PD and The Challenge of AI in Health-Care

Gyldenkærne, Christopher; Simonsen, Jesper; Mønsted, Troels Sune; From, Gustav

Published in:

Proceedings of the 16th biennial Conference on Participatory Design, PDC'2020, June 15-19, Manizales, Colombia: Participation(s) otherwise

DOI:

[10.1145/3384772.3385138](https://doi.org/10.1145/3384772.3385138)

Publication date:

2020

Document Version

Early version, also known as pre-print

Citation for published version (APA):

Gyldenkærne, C., Simonsen, J., Mønsted, T. S., & From, G. (2020). PD and The Challenge of AI in Health-Care. In C. Del Gaudio, L. Parra, & S. Agid (Eds.), Proceedings of the 16th biennial Conference on Participatory Design, PDC'2020, June 15-19, Manizales, Colombia: Participation(s) otherwise (Vol. II, pp. 26-29). Association for Computing Machinery. P D C <https://doi.org/10.1145/3384772.3385138>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain.
- You may freely distribute the URL identifying the publication in the public portal.

Take down policy

If you believe that this document breaches copyright please contact rucforsk@kb.dk providing details, and we will remove access to the work immediately and investigate your claim.

PD and The Challenge of AI in Health-Care

CHRISTOPHER H. GYLDENKÆRNE, Department of People and Technology, Roskilde University, Roskilde, Denmark, chrgyl@ruc.dk

JESPER SIMONSEN, Department of People and Technology, Roskilde University, Roskilde, Denmark, simonsen@ruc.dk

TROELS MØNSTED, Department of People and Technology, Roskilde University, Roskilde, Denmark, moensted@ruc.dk

GUSTAV FROM, Digestive Disease Center, Bispebjerg Hospital, Copenhagen, Denmark, gustav.from@regionh.dk

In its promise to contribute to considerable cost savings and improved patient care through efficient analysis of the tremendous amount of data stored in electronic health records (EHR), there is currently a strong push for the proliferation of artificial intelligence (AI) in health-care. We identify, through a study of AI being used to predict patient no-show's, that for the AI to gain full potential there lies a need to balance the introduction of AI with a proper focus on the patients and the clinicians' interests. We call for a Participatory Design (PD) approach to understand and reconfigure the socio-technical setup in health-care, especially where AI is being used on EHR data that are manually being submitted by health-care personnel.

CCS Concepts: • **Human-centered computing** → **Participatory design**; *Empirical studies in collaborative and social computing*.

Additional Key Words and Phrases: Participatory Design, Artificial Intelligence, Electronic Health Record data, Primary- and Secondary Use data, precision medicine

ACM Reference Format:

Christopher H. Gyldenkærne, Jesper Simonsen, Troels Mønsted, and Gustav From. 2020. PD and The Challenge of AI in Health-Care. In *Proceedings of PDC2020: Participation(s) otherwise (16th Participatory Design Conference PDC2020 – Participation(s) otherwise)*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Large scale implementations of Electronic Health Record systems (EHR) are transforming the health-care sector into an arena for data capture used for health-care statistics and benchmarking, governance, decision support, efficient resource allocation, and cost savings in general [3, 12]. The increasing amount of EHR data opens for unprecedented potentials for Artificial Intelligence (AI). In this paper, we focus on distinguishing between primary and secondary use of EHR data: Primary use of data here covers data's intended use by the health care worker, such as treatment and delivery of direct care. Secondary use covers statistical use-cases, and other insights made through AI technologies. The prospects of secondary use of big data by AI are, above all, visible and tempting at administrative, economic, and decision maker levels [17, 19]. Hospitals around the world acquire AI systems pre-developed to fit the hospital sector,

Pre-print of final manuscript submitted to ACM. © 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Please **Cite this article as:**

Gyldenkærne, C.H., J. Simonsen, T. Mønsted, and G. From (2020). PD and The Challenge of AI in Health-Care, *Proceedings of the 16th biennial Conference on Participatory Design, PDC'2020, June 15-19, Manizales, Colombia: Participation(s) otherwise*, Vol. II, ACM. <https://doi.org/10.1145/3384772.3385138>

to be re-engineered to different national contexts, and furthermore to be adapted to local needs at the department level of hospitals.

As an example, the largest EHR vendor in the world, Epic, offers a suite of applications for AI and analytics, including a cloud-based and highly configurable platform of machine learning algorithms (epic.com/software). Epic was recently implemented as a hospital-wide EHR system in two out of five regions in Denmark. The transformation from paperbased records to EHR also implied the clinicians having to meticulously enter vast amounts of EHR data through a complicated interface to Epic [9]. Physicians, nurses and health-care staff at all hospital levels are experiencing the data-capturing as a burden to their work focusing on patient treatment and care [6]. After several years of highly voiced critique from the clinicians, hospital management are eager to create success projects that utilize the growing amount of EHR data and demonstrate the value of the clinicians' recording efforts: This forms a starting point for introducing AI in health-care.

An initial project in Denmark was established applying a predictive AI model for patients not showing up for their outpatient surgery at an out-patient clinic. The AI algorithm was applied to historical EHR data and proved being able to predict 98 % of patients showing up and 70 % not-showing up [8]. This is an impressive result that potentially may improve patient treatment, streamline surgery planning, and result in cost reductions, both for the specific out-patient clinic and, with a general prospect, for various outpatient treatments. However, applying AI to these EHR data also discovered that 80 % of the data had to be removed from training-data (i.e. the prediction could only be completed for 20% of the patients) due to incomplete patient records, lack of consistency in the recordings, redundancy, and proximity to events in fatal variables. This ambiguity in the primary use data – often labeled as a “data quality issue” [e.g. [1]] – is a result of different local documentation practices and work-arounds established to cope with a poorly designed and introduced system with cumbersome and ambiguous interfaces.

In this paper we demonstrate how applying AI to EHR data introduce a dual purpose and a potential conflict between primary and secondary use. On one hand, the clinicians' struggle to use EHR as a workable tool to communicate, coordinate, document, and support their treatment and care of the patients. On the other hand, the managerial and administrative interests in utilizing the potential of big data analytics push a demand for detailed, plausible, and proximate EHR data. This calls for a PD agenda resembling the heritage of PD [13]: PD should support clinicians to gain a voice in the technologies that affect their patient treatment and care, introduce a sensitivity to existing work practices and power relations, and contend that the work to obtain AI requirements are not obtained blindly at the expense of the clinicians' patient focus.

This paper is co-authored by researchers and the quality responsible chief physician who have been deeply involved in the project since 2017. The first author conducted the initial project as part of a master's thesis [8], and, since 2018, as a Ph.D.-project. The analysis is additionally informed by another master's thesis project [11]. The empirical activities comprise 4 months of situated work with hospital data management team (administration), 8 work-observations, 6 experience exchange group meetings (department staff, administration and data management team) and 4 PD workshops. In the following, we present the project as an exemplary case of applying a predictive AI model to primary-use EHR data. We discuss the different characteristics of primary- and secondary use of EHR data and conclude by outlining implications for PD.

2 APPLYING AI TO PRIMARY-USE EHR DATA

Our case takes place at the Digestive Disease Centre, named in Danish “Abdominalcenter K” (AK), an out-patient clinic comprising the specialties surgery and gastroenterology/hepatology. 10 % of the patients are not showing up for their surgical appointments at AK. This challenges staff, economy, and patients' access to care. Interventions such as

reminder letters and phone calls have shown to bring down the number of patients failing to show up but this is also very resource intensive to staff, taking away time from clinical work. If AK could predict the patients with high risks of not showing up, they could prioritize and target interventions to help these patients, improve operations planning, and reduce unoccupied operating theatres and empty beds at AK.

In response to this challenge, a predictive AI model was developed at AK and trained on local historical data of patient no-shows extracted from the EHR. Through a process known as feature engineering [15], this data was modelled into 13 variables known from public health literature to be predictive for patient no-show rates [5, 10]. As the most important output, the resulting AI model was able to differentiate between factors related to the health-care provider and factors related to the patient as these would initiate different responses from the clinical staff. During the training of the AI model it was revealed that the history of the patient’s former no-shows critically affected its ability to predict future no-shows. It was important that clinicians reported the reason for why a patient did not receive the planned operational procedure, referred to as the cancellation-reason, through a drop-down menu with 19 different options. In order to predict all patient no-shows, the AI model required the cancellation-reason variable to (1) be completed (that is, entered with a correct value) for all patient cases and (2) that each of the 19 different cancellation-reasons was clearly binary distinguished as either a reason related to the behavior of the patient (e.g. the patient neglected the appointment) or caused by an event from the health-care system, i.e. the hospital or AK (e.g. the patient’s elective surgery was postponed due to an acute patient). While the current practice of recording the cancellation-reason satisfied the clinicians’ primary use, it did not comply with the AI models’ requirements causing the predictions to only be possible to calculate for 20 % of the patient population. In the following, we elaborate the analysis of this specific recording issue. Specifically, we have found that this is primarily due to the different documentation practices at AK and to ambiguities among the cancellation-reason categories.

2.1 Different documentation practices

In the years following the deployment of the EHR at the hospital several updates and re-configurations were released, including the functionality for reporting cancellation-reason. This effectively rendered the guidelines for use of the EHR partially obsolete, which left the clinicians in a situation where they had to continually reinvent their documentation practices [8, 11]. This was particularly visible at three departments that refer patients to surgical procedures, where the documentation practices of cancellation-reasons developed in different ways.

In department A the physicians recorded a cancellation on a hand-written note that they later, by the end of their shift, handed over to the medical secretaries. The medical secretaries then recorded all cancellation-reasons based on translating these notes and with limited additional context information.

In contrast, clinicians at department B refrained from recording cancellation-reasons as this was not relevant to their primary use of the EHR. The reason was a procedure in this department to re-book patients immediately after a cancellation, through another interface-part of the EHR. As a result, the cancellation-reason variable was left uncategorized and therefore unusable to the AI model.

Finally, at department C, the staff deliberately did not record cancellation-reasons: The reason was that the clinicians learned that the system automatically called in patients if no re-booking has been scheduled and the cancellation reason has not been recorded within three days. The EHR then automatically informed the patient and asked for a re-booking through a national e-postal system, and set the cancellation reason to ‘annulled’. The annulled status was unusable for the AI model as it did not enable it to distinguish between patient-related or hospital-related reasons for the cancellation.

2.2 Ambiguity of categories

The 19 different available categories for recording a cancellation- reason was likely a result of the updates continuously introduced in the EHR, to accommodate the various needs of different departments. Some of the categories were irrelevant at AK, including, for example, the category “Birth” (that is, the patient has been giving birth to a child).

The development of these categories was, however, not aligned with the requirements imposed by the AI model: Seven of the 19 categories did not distinguish between binary criteria of patient or hospital related reason for cancellation. For example, one category option was labelled “Rebooked, postponed” which indicated a postponed booking without specifying a reason that could be related to either the hospital or to the patient. Of the 19 categories, “Other” was frequently used. This particular option enabled the user to report an explanation in an adjacent free-text field. Due to the absence of text-mining capability and classification algorithms [4], this data was, however, not useful for the AI. The reason for the extended use of the “Other” option at AK is partly because of ambiguous categories and partly in order to produce better primary data, as the following example demonstrates.

3 DISCUSSION

A fundamental promise for using AIs in health-care is their ability to harness the tremendous amounts of data stored in EHRs and efficiently elicit unique insights. This secondary use entails that data is analyzed in a different context and for a different purpose than the primary uses it was originally produced for. While AI technologies offer unprecedented potential for secondary data analysis in health-care, they are also vulnerable to the availability and quality of data. According to Raghupati and Raghupati (2014) a fundamental challenge for secondary analysis of big health-care data is that the algorithms and system must be able to handle the volume (the sheer amount of data), velocity (the rapid speed by which new data is produced), variety (the various data sources and formats), and veracity of the data (the quality of the data and extent to which it represents the intended phenomenon). While these challenges may in part be solved by developing robust and sophisticated infrastructures for data retrieval and management, our study suggests that the requirements to the data and the data infrastructure imposed by the AI in some cases are altering the conditions for, the work involved in producing and using the data performed by clinicians. In other words, to ensure an effective deployment and use of AI in health-care it is necessary to align the primary and secondary uses of the data. The velocity and variety of the data required by the AI in our case proved to be of little significance for the challenge of producing and retrieving the data necessary for the AI. While the velocity of data production may be an important challenge for AIs used to perform real-time analytics, for instance as part of diagnostics processes [16], the prediction performed by the AI at AK was not time-critical and did not depend on a constant inflow of clinician-produced data. This enabled the clinicians to maintain a great deal of flexibility to align data production with the rhythms of their local practices. In similar lines, the variety of data is in many cases a significant challenge for AIs in health-care. This implies that data is produced and stored in various formats, often in different data repositories [2]. In particular unstructured or semi-structured data, including progress notes and admission and discharge letters, require that AIs are able to cover and convert a variety of formats to enable the secondary analysis to take place [16]. In our case, however, the variety of data was reduced to a bare minimum and the data retrieval was efficiently supported by the existing EHR system that is designed as a platform for data-sharing among different sources in the Danish health-care system (i.e. brings in data from central governmental registration of citizens such as date of birth, gender, etc.). In contrast to velocity and variety, the volume of data was a challenge for the AI. This was essentially not a technical challenge, as the limited amount of data needed to produce precise prediction of patient no-shows was well handled by the technical infrastructure. For

the clinicians, an extra workload imposed to increase this data production, caused challenges. First, it introduced yet another type of work in an already busy documentation practice, hereby taking away time from patient-related activities. Furthermore, since registration of reasons for no-shows had no immediate benefit for many clinicians, this effectively created a disparity between work and benefit – between primary and secondary use – closely resembling what is known to be a challenge when designing collaborative systems [7]. Because of this, clinicians at all three departments found ways to manage patient no-shows that were meaningful in their clinical practice. While department A maintained registration of the data, albeit in a way that reduced quality of the data, department B and C found work-arounds, either by completely avoiding registration of no-shows or by resorting to the standard functionality in the EHR that automatically report a no-show three days after the appointment, i.e. without requiring human intervention. Hereby, two out of three departments did not contribute to the accumulation of a volume of data necessary for the AI, effectively resulting in a less precise performance of the algorithm. Furthermore, if 80 % of the data is useless for AI purposes, there is a risk of the model for prediction being heavily biased towards conclusions that does not represent the general patient population. The different procedures of data production were adapted to fit with the documentation practices at the departments driven by the interests of primary uses of the data. This adaptation has evolved without a clear conception of how workarounds and local practices affect the intended secondary use of data. An implication for PD would be to enable clinicians to take the interests of the AI into account. The AI model's requirement of a sufficient volume of data needs to be acknowledged and accepted by the clinicians. The requirement for data volume should be aligned with primary use by (1) efforts to minimize the work involved in this data production and (2) a documentation procedure designed to provide concurrent benefits of the data for clinical practice. The veracity of the data surfaced inherent tensions between, most notably, the AI's need for detailed and unambiguous data and the heterogeneity of information practices. The need for veracity of data for AIs entail that the data on which the analysis is based must be error-free, credible and offer a precise representation of the phenomena that the AI is assumed to offer insights into. The notion of veracity assumes that the meaning or semantics of the data is shared by all producers. This is, however, often not the case as categories sometimes produces different meanings for different groups of health professionals as well as for patients [14]. As a practical consequence, different health professionals may document identical activities in different ways. Traditionally, this has not been a significant issue as data in health-care is predominantly produced to serve its primary use, such as support communication between nurses, document specific examinations of the patient, etc. The precision of the data analysis performed by AIs typically depends on the granularity of the data as they generally operate better with a high number of unique categories. This may impose restrictions on how health professionals produce data, for instance by requiring them to document using aggregated, different, and possibly more specific categories than they would otherwise do. In our case, the AI model introduced new requirements to the use of the 19 existing categories for cancellation-reasons. From the perspective of primary data use, the 19 categories were partially overlapping and, as shown in the case, different uses of these had developed in the three departments. To alleviate the need to synthesize some of the categories to express a cancellation-reason that was meaningful from a clinical perspective, clinicians would often resort to using the input field 'other'. This, however, effectively decreased the veracity of the data from the perspective of secondary data use. The inconsistent use of the 19 categories, the fact that seven of the categories did not differentiate between the two factors required by the AI model, and the extensive use of the 'other' category crippled the ability of the AI model to differentiate between the if a no-show was caused by the patient or by the hospital.

The AI model in our case require a reconfiguration of the categories to unequivocally distinguish cancellation-reasons as caused by patient behavior or hospital events. To PD this presents an obligation to align a reconfiguration of

categories needed by the AI with an efficient and meaningful documentation practice for clinicians. In other words, the participation of clinicians' is important in order to (1) clarify the extent to which the AI-required binary distinction can differentiate cancellation-reasons, and (2) to design a consistent category system serving both primary and secondary data use.

4 CONCLUSION

The dissemination of AI is pushed for many reasons and from many stakeholders. AI alter the conditions for human entered primary-use EHR data. AI introduces a dual purpose of EHR data where the primary-use from clinicians treating patients may compete with a secondary-use from AI governed by administrative and economic concerns. The role of PD is to ensure the voice of the user and end-user, that is, balance the introduction of AI by maintaining a proper focus on the patients and the clinicians closely involved in their treatment and care. The premise for PD is to establish "shared and agreed-upon goals" [18, p. 5] and to enable a tight alignment between primary- and secondary-use of EHR data. As we have demonstrated, AI introduces an occasion where PD might contribute by designing re-configurations of work and technologies that benefit both types of use and purpose of EHR data.

REFERENCES

- [1] Taxiarchis Botsis, Gunnar Hartvigsen, Fei Chen, and Chunhua Weng. 2010. Secondary use of EHR: data quality issues and informatics opportunities. *Summit on Translational Bioinformatics 2010* (2010), 1–5.
- [2] Bendik Bygstad, Ole Hanseth, and Dan Truong Le. 2015. From IT silos to integrated solutions. A study in e-health complexity. (2015).
- [3] RD Canlas. 2009. Data mining in healthcare: Current applications and issues. *School of Information Systems & Management, Carnegie Mellon University, Australia* (2009).
- [4] Hsinchun Chen, Sherrilynne S Fuller, Carol Friedman, and William Hersh. 2005. Knowledge management, data mining, and text mining in medical informatics. In *Medical Informatics*. Springer, 3–33.
- [5] Richard A Deyo and Thomas S Inui. 1980. Dropouts and broken appointments: a literature review and agenda for future research. *Medical care* (1980), 1146–1157.
- [6] Atul Gawande. 2018. Why doctors hate their computers. *The New Yorker* 12 (2018), 1–31.
- [7] Jonathan Grudin. 1994. Groupware and social dynamics: Eight challenges for developers. *Commun. ACM* 37, 1 (1994), 92–105.
- [8] Christopher H. Gyldenkærne, Gustav From, and Jakob G. Simonsen. 2020. Predicting Patient No-Shows in Endoscopy: A Case of Situated Machine Learning Development with Imperfect Data. *Submitted Spring* (2020).
- [9] Morten Hertzum and Gunnar Ellingsen. 2019. The implementation of an electronic health record: Comparing preparations for Epic in Norway with experiences from the UK and Denmark. *International journal of medical informatics* 129 (2019), 312–317.
- [10] Ulf Kåre Jansbøl and Rasmus Trap Wolf. 2016. Reduktion af udeblivelser: Et systematisk review om effekt og besparelse ved påmindelser og gebyr.
- [11] Simon Rove Jensen and Bundgaard Silke. 2019. *Et studie af patientafbud og udeblivelsesregistrering i Sundheds-plattformen, på Bispebjerghospital, abdominalcenter K*. Master's thesis.
- [12] Fei Jiang, Yong Jiang, Hui Zhi, Yi Dong, Hao Li, Sufeng Ma, Yilong Wang, Qiang Dong, Haipeng Shen, and Yongjun Wang. 2017. Artificial intelligence in healthcare: past, present and future. *Stroke and vascular neurology* 2, 4 (2017), 230–243.
- [13] Finn Kensing and Joan Greenbaum. 2013. Heritage: Having a say. In *Routledge international handbook of participatory design*. Routledge, 21–36.
- [14] Troels Mønsted. 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 (2019), 491–499.
- [15] Fatemeh Nargesian, Horst Samulowitz, Udayan Khurana, Elias B Khalil, and Deepak S Turaga. 2017. Learning Feature Engineering for Classification.. In *IJCAI*. 2529–2535.
- [16] Wullianallur Raghupathi and Viju Raghupathi. 2014. Big data analytics in healthcare: promise and potential. *Health information science and systems* 2, 1 (2014), 3.
- [17] Joachim Roski, George W Bo-Linn, and Timothy A Andrews. 2014. Creating value in health care through big data: opportunities and policy implications. *Health affairs* 33, 7 (2014), 1115–1122.
- [18] Jesper Simonsen and Toni Robertson. 2012. *Routledge international handbook of participatory design*. Routledge.
- [19] Matthew A Waller and Stanley E Fawcett. 2013. Data science, predictive analytics, and big data: a revolution that will transform supply chain design and management. *Journal of Business Logistics* 34, 2 (2013), 77–84.