

ORIGINAL RESEARCH ARTICLE

Ethical Considerations in Algorithm Construction: Evidence from Two Human-Resources Algorithms

Clotilde Coron¹ and Simon Porcher²

¹RITM Research Centre in Economics & Management, Université Paris Saclay, Sceaux, France;

²Dauphine recherches en management (DRM), Université Paris Dauphine-PSL, Paris, France

Abstract

Organizations increasingly rely on algorithms for decision-making, which raises significant ethical issues. In this paper, we provide a detailed case study of the development and deployment of two human resources (HR) algorithms in a major French digital company. Our findings show that these ethical issues reflect the ethical considerations of the various stakeholders involved in the process, including data scientists, HR practitioners, and legal experts. We discuss how these considerations intervene during the decision-making process in algorithm design and usage, offering insights for both academics and practitioners into how ethical issues are approached by different actors.

Keywords: *Algorithm; Ethics; Human resource management; Artificial intelligence*

Handling editor: Wafa Ben Khaled; Received: 12 July 2023; Accepted: 7 May 2025; Published: 16 March 2026

Algorithms, as mathematical functions modelled by organizations and interacting with users, are increasingly employed in organizational contexts, shaping various aspects of decision-making and interaction (Neyland, 2015). While there is evidence that supports the positive effects of algorithms on human capabilities and decision-making (Murray et al., 2021; Wilson & Daugherty, 2018), scholars have raised concerns about ethical issues related to their use, particularly in terms of accountability (Martin, 2019b; Neyland, 2015), fairness (Kim et al., 2018), and transparency (Martin, 2019b). As algorithms become increasingly integral to organizational operations (Meijerink & Bondarouk, 2023), addressing these ethical issues is imperative.

Despite the growing body of research on algorithmic ethics (Orlikowski & Scott, 2016; Turner, 2009), a significant gap remains in understanding how various stakeholders – each with differing perspectives – interact during the algorithmic design process, particularly around ethical issues, and how these considerations shape decision-making. Discussions about algorithmic parameters and how their creators negotiate ethical concerns – essentially their ‘expression of a view on how things ought to be or not to be, what is good or bad, or desirable or undesirable’ (Kraemer et al., 2011, p. 252) – remain largely unexplored. Yet these discussions influence key

decisions in both the development and deployment of algorithms, ultimately shaping their outcomes. Given the far-reaching implications of algorithmic decisions on individuals’ lives (Kellogg et al., 2020; Leicht-Deobald et al., 2019), it is crucial to examine how ethical issues are considered among stakeholders, the decisions that emerge from these discussions, and how ethical considerations are either embedded or overlooked during development.

Designers build their technologies based on assumptions about the world and how they believe their technology will integrate into it (Martin, 2019b). For instance, they may deliberately prioritize system performance over transparency (Langer & König, 2023), reflecting value choices that often remain implicit. A rich body of ethnographic work on the technology sector has shown that, in Silicon Valley companies, competition among engineers fosters the development of algorithms optimized for scalability (Turner, 2009). This race for efficiency drives solutions that favour large-scale optimization, often at the expense of clarity and comprehensibility. However, understanding the organizational structures and cultures surrounding algorithm design is not enough to fully grasp the values embedded in these systems. To gain deeper insights into algorithmic ethics, it is essential to examine the motivations, tensions, and trade-offs faced by developers.

*Corresponding author: Simon Porcher Email: simon.porcher@dauphine.psl.eu

This is the perspective from which our study emerges. This research aims to understand better the way algorithms are developed (Ananny & Crawford, 2018; Pasquale, 2015) by analysing the interactions and ethical considerations that intervene in the development and tests of human resources management (HRM) algorithms. We examine how different actors within an organization – data scientists, human resources (HR) practitioners, and legal experts – navigate ethical issues in constructing two algorithms developed by a major French digital company, anonymized here as ‘Digix’. One algorithm provides personalized training course recommendations for employees, while the other automates the pre-screening of curricula vitae (CVs) for external recruitment. By comparing the simultaneous development of these two systems within the same organization, we uncover key insights into the ethical tensions inherent in algorithm design and deployment, as well as how ethical considerations among stakeholders intervene during the process.

To understand the perceptions of the different actors about the algorithms and to understand the choices they make, we used a combination of methodologies and materials. Participant observation and internal documentation provided us with unique access to the actors, documents, and data of the algorithms. Additionally, qualitative interviews with data scientists were also conducted to gain insights into how the algorithms function. The analysis enabled us to identify the key ethical considerations, around three ethical issues identified in the literature – accountability, fairness, and transparency – related to the construction of algorithms.

Our findings offer significant contributions to the literature on algorithms in organizational settings. Unlike many studies that emphasize a clear divide between designers who embrace algorithms and users who resist them (Kellogg et al., 2020), we find a more nuanced relationship in practice (Raisch & Krakowski, 2021). Even within the design process itself (Glaser et al., 2021), the various actors, such as data scientists, legal experts, and HR practitioners, do not form a unified group but rather offer varied perspectives, particularly regarding the selection of variables. Our research shows that algorithms reflect the differing priorities and values within each group (Christin, 2020). This paper contributes to the literature on algorithms in management by offering a detailed case study of the ethical considerations surrounding the development and test of two algorithms along the chain of stakeholders. Understanding these dynamics is central to acknowledging the various roles that stakeholders play in the development and test of algorithms.

Our study is also of interest to scholars in business ethics. The results reveal how various actors approach ethical issues differently in terms of algorithm development, shedding light on divergent visions of accountability, fairness, and transparency. For example, HR managers emphasize transparency as a

way to enable end users to evaluate and challenge algorithmic outputs, although this raises concerns about accountability and fairness, such as relying on more straightforward, less equitable algorithms. Legal experts concentrate on managing input variables to maintain fairness and reduce accountability risks. On the other hand, data scientists view the reliability of algorithmic classifications as a measure of fairness, using transparency to shift accountability to end users when recommendations are flawed.

The remainder of the paper is organized as follows. The next section reviews the literature. The subsequent sections present the methodology, the algorithms under study, and the results of the case study. A discussion and conclusion follow.

Literature review

The increasing prevalence of algorithms within organizations

Algorithms are increasingly present within organizations. They have different objectives: personalizing experience via personalized recommendations, automating some tasks such as workforce planning, increasing performance, improving decision-making by using algorithms to suggest better decisions, for example in the field of recruitment or finance (Coron, 2022). The development of algorithms and their implementation within organizations require the collaboration of various stakeholders.

Management scholars have studied algorithms, mainly focusing on how they are built and the representations affecting how algorithmic outputs are used (for a literature review, see Christin, 2020). A substantial body of literature analyses how organizational cultures and structures shape the construction of algorithms. For example, Neff and Stark (2003) and Irani (2015) illustrate how the ‘beta’ mindset of the world prevalent in technology companies and the precarious nature of employment that revolves around the development of algorithms influence these dynamics. These organizational influences shape the priorities embedded in algorithmic systems, but are also the source of tensions among stakeholders. Developers, legal experts, managers, and end users often have competing expectations about how transparent, accountable, and adaptable an algorithm should be. While technical teams may advocate for explainability, legal experts and end users demand fairness and oversight, creating friction over what aspects of an algorithm should be disclosed or controlled (Selbst et al., 2019). These discussions are rarely resolved, leading to compromises that directly impact algorithm design, often reinforcing opacity rather than resolving it.

The importance of personal values and context shapes algorithmic development and decision-making. Algorithms make choices about what information to display, hide, or use,

and reflect the 'conscious and subconscious assumptions and ideas of their creators' (Hodder, 2009), thus raising potential accountability issues. For example, an algorithm designed to screen job applicants might prioritize certain keywords or educational backgrounds based on the developers' assumptions about what constitutes a 'good' candidate. This brings up significant accountability concerns, as developers, organizations, and regulators struggle to determine accountability for biased or unfair outcomes.

The ethics of algorithms: an overview

Ethics in algorithmic design focus on determining what is morally right or wrong (Ananny, 2016; Merrill, 2011). In algorithmic contexts, ethics translate into considerations such as regulatory compliance, outcome optimization, and how ethical standards emerge from the personal values and beliefs of designers (Ananny, 2016; Hursthouse, 1999). At the construction stage, developers must navigate trade-offs between interpretability, accuracy, and commercial viability. Optimizing for predictive power often comes at the cost of transparency, embedding structural opacity into the algorithm from its inception (Selbst et al., 2019). For instance, stakeholders may prioritize efficiency over fairness, influencing which features are optimized and which biases remain unaddressed. These competing pressures create an ethical grey zone, where ethical principles often collide with the realities of algorithm development and where formal guidelines take a backseat to the subjective decisions and trade-offs made by designers and developers. During the development phase, these challenges raise issues of accountability, fairness, and transparency, as subjective decisions can result in ethically questionable outcomes.

Accountability: who is responsible for algorithmic decisions?

Accountability is a central concern for organizations that rely on algorithms (Pasquale, 2015). By deploying algorithms that operate in a value-laden¹ (Martin, 2019b) and context-specific manner (Mittelstadt et al., 2016), organizations voluntarily become active participants in the decision-making process, assuming accountability for the consequences, including any harm created, ethical principles violated and rights diminished by the decision system (Martin, 2019b). Algorithms have social consequences and moral implications, and organizations cannot hide behind these algorithms (Faraj et al., 2018). They must be held accountable, particularly when the rules of decisions

¹Martin (2019b) argues that algorithms are not neutral but rather value-laden with preferences for certain outcomes, specified (constructed) by individuals in their design, implementation, and use.

are hidden to users but well-understood by those who created them.

Algorithms, like human decision-makers, are susceptible to mistakes, which may stem from various factors, such as flawed inputs, poorly designed models, or execution errors (Martin, 2019a). While human errors are common and typically subject to straightforward accountability, the question of who should be held responsible for algorithmic errors is more complex. As algorithms take over more decision-making tasks from humans, determining how accountability is distributed becomes critical.

Martin (2019a) argues that the design of algorithms plays a crucial role in addressing issues of accountability. For example, designers can increase the degree of social embeddedness and capacity for reflection of the algorithm (Martin, 2019a). Social embeddedness involves acknowledging the impact of contextual factors on decision-making and incorporating sensitivity tests. For example, users can compare the classification they receive with that of other users who share similar characteristics, enabling them to identify and evaluate potential errors. Reflection in decision-making entails the ability to revisit and challenge previous results, allowing algorithms to incorporate mechanisms for re-evaluation to ensure accurate classification.

Accountability extends beyond merely identifying mistakes and encompasses broader considerations of accountability. Algorithms are constructed within specific governance frameworks, with developers determining how accountability is delegated within the decision process (Martin, 2019b). Some algorithms are designed to operate autonomously, thereby assuming full accountability for decisions and shielding individuals from understanding their internal operations (Introna, 2016; Martin, 2019b). This opacity in algorithms results in greater accountability for their designers, as external observers are unable to scrutinize their functionality.

Algorithms can lead to unfair decisions

Algorithms can lead to unfair outcomes, where fairness implies that decisions should not lead to unjust, discriminatory, or disparate consequences (Starke et al., 2022). Such decisions may discriminate against certain groups or shape individuals' perceptions of their situation (Mittelstadt et al., 2016; Pasquale, 2015). As Martin (2019a) highlights, mistakes themselves are not necessarily unethical or unfair, as errors are common in business operations. However, mistakes that remain unaddressed or are exacerbated are unfair.

These errors may occur as a result of preexisting social values, technical limitations, or judicial constraints inherent to a given context of use (Mittelstadt et al., 2016). Social bias may arise from cultural or organizational values. Biases related to technical and judicial constraints may stem from the datasets

used for learning, missing data or selection biases, or from attempts to minimize aggregated prediction errors. For example, using sensitive attributes, such as gender or race, to infer behavioural patterns may lead to discriminatory outcomes, which are difficult to challenge due to the algorithm's inherent opacity. Algorithms, which often learn from past decisions, have the capacity to propagate mistakes, impacting a multitude of decisions rapidly.

The problem of transparency: understanding algorithmic decisions

Transparency is an important yet elusive feature of algorithmic decisions. Algorithms may be perceived as cruel, impenetrable, opaque, and intentionally poorly accessible (Faraj et al., 2018; Mittelstadt et al., 2016). Researchers often describe algorithms as 'black boxes' due to their complex and opaque nature, which obscures the internal workings and decision-making processes (Ananny & Crawford, 2018; Pasquale, 2015). Algorithms may seem objective, but their inner workings are often unclear, making it difficult to understand how they produce results. While outputs from an algorithm are easy to observe, understanding and explaining them requires familiarity with the algorithm (Von Krogh, 2018). The variables that are used in algorithms and their inherent causality are often obscure to external parties, resulting in opacity (Ananny & Crawford, 2018; Martin, 2019b).

Modern algorithms are often inaccessible for thorough examination. Introna (2016, p. 25) writes that 'they seem to operate under the surface or in the background'. The lack of transparency of algorithms stems from their lack of openness to direct inspection. Additionally, their use of dynamic inputs means that examining their static code alone does not offer a clear understanding of their behaviour during execution (Orlikowski & Scott, 2016). Therefore, comprehending algorithms requires observing their dynamics in action, presenting significant challenges for studying their impact on company functions. Practitioners also express unease about algorithms that produce recommendations that may contradict conventional wisdom, especially when hidden biases are difficult to detect in the underlying code (Wilson & Daugherty, 2018). While academic studies generally provide detailed explanations of how algorithms work, describing the variables, data processing methods, and optimization goals, the complexity of algorithmic design often results in undocumented assumptions, proprietary constraints, and emergent behaviours that even developers struggle to anticipate (Cheng & Hackett, 2021).

Indeed, algorithmic opacity is often embedded during development, as training processes, data choices, and proprietary constraints lead to emergent behaviours that even developers struggle to interpret (Burrell, 2016). This opacity is not merely a result of complex code but stems from institutional priorities

that favour efficiency over transparency (Pasquale, 2015). If the criteria and processes behind these predictions are opaque, it becomes nearly impossible for individuals to challenge unfair decisions, thereby obscuring the decision-making process and potentially embedding historical biases and perpetuating inequalities (Christin, 2020). These biases can shape users' experiences and outcomes, further complicating efforts to address and correct them (Orlikowski & Scott, 2015). In addition to technical complexity, these layers of opacity are also shaped by competing stakeholder interests. Organizations often selectively disclose information about their algorithms. Without clear governance structures, these considerations remain fragmented, with different actors advocating for conflicting levels of transparency (Binns, 2018).

Transparency is crucial because algorithms are typically unpredictable and poorly explainable, making them difficult to control, monitor, and correct. Enhancing transparency can improve understanding of how algorithms function, which is key to effective management. Besides, the lack of transparency raises concerns about shifts in influence and power dynamics. Overall, the opacity of algorithmic systems presents a critical challenge: while ethical frameworks offer theoretical guidance, there remains a lack of concrete mechanisms to ensure accountability and mitigate bias in practice.

Many 'algorithmic imaginaries' (Christin, 2020) are shaped based on the interactions of individuals with algorithms or information shared between peers. For example, hotel cleaners' work practices are deeply connected to guest feedback on platforms such as Tripadvisor and how the platform's algorithm incorporates comments in hotel rankings (Orlikowski & Scott, 2016). The opaque nature of algorithms thus allows them to control individuals, who may be punished or rewarded by algorithmic decisions (Kellogg et al., 2020).

On the other hand, several studies of algorithms (Christin, 2017; Lowrie, 2017) suggest that individuals with a satisfactory understanding of the algorithms can 'game' their results. In the gig economy, workers adapt their behaviours to algorithm management by using anticipatory compliance practices, for example by keeping emotions in check, to ensure their continued participation on the platform (Bucher et al., 2021). Individuals are thus involved in algorithmic decisions even though they are not actors in the decision-making (Martin, 2019b). By using a set of tactics to resist the control of algorithms (Kellogg et al., 2020), workers try to regain autonomy. Thus, the opacity of the algorithm does not matter *per se* but becomes ethically significant in relation to others.

In summary, algorithms are increasingly embedded in critical decision-making processes, shaping everything from hiring to promotions and customer service. However, as these systems become more influential, the ethical issues of accountability, fairness, and transparency cannot be ignored. To address these challenges, it is vital to examine the collaborative dynamics

among stakeholders involved in algorithm design. Understanding how ethical issues are considered during the development phase is key to addressing the ethical implications of algorithm use in practice. In the following sections of the paper, we seek to address this gap by exploring how accountability, fairness, and transparency are considered and implemented in practice by those involved in algorithm development and deployment.

Methodology and context

Data collection

We used a combination of participant observation, internal documentation, and interviews to study the development and use of two algorithms at Digix from 2016 to 2017 through multi-sited ethnography (Marcus, 1995). Digix, a large French telecommunications company with approximately 90,000 employees, aimed to gain insights into the algorithms' development and use, and stakeholders' perceptions of them. While ethics were not the initial focus of the data collection, the collected material revealed how stakeholders involved in the development of algorithms considered ethical issues in practice.

Participant observation and internal documentation

One author of the paper worked as a big data and HR project manager at Digix from January 2016 to September 2017, leading a project on the training suggestion algorithm. In this context, she conducted 'covert research' (Roulet et al., 2017) by working with the training department, which was in charge of defining the expectations regarding one of the algorithms, data, the HR data management team in charge of overseeing the HR information systems and providing data, along with two legal experts responsible for ensuring compliance with French data privacy laws. This involved numerous types of communications, including meetings with minutes, emails, and phone calls, involving individuals from various sites. This data collection, based on covert research, was facilitated by the fact that the collected material was part of her job, ensuring that the research itself had no impact on the company or its employees. The whole internal documentation is a comprehensive collection of over 230 documents, encompassing meeting minutes, framework documents, presentations made at project meetings and at various stages of the project's updates, budget and schedule monitoring reports, draft emails sent to employees to solicit their participation, question-and-answer documents on the algorithm, and survey feedback, among others. Smaller-scale internal documentation on the CV preselection algorithm, which another project manager oversaw, could also be retrieved.

While this covert research was integrated into her job, minimizing the impact on the company and employees, maintaining analytical distance remained a challenge (Schouten & McAlexander, 1995). To address this fact, the research included external interviews and the analysis began after the author left the company, facilitating a sense of detachment and a more objective research stance. Finally, co-authoring with someone outside the company further strengthened the analytical objectivity.

Interviews

To complement our knowledge of the development of the algorithms, a 2-h semi-structured interview was conducted at the end of the project (in 2017) with one of the data scientists who worked on the training suggestion algorithm ('Data scientist A'), and another 1.5-h semi-structured interview was conducted with the data scientist who worked on the CV preselection algorithm ('Data scientist B'). The objective was to gain a deeper understanding of the inner workings of the algorithms and to explore the perspectives and opinions of the data scientists regarding these topics.

Triangulating materials enhances our understanding of algorithm development and use. Participant observation and internal documents reveal the roles and dynamics of the actors involved (Seaver, 2017). Semi-structured interviews provide detailed insights into the perspectives of the data scientists who build the algorithms. As Seaver (2017) noted, collecting all relevant materials, particularly corporate materials rather than press releases, is crucial for studying algorithms.

Data analysis

The data analysis followed an interpretivist approach (Sandberg, 2005) and consisted of a two-stage process, with a distinct role for each type of material. In the first stage, a thematic analysis was conducted with an abductive approach, iterating between the empirical data and the three ethical issues identified in the literature review. This analysis allowed us to identify the roles and visions of the three main stakeholder groups – data scientists, HR practitioners, and legal experts – regarding their perceptions and representations of the algorithm and their ethical issues.

In the second stage, the focus shifted to these ethical considerations, specifically accountability, fairness, and transparency. These three dimensions were chosen based on our knowledge and observation of the field. The goal was to explore the different stakeholder views on these issues. While ethics were not always explicitly stated by the actors, we discerned relevant discussions from the documents and interviews related to these ethical issues.

The materials address concerns raised by various studies of algorithms, such as the diversity of algorithm conceptualizations (Ananny & Crawford, 2018; Cheng & Hackett, 2021; Lange et al., 2019). They help us understand the ethical considerations surrounding the development and use of algorithms and analyse the interactions among decision-makers, data scientists, legal experts, and users.

Initially, we aimed to compare the two algorithms to highlight their unique features and challenges. However, focusing on ethical considerations in the second stage revealed that both algorithms shared similar ethical concerns. This shift led us to view the algorithms as cumulative cases (Garreau, 2020) rather than as purely comparative (Yin, 1981a, 1981b), emphasizing their common ethical issues.

Descriptions of the two algorithms

Training suggestion algorithm ('training algorithm')

The first algorithm provides personalized suggestions for training courses for employees, who are, therefore, the algorithm's users. This project was conducted in four stages, spanning a period from February 2016 to April 2017. In the initial 3-month stage, the project was framed with the goal of testing the algorithm with voluntary employees and focusing only on short e-learning courses.

The second stage, which also lasted 3 months, involved data collection. An email was sent to 10,000 employees, informing them about the experiment and the data collection. Approximately 1,700 employees out of 10,000 agreed to participate in the experiment. The algorithm development took 5 months. The final 3-month stage consisted of gathering feedback and suggestions from employees. The final algorithm was composed of three sub-algorithms: collaborative filtering, thematic analysis, and a matching algorithm.

During development, data that would have been useful for suggesting training courses were either missing or of poor quality. For example, the company lacked data on individual employee skills. The available data included training data, training history, identification data – such as managerial status, classification, and job field – and monitored social network community data (Figure A1 in the Appendix).

In addition, the lack of data on various topics had to be compensated for by generating new data from existing information. For example, the data scientists felt that the lack of information on the evaluations of the training courses that learners attended could affect the quality of the final recommendations. As a result, they developed a method to extrapolate the evaluations of the training course based on behavioural assumptions. For example, they assumed that someone who spent fewer than 5 min on a course did not like it and that, conversely, logging on several times to take a course showed a

satisfactory appreciation of the course. In the end, this resulted in a satisfaction indicator.

Suggestions were sent to all registered employees in January 2017. To evaluate the relevance of the algorithm, a non-anonymous questionnaire was distributed, which linked the responses to the received suggestions. The survey questioned them, among other things, on the usefulness of the suggestions, the likelihood of employees following them, and their consideration of training without suggestions. Such a survey aims to consider user-centric evaluations in improving algorithmic decisions. Additionally, a report tracked whether employees followed the suggestions, and the change in the average monthly number of training courses attended was reviewed to measure any increase in training participation.

CV pre-screening algorithm ('pre-screening algorithm')

The second algorithm is a prescriptive tool aimed at automating the pre-screening of CVs and suggesting candidates for job interviews. A discussion between the recruitment department and a data scientist led to the identification of the need for the automatic sorting of CVs during external recruitment. Digix receives more than 100,000 CVs each year to be analysed for a range of 6,000–10,000 advertised jobs. Recruiters preselect CVs and interview candidates. The project's goal was to free up recruiters' time for proactive tasks, for example headhunting. Initial studies conducted by data scientists revealed that approximately two thirds of the applications were not aligned with the job criteria. Consequently, the algorithm aims to streamline pre-screening and enhance HR efficiency by automating the exclusion of irrelevant applications (Langley & Simon, 1995).

The project unfolded in three phases. From February 2016 to July 2016, an initial algorithm for a small dataset was developed. This phase focused on handling unstructured textual data. In the second half of 2016, the algorithm was tested and refined with a larger dataset,² employing semantic analysis to create and compare word clouds from CVs and job advertisements. The algorithm scores each CV based on its relevance to the job advertisement in terms of both its presence and its position in both documents. In early 2017, the algorithm was tested alongside human recruiters and discrepancies were analysed to improve its accuracy by using more data. The final deliverable of the algorithm is basic: it is a score, an estimate of the relevance of a candidate to a desired profile. It says "78% match between the person's profile and what we're looking for". (Interview with Data scientist B)

²All job advertisements produced and CVs received over two months, approximately 1,000 job advertisements for 400 different jobs and 10,000 CVs.

Data scientists then manually improved the algorithm by discussing CVs with recruiters, understanding their preferences, and aligning the algorithm's results with human judgements. This adjustment improved the relevance score, ranging from 65 to 80%.

The two algorithms differ in their purpose and use. The training suggestion algorithm gives rise to an entirely new process, whereas the pre-screening algorithm automates an existing process, impacting applicants who have no control over its use. Both algorithms involve data scientists, the HR department, and legal experts in their development. Data scientists programme the algorithms, HR ensures that they meet their needs, and legal experts ensure compliance. The two algorithms differ in terms of users. In the training algorithm, the end users are employees, whereas they are recruiters in the pre-screening algorithm. In the training algorithm, the input data come from the users, including their feedback, whose data are used to train the algorithms to make recommendations for other employees. The latter are not involved in the design of the algorithm.

In the pre-screening algorithm, recruiters make decisions on the recruitment process based on the suggestions of the algorithm, but the data processed are based on the CVs of applicants. While recruiters are both developers – in the sense that they are involved in the design of the algorithm – and users of the algorithm, applicants have no information on how the algorithm uses their data.

Ethical issues in the development and use of algorithms

The main findings are structured around the stages of development and use of the algorithms: data selection, construction of the algorithm, assessment of the algorithm's accuracy, and use of the results produced by the algorithm. For each stage, we focus on the ethical considerations that emerged among the various stakeholders involved and, on the decisions, made following these considerations.

Data selection

The data selection stage raises ethical considerations, notably surrounding fairness. Data scientists advocate that the more data there are and the better the quality of the data, the fairer the results produced by the algorithm – fairness being understood as increased objectivity and reduced bias. To achieve this, they prioritize what they term 'data quality', which refers to two criteria. The first is a high completeness rate, referring to the percentage of individuals for whom the data are provided. The second is a high degree of differentiation in the data. For example, if all individuals have the same skills, such as

proficiency in French, then the data lack the diversity needed for personalized recommendations.

On data retrieval, it was slightly long, but in the end, people were quite cooperative. Afterwards, in the construction of the algorithm, we had a considerable difficulty: we had the data, but we did not qualify the data enough beforehand. Two difficulties arose: data from the internal social network that was very hard to exploit and a training catalogue that was constantly changing. With very few notes, few uses, an ever-changing catalogue . . . (Interview with Data scientist A, conducted in 2017)

Data scientists also reported a lack of specific key data deemed essential for training algorithm effectiveness, such as employees' evaluation of training courses. In effect, collaborative filtering algorithms, in particular, yield better results when they incorporate not only data on whether users have accessed specific content but also their feedback on the quality of that content.

The data scientists believe it would be necessary to have information on whether employees appreciated the training they received in the past (in addition to their training history). For e-learning training, they propose the construction of a variable that approximates satisfaction with training based on the number of connections and the completion of the training. (Observation notes – internal meetings held in 2016 within the project team)

Overall, data scientists believe that as much data as possible should be retrieved upstream, as it is challenging to know in advance which variable will improve the quality of the algorithm and make it fairer, that is more objective and less biased.

This perspective contrasts with the view of legal experts, who argue that it is essential to predefine and justify the data used during the construction of the algorithm from the outset. They believe that avoiding discrimination and guaranteeing transparency requires careful consideration of which data will be used even before developing the algorithm. For example, a document that presents the objective of the project and the data to be used should be produced before the project is launched in order to obtain its approval.

Primary objective: experiment regarding the automated emailing of personalized suggestions for online training to employees.

[...]

Categories of data processed: internal personnel number, data from the internal social network, training history and administrative data (e.g., profession, being a manager). (Excerpt from the 'data protection correspondent' file, an internal document of the company produced by the legal team in 2016)

In the dialogue between legal experts and data scientists, disagreements also arose over the inclusion of certain specific variables. Legal experts emphasized the need to avoid discriminatory practices, and therefore advocated discarding variables

such as gender because of legal prohibitions against discrimination based on such characteristics. For legal experts, an essential aspect of ethics is embodied in the protection of personal data and in preventing direct discrimination through the exclusion of certain variables as inputs for algorithm decision-making. In contrast, data scientists argue that including these variables could help control or correct possible preexisting biases in the data and could thus increase fairness.

These divergent viewpoints illustrate the contrasting priorities in algorithmic modelling among different stakeholder groups. The final decision at this stage, made by HR practitioners and legal experts, was to exclude gender-based data and limit data collection to the strict minimum to comply with French regulations. Legal compliance with non-discrimination laws is a major concern for HR practitioners, who seek to mitigate the risk of lawsuits or sanctions. As a result, they tend to rely on legal experts for guidance and legal safeguards. However, restricting the data collected has direct implications for the algorithm's construction.

Construction of the algorithm

In the algorithm development phase, data scientists seem to be primarily responsible for constructing the algorithms and selecting computation methods. However, the available data, selected by legal experts and HR practitioners, as underlined earlier in the text, strongly influence the choice of algorithms. For example, the training algorithm ultimately comprises three sub-algorithms. The first sub-algorithm, collaborative filtering, performs well only for employees who have a sufficiently long training history for comparison with one another. The second sub-algorithm, thematic filtering, complements the first algorithm but is effective only for those with at least one training course in their history. For employees without a training history, a third algorithm is employed to match them with employees who have such histories based on their HR information. This approach reflects data scientists' focus on selecting algorithms suitable for the available data. For example, limited data availability prevents training course suggestions based on qualifications or employees' wishes.

However, HR practitioners require data scientists to explain their choices and to be able to describe the algorithm and how it works; therefore, underlining the need for transparency. In particular, HR practitioners require that the algorithm's inner workings be understandable to employees, leading to the choice of relatively simple algorithms that can be easily explained to HR practitioners and employees. 'Following feedback from the internal social network community on the lack of relevance of some training suggestions, an explanatory document presenting the three sub-algorithms used is sent to all participants' (Summary of internal email exchanges, beginning of 2017).

Similarly, for the pre-screening algorithm, the choice was influenced by the characteristics of the available data as well as the desire to be able to explain the algorithms to recruiters and employees. As the data scientists only had data on CVs and job advertisements, they selected an algorithm that builds word clouds and measures word frequencies for CVs and job advertisements, methods that can be comprehensible to both recruiters and employees.

These elements show that the need to be able to communicate about the algorithms to non-data-experts, which refers to the transparency issue, is a key factor in the development of the algorithm.

Assessment of the algorithm's accuracy

In general, human judgement serves as the benchmark for determining whether the algorithm's decisions are 'correct'. However, this criterion is not without contention, particularly between HR practitioners and data scientists, who often disagree on whether human judgement is an adequate standard. For example, in the pre-screening algorithm, a comparison was made between the first 20 CVs selected by the algorithm and those selected by the recruiter for each job advertisement. The percentage of agreement between the two sets of selections was then measured to assess algorithmic performance.

We got an 84% match. I did two exercises in parallel: I took a job advertisement and the applications, and the recruiters did the scoring (the ranking of the CVs) by hand. They said: out of all the applications received, here are the top twenty. We also compared that with the algorithm. There was general agreement; when the applications did not match at all, it was the same for the machine and the human. However, there were a few CVs that the machine discarded and the human retained, and vice versa. However, it still meant that out of ten CVs, there were one or two that the machine did not retain, and the human retained, which is socially acceptable ... We could not really explain it, however, because there is a lot of intuition involved. (Interview with Data scientist B, conducted in 2017)

The data scientist who was interviewed highlights that HR practitioners and recruiters generally accept the use of human judgement as a benchmark for evaluating algorithms, highlighting the social acceptance of human arbitrariness. This acceptance stems from the fact that humans can be held accountable for their decisions. However, a similar level of arbitrariness is not tolerated in algorithmic decisions. This stance sidesteps the considerations regarding the biases inherent in human decisions or the potential for algorithms to make 'better' decisions, in other words, that are less biased, more objective and more innovative than humans (Jago, 2019). Emphasizing this point, the other data scientist underlines the fact that the ultimate

achievement of a recommendation engine lies in generating insights that humans would not have thought of.

Everyone is free to do what they want, but if someone recommends obvious training courses to you, I don't think there's much point. It saves you time in the training catalogue, but you don't need that ... The point of recommendation engines is to correct mistakes or get back things you didn't have. Recommendations have worked in entertainment to create usage, and they're trying to create diversity. When you've bought five Beatles records, there's no point in Amazon recommending you a sixth. It's all about predictive power. What's great is when you can predict things you do not already know exist. (Interview with Data scientist A, conducted in 2017)

Finally, the quality of both algorithms is evaluated through the collection of user feedback gathered from employees or recruiters. For the pre-screening algorithm, feedback is collected from a relatively small user base, allowing for more direct and frequent inputs to refine the system. Conversely, for the training algorithm, feedback is obtained through a questionnaire due to the large number of users. Although employees appreciate the ability to receive training suggestions, their opinions on the quality of these suggestions are lukewarm (Figures A2 and A3 in the Appendix). Data scientists attribute the lukewarm opinion to strong technical constraints associated with the lack of data, therefore refusing to be held accountable for this. For example, in the internal social network community, employees express dissatisfaction with the predominance of English-language training courses, despite their limited proficiency in English. Data scientists, in turn, attribute this issue to the data, as the selected catalogue from which the suggestions were generated predominantly consists of English-language content.

Feedback on the use of the results produced by the algorithm

The initial feedback on the usage test of the results produced by the algorithm raises ethical issues related to human agency and autonomy. Concerning the training suggestion engine, the training department (HR practitioners) decided that each employee would receive their own training suggestions without informing the manager, thereby giving employees the autonomy to follow or disregard the suggestions as they saw fit.

Proposal for the pilot:

Only open-access training, without managerial authorization (short e-learning courses), is recommended (Skillsoft catalogue). This makes it possible to recommend training courses that are not specific to business areas (cross-functional training) and to give employees greater autonomy in their learning. (Excerpt from a

document presenting the training algorithm, an internal document of the company produced in 2016)

For the pre-screening algorithm, the recruitment department decided that the ranked results generated by the CV preselection algorithm would be sent only to recruiters, allowing them autonomy to either follow or disregard the recommendations.

In both algorithms, the end users who receive the results of the algorithms have the agency to decide whether to act on these suggestions, making them accountable for the final decision. The question of accountability is particularly complex in both cases and raises ethical considerations. When data scientists' accountability is questioned, for example regarding the quality of the results produced by the algorithm, they shift accountability to the final users, who make the final decisions. They argue that it is ultimately the human users who should be held accountable, as the algorithms merely provide recommendations, leaving the final choice to human judgement. This perspective aligns well with HR managers, who prioritize a social contract with employees. On the other hand, legal experts view accountability as a critical concern and would rather compromise algorithm effectiveness than risk using variables that could expose the organization to claims of unfair recommendations, as underlined above. The results illustrate the complexity of defining accountability among the various actors involved. Each group has different perspectives on accountability, offering insights that extend the existing literature on algorithmic accountability, notably Martin's (2019b) article. The various actors usually prefer to delegate accountability rather than accept it themselves, contrary to what has been assumed in previous research (Martin, 2019b). In the Digix case, as the algorithm is developed internally, the organization should be held accountable. While certain ethical considerations may not be codifiable in law, there should be ways to extend the current legal frameworks to recognize the responsibilities of designers and owners of algorithms.

In the case of the training algorithm, the user who receives the results is the same as the individual whose data are being processed (employee). For the pre-screening algorithm, the user (recruiter) who receives the results is distinct from the data subject (applicant). This means that for the training algorithm, the algorithm was explained to employees whose data were processed. However, for the pre-screening algorithm, while the functionality was explained to recruiters, applicants whose data were processed were not informed about this use. Interestingly, the various stakeholders were aware of this lack of transparency but did not mention it as a fundamental ethical issue, suggesting that the focus on transparency was primarily directed towards current employees rather than applicants. The question of social acceptability concerned the recruiters. As far as the applicants were concerned, the only issue was to

ensure non-discrimination in the event of an appeal' (Interview with Data scientist B, conducted in 2017).

The results underscore divergent definitions of transparency. Stakeholders, and notably HR practitioners, prioritize algorithmic transparency by emphasizing explanations that are understandable to both managers and employees, even if doing so entails sacrificing some performance. This contrasts with findings from other studies (Langer & König, 2023; Turner, 2009), where stakeholders focused more on performance. Moreover, HR practitioners emphasize transparency specifically for the benefit of employees within the organization. However, it remains unclear whether transparency in the functioning of algorithms enhances the decision-making of HR managers or users.

Discussion

Drawing on various qualitative materials, we identified key ethical considerations among the actors involved in the development and use of algorithms within a major digital company. Our research, which leverages the unique insider perspective of one of the co-authors, provides an in-depth analysis of the development and use of two algorithms, offering access to rarely examined internal dynamics. Our study shows that the discussion, creation, and adjustment of algorithms generate ethical considerations, particularly regarding accountability, fairness, and transparency. Our case study illustrates the ethical considerations raised by various stakeholders, such as data scientists, HR managers, legal experts, and different end users, which include recruiters and employees. Our results reveal that their goals and moral perspectives often diverge during algorithmic development.

Our results make three key academic contributions. First, they provide a detailed empirical exploration of the internal dynamics and ethical considerations that occur during algorithm construction – a process largely overlooked in previous research. Second, they demonstrate that the designers of algorithms are a heterogeneous group whose differing perspectives significantly shape outcomes related to accountability, fairness, and transparency. Third, they bridge the gap between theory and practice by linking these internal considerations to pressing managerial issues.

Unlike many studies on algorithms (e.g., Kellogg et al., 2020), which highlight a stark divide between designers embracing the algorithm and users actively resisting it, our findings suggest a more complex and nuanced relationship in practice (Raisch & Krakowski, 2021). Even in the construction of the algorithm (Glaser et al., 2021), we observe that their designers (data scientists, legal experts, and HR practitioners) are not a homogenous group. Instead, they offer different perspectives, particularly when selecting the variables to be considered in the algorithms. Our research shows that algorithms “function

as prisms that can reveal existing priorities within groups” (Christin, 2020, p. 906).

Our findings also shed light on how different actors leverage ethical issues to assert their points of view. For legal experts, accountability focuses primarily on controlling the inputs used in algorithmic decisions rather than on performance or transparency for end users. As a result, they treat fairness as a mechanism to mitigate accountability issues. On the other hand, data scientists interpret fairness in terms of the quality and accuracy of algorithmic classifications. For them, transparency serves to elucidate the limited relevance of certain algorithmic decisions, which helps shift accountability away from themselves in cases of poor recommendations and transfer it to end users. In contrast, HR managers prioritize transparency because it allows end users to critically evaluate and challenge algorithmic outputs through established procedures (Martin, 2019a). However, while transparency can empower users, it also raises concerns regarding accountability and fairness in final outcomes. For example, transparency increases the potential for ‘gaming’ the algorithm, resulting in unequal treatment of information among users (Christin, 2017), and may require simpler algorithms that sacrifice sophistication and fairness. In particular, the accountability of HR managers in constraining algorithmic choices for the sake of transparency remains unchallenged.

Empirically, our study contributes by facilitating a systematic comparison of two algorithms serving different end users within the same organization. This comparative approach reveals that transparency is typically more limited for external users than for internal users, but it also provides a robust analytical framework to think across cases.

Our study also builds a bridge between academia and practice by emphasizing three managerial issues. First, the empirical results underscore the fact that human choice remains the dominant criterion, with human accountability being paramount. Nevertheless, our case study shows that the use of algorithms must be accompanied by parallel investment in human expertise to enhance collaborative intelligence and mitigate ethical risks (Gregory et al., 2021; Leicht-Deobald et al., 2019) regarding these technologies, in order to enhance collaborative intelligence (Wilson & Daugherty, 2018) and minimize their potential risks from an ethical perspective. Companies can put in place co-creation processes and a commitment to developing employees’ ‘fusion skills’ for working at the human-machine interface (Wilson & Daugherty, 2018).

Second, while our results highlight the different visions of the world of the algorithms’ designers, we believe that our work also underscores the psychological distance between the designers of the algorithms (Donaldson & Neesham, 2020). Psychological distance refers to the mental distance between individuals: individuals care more about those to whom they feel close in terms of time or personal characteristics

(Donaldson & Neesham, 2020). Typically, HR managers are psychologically closer to the end users of algorithms than data scientists or legal experts. Psychological distance affects judgements about fairness – the greater the psychological distance, the higher the risk of unfair judgements – so it is important to account for this in the design of algorithms.

Finally, transparency also raises ethical considerations around the need to be comprehensive and the effectiveness of algorithms. However, the dynamics of algorithms were not considered. As algorithms are self-learning systems and evolve with the characteristics of the users, transparency must also signal that the system evolves over time. This observation invites future research to explore how dynamic transparency can be integrated into algorithm design to ensure both fairness and adaptability.

Conclusion

This paper aimed at opening the 'black box' of two algorithms developed by a major company in France. Our results emphasized the decision-making process in the creation of algorithms and the key ethical issues – accountability, fairness, and transparency – involved in the construction and operationalization of algorithms.

Our study has, of course, several limitations. The broad scope of materials studied using alternative methodologies may be considered one such limitation. However, providing the best overview of both algorithms studied was necessary (Christin, 2020; Seaver, 2017). We also conducted a limited number of interviews, with a focus on data scientists, whereas it would have been interesting to conduct more interviews with practitioners and users. The main limitation of the study is that it focuses on a single company. While studying a single company allows the comparison of the two algorithms without having to consider the effects of the organizational context, it prevents us from drawing general conclusions based on our results. However, both types of algorithm studied in this paper are becoming increasingly common within companies. Our case study may therefore be of interest to practitioners and researchers studying different organizational or cultural contexts.

Further studies could focus on comparing the construction of algorithms across different companies or industries. This would enhance our understanding of the different ethical issues at stake in various industries and the considerations of the various actors involved in the construction of algorithms.

Acknowledgements

The authors thank associate editor Wafa Ben Khaled and two anonymous reviewers for their comments throughout the editorial process. They also thank Olivier Cristofini, Jean-Loup

Richet, and Arun Rai for their comments on the previous versions of the paper. All errors remain that of the authors.

References

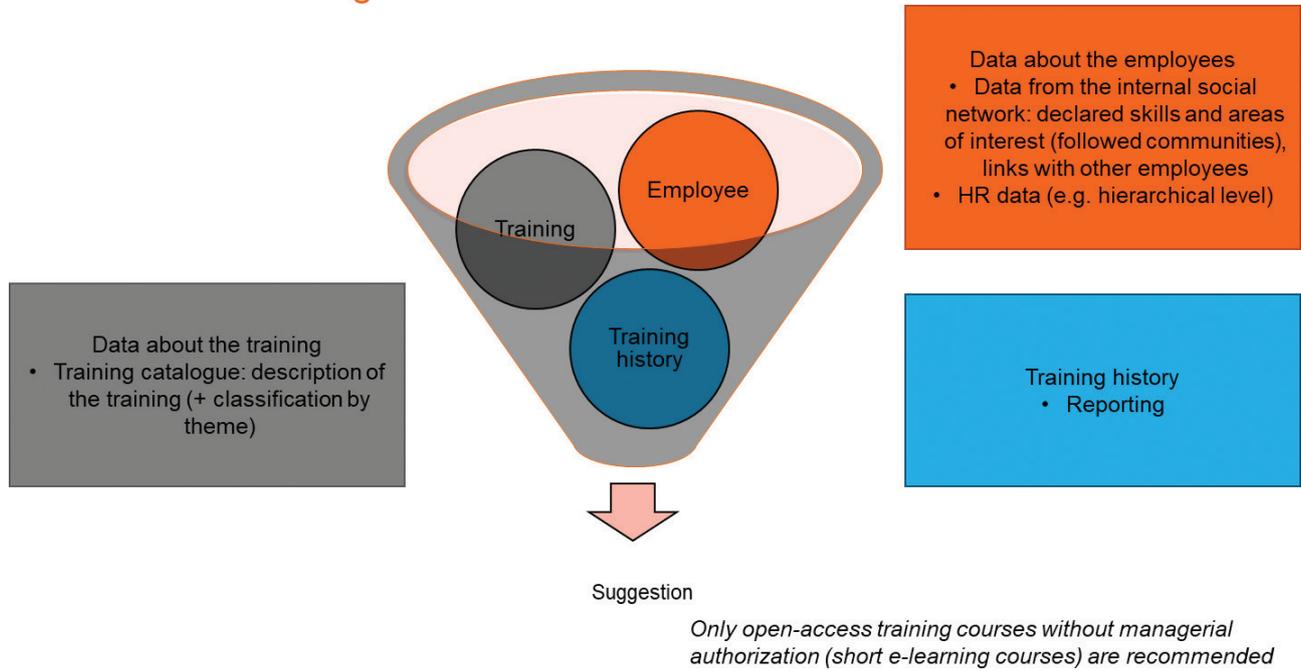
- Ananny, M. (2016). Toward an ethics of algorithms: Convening, observation, probability, and timeliness. *Science, Technology, & Human Values*, 41(1), 93–117. <https://doi.org/10.1177/0162243915606523>
- Ananny, M. & Crawford, K. (2018). Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society*, 20(3), 973–989. <https://doi.org/10.1177/1461444816676645>
- Binns, R. (2018). Algorithmic accountability and public reason. *Philosophy & Technology*, 31(4), 543–556. <https://doi.org/10.1007/s13347-017-0263-5>
- Bucher, E. L., Schou, P. K., & Waldkirch, M. (2021). Pacifying the algorithm—Anticipatory compliance in the face of algorithmic management in the gig economy. *Organization*, 28(1), 44–67. <https://doi.org/10.1177/1350508420961531>
- Burrell, J. (2016). How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1), 1–12. <https://doi.org/10.1177/2053951715622512>
- Cheng, M. M. & Hackett, R. D. (2021). A critical review of algorithms in HRM: Definition, theory, and practice. *Human Resource Management Review*, 31(1), 100698. <https://doi.org/10.1016/j.hrmr.2019.100698>
- Christin, A. (2017). Algorithms in practice: Comparing web journalism and criminal justice. *Big Data & Society*, 4(2), 205395171771885. <https://doi.org/10.1177/2053951717718855>
- Christin, A. (2020). The ethnographer and the algorithm: Beyond the black box. *Theory and Society*, 49, 897–918. <https://doi.org/10.1007/s1186-020-09411-3>
- Coron, C. (2022). Quantifying human resource management: A literature review. *Personnel Review*, 51(4), 1386–1409. <https://doi.org/10.1108/PR-05-2020-0322>
- Donaldson, T. J. & Neesham, C. (2020). The problem of value alignment in business decision making: Humans vs. artificial intelligence. *Academy of Management Proceedings*, 2020(1), 14706. <https://doi.org/10.5465/AMBPP.2020.14706abstract>
- Faraj, S., Pachidi, S. & Sayegh, K. (2018). Working and organizing in the age of the learning algorithm. *Information and Organization*, 28(1), 62–70. <https://doi.org/10.1016/j.infoandorg.2018.02.005>
- Garreau, L. (2020). Petit précis méthodologique. *Le Libellio d'AEGIS*, 16(2), 51–64.
- Glaser, V. L., Pollock, N. & D'Adderio, L. (2021). The biography of an algorithm: Performing algorithmic technologies in organizations. *Organization Theory*, 2(2), 1–27. <https://doi.org/10.1177/26317877211004609>
- Gregory, R. W., Henfridsson, O., Kaganer, E. & Kyriakou, H. (2021). The role of artificial intelligence and data network effects for creating user value. *Academy of Management Review*, 46(3), 534–551. <https://doi.org/10.5465/amr.2019.0178>
- Hodder, M. (2009, April 14). Why Amazon didn't just have a glitch. *TechCrunch*. Retrieved from <http://techcrunch.com/2009/04/14/guest-post-why-amazon-didnt-just-have-a-glitch/>
- Hursthouse, R. (1999). *On virtue ethics*. Oxford University Press. <https://doi.org/10.1093/0199247994.001.0001>
- Introna, L. D. (2016). Algorithms, Governance, and Governmentality: On Governing Academic Writing. *Science, Technology, & Human Values*, 41(1), 17–49. <https://doi.org/10.1177/0162243915587360>

- Introna, L. D. (2016). Algorithms, governance, and governmentality: On governing academic writing. *Science, Technology, & Human Values*, 41(1), 17–49. <https://doi.org/10.1177/0162243915587360>
- Irani, L. (2015). Difference and dependence among digital workers: The case of Amazon Mechanical Turk. *South Atlantic Quarterly*, 114(1), 225–234. <https://doi.org/10.1215/00382876-2831665>
- Jago, A. S. (2019). Algorithms and authenticity. *Academy of Management Discoveries*, 5(1), 38–56. <https://doi.org/10.5465/amd.2017.0002>
- Kellogg, K. C., Valentine, M. A. & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366–410. <https://doi.org/10.5465/annals.2018.0174>
- Kim, M. P., Reingold, O. & Rothblum, G. N. (2018). Fairness through computationally-bounded awareness. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, et al. (Eds.), *Advances in neural information processing systems* (Vol. 31, pp. 4842–4852). Curran Associates. Retrieved from https://papers.nips.cc/paper_files/paper/2018/file/c8dfce5cc68249206e4690fc4737a8d-Paper.pdf
- Kraemer, F., Van Overveld, K. & Peterson, M. (2011). Is there an ethics of algorithms? *Ethics and Information Technology*, 13(3), 251–260. <https://doi.org/10.1007/s10676-010-9233-7>
- Lange, A.-C., Lenglet, M. & Seyfert, R. (2019). On studying algorithms ethnographically: Making sense of objects of ignorance. *Organization*, 26(4), 598–617. <https://doi.org/10.1177/135050841880823>
- Langer, M. & König, C. J. (2023). Introducing a multi-stakeholder perspective on opacity, transparency and strategies to reduce opacity in algorithm-based human resource management. *Human Resource Management Review*, 33(1), 100881. <https://doi.org/10.1016/j.hrmr.2021.100881>
- Langley, P. & Simon, H. A. (1995). Applications of machine learning and rule induction. In J. Cohen (Ed.), *Communications of the ACM* (Vol. 38, no. 11, pp. 54–64). Association for Computing Machinery. <https://doi.org/10.1145/219717.219768>
- Leicht-Deobald, U., Busch, T., Schank, C., Weibel, A. et al. (2019). The challenges of algorithm-based HR decision-making for personal integrity. *Journal of Business Ethics*, 160(2), 377–392. <https://doi.org/10.1007/s10551-019-04204-w>
- Lowrie, I. (2017). Algorithmic rationality: Epistemology and efficiency in the data sciences. *Big Data & Society*, 4(1), 205395171770092. <https://doi.org/10.1177/2053951717700925>
- Marcus, G. E. (1995). Ethnography in/of the world system: The emergence of multi-sited ethnography. *Annual Review of Anthropology*, 24, 95–117. <https://doi.org/10.1146/annurev.an.24.100195.000523>
- Martin, K. (2019a). Designing ethical algorithms. *MIS Quarterly Executive*, 18(2), 129–142. <https://doi.org/10.17705/2msqe.00012>
- Martin, K. (2019b). Ethical implications and accountability of algorithms. *Journal of Business Ethics*, 160(4), 835–850. <https://doi.org/10.1007/s10551-018-3921-3>
- Meijerink, J. & Bondarouk, T. (2023). The duality of algorithmic management: Toward a research agenda on HRM algorithms, autonomy and value creation. *Human Resource Management Review*, 33(1), 100876. <https://doi.org/10.1016/j.hrmr.2021.100876>
- Merrill III, D. G. (2011). *Allocation-oriented algorithm design with application to gpu computing*. University of Virginia Press.
- Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S. et al. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 2053951716679679. <https://doi.org/10.1177/2053951716679679>
- Murray, A., Rhymer, J. & Sirmon, D. G. (2021). Humans and technology: Forms of conjoined agency in organizations. *Academy of Management Review*, 46(3), 552–571. <https://doi.org/10.5465/amr.2019.0186>
- Neff, G. & Stark, D. C. (2003). Permanently beta: Responsive organization in the Internet era. In P. N. Howard & S. Jones (Eds.), *Society online: The internet in context* (pp. 173–188). Sage.
- Neyland, D. (2015). On organizing algorithms. *Theory, Culture & Society*, 32(1), 119–132. <https://doi.org/10.1177/0263276414530477>
- Orlikowski, W. & Scott, S. V. (2015). The algorithm and the crowd: Considering the materiality of service innovation. *MIS Quarterly*, 39(1), 201–216. <https://doi.org/10.25300/MISQ/2015/39.1.09>
- Orlikowski, W. & Scott, S. V. (2016). Digital work: A research agenda. In B. Czarniawska (Ed.), *A research agenda for management and organization studies* (pp. 88–95). Edward Elgar. <https://doi.org/10.4337/9781784717025.00014>
- Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.
- Raisch, S. & Krakowski, S. (2021). Artificial intelligence and management: The automation-augmentation paradox. *Academy of Management Review*, 46(1), 192–210. <https://doi.org/10.5465/amr.2018.0072>
- Roulet, T. J., Gill, M. J., Stenger, S. & Gill, D. J. (2017). Reconsidering the value of covert research: The role of ambiguous consent in participant observation. *Organizational Research Methods*, 20(3), 487–517. <https://doi.org/10.1177/1094428117698745>
- Sandberg, J. (2005). How do we justify knowledge produced within interpretive approaches? *Organizational Research Methods*, 8(1), 41–68. <https://doi.org/10.1177/1094428104272000>
- Schouten, J. W. & McAlexander, J. H. (1995). Subcultures of consumption: An ethnography of the new bikers. *Journal of Consumer Research*, 22(1), 43–61. <https://doi.org/10.1086/209434>
- Seaver, N. (2017). Algorithms as culture: Some tactics for the ethnography of algorithmic systems. *Big Data & Society*, 4(2), 205395171773810. <https://doi.org/10.1177/2053951717738104>
- Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S. & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 59–68), January 29–31, 2019, Atlanta, GA: ACM.
- Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S. et al. (2019). Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 59–68). <https://doi.org/10.1145/3287560.3287598>
- Starke, C., Baleis, J., Keller, B. & Marcinkowski, F. (2022). Fairness perceptions of algorithmic decision-making: A systematic review of the empirical literature. *Big Data & Society*, 9(2), 20539517221115189. <https://doi.org/10.1177/20539517221115189>
- Turner, F. (2009). Burning man at Google: A cultural infrastructure for new media production. *New Media & Society*, 11(1–2), 73–94. <https://doi.org/10.1177/1461444808099575>
- Von Krogh, G. (2018). Artificial intelligence in organizations: New opportunities for phenomenon-based theorizing. *Academy of Management Discoveries*, 4(4), 404–409. <https://doi.org/10.5465/amd.2018.0084>
- Wilson, H. J. & Daugherty, P. R. (2018, July–August). Collaborative intelligence: Humans and AI are joining forces. *Harvard Business Review*. Retrieved from <https://hbr.org/2018/07/collaborative-intelligence-humans-and-ai-are-joining-forces>
- Yin, R. K. (1981a). The case study as a serious research strategy. *Science Communication*, 3(1), 97–114. <https://doi.org/10.1177/107554708100300106>
- Yin, R. K. (1981b). The case study crisis: Some answers. *Administrative Science Quarterly*, 26(1), 58–65. <https://doi.org/10.2307/2392599>

Appendices Lukewarm opinion on the

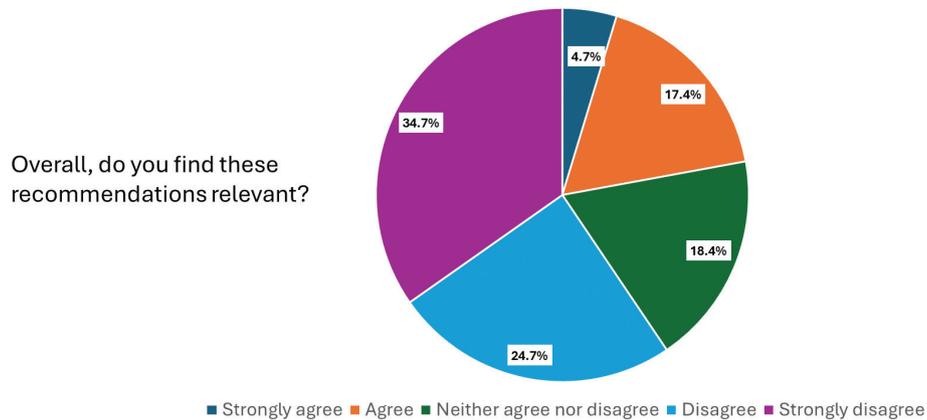
relevance and usefulness of suggestions

Data used in the algorithm



Source: Internal document that was produced and presented during an internal meeting in 2016; own translation.

Figure A1. Presentation of the data.



Source: Data from the 2017 questionnaire; own translation.

Figure A2. Chart appearing in the internal report on the training algorithm.

Analysis of the questionnaire

1. Lukewarm opinion on the relevance and usefulness of suggestions
2. These suggestions, however, seem to open new horizons for respondents, and just over a third of respondents think they will follow the suggestions
3. They find the user experience intuitive
4. Finally, receiving training suggestions is an interesting feature for a very large majority of respondents

Source: Data from the 2017 questionnaire; translated by the authors.

Figure A3. Excerpt from the internal report on the training algorithm.