
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Pääkkönen, Juho; Nelimarkka, Matti; Haapoja, Jesse; Lampinen, Airi
Bureaucracy as a Lens for Analyzing and Designing Algorithmic Systems

Published in:
CHI 2020 - Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems

DOI:
[10.1145/3313831.3376780](https://doi.org/10.1145/3313831.3376780)

Published: 21/04/2020

Document Version
Publisher's PDF, also known as Version of record

Please cite the original version:
Pääkkönen, J., Nelimarkka, M., Haapoja, J., & Lampinen, A. (2020). Bureaucracy as a Lens for Analyzing and Designing Algorithmic Systems. In *CHI 2020 - Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* Article 3376780 ACM SIGCOMM. <https://doi.org/10.1145/3313831.3376780>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Bureaucracy as a Lens for Analyzing and Designing Algorithmic Systems

Juho Pääkkönen^{1,2} Matti Nelimarkka^{1,2} Jesse Haapoja^{1,2} Airi Lampinen³

¹ University of Helsinki, Helsinki, Finland ² Aalto University, Espoo, Finland

³ Stockholm University, Stockholm, Sweden

¹ first.last@helsinki.fi, ² first.last@aalto.fi, ³ airi@dsv.su.se

ABSTRACT

Scholarship on algorithms has drawn on the analogy between *algorithmic systems* and *bureaucracies* to diagnose shortcomings in algorithmic decision-making. We extend the analogy further by drawing on Michel Crozier’s theory of bureaucratic organizations to analyze the relationship between algorithmic and human decision-making power. We present algorithms as analogous to impartial bureaucratic rules for controlling action, and argue that discretionary decision-making power in algorithmic systems accumulates at locations where *uncertainty* about the operation of algorithms persists. This key point of our essay connects with Alkhatib and Bernstein’s theory of ‘street-level algorithms’, and highlights that the role of human discretion in algorithmic systems is to accommodate uncertain situations which inflexible algorithms cannot handle. We conclude by discussing how the analysis and design of algorithmic systems could seek to *identify and cultivate important sources of uncertainty*, to enable the human discretionary work that enhances systemic resilience in the face of algorithmic errors.

Author Keywords

Algorithmic systems; Algorithmic power; Uncertainty; Bureaucracy; Street-level algorithms; Street-level bureaucracies; Automated decision-making.

CCS Concepts

•Human-centered computing → Collaborative and social computing theory, concepts and paradigms; •Applied computing → Law, social and behavioral sciences;

INTRODUCTION

Research on the social implications of *algorithmic systems* (e.g., [1, 14]) invites us to consider automated information processing as analogous to decision-making in *bureaucracies*. In this essay, we elaborate on this analogy between algorithmic systems and bureaucratic organizations. We contribute an analytical approach that emphasizes *uncertainty*

as a key concept to investigating how human and automated decision-making interact to produce powerful consequences. Understanding this interaction is important, given the increasing societal role of algorithmic processes in supporting and automating decisions that have been traditionally made by humans – across contexts as varied as governance [35], media [24, 3], health [60], and everyday life [72].

We develop the bureaucracy analogy to make it compatible with the well-established view that algorithmic systems are *socio-technical*. That is, we treat algorithmic systems as settings where implementations of algorithms intermesh with human action, practices, institutional conditions, and other technological artefacts. This view of technological systems as socially situated has been central in both HCI literature (e.g., [68, 33, 74]) and critical algorithm studies (e.g., [3, 4, 51, 15]). Our contribution is to build a bridge between this long-standing work and organizational theory on bureaucracy, as it is applied to analyze algorithmic systems. By doing so, we respond to calls for conceptual work in HCI [54], exploring conceptually how humans interact with algorithmic systems.

We draw on Michel Crozier’s [18] *theory of bureaucratic systems of organization* and propose to treat algorithmic systems as analogous to dysfunctional organizations. Both kinds of systems seek to subsume decision-making under rigid rule-governed procedures, but risk becoming too inflexible to adapt to unpredictable errors. Crozier’s theory holds that discretionary decision-making in such inflexible systems happens at places where there is *uncertainty* about rule application. Analogously, when it comes to algorithmic systems that involve inflexible automated processes, we argue that human discretionary decision-making works at *locations of uncertainty*. These are points in the decision-process where formal rules cannot be turned into actionable consequences without human judgment (cf., [18]). This analogy highlights that discretionary decision-making power works even in extensively automated systems, drawing on persisting uncertainties to influence decision-making. While this point in itself might not strike as novel to HCI [74], we deem it worth reiterating in relation to algorithmic systems that often take the appearance of objectivity or eliminating human decision-making [24].

Overall, the algorithms we focus on can be described as recommendation systems — a category commonly explored in the context of algorithmic power both within HCI [1] and critical studies on algorithms [4]. In particular, we examine supervised machine learning and rule-based systems for classi-

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 the author(s) 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.

CHI '20, April 25–30, 2020, Honolulu, HI, USA.

© 2020 Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-6708-0/20/04 ...\$15.00.

<http://dx.doi.org/10.1145/3313831.3376780>

fyng data. These algorithms can be said to be inflexible in the sense that their decisions are based on previously collected training data or predetermined criteria that cannot be reflexively modified at the moment of decision [1]. Importantly, algorithms that are inflexible in this sense can be retrained or modified to account for errors *after decisions*. Further, their decision boundaries can be very flexible in the sense of fitting to complex data. However, they are inflexible in the sense of being unable to reflect on novel features of unseen cases *at the moment of decision-making*.

In prior work, Alkhatib and Bernstein [1] have used the bureaucracy analogy to propose that this kind of inflexibility constitutes an essential difference between human and algorithmic decision-making. They link the dysfunctions of algorithmic systems to human discretion being replaced with inflexible algorithms. However, as the authors themselves note, their theory does not account for situations where algorithms interact with human discretion. Our approach, building on the socio-technical view of algorithmic systems, emphasizes that algorithms rarely work in solitude. As such, the dysfunctions of many algorithmic systems cannot be properly understood by focusing solely on algorithmic inflexibility. Crozier’s theory enables us to develop the bureaucracy analogy to account for this insight, while retaining its main point that algorithmic inflexibility is a source of error. Uncertainty becomes the key concept for analyzing how human discretion manages inflexibility to influence decisions in algorithmic systems.

To unpack our argument, we discuss three illustrative cases from prior literature: the use of risk prediction models in criminal court [15], an automated housing allocation system for the homeless [22], and automated worker-to-task matching on crowdwork platforms [28]. These cases correspond to increasing levels of automation in the system. While in the first case humans have leeway to exercise discretion, in the latter two decision-making is rigidly automated. However, in each case human action finds ways to influence the consequences of algorithmic processes. We demonstrate our analytical approach by identifying in each case how humans use discretion to deal with uncertainties within the system. Finally, on the basis of the cases, we discuss conceptual implications and explore how our perspective relates to design directions that seek to account for uncertainty in human-algorithm collaboration.

In the following, we first present the relevant theoretical background, considering how bureaucracies have been defined in the past and what is already known about decision-making in algorithmic systems. After presenting our analytical approach, we turn to the three cases to illustrate its usefulness. We contribute a theoretically-grounded perspective that highlights, first, how humans and algorithms collaborate to form a bureaucracy-like system (this does not happen by algorithms alone, cf. [1]), and second, that algorithmic inflexibility is managed at locations of uncertainty. To conclude, we offer some initial thoughts on how design work could account for and even emphasize locations of uncertainty.

BUREAUCRACY AND ALGORITHMIC SYSTEMS

We now introduce three different accounts of bureaucracy as a foundation for considering the analogy between bureaucracies and algorithmic systems. We start with the classical ideal type of bureaucracy as a machine-like system governed by formal rules. Second, we move to Michael Lipsky’s [44] theory of street-level bureaucracy, which emphasizes the role of human discretion in rule application. Finally, we introduce Crozier’s [18] theory of bureaucracies as dysfunctional systems of organization where human discretion and formal decision-making interact. Each of these accounts belongs to a long theoretical tradition with extensive literature. We present the theories in a schematic fashion in order to highlight their main points and contrasts. In particular, our intention is to focus on how uncertainty, emphasized by Crozier, enables analyzing the role of discretion in bureaucratic systems. Table 1 summarizes these three accounts of bureaucracy and highlights how they differ in their analyses of decision-making power.

Classical views of bureaucracy

In management and organization sciences, political science and public administration research, much emphasis has been put on decision making and bureaucracy – not only in public organizations but in all kinds of organizations. According to Weber [71] and Jaques [38], organizational decision making should be structured in specialized and hierarchical functions, with clear and centralized management, and personnel development should be focused on professional qualifications. Here, the underlying idea is that of organization as a machine, which takes commands from management and mechanically executes them [49]. From this classical perspective, the decision-making criteria of bureaucratic organizations are seen to be rational and the goal of the organization is to maintain the overall rationality and handle all cases in a consistent fashion. The aims of rationality and robustness make the connection between bureaucracy and algorithmic systems clear: bureaucracies, too, are meant to function like machines, taking inputs and systematically delivering outputs based on them.

This classical view of bureaucratic rationality was put forward by Weber as an ideal type that itself is never fully realized but that provides a point of contrast for studying real organizations: for example, bounded models of rationality [62] and the different functions organizations have [49] constrain their operation in the idealized manner. Due to these concerns, later social scientific accounts have sought to characterize bureaucratic decision-making by focusing on the *practices* through which formal rules are applied in concrete situations. This is the starting point for the theory of *street-level bureaucracy* [44] that Alkhatib and Bernstein [1] resurfaced as a resource for analyzing algorithmic decision-making power.

Street-level bureaucracy and street-level algorithms

Michael Lipsky’s [44] theory of street-level bureaucracy proposes that decision-making in public policy does not mechanically follow abstract policy rules. Rather, effective policy application necessarily involves the reflexive discretion of

	Classical bureaucracy	Street-level bureaucracy	Crozierian bureaucracy
<i>Definition</i>	Machine-like system governed by hierarchically organized rules.	Rule-governed system which depends on the discretion of individual officials.	Rule-governed system which is unable to flexibly correct its errors.
<i>Location of power</i>	Supervisors in command of formal decision-making power.	Individual officials with discretionary power.	Interplay of formal and discretionary power, both seeking to control uncertainty.
<i>Level of automation</i>	High.	Low, case-by-case.	Aspires to high automation but fails to control discretion.

Table 1. Schematic conceptions of bureaucracy and their dimensions.

street-level bureaucrats (on-the-ground human officials, such as police officers, judges, social workers, and teachers) who are responsible for adapting abstract rules to concrete situations. Although the decisions of street-level bureaucrats are directed and constrained by policy rules, in Lipsky’s analysis it is the discretionary application of rules to particular cases that really comes to constitute organizational policy. Thus, in contrast to the ideal of classical bureaucracy, described schematically above, street-level bureaucracy does not locate decision-making power at central management positions in formal hierarchy. In this account, the focus is rather on the *discretionary power* of individual officials, emphasizing their practices and routines of rule application (see Table 1).

Recent discussion in HCI has connected this perspective to algorithmic decision-making. Building on Lipsky’s theory, Alkhatib and Bernstein [1] propose a theory of *street-level algorithms*, or algorithms that “directly interact with and make decisions about people in a sociotechnical system” [44, p. 2]. Examples they draw on include YouTube content moderation algorithms (e.g., [26, 40]), quality assessment algorithms on crowdwork platforms [2], and recidivism risk prediction models in criminal justice [15]. Analogously to human officials, responsible for bridging the gap between high-level policy rules and concrete situations, street-level algorithms are the layer of information processing in algorithmic systems that is “specifically responsible for making decisions that affect the lives of the users and stakeholders” [1, p. 3].

However, Alkhatib and Bernstein argue that, unlike in street-level bureaucracies, in algorithmic decision-making reflexive discretion plays no role. Algorithmic decisions must follow predetermined design features which govern the systems’ operation, based on necessarily limited training data. Although algorithms can be retrained or modified after detecting errors in their decisions – leaving room for instance for recourse processes to mitigate biases [1, pp. 8-9] – they cannot flexibly adapt their decision boundaries *at the moment of decision* to treat novel or marginal cases correctly. This leads to the potential of mistreating users or clients. The inflexibility of algorithms and digital systems has been recognized to limit reflexive policy making also in public administration literature [9, 37].

We view this conceptualization as an important contribution to addressing the shortcomings of algorithmic decision-making. The theory of street-level algorithms connects the issue of algorithmic power to longstanding literature on organizational decision-making but it has the limitation of being geared mainly to deal with cases of fully automated decision-

making where human discretion plays little role (cf., [1, p. 10]). By definition, the theory does not cover situations where algorithms are used in roles other than that of replacing human agents. Yet, as a long lineage of studies in both HCI [33, 43, 74, 28] and critical algorithm studies [3, 4, 51, 15, 61] argue, many algorithmic systems are based on humans and algorithms collaborating and interacting. In such socio-technical systems, power and the associated dysfunctions cannot be analyzed merely by focusing on decisions made by automated systems. Rather, the analytic focus must *expand to include wider aspects of the socio-technical setting*. To achieve such theoretical expansion in the bureaucracy analogy, we turn to Crozier’s theory of bureaucratic systems of organization, and in particular its focus on uncertainty.

Crozier’s theory of bureaucratic systems of organization

According to Crozier’s theory [18], any system of organization that cannot flexibly correct its behavior in the face of errors is a bureaucratic system. Exemplary cases discussed by Crozier included a Parisian clerical agency and a state-owned French industrial monopoly, both of which were unable to cope with human relations issues and technical problems, to the detriment of their efficient functioning. Crozier’s theory views bureaucratic organizations not as rational machines, but rather as *complex rule-governed socio-technical setups, where systemic inflexibility is a defining feature*. Similar issues of inflexibility in the face of errors have been documented to occur with algorithmic systems, too. For instance, Google has struggled to react to what appeared like coordinated manipulation of search results because of pressures related to neutrality and objectivity [25].

Crozierian bureaucracies are born through the establishment of rigid sets of rules, which seek to meticulously govern all human action within the organization. The aim in establishing such rules is to centralize decision-making power and to eliminate as many sources of *uncertainty* as possible in the organization’s operations. For achieving this aim, it is essential that bureaucratic rules aim towards *impartiality* – in the sense that no human discretion in applying them is required – and *totality*, so that ideally every aspect of the system’s operation can be predicted and controlled. That is, the ideal model of a Crozierian bureaucratic organization is a procedurally fully automated system, akin to the classical ideal type of bureaucratic rationality.

However, Crozier illustrated with his case studies how this model of organization remains an unattained ideal. For instance, the Parisian clerical agency was a large organization

with poor working conditions and low staff morale. The agency was unable to rectify these problems due to strict national control of its operation. Work in the agency alternated between periods of routine and crisis, caused by occasional spikes in the volume of service calls. During crisis, the organization could only respond by policing its rules ever more strictly. This led to further lowering of work morale – a prime example of what Crozier called the *bureaucratic vicious circle* [18, p. 111]. In the case of the industrial monopoly, factories constrained by hierarchical rules were struggling to make necessary reforms in production technology and staff. This situation hindered productivity and enabled the development of uncontrolled power within the factories.

A central insight that Crozier demonstrated through his case studies was that, instead of eliminating human discretion, the establishment of rigid rules in bureaucracies leads to discretionary power accumulating in the hands of actors *who have the capability or skills to control the remaining sources of uncertainty* [18, Pt. I-II]. A case in point was the industrial monopoly, where occasional malfunctions of production machinery constituted a significant unpredictable factor in the organization's operations. This source of uncertainty could only be controlled by the plants' maintenance personnel, who possessed the skills necessary for repairing the machines – an expert status which gave them crucial strategic leverage in the organization. As long as the maintenance personnel superficially aligned their work with the formal rules of the organization, they were able to retain leeway in deciding how and when repair takes place. Thus, bureaucratic rules – which could not strictly determine how repair work was to be carried out – also effectively worked to make the organization too rigid to establish centralized control over significant failures in functioning.

The implication of Crozier's theory is that the bureaucratic tendency towards centralized control is a counterproductive aim: It sacrifices flexibility and adaptability for systemic predictability and optimization. According to Crozier, controlling for all sources of uncertainty in complex organizations is impossible [19], and thus bureaucratic rules – rather than achieving machine-like rationality – effectively redistribute discretion to the locations of remaining uncertainty within the system. In this sense, Crozier's account aligns with the street-level bureaucracy's idea that discretion is necessary for rule application. However, the import of Crozier's theory is to identify *how* discretionary power manages to operate within the confines of rigid formal rules. The key locations at which discretion finds leeway are the uncertainties that cannot be accommodated in terms of bureaucratic rules. The discretionary work at these locations can easily go unrecognized, because formal rules give the system an appearance of effectiveness and impartiality [56]. A focus on uncertainty allows us to see through these appearances, and analyze how discretion works jointly with formal rules to produce decisions.

In what follows, we treat algorithmic systems as analogous to Crozierian bureaucratic systems. We pose the analogy between rigid bureaucratic rules and inflexible automated infor-

mation processing to highlight uncertainty as key to analyzing how such inflexibility is managed in algorithmic systems.

ANALYTICAL APPROACH

Our analytical approach draws on Crozier's theory of bureaucratic systems of organization to provide a systematic way of studying how human discretion and algorithmic processes relate. In this section, we first describe the key components of our approach, which proceeds by 1) identifying locations of uncertainty in algorithmic systems and 2) analyzing how discretion works at those locations. Second, we discuss how the bureaucracy analogy based on Crozier's theory relates to the well-established view of algorithmic systems as socio-technical. We will proceed to the illustrative cases in the next section to demonstrate our approach.

Analyzing power through locations of uncertainty

Uncertainty in Crozierian bureaucracies is defined in terms of *incapability to predict and control events which influence the organization's operations*. Uncertain future events are unknowable, and as such they are different from risk, understood as probabilistic evaluations of different possible outcomes of a procedure (c.f., [7, 42]). Calculating the probability of future events under uncertainty is impossible, and consequently decisions will have to include an element of discretion.

The import of Crozier's theory for analyzing power in algorithmic systems is twofold. First, adapting this perspective to algorithmic systems suggests, by analogy, that *power in such systems accumulates at locations where uncertainty about the operation of algorithms persists*. Therefore, as the degree of automation in algorithmic systems grows, power in the form of reflexive discretion accumulates more and more tightly around tasks which cannot be automated, that is, to human agents. These are the locations of uncertainty where the operation of algorithms cannot yield actionable outcomes without intervening human judgment. Second, this perspective enables us to distinguish between *different locations at which humans can come to have discretionary power*. As our case discussion will illustrate, uncertainty in algorithmic systems can variously persist in interpreting algorithmic outputs, maintaining algorithms in operation, or coordinating human actions.

Supporting vs replacing human decision-making

To make sense of how humans interact with algorithmic systems, we distinguish between two roles through which algorithms can have power: automating or supporting human decision-making. Algorithms can be used to *replace* humans in certain tasks, through "encoding of human agency" [34] in programmatic and automated decision structures. For instance, the Facebook newsfeed [12, 57, 21] or Google PageRank [25, 75] algorithms are examples where automated technology replaces more traditional editorial work in deciding what information should be displayed to users (c.f. [24, p. 192]). In these cases the discretionary power is located in various stages of developing and maintaining the algorithmic system. By contrast, the use of risk assessment models in court by judges [15], or doctors' use of predictive modeling

of infections [46] are cases where algorithms serve as *support systems* that can extend or steer human action but ultimately leave room for discretion.

We argue that algorithms enter into complex interactions with humans both in their supporting and replacing roles. These insights highlight the need to subsequently examine the practices and institutional conditions which lead to particular power configurations. Crozier’s theory provides us with a systematic approach to achieve this, as it is specifically fleshed out to deal with systems where human discretionary decision-making is at play, despite efforts at automation.

Power in socio-technical systems

As has been long argued in both HCI and the critical algorithm studies literature, automated information processing is always situated in and shaped by features of the socio-technical systems within which it operates [68, 33, 51, 3]. Accordingly, while there is broad consensus that algorithmic systems can have power – in that they make automated decisions [69, 45] or act as information gatekeepers [25, 13, 41] – the consequences of algorithms have been recognized to depend for their meaning and efficacy on a *larger system* of human interpretative work, institutional context, and work practices [51]. Thus, recent literature on algorithmic systems continues well-established discussions on power and values in technical systems, following in the footsteps of, for example, Suchman’s work on human-machine reconfigurations [68], Bowker and Star’s insights into classification [10], and Friedman et al.’s and Nissenbaum’s sensitivity to values and differing stakeholder views [23, 52].

The bureaucracy analogy adds to the socio-technical view of algorithmic systems the idea that algorithms may be viewed as extensions of bureaucratic administration – as *tools for governance and management* – analogous to more traditional mechanisms for administering and classifying people, such as forms and questionnaires [14]. This characterization leads to framing discussions about the shortcomings of algorithms in a fresh way. For instance, discussions about problems with decision-making algorithms have tended to emphasize the discrimination their use may cause (see, e.g. [27, 63, 76]), focusing on technical solutions that aim to increase the *fairness, accountability* and *transparency* (FAT) of algorithms and machine learning systems. From the perspective of the bureaucracy analogy, technical solutions to problems in algorithmic systems constitute new administrative mechanisms, which potentially lead to redistributions of power that merit further analysis. Our analytical approach implies that focusing on persistent uncertainties is key to analyzing the way in which power works in the resulting systems.

ILLUSTRATIVE CASES

We use three cases from prior literature to demonstrate the usefulness of our analytical approach. These cases were selected to highlight that human actors can purposefully leverage uncertainty in even extensively automated algorithmic systems. To illustrate how this happens in different kinds of systems, we first examine a case where algorithms are used to *support* the work done by human decision-makers, namely, in

the use of risk prediction models in criminal justice (Christin, [15]). Second and third, we examine cases where the aim is to automate human-decision making. We examine an automated system for ranking and matching homeless people with available housing (Eubanks, [22]) and worker-to-task matching on crowdwork platforms (Gray and Suri [28]). In both of these, we will see how discretion finds leeway to influence the system’s operation, respectively through uncertainty involved in *repairing or maintaining* the ranking algorithm’s operation, or through *coordinating human work* in relation to inflexible crowdwork assessment algorithms. We conclude each section by summarizing the key observations.

Case 1: Risk prediction support systems in criminal court

Our first case is based on a study by Christin [15] on the use of algorithms for modeling criminal recidivism in U.S. courts. This case is set against the background of an exponential increase in U.S. jail and prison population during the past 50 years, along with a growing awareness of racial bias at “every step of the process” [15, p. 5]. The stated aim behind the recent widespread adoption of risk prediction algorithms in U.S. courts is to reduce incarceration rates and alleviate bias in judicial decision-making, by making judges’ decisions more objective and better grounded in data (see e.g. the 2013 TED talk by Anne Milgram [48]). However, as noted by Christin [15, p. 6], this stated aim has been strongly contested in recent debate (see, e.g. [5]), with the argument that prediction algorithms end up amplifying and reinforcing biases instead of reducing them [53]. Risk prediction in court, thus, is a case where algorithms are used with the explicit intent to support the work of human decision-makers, but where a conflicting relationship persists between human and algorithmic agents.

The bias-reinforcing tendency of predictions has been diagnosed to result from necessarily limited training data [1, pp. 7-8]. Given marginal or novel cases, including for instance intersectional combinations of traits that together have no antecedent in previously collected data, risk prediction algorithms are bound to produce incorrect results. At worst, these end up replicating biased decisions reified in past cases. Consequently, Alkhatib and Bernstein argue that the commonly advocated principles of fairness, accountability, and transparency in algorithmic design are an insufficient goal in judicial decision-making, and that retaining an element of reflexive discretion in the process is necessary for the just treatment of novel and marginal cases. The judges interviewed by Christin [15, p. 9] agree, holding that subjective discretion in court decision-making is constitutive of justice. Indeed, as Christin’s study shows, judges in U.S. courts often oppose algorithms, maintaining their discretionary power through the strategies of ignoring, manipulating, and openly criticizing the information produced by prediction models [15, pp. 9-10].

Given our analytical approach, our concern here is to understand how the judges’ authority to criticize the prediction algorithm is grounded in the wider social and organizational setting within which the algorithms are used. Importantly, Christin [15, pp. 10-11] locates the judges’ authority to criti-

cize algorithms to stem from the *uncertainty associated with interpreting the results of prediction algorithms as evidence in the context of justice*. More particularly, this source of uncertainty has to do with the lack of precedents in the legal tradition of relying on algorithmic information processing as evidence:

In criminal justice, innovation does not come with the glitter and appeal that it has in other sectors: it is often a source of uncertainty, because by definition an innovation arrives without the vetting of precedent. Any incentive to use new tools must be balanced against the clear motivations for relying on legal tradition. [15, p. 10]

Here, we observe how the institutional context influences who has the power to interpret algorithmic outcomes. The uncertainty associated with algorithmic predictions in judicial decision-making thus derives from the traditional practices of the criminal justice institution. Judges – who traditionally represent the legitimate source of decision-making power in courts – are able to retain their discretionary power and reinforce it through openly attacking risk prediction algorithms. As Christin maintains, in the context of criminal justice, algorithmic predictive modeling needs to “become part of the tradition before it can be trusted” [15, p. 11]. Before the source of this interpretive uncertainty can be eliminated, then, the judges’ discretion will likely continue to command decision-making power in criminal court – albeit this power might be masked “*under the patina of objectivity*” [15, p. 10] that surrounds algorithmic technologies [24, 14, 8, 56]. As noted above, formal rules can give bureaucracies the appearance of impartiality. The use of algorithms can work toward the same effect, making decisions seemingly independent from the effect of human judgments.

To sum, this case suggests that when algorithms are used to support (rather than automate) decision-making, the *interpretation of algorithmic outcomes can be a significant source of uncertainty*. Humans can interpret algorithmic outcomes if they have an appropriate institutional position or other kind of institutional credentials for doing so. In other words, the power of algorithmic systems is a *mutual product* of automated and discretionary decision-making. This means that the use of automated procedures for guiding human decisions leads *discretionary decision-making power to accumulate in the hands of people*, who are by training, expertise, or institutional position/credentials legitimated to decide what the outputs of algorithms should be taken to mean, and whether they should be taken seriously and acted upon.

Case 2: Automated housing allocation for the homeless

Our second case – an automated matching algorithm developed for evaluating the housing needs of homeless people in Los Angeles, as documented by Eubanks [22] – diverges from the first one in that here, the purpose of using algorithms was not to support, but rather to *automate* the decision-making process. The express motivations for using algorithms were similar to the previous case, however. First, the system was intended to improve efficiency, in response to a boom in the number of people living on the street in Los Angeles. Second, algorithmic allocation sought to enact impartiality, through

an automated ranking system based on a standardized survey questionnaire administered to the homeless as part of their housing application procedure. On the basis of survey responses, a ranking algorithm first calculated a score for each applicant and ranked them according to their need for housing. Then, another algorithm was run to match available housing opportunities with people who meet their eligibility criteria. The underlying aim of this process was to provide housing as quickly as possible to the people who need it the most, and who are most suitable for the available apartments. Before automated housing allocation, unhoused people in Los Angeles had to go through a complex and time-consuming system of street-level bureaucracy, where waiting times were long due to the constant resource scarcity of service providers [22]. Thus, the stated aim of establishing the automated system was to move decision-making power away from places where it was inefficient, and to reinstall it to work through an efficient and impartial algorithmic process.

However, on the basis of her interviews with homeless applicants, the system’s designers, and data gathering workers, Eubanks argued that what the system really amounted to was a tool for surveying and sanctioning the homeless. The system’s working depended on establishing a survey database of homeless people, which – in addition to their personal information – contained entries about private issues such as drug abuse and sexual behavior. As such, Eubanks argued that the system also represents a shift from community police work, carried out by street-level bureaucrats in interaction with neighborhoods familiar to them, to data-based surveillance classifying people socially on the basis of predetermined criteria.

Most importantly for our argument, *algorithmic allocation redistributed discretionary power* to work at a location of uncertainty where decisions could not be automated: the gathering of surveillance data. Eubanks followed several cases of homeless applicants and showed how different application agencies had varying interpretations about how to record survey responses into the allocation system [22]. While the algorithmic allocation system carried formal power for making decisions, this power only became efficacious through various interpretations of the data gathering instruments. Moreover, Eubanks describes how these interpretations worked on the basis of criteria extraneous to the application process. Data collectors in different agencies judged the same applicants differently, on the basis of the applicants’ attitude and the collectors’ diverging ideas of how the allocation system should work. Thus, uncertainty about how responses should be recorded in the system led to the data collectors gaining power to impose their views of how the homeless should live on the system’s operation.

In accordance with Crozier’s analysis of bureaucratic organizations, in this case we see how controlling for all possible sources of uncertainty is infeasible. Despite extensive automation, the system depended on data collection that could not be standardized extensively enough to eliminate discretion. Thus, this case demonstrates the point previously argued by Dourish [20] that *algorithmic processes cannot be*

properly understood in isolation from their particular implementations as programs and the data that they draw on.

Analogously to Crozier’s analysis of the power of factory maintenance personnel, we see how the dependency of algorithmic decision-making on implementation leads discretionary power to accumulate in the hands of those involved in *maintaining* algorithms. As has been discussed in the literature, the work involved in implementing algorithms may include specifying criteria used in classification [51] and iterative maintenance to keep particular implementations operational. Data collection and curation can be regarded as part of such *repair work* [36] involved in maintaining algorithmic information processing [55]. This is not to say that repair would not play a role when algorithms are used as support systems (see e.g., [59]). However, recognizing this source of uncertainty becomes pivotal in cases of extensive automation, where human actors ostensibly lack decision-making power.

To sum, when algorithms are used to automate decision-making, interpretation of outcomes can no longer be controlled by discretion. Instead, significant uncertainty can persist in relation to *implementing algorithms as programs*. This uncertainty can be controlled by those who have *access* to designing or maintaining algorithmic technologies, such as data collectors or people responsible for correcting program implementations. Thus, especially in the case of extensively automated algorithmic systems, discretionary power can end up in the hands of unexpected groups of people. Automation of decision-making does not eliminate discretion, but rather redistributes it to work at locations where the system’s operation is maintained. This implies that also the dysfunctions of extensively automated systems can be due to unrecognized discretionary work.

Case 3: Automated task matching in crowdwork

The final case we discuss – worker-to-task matching on crowdwork platforms such as Amazon Mechanical Turk – is also discussed by Alkhatib and Bernstein [1], as a form of piecework [2] that is essentially managed by automated quality assessment procedures. According to Alkhatib and Bernstein, the relationship of crowdworkers to quality control algorithms is similar to that of factory workers to their foremen, whose responsibility it was to assign work tasks to the workers who could best carry them out, and to evaluate the workers’ performance in order to direct their skill development. This task essentially depended on discretion in assessing the match between each idiosyncratic task, worker skills, and different styles of working. It is this crucial discretionary element that quality assessment algorithms have come to replace on crowdwork platforms. Consequently, research has found that crowdworkers must deal with a constant fear of unnecessary rejection by inflexible automated processes [47]. Algorithmic evaluation tied by strict definitions of successful task completion can at worst end up rejecting work that diverges from gold-standard solutions only in irrelevant or marginal ways [39].

Alkhatib and Bernstein’s analysis identifies a central problem in automated worker-to-task matching. Rigid automated control displaces decision-making power from human supervi-

sors and consequently places crowdworkers under pervasive uncertainty concerning appropriate task completion. To be sure, the underlying intent behind algorithmic work assessment is to provide efficient procedures for coordinating distributed cooperative work at large scales. However, this does not mean that the discretionary work of quality assessment disappears. This is one of the central messages of Gray and Suri’s [28] recent analysis of *ghost work* – that is, the invisible crowdwork involved in the operation of many technological systems. Gray and Suri argue that, despite the automation of worker-to-task matching, the discretionary efforts involved in these procedures – the “transaction costs of ghost work” – do not disappear, but instead are “shifted to the shoulders of requesters and workers” [28]. Thus, paradoxically, automated management all but eliminates human discretion in work assessment. Instead, what algorithmic matching achieves is a redistribution of the *responsibility* of handling this workload, from the hands of platform employees to on-demand workers and task providers. By fixing rigid criteria for task evaluation, algorithmic matching glosses over the uncertainty involved in coordinating workers’ skills with the requested tasks. Consequently, the discretionary work required to control this uncertainty is redistributed to workers and task providers. In this case we see how algorithms can redistribute the negative discretionary duties involved in maintaining system operation.

The essential difficulty involved in automating worker-to-task coordination is that it is practically impossible to enumerate exhaustively the multiple different ways in which all but very simple tasks can be completed (c.f., [16]). Thus, in this case, uncertainty stems from the inherent difficulty of pre-determining an interpretation for all possible actions that workers can undertake in completing a given task [68]. Of course, the difficulty of evaluation will vary from task to task. However, the problem of automated quality control in the case of crowdwork remains that of classifying a potentially unlimited number of different solutions using criteria which by necessity can capture only a limited number of cases [1].

As Gray and Suri note [28], for all relevant practical purposes, the quality assessment on crowdwork platforms works in an automated fashion. In automated worker-to-task matching in crowdwork, humans do not in practice have the possibility to directly override algorithmic decisions. Also, platform design does not permit requesters or workers to influence the operation of quality control algorithms. Instead, crowdworkers engage in efforts to *coordinate their work* through various informal channels in order to determine what kinds of solutions will be accepted for different tasks. For instance, workers have been found to engage in back-channel communication [29] to discuss different solutions and the reliability of task providers. Such invisible coordination efforts have been discussed in the CSCW literature as *articulation work*, or discretionary work which maintains the efficiency of cooperation in the face of unexpected difficulties [66, 67]. In ghost work, articulation work serves the crucial function of mitigating the risk of unnecessary rejections by improving the fit between automated evaluation and available solutions. Further, research has found that automated management of distributed work can lead to workers developing work-arounds

to circumvent assessment [43]. The literature on algorithmic systems has recently started to discuss these latter kinds of activities under the umbrella notion of “gaming” [6], or acting strategically in order to intervene with established system behavior [30, 17].

As in the case of automated housing allocation, it is important to recognize these roles of discretion. As emphasized in prior literature, articulation work tends to be *invisible* by default [66]. Furthermore, gaming algorithms has been noted to be so ubiquitous that it is difficult to identify algorithmic systems which do not involve gaming in some form [6]. This implies that even in cases of extensive automation, some uncertainty always persists concerning the system’s operation.

To conclude, the potential for users to coordinate their actions is a significant source of uncertainty in extensively automated algorithmic systems. The implications and dysfunctions of automated systems thus depend partly on user activities and their coordination. For instance, biased outcomes might be due to an imbalance in knowledge about algorithms, which benefits only some groups of users, while disadvantaging others.

CONCEPTUAL IMPLICATIONS

The cases we have discussed illustrate the key insight of our analytical approach: when it comes to analyzing algorithmic systems, powerful consequences are a mutual product of automated and human discretionary decision-making. Uncertainty is a key concept for understanding the process that leads to them. We now elaborate on this insight to draw conceptual implications from our analysis.

Algorithmic inflexibility in socio-technical systems

In all the three cases we discussed, discretion worked to control sources of uncertainty in order to accommodate unanticipated circumstances. These cases demonstrate that the shortcomings and consequences of algorithmic systems cannot be understood by focusing solely on algorithmic inflexibility, not even when studying cases where human-decision making is *replaced* – rather than just *supported* – by automated decision-making. Human discretion can conflict with the formalized decision-making power of algorithms, as illustrated in Cases 1 and 3, or it can draw on automated procedures to impress its own agenda on the system’s operation, like in Case 2.

Thus, the well-established insight that algorithmic systems are socio-technical merits restatement in the context of the bureaucracy analogy. The theory of street-level algorithms highlights that algorithmic inflexibility is the source of dysfunctions in algorithmic systems. Following this insight, Alkhatib and Bernstein [1] discuss potential avenues for managing inflexibility, through recourse and tools that enable users to detect and report errors. Inflexibility is not an insurmountable obstacle for their theory, as it recognizes the potential of retraining and corrective design. Our approach supplements this account with the insight that recourse and user error management tools are effectively ways of *enabling discretionary control of uncertainties*. The way that these solutions are implemented, then, should take into account what

kinds of uncertainties persist in the system and how they are managed. In particular, we highlight that in many algorithmic systems, dysfunctions might not be due to algorithmic inflexibility alone. Rather, they might also involve discretionary power that accumulates at persistent locations of uncertainty. Algorithmic inflexibility is only one factor that can cause dysfunctions in algorithmic systems, and this might go unnoticed in recourse processes that target inflexibility as a source of error.

In addition to seeking to reapply the wealth of prior work on socio-technical systems, we draw a link to organization studies to highlight the societal importance and functioning of algorithmic systems. Here, issues created by bureaucratic systems are well established in prior literature. Algorithmic information processing provides the possibility to handle massive amounts of digital data at speeds beyond human capabilities. As such, algorithmic decision-making allows organizations to go further than traditional bureaucracies could in terms of automating processes in a rule-based manner. In extensively automated algorithmic systems, the gap between formal rules and their execution has become more streamlined than ever before. Consequently, as Caplan and boyd note [14], human control and grasp over the the systems’ actions is diminishing. Given these developments, it is more important than ever to develop theoretically-grounded accounts for analysing and understanding the “soft spots” where streamlined algorithmic processes can go awry. Our approach contributes such an account, basing on the Crozierian focus on systemic uncertainties.

Discretionary work adapts systems to uncertainties

While we have discussed well-known cases of algorithmic systems – risk prediction in courts, housing allocation, and worker-to-task matching in crowdwork – our approach contributes a fresh, theoretically grounded perspective on them by underlining *the role of reflexive discretion in adapting inflexible algorithmic processes to unanticipated circumstances*. From Crozierian perspective, the aim in automating decision-making is to establish control over uncertainties by rendering them calculable and predictable. As such, algorithms are used with the purpose of turning uncertainties in decision-making into risks [42] so they could be evaluated in an exact manner. Yet, it is often impossible to convert all uncertainties in decision-making into risks without some residue. Unanticipated, marginal, and novel cases always turn up. As such, attempts to assimilate them under a single, rigid decision-making logic are bound to fail in complex socio-technical contexts.

Our contribution to the bureaucracy analogy lies in highlighting that *the residue of uncertainty* – although often difficult to recognize – does not vanish in algorithmic systems, and that it is purposefully leveraged by human actors. Rather than eliminating issues of power, efforts at extensive automation *shift* the division and balance of power. Knowingly or not, automation may center power into new and sometimes unexpected hands.

In arguing that *discretionary work adapts algorithmic systems to uncertainties*, we highlight how human action works

to steer the system’s operation in accordance with its own purposes. It can either align with or go against the aims and intentions of designers or managers. Another way of saying this is that while automated processes can increase the *stability* of the system – the capability to maintain fixed functioning in the face of unanticipated events – discretion works to make the system *resilient*, in the sense of accommodating unanticipated events *as part of the system’s functioning* (c.f., [31, 32]). As we saw above, such resilience can be built into algorithmic systems in different ways, and can easily go unrecognized, especially when extensive automation is the aim.

As already noted above, Crozier argued that organizations can plummet into a bureaucratic vicious circle, where errors due to inflexibility are responded to with increasingly rigid standardization. Similarly, we posit in the case of algorithmic systems that the drive towards increasing systemic stability can result in an *algorithmic vicious circle*, where the work done by discretion in influencing system behavior is hidden under seemingly extensive automation. Based on our analysis, the design of algorithmic systems should seek to identify the work done by discretion in adapting automated processes to unanticipated circumstances. Instead of trying to eliminate sources of uncertainty in algorithmic systems, design could, instead, seek to *cultivate relevant uncertainties*, so that resilience is built into the system in a way that mitigates bias and discrimination in decision-making.¹

DESIGN DIRECTIONS TO CULTIVATE UNCERTAINTY

In this section, we discuss initial ideas which design could explore to account for uncertainty in algorithmic systems. These ideas are not tried recommendations, but rather potential directions for further inquiry. Our analytical approach highlights that uncertainty is a resource that discretion uses to manage inflexibility in even extensively automated systems. Future design of algorithmic systems could explore how this role of discretion can be supported in different contexts. To be sure, this idea is not new to HCI. However, the point is worth reiterating in the context of algorithmic systems which are often built with the express aim of automating away inefficient human discretion in decision processes. The bureaucracy analogy can help shift focus in the development of these systems to the way algorithms redistribute rather than eliminate discretion in decision-making. Below, we discuss this idea for the use of algorithms as support and replacement of human decision-making, respectively.

Design directions for supporting human-decision making

HCI has for long developed solutions, such as mixed-initiative interfaces [33], that learn iteratively from contextual information to optimize how automated processes communicate with users. Researchers have explored human-centered design frameworks to develop autonomous systems whose

¹We recognize that not all tensions in algorithmic systems are between humans and automated processes. Tensions may also emerge between different stakeholders like users and developers. Crozier’s theory highlights that uncertainties are a strategic resource that actors can draw upon in organizational power struggles. An analysis of uncertainties in algorithmic systems could potentially shed light into tensions between different groups of people, too.

behavior can be dynamically adjusted to accommodate unanticipated situations in teamwork [11]. Yet, decision-support systems can fail in many situations, due to the difficulty of controlling all relevant uncertainties involved in complex decision-making contexts. For instance, as we saw in Case 1, highly trained experts such as judges or doctors may be skeptical toward the usefulness of algorithmic recommendations (c.f., [74]). Further, research has also shown that support systems have a tendency to bias human decision-making [64, 65]. The expertise of a judge to oppose the algorithmic prediction may not be the case for all users, and thus human agents can be primed by the results of algorithmic systems.

In the context of support systems, our focus on cultivating uncertainties aligns with Yang et al’s [74] suggestion to build on the users’ established routines in designing algorithmic systems. This approach, based on what Tolmie et al. [70] called “unremarkable computing”, seeks to support already existing work practices rather than alter them through remarkable automated solutions. However, what this means in different contexts is an open question that requires further research (c.f., [73]). As our discussion of recidivism prediction in court illustrated, investigating institutional context is key to grasping the relevant uncertainties involved in discretionary decision-making. Again, this insight in itself is not novel to HCI [74], but merits repeating with respect to algorithmic systems that easily get criticized for inflexibility or praised for objectivity and efficiency.

With respect to cultivating uncertainty, one possible design direction that research could explore are recommender systems that implement a *negotiation style interaction process*. For example, when it comes to Case 1, the judge would be forced to conduct a risk prediction before seeing the recommendation from the system. If differences are large enough, both the judge and the system could be asked to list justifications for the prediction to negotiate the “appropriate” outcome. Such a system would, thus, set the judges’ evaluations against algorithmic predictions. Doing so could put greater emphasis on the discretionary work of interpreting the algorithms’ outputs, and might help offset the tendency to align decisions with computational predictions. Therefore, in algorithmic systems the interface could provide opportunities for mixed interaction approaches [33]. However, as said, this is one possible direction to explore in future research that follows from our focus on uncertainty, and it should not be taken as a tried design recommendation.

Finally, not all design implications are technical in nature – which, given the socio-technical nature of algorithmic systems, is not surprising. Strict rule-governed digital systems and institutional practices forced through them have been found to restrict reflexive discretion [37] and with information systems users often create workaround-solutions [58]. In this regard, our analytical approach inspires further research to explore different institutional arrangements which give rise to divergent uncertainties in interaction with algorithms. For instance, should institutions have clear approval processes for cases where a human agent decides to override algorithmic recommendations? What level of expertise should be re-

quired for such actions in any given context? We believe documenting these practises helps coordinating human discretion with respect to automated processes.

Design directions for replacing human work

Cases where algorithms seek to replace human discretion have been extensively discussed in the literature (e.g., [1, 12, 50]). Our analytical approach seeks to expand these discussions through emphasizing the work done by reflexive human discretion also in cases of extensive automation. This insight was illustrated by our discussion of Cases 2 and 3: even when the express aim is to subvert human discretion, decisions made within the algorithmic system often depend on a hybrid collaboration between humans and automated processes. Attempts at automation do not eliminate discretionary decision-making power in the system, but rather *redistribute it to work at locations of uncertainty* which may go unrecognized. That said, we acknowledge that there are good reasons to aim at automating decision-making: These systems are used to process large numbers of mundane decisions. In many contexts, it is impractical to demand increased leeway for discretionary work.

Nevertheless, we maintain that recognizing the invisible work done by discretion is crucial for designing algorithmic systems that replace human work. That is, in many cases while the *moment of decision* [1] is fully automated, the functions – and dysfunctions – of algorithmic systems emerge as a hybrid product of algorithmic and human action. Due to this, we believe that providing recourse or developing algorithms for increased reflexivity might not be sufficient to mitigate shortcomings in decision-making. In these solutions, the potentially unrecognized *hidden work* of discretion is not accounted for.

A design direction which could help bring the hybrid nature of algorithmic systems in view might be to focus on how such systems are presented. For example, as we saw in Case 2, the express motivations for deploying extensive automation in Los Angeles housing allocation were efficiency and impartiality in decision-making. The case illustrated that, in fact, the system did not unambiguously serve these purposes, while at the same time serving many that were not explicitly stated. Focusing on *the language and form in presenting algorithmic decision-making could help people to understand (and even challenge) algorithmic systems*. Already at the time of this writing, the opaqueness of algorithmic systems leads to challenges. In Fall 2019, there was a public outcry related to privacy concerns with voice assistants: the public was surprised to learn that part of the audio files necessary for operating voice assistants were listened and analysed by human workers tasked to do hidden work for quality control purposes. In other words, if laypersons are not aware of the hidden human work in algorithmic systems, they may experience a sense of violation upon finding out about it.

Our analytical approach suggests that an analysis of uncertainties could help identify how humans are involved in maintaining and coordinating decisions in highly automated systems. In such cases, a possible design direction to explore is the way in which discretion could be supported in stages like

data collection and entry. One possibility is to include information about the actors who make discretionary decisions as an additional input for the algorithmic decision process. First, this could be an additional field for data collection and entry staff. Second, algorithmic processes could also use the identity of different persons doing data collection and entry as an additional input field, thus allowing analysis on systemic differences. Again, we want to emphasize that further research is required to explore the potential pitfalls and advantages of this initial idea in different contexts.

CONCLUSION

In this essay, we have elaborated on bureaucracy as an analogy for studying algorithmic systems. Our account builds on Crozier's [18] theory of bureaucratic systems as dysfunctional organizations which cannot adapt to unanticipated situations. With the help of three illustrative cases documented in prior literature, we demonstrated how decision-making in algorithmic systems should be understood as the mutual product of automated processes and discretionary human action. Analogously to rigid rules in Crozierian bureaucracies, we argued that automation in algorithmic systems does not eliminate discretion. Instead, discretion is redistributed to work at locations where it can leverage persistent uncertainty to influence decisions made in the system.

While the cases we discussed illustrate some of the analytical benefits that a focus on uncertainties can provide, we also see a need for in-depth empirical research to document the different ways of controlling uncertainty in various algorithmic systems. We have focused on ambiguous cases where consequences cannot be attributed to algorithms or human discretion alone. In our view, extending the bureaucracy analogy to account for such cases can help analyze power in algorithmic systems as the joint outcome of human discretion and automated processes. Most importantly, the resulting account identifies *uncertainty* as the key concept in analyzing how inflexibility is managed in algorithmic systems.

ACKNOWLEDGMENTS

This research was supported by the KONE Foundation project *Algorithmic Systems, Power and Interaction*. This article is, in part, based on ideas discussed at the NOS-HS Nordic Perspectives on Algorithmic Systems workshop in 2019 in Stockholm. We thank the workshop participants, along with further colleagues at Stockholm University, for their feedback and support. We also thank the anonymous reviewers for their helpful comments.

REFERENCES

- [1] Ali Alkhatib and Michael Bernstein. 2019. Street-Level Algorithms: A Theory at the Gaps Between Policy and Decisions. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*. ACM Press, Glasgow, Scotland Uk, 1–13. DOI: <http://dx.doi.org/10.1145/3290605.3300760>
- [2] Ali Alkhatib, Michael S. Bernstein, and Margaret Levi. 2017. Examining Crowd Work and Gig Work Through The Historical Lens of Piecework. In *Proceedings of*

- the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*. ACM Press, Denver, Colorado, USA, 4599–4616. DOI : <http://dx.doi.org/10.1145/3025453.3025974>
- [3] Mike Ananny. 2016. Toward an Ethics of Algorithms: Convening, Observation, Probability, and Timeliness. *Science, Technology, & Human Values* 41, 1 (Jan. 2016), 93–117. DOI : <http://dx.doi.org/10.1177/0162243915606523>
- [4] Mike Ananny and Kate Crawford. 2018. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society* 20, 3 (March 2018), 973–989. DOI : <http://dx.doi.org/10.1177/1461444816676645>
- [5] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine Bias. *ProPublica* (2016). <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- [6] Jane Bambauer and Tal Zarsky. 2018. The Algorithm Game. *94 Notre Dame Law Review 1; Arizona Legal Studies Discussion Paper No. 18-09* (2018). <https://ssrn.com/abstract=3135949>
- [7] Jens Beckert. 2016. *Imagined Futures*. Harvard University Press. Google-Books-ID: BhSSJgEACAAJ.
- [8] David Beer. 2016. *Metric Power*. Palgrave Macmillan Limited, London, UNITED KINGDOM.
- [9] Mark Bovens and Stavros Zouridis. 2002. From Street-Level to System-Level Bureaucracies: How Information and Communication Technology is Transforming Administrative Discretion and Constitutional Control. *Public Administration Review* 62, 2 (2002), 174–184. DOI : <http://dx.doi.org/10.1111/0033-3352.00168>
- [10] Geoffrey C. Bowker and Susan Leigh Star. 1999. *Sorting things out: Classification and its consequences*. MIT Press, Cambridge, Mass. http://www.worldcat.org/title/sorting-things-out-classification-and-its-consequences/oclc/41273009&referer=brief_results
- [11] Jeffrey Bradshaw, Maarten Sierhuis, Alessandro Acquisti, Paul Feltovich, Robert Hoffman, Renia Jeffers, Debbie Prescott, Niranjani Suri, Andrzej Uszok, and Ron Van Hoof. 2003. Adjustable Autonomy and Human-Agent Teamwork in Practice: An Interim Report on Space Applications. In *Agent Autonomy. Multiagent Systems, Artificial Societies, and Simulated Organizations (International Book Series), vol 7*, Henry Hexmoor, Cristiano Castelfranchi, and Rino Falcone (Eds.). Springer, 243–280.
- [12] Taina Bucher. 2012. Want to be on the top? Algorithmic power and the threat of invisibility on Facebook. *New Media & Society* 14, 7 (Nov. 2012), 1164–1180. DOI : <http://dx.doi.org/10.1177/1461444812440159>
- [13] Taina Bucher. 2017. The algorithmic imaginary: Exploring the ordinary affects of Facebook algorithms. *Information, Communication & Society* 20, 1 (Jan. 2017), 30–44. DOI : <http://dx.doi.org/10.1080/1369118X.2016.1154086>
- [14] Robyn Caplan and danah boyd. 2018. Isomorphism through algorithms: Institutional dependencies in the case of Facebook. *Big Data & Society* 5, 1 (2018). DOI : <http://dx.doi.org/10.1177/2053951718757253>
- [15] Angèle Christin. 2017. Algorithms in practice: Comparing web journalism and criminal justice. *Big Data & Society* 4, 2 (Dec. 2017), 2053951717718855. DOI : <http://dx.doi.org/10.1177/2053951717718855>
- [16] Harry M. Collins and Martin Kusch. 1998. *The Shape of Actions: What Humans and Machines Can Do*. MIT Press. Google-Books-ID: W2IYF9RWpQwC.
- [17] Kelley Cotter. 2018. Playing the visibility game: How digital influencers and algorithms negotiate influence on Instagram. *New Media & Society* 21, 4 (2018).
- [18] Michel Crozier. 1964. *The Bureaucratic Phenomenon*. Tavistock.
- [19] Michel Crozier and Erhard Friedberg. 1980. *Actors and Systems: The Politics of Collective Action*. University of Chicago Press.
- [20] Paul Dourish. 2016. Algorithms and their others: Algorithmic culture in context. *Big Data & Society* (2016). DOI : <http://dx.doi.org/10.1177/2053951716665128>
- [21] Motahhare Eslami, Karrie Karahalios, Christian Sandvig, Kristen Vaccaro, Aimee Rickman, Kevin Hamilton, and Alex Kirlik. 2016. First I "like" it, then I hide it. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*. 2371–2382. DOI : <http://dx.doi.org/10.1145/2858036.2858494>
- [22] Virginia Eubanks. 2018. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St. Martin's Press, New York, NY.
- [23] Batya Friedman, Peter H Kahn, and Alan Borning. 2008. Value sensitive design and information systems. *The handbook of information and computer ethics* (2008), 69–101.
- [24] Tarleton Gillespie. 2014. The Relevance of Algorithms. In *Media technologies: Essays on communication, materiality, and society*. The MIT Press. <http://mitpress.universitypressscholarship.com/view/10.7551/mitpress/9780262525374.001.0001/upso-9780262525374-chapter-9>
- [25] Tarleton Gillespie. 2017. Algorithmically recognizable: Santorum's Google problem, and Google's Santorum problem. *Information, Communication & Society* 20, 1 (Jan. 2017), 63–80. DOI : <http://dx.doi.org/10.1080/1369118X.2016.1199721>

- [26] Tarleton Gillespie. 2018. *Custodians of the Internet: Platforms, content moderation, and the hidden decisions that shape social media*. Yale University Press.
- [27] Bryce Goodman. 2016. A Step Towards Accountable Algorithms?: Algorithmic Discrimination and the European Union General Data Protection. *29th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain*. Nips (2016), 1–7.
- [28] Mary L. Gray and Siddharth Suri. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Houghton Mifflin Harcourt.
- [29] Mary L. Gray, Siddharth Suri, Syed Shoaib Ali, and Deepti Kulkarni. 2016. The Crowd is a Collaborative Network. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. ACM, New York, NY, USA, 134–147. DOI: <http://dx.doi.org/10.1145/2818048.2819942>
- [30] Jesse Haapoja and Airi Lampinen. 2018. 'Datafied' reading: Framing behavioral data and algorithmic news recommendations. In *Proceedings of the 10th Nordic Conference on Human-Computer Interaction - NordiCHI '18*. ACM Press. DOI: <http://dx.doi.org/10.1145/3240167.3240194>
- [31] Orit Halpern, Robert Mitchell, and Bernard Dionysius Geoghegan. 2017. The Smartness Mandate: Notes toward a Critique. *Grey Room* 68 (Sept. 2017), 106–129. DOI: http://dx.doi.org/10.1162/GREY_a_00221
- [32] C S Holling. 1973. Resilience and Stability of Ecological Systems. *Annual Review of Ecology and Systematics* 4, 1 (1973), 1–23. DOI: <http://dx.doi.org/10.1146/annurev.es.04.110173.000245>
- [33] Eric Horvitz. 1999. Principles of Mixed-initiative User Interfaces. *Proceedings of the 1999 CHI Conference on Human Factors in Computing Systems - CHI '99* (1999), 159–166. DOI: <http://dx.doi.org/10.1145/302979.303030>
- [34] Lucas D. Introna. 2011. The Enframing of Code: Agency, Originality and the Plagiarist. *Theory, Culture & Society* 28, 6 (Nov. 2011), 113–141. DOI: <http://dx.doi.org/10.1177/0263276411418131>
- [35] Lucas D. Introna. 2016. Algorithms, Governance, and Governmentality: On Governing Academic Writing. *Science, Technology, & Human Values* 41, 1 (Jan. 2016), 17–49. DOI: <http://dx.doi.org/10.1177/0162243915587360>
- [36] Steven J. Jackson. 2014. Rethinking Repair. In *Media technologies : essays on communication, materiality, and society*. The MIT Press.
- [37] Gabriella Jansson and Gissur Ó Erlingsson. 2014. More E-Government, Less Street-Level Bureaucracy? On Legitimacy and the Human Side of Public Administration. *Journal of Information Technology & Politics* 11, 3 (July 2014), 291–308. DOI: <http://dx.doi.org/10.1080/19331681.2014.908155>
- [38] Elliott Jaques. 1976. *A general theory of bureaucracy*. Heinemann; New York: Halsted Press.
- [39] Sanjay Kairam and Jeffrey Heer. 2016. Parting Crowds: Characterizing Divergent Interpretations in Crowdsourced Annotation Tasks. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. ACM, New York, NY, USA, 1637–1648. DOI: <http://dx.doi.org/10.1145/2818048.2820016>
- [40] Eugene C Kim. 2007. YouTube: Testing the safe harbors of digital copyright law. *S. Cal. Interdisc. LJ* 17 (2007), 139.
- [41] Ulrike Klinger and Jakob Svensson. 2018. The End of Media Logics? On Algorithms and Agency. *New Media & Society* 20, 12 (12 2018), 4653–4670. DOI: <http://dx.doi.org/10.1177/1461444818779750>
- [42] Frank Knight. 1921. *Risk, Uncertainty, and Profit*. Dover Publications.
- [43] Min Kyung Lee, Daniel Kusbit, Evan Metsky, and Laura Dabbish. 2015. Working with Machines. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*. ACM Press, New York, New York, USA, 1603–1612. DOI: <http://dx.doi.org/10.1145/2702123.2702548>
- [44] Michael Lipsky. 1983. *Street-Level Bureaucracy: The Dilemmas of the Individual in Public Service*. Russell Sage Foundation.
- [45] 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 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16)*. ACM, New York, NY, USA, 1057–1062. DOI: <http://dx.doi.org/10.1145/2851581.2886426>
- [46] Claire Maiers. 2017. Analytics in action: users and predictive data in the neonatal intensive care unit. *Information Communication and Society* 20, 6 (2017), 915–929. DOI: <http://dx.doi.org/10.1080/1369118X.2017.1291701>
- [47] Brian McInnis, Dan Cosley, Chaebong Nam, and Gilly Leshed. 2016. Taking a HIT: Designing Around Rejection, Mistrust, Risk, and Workers' Experiences in Amazon Mechanical Turk. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 2271–2282. DOI: <http://dx.doi.org/10.1145/2858036.2858539>

- [48] Anne Milgram. 2013. Why smart statistics are the key to fighting crime. (2013).
https://www.ted.com/talks/anne_milgram_why_smart_statistics_are_the_key_to_fighting_crime
- [49] Gareth Morgan. 2006. *Images of organization: The executive edition*. SAGE Publications.
- [50] Judith Möller, Damian Trilling, Natali Helberger, and Bram van Es. 2018. Do not blame it on the algorithm: An empirical assessment of multiple recommender systems and their impact on content diversity. *Information, Communication & Society* 21, 7 (July 2018), 959–977. DOI :
<http://dx.doi.org/10.1080/1369118X.2018.1444076>
- [51] Daniel Neyland and Norma Möllers. 2017. Algorithmic IF ... THEN rules and the conditions and consequences of power. *Information, Communication & Society* 20, 1 (Jan. 2017), 45–62. DOI :
<http://dx.doi.org/10.1080/1369118X.2016.1156141>
- [52] Helen Nissenbaum. 2005. Values in Technical Design. In *Encyclopedia of Science, Technology, and Ethics*, Carl Mitcham (Ed.). MacMillan, New York, lxxi–lxx.
- [53] Cathy O’Neil. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown/Archetype.
- [54] Antti Oulasvirta and Kasper Hornbæk. 2016. HCI Research as Problem-Solving. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*. ACM Press, New York, New York, USA, 4956–4967. DOI :
<http://dx.doi.org/10.1145/2858036.2858283>
- [55] Sarah Pink, Minna Ruckenstein, Robert Willim, and Melisa Duque. 2018. Broken data: Conceptualising data in an emerging world. *Big Data & Society* 5, 1 (Jan. 2018), 2053951717753228. DOI :
<http://dx.doi.org/10.1177/2053951717753228>
- [56] Theodore M. Porter. 1995. *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton University Press.
<http://www.jstor.org/stable/j.ctt7sp8x>
- [57] Emilee Rader and Rebecca Gray. 2015. Understanding User Beliefs About Algorithmic Curation in the Facebook News Feed. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15* (2015), 173–182. DOI :
<http://dx.doi.org/10.1145/2702123.2702174>
- [58] David Randall, Richard Harper, and Mark Rouncefield. 2007. *Fieldwork for design: theory and practice*. Springer Science & Business Media.
- [59] Luis Ignacio Reyes-Galindo. 2016. Automating the Horae: Boundary-work in the age of computers. *Social Studies of Science* (2016).
<https://arxiv.org/pdf/1603.03824.pdf>
- [60] Minna Ruckenstein and Natasha Schüll. 2017. The Datafication of Health. *Annual Review of Anthropology* 46, 1 (2017).
- [61] Nick Seaver. 2018. What should an anthropology of algorithms do? *Cultural anthropology* 33, November (aug 2018), 1–8. DOI :
<http://dx.doi.org/10.14506/ca33.3.04>
- [62] Herbert A Simon. 1972. Theories of bounded rationality. *Decision and organization* 1, 1 (1972), 161–176.
- [63] Rashmi Sinha and Kirsten Swearingen. 2002. The Role of Transparency in Recommender Systems. In *CHI '02 Extended Abstracts on Human Factors in Computing Systems (CHI EA '02)*. ACM, New York, NY, USA, 830–831. DOI :
<http://dx.doi.org/10.1145/506443.506619>
- [64] Linda Skitka, Kathleen Mosier, and Mark Burdick. 1999. Does automation bias decision-making? *International Journal of Human-Computer Studies* 51, 5 (1999).
- [65] Linda Skitka, Kathleen Mosier, and Mark Burdick. 2000. Accountability and automation bias. *International Journal of Human-Computer Studies* 52, 4 (2000).
- [66] Susan Leigh Star and Anselm Strauss. 1999. Layers of Silence, Arenas of Voice: The Ecology of Visible and Invisible Work. *Computer Supported Cooperative Work (CSCW)* 8, 1 (March 1999), 9–30. DOI :
<http://dx.doi.org/10.1023/A:1008651105359>
- [67] Anselm Strauss. 1988. The Articulation of Project Work: An Organizational Process. *The Sociological Quarterly* 29 (1988).
- [68] Lucy A. Suchman. 1987. *Plans and Situated Actions: The Problem of Human-Machine Communication* (2 edition ed.). Cambridge University Press, Cambridge Cambridgeshire ; New York.
- [69] Suzanne L Thomas, Dawn Nafus, and Jamie Sherman. 2018. Algorithms as fetish: Faith and possibility in algorithmic work. *Big Data & Society* 5, 1 (6 2018), 205395171775155. DOI :
<http://dx.doi.org/10.1177/2053951717751552>
- [70] Peter Tolmie, James Pycocock, Tim Diggins, Allan MacLean, and Alain Karsenty. 2002. Unremarkable Computing. *Proceedings of the 2002 SIGCHI Conference on Human Factors in Computing Systems - CHI '02* (2002). DOI :
<http://dx.doi.org/10.1145/503376.503448>
- [71] Max Weber. 2013. *Bureacracy*. Routledge, 196–244.
- [72] Michele Willson. 2017. Algorithms (and the) everyday. *Information, Communication & Society* 20, 1 (Jan. 2017), 137–150. DOI :
<http://dx.doi.org/10.1080/1369118X.2016.1200645>

- [73] Qian Yang. 2018. Machine Learning as a UX Design Material: How Can We Imagine Beyond Automation, Recommenders, and Reminders? *2018 AAAI Spring Symposium Series* (2018). <https://www.aaai.org/ocs/index.php/SSS/SSS18/paper/view/17471>
- [74] Qian Yang, Aaron Steinfeld, and John Zimmerman. 2019. Unremarkable AI: Fitting Intelligent Decision Support into Critical, Clinical Decision-Making Processes. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19* (2019). DOI: <http://dx.doi.org/10.1145/3290605.3300468>
- [75] Malte Ziewitz. 2019. Rethinking gaming: The ethical work of optimization in web search engines. *Social Studies of Science* (aug 2019), 030631271986560. DOI: <http://dx.doi.org/10.1177/0306312719865607>
- [76] Indrė Žliobaitė. 2017. Measuring discrimination in algorithmic decision making. *Data Mining and Knowledge Discovery* 31, 4 (2017), 1060–1089.