Shifting Concepts of Value

Designing Algorithmic Decision-Support Systems for Public Services

Naja Holten Møller
Department of Computer Science,
University of Copenhagen,
Copenhagen, Denmark
naja@di.ku.dk

Irina Shklovski
Department of Computer Science and
Department of Communication,
University of Copenhagen,
Copenhagen, Denmark
ias@di.ku.dk

Thomas T. Hildebrandt
Department of Computer Science,
University of Copenhagen,
Copenhagen, Denmark
hilde@di.ku.dk

ABSTRACT

Calls for responsible design in algorithmic decision-support systems, especially those used in public services, are increasingly common. While an algorithmic system might promise greater precision and efficiency in domains such as critical care, the same efficiency is difficult to replicate in the public service domain, where caseworkers must exercise discretion in applying complex legal frameworks that directly affect individual lives. In this paper we examine the challenges in responsibly designing such an algorithmic decision-support system. We report findings from a large cross-disciplinary research project, aiming to develop an algorithmic component for municipal decision-support systems for job placement in Denmark. Our data showcases insights from how a team of data scientists, caseworkers, and system developers negotiated notions of value metrics and usefulness in a participatory design set-up. Whereas data scientists expected to focus on profiling individuals, the caseworkers instead pushed for systems that could help mitigate organizational contradictions and obscured processes in casework. We close with a discussion of challenges and future directions for participatory algorithmic systems design in municipal contexts.

CCS CONCEPTS

• **Human-centered computing** → Collaborative and social computing; Empirical studies in collaborative and social computing.

KEYWORDS

Algorithmic decision-support systems, Public services, Responsible design, Casework, Job placement, Participatory design methods

ACM Reference Format:

Naja Holten Møller, Irina Shklovski, and Thomas T. Hildebrandt. 2020. Shifting Concepts of Value: Designing Algorithmic Decision-Support Systems for Public Services. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society (NordiCHI '20)*, October 25–29, 2020, Tallinn, Estonia. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3419249.3420149>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

NordiCHI '20, October 25–29, 2020, Tallinn, Estonia

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7579-5/20/10...\$15.00

<https://doi.org/10.1145/3419249.3420149>

1 INTRODUCTION

Algorithmic decision-support in public services is gaining increased interest in the broader field of Human-Computer Interaction (HCI) [49] as a phenomenon fundamentally shaping society and how decisions are made. Scholarship in this area thus far has produced many insights. Algorithmic decision-support systems are considered in safety-critical domains such as diagnostics and critical care [15], crisis management in due time [39], and containment of disease [27]. Yet, the value of algorithms for decision-support is disputed across domains. This is especially acute in the domain of public services, where the focus is on improving an individual's status or situation as caseworkers and recipients of support and benefit interact. We see decision-support in public service as a domain in its own right [2, 12, 44, 52]. For systems operating in a context of critical decision-support (e.g. critical care of patients), it is more straightforward to “reliably” establish metrics for determining if an algorithmic intervention has improved an individual's situation. The processual and long-term nature of much public services casework, on the other hand, makes algorithmic decision-support tools challenging for this domain, because what constitutes improvement and successful intervention and other such value metrics can be difficult to establish due to the complexity of the kinds of processes that characterize casework. The value metrics for different types of casework (e.g. counseling) are challenging to identify, since some people can remain part of the system for years, and an individual's situation can stabilize or de-stabilize many times. Metrics can show an improvement in the public registered status of a person, the number of complaints over a process, or if criteria were met for due process [3, 34, 36]. However, disentangling the role of algorithms and such metrics is more complicated. Consequently, it is critical that both design teams and policymakers understand the practices of caseworkers, what Lipsky has termed “street-level bureaucrats” [30], and how they leverage discretion in their decision-making. This understanding is critical for ensuring responsible implementation of algorithms for decision-support in public service processes that are often long-term.

The idea of “street-level bureaucracy” as a way to ensure the adaptation of bureaucratic rules into practical action has been previously explored in HCI and related areas of research, with a focus on developing recommendations for the design of public service algorithmic decision-support systems [7, 10, 19, 31]. Scholars have grappled with the problem of inflexibility in the “street-level algorithm” [1], where discretion can (still) ensure proportionality in legal decisions that affect individual lives. For algorithmic systems to have a positive impact, they must be flexible [1]. This has led

others to argue that algorithms can be treated as analogous to impartial bureaucratic rules for controlling action [46]; however, a need for human discretion persists in decision-making processes where the most uncertainty about the potential operation of an algorithm exists [46]. According to Pääkkönen et al., algorithms can be useful even if they are inflexible, but it is the uncertainty about the operation of algorithms that can guide us in determining the parts of a process where algorithmic decision-support systems can be most valuable [46].

A key concern underlying this research is the processes through which algorithms are assigned power over human decisions and the social implications of algorithms [32]. Public service organizations increasingly seek to use the large, interlinked data sets at their disposal to make informed decisions in the spheres in which they operate. The scale of these datasets makes it impossible for such decisions to be made by humans alone [32 p. 1058]. From this perspective, value is already ascribed to the use of algorithmic decision-support in public services through the expectation that organizations will transform data into knowledge and informed decisions will lead to greater efficiencies in the use of resources. Data then becomes a promise of future bureaucratic efficiencies [37]. Metrics such as “performance” of an organization demonstrate how value is ascribed to algorithms following logics of bureaucracy. Pine & Liboiron [42] suggest that metrics are harnessing data science to govern both caseworker and organizational performance and fundamentally engineer work processes to fit programmatic decisions. Boulus-Rodje finds that when the caseworker’s mandate is to promote individuals’ welfare, public services may have various performance measures and clear end-goals, but in practice these come to be interpreted and measured in many different ways ([7] *following* [43]).

Despite the wealth of research on algorithmic decision-support, the complexity of the issues at hand and the difficulty of gaining necessary access to study the empirical processes where value is assigned to algorithms [41] has limited opportunities for research. In order to disentangle exactly what kind of role algorithmic decision-support systems can and should play in such critical societal contexts, we must understand the underlying negotiations of the value metrics that lend power to algorithms.

This paper presents findings from a large, cross-disciplinary research project in Denmark working to develop an algorithmic component for municipal decision-support systems, focusing in particular on municipal job placement services. In Denmark, caseworkers work with individuals, support them to advocate for their perspective and provide relevant information on the law that apply in their case, while also assembling the case documentation. We were part of a project creating decision-support for caseworkers. Our role as HCI scholars was to amplify the participatory and ethical aspects [9, 48] of the development of the algorithmic component by involving caseworkers in the design process. In particular we sought to investigate how caseworkers, data scientists, and system developers negotiate value metrics in the design of the algorithmic component. The main focus driving our research was to identify the key challenges for a responsible design practice suitable for the specific context of public services and casework.

At the time of the research project, the national Agency for Labour Market and Recruitment designed and implemented a different algorithmic component focusing on risk prediction for long-term unemployment—a concrete example of how algorithms are expected to support decision-making in job placement [41]. The municipality with which we worked on the research project articulated a similar interest at the outset of the project. The legality of the national project in Denmark was criticized for not being passed by the parliament at the time of its initiation and for being hasty [33], and this also impacted the discussions within the design team about value metrics in the case we studied. We found that when asked to describe where algorithmic decision-support could be valuable, caseworkers in public services tended to upend the expectations and assumptions of risk prediction connected to the individual in their discussion with our team of data scientists and system developers. Within our team this contestation by the caseworkers spurred an important debate about the establishment of different value metrics that could help steer the large research project. Where data scientists and the municipality expected to focus on profiling individuals, the caseworkers instead pushed for systems that could help mitigate organizational contradictions and clarify casework processes. This suggested a rather different notion of value for algorithmic decision-support systems in this context than the one outlined by the municipality – as well as the one adopted by the national Agency for Labour Market and Recruitment. Thus in this paper our main contribution is providing empirical insights into how value metrics for algorithmic systems come to be negotiated in a participatory design set-up in a politicized context. We provide suggestions for a pragmatic approach to responsible design of algorithmic decision-support systems.

2 RELATED WORK

Leaving aside the general hype around algorithms [35], a particular challenge for HCI-researchers interested in public service design concerns our commitment to professionals, in our case the caseworkers, as active participants in the shaping of the algorithmic decision-support systems that impact them [25, 48]. Research in HCI, and in particular Participatory Design (PD), has focused on shaping emergent technologies to support the practices and needs of those whose job it is to accomplish a given task or activity [13, 14]. Recent studies in this area consider “acceptance” of algorithmic decision-support systems by practitioners [15]. Researchers looked at what information the practitioners desired when asked to make sense of an algorithmic diagnostic agent for decision-support in cancer diagnostics. Acceptance in this study was considered in terms of how a process can provide practitioners the necessary understanding of algorithms at the *outset* of their introduction to them.

In the Cai et al. [15] study, when asked about their desires the medical practitioners negotiated metrics of value and success for the diagnostic algorithmic agent. Some of these metrics would also apply to a human and follow the same logic as performance metrics. For example, the medical practitioners wanted to know the intended utility over the user’s status quo, measured as the algorithmic agent’s efficiency, accuracy, consistency, etc. The medical

practitioners thus insisted on anchoring metrics for the algorithmic agent’s performance in human-relatable terms so they could make sense of how it improves their practice. However, they also formulated metrics that do not easily translate into performance measurement. For example, they wanted to know whether the algorithmic agent was tuned to work with a human partner and to what extent.

These recent studies of public services make clear that the concerns of PD, once articulated, are again relevant to consider. We observe how the relationship between the workplace and the employee is constantly re-negotiated - with employees having less influence regarding the various kinds of outcomes [26, 50]. In public services, practices are not merely negotiated for the individual workplace, but also in and through a paradigmatic shift to digital-ready law, initiated to make way for the new forms of digitalization that we have come to recognize as algorithmic decision-support systems [41]. Several European countries, including Denmark, Norway, and Sweden, apply a principle of digital-ready law. In Denmark the reform ensuring digital ready legislation explicitly called to reduce the role of discretion and subjective criteria in decision-making [23, 45], paving the way for algorithmic decision-support as well as automated decisions. What is not yet clear from more recent scholarship on the social impact of algorithms on professional work is how to involve those whose work is subject to change [29].

The issue at hand is thus how we as scholars working closely with data scientists and system developers can inform the processes that are critical to tempering the technological possibilities of algorithms and highlighting the concerns for practitioners, whose work practices are being re-negotiated as part of an overall change of how decisions are being made. From this perspective, it becomes even more critical to identify the unique characteristics of casework and discretion to allow for a reasonable worker contestation of the conclusions being drawn by design teams from a participatory setup.

2.1 Casework in Public Services

Caseworker practices have developed over time and, as with many other professions (e.g. clerks), have developed into occupations with little or no protection of their formal title [6, 35]. The role of caseworker varies across specializations, but overall the caseworker takes on individuals’ cases and support them to advocate for their perspective and provide relevant information on the law that applies to their case. This highly politicized context means that casework historically has reflected debates on who are the “deserving poor” in society – a group that has expanded in some periods and contracted in others [20]. The flux in the role of caseworkers in this sense reflects the changing politics in public services. Recent studies of casework in Denmark show the more problematic aspects when only some caseworkers are trained as social workers [7]. This understanding of casework as more or less an arbitrary function, rather than a fully-fledged profession, has come to fuel public debates on algorithmic decision-support in public services as more accurate and therefore more just in bureaucratic terms. In recent years, policymakers increasingly have been trying to resolve this ongoing negotiation and conflict over who falls within the group of

deserving poor through algorithmic decision-support, thus limiting the role of caseworkers.

Lipsky [30] coined the notion of discretion to describe the role-specific expertise and authority that must be leveraged when professionals such as caseworkers apply the law and take decisions that affect individuals. In Denmark, caseworkers traditionally have been given a great deal of autonomy to exercise discretion [34, 41]. In public services, the large datasets, amassed by municipalities and made available for algorithmic systems, are the outcome of many hours of work by public officials, such as municipal caseworkers, as they produce data about individuals and apply the law through the ongoing eligibility categorization that characterizes caseworker systems. The exercise of discretion is not only a matter of ‘capturing’ the individual situations, but also reflects the fact that the prescribed categories are not always meaningful - while a legal decision has to be made [4, 8, 36, 51]. After all, categorizations are never perfect and those people that tend to fall between the categories are often the ones that require more support through municipal services. As Bowker and Star have amply demonstrated, ‘torque’ coming from not fitting into categories can be devastating [8]. The resulting database structures in public services can be read as manifestations of the different kinds of categorization that take place as part of discretion, authorizing certain events [14]. However, without involving the domain experts (the caseworkers), the explanations behind categorizations cannot be well understood in the design of systems for decision-support. As algorithmic decision-support systems are becoming increasingly common, the practices that shape the available dataset are important to understand in order for systems designers and data scientists to work through the structuring of data in a way that is sensitive to practice and how the data was produced in the first place.

Caseworker jobs require balancing discretion and programmatic decisions, based in the dominant understanding of the society within which they operate, on who qualifies as the deserving poor. As such decisions, and the resulting dataset, are often contingent on a range of concerns, it is easy to see how in public services algorithms can come to represent impartiality in these kinds of decisions [20]. This points to the long-term debate in HCI on the role of domain experts in the design process. While domain expertise can often be a requirement for the design process, their involvement can tend to become a box-checking exercise without any real agency or influence on the design. Legitimacy, however, cannot simply be obtained from the fact of having the domain experts involved. Dourish [18] calls for HCI to move beyond this kind of legitimacy trap (what the PD community has termed a “false consensus”¹). Algorithmic components and agents have the potential to fundamentally re-configure casework, and for this reason “acceptance” by the domain experts is insufficient as we set criteria for design. If an algorithmic support system is to fundamentally alter how a particular work process is completed, then what constitutes domain expertise will likely change as well, in ways that would be difficult to predict for the current domain experts. Thus, their acceptance of a potential algorithmic system ultimately is unlikely to be particularly informative. Along these lines, Lustig, Pine, et al. [32] call for more research in the processes that ascribes power

¹<https://pdc2020cpsr.wordpress.com>

and thus value to algorithmic decision-support and explorations of the social implications of algorithms.

As large datasets become available across contexts, which in turn means that organizations are expected to make use of them, the role of algorithms for decision-support emerges at the forefront of scholarly debates on public services [32, 49, 52]. With the increasing push towards automation, the purpose and value of discretion is difficult to grasp. One of the more notable ways that scholars have re-stated this issue is Light & Seravalli's questioning of the municipality as a caring platform [28]. Writing in the context of public services, they point to the shift where instead of investing in civil society, policymakers are downsizing through the older structures of the "public" and replacing politics with metrics. "Care" and caring features oppose metrics from this perspective. It is of course possible to have care duties, but there appears to be little interest in conducting them. Caring through data [24] has become a phenomenon with increasingly data-driven services. In casework, discretion provides an important entry point for how relatives and support persons can leverage care by contesting or providing supplementary data as large datasets are rolled up for programmatic decisions in public services [33].

Drawn together, we learn from prior research that the concept of value metrics, so important in the design of algorithmic systems, is not monolithic and tends to be oversimplified. **We cannot and should not reduce the concept of value metrics to performance measurements alone.** Instead, in this paper we explore how we might qualify value metrics by making clear 1) the reasonable influence of caseworkers on the advancement of their occupation through 2) contesting the design of algorithmic decision-support of public services in relation to 3) the ethos of municipalities as platforms for the exercise of collective responsibility for each other.

3 CASE AND METHOD

This study was conducted as part of a large cross-disciplinary research project between 2017–2020 (ongoing) with the goal of designing an algorithmic component or "agent" for municipal decision-support systems. We conducted our studies at a job center, which serves a mid-sized municipality of approximately 69,000 individuals. In 2019 5,655 individuals in the municipality were unemployed. The municipality's ambition is to be "a valuable partner for skill development, recruitment and job placement, so that all unemployed individuals are assisted in finding a new job"². Furthermore, the municipality aims to "cooperate with businesses on social investments, so that more marginalized individuals can find meaningful jobs, training or practical work experience which matches their resources and skills" [21, 22]. This ambition is defined through the unemployment rate reduction goal of less than 11%, which was the 2019 level of individuals receiving financial support (in Danish *offentlig forsørgelse*) in this municipality. Similarly, according to the municipality's Employment Strategy 2020, the percentage of long-term unemployed people receiving financial support should be reduced from the 2019 level of 6.8% and the percentage of unemployed non-western immigrants and descendants receiving financial support should be reduced from the 2019 level of 16% [21].

The job center that we studied employed 107 caseworkers. 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. This paper reports on the part of the study organized around participatory design workshops between 2018–2019 that engaged with the relevant caseworkers in job placement offices. As many other public officials, caseworkers are continuously measured on productivity, and taking any time away from their regular tasks interferes with how they make themselves accountable to their management. As a result, we agreed with the job placement management that caseworkers could participate in three officially sanctioned workshops (so far) of 2 hours each. 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.

Due to the highly complex context of job placement, the study took extra consideration of research ethics:

- Caseworkers involved in the study were asked permission whenever the field researcher recorded meetings or design workshops.
- Contact information on researchers was provided to allow the caseworkers to correct, add, or delete information or ask questions about the research project.
- Access to data on unemployed individuals (16,000) was restricted to the part of the design team with a special legal and ethical clearance, through working directly with the database and municipality

To gain the necessary understanding of the context and case management system in use, the first author conducted a 2-day field study (8h) of casework in the municipal job placement (A), prior to the participatory design workshops with caseworkers in April 2019. A prior field study (84h) of job placement (B) in a different municipality between 2015–2016 [34] also formed the background for the participatory design workshops. For example, prior findings allowed us to use the concepts derived from caseworkers in our workshop design [47]. Having two studies meant that we were able to validate and deepen our understanding of casework and its associated processes across several job centers. Before and after the participatory workshops we conducted 9 in situ and semi-structured interviews with four caseworkers, a manager an IT-specialist, two data scientists and one coordinator between 2018 and 2019. We did this in order to understand the context and case management system in use. We also held regular meetings with the job placement management throughout the project.

The job placement caseworkers that participated in the workshops had several years of experience (only one was newly educated). Some were trained social workers, but others had diverse backgrounds, originally trained as, for example, teachers. All interviews were recorded and transcribed for later coding and analysis, based on our discussions across the team of data scientists, system developers, the municipality, and HCI scholars. In particular, the data analysis developed from a focus on the "usefulness" of the algorithmic component toward a more specific focus on how caseworkers are brought into the design process in a way that addresses and defines responsibility.

²<https://rar-bm.dk/media/9719/beskaeftigelsesplan-gladsaxe.pdf> (accessed May 2020).

Figure 1: Fieldwork and workshop progression in the study

4 JOB PLACEMENT

In the context of job placement, casework is the main point of human contact, providing the opportunity for principles such as individual agency and accountability to enter an otherwise highly regulated bureaucratic process. Caseworkers and individuals seeking work engage to match the individual's need for support with the legal and administrative rules for eligibility that authorize that support. The decisions are then registered in a digital case management system. The data generated as a result of this process produces large datasets that are available for the design of algorithmically based decision-support systems. However, it poses potential challenges and risks to how caseworkers concern for the individual's situation is expressed when data are rolled up for algorithmic decision-support. In the following we provide a brief overview of the context of the study and consider challenges and risks associated with making job placement data algorithm-ready.

In Denmark a municipal caseworker is formally responsible for documenting each case for job placement, including categorizing an unemployed individual according to national criteria (target groups), and prescribing who is eligible for unemployment benefit programs. The decision on how to categorize an individual within the caseworker management system has implications for the particular type of unemployment benefits to which an individual is entitled. It is a highly regulated process, but one that has been shown to sometimes vary across the local municipalities enacting the national policy.

In practice, the categorization of an individual is initiated at the time of application for job placement support and is determined based on general criteria for eligibility (e.g. residency), reason for unemployment (e.g. personal- and health issues), and age (e.g. < 30y). As the individual either "leaves" job placement or "remains," the target group shifts to reflect the kinds of support to which they are entitled and need. The national Agency for Labour Market and Recruitment defines the overall criteria for categorizing individuals for ordinary benefit eligibility:

- Target group 6.1 Individuals with unemployment insurance
- Target group 6.2 Job ready individuals >30y
- Target group 6.3 Activity ready individuals >30y
- Target group 6.4 Individuals eligible for education <30y
- Target group 6.5 Activity ready individuals <30y.

It is the caseworker's job to assess whether the right category has been assigned to the individual as part of documenting the case, as regulated by the Active Labour Market Policy Act. This

also intersects with the Consolidation Act on Social Services' rules and regulations guiding the benefits for which an individual can apply. Thus, another set of categories apply for individuals that require resource development (target group 6.8), a job with reduced hours (target group 6.9), re-training (target group 6.10) or early pension (target group 6.11). In particular, target group 6.3 reveals the patterns across individuals that are later categorized as one of these extra target groups. These individuals tend to be at risk of long-term unemployment and so became a special focus in the municipality with which we worked.

4.1 The Job Placement Dataset as a Basis for Algorithmic Design

Our overall research project goal was to design an algorithmic component, or "agent," for municipal decision-support. The primary source of data in the development of this algorithmic component was the municipal case management system described above. This case management system in job placement retains data from the entire case history of an individual as well as information such as gender, education, ethnicity, and other demographics. Key "events" registered by caseworkers include: assigning the unemployed individual to a specific target group (e.g. 6.2 "Job ready", 6.3 "Activity ready" etc.); "offers" or services (job training etc.) agreed on with the individual; illness; planned, cancelled, and completed meetings between the individual and the caseworker; and reason for closing a case.

The data on individuals are not readily available for algorithmic use, what we call "algorithm-ready", due to several challenges produced by the complex and shifting political context of job placement. Furthermore, there are pending legal challenges for engaging available data on individual job placement, illustrating why the construction of algorithmic decision-support in public services is by no means trivial when deciding on relevant value metrics. Understanding the data requires knowledge about its context; in this case, how politics, and hence legal rules and regulations, have shifted over the years with concrete impacts on caseworkers' practices and resulting database structure.

Challenges that we experienced in the design team when deciding how to engage the dataset on unemployed individuals included:

- Differences in registration and documentation practices, as well as types of "offerings," reflecting the changes in the law and politics, which could mean that the historical traces will no longer correspond to current traces in the data.

- Difficulties interpreting the assignment of target groups for the same individual, as well as the start and end time stamp, denoting the period for which the person is assigned this target group and when it reflects the continuation of an earlier target group registration or a new period in the database.
- Translation between the database structure generated from the case management system code in English and the corresponding interface in Danish used by caseworkers for data registration.

The practical work involved in making sense of the database structure required the team of data scientists and system developers to work closely with the caseworkers who are responsible for the day-to-day construction of these data. The municipality and data scientists initially wanted to explore profiling individuals based on their risk of long-term unemployment, which is a typical request in other contexts for the development of decision-support systems. Yet our data demonstrate that, from a caseworker perspective, the benefits of profiling individuals were less clear. We now turn to the findings from our process with caseworkers and how value metrics were shaped through our interaction with caseworkers.

5 FINDINGS: SHIFTING CONCEPTS OF VALUE IN PUBLIC SERVICE DESIGN

Algorithms for decision-support in public services can be used by municipalities and national authorities for profiling individuals seeking support. The Agency for Labour, Market, and Recruitment were already experimenting with algorithms for the prediction of long-term unemployment at the time of our participatory workshops with the municipal caseworkers. The discourse in Denmark, as in many other countries, is often divided over questions about the use of individual data. It was critical for us across the team of data scientists, system developers, and HCI scholars to emphasize that our common goal for the workshops was to give the caseworkers a say over the development and use of algorithmic decision-support technologies in job placement.

Identification and selection of appropriate value metrics became a primary concern for the design team from the beginning of the research project. Defining appropriate value metrics delineated the goals of the algorithmic system. However, defining a common value metric for the research project became complicated despite our agreement on the importance of getting the process “right” by taking a participatory approach that involved caseworkers. Our process with caseworkers demonstrated that the concept of value metrics, so important in the design of algorithmic systems, is not monolithic and tends to be oversimplified. We show how and why we cannot and should not reduce the concept of value metrics to performance measurements alone. Prior research shows that such oversimplification in classification can materialize within algorithmic decision-support systems as discrimination against specific people [11].

In the following example from one of the participatory workshops with caseworkers, the lead data scientist attempted to engage the municipal caseworkers by taking a participatory approach through staging “usability” as a main driver for design. At this point of the design process, the design team was still discussing the data structuring method. A concrete challenge was how to make sense

of the dataset of 16,000 individuals and to understand how we could “dig out” valuable insights. On this particular day, we convened in the usual meeting room of the job center. The lead data scientist had prepared a presentation in advance, which attempted to explain his concept of value metrics to the caseworkers and the concrete input we were looking for from their side in this 2h workshop:

Data scientist: “What I would like your input on ... including also what are the questions we should be asking [to the dataset] in order to understand more about this ... your systems and data ... what are the processes, what are your concepts of “value” in this context [casework].. and how can algorithms be useful – and not useful. In my world ... we [data scientist] simplify things. And, this [pointing to screen dump of database] is the data world that I see.”

(Workshop 27.05.2019)

Caseworkers’ complicated processes were important to understand in order to move forward in the research project. The data scientist was trying to explain that for the design team it was important to identify the points of uncertainty in the decision-making processes from the point of view of caseworkers. This would provide a concrete way to move forward in the structuring of data, necessary for designing an algorithmic component for the case management system. Structuring data is fundamentally a sense-making process. The data scientist was looking for the subprocesses in casework where decision-support would make a difference. It was clear to him from the beginning that a data scientist’s concept of value metrics needs translation. In an earlier workshop he explained this need in the following way:

Data scientist: “Let’s say I can choose between [route] A or B.. but it does not have a “cost” in any value metric ... A value metric could be money or it could be pleasure of working or something different.. but there has to be a consequence of taking the less optimal decision.. these are the subprocesses we are looking for ... where decision-support has a consequence.”

(Workshop 22.03.2019)

Rather than imposing the expected ideals of efficiency and cost reductions, the data scientist attempted to broaden the idea of performance metrics. For example, the municipality performance metrics for job placement follow a typical understanding of value as “money” or cost that can be saved if fewer individuals are long-term unemployed as laid out by the official municipal Employment Strategy. The data scientists in the project were looking to understand the caseworker point of view and make fewer assumptions about the ‘right’ kinds of metrics, but the focus remains on “usable” as a proxy for “useful”. What we soon discovered in the participatory workshops with caseworkers was how value metrics are a concept that not only requires translation, but takes on an entirely different form that is closely tied to caseworkers’ discretion as part of how they apply bureaucratic rules across individuals. As we show in the following section, the dialogue with caseworkers goes far beyond the idea of “usability” as a shared discussion point of the value metrics to steer by in the process of designing an algorithmic component for case management systems that follow a responsible approach.

5.1 Setting Criteria for Legitimacy

Our dialogue with caseworkers soon demonstrated that the negotiation of a shared value metric tended to quickly turn into a discussion about the legitimacy of an algorithmic component itself. Caseworkers did not immediately engage in a discussion of usefulness of the algorithmic component for their occupation. Instead they questioned the legitimacy of algorithmic decision-support of public services and what may be considered as the ethos of municipal job placement.

In the semi-structured interviews conducted prior to the workshop, caseworkers frequently contested the idea that risk prediction was useful in their work (Interview w/ caseworkers 10.04.2019). Working with target group 6.3, who are not considered ready to take a job on ordinary conditions, this metric was not suitable for the few individuals that could perhaps improve their situation and transition to target group 6.2, who are considered ready to take a job on ordinary conditions. Caseworkers suggested the alternative metric of the number 13-week internships (6h per week) an individual completed, which they argued demonstrates the motivation of individuals to acquire and hold down a job. Our interviews suggested that an algorithm for prediction of risk, such as the one provided by the national Agency for Labour, Market, and Recruitment, did not fit the practices of caseworkers working with individuals who have been unemployed for longer periods of time. Not surprisingly, when the data scientists looked for risks that an algorithm could identify in the initial workshops, the caseworkers instead wanted to discuss the legitimacy of this kind of prediction in their work, as illustrated in this example:

Data scientist: If this information [pointing] is available in the first meeting with the individual.. a red, yellow or green indication of how likely it is that a particular individual will be here [in job placement] for a longer time.. would that be considered valuable information for you?

Caseworker1: Our work is to have “faith” in the individual.. as a strategy for how we can change the outcome for the better.. and maybe I have misunderstood you.. when I am assigned to an individual.. my job is slightly different from what you have described here.. because my “unemployed” have been here for a long time already to be candidates for resource processes, which is the area that I work with.. If I had this “red” indication that most likely this individual will be in job placement for a long time.. then this [information] would do something to me..

Data scientist: Yes

Caseworker1: I consider all my individuals as “blank”.. and I value the history of the individual overall.. This kind of information would be annoying from my perspective.. because there is a reason why they have been assigned to me..

(Workshop 27.05.2019)

This example here illustrates an important point: The caseworker in this situation immediately understood the algorithmic component in terms of how its insights are derived from historical data. Her understanding of the role of a caseworker starts from the individual’s situation at hand, and not the history of other individuals. Algorithms train on historical data, and even with little understanding of the kind of predictions algorithms can make, the caseworker arrives at a central point– that the description of what her job

entails in the data scientist’s view and what she thinks it actually constitutes diverge in alarming ways. The conversation continued:

Data scientist: I fully understand this.. this information should only be presented to you if it makes sense for you [as a caseworker]. That is the entire premise.. that it makes sense. You and the algorithm agree in this case that it is the history of the individual that is the most important thing here..

Caseworker1: Not necessarily.. In job placement caseworkers are not necessarily agreeing.. and a red flag could also be interpreted into.. something that will not necessarily help this particular individual.. but just be interpreted as.. let me get this particular individual out the door.. [...] Sometimes other caseworkers were simply too busy.. to take the required action.

Data scientist: No, but could you imagine.. [...] if it’s a person that needs extra care.. and let’s say that you had known in advance.. the individual has this case history.. and you could have known in advance and taken another direction.. it does not have to be a red flag..

Caseworker: Well, that combination.. and system is of course interesting..

(Workshop 27.05.2019)

The caseworker focused on the role of caseworkers as a measure for the usefulness and legitimacy of a potential algorithmic component. She also raised caution about the implications of a “red flag,” which could prompt a caseworker to avoid a concrete individual due to the perceived extra work involved. A red flag combined with a metric that mainly targets the individual’s improvement, the caseworkers seems to suggest, is not aligned with the role of a caseworker. She needs to make her own judgment about each individual she encounters precisely because she understands that her colleagues are also making subjective, potentially implicitly biased judgments that should be approached with caution. At the same time, she also recognized the potential for this kind of information to perhaps help the caseworker make an argument for additional resources. The caseworker seemed open to the potential usefulness of the algorithmic output that the data scientist advocated for, but not without a degree of caution and concern.

5.2 Shifting the Concept of Value to the Organizational Level

At this point in the workshop the discussion on value metrics shifted to the organizational level. In the design team, we had already discussed the option of focusing on “lay time” – the period of time an unemployed individual is required to wait while various system processes run their course. Changes to lay time implied enhancing organizational efficiencies.

In the interviews prior to the participatory workshops, caseworkers were interested in the overall processes of the organization that can go on for years. For example, one caseworker wanted a better overview of the most recent medical documentation available in the individual case (Interview w/ caseworkers 10.04.2019).

Until this point, value metrics were mostly discussed within the design team in terms of traditional performance measurements, for example, risk prediction as one key interest expressed by the municipality. The push by caseworkers for an algorithmic decision-support system that could help mitigate organizational contradictions and

obscure processes, such as those caused by lay time, shifted the point of discussion. The caseworker reflected:

Caseworker2: Precisely the case of lay time can actually be caused by us. I have a person right now where there is a significant wait time before a decision can be made by the municipal rehabilitation team.

IT-specialist: But the wait time cannot exceed 4 weeks by law

Caseworker2: I have never experienced.. that the wait time was only 4 weeks..

Data scientist: Or if you could predict [lay time] for particular types of individuals with more absence?

Caseworker2: I think it makes more sense if this information on lay time could be used for other things as well..

Caseworker3: When we request the relevant case documentation we are also notified about absence

Data scientist: It is actually possible to predict if an individual will be more or less absent than the average unemployed individual

Caseworker4: But actually information on lay time could be more useful in relation to statistics... on how long time I should be expecting it to take. . . when I request [medical] documentation from psychiatry... That could be really interesting. . . Or if I could know that it will take 3-4 months for the individual's doctor to send the requested medical documentation... Then, I could continue the "offer" [e.g. therapy or job training] that this individual have already started. . .

Caseworker3: I think this could also prevent that a person lose faith that they can be part of... This is starting to be interesting, I think... because when we started out today I just couldn't tell how this [algorithms] could be useful.. but this is actually useful, this bit. As caseworkers we also like to be more efficient... because lay time is exhausting. . . and it's not always that you are able to take action in the time between. . .

(Workshop 27.05.2019)

What happened at this point in the dialogue with the caseworkers was important. As the focus shifts to lay time (the waiting time before a process can move forward, for example, as required documentation is being gathered or other external processes take place) the concept of value away shifts away from algorithmic profiling of individuals. Instead, the caseworkers become interested in the possibilities of improving the organization and optimizing processes so that individuals can coherently experience them. Knowing how long it would take to receive a psychiatrist report would allow an individual to continue therapy until documents are available from the relevant parties. This spurred new discussion within the design team as well. Caseworkers wanted the algorithmic gaze to turn towards the contradictions and inefficiencies of their own organization, assessing lay time and dryly commenting on the vast differences between legal expectations and the practical realities of public service functions.

5.3 Devaluating Concepts of Value at the Individual Level

In the backdrop of the workshop with the caseworkers, an alternative concept of an algorithmic component for case management systems started to take shape, shifting the potential goal of the system towards improvement to the organization. The design team

was encouraged by this unexpected shift in the dialogue that paved the way for a research agenda with more emphasis on organizational accountability and agency for the unemployed individual. These aspects are particularly important for re-balancing the power and value of algorithms with respect to human decisions. As the workshop ended the design team tried to reason through the experience:

Data scientist: I have not yet encountered anyone [caseworkers] that sees the classic information on long-term unemployment as relevant information

Lead system developer: What we can do is to continue working with the [qualitative] descriptions of workflows as well as lay time for different situations.. and also continue our work with the database. The interesting thing about lay time is that it shifts the focus away from the unemployed individual..

Data scientist: We can say if the wait time [for an unemployed individual] is longer than the normal lay time.. but I'm not sure that we can necessarily act on it..

Lead system developer: This could be the "mirror" that we offer to caseworkers.. We could use lay time to make predictions on the expected total wait time for a particular individual..

Data scientist: The algorithm will not be able to make predictions at this level.. It can profile an unemployed individual and tell us about the kind of pathway that an individual seems to be following..

Lead system developer: Maybe it is not necessary that we say to the unemployed "you look like.." but we could simply say "this case looks like.." My point is to keep focus on the pathways..

Data scientist: The essential thing here is if we want to do profiling or not.. I don't believe that we can make improvements on the level of pathways..

Lead system developer: But maybe if we focus more on the context.. if the hypothesis is that we are about to reduce lay time, then it may be possible for us to minimize time that is otherwise wasted..

(Design Team meeting 27.05.2019)

The example illustrates how the shift in focus away from the unemployed individual is not simple but requires the design team to be ready to restrain initial design expectations. In particular, the data scientist responsible for developing algorithms questioned this shift. In the situation, it required not only a broader concept of value metrics, but also a reorientation of what kind of information could be useful, to whom, and for what purposes. The municipality that we collaborated with states the formal performance metrics such as reduction in the percentage of individuals that are unemployed long-term ([21] p.8). These measurements are qualified by the statement that individuals are experts in their own lives and that their responsibility for themselves should be central for their own case.

The measure for when this goal is achieved is described by the municipality as the individual experience of being acknowledged and having agency through the caseworker, including the individual's cooperation around setting targets and sub-targets [21]. Algorithmic components can be designed to support performance measurement, where individual profiling became a logical move given the focus on individual responsibilities of people to ensure their own employment. The involvement of caseworkers resulted in a broader concept of value metrics, targets for predicting lay time in organizational processes, and a call for new forms of organizational

accountability. The caseworkers could use this information to align processes to support the individual and to guard them from, in the words of one caseworker, “losing faith” in the possibility of exercising their agency.

Rather than placing all responsibility onto the individual who has trouble staying employed or seeking employment, the caseworkers recognized the limitations of the system within which they operated and acknowledged the structural conditions that could exacerbate the challenges unemployed individuals already faced. They then looked to the algorithmic system to address the organizational challenges rather than shifting ever more pressure onto the already struggling individual.

6 DISCUSSION: SHIFTING CONCEPTS OF VALUE AS DESIGN

Decision-support in public service is a complex area because of the often overlapping and shifting legal and regulatory context and the divergent goals of governance mechanisms focused on minimizing costs and caseworkers engaging with people who face barriers to fitting into the orderly systems of classification designed to organize them. Our study demonstrates that the processual and long-term nature of much public services casework makes algorithmic decision-support tools challenging for this domain, because what constitutes improvement and successful intervention and other such value metrics can be difficult to establish. Determining value metrics can be a difficult and often counter-intuitive process. In our research caseworkers defined their position and professionalism in discussions of potential algorithmic decision-support even as they contested design decisions in relation to what constitutes value in job placement operations. The caseworkers negotiated the conflict between the ethos of responsibility and care for those needing support, and the municipality’s request for efficiency in its operations. We discuss these issues in turn below.

6.1 Caseworker as a professional

Early participatory design research has taught us that how we set up the design process affects how or whether we arrive at technological alternatives [25, 48]. Participatory design concepts have an element of politics and democratic interest. The intent is to equalize power relations; yet “having a say” requires that workers are educated, engaged, and act both on their own interests as well as in the interest of public good [48]. As we see a push towards public service design with large dataset and algorithms, new questions arise not only about how we can enable employees to participate, but also how to respect and involve them as professionals. Prior research has proposed the concept of data vision [40] to describe the process of learning to see through algorithmic abstraction with data science. In other areas, non-data scientists develop similar skills [26]. However, because they do so as an integrated part of the digitization of their practice, these skills are rarely recognized as part of what constitutes professional work.

Our case shows how data vision can be achieved by setting up a design practice around mutual learning, where caseworker expertise and professional positioning are respected and allowed to influence the fundamental design of the resulting algorithmic decision-support system. Caseworkers must be able to work through concepts of key value metrics that form the basic requirements for algorithmic system development. At stake is not simply public service design with large dataset, but rather a fundamental change in how society makes decisions. This requires public managers and designers to think through deeper motivations for how we set up a design practice if we intend for core community care institutional platforms, such as the municipality [28]. The participatory process that we have described provides opportunities to recognize the importance of building opportunities for flexibility and discretion into algorithmic decision-support systems. Rather than presuming to de-professionalize the jobs that, from a public management perspective might seem ripe for automation, our approach offers ways to treat caseworkers with the respect due to them as professionals and to design systems that can support treating those struggling and most vulnerable in society with dignity.

6.2 Contesting notions of value

Setting up the collaboration with caseworkers was challenging on several levels given the shift towards design with large dataset, both in terms of enabling data vision [40] and placing human and democratic values at the center of our research [29]. The process of arriving at a “value metric” that was considered useful across the caseworkers and design team turned into a thorny and complex process. Prior research on values in design [17] demonstrates that “thorny tensions” can manifest in the design process if particular values are prescribed in advance. Similar to Le Dantec and others [e.g. 15, 17], we set up the design process through workshops with caseworkers as an open-ended process, exploring the possible value metrics of the project. We found that it was not merely a driver for participatory “grounding” of design in practice; rather, the participatory workshops also shifted our object of inquiry as ideas took shape about how we could “profile” lay time in the “organization” instead of profiling the “individual,” as originally expected by the municipality and data scientists.

This insight cannot be ascribed to either the caseworkers or the design team alone. Instead, this shift emerged through the collective process of thinking through notions such as value metrics and lay time. In our case, caseworkers contested the idea that risk prediction could be useful in their work (Interview w/ caseworkers 10 04 2019). This finding aligns with prior research on caseworker and community concerns around risk prediction as fundamentally deficit-based and liable to divert attention away from positive outcomes [10]. The data scientists in our study actively sought a consensus, and there was some tentativeness in the caseworker’s response to the suggestions for how algorithmic decision-support could be useful in job placement. The diverse understanding within the design team was also informed by critical discussions revolving around the national Agency for Labour Market and Recruitment initiative, and by pushback from caseworkers that formed the ground for an open exchange in the participatory workshops. HCI research that

considers public service design often engages concepts of citizenship [5, 13, 17, 19], yet the question that stands out in our case is what happens when data and algorithms come to mediate the relationship between “street level bureaucrats” and the individuals they encounter.

Participatory workshops can be a particularly useful strategy for designers and users to find a common ground and to arrive at a sustainable design based on the design team’s fundamental knowledge of existing practices [13, 14]. As we shift our object of inquiry towards design with large dataset, however, we must be conscious of the dangers of false consensus, in which stakeholder or domain expert involvement is taken to signal legitimacy of design decisions regardless of actual contributions. We worked to understand the consequences in relation to the marginal cases of job placement by staying focused on the case. The caseworkers we engaged with were mostly assigned to individuals who are not new to job placement (e.g. target group 6.3 who are not deemed capable of working or have not been able to keep a job for a variety of reasons). This focus shifted the level of complexity and challenged the “common sense” idea of job placement as a practice that always results in a job. In the database, the contradicting values and logics [53] become clear in the time stamps of when a target group was assigned to an unemployed individual – as it is not unusual to see earlier target group registrations that differ from the final registration.

One important limitation of this design practice is that many individuals were only given voice in the project through prior observational studies of interaction of caseworkers and unemployed individuals in meetings in the job center [34]. The questions that remain are: Who are the individuals that we will never encounter in this kind of study due to anxiety and other factors preventing them from showing up in the job placement the first place - or in a participatory workshop? What are the practices by which we identify individual cases that come to “prototype” the cases that algorithms can train on, and will we then necessarily overlook the truly marginal ones?

6.3 Whose responsibility is it anyway?

Recent HCI-research questions whether it is possible for municipal public services to serve as the care platforms of society today [28]. A similar argument could be made for job placement, which is a key municipal service in Denmark in the move towards an increased use of algorithmic systems. Municipality job placements are set up as a care platform but at the same time they are measures for control. As our case illustrates, caseworkers often work to balance legal demands with a “human” or “caring” approach [24, 28]. The notion of human and democratic values is an interesting starting point, as our case demonstrates: Caseworkers do not explicitly refer to or mention care or human values. Rather, they emphasize having faith in an unemployed person as the starting point for their work. In our discussions, the design team brought different perspectives to bear. The design team also experienced challenges in working across the different ontologies of data science and critical data studies [38]. Participation of caseworkers [25] critically shaped our design practice, broadening the concept of value by shifting the focus of the proposed algorithmic decision-support towards negotiating

the need for additional resources with management in a particular individual’s case.

There is a promise that algorithmic decision-support could help caseworkers identify useful patterns across cases [16]. Thinking through concepts of “human values” as a basis for a caring perspective, we sought to leverage the design process for data and algorithms through a more responsible approach to design. As caseworkers pushed against profiling individuals in their care, they reacted to the idea of ‘flagging’ those that already struggle. This in part had to do with a recognition that not all struggles come from personal failings and that the structures within which we operate are often just as implicated. Caseworkers clearly recognized their own limitations and that they sometimes might act from a place of bias or carelessness in their work, but there was no clear route for an algorithmic system to mitigate these issues. Instead, caseworkers pointed to the unnecessary problems that the institution of job placement itself created for their charges, and here they were happy to look to the algorithmic system for support.

7 CONCLUSION

How can our design practice allow for the development of algorithmic decision-support systems but not lose sight of democratic and human values? Through an empirical study of designing responsible algorithmic systems together with caseworkers in the context of job placement, we consider caseworkers’ discretion as the main access point for human values to enter society’s decision-making practices. Job placement provides a strong case for learning about public service design with large datasets and how we can take a participatory and responsible approach as we develop algorithmic components or “agents.” However, we lack a full understanding of how data can be leveraged to support caseworkers when they apply legal regulations. Working with caseworkers, our open-ended discussion of value metrics in the project shifted the focus away from profiling the unemployed “individual” to profiling lay time in the “organization.” This enabled a more nuanced value metric essential for the design of algorithmically based systems. Our main contributions are insights into how value metrics are negotiated in a participatory design set-up. Whereas the data scientists in our case and municipality initially expected to focus on profiling individuals, the caseworkers instead pushed for systems that could help mitigate organizational contradictions and obscured processes in casework. As legal rules and requirements and algorithmic decision-support systems converge through public service design, questions arise for future research about shifting concepts of value in this context, redefining existing relationships and identifying the unique qualitative aspects of discretion as valuable for casework.

ACKNOWLEDGMENTS

We thank the caseworkers that cared to discuss with us their practices, the municipal job placement management, the union representing a significant part of the Danish caseworkers (Socialrådgiverforeningen), and our industry partners for their openness and support of our work. Special thanks to Dr. Maria Ménendez-Blanco and colleagues in the EcoKnow project for engaged discussions. This work has been supported by the Innovation Fund Denmark (award number 7050-00034A).

REFERENCES

- [1] Ali Alkhatib and Michael Bernstein. 2019. Street-level algorithms: A theory at the gaps between policy and decisions. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, Article 530.
- [2] Karla Bardillo-Urquiola, Xinru Page and Pamela J. Wisniewski. 2019. Risk vs. restriction: The tension between providing a sense of normalcy and keeping foster teens safe online. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, Article 267.
- [3] Reuben Binns, Max Van Kleek, Michael Veale, Ulrik Lyngs, Jun Zhao, and Nigel Shadbolt. 2018. 'It's Reducing a Human Being to a Percentage': Perceptions of Justice in Algorithmic Decisions. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, Article 377.
- [4] Pernille Bjørn & Ellen Balka. 2007. Healthcare Categories Have Politics Too: Unpacking the Managerial Agendas of Electronic Triage Systems. In *Proceedings of The European Conference on Computer-Supported Cooperative Work (ECSCW)*, 371-390.
- [5] Nikolaj Gandrup Borchorst & Susanne Bødker. 2011. "You Probably Shouldn't Give Them Too Much Information". Supporting Citizen-Government Collaboration. In *Proceedings of The European Conference on Computer-Supported Cooperative Work (ECSCW)*, 173-192.
- [6] Claus Bossen, Yunan Chen & Kathleen H. Pine. 2019. The Emergence of New Data Work Occupations in Healthcare: The Case of Medical Scribes. *International Journal of Medical Informatics* 123, March 2019, 76-83.
- [7] Nina Boulus-Rødje. 2018. In Search for the Perfect Pathway: Supporting Knowledge Work of Welfare Workers. *Computer Supported Cooperative Work: An International Journal (JCSCW)*, 27(3-6), 841-874.
- [8] Geoffrey C. Bowker & Susan Leigh-Star. 1999. *Sorting Things Out. Classification and Its Consequences*, MIT Press.
- [9] Tone Bratteteig and Ina Wagner. 2014. *Disentangling Participation: Power and Decision-Making in Participatory Design*. Springer.
- [10] Anna Brown, Alexandra Chouldechova, Emily Putnam-Hornstein, Andrew Tobin, and Rhema Vaithianathan. 2019. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, Article 41.
- [11] Jenna Burrell. 2016. How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society*, January-June 2016, 1-12.
- [12] Peter André Busch and Helle Zinner Henriksen. 2018. Digital discretion: A systematic literature review of ICT and street-level discretion. *Information Polity*, 23, 3-28.
- [13] Susanne Bødker, Christian Dindler & Ole Sejer Iversen. Tying Knots. *Participatory Infrastructuring at Work. Computer Supported Cooperative Work: An International Journal (JCSCW)*, 26(1-2), 245-273.
- [14] Susanne Bødker, Pelle Ehn, Dan Sjögren & Yngve Sundblad. 2000. *Cooperative Design – Perspectives on 20 Years With 'the Scandinavian IT Design Model'*. Invited paper in *Proceedings of the 11th Nordic Conference on Human-Computer Interaction (NordiCHI)*, 1-11.
- [15] Carrie J. Cai, Samantha Winter, David Steiner, Lauren Wilcox, and Michael Terry. 2019. "Hello AI": Uncovering the Onboarding Needs of Medical Practitioners for Human-AI Collaborative Decision-Making. *Proceedings of ACM Human Computer Interaction* 3, CSCW, Article 104.
- [16] Alexander Campolo and Kate Crawford. 2020. Enchanted determinism. *Power without responsibility in Artificial Intelligence. Engaging Science, Technology & Society*, 6(1), 1-19.
- [17] Christopher A. Le Dantec. 2016. *Designing Publics*. The MIT Press.
- [18] Paul Dourish. 2019. User Experience as Legitimacy Trap. *ACM Interactions*, Nov-Dec. 2019, 46-49.
- [19] Andy Dow, Rob Comber, and John Vines. 2018. Between grassroots and the hierarchy: Lessons learned from the design of a public services directory. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, Article 442.
- [20] Virginia Eubanks. 2018. *Automating Inequality: How High-Tech Tools Profile, Police and Punish the Poor*. St. Martins Press.
- [21] Gladsaxe Kommune. 2020. *Beskæftigelsesplan 2020 – Vækst, Velfærd og Unge i Uddannelse og Job*, 1-10 <https://rar-bm.dk/media/9719/beskaeftigelsesplan-gladsaxe.pdf> (accessed May 2020).
- [22] Gladsaxe Kommune. 2020. *Gladsaxestrategien 2018-2022. Bæredygtig Vækst og Velfærd*, 1-24 <https://gladsaxe.dk/Files/Files/Faelles-dokumenter/Planer-politikker-visioner/glx/Gladsaxestrategien-2018-2022-web-5.pdf> (accessed May 2020).
- [23] Lise Justesen, & Ursula Plesner. 2018. *Fra Skøn til Algoritme: Digitaliseringsklar Lovgivning og Automatisering af Administrativ Sagsbehandling*. *Tidsskrift for Arbejdsliv*, 20(3), 9–23.
- [24] Elizabeth Kazunas, Mark S. Ackerman, Silvia Lindtner & Joyce M. Lee. 2017. Car-ing through Data: Attending to the Social and Emotional Experiences of Health Datafication. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*, 2260-2272.
- [25] Finn Kensing & Joan Greenbaum. 2013. *Heritage. Having a Say*. Simonsen, J. & Robertson (Eds), T. Routledge *International Handbook of Participatory Design*, Routledge, 21-36.
- [26] Kristian Helbo Kristiansen, Mathias A. Valeur-Møller, Lynn Dombrowski & Naja Holten Møller. 2018. Accountability in the blue-collar data-driven workplace. In *proceedings of the ACM Conference on Human Factors in Computing (CHI)*, Paper 332.
- [27] David Lazer, Ryan Kennedy, Gary King & Alessandro Vespignani. 2014. The para-ble of Google flu: Traps in Big Data analysis. *Science* 343: 1203– 1205. Available at: www.sciencemag.org/content/343/6176/1203.full/ (accessed May 2020).
- [28] Ann Light & Anna Seravalli. 2019. The breakdown of the municipality as caring platform: Lessons for co-design and co-learning in the age of platform capitalism. In *CoDesign. International Journal of Co-creation in Design and the Arts*, 192-211.
- [29] Ann Light, Alison Powell & Irina Shklovski. 2017. Design for existential crisis in the anthropocene age. In *proceedings of the ACM Conference on Human Factors in Computing (CHI)*, 722-734.
- [30] Michael Lipsky. 1980. *Toward a Theory of Street-Level Bureaucracy*. Russell Sage.
- [31] Thomas Lodato and Carl DiSalvo. 2018. Institutional constraints: The forms and limits of Participatory Design in the public realm. In *Proceedings of the ACM Participatory Design Conference (PDC)*, Article 5.
- [32] Caitlin Lustig, Katie Pine, Bonnie Nardi, Lilly Irani, Min Kyung Lee, Dawn Nafus, and Christian Sandvig. 2016. Algorithmic authority: The ethics, politics, and economics of algorithms that interpret, decide, and manage. In *Proceedings of the ACM Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA)*, 1057–1062.
- [33] Hanne Marie Motzfeldt. 2019. Socialrådgivere og Fagforeninger: Grib dog redningskransen! *Uden for Nummer*, 39(19), 32-41.
- [34] Naja Holten Møller; Geraldine Fitzpatrick & Christopher A. Le Dantec. 2019. Assembling the case: Citizens' strategies for exercising authority and personal autonomy in social welfare. In *Proceedings of the ACM on Human-Computer Interaction* 3, GROUP, Article 244.
- [35] Naja Holten Møller. 2018. The future of clerical work is precarious. *ACM Interactions* 25(4), 75-77.
- [36] Naja Holten Møller & Pernille Bjørn. 2011. Layers in sorting practices. *Sorting out patients with potential cancer. Computer Supported Cooperative Work: An International Journal (JCSCW)*, 20(3), 123-153.
- [37] Klaus Lindgaard Høyer. 2019. Data as promise: Reconfiguring Danish public health through personalized medicine. *Social Studies of Science*, 49(4), 531-555.
- [38] Gina Neff, Anissa Tanweer, Brittany Fiore-Gartland & Laura Osborn. 2017. Critique and contribute. A practice-based framework for improving critical data studies and data science. *Big Data*, 5(2), 85-97.
- [39] Wendy Norris, Amy Volda, Leysia Palen & Stephen Volda. 2019. 'Is the time right now?': Reconciling sociotemporal disorder in distributed team work. In *Proceedings of the ACM Conference on Human Factors in Computing (CHI)*, Article 98.
- [40] Samir Passi & Steven J. Jackson. 2017. Data vision. Learning to see through algorithmic abstraction. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*, 2436-2447.
- [41] Anette C. M. Petersen, Lars Rune Christensen & Thomas T. Hildebrandt. 2020. The role of discretion in the age of automation. *Computer Supported Cooperative Work: An International Journal (JCSCW)*.
- [42] Kathleen H. Pine & Max Liboiron. 2015. The politics of measurement and action. In *proceedings of the ACM Conference on Human Factors in Computing (CHI)*, 3147-3156.
- [43] David Pinelle & Carl Gutwin. 2006. Loose coupling and healthcare organizations: Deployment strategies for groupware. *Computer Supported Cooperative Work (CSCW): An International Journal (JCSCW)*, 15(6), 537–572.
- [44] Anthony T. Pinter, Pamela Wisniewski, Heng Xu, Mary Beth Rosson, and John M. Carroll. 2017. Adolescent online safety: Moving beyond formative evaluations to designing solutions to the future. Note in *Proceedings of the ACM Interaction Design and Children Conference (IDC)*, 352-357.
- [45] Ursula Plesner, Lise Justesen & Cecilie Glørup. 2018. The transformation of work in digitized public sector organizations. *Journal of Organizational Change Management*, 31(5), 1176–1190.
- [46] Juho Paäkköinen, Matti Nelimarkka, Jesse Haapoja & Airi Lampinen. 2020. Bureaucracy as a lens for analyzing and designing algorithmic systems. In *Proceedings of the ACM Conference on Human Factors in Computing (CHI)*, Article 651.
- [47] Dave Randall, Richard Harper & Mark Rouncefield. 2007. *Fieldwork for Design. Theory and Practice*. Springer.
- [48] Toni Robertson & Ina Wagner. 2013. Simonsen, J. & Robertson (Eds), T. Routledge *International Handbook of Participatory Design*, Routledge, 64 85.
- [49] Devansh Saxena, Karla Badillo-Urquiola, Pamela J Wisniewski, and Shion Guha. 2020b. A human-centered review of the algorithms used within the U.S. child welfare system." In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, Article 102.
- [50] Irina Shklovski, Emily Troshynski & Paul Dourish. 2015. Mobile technologies and spatiotemporal configurations of institutional practice. *Journal of the Association for Information Science and Technology*, 66(10), 2098–2115.
- [51] Lucy Suchman. 1993. Do categories have politics? The language/ action perspective reconsidered. *The European Conference on Computer Supported Cooperative*

- Work (ECSCW), 1-14.
- [52] Michael Veale, Max Van Kleek, and Reuben Binns. 2018. Fairness and accountability design needs for algorithmic support in high-stakes public sector decision-making. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, Article 440.
- [53] Amy Volda, Lynn Dombrowski, Gillian R Hayes & Melissa Mazmanian. 2014. Shared values / conflicting logics: Working around e-government systems. In *Proceedings of the ACM Conference on Human Factors in Computing (CHI)*, 3583-3592.