



A Framework for Exploring the Consequences of AI-Mediated Enterprise Knowledge Access and Identifying Risks to Workers

Anna Gausen
Imperial College London
London, UK
anna.gausen16@imperial.ac.uk

Bhaskar Mitra
Microsoft Research
Montreal, Canada

Siân Lindley
Microsoft Research
Cambridge, UK

ABSTRACT

Organisations generate vast amounts of information, which has resulted in a long-term research effort into knowledge access systems for enterprise settings. Recent developments in artificial intelligence, in relation to large language models, are poised to have significant impact on knowledge access. This has the potential to shape the workplace and knowledge in new and unanticipated ways. Many risks can arise from the deployment of these types of AI systems, due to interactions between the technical system and organisational power dynamics.

This paper presents the Consequences-Mechanisms-Risks framework to identify risks to workers from AI-mediated enterprise knowledge access systems. We have drawn on wide-ranging literature detailing risks to workers, and categorised risks as being to worker value, power, and wellbeing. The contribution of our framework is to additionally consider (i) the consequences of these systems that are of moral import: commodification, appropriation, concentration of power, and marginalisation, and (ii) the mechanisms, which represent how these consequences may take effect in the system. The mechanisms are a means of contextualising risk within specific system processes, which is critical for mitigation. This framework is aimed at helping practitioners involved in the design and deployment of AI-mediated knowledge access systems to consider the risks introduced to workers, identify the precise system mechanisms that introduce those risks, and begin to approach mitigation. Future work could apply this framework to other technological systems to promote the protection of workers and other groups.

CCS CONCEPTS

• Information systems → Enterprise applications; • Computing methodologies → Artificial intelligence.

KEYWORDS

enterprise knowledge access systems, risks, large language models, workers

ACM Reference Format:

Anna Gausen, Bhaskar Mitra, and Siân Lindley. 2024. A Framework for Exploring the Consequences of AI-Mediated Enterprise Knowledge Access



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

FACCT '24, June 03–06, 2024, Rio de Janeiro, Brazil

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0450-5/24/06

<https://doi.org/10.1145/3630106.3658900>

and Identifying Risks to Workers. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FACCT '24)*, June 03–06, 2024, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3630106.3658900>

1 INTRODUCTION

People are increasingly interacting with, and being affected by, the deployment of AI systems in the workplace. How AI systems are impacting workers is a pressing matter for system designers, policy-makers, and workers themselves [79]. This paper will focus on how AI-based systems for enterprise knowledge access (EKA) can specifically impact workers. We propose a framework for identifying and understanding mechanisms that lead to systemic consequences of moral import with the goal of safeguarding worker power, value, and wellbeing.

Organisations generate huge amounts of information that raise challenges associated with the maintenance, dissemination, and discovery of organisational knowledge [71]. This has led to a long-term research effort into how knowledge access systems can be developed for enterprise settings. Recent examples of these systems can automatically extract knowledge that an organisation has produced and store it in a knowledge base, from where it can be surfaced to end-users [116] [76]. Recent developments in artificial intelligence (AI), notably large language models (LLMs) like OpenAI's Generative Pre-trained Transformer (GPT4) [85], present a shift in what is possible in this domain. These models are trained on exceptionally large datasets and can be applied to a broad set of downstream tasks [106]. Their capabilities could enable more extensive mining, knowledge synthesis, and natural language interaction in relation to knowledge. They could make EKA systems an integral part of people's work, by surfacing knowledge in new ways, including through interactions that are implicit and proactive [65] [99].

At this moment, where AI-mediated EKA is enabled in commercially available enterprise technologies, and where this may be expedited and transformed by LLMs, there is a need to consider implications for workers. For the benefits of EKA to be fully realised, worker adoption, organisational responsibility, and high-quality knowledge sharing are all necessary. All three require the system to be designed in a way that minimises negative outcomes for workers. This necessitates looking beyond concerns relating to privacy and surveillance [12] [2], and additionally considering the many other risks that can arise from the deployment of these types of AI systems due to the "co-productive" [60] [100] interactions between the technical system and organisational power dynamics [9]. The research and engineering community is missing a structured and actionable framework for critically examining the consequences of

AI-mediated EKA systems in use and identifying potential risks to workers from these systems.

In this paper, we propose a framework to identify the risks to workers associated with deploying AI-mediated EKA systems. We call this the *Consequences-Mechanisms-Risks Framework*. The framework is aimed at supporting those involved in the design and/or deployment of such systems to identify the risks they introduce, the specific system mechanisms that introduce those risks, and the actionable levers to reduce those risks. Many existing papers identify risks. The contribution of our framework is to also consider consequences, which are systemic changes affected by the system deployment which may correspond to specific risks, and the mechanisms, which cause these consequences and may contribute to certain risks. The mechanisms also help contextualise risk within specific system processes. As the technical system cannot be separated from organisational socio-cultural dynamics [13] [37], we consider potential risks to be sociotechnical and system-level rather than by adopting a model-centric approach [114].

This work will have the following research contributions:

- We develop the Consequences-Mechanisms-Risks Framework to explore the consequences of moral import of a system, the system mechanisms that introduce risk, and the risk manifestations.
- We apply the Consequences-Mechanisms-Risks Framework to AI-mediated EKA systems and identify a set of consequences, mechanisms, and risks to workers.
- We highlight certain considerations to help practitioners, involved in the design and deployment of such systems, to reduce risk to workers.

2 LITERATURE SURVEY

2.1 Knowledge and the Enterprise

The study of knowledge in the enterprise has a long history in the fields of human-computer interaction (HCI), computer-supported cooperative work, organisational studies, and more [73]. Organisational knowledge entails both knowledge that is generated through work (i.e., the outputs of “knowledge work”) as well as the “know how” that enables people to get their work done, be this by following workplace policies or by adhering to “unofficial” codes of behaviour [89], and knowing how to work with and draw on the expertise of others [93]. Scholars have differentiated between tacit, implicit, and explicit knowledge, which differ in terms of whether knowledge has been, or can be, articulated, as well as between declarative and procedural knowledge, which can also be framed as the difference between “know about” and “know how” [82]. Prior research has highlighted that knowledge development and dissemination are inherently social processes [82] [84] [83] that are dependent on shared norms, understanding, and connections [5].

Our paper focuses on AI-mediated EKA specifically. EKA systems include knowledge management systems [98] [116], enterprise search systems [35], and enterprise chat interfaces powered by conversational agents. Examples include Amazon’s Kendra, which uses generative AI for conversational agents that sit on top of enterprise content [15], and Microsoft’s VIVA Topics, which uses AI to automatically compile and surface topic-based information to workers [76]. Research has explored the potential for risks to arise from the

deployment of knowledge systems in enterprise settings. Examples include a discussion of social and ethical considerations relevant to expert identification and recommendation via workplace technologies [71], and considerations for the responsible design of enterprise knowledge bases produced through implicit interactions between organisation members and technologies [73]. AI-mediation can both introduce new challenges, such as system opacity which may make recourse harder, and worsen existing challenges, such as reproducing the historic biases within an organisation. AI-mediated EKA systems have a unique set of challenges, compared to other EKA systems, due to the workplace power asymmetries they are deployed under. This paper advances prior work through the development of a formalised framework to address risk to workers from AI-mediated EKA.

2.2 Risk to Workers

This section will survey literature on risks and harms to workers, with a particular focus on technological systems. Our perspective on risks to workers has been strengthened by wide ranging epistemologies including feminism [58] [110] [62], post-colonial theory [38] [107] [77], labour rights [90] [80], and design justice [37]. We identified three high-level risk areas: (1) reduced worker value, (2) reduced worker power, and (3) reduced worker wellbeing. These risk areas are discussed in detail below.

The first identified risk area in the literature is *reduced worker value*, both perceived and economic [57]. Automation can shift economic benefit from workers to the organisation depending on who claims the surplus from efficiency gains [30]. This is not an inevitable outcome of automation, it could be designed in ways that shares the benefit with workers. Systems can also lead to automation bias, where there is over-reliance on a system. Systems that act as mediators between people “rearrange social contexts” [20] by controlling who and what is seen [70] [36]. Systems can heighten favourable visibility, heighten unfavourable visibility or lower visibility. This can cause allocative [49] [23] or representative harms. The use of workplace technologies can also result in an environment where workers are treated as “fungible human capital” [11], experience deskilling [94], and are seen as easily replaceable. Deskilling of workers may be accompanied by an expectation to reskill. This represents a shifting of power away from workers and puts the burden on them to adapt in the face of changes they have no power over.

The second identified risk area in the literature is *reduced worker power*. AI systems often reflect, reinforce, and codify the inequities built into the power structures into which they are deployed [22]. Due to the power asymmetries that already exist in the workplace, workers will be particularly vulnerable to this [9] [42]. Loss of agency happens when an AI-system results in the reduction of autonomy and decision-making power for users [100]. Systems that mediate interactions and knowledge sharing may lower inter-worker relations, reduce their social capital [21], and cause alienation, impacting their ability to organise and engage in collective action. Collective action is a fundamental defence for workers [59] [94], enabling them to negotiate other working conditions such as wages or privacy concerns [80]. Data leverage, the influence individuals or groups have due to the reliance of computing technologies

on their data contributions, is one way workers could bargain with the aid of such systems [109].

The final identified risk area is *reduced worker wellbeing*. This can occur through increased workloads, feelings of surveillance, and less meaningful work. A system may create "work intensification" [23] if workers are required to maintain and feed data to the system in addition to their existing responsibilities [71]. Meaningful work is a key part of Rawlsian Justice in the workplace [8], yet a survey of the last century of automation suggests it has reduced meaningful tasks [40] [102]. Workplace systems can result in the unwanted collection of personal data [92] and the collection of data without informed consent [33]. Data collection can take the form of relying on workers to provide the system with continuous data, a form of participatory surveillance [10], or it could mine worker data without their awareness, a form of discreet surveillance [57]. This can lead to harms including loss of desired anonymity, non-consensual data collection, loss of the right to be forgotten, non-consensual representation or classification of individuals, and data used for unexpected purposes [100]. These three identified risk areas interact: reduction of worker value reduces worker power as it diminishes their bargaining power. The diminished power will, in turn, lower worker value and wellbeing, as it reduces their ability to negotiate better pay or working conditions.

Structured frameworks, including taxonomies and landscape analyses, encourage practitioners to anticipate risks, and formalise social and ethical considerations around deployment. Structured frameworks have been developed to analyse harms, risks, and failures associated with AI and other technological systems, from different perspectives. These frameworks can be oriented by specific types of models [69] [115] [106] [26], by particular harm types [67] [108] [7], by system failures [16] [91] [101] or by domain of use [97] [17] [111] [108] [67] [14]. Our analysis will be oriented by domain, focusing on risks to workers from AI-mediated EKA. We have not chosen a model-oriented approach because the risks arise based on how the model is entangled with the socio-technical environment that it operates within. Similarly, we have not narrowed our analysis to a specific type of risk, as many different types of risk can arise from a system of interest. Finally, we are interested in risks that result from functioning systems, which may be unintentional or unanticipated. System failures are out of the scope of this analysis.

3 CONSEQUENCES-MECHANISMS-RISKS FRAMEWORK

The aim of this research is to formulate an actionable framework to help practitioners identify and understand risks to workers associated with the deployment of AI-mediated EKA systems. In this paper, we use the term *AI-mediated enterprise knowledge access (EKA)* to describe a system that has predictive capabilities (AI), and that extracts and surfaces knowledge (knowledge access) used in an enterprise setting, such as retrieval systems or chat interfaces. This analysis consists of identifying the high-level *consequences* of moral importance of these systems and the specific *system mechanisms* that introduce risks to workers, which represent how these consequences may take effect. These consequences *risk* reducing the power, value, and wellbeing of workers. *Risk manifestations* are the concrete ways in which risks to workers could materialize. A

critical aspect of the framework is the separation between high-level systemic *consequences* and granular *mechanisms*, the latter contextualising analysis to a specific system and serving to identify actionable points of mitigation.

In this research, we identify four potential consequences: (1) Commodification; (2) Appropriation; (3) Concentration of Power; (4) Marginalisation. It is expected that these will interact. A first interaction is between (1) and (2). Commodification of knowledge enables the appropriation of that knowledge by the system. In turn, appropriation incentivises further commodification. A second interaction is between (3) and (4). Systems can push certain groups or individuals to the margins of an organisation whilst simultaneously causing power to be concentrated in the centre, worsening the effects of marginalisation.

We identify the mechanisms within an AI-mediated EKA system that introduce risk to workers within each consequence. The inclusion of mechanisms contextualises this work to AI-mediated EKA and enables the analysis to be applied to a specific deployed system. Practitioners could apply this analysis by reflecting on how mechanisms are exhibited within their specific system. The implications for workers of this introduced risk are based on the review in Section 2.2. The following sections will define each potential consequence, the related system mechanisms that introduce risk to workers, and the possible risk manifestations. Finally, we will provide considerations for practitioners working on these systems to help inform decisions about system design and deployment.

3.1 Process of Construction

The framework was constructed initially via a bottom-up literature review, see Section 2. The literature review process began with defining certain keywords and terms related to both enterprise knowledge systems and risks to workers. To identify relevant literature we used academic databases, such as Google Scholar, and then we reviewed the reference lists of relevant articles to identify further papers not captured in the database searches. We then performed a grouping and mapping exercise to extend this analysis to identify high-level consequences and map these to system mechanisms. For refinement, we conducted two workshops with two groups of our colleagues, who work as technology researchers and developers. The workshop attendees were presented a version of the framework, applied part of the framework to an EKA system, and we made some minor refinements to the framework based on their feedback on this process. The identified consequences, mechanisms, and risk manifestations in the framework are not intended to be exhaustive, but to provide guidance.

The scope of the framework is to consider consequences of moral import of AI-mediated EKA with corresponding mechanisms that introduce risk to workers. This framework does not explore other consequences of these systems that do not introduce risk, such as the potential for increased availability or implications for collaboration. We use the term workers to describe knowledge workers specifically, other types of workers are not in scope for this research. Most papers examine the point of origin of risk from the perspective of the AI model life cycle [113]. However, we examine this from the perspective of a deployed system, through system mechanisms, to account for the sociotechnical nature of the risks.

3.2 Commodification

The first potential consequence of AI-mediated EKA that is discussed is commodification. Commodification of knowledge refers to the acts of transforming knowledge artifacts into commodities, defined as objects of economic value whose instances are treated as equivalent, or nearly so, with no regard to who produced them [104]. Related, the commodification of workers has been examined previously in the context of the gig economy [90] and has parallels with the “datafication of employment” [6].

3.2.1 System Mechanisms that Introduce Risk. The process of commodification of knowledge may involve:

Disturbed Relationality Disturbed relationality refers to the acts of moving knowledge artifacts out of the relational context in which they exist or are produced, including who produced them, the social and procedural context in which they are produced, and the context of other knowledge artifacts in which they exist or were produced [31]. An example of disturbed relationality would be if a system did not maintain the state of a document (i.e. whether the document is a draft) or if it surfaced knowledge without the context of who authored it. HCI scholars question whether knowledge can be meaningfully surfaced, shared, and used in contexts outside of how it was created [4] [5] [3]. Orlikowski highlights that knowledge is not a “static embedded capability” but is ongoing and a part of those that engage with it [86]. The utility of knowledge is arguably contextual [87].

Changing Value of Certain Types of Knowledge Changing value of certain types of knowledge refers to the acts that lead to systemic differences in valuation of knowledge artifacts compared to their existing and historical valuation. For example, a system may value types of knowledge that rely on qualitative documentation, which it is built to mine, over visual knowledge.

Shifting the emphasis from praxis to proxies Shifting the emphasis from praxis to proxies refers to the acts of creating an environment where workers are encouraged to focus more on optimizing towards, often top-down, pre-stated quantitative measures of outcomes (proxies) over reflection and action directed at the outcome to be transformed. This process shifts away from reflection and action in the context of one’s work to optimizing for specified value measures in the context of capital production [34]. An example would be if a system uses ‘number of documents authored on a subject’ as a proxy metric for ‘expertise’, instead of a holistic view of a worker’s experience and long-term contributions to an area of expertise.

Standardisation and Commensuration Standardisation refers to the acts of enforcing conformity over things that are not strictly similar. The system can cause commodification if all knowledge is transformed into a fixed schema so that they (and correspondingly, knowledge experts) become interchangeable and standardised. This creates a situation where distinctiveness and differences are seen as annoyances rather than value generation [27]. An example would be, if a system adopts a standard schema to represent all properties of a specific type of knowledge, disregarding other information or

properties that cannot fit into this schema. Commensuration refers to the acts of transforming different qualities into a common metric [51]. The decision of what properties should be included may be commensurate with some notion of their usefulness, which disregards how properties may be differently relevant or important for a given knowledge artefact, context, discipline or individual. The concepts of standardisation and commensuration are inter-related. Commensuration may erase the value derived from distinctiveness and diversity making it harder to critique standardisation. Similarly, standardisation may aid the process of commensuration by enforcing conformity.

Homogenisation Homogenisation refers to the process of making things uniform or similar. Both standardisation and commensuration may lead to homogenisation because the system may enforce conformity of knowledge, such as fixed schemas, which can create a sociotechnical feedback loop that homogenises style and behaviour [27]. These systems may nudge creators towards adopting a similar style as more and more content is developed by and for AI as opposed to humans, which could change the quality and nature of knowledge [103]. An example of homogenisation would be if workers are incentivized to make their content more easily extractable by an automated system that may lead to homogenisation of their authoring style towards what the machine can best extract. Another way homogenisation can occur is through recommendation, if the model recommends the same extract to everyone searching for information on a particular subject, this will homogenize knowledge to that single expression of it [69].

3.2.2 Risk to Workers. Commodification introduces risks to workers. We outline how these risks could manifest.

Reduced Worker Value Commodification can have the following manifestations of risk to worker value. If a system takes the knowledge artefact out of the context of who authored it, an example of *disturbed relationality*, this may erase the author’s labour and expertise. This could also reduce workers’ feelings of ownership over their knowledge artefacts and recognition for their expertise. A system can *change the value* of certain types knowledge and knowing. For example, a system may not be able to ingest certain types of knowledge, such as visual knowledge in the form of graphs, which will reduce its visibility and perceived value. Workers who practice that knowledge will experience fewer opportunities through the system, such as being contacted as an expert. This shift in value of certain types of knowledge can be at odds with a worker’s area of expertise and could lead to deskilling. If a system *shifts emphasis from praxis to proxies*, workers’ contributions or expertise may not be fully captured by the chosen proxies, leading to erasure, fewer opportunities and less growth as the workers shift their focus to optimize the proxy instead of practicing their skills and learning from reflection. If certain types or parts of knowledge are not accommodated by the fixed system schemas, through *standardisation*, these will not be captured by the system and therefore erased.

Reduced Worker Power A system can result in loss of worker power through the following risk manifestations. This risk could manifest through loss of worker agency. If a system determines

which properties are useful for a given knowledge type, *commensuration*, this takes agency away from workers. Technological *standardisation* of work practices can prevent localised or team specific approaches [27]. Properties of a system, such as having a fixed schema, may nudge workers towards adopting specific style and producing *homogenized* content that is easily extractable. This may prevent workers from being able to work and create knowledge in individual and distinctive ways. If a system moves knowledge out of the context of its author, the system acts like an intermediary obfuscating the actual content creator [9]. This could remove the social process in which workers share knowledge and identify one another as experts. Instead of contacting colleagues for their expertise, workers communicate directly with the AI, which becomes the “expert”. These changes could impact social capital [21], alienate workers from their work community and reduce trust between workers, as seen with gig workers [90]. This risks weakening labor relations and the ability to organise.

Reduced Worker Wellbeing Commodification can have the following implications for worker wellbeing. If a system optimizes for *proxies*, this may change the nature of work. Proxies can result in additional workloads for workers to have their contributions captured by the system and lead to “perverse incentives” [71], where workers try to maximise the proxy (i.e. document production) as opposed to the true metric (i.e. quality work). A system could *change the value* of certain types of knowledge based on what it can mine and capture. This may push workers away from their areas of interest and expertise. Workers may experience less fulfillment as the system changes the nature of their work and responsibilities [57] [94]. Finally, a system may nudge workers towards certain knowledge types or *homogenized* authoring styles, making work less diverse and meaningful.

3.2.3 Considerations to Reduce Risk. Systems should be designed in a way that reduces the risks introduced by the mechanisms of commodification outlined above. One consideration would be to design a system that preserves context when surfacing knowledge, to protect worker recognition and value. Context preservation must be balanced with the risk of hyper-exposure or diminished privacy. There may be contexts in which authors do not want explicit attribution, therefore context preservation should be applied with either consent mechanisms or controls, such as privacy tags. It is also important to consider the proxies used within a system: how these are measured, what assumptions are embedded in their use and what their implications are. A system should meaningfully capture and value multi-modal knowledge types. A system should support workers to record knowledge in flexible structures and in varied styles. It is important to note that these considerations are a starting point and not exhaustive. For a specific system, which considerations are possible and impactful will depend on the data types, model choice, and deployment context of that system.

3.3 Appropriation

The second potential consequence of AI-mediated EKA is appropriation. Appropriation refers to the acts of taking something for your own use, usually without permission. We use appropriation as a

framing to discuss how a system may explicitly extract and implicitly infer workers’ knowledge so that it can co-opt this knowledge as the system’s own. Our use of appropriation has parallels with data colonialism, defined as “the predatory extractive practices of historical colonialism with the abstract quantification methods of computing” [38]. However, appropriation encompasses both extracting workers’ knowledge and co-opting it.

AI-based automation that explicitly tries to mimic the creative and work processes of workers may be seen as a distinct form of automation-by-appropriation [28]. There are parallels between how these systems could appropriate worker’s knowledge artefacts, obscuring workers’ labour and expertise in their creation, with the concept of ghost work and the hidden labour behind AI [55].

3.3.1 System Mechanisms that Introduce Risk. Appropriation may involve:

Knowledge Extractivism We use knowledge extractivism to refer to the process by which a system extracts individual and organisational knowledge, and appropriates it. This process is also known as “accumulation by dispossession” in data colonial theory, which describes data extraction within asymmetric power relations [107]. Our definition is distinct from definition by Pasquinelli et al. which refers to the accumulation of open source data, also known as Big Data [88].

Knowledge extractivism can occur through creating an environment where workers are incentivized to produce knowledge for a system or where a system mines knowledge without worker participation [10] to feed the system’s “enormous appetites” [39]. Due to workplace power imbalances, this extraction can occur without meaningful worker consent or control over the process [33]. The ability to co-opt the extracted knowledge depends on disturbed relationality, as the commodification of knowledge enables appropriation. An example would be if a system surfaces a paragraph from a document without recognition of who authored it, obscuring the author’s labour and co-opting the knowledge as the system’s own.

Capture of In-Use Knowledge Capture of in-use knowledge refers to acts of establishing exclusive control over implicit knowledge. This implicit knowledge is produced by and situated in the context of how people interact with these systems, such as user behaviour signals or curation. This can enable a system to generate new forms of knowledge [72]. For example, a system may recognize relationships between workers based on patterns of how users interact with each other and with artefacts. This has strong parallels to click signals leading to system improvements in web search [61]. This knowledge, that only the system has access to, created a “stickiness” factor and enables it to develop special strategic moats as well potentially leveraging that moat to further create walled-gardens and anti-competitive practices [75].

3.3.2 Risk to Workers. Here we outline the ways in which risk introduced by appropriation could manifest.

Reduced Worker Value Systems that appropriate knowledge can have implications for worker value. If a system *appropriates* a

worker's knowledge and expertise as its own by surfacing a paragraph they wrote without recognition of who authored it, this will reduce the worker's feeling of ownership over their work and the recognition they receive for it. Through *knowledge extractivism*, a system could co-opt workers' expertise and make them more replaceable, reducing incentives for organisations to invest in further practice or skills development.

Reduced Worker Power Appropriation can have negative implications for worker power. Data-based technologies can pose the opportunity for “data leverage”, which is the influence one has over computing technologies as they rely so heavily on our data contributions [109]. However, there is a risk that some systems will not enable workers to meaningfully remove their contributions from the system once it has been *extracted*. This prompts questions about who should own the knowledge extracted by these systems. Additionally, if a system *appropriates* workers' knowledge and skills, it will make them more replaceable, reducing their power in negotiations. The system could become the expert, centralising organisational expertise whilst removing recognition for individual workers. This could also reduce the need to contact colleagues for information, which could lower the sense of community between workers, and their participation in collective organising and bargaining.

Reduced Worker Wellbeing Appropriation can have negative implications for worker wellbeing by leading to surveillance, reduced privacy, and work intensification. A system can create an environment of *extraction* where the worker must explicitly or passively provide it with knowledge, including personal or private information. A system may *capture* and have proprietary control over data from user behavioural signals, which the worker is unaware of or has not meaningfully consented to. This can lead to feelings of surveillance or reduced privacy. A system can result in work intensification by creating an environment where the worker must explicitly provide it with knowledge to *extract* and must maintain that knowledge, in addition to their existing responsibilities. In addition, this appropriation of workers' expertise could make them feel alienated from their work identity and expert status.

3.3.3 Considerations to Reduce Risk. The appropriation of knowledge by these systems can introduce risk to workers through a number of mechanisms, as outlined above. The following considerations when designing and deploying AI-mediated EKA could aid in reducing these risks. Firstly, maintaining the context of the author when surfacing knowledge can help ensure workers get recognition for their expertise. Secondly, it is crucial that workers can meaningfully consent to both explicit data collection, such as from documents and emails, and implicit data collection, such as from behaviour signals. It is also important to consider how much data is being mined by the system and its relationship to utility. If the creation and maintenance of the supply of knowledge to the system creates more work than it removes for certain groups of workers, it is important for this to be reflected in a form of compensation. Finally, a critical consideration when developing systems for EKA, and beyond, is whether a system is being developed to support users in expanding the tasks they can do or developed to do the

tasks that users currently do. The latter creates strong incentives to appropriate workers' knowledge and expertise. Research has shown that the greatest productivity enhancers are technologies that support users in being able to do new tasks [28]. This approach to developing system functionality will reduce the risks discussed in this section.

3.4 Concentration of Power

The third potential consequence of AI-mediated EKA is concentration of power. Concentration of power refers to acts of worsening inequities in how power and control are distributed within an organisation. Shifting power away from workers is outlined as a risk area, but we propose that the mechanisms that concentrate power introduce risk across all three risk areas.

3.4.1 System Mechanisms that Introduce Risk. There are several system mechanisms related to the concentration of power:

Reduced space for negotiation Reduced space for negotiation refers to acts that shrink the set of decisions that workers can meaningfully contribute to or contest [63]. For example, in the problem of expert identification, these systems may take power away from workers to decide who is an expert on a topic and give that power to a system without means for contestation. Similarly, questions of who is allowed to curate, how curated knowledge is maintained, and how curations can be contested, raise questions about power and responsibility. A system may also enforce certain norms or perspectives, influencing the workplace, which will feed back into the system. This can cause “value lock-in” [113] where it can become harder for organisational norms to evolve because they are being reinforced by a system [52].

Worker Dispossession Worker dispossession refers to the acts of depriving workers of their job opportunities, skills, and expert status [107]. This relates to “labour distancing” [96], a concept that exists within Marxist theory where workers were distanced from their products by industrialisation, stripping them of their skills [74]. Under automation, certain job functions may disappear or be drastically reshaped, the corresponding workers over time may get deskilled due to lack of opportunities to practice their expertise, and their expertise in the original function rendered of low value. For example, a system could change a developer's role to spend more time writing documentation instead of writing code, as the system is better at recognizing word documents than python scripts. As well as dispossession from their expertise status, workers may experience it from the fruits of their labor.

System Opacity System opacity refers to the acts of intentional and unintentional obfuscation of how systems work under the hood. System opacity can further concentrate power by making it harder for workers and data subjects to meaningfully critique it. It may not be transparent what data a system is extracting, how individuals are represented in a system, and what proxies a system will optimise for. For example, workers do not know if a document they are working on can be mined by the system and seen by their colleagues. This can lead to workers forming folk theories, or algorithmic imaginaries [29], around how a system works and

how it represents them to improve their visibility, recognition, or other metrics [45]. This phenomenon has been seen in gigwork and social media [46].

3.4.2 Risk to Workers. The introduced risk to workers from concentration of power could manifest in the following ways.

Reduced Worker Value A system can reshape roles and *dispossess workers*, preventing them from practicing their area of expertise as they are nudged towards new areas of work, reducing their status as an expert. This can create precarity for the workers that may result in both their deskilling as well as raise expectations to learn new skills. A system may determine which types of knowledge and expertise are valued, without transparency about what values this is being driven by and how these can be *negotiated*.

Reduced Worker Power EKA systems can be *opaque*, meaning workers do not understand what data is being collected, how it is being used, and where it is being surfaced [8]. This makes it difficult for workers to meaningfully critique the system. A system could demand a high level of transparency from workers whilst its black-box nature means there is little transparency for workers in terms of data collection and decision-making. This could lead to high transparency and knowledge asymmetries between employer and worker, which will concentrate power in the hands of the organisation [8]. This opacity and the power asymmetries in the workplace make workers' consent to system practices "meaningless" as it may not be voluntary or informed [33]. Finally, a system can *reduce opportunity for negotiation* and debate between organisation and worker by enforcing values and processes. Workers may experience a loss of control over self-identification or representation [67] [73] which can lead to misrepresentation, stereotyping, and representative harms [115]. This can reduce workers' ability to bargain collectively and individually.

Reduced Worker Wellbeing A system can lead to feelings of surveillance or reduced privacy through the following mechanisms. Firstly, these *opaque* systems could mine data that workers are not aware of and use it in ways that workers do not understand. Workers may feel surveilled based on imaginings of how a system works. Finally, a system may not provide an opportunity for workers to *negotiate* or meaningfully consent to data collection practices.

3.4.3 Considerations to Reduce Risk. There are a number of considerations to reduce the risk from the system concentrating power. Firstly, consider designing the system to give workers control through the ability to curate or edit. This is particularly important when workers are implicitly or explicitly represented in the system. Alternatively, ensure there are clear paths of recourse for misrepresentation or inaccurate information produced by the system. Secondly, consider how the functionality of the system is communicated to workers. Do they understand what data is collected, how it is collected, and how it will be used? Relatedly, it is important to design the system to enable meaningful and informed consent; approaches to this are explored by Chowdhary et al. [33]. Finally, it is important to consider whether the design of the system changes the nature of workers' current roles and distances them from their areas of expertise.

3.5 Marginalisation

The fourth potential consequence, or by-product, of AI-mediated EKA in the enterprise is marginalisation. Marginalisation refers to the process of relegating certain individuals and groups to the fringes of society or of an organisation, and their corresponding discrimination. The sociotechnical feedback loop between a system and an organisational environment may also worsen inequities between groups [43].

3.5.1 System Mechanisms that Introduce Risk. Marginalisation may involve:

Systematic Reproduction Systematic reproduction refers to the acts that lead to marginalisation of the same groups who have been historically discriminated against by society (or an organisation). If this marginalisation occurs across multiple decision points in an organisation, it will become systemic. AI models have been shown to cause unfair discrimination and promote harmful stereotypes [1] against those on the social margins [41]. These systems are not neutral [58] and can perpetuate exclusionary norms by not capturing representations that exist outside of these norms [44]. This can lead to both representational harms and allocative harms [18]. For example, a knowledge extraction system may reproduce historical marginalisation by race or geography if it fails to adequately handle linguistic differences between groups because it has access to less training data from certain communities [64].

Demographic Blindness Demographic blindness refers to the acts of treating different individuals and groups uniformly when that uniformity is not warranted. This framing is inspired by color-blind racial ideology [48]. This relates to certain types of commensuration that result in marginalisation of groups. For example, ignoring linguistic differences when such differences exist may further intensify disparities in system outcomes for different language variants, such as sociolects or dialects [24], while making it difficult to acknowledge the resulting gaps. In some cases, this may map to blindness of groups that are not demographic.

Bias Amplification Bias amplification refers to the acts of amplifying pre-existing inequities. Algorithms themselves can be biased and amplify bias [105]. For example, an expert identification system may have certain demographic biases [66], but these biases may be further amplified in terms of exposure bias when a system's predictions are presented as a ranked list under heavy position bias in how users inspect the presented results [50] [117] [47].

3.5.2 Risk to Workers. The introduced risk to workers from marginalisation could manifest in the following ways.

Reduced Worker Value

A system can exhibit explicit discrimination of certain groups based on *historical marginalisation* [19], which has been seen in hiring algorithms [25]. Bias can also result from a lack of representation in a system's training set. For example, organisational geographic power asymmetries can lead to geographical performance disparities as a system may work better for more dominant languages [26]. By treating individuals and groups uniformly through

demographic blindness, certain communities may have their identities and needs erased. This can be worsened through *bias amplification*. If the same system can be used to control multiple aspects of working life, such as knowledge surfacing, expert identification, and hierarchy inference, risk will be amplified. This is because the strengths, weaknesses, limitations and, critically, the biases in the system will be standardised leading to *systemic marginalisation*.

Reduced Worker Power If certain groups are not represented in the system, through *demographic blindness*, it will be harder for them to form communities and organise. If individuals or groups of workers experience erasure, lack of representation, and poor service from a system due to *historical marginalisation* or *bias amplification*, this could reduce their value, discussed above, and therefore their power in both individual and collective bargaining.

Reduced Worker Wellbeing If a system suffers from *demographic blindness*, it may not capture the true diversity of representation and identity. Workers whose representations are not (fully) captured will experience alienation from their self-identity. Workers may also experience alienation from their work-identity and expertise if they experience *systemic* erasure and barriers to opportunity. Worker from historically marginalised groups may experience work intensification, compared to their peers, as they must invest more time curating content, contesting representation, or engaging in other efforts to improve their representation in the face of *marginalisation* from the system.

3.5.3 Considerations to Reduce Risk. There are a number of useful considerations for the design and deployment of an EKA system to reduce risk to workers from marginalisation. Marginalisation is embedded in social, organisational, and historical factors, therefore these considerations may be more nuanced. Firstly, it is important to consider what biases may exist in your context. In the example of bias in hiring algorithms [25], the historic marginalisation of women in hiring should have been considered when that system was designed. Secondly, one should consider how the ways in which a system mines knowledge or surfaces knowledge could amplify these biases. If a system is unable to mine knowledge from certain groups this could lead to algorithmic erasure [100] and the surfacing of knowledge could suffer from exposure bias [47]. One approach could be to measure and monitor metrics on how much exposure different groups get based on log data, however it is critical to ensure that monitoring in the name of equity does not lead to increased surveillance of users of the system. Finally, it is important to consider assumptions that a system may be making by treating groups of workers uniformly, if that is not in fact the case. These considerations are a starting point, there is a wealth of research that could enrich these suggestions.

4 DISCUSSION

4.1 Using the Framework in Practice

The aim of the framework is to support practitioners, who are involved with the design and/or deployment of AI-mediated EKA systems, towards understanding the introduced risk to workers and the paths to mitigation of those risks. The framework provides

information for general reflection on these systems and their associated risks. It can also be used to perform a more structured assessment of the risks in specific systems. The framework will be applied differently based on the system, organisational structure, and other factors. We hope to see best practices around the framework evolve over time with more practitioner use.

Practitioners can use the framework with their specific expertise to achieve an extensive risk assessment of their EKA system of interest. Here we propose one way in which the framework could be applied in practice, based on an exercise in our workshops. This framework is tailored to each system of interest through the mechanisms. For each consequence, practitioners could first identify which mechanisms are present in their system. It is not expected that all mechanisms will be present in all AI-mediated EKA systems. Then, the practitioner could consider all the ways in which each mechanism manifests in their system, reflecting on the definitions in the overview presented in Table 1 in Appendix A. Based on these mechanisms, practitioners could consider how the mechanisms could lead to risks to worker value, power, and wellbeing, referring to the examples in Section 3. This analysis would provide an overview of the mechanisms within their system that introduce risk and an understanding of how these risks could manifest, for each consequence of the EKA system of interest. This could enable practitioners to prioritise which risks to mitigate and, therefore, which system-level mechanism they should target.

4.2 Considerations for Large Language Models

In this paper we use the term AI-mediated EKA to generalise our framework across models so we can focus on the socio-technical system instead. Whilst we do not want to restrict this analysis to a single model, in this section we will outline how the specific properties of LLMs could impact each potential consequence of AI-mediated EKA, as LLMs present a paradigm shift in AI and extend what has been possible in this context. It should be noted that LLMs can enable new scenarios that posit risk to workers, not because of the LLMs themselves but because of the socio-technical systems they enable. In these cases, the LLM model could attract heightened and disproportionate attention from the FATE community, distracting from harms emanating from other parts of the system, thus providing cover to these alternate sources of risk.

4.2.1 Commodification. Commodification of knowledge can happen with any AI-mediated EKA system. However, LLMs may uniquely exacerbate commodification in a few ways. LLMs trained directly on the enterprise data that are interacted with directly will not maintain provenance when surfacing knowledge, leading to disturbed relationality. As well as standardisation of input, LLMs could lead to standardisation of output. LLMs are very effective at creating human-like text, which could lead to a situation where more and more content is developed by AI and fed back into the system, changing the quality and nature of knowledge [103].

4.2.2 Appropriation. In terms of appropriation, the automation from this technological advancement differs from previous iterations due to two properties; these models are general purpose and rapidly adaptable. This means there is a greater pace of change and reach of automation. There will be automation creep (e.g. starts

with a developer writing code with the help of AI, ends with a developer debugging AI's code). This is also known as "mission creep" [8]. In terms of data capture, the loss of provenance of where knowledge comes can further obscure the labour behind the knowledge artefact, enabling the system to appropriate it as its own. Finally, a unique aspect of LLM-based systems is their ability create human-like interactions, which will impact how they can capture information. The People and AI Research Group at Google found that human-like interactions can cause people to "disclose more information than they would otherwise, or rely on the system more than they should" [56]. There is also a risk that this could lead to communication fatigue, which has been seen with the rise of online meetings [63].

4.2.3 Concentration of Power. LLMs have certain properties that can lead to greater concentration of power. LLMs' generality and adaptability mean there is a high pace of change and adoption. This concentrates control with those deploying the system as opposed to those who the system is deployed on. LLMs trained directly on task-specific data will not maintain the provenance of data, which distances workers from their labour and removes the autonomy they have over how it is used. Additionally, there is a rise in LLM co-auditing tools, which encourage "skill-transfer" [54] from humans to the model. These systems will require less human input over time, which in the enterprise setting will distance workers from their expertise, a form of dispossession. If curation is not enabled, negotiation or recourse can be particularly difficult in the case of LLMs, due to the abstraction between the model developers and downstream applications [26]. In addition, there is a risk that LLMs could leak sensitive information present in training data [32] or could be used to infer sensitive information [113], such as age [78] [81]. These cases would be difficult for users to dispute. Finally, model opacity is often higher for LLMs, as their abstraction and sheer size makes documentation particularly complex. Gebru et al. described them as having "unfathomable training data" [53].

4.2.4 Marginalisation. In most AI-mediated EKA systems, there will be a risk of reproducing historic marginalisation. Language models in particular can perform differently across different languages [64] [95] and dialects or sociolects within languages [24]. LLMs represent and perpetuate norms present in training data by attempting to faithfully encode the patterns present in that data. This can lead to enforced exclusionary norms and harmful stereotypes [1]. In particular, language models are known to amplify biases [68] because they can often overrepresent the biases that appear in the training data [112] [117]. In the case of LLMs this can occur for biases both in the upstream model training data and the downstream task-specific data. In addition, as these models are general-purpose and highly adaptable, they can be used across multiple downstream tasks and decision points. This means any biases within the model could become systemic.

5 CONCLUDING REMARKS

In this paper, we present the Consequences-Mechanisms-Risks Framework to explore the consequences of AI-mediated EKA systems and identify risks to workers. Our framework provides a mapping from risks to the specific mechanisms within the system

that introduce that risk, and to the high-level consequences of AI-mediated EKA. The intention is that this conceptualisation will help practitioners apply this framework to their system of interest. The framework should provide a comprehensive overview of the potential risks to workers from these systems, however, it should not be seen as exhaustive or static, as both the technology and the resulting risks continue to evolve.

Future work could extend this framework to consider other identified consequences of moral import. The framework could also be extended to incorporate potential pathways for mitigation. A critical goal of our framework was to provide a mapping between the high-level consequences of the system, the specific system mechanisms, and potential introduced risks to workers. To evolve the framework, further research should be carried out into its evaluation. This could include detailing the prevalence of the different mechanisms and how the framework could be adapted across different AI-mediated EKA systems. The framework itself is highly flexible and could be widely applicable across domains. Future work could apply this novel Consequences-Mechanisms-Risks Framework to other deployed AI systems, working towards protecting workers and other groups.

ACKNOWLEDGMENTS

The authors gratefully acknowledge feedback on the framework and early drafts from Ida Larsen-Ledet, Solon Barocas, Mary L. Gray, and others. We also thank all of our workshop participants for their valuable contributions. Anna Gausen is supported by UKRI (grant number EP/S023356/1).

ETHICAL CONSIDERATIONS STATEMENT

The authors acknowledge that frameworks are both incomplete and not value neutral. A framework cannot be claimed to be exhaustive. Therefore the process involves the authors making normative and value-laden decisions about what is in scope, what is prioritised, and how to classify.

POSITIONALITY STATEMENT

Our research was carried out with an awareness of our positionality. As a group of authors, we encompass different genders, races, and cultural backgrounds. However, our insights will still be limited; we are a small group of authors, and all authors are highly educated, trained in computer science related fields, work in a western context, and at a large technology company. Additionally, the overwhelming majority of both our referenced sources and workshop participants are highly-educated and, the latter, work at the same technology company as the authors. This will inform the authors' understanding of what it is to be a worker, how an organisation works, how we understand risk, and other normative decisions made in this work.

ADVERSE IMPACT STATEMENT

The authors acknowledge that the development of this framework involved normative and value-laden decisions about what is in scope and what is prioritised, as mentioned. This process could have overlooked or de-prioritised certain risks. The framework could have adverse impacts on certain workers or groups of workers if they experience manifestations of these risks, which are not considered by the system designers or deployers.

REFERENCES

- [1] Abubakar Abid, Maheen Farooqi, and James Zou. 2021. Large language models associate Muslims with violence. *Nature Machine Intelligence* 2021 3:6 3 (6 2021), 461–463. Issue 6. <https://doi.org/10.1038/s42256-021-00359-2>
- [2] Danielle Abril. 2023. *Companies want to use AI tracking to make you better at your job*. The Washington Post. <https://www.washingtonpost.com/>
- [3] Mark S. Ackerman, Juri Dachtera, Volkmar Pipek, and Volker Wulf. 2013. Sharing knowledge and expertise: The CSCW view of knowledge management. *Computer Supported Cooperative Work: CSCW: An International Journal* 22 (8 2013), 531–573. Issue 4-6. <https://doi.org/10.1007/S10606-013-9192-8/METRICS>
- [4] Mark S Ackerman and Christine Miller. 1999. Organizational memory: processes, boundary objects, and trajectories. In *Proceedings of the 32nd Annual Hawaii International Conference on Systems Sciences*. 1999. HICSS-32. Abstracts and CD-ROM of Full Papers. IEEE, Hawaii, 12–pp.
- [5] Mark S. Ackerman, Volker Wulf, and Volkmar Pipek. 2002. *Sharing Expertise: Beyond Knowledge Management*. MIT Press, Cambridge, MA, USA.
- [6] Sam Adler-Bell and Michelle Miller. 2018. *The datafication of employment: How surveillance and capitalism are shaping workers' futures without their knowledge*. The Century Foundation, New York, NY.
- [7] Ioannis Agraftiotis, Jason RC Nurse, Michael Goldsmith, Sadie Creese, and David Upton. 2018. A taxonomy of cyber-harms: Defining the impacts of cyber-attacks and understanding how they propagate. *Journal of Cybersecurity* 4, 1 (2018), 1–15.
- [8] Ifeoma Ajunwa. 2020. The “Black Box” at Work. *Big Data & Society* 7, 2 (2020), 1–5.
- [9] Ifeoma Ajunwa. 2023. *The Quantified Worker: Law and Technology in the Modern Workplace*. Cambridge University Press, Cambridge, UK. <https://doi.org/10.1017/9781316888681>
- [10] Ifeoma Ajunwa, Kate Crawford, and Jason Schultz. 2017. Limitless Worker Surveillance. *California Law Review* 105 (2017), 735. <https://doi.org/10.15779/Z38BR8MF94>
- [11] Ifeoma Ajunwa and Daniel Greene. 2019. Platforms at work: Automated hiring platforms and other new intermediaries in the organization of work. In *Work and labor in the digital age*. Vol. 33. Emerald Publishing Limited, Bingley, UK, 61–91.
- [12] Saima Akhtar. 2021. *Employers' new tools to surveil and monitor workers are historically rooted*. Washington Post. <https://www.washingtonpost.com/>
- [13] Ali Alkhatib. 2021. To Live in Their Utopia: Why Algorithmic Systems Create Absurd Outcomes. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 95, 9 pages. <https://doi.org/10.1145/3411764.3445740>
- [14] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. *Machine bias: There's software used across the country to predict future criminals. And it's biased against blacks*. ProPublica.
- [15] AWS. 2024 (Accessed). Amazon Kendra. <https://aws.amazon.com/kendra/>
- [16] Jack Bandy. 2021. Problematic machine behavior: A systematic literature review of algorithm audits. *Proceedings of the acm on human-computer interaction* 5, CSCW1 (2021), 1–34.
- [17] Michele Banko, Brendon MacKeen, and Laurie Ray. 2020. A Unified Taxonomy of Harmful Content. In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, Seyi Akiwowo, Bertie Vidgen, Vinodkumar Prabhakaran, and Zeerak Waseem (Eds.). Association for Computational Linguistics, Online, 125–137. <https://doi.org/10.18653/v1/2020.alw-1.16>
- [18] Solon Barocas, Kate Crawford, Aaron Shapiro, and Hanna Wallach. 2017. *The problem with bias: Allocative versus representational harms in machine learning*. 9th Annual conference of the special interest group for computing, information and society, SIGCIS.
- [19] Solon Barocas and Andrew D Selbst. 2016. Big Data's Disparate Impact. *California Law Review* 104, 3 (2016), 671–732.
- [20] Nancy Baym and Nicole B Ellison. 2023. Toward work's new futures: Editors' Introduction to Technology and the Future of Work special issue. *Journal of Computer-Mediated Communication* 28 (6 2023), 1–5. Issue 4. <https://doi.org/10.1093/JCMC/ZMAD031>
- [21] Nancy Baym, Jonathan Larson, and Ronnie Martin. 2021. *What a Year of WFH Has Done to Our Relationships at Work*. Harvard Business Review. <https://hbr.org/2021/03/what-a-year-of-wfh-has-done-to-our-relationships-at-work>
- [22] Ruha Benjamin. 2019. *Race After Technology: Abolitionist Tools for the New Jim Code*. Polity, Cambridge, UK.
- [23] Annette Bernhardt, Lisa Kresge, and Reem Suleiman. 2021. *Data and Algorithms at Work: The Case for Worker Technology Rights*. UC Berkeley: Center for Labor Research and Education. <https://escholarship.org/uc/item/9831k83p>
- [24] Su Lin Blodgett, Lisa Green, and Brendan O'Connor. 2016. Demographic Dialectal Variation in Social Media: A Case Study of African-American English. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Jian Su, Kevin Duh, and Xavier Carreras (Eds.). Association for Computational Linguistics, Austin, Texas, 1119–1130. <https://doi.org/10.18653/v1/D16-1120>
- [25] Miranda Bogen and A. Rieke. 2018. *Help wanted: an examination of hiring algorithms, equity, and bias*. Upturn, Los Angeles, CA.
- [26] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, et al. 2022. On the Opportunities and Risks of Foundation Models. arXiv:2108.07258 [cs.LG]
- [27] Geoffrey C Bowker and Susan Leigh Star. 2000. *Sorting things out: Classification and its consequences*. MIT press, Cambridge, MA.
- [28] Erik Brynjolfsson. 2022. The Turing Trap: The Promise and Peril of Human-Like Artificial Intelligence. *Daedalus* 151 (1 2022), 272–287. Issue 2. https://doi.org/10.1162/DAED_a_01915
- [29] Taina Bucher. 2019. The algorithmic imaginary: Exploring the ordinary affects of Facebook algorithms. In *The Social Power of Algorithms*. Routledge, London, UK, 30–44.
- [30] Joy Buolamwini and Timnit Gebru. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A. Friedler and Christo Wilson (Eds.). PMLR, online, 77–91. <https://proceedings.mlr.press/v81/buolamwini18a.html>
- [31] Jodi A Byrd, Alyosha Goldstein, Jodi Melamed, and Chandan Reddy. 2018. Predatory value: Economies of dispossession and disturbed relationalities. *Social Text* 36, 2 (2018), 1–18.
- [32] Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. 2021. Extracting training data from large language models. *30th USENIX Security Symposium (USENIX Security 21)* 30, 21 (2021), 2633–2650.
- [33] Shreya Chowdhary, Anna Kawakami, Mary L Gray, Jina Suh, Alexandra Olteanu, and Koustuv Saha. 2023. Can Workers Meaningfully Consent to Workplace Wellbeing Technologies?. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) (FAccT '23). Association for Computing Machinery, New York, NY, USA, 569–582. <https://doi.org/10.1145/3593013.3594023>
- [34] Vanessa Ciccone. 2023. Transparency, openness and privacy among software professionals: discourses and practices surrounding use of the digital calendar. *Journal of Computer-Mediated Communication* 28 (6 2023), 1–10. Issue 4. <https://doi.org/10.1093/JCMC/ZMAD015>

- [35] Paul H Cleverley and Simon Burnett. 2019. Enterprise search and discovery capability: the factors and generative mechanisms for user satisfaction. *Journal of information science* 45, 1 (2019), 29–52.
- [36] Karina Cortiñas-Lorenzo, Siân Lindley, Ida Larsen-Ledet, and Bhaskar Mitra. 2024. Through the Looking Glass: Transparency Implications and Challenges in Enterprise AI Knowledge Systems. arXiv:2401.09410.
- [37] Sasha Costanza-Chock. 2020. *Design Justice: Community-Led Practices to Build the Worlds We Need*. The MIT Press, Cambridge, MA. <https://doi.org/10.7551/MITPRESS/12255.001.0001>
- [38] Nick Couldry and Ulises A. Mejias. 2018. Data Colonialism: Rethinking Big Data's Relation to the Contemporary Subject. *Television and New Media* 20 (9 2018), 336–349. Issue 4. <https://doi.org/10.1177/1527476418796632>
- [39] Kate Crawford. 2021. *The atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press, New Haven, CT.
- [40] Kate Crawford. 2023. Unpacking AI: "An Exponential Disruption". <https://www.msnbc.com/msnbc-podcast/>
- [41] Kimberlé Crenshaw. 2017. On Intersectionality: Essential Writings. *Faculty Books* 255 (2017), 320.
- [42] Roderic Crooks and Morgan Currie. 2021. Numbers will not save us: Agonistic data practices. *The Information Society* 37 (5 2021), 201–213. Issue 4. <https://doi.org/10.1080/01972243.2021.1920081>
- [43] Alexander D'Amour, Hansa Srinivasan, James Atwood, Pallavi Baljekar, D. Sculley, and Yoni Halpern. 2020. Fairness is Not Static: Deeper Understanding of Long Term Fairness via Simulation Studies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (FAT* '20). Association for Computing Machinery, New York, NY, USA, 525–534. <https://doi.org/10.1145/3351095.3372878>
- [44] Dipto Das, Carsten Østerlund, and Bryan Semaan. 2021. "Jol" or "Pani"? How Does Governance Shape a Platform's Identity? *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–25.
- [45] Michael A. DeVito. 2021. Adaptive Folk Theorization as a Path to Algorithmic Literacy on Changing Platforms. *Proceedings of the ACM on Human-Computer Interaction* 5 (10 2021), 38. Issue CSCW2. <https://doi.org/10.1145/3476080>
- [46] Michael A. DeVito, Jeffrey T. Hancock, Megan French, Jeremy Birnholtz, Judd Antin, Karrie Karahalios, Stephanie Tong, and Irina Shklovski. 2018. The Algorithm and the User: How Can HCI Use Lay Understandings of Algorithmic Systems?. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI EA '18). Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/3170427.3186320>
- [47] Fernando Diaz, Bhaskar Mitra, Michael D. Ekstrand, Asia J. Biega, and Ben Carterette. 2020. Evaluating Stochastic Rankings with Expected Exposure. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management* (Virtual Event, Ireland) (CIKM '20). Association for Computing Machinery, New York, NY, USA, 275–284. <https://doi.org/10.1145/3340531.3411962>
- [48] Ashley Doane. 2017. Beyond color-blindness: (Re)theorizing racial ideology. *Sociological Perspectives* 60, 5 (2017), 975–991.
- [49] Veena Dubal. 2023. On Algorithmic Wage Discrimination. <https://doi.org/10.2139/SSRN.4331080>
- [50] Michael D Ekstrand, Anubrata Das, Robin Burke, Fernando Diaz, et al. 2022. Fairness in Information Access Systems. *Foundations and Trends® in Information Retrieval* 16, 1-2 (2022), 1–177.
- [51] Wendy Nelson Espeland and Mitchell L Stevens. 1998. Commensuration as a social process. *Annual review of sociology* 24, 1 (1998), 313–343.
- [52] Iason Gabriel and Vafa Ghazavi. 2021. The Challenge of Value Alignment: from Fairer Algorithms to AI Safety. <https://arxiv.org/abs/2101.06060v2>
- [53] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé Iii, and Kate Crawford. 2021. Datasheets for datasets. *Commun. ACM* 64 (12 2021), 86–92. Issue 12. <https://doi.org/10.1145/3458723>
- [54] Andrew D Gordon, Carina Negreanu, José Cambronero, Rasika Chakravarthy, Ian Drosos, Hao Fang, Bhaskar Mitra, Hannah Richardson, Advait Sarkar, Stephanie Simmons, et al. 2023. Co-audit: tools to help humans double-check AI-generated content. arXiv preprint arXiv:2310.01297.
- [55] Mary L Gray and Siddharth Suri. 2019. *Ghost work: How to stop Silicon Valley from building a new global underclass*. Harper Collins, New York, NY.
- [56] People + AI Guidebook. 2019. *Mental Models*. Google PAIR. <https://pair.withgoogle.com/chapter/mental-models/>
- [57] Jessie Hammerling, Chris Benner, Jeffrey Buchanan, Françoise Carré, Beth Gutelius, Sara Hinkley, Ken Jacobs, Lisa Kresge, Adam Seth Litwin, Jenifer MacGillivray, Chris Tilly, and Steve Viscelli. 2022. *Technological Change in Five Industries: Threats to Jobs, Wages, and Working Conditions*. Berkeley Labor Center. <https://laborcenter.berkeley.edu/>
- [58] Donna Haraway. 1998. Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective (1988). *Feminist Studies* 14, 3 (1 1998), 236–240. <https://doi.org/10.2307/3178066>
- [59] Moritz Hardt, Eric Mazumdar, Celestine Mender-Dünner, and Tijana Zrnic. 2023. Algorithmic Collective Action in Machine Learning. <https://arxiv.org/abs/2302.04262v2>
- [60] Sheila Jasanoff. 2004. *States of knowledge: The co-production of science and the social order*. Routledge Taylor and Francis Group, London, UK. 1–317 pages. <https://doi.org/10.4324/9780203413845>
- [61] Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, Filip Radlinski, and Geri Gay. 2007. Evaluating the Accuracy of Implicit Feedback from Clicks and Query Reformulations in Web Search. *ACM Trans. Inf. Syst.* 25, 2 (Apr 2007), 7–es. <https://doi.org/10.1145/1229179.1229181>
- [62] Gabbrielle M. Johnson. 2020. Are Algorithms Value-Free? Feminist Theoretical Virtues in Machine Learning. , 35 pages. <https://doi.org/10.1163/17455243-20234372>
- [63] Jeffrey Alan Johnson. 2014. From open data to information justice. *Ethics and Information Technology* 16 (12 2014), 263–274. Issue 4. <https://doi.org/10.1007/S10676-014-9351-8/METRCS>
- [64] Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. The State and Fate of Linguistic Diversity and Inclusion in the NLP World. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 6282–6293. <https://aclanthology.org/2020.acl-main.560>
- [65] Wendy Ju and Larry Leifer. 2008. The design of implicit interactions: Making interactive systems less obnoxious. *Design Issues* 24, 3 (2008), 72–84.
- [66] Rikke Frank Jørgensen and Anja Bechmann. 2019. *Human Rights in the Age of Platforms*. The MIT Press, Cambridge, MA, Chapter Data as Humans: Representation, Accountability, and Equality in Big Data, 73–94. <https://doi.org/10.7551/MITPRESS/11304.003.0008>
- [67] Jared Katzman, Angelina Wang, Morgan Scheuerman, Su Lin Blodgett, Kristen Laird, Hanna Wallach, and Solon Barocas. 2023. Taxonomizing and Measuring Representational Harms: A Look at Image Tagging. arXiv preprint arXiv:2305.01776.
- [68] Hannah Rose Kirk, Yennie Jun, Haider Iqbal, Elias Benussi, Filippo Volpin, Frederic A. Dreyer, Aleksandar Shtedritski, and Yuki M. Asano. 2021. Bias Out-of-the-Box: An Empirical Analysis of Intersectional Occupational Biases in Popular Generative Language Models. *Advances in Neural Information Processing Systems* 4 (2 2021), 2611–2624. <https://arxiv.org/abs/2102.04130v3>
- [69] Hannah Rose Kirk, Bertie Vidgen, Paul Röttger, and Scott A Hale. 2023. Personalisation within bounds: A risk taxonomy and policy framework for the alignment of large language models with personalised feedback. arXiv preprint arXiv:2303.05453.
- [70] Ida Larsen-Ledet and Siân Lindley. 2022. Ways of seeing and being seen: People in the algorithmic knowledge base. Workshop at the 20th Eur. Conf. Comput.-Supported Cooperative Work.
- [71] Ida Larsen-Ledet, Bhaskar Mitra, and Siân Lindley. 2022. Ethical and Social Considerations in Automatic Expert Identification and People Recommendation in Organizational Knowledge Management Systems. In *Proceedings of the 5th FAccTRec Workshop on Responsible Recommendation at RecSys 2022*. ACM, New York, NY, 1–6.
- [72] Siân Lindley, Denise Wilkins, and Britta Burlin. 2021. Actions and their Consequences? Implicit Interactions with Workplace Knowledge Bases. , 6 pages. AutomationXP CHI.
- [73] Siân E. Lindley and Denise J. Wilkins. 2023. Building Knowledge through Action: Considerations for Machine Learning in the Workplace. *ACM Trans. Comput.-Hum. Interact.* 30, 5, Article 72 (Sep 2023), 51 pages. <https://doi.org/10.1145/3584947>
- [74] Karl Marx. 1844. Comments on James Mill, *Éléments D'économie Politique*.
- [75] Salil K. Mehra. 2011. Paradise is a Walled Garden? Trust, Antitrust and User Dynamism. *George Mason Law Review*, doi: 10.1016/B978-0-407-76001-1.50020-X.
- [76] Microsoft. 2023. *Microsoft Viva Topics Overview*. Microsoft. <https://learn.microsoft.com>
- [77] Shakir Mohamed, Marie-Therese Png, and William Isaac. 2020. Decolonial AI: Decolonial Theory as Sociotechnical Foresight in Artificial Intelligence. *Philosophy and Technology* 33, 1 (2020), 1–26. Issue 405. <https://doi.org/10.1007/s13347-020-00405-8>
- [78] Antonio A. Morgan-Lopez, Annice E. Kim, Robert F. Chew, and Paul Ruddle. 2017. Predicting age groups of Twitter users based on language and metadata features. *PLOS ONE* 12 (8 2017), e0183537. Issue 8. <https://doi.org/10.1371/JOURNAL.PONE.0183537>
- [79] Sarah Myers West, Veena Dubal, Zephyr Teachout, and Zubin Soleimany. 2023. *Algorithmic Management*. AI Now Salons. <https://ainowinstitute.org/series/ai-now-salon-series>
- [80] Nathan Newman. 2016. UnMarginalizing Workers: How Big Data Drives Lower Wages and How Reframing Labor Law Can Restore Information Equality in the Workplace. <https://doi.org/10.2139/SSRN.2819142>
- [81] Dong Nguyen, Rilana Gravel, Dolf Trieschnigg, and Theo Meder. 2013. "How Old Do You Think I Am?" A Study of Language and Age in Twitter. *Proceedings of the International AAAI Conference on Web and Social Media* 7 (2013), 439–448. Issue 1. <https://doi.org/10.1609/ICWSM.V7I1.14381>

- [82] Ikujiro Nonaka. 1994. A dynamic theory of organizational knowledge creation. *Organization science* 5, 1 (1994), 14–37.
- [83] Ikujiro Nonaka and Ryoko Toyama. 2015. The knowledge-creating theory revisited: knowledge creation as a synthesizing process. *Knowledge management research & practice* 1, 1 (2015), 95–110.
- [84] Ikujiro Nonaka, Ryoko Toyama, and Noboru Konno. 2000. SECI, Ba and leadership: a unified model of dynamic knowledge creation. *Long range planning* 33, 1 (2000), 5–34.
- [85] OpenAI. 2023. *GPT-4 Technical Report*. OpenAI.
- [86] Wanda J Orlikowski. 2002. Knowing in Practice: Enacting a Collective Capability in Distributed Organizing. *Organization Science* 13 (2002), 249–273. Issue 3.
- [87] Wanda J Orlikowski. 2006. Material knowing: the scaffolding of human knowledgeability. *European Journal of Information Systems* 15 (2006), 460–466.
- [88] Matteo Pasquinelli and Vladan Joler. 2021. The Nooscape manifested: AI as instrument of knowledge extractivism. *AI and Society* 36 (12 2021), 1263–1280. Issue 4. <https://doi.org/10.1007/S00146-020-01097-6>METRICS
- [89] Phillip H Phan and Theodore Peridis. 2000. Knowledge creation in strategic alliances: Another look at organizational learning. *Asia Pacific journal of management* 17 (2000), 201–222.
- [90] Jeremias Prassl. 2018. *Humans as a Service: The Promise and Perils of Work in the Gig Economy*. Oxford University Press, Oxford, UK. 1–199 pages. <https://doi.org/10.1093/OSO/9780198797012.001.0001>
- [91] Inioluwa Deborah Raji, I. Elizabeth Kumar, Aaron Horowitz, and Andrew Selbst. 2022. The Fallacy of AI Functionality. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 959–972. <https://doi.org/10.1145/3531146.3533158>
- [92] Divya Ramesh, Vaishnav Kameswaran, Ding Wang, and Nithya Sambasivan. 2022. How Platform-User Power Relations Shape Algorithmic Accountability: A Case Study of Instant Loan Platforms and Financially Stressed Users in India. *ACM International Conference Proceeding Series* 22 (6 2022), 1917–1928. <https://doi.org/10.1145/3531146.3533237>
- [93] David Randall, John Hughes, Jon O'Brien, Mark Rouncefield, and Peter Tolmie. 2001. 'Memories are made of this': explicating organisational knowledge and memory. *European Journal of Information Systems* 10 (2001), 113–121.
- [94] Rowena Rodrigues. 2020. Legal and human rights issues of AI: Gaps, challenges and vulnerabilities. *Journal of Responsible Technology* 4 (12 2020), 100005. <https://doi.org/10.1016/J.JRT.2020.100005>
- [95] Ruder. 2020. Why You Should Do NLP Beyond English. <https://www.ruder.io/nlp-beyond-english/>
- [96] Advait Sarkar. 2023. Enough With "Human-AI Collaboration". In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI EA '23). Association for Computing Machinery, New York, NY, USA, Article 415, 8 pages. <https://doi.org/10.1145/3544549.3582735>
- [97] Morgan Klaus Scheuerman, Jialun Aaron Jiang, Casey Fiesler, and Jed R Brubaker. 2021. A framework of severity for harmful content online. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–33.
- [98] Stefan Schmid, Cory Henson, and Tuan Tran. 2019. Using knowledge graphs to search an enterprise data lake. In *The Semantic Web: ESWC 2019 Satellite Events: ESWC 2019 Satellite Events*. Springer, online, 262–266.
- [99] Bariş Serim and Giulio Jacucci. 2019. Explicating "Implicit Interaction": An Examination of the Concept and Challenges for Research. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–16. <https://doi.org/10.1145/3290605.3300647>
- [100] Renee Shelby, Shalaleh Rismani, Kathryn Henne, AJung Moon, Negar Rostamzadeh, Paul Nicholas, N'Mah Yilla, Jess Gallegos, Andrew Smart, Emilio Garcia, et al. 2022. Sociotechnical harms: scoping a taxonomy for harm reduction. arXiv preprint arXiv:2210.05791.
- [101] Hong Shen, Alicia DeVos, Motahhare Eslami, and Kenneth Holstein. 2021. Everyday algorithm auditing: Understanding the power of everyday users in surfacing harmful algorithmic behaviors. *Proceedings of the ACM on Human-Computer Interaction* 5 (5 2021), 29. Issue CSCW2. <https://doi.org/10.1145/3479577>
- [102] Divya Siddharth, Daron Acemoglu, Danielle Allen, Kate Crawford, James Evans, Michael Jordan, and E Weyl. 2021. How AI fails us. arXiv preprint arXiv:2201.04200.
- [103] Matthew Sparkes. 2023. *Filling the internet with AI-created images will harm future AIs*. New Scientist. <https://www.newscientist.com>
- [104] Valerio De Stefano. 2018. 'Negotiating the Algorithm': Automation, Artificial Intelligence and Labour Protection. <https://doi.org/10.2139/SSRN.3178233>
- [105] Catherine Stinson. 2022. Algorithms are not neutral: Bias in collaborative filtering. *AI and Ethics* 2, 4 (2022), 763–770.
- [106] Alex Tamkin, Miles Brundage, Jack Clark, and Deep Ganguli. 2021. Understanding the capabilities, limitations, and societal impact of large language models. arXiv preprint arXiv:2102.02503.
- [107] Jim Thatcher, David O'Sullivan, and Dillon Mahmoudi. 2016. Data colonialism through accumulation by dispossession: New metaphors for daily data. *Environment and Planning D: Society and Space* 34, 6 (2016), 990–1006.
- [108] Thi Tran, Rohit Valecha, Paul Rad, and H. Raghav Rao. 2020. An investigation of misinformation harms related to social media during humanitarian crises. *Communications in Computer and Information Science* 1186 CCIS (2020), 167–181. https://doi.org/10.1007/978-981-15-3817-9_10
- [109] Nicholas Vincent, Hanlin Li, Nicole Tilly, Stevie Chancellor, and Brent Hecht. 2021. Data Leverage: A Framework for Empowering the Public in Its Relationship with Technology Companies. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) (FAccT '21). Association for Computing Machinery, New York, NY, USA, 215–227. <https://doi.org/10.1145/3442188.3445885>
- [110] Judy Wajcman. 1991. *Feminism confronts technology*. Penn State Press, University Park, PA. 184 pages.
- [111] Ashley Marie Walker and Michael A. DeVito. 2020. "More Gay" Fits in Better": Intracommunity Power Dynamics and Harms in Online LGBTQ+ Spaces. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3313831.3376497>
- [112] Angelina Wang and Olga Russakovsky. 2021. Directional Bias Amplification. *Proceedings of Machine Learning Research* 139 (2 2021), 10882–10893. <https://arxiv.org/abs/2102.12594v2>
- [113] Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, et al. 2021. Ethical and social risks of harm from Language Models. arXiv preprint arXiv:2112.04359v1.
- [114] Laura Weidinger, Maribeth Rauh, Nahema Marchal, Arianna Manzini, et al. 2023. Sociotechnical Safety Evaluation of Generative AI Systems. arXiv preprint arXiv:2310.11986.
- [115] Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, et al. 2022. Taxonomy of Risks posed by Language Models. *ACM International Conference Proceeding Series* 22 (6 2022), 214–229. <https://doi.org/10.1145/3531146.3533088>
- [116] John Winn, John Guiver, Sam Webster, Yordan Zaykov, Martin Kukla, and Dany Fabian. 2018. Alexandria: Unsupervised high-precision knowledge base construction using a probabilistic program. Automated Knowledge Base Construction (AKBC).
- [117] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai Wei Chang. 2017. Men Also Like Shopping: Reducing Gender Bias Amplification using Corpus-level Constraints. In *Proceedings of EMNLP 2017 - Conference on Empirical Methods in Natural Language Processing*.

A APPENDIX

A.1 Overview Table

Here we present an overview table of the consequences of moral import of AI-mediated EKA systems and the related system mechanisms. This table summarises the consequences and related mechanisms presented in Section 3. This table is aimed at helping practitioners apply the Consequence-Mechanism-Risk framework to their EKA system of interest by providing a high-level summarisation of definitions and examples of each mechanism. We encourage practitioners to read the full paper before attempting to apply the framework in practice.

accepted 30 March 2024

Table 1: Overview of the consequences and the related system mechanisms that introduce risk to workers from AI-mediated EKA.

Consequence	Mechanism	Description	Specific Example
Commodification	Disturbed Relationality	The acts of moving knowledge artifacts out of the relational context in which they exist or are produced, including who produced them, the social and procedural context in which they are produced, and the context of other knowledge artifacts in which they exist or were produced.	A system may not maintain the provenance of information, such as author, if it is an LLM that is trained directly on enterprise data and interacted with directly.
	Changing Value	The acts that lead to systemic departure in valuation of knowledge artifacts compared to their existing and historical valuation.	If a worker is skilled at debugging, this type of knowledge may not be captured by the system as it is hard to infer from written documents and therefore will be valued less.
	Shifting Focus from Praxis to Proxies	The acts of creating an environment where workers are encouraged to focus more on optimizing towards (often top-down) pre-stated quantitative measures of outcomes (proxies) over (ideally bottom-up) reflection and action directed at the outcome to be transformed.	A system may use number of contributions on a subject as a proxy metric for expertise. This could incentivise workers to write documentation for the system as opposed to practising their expertise.
	Commensuration and Standardisation	The acts of transforming different qualities into a common metric and enforcing conformity over things that are not strictly similar.	A system may adopt a standard schema to represent all entities of a specific type, disregarding the knowledge about individual entities that can't fit into the schema.
	Homogenisation	The process of making things uniform or similar.	If workers are incentivized to make their content more easily extractable by an automated system that may lead to homogenization of their authoring style towards what the machine can best extract from.
Appropriation	Knowledge Extractivism	The process of extracting knowledge and co-opting it as the system's own.	A system may surface a summary of a document without recognition of the original author, obscuring their labour.

Dimension of Risk	Mechanism	Description	Specific Example
Appropriation	Capture of In-Use Knowledge	The acts of establishing exclusive control over knowledge from interaction with the system.	A system may recognize relationships between different topics based on patterns of how users interact with them, rather than based on what is documented in the content.
Concentration of Power	Reduced space for Negotiation	Reduced space for negotiation refers to the acts that shrink the set of decisions that workers can meaningfully contribute to or contest.	In the problem of expert identification, a system may replace social processes with alternative computational approaches that take away any say the workers previously had in this context and limited contestation.
	Worker Dispossession	Worker dispossession refers to the acts of distancing workers from their expertise and interests.	A system may change a developer's role to spend more time developing content for the KB instead of writing code, as the system is better at recognising documents than python scripts.
	System Opacity	System opacity refers to the acts of intentional and unintentional obfuscation of how systems work under the hood.	Workers do not know what information about them a system collects and what may be inferred from it.
Marginalization	Systemic Reproduction	The acts that lead to marginalization of the same groups who have been historically discriminated against and marginalized by society (or the organization).	A system may reproduce historical marginalization by race if it fails to adequately handle linguistic differences between groups because it has access to less training data from certain communities.
	Demographic Blindness	The acts of treating different individuals and groups uniformly when that uniformity is not warranted.	Ignoring linguistic differences when such differences exist may further intensify disparities in system outcomes for different language variants while making it difficult to acknowledge the resulting gaps.
	Bias Amplification	The acts of amplifying pre-existing inequities and marginalization.	An expert identification system may have certain demographic biases, but these biases may be further amplified in terms of exposure bias when the system's predictions are presented as a ranked list under heavy position bias in how users inspect the presented results.