADM) are increasing in private companies, governments and public authorities, automation processes are spreading from manufacturing products and services to making important decisions about people's lives. This workshop brought together an interdisciplinary group of scholars and experts to unpack and discuss **the shared challenge of re-humanising ADM**. We define ADM as procedures in which decisions are delegated to a public or corporate entity, which then uses automatically executed human-made decision-making models. Whereas earlier research explores problems arising from ADM systems, their connections to inequalities, lack of transparency and opaque understanding of people and their lives, our aim was to discuss complementary perspectives by focusing on ADM as it relates to infrastructures of everyday lives and shapes our imaginaries of the world, we want to live in.

With the move towards re-humanizing ADM we seek to make visible human forces and ideals that the dominant logics, defined by techno-optimism and political-economic aims of efficiency and optimization, conceal. Re-humanizing is a starting point for exploring the complexities of ADM by establishing the human as a critical and creative agent in human-machine relationships that are emerging in the wake of recent ADM technologies and related discourses. We deliberately focus on ADM rather than artificial intelligence (AI), which is currently the hyped term, because AI invokes connotations of machinic intelligence that operates without human involvement. Even if the cases discussed in the workshop might bring us to the realm of AI, we remain cautious of ascribing human-like autonomy and intentionality to machine-based procedures. Instead, we maintain our focus on how humans are contributors in the design of automated services, as objects of data collection and processing, in making sense of data and making decisions about data, and in implementing decisions. Here, we emphasize the importance of not merely highlighting human involvement in technological processes, but explore how humans are intentionally involved and thereby implicated in such processes. We specifically ask questions about the concrete operations of ADM systems, the divisions of labour between human and machinic operations and the forms of collaboration implied.

To achieve an empirically grounded conversation, the workshop departed from a set of case presentations about ADM that were freely chosen by the presenters. Three of the cases were from Finland, two from Sweden, and three from Denmark. The cases, while representing some of the current research that the members in the network are undertaking, were presented to the attending interdisciplinary group of scholars in descriptive terms, using a shared template to explain the case, why it was chosen as a relevant case of ADM, what ADM does and what questions this raises for the challenge of rehumanizing ADM. The purpose of this kind of guided presentation mode was to offer a bottom-up perspective to current developments, and to build a common ground for interdisciplinary inquiry. The presenters kindly agreed to write-up the case presentations for this report. We hope that the empirical cases offer material to push our thinking to develop themes and concepts that may cut across the ADM landscape.

The workshop proceeded as follows: Minna Ruckenstein (University of Helsinki) introduced the workshop and offered a brief igniting talk about Rehumanizing ADM to frame the case presentations. Selected network participants then presented their cases under the guided framework. To elicit common themes and questions emerging from the cases, workshop participants then worked in small groups, before the workshop was closed by a plenary discussion about the ways forward in researching and advancing the re-humanization of ADM.

The workshop was hosted by Minna Ruckenstein and Tuukka Lehtiniemi from the University of Helsinki.

Case presentation: "Algorithmic Decision-Support for Municipal Caseworker Systems"

Naja Holten Møller, University of Copenhagen

The case explained

What: The case we report is the design process of an algorithmic component for caseworker systems utilizing the Dynamic Condition Response (DCR) graphs tool for decision-support in municipal job placement in Denmark. DCR is a method and a full technology stack covering design, simulation, analysis, documentation and execution of declarative processes (see, dcrsolutions.net.) The nature of casework and decision-making may be understood in terms of the processes of governmental decision-making that are often combinations of stable, predicable sub processes and unpredictable events and changing rules in governmental decisions (Hildebrandt et al., 2020). In our case, the idea is not to design caseworker systems from anew. Instead, the goal for an algorithmic component is that can be integrated in existing commercial case management and workflow tools used widely in Danish local and central government – while being developed through use of participatory techniques bringing caseworkers into the design process.

How: Taking a participatory approach (Møller et al., 2020), we wanted to explore value metrics together with caseworkers as part of a cross-disciplinary research project in an open-ended manner. The project is a collaboration across data scientists, software developers and scholars specializing in ethnographic studies – with two municipalities partners and fieldwork sites in the project. The union representing a large part of the caseworkers in Denmark acted as an observer in the project, along with the organization representing the municipalities. National authorities are already experimenting with use of algorithms for prediction of *individual* risk in job placement. Showing a participatory approach to the design of algorithmic decision-support systems, our process with caseworkers suggested a rather different notion of value for algorithmic decision-support systems in this context.

Who: We conducted the study at a job centre, which serves a mid-sized municipality of approximately 69,000 citizens. In 2019, 5,655 individuals in the municipality were unemployed. Participatory design workshops (n=3) were organized between 2018-2019. These engaged with the relevant caseworkers in job placement offices – in addition to observational studies and in situ interviews (n=9) with caseworkers prior to the workshops. These workshops were for the purpose of making sense of the dataset of 16,000 currently and formerly unemployed individuals, negotiating a shared "value metric," and developing an algorithmic component for the existing case management system. Finally, a prior field study (84h) of job placement in a different municipality

(2015-2016) also formed the background for the participatory design workshops (Møller et al., 2019).

Why did you choose that case?

The cross-disciplinary research project is risky but also represents a real opportunity to broaden knowledge of how design processes can be set up responsibly to include all stakeholders. This research is particularly critical due to the strong interest and push for developing algorithmic decision-support tools for the public sector. The openness and support of our work from the municipal job centre and our industry partners have been instrumental for showcasing a participatory approach.

What does ADM do?

The DCR tool allow caseworkers to see a visual representation of the possible legal routes through job placement: Activities can be assigned roles (e.g. caseworkers with different specialties), indicating that an actor assigned this role can execute the activity:

(1) Within [days] after receiving a notification under sections X-Y, the municipal council shall acknowledge receipt of the notification to the notifier. (2) The municipal council shall inform the notifier under section Y above whether it has initiated an investigation or measures pertaining to the person to whom the notification relates. Notwithstanding the aforesaid, this shall not apply where special circumstances exist.

The DCR tool in this example is used for formalizing the rules that are considered to be needed for case management to be compliant with the law text above. Another potential use is through predicting risk of organizational "delay" in the processes ahead of this delay, and thus to enable better planning.

Case presentation: "Seeking efficiency – Migri's chatbots"

Sonal Makhija, University of Helsinki

The case explained

What: Migri (the Finnish Immigration Service) is a government agency responsible for making decisions in matters related to immigration, Finnish citizenship, asylum and refugee status. In 2019, Migri made a total of 104,000 decisions. The slow processing times for those seeking asylum, applying for citizenship, work or student permits, have been criticized by companies and universities, as they have led to non-EU students missing a term, or even year-long waiting times for work permits. Consequently, Migri has been exploring ADM to speed-up permit processes, including an ability to quickly respond to queries. 'Kamu', the Migri chatbot (meaning pal in Finnish), was designed to reply to day-to-day queries and to reduce customer calls. The goal is to automate all tasks that do not require human input. The goal is that by the end of 2020, 90 % of all interactions would take place via self-service channels. Prior to the introduction of the chatbot, only around 20 % of incoming calls were responded to. After chatbot automation, the share of responded calls already increased to around 80 %.

How: The chatbot can be accessed on the right-hand corner of Migri's website. By clicking on it, a menu of questions appears together with the recommendation of not revealing one's personal identification number. The chatbot takes care of 'routine' questions and issues, and redirects service-seekers to other key government agencies, like Business Finland, and the tax office.

Who: Primary customers of the chatbot are those who apply for work-based/family-based residence permit applications, and students, given the chatbot is only available in English and in Finnish.

Why did you choose that case?

The chatbot was created with the aim of responding to queries quickly and reducing staff workload. Yet, the chatbot is only available in English and in Finnish, and responds to routine questions that are already available in the form of FAQs on the website. Moreover, it only answers, if questions are broken into short sentences, as it picks on keywords and provides a menu of options based on those keywords. Thus, the chatbot almost requires adapting one's use of language, so that it is understood by the chatbot. It might also require training of the chatbot by going through customer questions and add to the question/answer pool that the chatbot can respond to. The aim of the chatbot was to reduce Migri's staff workload, redirect repetitive queries to the chatbot, and respond to maximum customers. The chatbot filters the queries that human staff has to respond to, and according to Migri, this has reduced their workload and improved customer response time. Here, an important question is whether the chatbot truly responds to queries that customers face, or if customers hit a wall leaving them to redirect their questions to immigrant forums or even giving-up on receiving an official response to their queries. In any case, the queries might need human training or filtering to evolve, since a large part of the chatbot's work entails mechanically drafting questions and responding to them.

My interest lies in understanding how the human and chatbot collaboration works out. Does it often mean humans doing more mechanized work, even though the aim is to reduce routine or repetitive queries that need human attention? Do a majority of customer queries get responded to? More so, how does this change how we adapt to chatbots in our need to get answers to our urgent questions?

What does ADM do?

In terms of ADM, the chatbot is based on a simple decision-tree model and performs a step-by-step process to assist the customer in finding the right information/answer to the question. The main aim of the step-by-step format is to discover the precise question, and break it down into smaller parts to respond to it. Thus, the process moves from broad themes to narrowing it down in order to reach the precise question, branching out and providing various options. For instance, if your question is on processing times of your residence permit application, the chatbot will provide a menu of different types of residence permit-based applications. In case the customer wants to switch the language, or if the customer has landed on the wrong language page, the chatbot easily switches the language to Finnish or English based on customer response. Additionally, based on keywords written by the customer, the chatbot redirects the customer to other agencies such as Vero or Business Finland if the residence permit or entrepreneurship-based residence permits are being processed.

In the Finnish version of the chatbot it is furthermore so that if the customer reaches a dead-end, so to speak, or the chatbot fails to respond to the question, the chatbot offers to redirect the customer to a human colleague even if the conversation switches to English later. This is not available in the English version (as checked on 16.10.2020).

Case presentation: "Eksote's automatic warning system for marginalisation"

Tuukka Lehtiniemi, University of Helsinki

The case explained

What: Eksote is a social and healthcare district in the south-eastern part of Finland, consisting of nine municipalities with about 130 000 citizens in total. Eksote has for a few years investigated machine-learning-based identification and warning systems for risks of individuals experiencing different kinds of problems in their lives, such as forms of marginalization and exclusion. Such risks have been identified for groups such as children, young adults, and customers of mental health services.

How: Risk factors are derived from machine learning models that analyse which variables statistically predict negative "endpoints" defined by social and healthcare professionals. Both the endpoints and potential risk factors are derived from data in social and healthcare registers. The use of registers of other administrative branches, such as education, has also been investigated. In one example involving predicting risks of problems with children and youth, six endpoints were defined by professionals (ranging from substance abuse to child custody) and hundreds of risk factors were identified by the machine learning models (ranging from bad dental health, to parents missing maternal clinic appointments, to an increased number of x-rays to the ankle, foot or wrist). The aim is not only to identify these risk factors in a statistical sense, but to also develop an information system that automatically issues a warning when a high number of risk factors are associated with an individual.

Who: In the system now under development, warning signs are issued for people who are in an existing customer relationship with social and healthcare professionals. Warnings would be displayed for social and healthcare workers who interact with the customer or patient, and would be shown on-screen in the information system used during these interactions. Investigation of the involved risk factors is subject to explicit consent provided by the customer or patient.

Why did you choose that case?

Data studies literature provides more than enough examples of problematic uses of ADM systems, predictive policing systems being one well-known example. From an imaginary informed by this literature, early warning systems for risks carry the potential of enforcing a surveillance-based logic onto care relations, introducing arbitrary punitive outcomes into the social and healthcare system, and exacerbating existing inequalities.

In contrast, the imaginary underlying the development of systems that produce such warning signs seems to be grounded in more optimistic beliefs about data and automation: that data contains warning signs that we should not neglect, that humans either cannot and do not have the time to spot these warning signs without automation, and that providing warning signs early on leads to possibilities for early professional intervention.

It would be easy to conclude that the Eksote case is a problematic and potentially punitive example of an ADM system, especially as it involves potentially vulnerable individuals in risk of marginalization. However, both the fears and the hopes outlined above are speculative in nature; they reflect our assumptions about what happens when risk factors are identified and warnings issued and displayed. Signals brought up by AI-branded systems could be, for instance, deemed untrustworthy by professionals, or simply ignored for other reasons. Alternatively, they could become something that professionals cannot ignore, for fears of making a decision that looks bad in retrospect, which could result in focusing resources according to the machinic logic. Displayed warnings could become a useful support for human professionals in deciding which cases should be given more attention early on. Instead, they could be treated as inputs for a different automated system. The Eksote case, then, highlights that ADM is not just a technical but a socio-technical system, and projecting either imagined fears or imagined hopes into the technical part can be misleading.

The case brings together a number of interconnected aims, for instance: providing early help for people, avoiding serious and drastic consequences for their lives; cost savings for the social and healthcare sector by making possible intervention before cases more difficult and more resource-consuming to resolve; and the societal aim of reducing different forms of exclusion and marginalization. At the face of it, this is a technically straightforward, obviously well-intentioned, but nevertheless ethically complex case of ADM.

What does ADM do?

Identification of risk factors based on machine learning perhaps does not as such qualify as ADM. However, there are ADM qualities in a system that detects risk factors for individuals based on the machine learning models, and issues warnings when risk factors are detected according to the model. Whether this is full ADM, partial decision-aid system or something else depends on what the system becomes in a socio-technical sense; how the ADM-like technologies are connected with social and health care processes and practices.

Case presentation: "Automated decisions on social benefit applications – Trelleborg municipality"

Anne Kaun, Södertörn University

The case explained

What: Since 2017, Trelleborg municipality with roughly 46,000 residents has introduced fully automated decisions on applications for social benefits. This case of automated decision-making is one of the most well-known and most widely discussed in Sweden. Trelleborg municipality prides itself for being at the forefront of automation efforts leading the piloting innovation program from *Rebel to Model* (Rakar, 2018). It often serves as a reference point to explore both algorithmic culture as well as the implications of the digital welfare state (Choroszewicz & Mäihäniemi, 2020).

How: More specifically the system is based on a rather simple decision tree model that cross-checks certain variables with databases by for instance the Tax Agency including income or payments by the state health insurance. All initial applications are processed manually. However, follow-up applications that have to be submitted once a month are processed automatically. One of the most discussed aspects of the introduction of this ADM system is the reduction of civil servants working with social benefit applications from eleven to three case workers.

Who: The ones affected by the ADM system are residents of the municipality that are applying for economic support (*ekonomisk bistånd*) including welfare benefits (*försörjningsstöd*) that either fully or partially cover costs for housing, food, clothing, telephone and internet access. The number of residents who no longer rely on social benefits increased during 2017, the year the automation algorithm was introduced, to 450. The same number of residents moving away from social benefits into other ways of securing an income was around 168 five years earlier. This decrease in beneficiaries is not merely attributed to automation by the municipality and journalists, but to a holistic program of re-integrating long-term unemployed into the job market (SVT, 2018).

Why did you choose that case?

The Trelleborg case has been controversially discussed early on. Critique ranged from its illegal delegation of decisions to algorithmic systems that is not supported by laws for municipalities, to questions of transparency as well as the future of civil servants more generally. I draw on the case to study the process of mundanisation of technologies, i.e. the integration of complex technical systems into our everyday lives that is also based on controversial negotiations between different societal actors, in this case journalists, unions and lawyers.

What does ADM do?

ADM in the context of Trelleborg is a fairly simple decision tree that cross-checks applications with external databases by other public agencies. At the same time, it is contested by different parties whether this is actually fully automated decision-making or not. As the project manager from Trelleborg argues, the final decision whether a person is fit for the job market or not – which ultimately is the decision that needs to be taken – always lies with a civil servant. There are however steps that lead up to this decision and these steps are partly automated. Hence, the municipality prefers to speak of automated decision-support systems, mainly to keep within a legal grey zone that does not allow for fully automated decisions on the municipal level yet (Velkova & Kaun, 2019).

Case presentation: "U-prevent – Personalized Drug Protocol"

Laetitia Tanqueray, Lund University

The case explained

What: Dutch researchers specialized in epidemiology have created and introduced to the public 'U-prevent', a tool to *support* the clinician and patient for personalized vascular medicine. The tool is a calculator, nested in a website, which is framed as user-friendly for both the patient and clinician. The tool predicts the risk of heart attack, or other predictions, or assessments of whether the medication is correct for the patient; its end goal is to calculate the patient's cardiovascular risk and the effect of preventative treatment on the patient. U-prevent has been based on 31 peer-reviewed medical studies, dating between 2011-2019, and conducted by the researchers that have created it.

How: The patient or the clinician assigned to the patient can input the details on a calculator that has been advised for the patient by U-prevent (there are 7 different calculators to choose from). However, the calculators used require very specific details about the patient's health: if the information is not filled out accurately, the tool may not correctly calculate the patient's cardiovascular risk and the effect of preventative treatment on the patient. The tool may use Electronic Medical Record System in order to obtain the relevant data. Once the calculator has some data inputted, it will match the patient's details to previous studies. From the website, it is unclear which studies have been used to calculate this and how representative the data is. Finally, the clinician can review the results with the patient to decide the next step for the patient's health.

Who: The patient does not need to show any signs of vascular issues. This tool can be used for patients of any age depending on the symptoms (usually ranges between 30 and 90 years old) from typically "Europe low risk", "Europe high risk", "Western Europe", "Netherlands" and "North America" only. It is not obvious who it is best suited for, maybe white men (and potentially white women) based on the small selection of 'nationality'/'area' U-prevent offers.

To bear in mind: It is important to contrast U-prevent to a current study about the use of personalized drug protocols. The study, titled "Use of an electronic decision support tool to reduce polypharmacy in elderly people with chronic diseases: cluster randomised controlled trial" (Rieckert et al., 2020) aims to demonstrate how general practitioners benefit from using support tools, i.e. if the tools actually help the patient and addresses her/his needs. The study shows that it is currently inconclusive whether this is the case. It is also key to point out that the study found that the support tool cannot deal with individual patient's needs (which is what U-prevent claims) and that doctors need to input the data accurately, otherwise the tool does not give reliable results. Hence, this study demonstrates how U-prevent should be used with scepticism by the clinician and patient.

Why did you choose that case?

U-prevent is an interesting case as it is currently available to be used by any clinician/patient. The lack of transparency of the calculation should be questioned: what studies did researchers use to train/build the tool? Was the input data to train/build the AI tool based on high quality data or disproportionate data of patients with rare diseases? Is the data helping train the tool despite the data being able to be erased if the patient wishes to do so? I decided to look into this case to offer critical reflection on the degree to which the public healthcare sector is equipped for these kinds of tools created by private entities to be used by the public healthcare sector.

What does ADM do?

The extent of ADM used in U-prevent is uncertain. There is a lack of transparency on how U-prevent calculates, especially regarding the previous studies used to calculate and whether it is weak-AI (very rule-based which thus produces results similar to humans) or whether it is more an in-between AI (informed by human reasoning but somewhat autonomous). Furthermore, usually within the healthcare sector, the literature steers away from the term 'ADM' and instead uses the notion of a "decision support" tool/system. However, in this instance it would seem that U-prevent falls under ADM due to calculating specific data input without much human intervention/understanding (especially if the Electronic Medical Record is used).

Case presentation: Symptom Checkers and Algorithmic Care

Sonja Trifuljesko, University of Helsinki

The case explained

What: Omaolo (eng. MyFeel) Symptom Checkers are web-based apps used to evaluate one's own symptoms or health condition. They are a part of the *Omaolo* digital service channel, which - besides symptom assessments - currently includes three other services: electronic health check-ups, well-being coaching programmes, and service assessments. In addition, there is a separate planning section. *Omaolo* was initially piloted within the national self-care and digital value services project, which ended in autumn 2018. After that, the development and distribution of *Omaolo* was transferred to a state-owned company dedicated to advancing the digitalisation of health and social services called *SoteDigi*, which, from the beginning of October 2020, is known as *DigiFinland*. There are currently 16 symptom checkers on the *Omaolo* website, ranging from lower back pain to cough. Since the end of October 2020, all symptom checkers are also available in English, besides the two national languages of Finland, Finnish and Swedish. Prior to that, however, only the coronavirus symptom checker was accessible in all three languages, and it is on this symptom checker – due to its topicality and popularity - that I focus here. The main purpose of the coronavirus symptom checker has been to assess the possibility of the COVID-19 infection, as well as to advise about preventing the spread of the infection and to determine the necessity for the involvement of health professionals. The coronavirus symptom checker is a CE marked medical device, indicating that it complies with the applicable EU regulation. Like other symptom checkers, the COVID-19 one was produced in cooperation with the Finnish Medical Society *Duodecim*. In addition, the Finnish Institute for Health and Welfare has been involved. Finally, application development and analytics are provided by two Finnish private companies.

How: To use the coronavirus symptom checker, one simply needs to go to the omaolo.fi website and start filling in the e-form. The questions in the e-form are mostly in the yes/no or multiple-choice format, but – depending on how one answers - there might be also some open-ended questions. Prior to using the symptom checker, one needs to accept the *Omaolo* service terms of use and privacy policy, to agree that the service would assess the need for treatment and make recommendations based on the provided information, and – finally – to consent to the use of cookies in the *Omaolo* service. From the end of August, when the Finnish exposure notification mobile app has been released, one can also switch directly from the app to the *Omaolo* website. In that case, the mobile app will automatically indicate if the person answering has received the notification about the COVID-19 exposure. As with other *Omaolo* services, the coronavirus symptom checker could be used anonymously, or one can log in by identifying oneself electronically. If a person uses the coronavirus symptom checker anonymously, they also get a chance to log in when they receive the results and to give a consent to transfer the information they have provided about symptoms to the relevant parties. Otherwise, the information is erased when the user closes the browser.

Who: The coronavirus symptom checker is freely available online for anyone to use it with skills in Finnish, Swedish or English. Interestingly, unlike corresponding apps in some other countries, the *Omaolo* coronavirus symptom checker is available also for people wanting just to try it out (it is only expected from them to tick a box in the form referring to a such use). On the other hand, people whose condition is very weak, or who suspect themselves to be seriously ill are advised not to use the symptom checker, but rather to call emergency services immediately. Those who opt to fully utilize *Omaolo* services on their own need to be minimum 16 years of age. In addition, they need to live in one of the municipalities and hospital districts where the service is already integrated. In fact, approximately 3,3 out of 5,5 million people living in Finland currently have the possibility to send the symptom report to a health professional (Lehtiniemi and Ruckenstein 2020), but there are plans to increase this number and cover the whole country. In terms of visits, from mid-March when it was introduced till the end of August, the coronavirus symptom checker has been used more than one million times.

Why did you choose this case?

Since April 2019, I have been following discussions around ethical issues related to data-driven algorithmic technologies, which have been very lively in Finland, as well as the rest of the Europe and North America. In this sense, it is interesting that this Finnish ADM case - unlike some others highlighted by our research group in the *AlgorithmWatch* report (Ruckenstein & Velkova, 2019) - has raised no concerns. This is even more curious if one bears in mind that the debates about the transformation of Finnish social welfare and healthcare services have been extremely vocal over the past years. It has not been critiqued publicly that the symptom checkers have not been available in the languages of dominant immigrant groups in Finland, while, for instance, the COVID-19 infection rate has been at least in the early stages of the pandemic significantly higher for Somali-speaking Helsinkians compared to other demographic groups. One of the explanations for this lack of ethical concerns regarding the language of the ADM system might be that the service – with its supposedly opensource code and people's consent-based participation - seems to be in alignment with the hegemonic ethics discussions, which have been configured around the "logic of choice" (Mol, 2008). People's autonomy and a sense of control have been the main targets. In this sense, it is symptomatic that the controversy about e.g. Spanish renderings of the coronavirus symptom checkers has been exactly about the coercion to submit personal data (Calatayud, 2020). The "logic of care", on the other hand, starts out from the fleshiness and fragility of life, which simply cannot be neglected in the time of pandemic. Shifting focus to care makes us pay attention to interdependence, to view the coronavirus symptom checker not as a transaction in which something is exchanged, but rather part of an interaction which goes back and forth between all parties involved in the process (human and otherwise). Caring is a question of tinkering with bodies, technologies, knowledge and people. The issue, thus, shifts from people deciding to choose using ADM or not, to asking how humans could shape good care jointly with the technology, while at the same time handling that technology with care (Mol, 2008).

What does ADM do?

In terms of ADM, this case seems to be extremely simple: questionnaire input - decision tree - output assessment and possibly scheduling a visit to the doctor. The assessment can point to some other useful pages. ADM is simply an instrument, means to an end, which is how technologies are understood within the logic of choice, and the more effective they are, the better (Mol, 2008, p. 57). Accordingly, the symptom checker has been updated at least 15 times since its publishing, following the state of pandemic and changes in regulation (Lehtiniemi & Ruckenstein, 2020).

Case presentation: "Re-defining elderly and frail patients through Electronic Health Record data – a case of patients not showing up for diagnostics and surgery"

Christopher Gyldenkærne, Roskilde University

The case explained

What: Patient no-shows in healthcare is a global problem and is known from literature to range from around five to twenty percent of all clinical appointments. In Danish hospitals, the average is five percent and there are currently no well-known interventions that have shown to significantly bring down this number. After the implementation of a big Electronic Health Record (EHR) system (EPIC) and the roll out of an overall digital approach to citizen communications in Denmark, several clinics are reporting an increase in patient no-shows.

Patients not showing up for clinical appointments are known to correlate with frail patients not being able to access healthcare on “equal terms” with the general population. Many of these frail patients are known to suffer from co-morbidity, lifestyle diseases and even premature death.
My project covers an initiative at Bispebjerg and Frederiksberg hospitals in Copenhagen, where Neural Networks (AI) and more simple algorithms are being used cover such technological possibilities to predict future patient no-show.

How: By letting algorithms and AI look through past appointments in Danish clinics in conjunction with variables found in the EHR-systems known from literature to be associated with, or directly predict, patient no-show’s it is theorized that such technologies can help to better understand local population patterns of no-shows and further help clinics prevent, plan and work with and around patients at risk of not getting the appropriate clinical attention and care through showing up at the hospital. The long-term goal is to create data-driven interventions where a new view of how medical resources are being used and on who work on a danish hospital sector that can navigate with more local knowledge of population frailty and get away from the broad-spectrum interventions so far characterizing public health efforts in Denmark.

Who: Data originates from 77.850 past observations of Danish patients visiting cardiology for diagnostics and care and 2880 past observations of Danish patients visiting gastro-endoscopy for screening of rectal cancer. The data is blinded and all variables GDPR-compliant.

Why did you choose that case?

First and foremost, I’m involved with the case as primary data engineer of the project. Working with this case has opened up to various dilemmas ranging from access to data, data protection, governance, design implications for health care workers, intervention design that does not mark citizens through use of data and considerations with regards to explainability.

What does ADM do?

The vision of the project is to pin-point healthcare workers and decision makers to patients at risk of not getting appropriate care (as understood from a Danish public health point of view. Denmark may be characterized as having a social democratic view on healthcare.) The ADM comes in, when the actual intervention (intervention not designed yet) happens powered by the AI. An example could play out like this:

A local specialized clinic plans the coming week of patient visits → AI is highlighting 5-10 % of patients in need of follow-up/intervention to encourage show-up → specific guidelines to deal with these patients are being followed → patient outcomes are being monitored and evaluated with the main goal of bringing down the percentage of no-shows.

Case presentation: "DIY Artificial Pancreas System"

Henriette Langstrup, University of Copenhagen

The case explained

In type 1 diabetes the holy grail is the artificial pancreas – a solution to the fact that your own pancreas doesn't provide you with insulin. This vision has been pursued with stem cell research and with medical device innovation for years. Here I will focus on the latter, and specifically recent developments in the treatment technology of automated insulin delivering (AID). This development is mainly driven by patient entrepreneurs and international online and off-line collectives of people with diabetes (PwD) and relatives to PwD and outside of the realm of regulatory bodies and health professionals' involvement. Hence it may be characterized as DIY automated decision-making. It is estimated that at least 2000 PwD worldwide are using these systems and the number is growing.

These DIY automated insulin delivery systems have algorithmic decision-making at their core. Diabetes management involves a lot of decisions and mental calculation to adjust your carb intake, activities and insulin delivery. For PwD a central issue is the mental load involved in managing diabetes, and automation is seen as central to achieving better quality of life and better patient outcomes. PwD and parents of children with diabetes want to "share" some of this burden or responsibility with technology, allowing software to make continues decisions on the level of insulin provided through an insulin pump based on continuous feedback from a sensor measuring their glucose levels. But responsibility in healthcare decision-making is a legally, ethically and epistemologically complex issue. These patients make (anything but automated) decisions on technology use without formal (even against) medical advice and program their own devices with community help to automate decisions. The technology is DIY and open source and not FDA approved or CE marked. In "regular" treatment procedures, treatment decisions are delegated to the diabetes patients as advice and instructions from health professionals. For the community members – often referred to as "loopers", because they are "closing the loop" between insulin pump, sensor and body with an algorithm – there is great confidence in the automation, while always explained with reference to the help of the community and acknowledging the constant work of attuning the algorithm to local circumstances.

Why did you choose this case?

The case raises question about the humanizing potential of automation, about what constitutes caring infrastructures in disease management – how respons-able sociotechnical arrangements are crafted bottom-up and with what implications for whom. In this case, the risks and benefits of running on autopilot are both present. Automation is mobilized as a way to re-humanize diabetes technology – to get it to serve PwD before anyone else, be it device industry or health authorities. Here automation is part of an activist and emancipatory agenda. At the same time, the practice of "looping" involves many, also many new, decisions and types of responsibility, individually and collectively. It may differ from other types of ADM, where the end-user is not involved in setting up and monitoring the algorithm and may even be unknowledgeable about the automation and its settings. However, even in this highly elitist group some do experience some degree of "algorithmic vertigo" when applying an algorithm found on an online forum – and not fully understanding its workings.

I am fascinated with the way the users and PwD entrepreneurs talk about the care of the computer and the collective around "looping". One example is from the "origin story" of one of the initial entrepreneurs, Dana Lewis, who writes about the moment when they "closed the loop" – "I never want to turn it off or let anyone take it from me". Automation is described as caring and responsible: "A computer will watch carefully, constantly, and be able to respond more quickly than a human does in most situations" (Lewis, 2019, p. 35). Also, a concept such as "autosensitivity" (Lewis, 2019, p. 82), used to describe how the algorithm can be trained to be more and more sensitive, I find fascinating. Tensions between embodied and disembodied sensation is at the core of automation in personal health technology.

What does ADM do?

Based on direct sensor-data input from a continuous blood glucose monitor and data from insulin pump the algorithm predicts insulin need and adjusts insulin delivery. The system can be hybrid or fully automated depending on whether the user will input data on carbs in relation to meals or not. Most are using hybrid versions.

Discussion and ways forward

Our workshop consisted of a variety of cases, discussed under the umbrella term of ADM systems. The current ADM debate covers a variety of technical tools and systems, and a plethora of ADM definitions. Designers, legal scholars, policy makers, ethnographers, and data scientists can rely on different notions of ADM when they discuss the decision-making qualities of machines. The messiness of the ADM field might be seen as a problem, and one way forward is to engage in a cross-disciplinary mapping of ADM definitions in order to produce taxonomies and classifications for a shared vocabulary. Based on the cases presented in the workshop, however, there is also an alternative way forward: one can depart from the techno-centricity of the debate, side-line the technical arrangements of ADM after having carefully explored them, and focus on what kinds of social and societal arrangements are currently being built with ADM systems, and what kinds of problems they are seen to solve. As suggested by the Trelleborg case in Sweden, the exploration of an ADM case can aim to clarify the diverging and partly contradictory notions of what ADM does, and to whom, and highlight the related legal and political tensions, struggles, and consequences. The meanings and values attached to ADM systems are negotiated, and they might not be stable in any way, as they are connected with the changing imaginaries of benefits and harms of implementing ADM systems.

Focusing on what ADM does, rather than what ADM is, suggests that in order to humanise and rehumanise the field, we need **situational ADM studies** that offer a more dynamic and processual view to ADM systems. For this, empirical analyses of actual ADM cases are essential, as they can foreground who designs ADM systems, with whose decision-trees, data and algorithms, and with what kinds of implications. It makes a big difference whether a system is designed to solve organisational inefficiencies, as in the case of the Migri chatbots in Finland and in the case of predicting and preventing patient no-shows in Denmark, or whether the system seeks to promote better self-care in the context of a chronic disease, as in the case of the U-prevent tool. Alongside the design aims, particular attention should be paid to the changing nature of ADM systems over time, as these systems develop with their uses.

Furthermore, a historical awareness allows us not only to discern the kinds of continuity and change in ADM developments over a longer temporal scale, but also to understand the different regulatory and governance regimes under which ADM systems in different sectors and domains have developed, and are developing. With a historical sensibility, promoted with a situational exploration of ADM systems, we can witness the strengthening of existing infrastructures and efforts to build new ones. This is one way to advance ADM studies: to focus on the different infrastructural arrangements, including the stakeholders involved in the building of ADM systems, and their present or future uses.

Since most of the cases presented in the workshop deal with health and social welfare, it appears that the current healthcare and welfare infrastructures are particularly suited for observing changes in the ADM field in the Nordic countries. For researchers seeking out empirical cases for studying ADM systems, public sector cases and public-private partnerships in the healthcare sector, and in social services, appear not only timely, but raise questions concerning broader developments in the healthcare sector and beyond. For instance, the Finnish Eksote case can be used as an example of what gets developed, as similar projects are also ongoing with other authorities, with attempts to optimize and make current infrastructures more efficient with ADM systems. Here, ADM studies point towards prior organizational arrangements that need to be addressed in order for the ADM to work well. On the other hand, however, the workshop points to the direction of new kinds of digital health infrastructures, as the expansion of digital symptom checkers, and the DIY artificial pancreas system suggest. Separating between ADM systems that are add-ons to existing infrastructures, and those that suggest new kinds of infrastructural arrangements is fundamental to understanding what the ADM aims to do in society.

The presentations at the workshop demonstrate many ways for researchers to contribute to a better understanding of ADM systems. For instance, in terms of a municipal caseworker system in Denmark, empirical insights reveal how value metrics for algorithmic systems come to be negotiated in a participatory design set-up in a politicized context (Møller et al., 2020). Here, researchers provide practical suggestions for responsible design of decision-support systems. Overall, the empirical cases offer opportunities to reflect on the strengths and weaknesses of using ADM in various organizational contexts. Researchers can highlight trust in machines, and evaluate new divisions of labour. From this vantage point, we can start recognising in a more in-depth manner what collaboration between human and machine entails.

Many of the ADM systems, presented and discussed in the workshop, are not completed, but they are still in the making. They exemplify ADM visions and imaginaries, and narratives of what ADM will do, rather than what it already does. This leads questions of power: whose ADM narratives and imaginaries are currently being designed into actual systems. What is seen as worth developing and promoting? The presentations of our workshop suggest that we need to engage with a wider range of imaginaries, in order to develop new kinds of conceptual frameworks for capturing the processual nature of ADM systems. This kind of work is also helpful in terms of rethinking questions of governance.

Seeing alternative and other possible ADM futures requires a shift in perspective, promoted in our workshop, but perhaps not equally endorsed across the board. Yet, in order to renew the conversation, ADM debates need to let go of the techno-centricity that treats automated decision-making as a stand-alone product, innovation, or a solution to existing infrastructural inefficiencies and gaps. Instead, ADM needs to be treated as a complex socio-technical system that develop over time and need ongoing stabilisation of human-technology relations.

References

Calatayud, J. M. (2020). Country analysis: Spain. In F. Chiusi, S. Fischer, & M. Spielkamp (Eds.), *ADM Systems in the COVID-19 Pandemic: A European Perspective*. AlgorithmWatch.

Choroszewicz, M., & Mäihäniemi, B. (2020). Developing a Digital Welfare State: Data Protection and the Use of Automated Decision-Making in the Public Sector across Six EU Countries. *Global Perspectives*, *1*(1), 12910.

Hildebrandt, T. T., Andaloussi, A. A., Christensen, L. R., Debois, S., Healy, N. P., López, H. A., Marquard, M., Møller, N. H., Petersen, A. C., Slaats, T., & Weber, B. (2020). EcoKnow: Engineering Effective, Co-created and Compliant Adaptive Case Management Systems for Knowledge Workers. *Proceedings of the International Conference on Software and System Processes*, 155–164.

Lehtiniemi, T., & Ruckenstein, M. (2020). Country analysis: Finland. In F. Chiusi, S. Fischer, & M. Spielkamp (Eds.), *ADM Systems in the COVID-19 Pandemic: A European Perspective*. AlgorithmWatch.

Lewis, D. (2019). *Automated Insulin Delivery: How artificial pancreas "closed loop" systems can aid you in living with diabetes*. https://luckyloop.koeln/trlababa/wp-content/uploads/2019/10/Automated_Insulin_Delivery_by_Dana_M_Lewis_PDF_v0.0.4_May_31_2019.pdf

Mol, A. (2008). *The logic of care: Health and the problem of patient choice*. Routledge.

Møller, N. H., Fitzpatrick, G., & Le Dantec, C. A. (2019). Assembling the Case: Citizens' Strategies for Exercising Authority and Personal Autonomy in Social Welfare. *PACM on Human-Computer Interaction*, *3*(GROUP), Article 244.

Møller, N. H., Shklovski, I., & Hildebrandt, T. T. (2020). Shifting Concepts of Value. Designing Algorithmic Decision- Support Systems for Public Services. *Proceedings of the Nordic Conference on Human-Computer Interaction (NordiCHI)*. NordiCHI'20. Forthcoming, https://nordichi2020.org/schedule#/

Rakar, F. (2018). *Lärprojekt Trelleborgsmodellen – från rebell till modell* (pp. 1–21). Rhetikfabriken. https://moten.trelleborg.se/welcome-sv/namnder-styrelser/arbetsmarknadsnamnden/arbetsmarknadsnamnden-2018-06-11/agenda/larprojekt-trelleborgsmodellen-002pdf-1?downloadMode=open

Rieckert, A., Reeves, D., Altiner, A., Drewelow, E., Esmail, A., Flamm, M., Hann, M., Johansson, T., Klaassen-Mielke, R., Kunnamo, I., Löffler, C., Piccoliori, G., Sommerauer, C., Trampisch, U. S., Vögele, A., Woodham, A., & Sönnichsen, A. (2020). Use of an electronic decision support tool to reduce polypharmacy in elderly people with chronic diseases: Cluster randomised controlled trial. *BMJ*, *369*(1822), 1–10.

Ruckenstein, M., & Velkova, J. (2019). Finland. In M. Spielkamp (Ed.), *Automating Society. Taking Stock of Automated Decision-Making in the EU*. AlgorithmWatch.

SVT. (2018). *Omdiskuterad robot avgör vem som får stöd i Trelleborg [Contested Robot decides who gets benefits in Trelleborg]*. https://www.svt.se/nyheter/lokalt/skane/robot-avgor-vem-som-far-stod

Velkova, J., & Kaun, A. (2019). Algorithmic resistance: Media practices and the politics of repair. *Information, Communication & Society*, 1–18.

Annex: List of participants

Network participants, September 2020:

Minna Ruckenstein, University of Helsinki

Sonja Trifuljesko, University of Helsinki

Sonal Makhija, University of Helsinki

Tuukka Lehtiniemi, University of Helsinki

Maiju Tanninen, Tampere University

Yu-Shan Tseng, University of Helsinki

Laura Savolainen, University of Helsinki

Kirsikka Gron, University of Helsinki

Santeri Räisänen, University of Helsinki

Mea Lakso, University of Helsinki

Stefan Larsson, technology and society (ToS), Lund University

Stine Lomborg, communication and it, University of Copenhagen

Fredrik Heintz, computer science, Linköping University/Lund University

Katja de Vries, law, tech, philosophy, Uppsala University

Laetitia Tanqueray, master's student, sociology of law, Lund University, project assistant

Charlotte Högberg, PhD student Lund University, medical AI project (AIR Lund)

Anne Henriksen, PhD student Aarhus, digital social research, visiting Lund University

Lupita Svensson, social work, Lund University

Riikka Koulu, law and digitalization, University of Helsinki

Jacob Livingston Slosser, law/cognitive science, University of Copenhagen

Katarzyna "Kasia" Söderlund, PhD student ToS, law, AI and consumer trust

Jonas Ledendal, business law, Lund University

Anne Kaun, media and communication studies, Södertörn University

Sne Scott Hansen, communication and it, University of Copenhagen

Dorthe Brogård Kristensen, human health, University of Southern Denmark

Léonard Van Rompaey, PhD law, postdoc University of Copenhagen / Synch Law Firm

Marta Maroni, PhD student, Law, University of Helsinki

Christopher Gyldenkærne, PhD student, humans and technology, Roskilde University

Henriette Langstrup, public health, section of health services research, University of Copenhagen

Finn Kensing, computer science, University of Copenhagen

Naja Holten-Møller, computer science, University of Copenhagen

Sofie Flensburg, communication and it, University of Copenhagen

Sille Søe, information studies, University of Copenhagen

Liu Jun, communication and it, University of Copenhagen

Klara Jangaard, communication and it, University of Copenhagen