

Accessibility statement

This is an accessibility statement for the journal: Encounters.

Conformance status

The Web Content Accessibility Guidelines (WCAG) defines requirements for designers and developers to improve accessibility for people with disabilities. It defines three levels of conformance: Level A, Level AA, and Level AAA. This statement is relevant for volume 10, number 5, 2018 through volume 12, number 1, 2021. Encounters is partially conformant with WCAG 2.1 level AA. Partially conformant means that some parts of the content do not fully conform to the accessibility standard. Despite our best efforts to ensure accessibility, footnotes and graphs may not be accessible for screen readers at this point in time.

Feedback

We welcome your feedback on the accessibility of the journal. Please let us know if you encounter accessibility barriers. You can reach us at:

E-mail: imvko@cc.au.dk

Address: STS Center, Helsingforsgade 14, 8200 Aarhus N



STS Encounters

Research papers from DASTS

SPECIAL ISSUE

Volume 12 • Number 1 • 2021

Experimenting on the Enactment of Predictive AI: The Quest for a Future Proactive Healthcare Sector

Anne Henriksen

PhD fellow, Science-Technology-Society Centre, Aarhus University

Finn Olesen

Associate professor, Department of Digital Design & Information
Studies, Aarhus University

DASTS is the primary academic association for STS in Denmark. Its purpose is to develop the quality and breadth of STS research within Denmark, while generating and developing national and international collaboration.

Abstract

Currently, a large number of AI projects are experimenting with the use of AI and big data for various purposes, especially in the public sector. In this article, we explore one such AI project. Specifically, we study a group of developers in Scandinavia and their efforts to enact predictive AI through the development of a clinical decision support system (CDSS) in pursuit of a future proactive healthcare sector. This yet-to-be system was envisioned to prevent unplanned hospitalizations by ‘turning’ what we term ‘potential patients’, i.e. the effective management of patient trajectories, in pursuit of a proactive healthcare sector. In the article, we investigate this particular project as an ‘experiment’ and conceptualize the developing CDSS as a ‘partially existing object’ with an uncertain ontological status. By studying the gradual enactment and emergence of the CDSS, we illuminate how this fuzzy data-driven object is performed and gradually attributed with solid reality: during its creation process, it advances from being a proactive device imagined to be used in primary healthcare to becoming a triage tool embedded in the prehospital emergency department. Along the way, the project developers are also transformed, learning what ‘moves’ and ‘actions’ to make, and, thereby, becoming skillful CDSS-operators. By using ‘experiment’ as our analytical lens, the article renders visible how persons, locations, and procedures have to be changed, revoked, and suspended in order for the AI project to succeed. Thus, the article contributes to showing how ‘social mangling’ is an essential precondition for predictive AI to succeed as a prolific solution to specific healthcare challenges, along with developers’ learning and transformation.

Introduction

With several recent technological advances and the explosion in digital data (Babak, 2015), artificial intelligence (AI) seems to pose a new “promissory technology” (cf. Hoeyer, 2019; Tupasela & Di Nucci, 2020) imagined to solve all sorts of challenges outside confined laboratory spaces. In Amsterdam and Helsinki alone, more than thirty AI projects are currently running (Olsen, 2021). Similarly, a recent report on automated decision-systems finds that a large number of countries are “experimenting” with the use of AI for various purposes, especially in the public sector (Chiusi et al., 2020: 6ff.). Precisely the word ‘experimenting’ seems apt to describe the situation, even though modern AI has been with us for some time through various applications (Bryson, 2019). As Stilgoe (2018: 26) suggests, modern AI is still very much “a work-in-progress laden with promises for what it might become”. Self-driving cars relying on big data and AI sometimes crash (Stilgoe, 2018), and development projects, which have otherwise been championed, are occasionally discontinued because the use of algorithms turns out to nullify the expectations (see e.g. Hao, 2020; Heaven, 2020).

It is not new that particular capacities and reality are attributed to technological objects *even before* they have a stable existence. By virtue of expectations, aspirations, and imaginaries (Jasanoff & Kim, 2015), they are enacted as *partially existing objects* (Latour, 1999; Jensen 2010). Following Bruno Latour, we explore AI algorithms and the computer systems in which they are embedded as partially existing objects; objects which have an uncertain ontological status, and which exist and are defined only relatively to their networks of construction (Latour 1999; Latour & Weibel 2005; Jensen 2010).

How do partially existing objects as fuzzy as developing AI algorithms, with only limited materiality, garner existence? How do they and the computer systems they are built into go from being ‘weak’ objects to gradually becoming more and more ‘real’ material devices that have particular uses and are woven into well-established practices? How are they enacted as specific versions and realities? By means of

what politics? And as part of which imaginative spaces of opportunity enabled by particular actors?

These questions motivate our study and are explored on an empirical level based on one particular Scandinavian AI project, studied ethnographically by the first author. This project strived to design and develop a *clinical decision support system* (CDSS), based on modern AI techniques and big data. In the article, we investigate this particular project as an experiment. We do so to better grasp the emergence and enactment of the CDSS, and the dynamics played out along the way, i.e. how and why the AI project developed in a certain manner. This strategy reveals the preliminary and predetermined existence of CDSS before it's actual development, and yet it suggests that its ontological status is uncertain due to a lack of tangible qualities, materiality, and embeddedness into specific practices. By conceptualizing the AI project as an experiment with a partially existing object – the CDSS – we accentuate the question of how such objects are performed and attributed with existence in highly local design spaces, and how they attain stability.

Our article does not provide a detailed analysis of how the AI project progressed during the period of ethnographic inquiry. Instead, it seeks to understand how the project and the developing CDSS ties to larger societal transitions and the social and political shaping of society in virtue of their immersion into particular socio-technical settings. For the sake of the anonymity of the informants in the AI project studied, we do not refer to documents etc. that may disclose information about the particular project.

While much STS research has studied developing technologies and technological futures not yet 'boxed in' (cf. Latour, 1987), especially in the healthcare domain (see e.g. Jensen, 2010), only a limited number of STS studies have focused on modern AI systems and furthermore studied them empirically. This article contributes to the literature by filling this gap. In particular, it aims to show how an experiment-based analysis can contribute to illuminating the ongoing construction process in AI development projects, or the *social mangling* (cf. Pickering, 1995), through which AI algorithms and computer systems are enacted.

Experiments, Experimentation & Performativity

What is an experiment? Historically, a scientific experiment is a particular step in epistemological inquiries to create knowledge about a delimited phenomenon in nature. Reality must be manipulated so we can learn about it; we must "twist the lion's tail", as Francis Bacon taught us (Hacking, 1983: 149). The classical experiment is thus modelled after natural science ideals *as a method to test hypotheses about a delineated natural phenomenon in a controlled manner within well-defined laboratory spaces*. This classical model of experiments and experimentation, however, does not help to explain the big data and AI-based experimental practices in our study.

The philosopher Ian Hacking argues that experimentation needs to be investigated in its own right *as a practice*. It is not just a "step on the royal road" to knowledge. Experimentation is *doing* rather than *thinking*, and the experimental method is not just another name for "the scientific method" (Hacking, 1983: 14f). Hacking's arguments offer a fruitful entry to conceptualize experimentation with AI and big data.

Experimentation regards "the creation of phenomena", not their discovery, Hacking argues (Hacking, 1983: 220). It is an extremely complicated task to refine and stabilize phenomena as sources of relevant data. Not least, it is difficult to refine what should count as 'data' in an experiment. This difficulty traverses a long road from talking about data in a specific context to presenting universal phenomenal statements about the world. The task involves a significant learning process, requiring practical rather than theoretical abilities, where the experimentalists must patiently *train a range of skills* before they are able to make reliable observations. These are, for instance, the skills of turning, cutting, extracting, preserving, pressing, and repeating. In the process, they must *learn when the experiment has succeeded*, i.e. when 'nature' has spoken. Only when the whole setting and the apparatus work in the 'right' way is it possible to observe specific phenomena (Hacking, 1983: 230). It follows that observation plays a relatively modest role in experimental science compared to other

tasks. Endurance and practice create the experimentalists' ability to distinguish *artefacts* produced by the instrument from the *effects* produced by the observed entity.

A widespread paradox in the sciences is that "[...] most scientific experiments don't work most of the time [...]" (Hacking, 1983: 230). Because of the complexities and tough learning process involved in creating phenomena in an experiment, the risk of failure is high. The sociologist of science Andy Pickering agrees and suggests that we consider experiments as complex events in which 'dances of agency', or dialectics of resistances and 'accommodation', happen:

My basic image of science is a performative one, in which the performances, the doings of human and material agency come to the fore. [...] The dance of agency, seen asymmetrically from the human end, thus takes the form of a *dialectic of resistance and accommodations*, where resistance denotes the failure to achieve an intended capture of agency in practice, and accommodation an active human strategy of response to resistance, which can include revisions to goals and intentions as well as to the material form of the machine in question and to the human frame of gestures and social relations that surround it.

(Pickering, 1995: 21f.)

We will return to the image of the dances of agency later. It appears that both Hacking and Pickering exclusively write about experiments situated in concrete settings. Bruno Latour has pulled the experiment out of designated spaces and buildings, and argued that experiments and laboratories are *movable devices*. This is a significant analytic suggestion because it allows us to follow the 'CDSS-experimentalists' as they move around with their partially existing object: the developing CDSS.

Nothing extraordinary or distinctly 'scientific' happens inside the

walls of the laboratory, Latour (1983: 141) claims. So, why are laboratories considered to be extraordinary places? The explanation is simple. Any notable laboratory has run through a series of displacements in order to achieve its current status. A 'displacement' is here understood as a semiotic movement from one position to another on a flat surface. It is misleading to ask, *where is the laboratory* and *where is society*? The lab and the society are mixed up from the beginning.

In a dynamic process of displacements, things and humans are transformed. Latour illustrates this through his material-semiotic interpretation of a famous historical event (Latour, 1988). The agricultural system in France was transformed when the microbiologist Louis Pasteur in May 1881 displaced his lab, moving it from École Normale Supérieure in Paris to the village Pouilly le Fort, thereby moving it into the center of French farmers' interests. When Pasteur returned to Paris, he brought with him two things of utmost importance to the farmers: a cultivated specimen of the anthrax bacterium *and* the interest of the farmers, who wanted a cure to save their cattle. Thereby, Pasteur's laboratory was transformed from being a rather secluded setting in Paris to becoming a nationally significant experiment to save farmers' livestock. Latour's reading of Pasteur's achievements is a semiotic overruling of the contrast between text and biological material. Latour asks how the laboratory was made relevant. The answer is that the village was turned into a laboratory: "The only terrain in which a laboratory scientist is a master is that of experiments, of laboratory logbooks, test tubes and dogs" (Latour, 1988: 61). Hence, society must be transformed into such terrain if the scientists are to have relevance in society. We bring Latour's analysis of the laboratory-society transformations into our investigation as a remedy to learn how the CDSS-experiment was displaced to become a meaningful device in healthcare contexts.

We use Hacking's detailed exposition of experimenters' learning approach, combined with Pickering's evocative notions of the mangle, and Latour's semiotic analysis of transformative displacements in our own study of the AI-based CDSS. We additionally apply further theories foregrounding performativity in technology development

and experimental practice.

Decision Support Systems & AI in Healthcare

What is a clinical decision support system, and what is experimental about the development of such a system and the use of AI for this purpose? The history of clinical decision support systems tells us that they are “a class of computer-based systems that aids the process of decision making” (Ozaydin et al., 2016: 46) and includes some kind of “decision support capabilities” (Berner & La Lande, 2016: 2). Reading this, one quickly realizes that such systems are not new. In fact, they have been used for more than 50 years as parts of healthcare information systems with a view to “change the way medicine has been taught and practiced” and, in particular, prevent medical errors and improve diagnoses (Berner & La Lande, 2016: 2). Not least, clinical decision support systems have played a crucial role in making electronic healthcare records (EHRs) useable in practice (ibid.). Probabilities and probabilistic knowledge are also not new to clinical practice (Spooner, 2016). On the contrary, they have been instrumental in shaping medical science, especially epidemiology (Tversky & Kahneman, 1974). Finally, it is not novel to use AI approaches to develop computer systems for clinical use, including decision support systems. Several earlier AI approaches, e.g. Bayesian networks forming a part of the artificial neural networks characteristic of AI (see e.g. Press, 2016), were developed in relation to the medical domain through work on knowledge-based systems (Spooner, 2016; Liu et al., 2020).

What is new, however, is the use of *automation and big data* as means for building probabilistic knowledge and, more specifically, predictions used for making decisions and actions (Mackenzie, 2015; Mackenzie, 2017). In the current new ‘era of AI’, the “inference engine”, as it is called in knowledge-based systems (Spooner, 2016: 31), is learned through data rather than programmed by humans. The result is a machine learning model, or ‘algorithm’, that can automatically process and interpret huge volumes of data, by recognizing patterns for the purpose

of predicting future behavior of “entities”¹ (Ozaydin et al., 2016: 46, 50). In this way, data items representing humans (Bechmann, 2019) are automatically classified and mapped into predesigned categories (Bowker & Star, 1999), e.g. sick/not sick, which become pivotal in predictions of, for instance, which people need treatment. This new wave of AI techniques draws on the branch of machine learning methods called *deep learning* (Liu et al., 2020; see also Alpaydin, 2016), and is also referred to as *data mining* or *predictive modeling*. Here, the key objective is to “infer from a collection of data/measurements mechanisms to facilitate decision-making processes” (Ozaydin et al., 2016: 48). Hence, as this quote suggests, data are used as proxies for certain behaviors put under scrutiny.

Today, clinical decision support systems, fueled with big data and algorithms, are often envisioned to improve the management of treatment, medication, and screening of patients (see e.g. Galetsi & Katsaliaki, 2020; Raghupathi & Raghupathi, 2014), for instance by providing and supporting preventive care to individuals through predictions, patient profiling, and segmentation (ibid.; see also Mønsted, 2019). In this sense, modern clinical decision support systems come with a particular *ontology*. This ontology may be characterized by a vision of the world as utterly stable, determinate, and knowable (cf. Pickering, 2016; Law, 2004), composed of data ‘out there’ that merely have to be mined and processed in order for behaviors to be predicted and subsequently tamed and controlled (cf. Berg, 1997). Yet, research shows that extensive work is necessary in order to make data ‘ready’ for uses other than those they were originally produced for and as parts of (Bonde et al., 2019; Møller et al., 2020; see also Loukissas, 2019). Such work and the demand for high-quality ‘reusable’ data have to be viewed in the context of the increasing need for proper data infrastructures, thanks to the gradual turn in the healthcare sector and society to data and data-driven computer systems (Kaun & Dencik, 2020; Bossen & Piras, 2020). In the case of data produced in

¹ This is referred to as supervised learning which draws on target variables in the training of learning algorithms (Russell & Norvig, 2016).

medical contexts, however, it may be that, no matter how much such data are ‘recooked’ (cf. Gitelman, 2013), they will still be ingrained by “uncertainty” due to the approximation, inadequacy, complexity, and ambiguity that generally permeate clinical practice, especially in primary healthcare (Cabitza et al., 2018). This intrinsic uncertainty can create problems for AI developers to achieve veracity of, and ‘truth’ in, algorithmic outputs (Mønsted, 2019; Henriksen & Bechmann, 2020).

In the following, we focus more closely on the empirical basis of our study – the emerging CDSS – by accounting for how it was studied ethnographically.

The Empirical Basis: Studying the AI Project & the Developing CDSS

The AI project we study took place in Scandinavia over a three year-period in the form of a research project, initiated and managed by a publicly funded regional innovation incubator; this is described on the company website as an institution with an aim to bridge public and private organizations in order to “develop, test, and implement welfare and healthcare technology, solutions, and services in the national healthcare system”. Furthermore, the project involved a private AI company, a municipality, a state university, and a regional hospital that had already been engaged in a number of technology development projects and was generally noted for its ‘innovation-readiness’. In this way, the project formed a public-private partnership which is characteristic of how digitalization of public services is achieved today in Scandinavia (see e.g. Hockenhuil & Cohn, 2021). By asking the steering committee of the project for permission to follow the AI developers’ work on the design and development of the CDSS, the first author was allowed to study this work close-up. She primarily followed the activities inside the AI company. Here, a small team of data engineers, designers, and a physician met on a daily basis, often with the involvement of the project manager from the innovation incubator and the managing director of

the AI company.

All in all, the first author *followed* (Latour, 1987) the developers’ work and thereby the developing CDSS from late 2018 until early 2020, when the computer system was at its infant stage. She investigated the developers’ work by drawing on an ethnographic approach as in earlier ethnographic studies within STS that focus on AI development (Agre, 2016; Forsythe & Hess, 2001; Suchman, 1987). Various methods were used to generate “thick descriptions” of the work (Geertz, 1973). Participant observation was conducted at numerous meetings and workshops held in relation to the AI project during the entire period. Furthermore, the first author stayed with the company on an everyday basis from March to August September 2019, where she performed several spontaneous on-the-spot interviews with developers and conducted day-to-day observations of their everyday work. Additionally, she conducted more than 20 semi-structured interviews (Kvale, 2008), for instance with managers, business developers, data scientists, data modelers and so forth. These interviews were primarily conducted during two periods: August-September 2019 and January-February 2020. Two interviews were co-conducted with the second author of this article. Almost all developers participating in the development of the CDSS were interviewed. As the first author had agreed to assist the CDSS development team in exchange for access, she made the observations and spontaneous interviews in the role as a *participant observer* (Hammersley & Atkinson, 2007). For instance, she helped to investigate work procedures among general practitioners (GPs) in the very beginning of the project, and to take notes at meetings and participate in discussions. However, she never played any crucial role in the project or in the AI company in general. In fact, she found it difficult sometimes to be *truly* involved in the daily work on the design and development of the CDSS. Such difficulties may be seen as a fundamental premise of ethnographic research (Hammersley & Atkinson, 2007). All the data used in this article have been fully transcribed and subjected to an iterative analysis process, open to new themes emerging from the data and yet informed by our research interest. More specifically,

the data have been analyzed by means of initial categorization based on “the participants’ voice” resulting in preliminary themes and topics (Malterud 2012, p. 796), repeated readings to generate more condensed meaning units (Davies, 2008), and simultaneous writing and thinking to generate more well-found interpretations of the data (Denzin, 2013; St. Pierre, 2011). Furthermore, the first author analyzed documents collected in the field with the aim to gain a better understanding of the AI project and the developing CDSS. These documents included different descriptions, e.g. the project description in the application for *The National Innovation Fund*, which granted the AI project its initial funding. Latour (1986) argues that such material objects actively take part in the construction of new ‘things’, both in terms of knowledge and material. Using technical drawings as an example, he contends that material objects serve as visualizations of ‘the future’ because they are used by scientists and innovators in their attempts to convince audiences how their proposed ‘thing’ functions like a roadmap to the future, i.e. the one and only way. From that perspective, the above-mentioned project description would be considered a crucial vehicle in the attempt to mobilize funding for the AI project.

In the following section, we explore which future this project description suggested was engendered by the CDSS and how. In doing so, we make references to this description.

The CDSS Envisioned

According to the project description in the application to The National Innovation Fund, the stated goal of the AI project was to “predict unplanned admissions including readmissions” and thereby “identify” individuals in the risk of such admissions “before they require acute treatment”. In that sense, unplanned admissions were viewed as *avoidable* and, consequently, *manageable* admissions. The project was to achieve this goal by means of the proposed CDSS. This yet-to-be system was defined as “a machine learning-based clinical decision support system for proactive healthcare”. It was furthermore denoted a

“predictive system” – a system building on machine learning-processed “predictions” generated by an AI algorithm. This was “an algorithm for early identification of unplanned admissions” developed with the use of deep learning methods, namely “deep neural networks (DNN)”. It was the development of this algorithm and an appertaining “explanation engine” and “simulator” which made up the “primary research goal” of the proposed project. In other words, the developers proposed an applied research project highly focused on technology development. Furthermore, the final CDSS was aimed to be marketed as a product; thus, “commercialization” was also to form a significant aspect of the project. In consideration of these various different parameters, ‘experimentation’ seems a fitting headline for the complex line of work in this AI project.

The predictions generated with the CDSS should additionally be coupled with “clinical aspects”, e.g. “early screening, preventive care, and ultimately diagnosis”, which could be put in place by health professionals in order to prevent detected citizens from potentially being (re)admitted to the hospital as emergencies. Hence, with the CDSS, it would be possible to “screen each individual citizen at very high intervals, determining which people require care and help” and “predict and change patient trajectories”. The CDSS would thus not only help bring down a big bulk of admissions, expected to increase even further, and reduce the great costs of such admissions; it would also, and perhaps more importantly, “position the healthcare sector in a proactive role instead of a reactive role”, thereby supporting new national strategies for the healthcare sector. In this way, the CDSS was envisioned as a key ingredient and future-generating device (cf. Jensen, 2010) in the creation of a specific healthcare future which several actors imagined to become real: *the proactive healthcare sector*. Researchers in the sociology of expectations have argued that it is by articulating such futures through visions and expectations as forms of “wishful enactments of a desired future” (Borup, 2006: 286) that entrepreneurs contribute to the *materialization* and *performance* of such futures.

We suggest that a particular configuration, or version, of the patient

was woven into the CDSS in this envisioned future, namely *the potential patient*, i.e. a patient who closely resembles a rational citizen with a moral standing who had not (yet) been hospitalized and was willing to do as prescribed in order not to be so, for instance attending smoking cessation courses. Thereby, potential patients were assigned an important role which they – in the envisioned future – had to fulfil *in order for the CDSS-experiment to succeed*, and for the imagined new automatized data-driven procedures to become real.

Data also had a significant role to play – a rather fundamental requirement in order for AI algorithms to function: “Algorithms are inert, meaningless machines until paired with databases upon which to function” (Gillespie, 2014: 169). Different approaches and methods to preventing (re)admissions and improving the cross-sectoral collaboration for this purpose have been tested over the years (see e.g. Wadman et al., 2009), also including the use of statistical methods (see e.g. Data Study Group team, 2019; National Services Scotland, 2008). Seemingly, it was the commitment and endeavor to take big data and the data-intensive method of deep learning as the starting point for knowing and intervening which constituted the novelty of the AI project and its ‘innovation power’. By doing so, the project provides yet another example of how public services are increasingly organized around data and data-driven computer systems (Kaun & Dencik, 2020), not to mention automated data analysis-methods. It is argued that this data-centered ordering causes a huge transformation of public services, the relation(ship) between citizens and state institutions, and the welfare state in general (ibid.: 2; see also Ruppert, 2019).

The data made available to the AI project in order to construct the *ground truth* (Jaton, 2017) of the algorithm came from a research database. This database was managed by the regional hospital participating in the AI project and was developed within a framework of *another* publicly funded research project. The broad aim of this project was shared by the AI project in our study: to use the vast quantities of data collected by national and local authorities – a hallmark of the Nordic welfare state (Tupasela et al., 2020) – to build knowledge on

the causes for unplanned (re)admissions. This was in order to focus and strengthen the cross-sectoral collaboration and identify the ‘ideal’ coherent continuity of care across primary and secondary healthcare. According to a description of the database, it consisted of four sources of data: (1) clinical and administrative data from the regional hospital; (2) prescription, telephone/mail, and attendance services-data from GPs; (3) data from municipalities containing information about health and social services provided for citizens as well as address register-data; and (4) data from national records such as CPR-data, electronic patient record-data, and socioeconomic data. Field observations show that the database was continually supplied with more data over the years.

The Enactment & Emergence of the CDSS

We can conceive the CDSS as a “bounded *something*” (Jensen, 2010: 24) in virtue of the opportunities, aims, and methods stated in the application for the National Innovation Fund, the money granted by this fund, and the data made available to the developers. It was a ‘thing’ that partially existed by virtue of the conceptions that had been made of how it should be performed and thereby yield desirable outcomes and changes. Yet, it was also still a ‘thing’ that had to be further constituted through situated contexts of creation (Hacking, 1983), as it only made up a somewhat *diffuse* part of the healthcare sector. Its ontology had fairly blurred contours, so to speak.

When the first author began her study of the developers’ work in the AI project, the CDSS had just recently been featured in the national news with the message that engineers at the AI company had made a discovery. In their search for signals in the data, they had run some machine learning experiments and discovered that it was possible to predict, with great accuracy, the likelihood of citizens being hospitalized as emergencies within the next 100 days. Hence, the formation of the CDSS had started in a laboratory manner by examining data for patterns and, consequently, statistical correlations. This was in line with the work plan described in the project description, proclaiming

that the first step in the AI project would be to “develop algorithms to get early results”.

While the first author began to participate in meetings and workshops, it became clear to her that it still remained an open question to the developers regarding how to *use* predictions to stave off unplanned admissions. Arguably, predictions alone would not prevent potential patients from being hospitalized; they needed to be enacted and performed as a part of well-established practices. Hence, the discovery that it would be possible to predict potential patients within the next 100 days only posed a success *on paper*. It was still quite unclear how to apply the predictions *in practice*. In this way, the situation was different from Pickering’s (2018, p. 7) Glen Canyon Dam example, where the engineers’ calculations of the future imply that “there is nothing left to find out in dam design”. In this case, the developers had learned that there was yet much to find out. What they particularly needed to explore was how to “turn what by engineers is regarded as a great algorithm into an algorithm that is applicable in a clinical setting”, as the project manager contended at a meeting. A little later, when the first author started to spend time with the AI company on an everyday basis, it turned out that the developers had begun to consider their *data-driven* approach somewhat mistaken or at least insufficient. The reconsideration of their standard modus operandi was later expressed by an engineer in an interview:

Informant: You see a lot of stuff being published on what one can do with AI... and 99 percent of these articles come from engineers who *start* the problems, or, I mean, the projects *themselves*, without considering: “Is this an actual problem, and is there a sensible intervention or action?” This is why we see that PubMed and ArXiv.org are *crammed* with trivial AI studies.

Interviewer: Which spring from data because *these data exist*?

Informant: Yes, technical fascination with problems that

can be solved *only because data exist*. [...] Engineers and data scientists have a tendency to [...] just look at data first because it is *data-driven* [...] And then we don’t really consider, for instance, what sort of outcome do we actually want to look at. [...] Statisticians cannot just run multiple tests of various kinds because then they can detect correlations in *all data*. But this is the approach you use as an engineer, in principle, because you just look at data first. And that’s why you find: “Oh, there is a problem here that I can solve! Let’s do an article on that”. But starting with the data is just the wrong place to start out.

(Chief engineer, interview, Feb. 2020)

In order to understand how to prevent unplanned admissions predictively in practice and make the experiment succeed (Hacking, 1983), the project manager and the managing director of the AI company made a guess that there was a need for means and methods centered more directly around *users and working practices*. “Where in the healthcare sector are the good AI use cases?”, the project manager asked at a meeting, as if such cases merely had to be discovered. Based on her quest for such cases, she arranged a workshop to identify potential users of the CDSS. The first author attended the workshop and became aware that the majority of the participants² had begun to hypothesize that unplanned admissions would be avoidable *if the CDSS was used by GPs*. This was based on the discovery that potential patients could be predicted 100 days prior to a hospitalization. Consequently, it was surmised that potential patients assumedly needed to be “turned” *before* being admitted to the hospital as emergencies. Hence, the CDSS should be integrated in primary healthcare, and most likely general practice, since GPs would typically be the first to see potential patients. Furthermore, GPs had the authority to refer patients to hospitals and

² Participants came from both the AI company and the regional hospital, and, furthermore, included the project manager from the innovation incubator.

municipality services, and to prescribe medicine, unlike other clinicians in primary healthcare. Again, this idea of ‘turning’ potential patients on the basis of predictions and actions builds on the presupposition that *potential patients would act as rational citizens and participants in the proactive healthcare sector*.

Soon after the workshop, while the first author began to conduct day-to-day observations at the AI company, a number of new employees suddenly appeared: a design director, a UX designer trained as an anthropologist, and a physician. This marked the beginning of a new phase in the AI project, characterized by another more *explorative* “mode” of experimentation (Pickering, 2016). This mode was informed by *design thinking* – a methodology known as a user-centric and iterative process for development and innovation (see Brown, 2008), which has also been conceptualized as a practice of “opportunity creation” (Nielsen et al., 2017). The design director explained at initial meetings that the aim with design thinking was to *be open* to what the CDSS might become and avoid going into “solution mode”. “It could even be a mobile app”, he suggested. Hence, the idea was, in other words, to assume an “unpredictable becoming” of the emerging CDSS (Pickering, 2016: 4), as it was too early for it to be ‘black boxed’ (cf. Latour, 1987). We suggest that the CDSS now was to be regarded as a *multiple object* (Mol, 1999), in the sense that it could be many different ‘things’. As long as it would contribute to detecting and, in effect, reducing (un)planned hospital admissions as a marketable AI system, it was less important *how* – through which practices – it was performed and enacted. It is perhaps telling that the managing director of the AI company at one point noted the difficulty in adopting this almost *naive* view on the CDSS: “data *really* biases one’s ideas”, he claimed.

The design thinking process was introduced to the development team as a non-linear process involving five phases: emphasize, define, ideate, prototype, and test (see e.g. Garcia & Lähdesmäki, 2019). The idea was to repeatedly make observations of clinicians’ work, use such observations for identifying the clinicians’ presumed “pains”

and “gains”³, and then brainstorm on possible solutions, i.e. different enactments of the CDSS, followed by prototyping and testing (ibid.: 74). Hence in design thinking, it was *ethnographic observations* which posed the driving force for understanding the world and the emerging CDSS – not digital data and statistics. The integration of design thinking into the CDSS-experiment thus had two important implications. Firstly, it meant that the AI-based ontology for constituting and developing the CDSS was supplemented by other ontological ideas as a means for making the experiment succeed. Secondly, it meant that the project was now to be performed at other locations outside the AI company, thereby *distributing* the experiment (Latour, 1983) and *expanding* the innovation space, i.e. the space of ‘opportunity’ (cf. Nielsen et al., 2017; see also Kjærdsgaard et al., 2016).

The first place where observations were made by the anthropologist, and partially also by the physician and first author, was among GPs with the aim to test the initial hypothesis: potential patients predicted within the next 100 days could be prevented from being hospitalized if GPs used predictions. Hence, the mode of experimentation was not entirely explorative but perhaps more a ‘search for answers’ and, consequently, opportunity. Subsequent sessions to identify ‘pains’ and ‘gains’, brainstorm on possible solutions, and create initial prototypes followed on the basis of the observations brought back to the AI company, i.e. notes and statements from the doctor’s practice. Especially the brainstorming on possible solutions and prototyping played out in a manner we choose to describe as *virtual dances of agency* (Pickering, 1995). In its original sense, dances of agency are characterized as time-extended struggles between the scientist and obstinate machinery, where the resistance of the machinery is viewed as a kind of “material agency” that struggles with the human agency of the scientist in a dynamic process. But as the lack of so-called “live data” made it highly difficult for the developers to explore and test how the CDSS would perform – what the apparatus would do – they

³ “Pains” are explained as “the negative outcomes from current situations” and “gains” as “positive outcomes that users are trying to achieve” (Garcia & Lähdesmäki, 2019: 74).

had to *imagine* it. In this sense, the struggles rather took place in the form of thought experiments, where the developers tried to make the CDSS accommodable, or at least *figure out the opportunity for making it accommodable*, while constraints would turn up as resistances. For instance, during these experiments, it turned out to pose a significant constraint, which was difficult to get round, that GPs had no reasons *not* to refer potential patients to the hospital – not even profit.

The idea of the CDSS as a ‘thing’ used by the GP gradually vanished through these virtual struggles, while new hypotheses emerged. This fueled the continued exploration among, first, on-call GPs, and then prehospital emergency nurses and home care workers. Further constraints emerged as resistances while the processes of pains and gains-identification, solution brainstorming, and prototyping were repeated based on observations. For one thing, it was doubted whether it would be possible to ‘capture’, i.e. detect, potential patients on the basis of the data generated in primary healthcare which had turned out to be of poor quality. This was a significant hurdle and prompted the project manager to contend at a meeting that municipalities should require home care workers to work in a more “standardized” manner. That was to ensure they would use digital systems more identically and thereby produce more complete and consistent data, thus providing a proper data infrastructure for the CDSS (Kaun & Dencik, 2020; Bossen & Piras, 2020). Also, it was difficult to imagine how the CDSS could be used to detect potential patients and propose preventive care initiatives when citizens were generally evaluated against *individual baselines*:

It became evident that *habitual* aspects really count in primary healthcare; everything is measured against what is habitual. You may have some high, crazy values, but, if they are habitual for that patient, there is no need to react. So, it’s always, you know, measured against *what?* [...] And that’s what’s making it extra difficult because which actions should we then propose are put in place? (Anthropologist, interview, Feb. 2020)

Consequently, it became difficult to *change patient trajectories* and enable the ‘turning’ pursued, as the quote suggests. Hence, the idea to predict potential patients within 100 days and prevent them from being hospitalized by enacting the CDSS as a part of primary healthcare gradually vanished during the virtual dances of agency that the developers engaged in. The project manager, meanwhile, learned that a physician executive consultant from the regional hospital, who participated in the AI project, was working on problems that might pose an opportunity. The development team was motivated by this news and initiated observations at the prehospital emergency department of the hospital. The experiment was now further distributed and took a new turn (Latour, 1983). Additionally, the team invited the physician executive consultant to a meeting, where he showed great interest in the emerging CDSS and shared his visions with the developers. He envisioned that patients hospitalized as emergencies could have their condition predicted during preadmission evaluation. Statistics showed that the majority of such patients were hospitalized as patients with *unstable conditions*, although most of them were discharged as patients with *stable conditions*. Hence, he claimed that these patients would have had too many examinations carried out compared to their ‘actual’ condition. In this way, it was not only observations that were *imported* into the experiment (cf. Galison, 1987) but also the *interests* of other actors (Callon & Latour, 1981). Motivated by a declaration of interests from at least one external actor, the physician executive consultant, further meetings were conducted to retain and translate his interests in the project. Not least, further meetings were conducted to retain and *nurture the opportunity* to enact the CDSS as a part of the prehospital emergency department (cf. Nielsen et al., 2017). Thereby, the content and context of the experiment were stabilized for the time being.

At the time when the first author left the experiment in February 2020, the CDSS performed more and more as a *triage tool* with an ontology based on binary logics (stable/unstable), leaving the developers with the feeling that their experiment would soon succeed. In

this regard, one particularly factor to which the developers assigned much significance was *their own learning*. This was expressed by several developers, including the director of the AI company, but especially also by an engineer involved in the project on a regular basis:

I believe many AI projects start out at the wrong place; they begin by providing data to some engineers who then learn that they can do something with these data: make predictions or classifications. And then they measure how well they do it. If they do it really well, they'll think, "This is awesome; maybe we're just as great at this as clinicians are, or perhaps even better – we've solved a big issue!" But that process should start by considering where there are problems to be solved. What is it that doesn't work? [...] You know, begin with the clinical professionals or such an anthropological approach in order to observe problems [...] Only if you've found things that don't work well [among clinicians], can you start to consider if you can do something about this at all. If yes, then you can start to look in the toolbox, and *one* of the tools you'll have in that box; that's AI. And only then can you start to consider if AI could help to solve the problem *rather than starting with AI*. [...] Then the next step will be action... because if the prediction or classification wouldn't lead to an action which will *change* the workflow or care trajectory, it makes no difference.
(Chief engineer, interview, Feb. 2020)

As the quote suggests, the developers had learned that they might be more successful in enacting predictive AI in the healthcare sector by drawing on designers, user-centric approaches, and ethnographic observations *in addition* to data, statistics, and engineering, and by exploring the working practices first rather than the digital data produced through such practices. As previously noted, Hacking (1983)

contends that experiments are about learning how to use an apparatus or instrument in the right way, and knowing when the experiment succeeds. In our case, we might say that it was about learning *how to perform an experiment* in order for predictive AI, and thus automatized data-driven procedures, to succeed as a solution to specific healthcare challenges.

Conclusion

In this article, we have studied the experiment carried out in one particular AI project in Scandinavia. This particular project strived to enact a big data and AI-based clinical decision support system (CDSS) to prevent unplanned hospitalizations in pursuit of a specific healthcare future: the proactive healthcare sector. We have argued that the enactment of this CDSS relies on the effective 'turning' of what we have called 'potential patients', and the presumption that such patients will act as rational citizens. We chose to regard the developing CDSS as a 'partially existing object' (Latour, 1999; Latour & Weibel, 2005) with an uncertain ontological status. It was a 'thing' in virtue of its stated aims, data, methods, and funding, and the conceptions made of it beforehand, and yet a 'thing' that had to be further constituted through situated contexts of creation (Hacking, 1983). That creation was undertaken by a small-scale development team at an AI company followed by the first author in an ethnographic case study over the course of a year.

By studying the gradual enactment and emergence of the CDSS through its construction process, we have shown how it went from being a future proactive device to becoming a triage tool to be integrated into medical triage at the prehospital emergency department. Paradoxically, this was the department that it had previously been envisioned to *keep potential patients out of* by its imagined enactment in primary healthcare and especially general practice. This imagined enactment was based on a discovery made in data that it would be possible to predict, with great accuracy, the likelihood of citizens being

hospitalized within the next 100 days. It is important to note that the CDSS did not simply 'move' to the prehospital emergency department and thus the hospital by itself. The developers learned en route how to *perform* the CDSS-experiment so that it could succeed (Hacking, 1983), and, as they learned this, they and the AI company were *transformed*.

First, the developers learned how to import (cf. Galison, 1987) methods and ideas into the AI project that would change the horizon and modus operandi by which they had attempted to conceive and develop the CDSS initially. This became evident in that they began to draw on a more explorative and user-centric mode of experimentation (Pickering, 2016) informed by design thinking which incorporated ethnographic observations (Brown, 2006). Through so-called thought experiments, the developers engaged repeatedly in 'virtual' dances of agency (Pickering, 1995) with a focus on primary healthcare. The choreography of these dances, implying multiple resistances, made the developers see new opportunities, and, perhaps more importantly, it made the developers *revise* the apparently revisable aims and scope of the AI project in their attempt to accommodate the CDSS to healthcare contexts. Furthermore, the developers learned how to draw the experiment into a wider, situated context, and how to *displace* the CDSS-experiment (Latour, 1983) so that it could recruit new actors and interests and allow for new opportunities for 'enactment' of the CDSS to emerge (cf. Nielsen et al., 2017). Arguably, it was this *transformation* and *learning* that enabled the developers in the project to constitute the CDSS as a concrete device in medical triage. Through the transformation and learning processes, the developers changed from being beginners to becoming competent CDSS-operators, knowing what 'moves' and 'actions' to make in order to enact predictive AI in healthcare.

In the end, we can understand the CDSS-experiment in our study as a dynamic process through which actors strive to *re-enact* the healthcare sector by means of particular roles assigned to patients, new automated data-driven procedures based on modern AI techniques, and the involvement of private AI companies in decision-making processes, i.e. triage. By studying the AI project and emerging CDSS through

an analytical lens focused on 'experiment', the article has rendered visible how persons, locations, and procedures had to be changed, revoked, and suspended in order for the AI project to succeed. Thus, the article contributes to showing how 'social mangling' (cf. Pickering, 1995) is an *essential precondition* for succeeding in experiments on the enactment of predictive AI, along with developers' learning and transformation. Despite their presumed powers, it is not data, methods, and technology per se that make such experiments succeed as an instrumental stance would suggest. Rather, it is what is done *with* and *around* such machinery that matters, and through which fuzzy AI algorithms and data-driven computer systems become 'real' material devices with concrete uses in particular healthcare contexts.

References

Agre, P. E. (1997). *Toward a Critical Technical Practice: Lessons Learned in Trying to Reform AI*. In: G. C. Bowker, L. Gasser, S. L. Star and B. Turner (eds.), *Bridging the Great Divide: Social Science, Technical Systems, and Cooperative Work*. USA: L. Erlbaum Associates Inc.

Alpaydin, E. (2016). *Machine Learning: The New AI*. The MIT Press.

Bechmann, A. (2019). *Data as Humans: Representation, Accountability, and Equality in Big Data*. In: Jørgensen (ed.), *Human Rights in the Age of Platforms: 73-93*. Cambridge, MA: The MIT Press.

Babak, H. (2015). *The AI Resurgence: Why Now?* Wired. Available at: wired.com/insights/2015/03/ai-resurgence-now/ (Accessed May 2 2021).

Berg, M. (1997). *Rationalizing medical work: decision-support techniques and medical practices*. Cambridge, MA: The MIT Press.

Berner, E. S. and La Lande, T. J. (2016). *Overview of Clinical Decision Support Systems*. In: E. S. Berner (ed.), *Clinical Decision Support Systems*. Healthcare Informatics. Switzerland: Springer International Publishing. DOI: 10.1007/978-3-319-31913-1_3.

Bonde, M., Bossen, C., and Danholt, P. (2019). *Data-work and friction: Investigating the practices of repurposing healthcare data*. *Health Informatics Journal*, 25(3): 558-566. DOI: 10.1177/1460458219856462.

Borup, M. et al. (2006). *The sociology of expectations in science and technology*. *Technology Analysis and Strategic Management*, 18(3/4): 285-98.

Bossen C. and Piras, E.M. (2020). *Introduction to the Special Issue on 'Information Infrastructures in Healthcare: Governance, Quality Improvement and Service Efficiency'*. *Computer Supported Cooperative Work (CSCW)*, 29: 381-386. Springer Nature. DOI: 10.1007/s10606-020-09381-1.

Bowker, G.C. and Star, S.L. (1999). *Sorting things out: Classification and its consequences*. Cambridge, MA: The MIT Press.

Bryson, J. (2019). *The Past Decade and Future of AI's Impact on Society. Towards a New Enlightenment? A Transcendent Decade: 1-33*. Open Mind. Available at: bbvaopenmind.com/en/articles/the-past-decade-and-future-of-ais-impact-on-society/.

Brown, T. (2008). *Design Thinking*. Harvard Business Review. Available at: hbr.org/2008/06/design-thinking.

Cabitza, F., Ciucci, D., and Rasoini, R. (2018). *A Giant Feet of Clay: On the Validity of the Data that Feed Machine Learning in Medicine*. In: F.

Cabitza, C. Batini, and M. Magni (eds.) *Organizing for the Digital World, Lecture Notes in Information Systems and Organisation 28*, 121-136. Springer International Publishing. DOI: doi.org/10.1007/978-3-319-90503-7_10.

Callon, M. and Latour, B. (1981). *Unscrewing the big Leviathan, or how actors macro-structure reality and how sociologists help them to do so*. In: K. Knorr Cetina and A. Cicourel (eds.), *Advances in Social Theory and Methodology*, 277-303. London, UK: Routledge and Kegan Paul.

Chiusi, F. et al., (2020). *Automating Society Report 2020*. AlgorithmWatch gGmbH. Available at: automatingsociety.algorithmwatch.org
Data Study Group team. (2019). *Data Study Group Final Report: NHS Scotland*. Zenodo. DOI: doi.org/10.5281/zenodo.2539563.

Davies, C. A. (2008). *Reflexive Ethnography – A Guide to Researching Selves and Others*. London, UK; New York, USA: Routledge.

Denzin, N. K. (2013). Writing and/as Analysis or Performing the World. In U. Flick (ed.), *The SAGE Handbook of Qualitative Data Analysis*, 569-584. London, UK: Sage Publications.

Forsythe, D. and Hess, D. J. (2001). *Studying those who study us: An anthropologist in the world of artificial intelligence*. USA: Stanford University Press.

Galetsis, P. and Katsaliaki, K. (2020). Big data analytics in health: an overview and bibliometric study of research activity. *Health Information and Libraries Journal*, 37(1), 5-25. DOI: doi.org/10.1111/hir.12286.

Galison, P. (1987). *How Experiments End*. Chicago, IL, USA: University of Chicago Press.

Garcia, C. and Lähdesmäki, R. (2019). Long live design thinking: How to nurture and grow design thinking in your organization. *Idean and Cap Gagemini Invent*. London, UK: Calverts.

Geertz, C. 1973. *The Interpretation of Cultures*. New York, NY, USA: Basic Books.

Gillespie, Tarleton. 2014. "The Relevance of Algorithms." In: T. Gillespie, P.J. Boczkowski, and K.A. Foot (eds.), *Media Technologies: Essays on Communication, Materiality, and Society*, 167-193. Cambridge, MA, USA: MIT Press.

Gitelman L. (2013). "Raw data" is an oxymoron. Cambridge, MA: MIT Press.

Hacking, I. (1983). *Representing and intervening: Introductory topics in the philosophy of natural science*. Cambridge, CA: Cambridge University Press.

Hammersley, M. and Atkinson, P. (2007). *Ethnography: principles in practice*, 3rd Edition. New York, NY, USA: Routledge.

Hao, K. (2020). The UK exam debacle reminds us that algorithms can't fix broken systems. *MIT Technology Review*. Available at: technologyreview.com/2020/08/20/1007502/uk-exam-algorithm-cant-fix-broken-system/ (Accessed May 2 2021).

Heaven, W.D. (2020). Google's medical AI was super accurate in a lab. Real life was a different story. *MIT Technology Review*, April 27: technologyreview.com/2020/04/27/1000658/google-medical-ai-accurate-lab-real-life-clinic-covid-diabetes-retina-disease/ (Accessed 12 April 2020).

Henriksen, A. and Bechmann, A. (2020). Building truths in AI: Making predictive algorithms doable in healthcare. *Information, Communication & Society*, 23(6): 1-15. DOI: doi.org/10.1080/1369118X.2020.1751866.

Hockenhull, M. and Cohn, M.L. (2021). Hot air and corporate socio-technical imaginaries: Performing and translating digital futures in the Danish tech scene. In *New Media & Society*, 23(2): 302-321. DOI: doi.org/10.1177/1461444820929319.

Høyer, K. L. (2019). Data as promise: reconfiguring Danish public health through personalized medicine. *Social Studies of Science*, 49(4): 531-555. DOI: doi.org/10.1177/0306312719858697.

Jasanoff, S. and Kim, S. (2015). *Dreamscapes of Modernity: Sociotechnical Imaginaries and the Fabrication of Power*.

Jaton, F. (2017). We get the algorithms of our ground truths: Designing referential databases in digital image processing. *Social studies of Science*, 47(6): 811-840. DOI: doi.org/10.1177/0306312717730428.

Jensen, C. B. (2010). Ontologies for Developing Things: Making Health Care Futures Through Technology. In: Hunsinger, J. and Nolan, J. (eds.), *Transdisciplinary studies*, 3. Sense Publishers.

Tversky A. and Kahneman D. (1974). Judgment under uncertainty. *Science*, 185(4157): 1124-31. DOI: 10.1126/science.185.4157.1124.

Kaun, A. and Dencik, L. (2020). Datafication and the Welfare State: An Introduction. *Global Perspectives*, 1(1): 12912. DOI: doi.org/10.1525/gp.2020.12912.

Kjærdsgaard, M. G. et al., (2016). Introduction: Design Anthropological Futures. In: R. C. Smith (ed.), *Design Anthropological Futures: Exploring Emergence, Intervention, and Formation*, 1-16. Bloomsbury Publishing Plc.

Kvale, S. (2008). *Doing interviews*. Sage Publications.

Latour, B. (1983). Give me a laboratory and I will move the world. K. Knorr and M. Mulkay (eds.), *Science Observed*, 141-170. London, UK: Sage Publications.

Latour, B. (1986). *Visualisation and Cognition: Drawing Things Together*. *Knowledge and Society*, 6: 1-40.

Latour, B. (1987). *Science in action: How to follow scientists and engineers through society*. Harvard University Press.

Latour, B. (1988). *The Pasteurization of France*. Boston: Harvard University Press.

Latour, B. (1999). *Pandora's Hope: An Essay on the Reality of Science Studies*. Cambridge, MA, USA: Harvard University Press.

Latour, B. and Weibel, P. (2005). *Making Things Public: Atmospheres of Democracy*. MIT Press.

Law, J. (2004). *After Method: Mess in social science research*. Routledge.

Liu, R., Rong, Y., and Peng, Z. (2020). A review of medical artificial intelligence. *Global Health Journal*, 4(2), 42-45. DOI: doi.org/10.1016/j.glohj.2020.04.002.

Loukissas, Y. A. (2019). *All Data are Local. Thinking Critically in a Data-Driven Society*. The MIT Press.

Mackenzie, A. (2015). The production of prediction: What does machine learning want? *European Journal of Cultural Studies*, 18(4-5): 429-445. DOI: doi.org/10.1177/1367549415577384.

Mackenzie, A. (2017). *Machine Learners: Archaeology of a data practice*. The MIT Press.

Mol, A. (1999). Ontological Politics. A word and some questions. In: J. Law and J. Hassard, *Sociological Review Monograph Series: Actor Network Theory and after*, 47(1) (special issue): 74-89.

Møller et al., (2020). Who Does the Work of Data? *Forum: Health Matters*, 52-55. *Interactions ACM*. DOI: doi.org/10.1145/3386389.

Mønsted, T (2019). Achieving veracity: A study of the development and use of an information system for data analysis in preventive healthcare. *Health Informatics Journal*, 25(3): 491-499. DOI: 10.1177/1460458218796665.

National Services Scotland. (2008). Scottish Patients At Risk of Readmission and Admission: A report on development work to extend the algorithm's applicability to patients of all ages. Available at: isdscotland.org/Health-Topics/Health-and-Social-Community-Care/SPARRA/2008_06_16_SPARRA_All_Ages_Report.pdf.

Nielsen, S. L., Christensen, P. R., Lassen, A. H., and Mikkelsen, M. (2017). Hunting the Opportunity: The Promising Nexus of Design and Entrepreneurship. *The Design Journal* 20:5: 617-638, DOI: 10.1080/14606925.2017.1349983.

Olsen, B. (2021). Offentlige AI-registrer øger troværdigheden i vores smarte byer. *DataEthics*. Available at: dataethics.eu/da/offentlige-ai-registrer-oeger-trovaerdigheden-i-vores-smarte-byer/ (Accessed May 2 2021).

Ozaydin, B., Hardin, J. M., and Chieeng, D.C. (2016). Data Mining and Clinical Decision Support Systems. In: E. S. Berner (ed.), *Clinical Decision Support Systems*. Healthcare Informatics. Switzerland: Springer International Publishing. DOI: 10.1007/978-3-319-31913-1_3.

Pickering, A. (1995). *The mangle of practice: Time, agency, and science*. University of Chicago Press.

Pickering, A. (2016). Art, Science and Experiment. *MaHKUscript: Journal of Fine Art Research*, 1(1): 2, 1–6. DOI: doi.org/10.5334/mjfar.2.

Press, G. (2016). A Very Short History Of Artificial Intelligence (AI). *Forbes*. Available at: forbes.com/sites/gilpress/2016/12/30/a-very-short-history-of-artificial-intelligence-ai/ (Accessed May 2 2021).

Raghupathi W. and Raghupathi, V. (2014). Big data analytics in health-care: promise and potential. *Health Information Science and Systems*, 2(3).

Russell, S. J., and Norvig, P. (2016). *Artificial intelligence: A modern approach*, 3rd Edition. Pearson India Education Services Pvt. Ltd.

Ruppert, E. (2019). Different data futures: An experiment in citizen data. *Statistical Journal of the IAOS*, 35: 633–641. DOI: 10.3233/SJI-190538.

Spooner, S. A. (2016). *Mathematical Foundations of Decision Support Systems*. In: E. S. Berner (ed.), *Clinical Decision Support Systems*. Healthcare Informatics. Switzerland: Springer International Publishing. DOI: 10.1007/978-3-319-31913-1_2.

Stilgoe, J. (2018). Machine learning, social learning and the governance of self-driving cars. *Social Studies of Science*, 48(1): 25–56. DOI: doi.org/10.1177/0306312717741687.

Suchman, L. A. (1987). *Plans and situated actions: The problem of human-machine communication*. New York, NY, USA: Cambridge University Press.

Tupasela, A., and Di Nucci, E. (2020). Concordance as evidence in the Watson for Oncology decision-support system. *AI & SOCIETY*. DOI: doi.org/10.1007/s00146-020-00945-9.

Tupasela, A., Snell, K., and Tarkkala, H. (2020). The Nordic data imaginary. *Big Data & Society*, 7(1). DOI: doi.org/10.1177/2053951720907107.

Wadman, S., Strandberg-Larsen, M., and Vrangbæk, K. 2009. Coordination between primary and secondary healthcare in Denmark and Sweden. *International Journal of Integrated Care*. DOI: 10.5334/ijic.302.