

KARINA CORTIÑAS-LORENZO, Microsoft, Switzerland  
 SIÂN LINDLEY, Microsoft Research, United Kingdom  
 IDA LARSEN-LEDET, School of Applied Psychology, University College Cork, Ireland  
 BHASKAR MITRA, Microsoft Research, Canada

CCS Concepts: • **Human-centered computing** → HCI theory, concepts and models.

Additional Key Words and Phrases: Transparency, Explainability, Knowledge Management, Representational harms

**ACM Reference Format:**

Karina Cortiñas-Lorenzo, Siân Lindley, Ida Larsen-Ledet, and Bhaskar Mitra. 2025. Through the Looking-Glass: Transparency Implications and Challenges in Enterprise AI Knowledge Systems. In . ACM, New York, NY, USA, 30 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

*Oh, Kitty! how nice it would be if we could only get through into Looking-glass House! I'm sure it's got, oh! such beautiful things in it! Let's pretend there's a way of getting through into it, somehow, Kitty. Let's pretend the glass has got all soft like gauze, so that we can get through.*

*...Oh, what fun it'll be, when they see me through the glass in here, and can't get at me!*

– Lewis Carroll, “Through the Looking-Glass”

## 1 INTRODUCTION

Knowledge is created, used and shared by people. Inherently human-centric and relational, in organizations it becomes embedded not only in documents and repositories, but also in processes, norms, communications and actions [51, 107, 126, 130, 152]. To foster effective utilization of knowledge in an organization, *Artificial Intelligence (AI) knowledge systems* operate by mining vast volumes of work-related data, automatically extracting and surfacing knowledge to

help individuals progress work [118, 161]. Because the production and the use of knowledge is embedded in social phenomena [152], these systems are intrinsically sociotechnical and must be human-centered by design.

One way to extract knowledge from work-related data involves mining associations between knowledge areas, artefacts and individuals [120, 159]. For example, this could involve automatically identifying who knows about a given knowledge topic in a company or linking related documents to a given subject area. As AI knowledge systems extract and surface these associations, they enable and shape different types of seeing [96, 162], impacting how we see and perceive others in an organization and how we see ourselves and our own contributions to organizational knowledge. In the enterprise context, what information gets brought to the foreground by these systems and what information is pushed to the periphery can influence who we interact and collaborate with, what information is trusted and how social interactions are established, potentially leading to interpersonal harms if the system misses relevant associations, surfaces inaccurate information or exposes work-related data without explicit user consent and proper contextualization [66, 97, 144]. As others are seen, individuals might also wonder how they are seen by others *through* the system lenses, what aspects of their work get surfaced by the system, to whom and where. How these questions are addressed and what information is disclosed when doing so has the potential to impact self-concept beliefs, influencing the extent to which these systems can augment and amplify humans in creating, sharing and applying knowledge in practice.

As we design and develop systems aimed at improving knowledge workers' efficiency and productivity, what gets seen and how end-users interpret what they see can have deep implications in the workplace, potentially leading to worsened situations by hindering individual's efficacy and well-being. In this context, transparency can help end-users of AI knowledge systems understand, interpret and calibrate the importance of what they see *through* the system, helping mitigate potential negative effects on self-concept beliefs and perceptions of others [64, 99, 163]. Yet, implementing transparency in complex real-world scenarios remains a challenge [57, 104]. As novel AI approaches such as Large Language Models (LLMs) continue to be developed, several challenges arise, such as opaque and large architectures, complex and uncertain AI capabilities and new stakeholder groups with unique transparency needs [104]. In enterprise knowledge systems, emerging LLM-based approaches can also exacerbate and pose new risks [66], introducing new biases into existing information retrieval capabilities [29, 50] and magnifying existing risks such as exposure inequality [6, 11, 61, 145]. As AI capabilities continue to evolve, several voices have raised concerns about the widespread deployment of LLM-infused applications before appropriate safeguards and risk mitigations are in place, highlighting the importance of taking additional measures to ensure that transparency and other responsible AI considerations remain a priority [82, 93, 104]. In this context, while past research has recognized transparency as a crucial aspect of responsible design [66, 97, 107], it's unclear how transparency should be implemented in enterprise AI knowledge systems, which forms of transparency should be prioritized and what the potential impacts and challenges related to their implementation might be. As AI systems continue to grow in complexity, reflecting on how to implement AI transparency in enterprise knowledge systems and the ethical implications of doing so is crucial to ensure these systems are human-centered by design and can augment both human labor and well-being.

To this end, this paper presents a reflective analysis of the potential requirements and impacts of AI transparency in enterprise AI knowledge systems. By *AI knowledge system*, we understand any system that extracts and surfaces knowledge by using predictive algorithms that learn from data. To inform our analysis, we conducted a multidisciplinary literature review drawing from different fields including computer science, psychology, organizational studies and human-computer interaction. The findings highlighted the complexities of incorporating transparency in enterprise settings, helping us contextualize the impacts of AI transparency in knowledge systems. We then hypothesized that

transparency in these systems can shape different *ways of seeing* in the workplace. Drawing on Lewis Carroll’s depiction of a magical mirror [33], we introduced the *looking-glass* metaphor as a conceptual tool for practitioners to critically examine the needs and implications of different transparency dimensions. As compared to other metaphors used to describe AI systems, the looking-glass metaphor conveys the idea that AI systems not only reflect, but also distort and reshape what is perceived, often in ways that are imperceptible to users, leading them to accept skewed perceptions as truth. The proposed metaphor is particularly valuable in sensitive contexts like workplace settings, where transparency can shape perceptions by creating a false sense of agency [8], impacting power dynamics in an organization [66] or introducing tensions between different stakeholder groups [132]. By broadening the scope of transparency, the looking-glass metaphor facilitates wider discussions of its potential effects, helping practitioners overcome “failures of imagination” [32] when anticipating transparency impacts and challenges. We demonstrated the use of the metaphor by applying it to our analysis of transparency implications in AI knowledge systems. This analysis highlighted the limitations of transparency initiatives solely focused on explaining the inner workings of AI knowledge systems (system transparency), emphasising the need to also consider how the system operates in the organizational context (procedural transparency) and the outcomes of system use in the long run (transparency of outcomes). As we reflected on the wider sociotechnical gap or divide between what we must support socially and what we can support technically [3], we used the term *AI transparency* to refer to a combination of decision-making and tooling aimed at revealing and concealing information about an AI system.

In summary, our contributions are four-fold:

- Building on past literature, we present the looking-glass metaphor and use it to conceptualize AI knowledge systems as systems that reflect and distort, enabling different types of seeing. This framework allows us to widen the scope of transparency needs, moving from a narrow system focus where the sole goal is allowing *seeing into* the system, to a wider view where what lies outside the system is brought to light.
- We use the looking-glass metaphor to explore the impacts of different types of seeing and identify key requirements for transparency in AI knowledge systems. Recognizing the sociotechnical nature of these systems, we identify three transparency dimensions that we believe are necessary to realize the value of AI knowledge systems, namely system transparency, procedural transparency and transparency of outcomes.
- We discuss key challenges that hinder the implementation of these forms of transparency, bringing to light the wider sociotechnical gap and tradeoffs to enable decision-making and support interdisciplinary discourse.
- We identify future areas of transparency research in AI knowledge systems, encouraging Computer-supported Cooperative Work (CSCW) researchers to further study transparency implications and first-order approximations to the problems we’re highlighting.

The analysis we contribute focuses on AI transparency requirements, impacts and challenges specifically in relation to enterprise knowledge systems. It is aimed at informing further empirical research and encouraging the CSCW community to reflect on the broader social, ethical and technical impacts of AI transparency in workplace settings.

## 2 RELATED WORK

### 2.1 Algorithmic Impact Assessments

The analysis of AI impacts is considered a critical step to understand the potential implications of AI systems in the world [32, 116, 136, 158]. Extensive research has shown that AI systems can lead to unintended consequences, having negative effects at the individual, collective and societal levels in ways that are often difficult to predict or control

[7, 18, 24, 141]. In response to these concerns, *algorithmic impact assessments* aim to anticipate adverse impacts *before* they occur (i.e., “*ex ante*” [158]), prompting interventions in the system’s design to mitigate any potential negative effects. Drawing directly from environmental protection, privacy and human rights policy domains [136], impact assessments are typically measured against alternative scenarios, making use of open-ended questions to uncover unknown unknowns [140] instead of “checklists” [32], and often relying on prior use cases of technology and narrative records of how the system was designed and iteratively developed [158]. By drawing attention to design choices and consequences [158], impact assessments can influence decision-making, encouraging more deliberate and thoughtful designs [158]. However, determining what constitutes an impact is not always clear [116]. Because impacts are not directly observed nor measured, the assessment of impacts is considered as much of a “thought experiment as an empirical endeavor” [116], being prone to “failures of imagination” [32], where possible impacts may be ignored or underestimated.

While impact assessments may not fully capture all possible implications, the discussions they generate are critical spaces for expressing dissent and facilitating negotiation [158]. Because foreseeing failures that have not been observed before is inherently a difficult task, wider and interdisciplinary discussions are needed [32], also for AI work aimed at making AI systems fair and transparent [129]. While transparency in AI systems is often considered a key mitigation strategy on its own and a fundamental dimension for responsible AI design [55, 90, 104], in contexts such as the enterprise it can have multiple ramifications and adverse consequences [66, 132], requiring a more in depth analysis of its potential impacts. In this paper, we fill this gap by critically reflecting on the needs for transparency in AI knowledge systems and the possible risks and challenges related to its implementation in the workplace.

## 2.2 AI Metaphors

Ideas about what technology is, or should do, can influence how its impacts are assessed and the effects it has on the world [68]. As cognitive tools [95], metaphors play a crucial role in shaping these ideas, influencing both design, development and evaluation choices. Considered a generative design tool [108] and pervasive in everyday life [95], metaphors can provide new ways of seeing, thinking about and experiencing technology, illuminating and hiding different aspects of it and creating opportunities to view objects or phenomena as something different [124]. Past research has shown that the practice of AI draws heavily on metaphors [124]. Some unnoticed foundational metaphors related to AI transparency include “explanation”, relating the computational information about a model operation to the social processes by which humans construct explanations, and “closed-box”<sup>1</sup>, implying that opening the opaque box would expose what’s hidden and divert attention from asking who lies outside the box or how the box came to be [124]. Metaphors like these can not only shape how AI systems are designed but also influence the expectations and evaluations of end-users when interacting with these systems [94]. For example, past work has conceptualized AI systems as “collaborators”, “teammates”, “assistants” or “co-pilots” [94, 137, 138, 142]. While some of these metaphors can personify AI systems with problematic consequences [94, 137], others can illuminate new design spaces, shaping the ways we interact with new technology [142].

Different metaphors have been used to describe AI personalization systems that implicitly or explicitly represent people. For instance, French [64] and Hess [79] conceptualize AI recommenders as mirrors reflecting the system’s approximation of the self back to the individual in the form of personalized content. Conceiving algorithmic personalization as a form of feedback about the self [64], interacting with AI recommenders is seen as “looking into a mirror

<sup>1</sup>We use the term “closed-box” instead of “black-box” to avoid perpetuating the narrative that “black” is undesirable, wrong or something to be avoided

rather than looking out a window” [79]. Personalized AI algorithms in digital platforms such as TikTok have also been conceptualized as “crystals” [99], conveying the idea that a crystal can be polished or refined by the user and inspiring new design choices that emphasize user agency and autonomy. The image of a crystal helps understand the capabilities of the system to support reflection of end-user self-concepts via personalized content, as well as refraction, shaping perspectives of others through the algorithm [99]. In this paper, we build on this work and propose the looking-glass metaphor to conceptualize AI knowledge systems. As our goal is to reflect on the needs for transparency in these systems and the challenges and impacts arising from its implementation, the metaphor highlights how AI systems can distort reality in subtle ways, skewing perceptions of the self and others and mediating transformative effects.

### 3 METHOD

#### 3.1 Literature Review

To guide our analysis, we reviewed multidisciplinary literature, focusing primarily on computer science, psychology, organizational studies and human-computer interaction. Due to the multidisciplinary nature of the reviewed works, we reviewed the literature iteratively, continuously refining our search and selection criteria as new insights emerged. To ensure multidisciplinary coverage, papers were drawn from multiple sources, including Google Scholar, ACM Guide to Computing Literature, PsycINFO and SAGE Journals. Weekly discussions among the authors helped synthesize the insights and refine the search, collaboratively iterating the emerging themes.

The literature review focused on three key aspects necessary for the assessment of transparency requirements and impacts in enterprise AI knowledge systems: the conceptualization of transparency, the use of AI methods in these systems and the impacts of AI outputs on self-perception and social processes. These aspects guided the search and the selection of relevant references across multiple disciplines. First, we reviewed the conceptualization of transparency, unpacking what is meant by the term in different contexts, both technical and social. By exploring different conceptualizations, we also sought to identify where tensions and trade-offs might arise. Second, we examined the use of AI methods in enterprise knowledge systems taking a narrative approach, reviewing their introduction over time and the different CSCW perspectives shaping their role in knowledge extraction and surfacing. Lastly, we explored the reported impacts of AI systems on self-perception beliefs and social processes to understand the extent to which AI transparency, and the lack thereof, can impact individuals in workplace settings. By reviewing these three interconnected areas, the literature review aimed to inform our understanding of the nuances of AI transparency in enterprise contexts, helping inform our assessment of requirements and potential impacts.

#### 3.2 Reflective Analysis and Positionality Statement

As researchers, we acknowledge that our backgrounds, life experiences and beliefs can shape our work. Indeed, the reflective analysis we contribute in this paper was motivated by our past work researching how AI technology can augment knowledge workers and be integrated in enterprise settings in a way that’s ethical and human-centered. Our analysis of potential transparency requirements and challenges was influenced by our own lived experiences as end-users of AI knowledge systems within a large technology company. As we reflected on the types of seeing enabled by these systems and how these can be shaped by transparency mechanisms, we combined reflections from the literature review with insights gathered via discussions with AI researchers working in AI knowledge systems in the same company, who provided valuable inputs on the practical challenges of implementing different forms of transparency. Acknowledging the role that the researcher’s position plays in the creation of knowledge, our research team engaged

with reflexive practice. Our educational background and collective experience in HCI and machine learning, along with our roles as knowledge workers in a large technology corporation, influenced our own understanding of what it means to be a knowledge worker, how an organisation works and what constitutes an adverse impact. We recognize, however, that these perspectives are not unique and are also shaped by our location in a Western context, influencing our interpretations. As normative decisions were made throughout the analysis, we acknowledge that the foresight of transparency requirements and impacts presented in this paper is necessarily incomplete and not value neutral.

## 4 LITERATURE REVIEW

We begin by reviewing the use of the term *transparency* in different disciplines and the tensions arising from different conceptualizations. Next, we characterize AI knowledge systems in work settings and review the use of AI algorithms to extract and surface organizational knowledge. We finish by exploring the self-concept psychology and sociology theories explaining the mechanisms that underlie self-perception and how the formation of beliefs about oneself can be shaped by AI systems.

### 4.1 Transparency: *seeing* and *knowing*

As AI systems continue to evolve and become integrated into several aspects of our lives, calls for greater transparency have been widespread. However, what is meant by transparency is often left implicit [9, 46]. While in fields such as optics and physics transparency denotes the property of a material to let light pass through, in others it's often conceptualized as a wide array of practices aimed at holding collectives accountable [98]. Regardless of the discipline, most works see transparency as being closely intertwined with *seeing*: a transparent object enables observation through it whilst a transparent organization makes aspects of its operations visible and, hence, open to scrutiny. Seeing, understood in a metaphorical sense as “being aware of”, doesn't automatically equate to the kind of knowing or understanding that enables action [8, 98]. If we understand transparency as informational transparency, that is, as information about an AI system that should be disclosed to enable appropriate understanding [104], then, considering the sheer and increasing complexity of AI systems, transparency necessarily requires a conscious decision and assessment of which information *should* be made visible, sometimes at the expense of concealing other information [46].

What to conceal, what to reveal and how much to conceal or reveal are timeless questions that have been examined in philosophy, computing sciences and Human Computer Interaction (HCI) communities. In philosophy, past works have differentiated between *transparency in use*, understood as technology being so easy to use that it allows individuals to *see through* it and focus on the activity at hand [28], and *reflective transparency*, interpreted as the extent by which a system allows a user to *see into* its mechanisms [9, 43, 134]. In CSCW, the discussion was initially framed from an awareness perspective [110], with a focus on the interpretation of computing-based representations of people, activities or artifacts. In representing people and their work, *we have to make a choice, as to what we exclude and include, what we represent in detail and what is abstract, what we actively interpret and what we leave to human interpretation* [37]. In the context of knowledge management, social translucence was defined as an approach to designing digital systems that emphasize sharing social information, supporting visibility, awareness and accountability [59]. Subsequent CSCW work proposed a framework for thinking about social transparency, highlighting potential tradeoffs and implications of making social information visible for collaborative work [148]. In HCI, the tension between revealing and concealing information was formalized as a decision on which seams or mismatches to make visible at the expense of what to hold invisible [38], initiating a debate between seamless design – emphasizing clarity, simplicity and ease of use – and seamful design, prioritizing configurability, user appropriation and revelation of complexity [84]. In the realms of



privacy research, the question has often been depicted as a conflict between selective control of access to the self, and the desire to see and be seen by others [65, 131]. In determining disclosure behavior, privacy works have highlighted the importance of knowing the audiences involved in the process of information sharing and consumption [131], as well as the nature of the relationship [91] and levels of trust [127] between the discloser and the recipient.

Like privacy [131], transparency is not one-sided. More transparency is not necessarily better [46, 125]. The optimal state along the spectrum of revelation and concealment is highly contextual. By shifting the focus towards the potential impact of making something visible, our work draws on welfarism (the maximization of overall benefits and welfare with consideration to different stakeholders) as a basis to evaluate conflicting transparency desiderata [125]. Building on past research, we understand transparency in the context of AI as a combination of decision-making and tooling aimed at revealing and concealing information, thus modulating different types of seeing in AI systems.

## 4.2 Enterprise AI Knowledge Systems

The vast scale of knowledge generated by organizations, coupled with the challenges of externalizing tacit knowledge and identifying those who hold it, poses several challenges to the efficient utilization of knowledge in work settings. As people work, large volumes of dispersed and siloed data might be generated, such as documents created using productivity tools, meeting transcripts or emails produced via collaboration platforms. In this context, the goal of AI knowledge systems is to support knowledge management and dissemination in an organization by automatically extracting, organizing and surfacing organizational knowledge. In this section, we review the history and role of AI in this type of system, focusing on how AI algorithms can represent individuals both implicitly and explicitly. In what follows, we refer to *AI knowledge systems* as systems that use predictive algorithms that learn from data to support knowledge extraction and sharing.

**4.2.1 Using AI to extract knowledge associations.** Early technical CSCW solutions for knowledge management and dissemination focused on the externalization of knowledge in the form of repositories, emphasizing *knowledge sharing* via the codification of knowledge in databases [4]. While these solutions aimed to centralize knowledge, they relied on rigid information structures, requiring users to manually input data and quickly becoming outdated [3, 165]. In this context, AI methods have offered a way to address these shortcomings by facilitating the automatic extraction of associations between knowledge areas and knowledge contributions [87, 105, 159]. As such, over the past years, Natural Language Processing (NLP) techniques like those using deep learning and probabilistic models have been applied in different tasks, including information extraction [80, 102, 133], automatic knowledge graph construction [169] and knowledge base population [36, 88, 161]. Yet despite continuous AI development and improvement of extraction performance via novel AI approaches such as LLM-based extraction [52, 103, 167], several voices in the CSCW community have warned against the decontextualization of knowledge, raising awareness of the need to support not only the externalization and categorization of knowledge, but also its social dimensions [3, 51, 107]. In this context, **implicit associations** between knowledge areas and knowledgeable people can help users interpret and contextualize knowledge contributions extracted via AI. For instance, displaying social metadata like the author of a document can allow users to gather more information about the knowledge source, facilitating its reuse in different contexts (Figure 1).

Recognizing that knowledge is situated in relation to people, *expertise sharing* efforts in CSCW emphasized the social aspect of knowledge use, placing less emphasis on the externalization of knowledge in favor of fostering collaboration between knowledgeable individuals [4, 107]. A key area of focus became how to identify who knows what in an organization [20, 165], finding ways to automatically connect individuals who need information with experts who

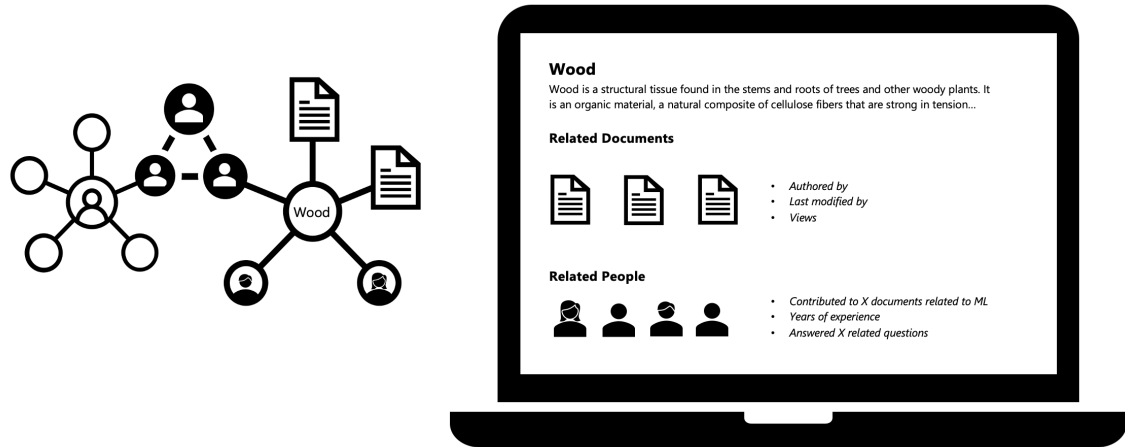


Fig. 1. AI knowledge management systems can *implicitly* associate people with topics or areas of knowledge via retrieved documents (e.g., by surfacing a list of related documents) or *explicitly* link people to areas of knowledge (e.g., by displaying a list of related people). When displaying linked documents, various metadata, such as the author, the last individual who modified the document, or the number of views, can be shown. When showcasing a list of related people, metadata such as the number of contributed documents, years of experience in the organization, or the number of questions answered can be displayed.

can provide it [107]. While an initial approach involved the use of manually curated expert databases [165], these solutions suffered from similar shortcomings as knowledge databases, requiring labor intensive maintenance and leading to rapid obsolescence [20, 165]. As a response to these challenges, different AI methods have been applied in past literature on tasks such as automatic expert profiling in organizations (also referred to as expert profiling or expert identification) [15, 62, 62], people search by expertise on a topic (known as expertise search, expertise retrieval or expert finding systems) [72, 83, 106] or algorithmic recommendations of knowledgeable people to help progress work [21]. Aimed at fostering collaboration, AI algorithms in these contexts use work-related data to mine **explicit associations** between areas of expertise and individuals (Figure 1). Despite being a construct highly influenced by the social context surrounding an individual [113], past literature in the AI domain commonly refers to “expertise” as a modelling target. As such, different sources of data have been used to quantify it, such as documents produced by individuals as they work, completed learning activities, social signal such as answered questions in a forum and endorsements, and/or communications in the form of emails or chats [72, 83, 106]. Mined associations between artefacts and knowledge areas can also be leveraged when modelling explicit linkages between individuals and knowledge topics. For example, heuristic associations between people and retrieved documents such as authorship or amount of edits can be used to determine who should be linked to a knowledge area. However, when linking individuals to documents, not all people appearing in a document might contribute equally to the topic [106]. Thus, how associations between people and content are established and the types of sources considered by the system are important transparency aspects that can help end-users understand system behavior.

**4.2.2 Surfacing knowledge associations.** Once extracted, associations between knowledge and people can be surfaced in a myriad of ways. For example, they can be integrated into search results [26, 69, 120], employee profiles [119] or served as suggestions in chat-based experiences to help progress work [5, 14, 121–123]. As several documents and people might be linked to a given topic of interest, ranking models are commonly used to surface the most associations



to end-users [92]. For instance, Viva Topics [120] presents a curated list of suggested people related to a knowledge topic, allowing end-users to expand the list via UI controls. To identify and rank the most knowledgeable individuals on a given topic, different AI ranking models have been used in past literature, including generative probabilistic models, supervised learn-to-rank methods, similarity approaches and voting methods [83, 106].

While not all commercialized knowledge systems implement all of these extraction and surfacing capabilities, different knowledge management products offer various combinations of them. For instance, in Microsoft Viva [118], AI technology is used to automatically extract and bring knowledge to users of Microsoft 365 apps via the Topics application, where areas of knowledge, people and knowledge resources are automatically linked and collected into a coherent repository [120]. Glean [69] builds a knowledge graph that relates people and content and surfaces this information via search in a personalized way. People are linked to knowledge areas based on whether they've written about a topic. If someone is in need of a knowledge resource that doesn't yet exist, Glean leverages these associations to suggest subject matter experts who could help progress work. Other AI knowledge systems that extract and surface knowledge include different capabilities for different roles in an organization. For example, Bloomfire [26] provides leadership solutions to help executives and managers drive strategic decision-making, and reporting tools to help knowledge management teams know what resources are being accessed and which ones aren't.

*4.2.3 Risks and ethical considerations.* Extracting and surfacing algorithmic representations of people and their work, whether implicitly or explicitly, raises multiple ethical concerns [97], as they can potentially lead to representational, allocative, quality of service and/or interpersonal harms [144]. Past research has shown that AI information extraction methods often underperform when applied on data that deviates from the training set [2, 42], such as languages with different morphology [81, 143] or low resource levels [52, 78], multilingual text [35, 42, 160], as well as text documents coming from different domains [12] or written using various perspectives and writing styles [2]. Additionally, several voices have warned against the naive use of social data, or any kind of digital trace produced by or about people [128], as methodological limitations such as the lack of *construct validity* (i.e., whether a specific measurement, like the amount of authored documents, actually measures the theoretical construct of interest, such as human expertise) and biases introduced along the analysis pipeline are often overlooked, leading to important negative implications [19, 25, 128]. More recent AI developments, such as LLM-based methods, have been shown to introduce additional risks via the fabrication of inaccurate information, generating factually incorrect data by extrapolating information from the biases in training data, misinterpreting ambiguous prompts or modifying the information to align superficially with the input [89, 154]. In enterprise knowledge systems, emerging LLM-based approaches can exacerbate and pose new risks [66], introducing biases into existing information retrieval capabilities [29, 50] and magnifying existing risks such as exposure inequality [6, 11, 61, 145]. As connections between knowledge, artefacts and people are automatically extracted and surfaced, important challenges at the intersection of work and the social also arise, such as the introduction of new workloads for those individuals linked to knowledge areas in the system or the reinforcement of existing power dynamics in an organization [97]. Therefore, as compared to non-AI approaches, the introduction of AI methods in sensitive contexts like knowledge management systems opens the space to new technical and societal challenges, as inaccuracies in knowledge extraction and surfacing can lead to disproportionate harms [66], requiring a holistic understanding of the AI's sociotechnical impacts [146].

While past work has highlighted the increased risks to workers posed by LLMs and people recommenders recognizing transparency as a key dimension of responsible design [34, 61, 63, 66, 97], to the best of our knowledge, related work on how to approach transparency in AI knowledge systems that represent people is scarce. In recommender and

LLM-based systems, past research has recognized the closed-box nature of deep AI methods as a key barrier towards achieving transparency [104, 168]. Motivating transparency from a trust and user acceptance angle, past works in social recommender systems have highlighted the importance of not only explaining why certain items are recommended by an AI algorithm, but also integrating controllability with transparency, increasing end-user satisfaction by enabling end-users to participate in the recommendation process [77, 156, 157]. As such, past works have integrated various kinds of controllability inputs such as adjustments of recommendation preferences [157], in combination with explainability measures of familiarity, social closeness and commonalities between people as a way to explain both people [76] and item recommendations [30]. Recent work has also proposed making visible users' algorithmic representations in recommendation and personalization systems by communicating to end-users how they are being perceived by the system using tag clouds [164], automatically informed profiles [17] or natural language explanations [40, 135]. In the realms of knowledge management, knowledge sources (that is, documents and other resources provided as "evidence" for a given association) and knowledge currency (whether an association is up to date) have been proposed as transparency requirements for enterprise knowledge bases [163].

### 4.3 Theories of self: self-concept and its impact on human motivation and learning

As AI knowledge systems get embedded into organizational contexts, different ways to approach AI transparency might lead to different impacts on self-perception and social processes. But to what extent can transparency modulate these effects? And to what extent can AI knowledge systems influence how individuals perceive their identity at work, or how they evaluate their contributions? Could these effects be detrimental to workers' well-being and productivity, going against the same goals for which these systems are developed and deployed? In this section, we start by reviewing different self-concept theories from psychology and sociology that explain how individuals form and maintain their self-perception and how the formation of beliefs about oneself can shape human behavior, motivation and well-being in the workplace. Next, we review the reported impacts of AI systems on self-perception beliefs and social processes.

*4.3.1 Identity and self-concept.* Identity and self are multifaceted. How someone thinks about, perceives or evaluates themselves is influenced by the imagined opinions and judgement of others [45, 86, 114], being constantly shaped by social interactions, perceived membership in a social group [49, 109, 150] and expected audiences [153]. Self-concept beliefs are also modulated by behavior [67]. As an individual learns how to best interact with an environment to achieve their goals, they form self-efficacy beliefs that influence their motivation and drive to further adapt and learn, shaping individual's confidence in their ability and capacity to accomplish a task [16]. How to best interact within an environment also means how to best interact with others and, in doing so, how to best present oneself. As if it was a continual performance [71], individuals are driven to match their representations to the imagined audience's expectations and to their ideal self. In organizational settings where goals are tied to economic incentives, the self becomes constrained and shaped by organizational discourses of power ("crystallized" as described by Tracy et al. [155]). These constraints might induce individuals to tone down disfavored attributes to fit into a given discourse, a phenomenon known as "covering" [70, 166]. At work, self-perceptions of abilities, skills and knowledge become key components of self-concept beliefs [101]. Such beliefs are fundamental attributes in identity management [41] and are constantly shaped by collected feedback. Organizational behavior, hence, becomes driven by a desire to collect feedback that confirms and enhances existing self-perceptions and by the resolution of unpleasant states of cognitive dissonance when receiving feedback that disconfirms existing perceptions of self [101]. As individuals receive feedback from others, their meta-perceptions – beliefs about how they are perceived by others – are also calibrated [74]. Extensive research

indicates that individuals' meta-perceptions can have implications on their affect, cognition, behavior and relationships [74]. In organizations, meta-perceptions of expertise, in particular, have been shown to have deep implications for employee motivation and performance, shaping productivity, intent to stay in an organization and psychological well-being [74, 75, 115].

*4.3.2 Impacts of algorithmic feedback on self-perception and identity formation.* Feedback doesn't have to be human to influence self-perceptions. Exposure to algorithmic bias in the form of personalized recommendations can shape self-perceptions of confidence and leadership skills [64] as well as people's understanding of their identities and their orientation to others [23, 85, 99, 149]. The extent to which algorithmic feedback can affect self-perceptions has also been shown to be dependent on how individuals perceive causes in their environment and the extent to which system outputs are perceived to reflect the individual's intrinsic characteristics, such as their personality or interests, as opposed to external circumstances such as the behavior of other people in the platform or software updates [64]. How algorithmic inferences are interpreted is heavily dependent on how people make sense of complex systems and the informal theories they establish to explain what they see and experience while using the technology. Known as *folk theories*, these theories are influenced by prior beliefs and expectations [54], are malleable and dynamic [53], can be influenced by the system's transparency [60] and evolve with direct experience, helping individuals navigate the complexity of ever-changing AI systems.

As AI-mediated technology gets embedded into social processes, folk theories and algorithmic imaginaries enable individuals to cope with social boundaries that become blurry and distorted, helping them navigate spaces where audiences aren't directly seen but rather imagined and collapsed into one [112]. For example, social media technologies such as X (formerly Twitter) flatten multiple audiences into one, a phenomenon known as "context collapse" that forces end-users to present themselves in the same way to different audiences or people in their network [112]. In online environments where the way an individual gets represented can impact both social capital and economic opportunities, user understandings of social media systems can influence their ability to adapt to them, driving self-presentation behavior [53] and mediating self-transformative effects [73]. As AI systems represent people, individuals may have little control over representations that are derived from having done something at a particular point in time and/or from sharing attributes with a particular demographic segment, prompting those being represented to adopt different strategies in an effort to resolve dissonances between their algorithmic selves and their perceived and/or ideal selves, such as creating distinct online identities or embracing sites likely to appear high in rankings in self search results [111], adjusting work strategies in the case of gig [100] and digital workers [39] and being deliberate in how to interact with technology in social media apps [99].

As we design and develop systems aimed at improving knowledge workers' well-being and productivity, thinking through the impacts of technology on self-concept beliefs becomes of paramount importance. As we assess different scenarios and ramifications below, we see that unintended transparency effects might counteract core system purposes, potentially leading to worse-off situations at the individual level.

## 5 TRANSPARENCY IN ENTERPRISE AI KNOWLEDGE SYSTEMS

The literature review highlighted a number of important considerations when addressing transparency in AI knowledge systems. First, transparency is multifaceted and closely linked to *seeing and perception*. *It determines what information is revealed or concealed, impacting what users perceive while interacting with AI systems.* Second, *the use of AI algorithms to represent individuals and their expertise in enterprise knowledge systems introduces ethical risks, as the opacity of these*

*algorithms, poor performance on unseen data and lack of construct validity can lead to representational, allocative and interpersonal harms. Transparency mechanisms are necessary not only to mitigate these risks but also to ensure end-users can understand and trust the system’s representations. Lastly, past research has shown that AI systems embedded in organizational and social contexts can influence individuals’ self-perception and identity formation, with these effects being shaped by transparency mechanisms and users’ interpretation of system outputs.*

Drawing from these insights, we hypothesize that transparency in AI knowledge systems can shape different ways of seeing in the workplace. Thus, we start our analysis of transparency requirements in AI knowledge systems by assessing the types of *seeing* enabled by these systems in organizational settings. To identify these, we use the looking-glass metaphor as a way to conceptualize AI knowledge systems, bringing to the foreground the social context surrounding their use. From this perspective, we identify different forms of understanding sought by individuals in organizational settings, motivating the need for different transparency dimensions: system transparency, procedural transparency and transparency of outcomes.

### 5.1 The looking-glass metaphor: seeing *into* and *through* AI knowledge systems

In the sequel to Alice’s Adventures in Wonderland [33], Carroll depicts a magical mirror or *looking-glass* through which Alice enters a strange world that is recognizable, yet subtly altered. When Alice sees herself and steps through this mirror, she finds herself in a parallel reality where things get distorted (Figure 2). Similarly, as AI knowledge systems get embedded into the workflows of knowledge workers, individuals might be confronted with different reflections of their own work. These reflections might not be accurate or aligned with people’s own perceptions of their contributions or the way they wish to portray themselves to others in the organization. As they engage with algorithmic representations of people and knowledge, they might step into new settings where views of others are simplified and distorted (Figure 3). Inspired by past literature depicting AI systems as mirrors [64, 79] and crystals [99], we use the looking-glass metaphor to conceptualize enterprise AI knowledge systems as systems that both reflect and distort, allowing individuals to see themselves and others via algorithmic lenses. As compared to the mirror metaphor [64, 79], the looking-glass metaphor conveys the idea that AI systems not only reflect, but also distort and reshape what is perceived, often in ways that are imperceptible to users. Similarly, while the crystal metaphor implies that AI systems can shape perspectives of others as refracted images of the self (e.g., by suggesting other people with similar interests in social media platforms) [99], the looking-glass metaphor emphasizes how AI systems distort reality in ways users are unaware of, leading them to accept skewed perceptions as truth. Using this conceptual frame, we identify different types of seeing enabled by these systems and reflect on the questions end-users may pose while *seeing*.

**5.1.1 Seeing into the system.** If Alice knew how the looking-glass worked, this understanding would have probably shaped the way she experienced the alternative world she stepped into. AI knowledge systems that represent people and their work often consist of multiple interacting models, services and components. While some of the AI models used to extract and surface associations might be opaque, like LLMs, other components relying on traditional AI approaches, such as enterprise knowledge graphs, allow for some degree of transparency and insight into system behavior. Thus, even if full algorithmic transparency might not be possible, some of the inner workings of the system and the decisions made during its design and development can be exposed through documentation and user interface explanations, allowing users to see *into* the system. For example, users might want to know what type of data is mined when extracting knowledge (such as written documents) or what metrics are used to rank people associated with a knowledge area (e.g., the number of authored documents). If we take the perspective of an end-user facing an AI




Fig. 2. Alice entering the looking-glass. Illustration by John Tenniel. (Source: [https://en.wikipedia.org/wiki/Through\\_the\\_Looking-Glass#/media/File:Aliceroom3.jpg](https://en.wikipedia.org/wiki/Through_the_Looking-Glass#/media/File:Aliceroom3.jpg))

knowledge system with opaque behavior, additional questions about the system might arise. For instance, how is the data used to extract knowledge? In which organizational contexts is this data being created? What aspects of work are impacted by these data extraction practices and how might these influence system drift and adaptations over time? As knowledge is surfaced to end-users, questions might also cover what are the contexts in which knowledge is surfaced and how the exposure to algorithmic representations shapes knowledge sharing practices in the organization over time.

*5.1.2 Seeing others individually through the system.* Because AI knowledge systems automatically extract and surface knowledge using data that is directly or indirectly produced by people as they work, individuals see both the extracted knowledge and, through that, the people who created the data. For example, in systems that mine associations between knowledge areas and knowledge artefacts [69, 120], users might be able to see who are the authors of artefacts being surfaced (Figure 1). In systems that surface explicit associations between areas of expertise and knowledgeable individuals [26, 69, 120], end-users can directly relate a given knowledge topic to an individual (Figure 1). As people use the system and step through the mirror, the exposure to knowledge created by others *and* extracted by the system can modulate how they see other individuals within their organization, allowing them to infer constructs such as expertise. As AI algorithms often under-perform in low-resource domains and can only mine explicit, externalized knowledge (Section 4.2.3), not all work generated in an organization might be mined and surfaced. This can lead to a distortion in how individuals are represented and seen in the organization, implicitly inducing proxies for expertise and indirectly shaping interpersonal trust.



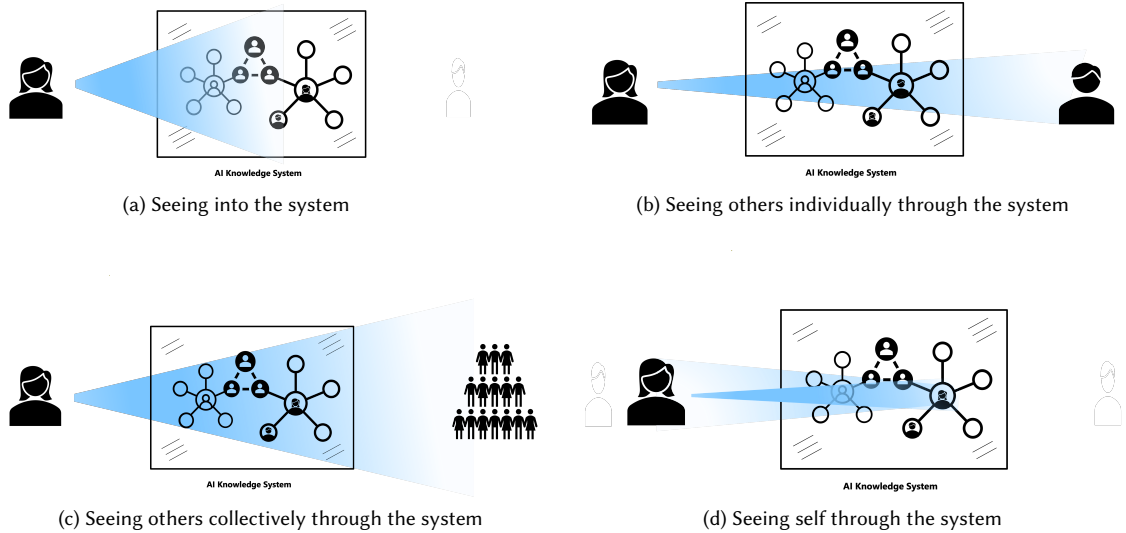


Fig. 3. When people interact with AI knowledge systems, they experience different ways of seeing. In the visual representations above, the blue lines illustrate these forms of perception. Transparency mechanisms can influence these types of seeing, shaping how users perceive the system, others and themselves.

When stepping through the mirror, the pieces of work within the system’s training domain might be given a better light, while unseen examples outside this domain, which are harder to extract, may be overlooked. In this way, people who create knowledge in forms suited for typical AI extraction, such as individuals authoring textual documents, might be more visible than those whose contributions are in a form less interpretable by current AI systems. As a simple example, the system may under-perform in low-resource languages due to the comparatively low number of examples it can draw on. In an English-speaking organization, a Japanese subject matter expert might not get exposure if their working documents are authored in Japanese, while another individual who is less knowledgeable about the topic but writes about it in English might be identified as an “expert”, potentially leading to new, undesirable workloads for this individual while overlooking the actual expert. This example illustrates a sociotechnical problem, where the technical limitations of the system in low-resource languages intersect with organizational dynamics, leading to unequal distributions of work and misrepresentation of expertise within the organization.

As end-users see others individually *through* the system, understanding how the system extracts the exposed associations can help them decide how to best act on the system outputs. For example, what data features and documents does the system consider when relating someone else to an area of knowledge? How reliable are the associations between a given individual and the knowledge identified by the system? What actions and knowledge contributions support these associations and how aligned are algorithmic representations of a person’s expertise with the way they see their own competences and skills? Are these individuals aware of these algorithmic representations and do they feel comfortable being contacted about the knowledge areas with which they are associated? As the system pushes to the foreground knowledge artefacts relating people and knowledge, how does this shape perceptions of expertise across the organization? And how does this influence how users interact and collaborate with others?



**5.1.3 Seeing others collectively through the system.** As end-users see others individually in the organization *through* the AI knowledge system, these experiences might accumulate over time, implicitly informing people's schemas of others and contributing to the establishment of patterns and categorizations that help them navigate organizational complexity and the ever-changing dynamics of work. As seen in Section 4.2.3, issues of construct validity when quantifying human expertise can cause AI knowledge systems to inadequately recognize different types of expertise, entirely overlooking the expertise of certain groups and potentially reinforcing stereotypes. When individuals step through the algorithmic *looking-glass* they might only see certain segments of people in their organization [31]. For example, demographic groups working with external collaborators or with lower tenure in the organization might not be seen, despite the knowledge they hold and help create. The variety of data sources considered by AI knowledge systems and the importance given to proxies that are easy to quantify, such as the amount of documents or edits, directly determines which segments get exposure to their work and which groups see their labor systematically erased.

As end-users step *through* the mirror and see others collectively through the system lenses, several questions might arise. For instance, who gets to be seen and which groups are systematically erased? How do the data types considered by the system and data features extracted influence which pieces of work are seen? Are there groups in the organization that contribute to knowledge not captured by the system? What incentives are in place driving knowledge contributions suitable for system extraction? How does the system influence the exposure to different segments of people over time and what biases and proxies of expertise are being introduced?

**5.1.4 Seeing self through the system.** As individuals see others *through* the system, they might wonder how they are seen themselves by others in their organization. In doing so, they might wish to know how they're seen by the system, which pieces of work are accessible and not accessible to the system and how their knowledge contributions are made visible to others. Drawing on the literature review, this type of seeing has the potential to impact self-concept perceptions, affecting self-efficacy beliefs when the system fails to correctly recognize knowledge contributions. As individuals see themselves through the system, the reflections might be distorted and not accurate of what others see. For instance, since the view of a person depends on what data sources the observer has access to, an individual may be represented based on documents they themselves don't have access to, leading to misalignments between how they perceive themselves and how they're seen by others through the system. As shown in Section 4.3.1, when these reflections aren't aligned with an individual perceived and/or ideal self, individuals might take action to solve potential dissonances. As an example, if the system surfaces a list of individuals related to a knowledge area indicating the amount of documents they've authored (as illustrated in Figure 1), when the system misses a given individual working on the topic, this individual might increase the amount of documents they author to get exposure in the system, aligning system outputs with the way they wish to be represented. As AI knowledge systems rely on quantifiable aspects of externalized knowledge, distorted reflections placing disproportionate weight on simple metrics can lead to perverse incentives, prompting changes in behavior that are not necessarily aligned with productivity nor employee well-being. In particular, missed knowledge and wrong surfaced associations can lead to representational harms, inducing feelings of alienation and lack of social control.

As end-users are confronted with these reflections, they might seek understanding about how they're seen by the system and how they're seen by others *through* the algorithmic lenses. From a system perspective, questions might include what data sources and data features are considered by the system when extracting a given association, why might algorithmic associations be incorrect or incomplete and how can end-users see and control these associations. As associations are surfaced to others, who can see these and in what way? How are system outputs used by others in

context? As the system allows individuals to know and control how they're seen, do people in the organization have different levels of visibility and/or control (e.g., depending on their role or hierarchical position)? As an outcome of seeing and interacting with system reflections, how is the system shaping an individual's behavior towards knowledge creation and sharing and how might these changes in behavior influence system adaptations over time?

## 5.2 Transparency requirements in AI knowledge systems

While the list of questions above isn't exhaustive, thinking about these questions can help us discern different forms of transparency needed to help end-users interpret what they see while interacting with AI knowledge systems. Some of these questions relate to *technical* information about the features implemented by the system and the behavior of AI mechanisms used to extract and surface knowledge. Other questions, however, relate to *social* information about the production and use of knowledge in the organization and the way the system is embedded into knowledge workflows and practices. As the *technical* and the *social* interact, different outcomes at the individual and organizational level might take place. Based on this, we categorize the questions and transparency requirements into three key dimensions:

- **System transparency:** information about the AI knowledge system targeting how the system works, including what data is consumed by the system and how knowledge gets extracted and surfaced via AI. This type of transparency helps end-users understand when they're seeing things *through* algorithmic lenses and how these ways of seeing might be incomplete, distorted and governed by system rules and behavior.
- **Procedural transparency:** this transparency dimension covers how the system operates *in context* and how the system is embedded in the organization. This includes how system outputs are used and consumed by people in different organizational contexts, the norms surrounding system use as well as the processes and practices established by the organization when adopting and deploying the technology.
- **Transparency of outcomes:** making visible the impacts and outcomes resulting from the *looking-glass* use is fundamental to help end-users understand how the system impacts themselves and their organization in the long run, how their behavior might be influenced as they use the system over time and how these changes might in turn affect system development and adaptations.

As individuals make use of AI knowledge systems, they see not only *into* the system, but also others and themselves *through* the system. When we think of AI knowledge systems as *looking-glasses*, opening a closed-box to reveal the system's inner mechanisms ceases to be the sole focus of transparency. The looking-glass metaphor thus complements the closed-box, moving us from a narrow system focus where the primary goal is allowing end-users to *see into* the system [56] to a wider view where what lies around the system is brought to light. As we extend the discussion of transparency to span system behavior, contextual use and outcomes, different forms of understanding can be provided, potentially mitigating representational, allocative and interpersonal harms by helping end-users better interpret and make sense of what they see. Table 1 outlines questions practitioners may ask themselves to interrogate the intersections between ways of seeing and the three dimensions of transparency. However, what we should implement from an end-user and social perspective might not be aligned with what we can technically implement in practice. In the following section, we bring light to this sociotechnical gap and identify key challenges arising in the implementation of these forms of transparency.

Table 1. Questions for interrogating the intersections between ways of seeing and the three dimensions of transparency.

	System transparency	Procedural transparency	Transparency of outcomes
<b>Seeing into the system</b>	<p>What type of data is mined by the system and how is it used to extract knowledge?</p> <p>How does the system surface knowledge? E.g., what metrics does the system use to rank knowledge contributors, and how does this influence where they are surfaced or how they are presented?</p>	<p>In which organizational contexts is data created? Which of the organizational data is the system using?</p> <p>In which organizational contexts is knowledge being surfaced?</p>	<p>What aspects of work are impacted by data extraction practices? How might altered work practices and outputs influence the system over time?</p> <p>How is the system shaping knowledge sharing practices in the organization?</p>
<b>Seeing others individually through the system</b>	<p>What data and features does the system consider when relating someone to an area of knowledge?</p> <p>How reliable are surfaced associations between a given individual and knowledge?</p>	<p>What actions and knowledge contributions in the organization support the associations inferred by the system? Do these reflect the way expertise is established socially in the organization?</p> <p>How comfortable are people with the way they are being associated with knowledge areas? E.g. are they open to be contacted about the areas they have been associated with?</p>	<p>How is the system shaping the way the expertise of individuals is perceived in the organization?</p> <p>How is the system shaping the way individuals interact with others in the organization over time?</p>
<b>Seeing others collectively through the system</b>	<p>How does the type of data considered by the system influences who gets to be seen?</p> <p>What features and metrics does the system consider, and how do these influence the pieces of work and contributors that end-users see?</p>	<p>Are there groups of people in the organization who contribute knowledge that is not captured by the system?</p> <p>What incentives in the organization drive knowledge contributions that are suitable for system extraction?</p>	<p>What are the impacts of disparate system performance across different demographic subgroups? How is the system shaping the exposure to different segments of people over time?</p> <p>How is the system shaping the establishment of objectives and measures of success in the organization?</p>
<b>Seeing self through the system</b>	<p>What data and features does the system use to represent an individual's work? Why might algorithmic associations be incorrect and/or incomplete?</p> <p>How can end-users know and control how they're represented by the system?</p>	<p>How well does the surfaced system representation match individuals' self-perception?</p> <p>How can end-users know and control how they're seen by others through the system? Do people in the organization have different levels of visibility into and/or control of how they're seen?</p>	<p>How are individuals adapting their behavior towards knowledge creation and sharing to bridge their self-perception and the system representation?</p> <p>Can individuals tell how changes in their behavior influence system behavior over time?</p>

## 6 TRANSPARENCY IN ENTERPRISE AI KNOWLEDGE SYSTEMS: KEY CHALLENGES AND IMPLICATIONS

How transparency gets implemented matters. In this section, we explore what challenges might arise when implementing different forms of transparency in AI knowledge systems. Thinking about these challenges is important to bring into light the gap between what we can technically implement and what we need to implement from a social perspective. When thinking about potential challenges, differentiating between system transparency, procedural transparency and transparency of outcomes can help us identify which aspects require a technical solution and which ones require a social approach, informing further interdisciplinary discourse and decision-making in work settings. As the use of emerging AI techniques in enterprise knowledge systems raises multiple challenges, introducing new risks and magnifying existing ones (Section 4.2.3), our discussion is focused on enterprise knowledge systems powered by AI.

Still, we recognize that some of the issues covered here are not unique to AI and can also occur with non-AI approaches such as heuristic rules.

## 6.1 System transparency

Several technical challenges arise in the implementation of system transparency in AI knowledge systems. While some AI methods such as deep learning techniques are inherently opaque, in this section we discuss the fundamental challenges arising with basic mechanisms of transparency, such as exposing the data used for inference and knowledge extraction. We acknowledge that in the presence of opaque algorithms, these mechanisms are not sufficient to make AI knowledge systems fully transparent and that the issues discussed in this section can still arise if transparent and non-opaque AI algorithms were used for knowledge extraction. In what follows, we explore some of the difficulties related to documentation and communication of system behavior, faithfulness, actionability and perverse incentives in work settings.

- Communication of system capabilities, mechanisms and limitations:** Especially in the age of LMMs, the way information gets communicated matters [104]. When communicating system capabilities and limitations, special care must be taken to account for different backgrounds and levels of exposure to AI. As different sources of information and media continue to shape folk theories about AI, public perception of AI technology is likely to include misconceptions, limiting the extent to which end-users can understand and act upon system documentation [22, 139]. Due to the proprietary nature of many commercialized AI systems and competitive factors, providers have incentives to withhold detailed information about how and where AI is being used, limiting the amount of information that is shared with the public. When detailed information is provided, system documentation might be ignored, misinterpreted or used for the benefit of certain stakeholder groups over others. For instance, when communicating non-intended uses, the exception can become the rule if the consumer is inadvertently prompted to misuse the system. In this way, documenting that the system might offer only a partial view of an employee's contributions could highlight the fact that the system does provide *some* insights into employee performance and productivity, directing the consumer's attention to that potential misuse. To prevent uses that shouldn't be supported by the technology, full transparency is, hence, not enough and can, in fact, be detrimental, spurring actors to behave in an adversarial manner and game the system to their advantage. Safeguards against these uses must be implemented and deployed. For instance, by preventing aggregated views and their use in reporting on people, or by incorporating inverse transparency by design [170].
- Faithfulness and access:** One way to improve the transparency of AI knowledge systems involves making visible the pieces of work used to relate people to knowledge (sometimes referred to as "evidence"). This data might be created by different individuals belonging to different teams and sub-organizations, and since many organizations have file access rules and privacy settings that vary depending on membership of departments and teams, people's access to sources of information may vary from person to person. In practice, this means that evidence for knowledge associations may not be accessible to the end-users who see them. This can have different implications for how individuals see *into* the system and how they see others and themselves *through* the system. To illustrate the impacts when seeing *into* the system, in global organizations, legal requirements and restrictions can prevent data transfer across geopolitical boundaries. This leads to differing configurations of access control that might prevent people from seeing knowledge contributions or documents authored by people in a different location than their own. If these knowledge contributions are used to explain system associations,

this results in inconsistent explanations across end-users: While those who can access the documents are in a better position to understand system behavior, this is not the case for those who can't see the documents supporting the associations. There are similar implications for seeing *others through* the system: The inability for individuals to be sure whether views of each other are shared might accumulate, so that the collective as a whole can't be certain of the degree to which they share a view of the organization. This could negatively impact organizational identity and ability to act collectively.

Finally, there are also implications for seeing *self through* the system: In principle, the resources used to make inferences about an individual can include resources that they do not have access to, such as human resources documents or other people's emails and chats that mention the individual. The inability to expose these pieces of data prevents the individual from getting the full picture of why they're represented in the system in the way they are. As an example, an individual may have been associated to a knowledge area based on, among other things, personal chat communication between their team members as they coordinate work. The right approach to transparency is not merely a matter of providing as much transparency as possible. In this case, informing the individual that some of the evidence for the association can't be exposed due to privacy constraints could have a deep impact on the beliefs the individual holds about how they're being seen by others in the organization. Is their continued employment being discussed in emails among management? Has someone made a complaint about them that has been documented in a human resources file? As seen in Section 4.3.1, how an individual imagines others perceive them can influence self-efficacy beliefs and self-esteem, directly affecting people's behavior and interpersonal relationships within the organization. The lack of transparency into the way others see the individual can add to this: Not only are there sources that contribute to how others see them that they can't themselves access, but certain others may actually be able to access those sources, with the effect that those other people's views of the individual are informed by resources that the individual can't access, let alone control.

These examples show that not only can discrepant access to organizational data result in an incomplete picture of the system, but negative consequences can follow both from incorrect assumptions of completeness and from attempts to fill in the gaps of the incompleteness. Furthermore, the use of work-related data as "evidence" enforces a layer of transparency in the social context of the organization, directly or indirectly exposing information that people might have chosen not to bring attention to. But failing to provide transparency can negatively affect the productivity and well-being of employees, as it hinders their ability to understand and assess the validity of system outputs and limits their agency to exercise control over how they are represented at work.

- **Actionability and perverse incentives:** When shared with end-users, different pieces of information might spur different actions and incentives, some of which might hinder rather than augment individuals, limiting their ability to create, share and make sense of knowledge. For example, simple proxies and metrics used by ranking algorithms such as the amount of edits or volume of documents authored by workers might incentivize those who don't get exposure to prioritize the quantity of work contributions, as opposed to their quality. Similarly, social signal such as views or likes might incentivize individuals to create content that's easily consumed, but which may only have short-lived value for them or the organization. Such adaptations of work, aimed at fitting into the system's extractive mechanisms, might lead to irrelevant associations being mined, in turn hindering the value of the system in the long run. As we think about what information to make visible, we should ask ourselves what behavior we wish to encourage when surfacing explanations, metrics or data sources in the UI. When designing the ways users can see themselves through the system, do we want to motivate them to

generate and share further contributions for knowledge extraction, contributing to system value and adoption? Do we want to help them reflect on their impact and past work? What are the behaviors we're nudging, and why are nudging those?

Other information, such as membership of demographic groups, could also be surfaced to help end-users understand system behavior. However, extreme care is needed when making these aspects transparent, as these may introduce new harms. As shown in Section 4.3.1, within an organization, self-identity might be shaped by the perceived membership to multiple groups at once, such as job role, seniority, as well as demographic segments defined by age, gender or location. Some of these groups might be historically marginalized groups and perceived as stigmatized. A vast volume of past psychological research has shown that beliefs of membership within different stigmatized groups and cognitions held by individuals about their stigma, including their awareness of stereotypes associated with their group, might result in unique combinations of discrimination and privilege (often referred to as intersectionality [48, 58]), and have important affective, motivational and interpersonal consequences [49, 109]. Making transparent how membership within a group influences system behavior could enforce stigma beliefs, leading to perverse incentives where employees alter their behavior to fit into group expectations, potentially affecting the well-being, motivation and sense of belonging in the workplace.

## 6.2 Procedural transparency

When making transparent how AI knowledge systems are embedded in a given organization, several dimensions related to system use in context come to play. These dimensions might fall outside the system boundaries, belonging to the *social*, rather than the *technical*, and might not be able to be tracked and summarized in a way that's actionable and valuable to end-users:

- **Tracking system use in context:** As the system automatically mines data and creates representations of individuals and their work, knowledge workers might care not only about how the system relates them to knowledge, but also about how system outputs are shown to others and used in practice. To help users achieve their goals, transparency needs to go beyond explaining the provenance of system outputs. It needs to make transparent what happens next: in which contexts are these inferences surfaced, to whom and how are these used in practice? Being transparent about how representations are used can give agency to individuals, allowing them to understand and control how they're being seen and how their work is presented to others. While tracking where representations are surfaced and to whom might be technically possible, the use of people's representations might not be explicitly documented and, hence, not within the scope of what can be made transparent. For example, whether a representation is actually being used to progress work, to establish a collaboration or to assess the work of others might not be tracked anywhere. From a symbolic interactionism perspective [27], the meaning ascribed to the associations being surfaced by AI knowledge systems, and therefore their use, might be a product of the social interactions happening outside of the scope of the system. Hence, being transparent on how representations are consumed and the norms surrounding system use in context requires the establishment of social processes in the organization, such as sharing clear guidelines for the ethical use of AI-generated representations in decision-making, developing organizational policies that prevent the misuse of these representations in certain contexts or implementing feedback channels for employees to raise concerns.
- **Curse of dimensionality:** As light passing through glass in infinite ways and creating different views *through*, people representations in AI knowledge systems might be surfaced in numerous ways, offering different views



depending on who is seeing and the context where seeing takes place. Being transparent about how people representations are presented to others and used by others opens a big space of possibilities for which information to reveal. As the number of users and contexts of use grows, the amount of information that can be made visible grows exponentially. Similarly to the exponential increase of computational complexity required in machine learning when data grows in dimensions (the curse of dimensionality), the multi-perspective nature of AI knowledge systems and the contextual variance in how extracted knowledge is (re-)presented and interpreted poses several challenges to summarizing the information in a faithful and accurate way. When summarizing and condensing information related to system use, we also need to preserve the privacy of those who see and those who are being seen. As we think about helping end-users understand how the system operates in context, communicating all the ways in which individuals are related to knowledge in the system and all the ways in which this information is surfaced to others in the organization might not be feasible nor helpful in practice, potentially leading to overwhelming amounts of information in detriment of focus and productivity.

- **Organizational power asymmetries:** Gausen et al. [66] highlight that system opacity can make it harder for workers and data subjects to meaningfully critique the system, further intensifying existing organizational power disparities. However, transparency alone is not a one-size-fits-all solution. In the backdrop of existing power asymmetries, transparency mechanisms themselves can be misused. For instance, making visible performance metrics can contribute to workplace surveillance [151], reinforcing power disparities and eroding workers' autonomy. Therefore, the design of any proposed transparency mechanism should carefully consider who is the subject and who is the consumer of the information that is being exposed, ensuring that transparency mechanisms do not disproportionately benefit those in positions of power, nor reduce workers to mere data points for organizational consumption [1].

### 6.3 Transparency of Outcomes

As we think about how AI knowledge systems could support transparency of outcomes, how to measure the impacts of using these systems in the long run involves several challenges:

- **Measuring and attributing actual outcomes:** While providers might be able to communicate intended outcomes and system goals, realized outcomes can be difficult to measure and attribute, being deeply entangled in social phenomena. For instance, changes in behavior at the individual level might respond to different incentives established within the organization as well as feedback received other than the information provided by the system. Similarly, how the system evolves over time might respond to new market needs and technological innovations, making it difficult to track those updates being influenced by changes in end-users' behavior and work-related data. Outcomes at the group level such as deviations in system performance across different demographic segments of users might be easier to track but difficult to communicate to individuals. Similarly to the negative downstream effects identified in system transparency, communicating that the system works worse for individuals belonging to a minority group could reinforce existing stigma beliefs in an organization, negatively affecting the well-being of knowledge workers.

## 7 DISCUSSION AND FUTURE DIRECTIONS

Transparency is multifaceted. As we think about which forms of transparency should AI knowledge systems support, several challenges and tradeoffs emerge. In what follows, we discuss future directions and areas of transparency research

in AI knowledge systems, encouraging CSCW researchers to further study transparency implications and ways to overcome implementation challenges.

### 7.1 The looking-glass metaphor

The looking-glass metaphor has been instrumental to our conversation about the reciprocal relationship between the *social* and the *technical* in AI knowledge systems. Whereas the closed-box metaphor directs the focus towards the system and the technical aspects that can be exposed to end-users, the looking-glass metaphor brings to the foreground the people and the social processes surrounding system use.

Considering how people become part of a system that both distorts and reflects back, and how this affects both individuals and collectives, is a vital part of understanding the personal and social impacts of AI systems. By aiding these reflections, the looking-glass metaphor has enabled our discussion about transparency to be about not just what the system does and how it does this, but also about how the system represents people and, in turn, how it causes people to see each other and themselves *through* algorithmic lenses. Similar to Plato’s allegory of the cave where prisoners mistake shadows for reality [47], the looking-glass metaphor calls attention to the risk of creating distorted perceptions of reality, where what we believe to be accurate may in fact be a skewed version of the truth. In the case of AI knowledge systems, the looking-glass metaphor can support wider discussions on how to design AI knowledge systems that augment reality in meaningful ways, rather than create distorted virtual realities that disturb knowledge relationality [66].

We encourage others to apply the looking-glass metaphor in other sociotechnical contexts such as social media or personalization systems and to develop reflexive questions such as the ones listed in Table 1, to inform further exploration of transparency requirements across different dimensions while bringing attention to both the *social* and the *technical* aspects surrounding system use.

### 7.2 System transparency

How system transparency gets implemented matters. How information is chosen to be shared might shape in different ways the types of seeing enabled by AI systems, potentially leading to unforeseen harms. As we think about what information to make visible to help end-users understand how AI systems work, we warn against making information visible for the sake of improving system transparency or claiming the system is made “transparent”. We particularly caution practitioners and researchers against making visible algorithmic representations of people that may be incomplete and out of context in a way that might presume ground truth, such as representations curated by personalization systems [40, 135], without thoroughly analyzing the implications of doing so. We recommend paying special attention to the language used when communicating algorithmic inferences, using specific labels to avoid misinterpretations and avoiding making value judgements or categorizations about end-users that might be flawed. Recognizing that these systems can shape how people see themselves and how they see others, we prompt the AI community to think through and further research the effects of transparency in sociotechnical contexts, carefully assessing the ramifications of algorithmic feedback on both self-concept beliefs and the ways people perceive individuals and collectives through algorithmic lenses.

### 7.3 Procedural transparency

In theory, AI knowledge systems can also serve the needs of leaders in organizations who need to make decisions on how to best allocate their resources for enhanced productivity and competitiveness in their market. Knowing areas of expertise

or work within different teams and identifying overlaps within organizations becomes critical for decision making on how to best utilize and capitalize the knowledge within an enterprise context, helping global enterprises take advantage of economies of scale while minimizing the likelihood of diseconomies of scale in growing and expanding organizations [10, 13, 44]. In this context, AI knowledge systems might be misused in unforeseen ways, leading to interpersonal harms and loss of agency for workers, while instilling feelings of surveillance, damaging worker’s well-being, productivity and enjoyment of work. In misusing the system, power dynamics in an organization might lead to withholding of information on the actual uses and norms surrounding the adoption of knowledge management technology. From a provider perspective, transparency via communication of intended and non intended uses isn’t enough to prevent misuse of technology. Rather, specific guardrails need to be in place to prevent product misuse and guarantee that both workers’ and organizational goals are prioritized, such as removing reporting of knowledge contributions and people analytics capabilities. System designers can also be challenged to provide power-aware transparency mechanisms that allow “studying up” [117] to reduce inequities of organizational power over time. A combination of guardrails with explicit communication is also needed to prevent individual harms at scale, informing end-users of the system about intended uses of the technology as well as any potential capabilities, if any, given to other members of the organization (such as administrators).

#### 7.4 Transparency of Outcomes

Even when actual outcomes might be difficult to measure, when individuals are confronted with algorithmic outcomes that are unfair and harmful, informed recourse should be provided so that those who feel adversely affected by the system know how to remediate potential unfair consequences. For example, when individuals believe system outputs are incorrect, impacting their social capital at work or generating increased labor to be reversed, clear documentation of ways to remediate potential negative outcomes should be made available. These recourse processes should be actionable and allow individuals to act on the information provided, giving them the agency to understand and control how to best present themselves through the system.

### 8 LIMITATIONS

The analysis of transparency requirements and impacts presented in this paper is reflective, value-laden and necessarily incomplete. Its aim is to encourage practitioners and researchers to critically reflect on the implications of AI transparency in sensitive contexts, rather than offering a complete account of transparency needs and potential challenges in AI knowledge systems. As foreseeing impacts that haven’t been observed before is inherently a difficult task [32], we introduced the looking-glass metaphor to support counterfactual thinking and encourage thoughtful transparency designs in sociotechnical contexts. While the metaphor might be applicable to other AI systems that implicitly and/or explicitly represent people, researchers might need to adjust the narrative or use the metaphor in combination with other metaphors as required by their use case.

The analysis of transparency requirements is intended to inform further empirical research, serving as a base ground of hypotheses that require further evaluation. When access to users is limited, practitioners can use the questions outlined in Table 1 to start exploring the design space of transparency, focusing on how different forms of transparency can impact various ways of seeing. Similarly, the analysis of challenges presented in this paper is not exhaustive. While the analysis can help guide decision-making and inform the design of transparency mechanisms in AI knowledge systems, the challenges identified may not extend to other AI systems. Furthermore, the assessment of transparency

impacts in AI systems must be iterative and ongoing, evolving alongside system development and regularly updated as new insights emerge from user interactions and feedback.

## 9 CONCLUSION

In this paper, we've taken a first step into identifying transparency needs, implications and challenges in enterprise AI knowledge systems. We've proposed the looking-glass metaphor as a way to conceptualize AI knowledge systems, expanding our view on transparency requirements by considering the implications of exposing information when seeing *through* the system. Grounded in past research, in this paper we've formulated potential effects on self-concept beliefs and a set of reflexive questions that require further empirical and user research in work settings. As we take a step back and zoom out, understanding AI systems as systems that reflect and distort puts into light their role in the economics of information and in shaping information asymmetries. With the rise of foundation models and continuous development of AI technology, we encourage cross-disciplinary research in AI transparency to help us step back and see the bigger picture to understand its societal implications. In doing so, economic views explaining the motivations of companies creating and consuming AI knowledge systems [147], a perspective less often sought in CSCW, might provide powerful ways to frame and put into perspective different transparency challenges and tensions arising in sociotechnical contexts. As AI systems shape new futures of work, we believe AI transparency will be critical to ensure these systems augment human labor in a way that is meaningful and equitable and we hope our discussion sparks further research.

## ACKNOWLEDGMENTS

## REFERENCES

- [1] 2020. The “Black Box” at Work. *Big Data and Society* 7 (2020). Issue 2. <https://doi.org/10.1177/2053951720938093>
- [2] Mohd Hafizul Afifi Abdullah, Norshakirah Aziz, Said Jadid Abdulkadir, Hitham Seddig Alhassan Alhussian, and Noureen Talpur. 2023. Systematic literature review of information extraction from textual data: recent methods, applications, trends, and challenges. *IEEE Access* 11 (2023), 10535–10562.
- [3] Mark S Ackerman. 2000. The intellectual challenge of CSCW: the gap between social requirements and technical feasibility. *Human-Computer Interaction* 15, 2-3 (2000), 179–203.
- [4] 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)* 22 (2013), 531–573.
- [5] Amazon. 2024. *Amazon Q*. Retrieved January 15, 2024 from <https://aws.amazon.com/q/>
- [6] Maddalena Amendola, Carlos Castillo, Andrea Passarella, and Raffaele Perego. 2024. Understanding and Addressing Gender Bias in Expert Finding Task. *arXiv preprint arXiv:2407.05335* (2024).
- [7] Dario Amodi, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. 2016. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565* (2016).
- [8] Mike Ananny and Kate Crawford. 2018. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *new media & society* 20, 3 (2018), 973–989.
- [9] Gloria Andrada, Robert W Clowes, and Paul R Smart. 2023. Varieties of transparency: Exploring agency within AI systems. *AI & society* 38, 4 (2023), 1321–1331.
- [10] Tatiana Andreeva and Aino Kianto. 2012. Does knowledge management really matter? Linking knowledge management practices, competitiveness and economic performance. *Journal of knowledge management* 16, 4 (2012), 617–636.
- [11] Jacy Anthis, Kristian Lum, Michael Ekstrand, Avi Feller, Alexander D'Amour, and Chenhao Tan. 2024. The Impossibility of Fair LLMs. *arXiv preprint arXiv:2406.03198* (2024).
- [12] Amelia Devi Putri Ariyanto, Diana Purwitasari, and Chastine Fatichah. 2024. A Systematic Review on Semantic Role Labeling for Information Extraction in Low-Resource Data. *IEEE Access* (2024).
- [13] Kenneth J Arrow. 1974. *The limits of organization*. WW Norton & Company.
- [14] Atlassian. 2024. *Atlassian Intelligence*. Retrieved January 15, 2024 from <https://www.atlassian.com/software/artificial-intelligence>
- [15] Krisztian Balog, Maarten De Rijke, et al. 2007. Determining Expert Profiles (With an Application to Expert Finding). In *IJCAI*, Vol. 7. 2657–2662.
- [16] Albert Bandura. 1977. Self-efficacy: toward a unifying theory of behavioral change. *Psychological review* 84, 2 (1977), 191.

- [17] Natã M Barbosa, Gang Wang, Blase Ur, and Yang Wang. 2021. Who Am I? A Design Probe Exploring Real-Time Transparency about Online and Offline User Profiling Underlying Targeted Ads. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 3 (2021), 1–32.
- [18] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2023. *Fairness and machine learning: Limitations and opportunities*. MIT press.
- [19] Solon Barocas and Andrew D Selbst. 2016. Big data's disparate impact. *Calif. L. Rev.* 104 (2016), 671.
- [20] Irma Becerra-Fernandez. 2000. The role of artificial intelligence technologies in the implementation of people-finder knowledge management systems. *Knowledge-Based Systems* 13, 5 (2000), 315–320.
- [21] Günter Beham, Barbara Kump, Tobias Ley, and Stefanie Lindstaedt. 2010. Recommending knowledgeable people in a work-integrated learning system. *Procedia Computer Science* 1, 2 (2010), 2783–2792.
- [22] Arne Bewersdorff, Xiaoming Zhai, Jessica Roberts, and Claudia Nerdel. 2023. Myths, mis- and preconceptions of artificial intelligence: A review of the literature. *Computers and Education: Artificial Intelligence* 4 (2023), 100143.
- [23] Aparajita Bhandari and Sara Bimo. 2022. Why's everyone on TikTok now? The algorithmized self and the future of self-making on social media. *Social Media+ Society* 8, 1 (2022), 20563051221086241.
- [24] Reuben Binns. 2018. Fairness in machine learning: Lessons from political philosophy. In *Conference on fairness, accountability and transparency*. PMLR, 149–159.
- [25] Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (technology) is power: A critical survey of "bias" in nlp. *arXiv preprint arXiv:2005.14050* (2020).
- [26] Bloomfire. 2024. *Bloomfire platform*. Retrieved January 15, 2024 from <https://bloomfire.com/platform/>
- [27] Herbert Blumer. 1986. *Symbolic interactionism: Perspective and method*. Univ of California Press.
- [28] Susanne Bodker. 2021. *Through the interface: A human activity approach to user interface design*. CRC Press.
- [29] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* (2021).
- [30] Philip Bonhard and Martina Angela Sasse. 2006. 'Knowing me, knowing you'—Using profiles and social networking to improve recommender systems. *BT Technology Journal* 24, 3 (2006), 84–98.
- [31] Geoffrey C Bowker and Susan Leigh Star. 2000. *Sorting things out: Classification and its consequences*. MIT press.
- [32] Margarita Boyarskaya, Alexandra Olteanu, and Kate Crawford. 2020. Overcoming failures of imagination in AI infused system development and deployment. *arXiv preprint arXiv:2011.13416* (2020).
- [33] Lewis Carrol. 2006. *Through the looking glass*. Read Books.
- [34] Carlos Castillo. 2019. Fairness and transparency in ranking. In *Acm sigir forum*, Vol. 52. ACM New York, NY, USA, 64–71.
- [35] Özlem Çetinoğlu, Sarah Schulz, and Ngoc Thang Vu. 2016. Challenges of computational processing of code-switching. *arXiv preprint arXiv:1610.02213* (2016).
- [36] Junyi Chai, Yujie He, Homa Hashemi, Bing Li, Daraksha Parveen, Ranganath Kondapally, and Wenjin Xu. 2021. Automatic Construction of Enterprise Knowledge Base. *arXiv preprint arXiv:2106.15085* (2021).
- [37] Matthew Chalmers. 2001. Information awareness and representation. *Computer Supported Cooperative Work (CSCW). An International Journal* (2001).
- [38] Matthew Chalmers and Areti Galani. 2004. Seamliness interweaving: heterogeneity in the theory and design of interactive systems. In *Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques*. 243–252.
- [39] Yoonseo Choi, Eun Jeong Kang, Min Kyung Lee, and Juho Kim. 2023. Creator-friendly Algorithms: Behaviors, Challenges, and Design Opportunities in Algorithmic Platforms. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–22.
- [40] Konstantina Christakopoulou, Alberto Lalama, Cj Adams, Iris Qu, Yifat Amir, Samer Chucuri, Pierce Vollucci, Fabio Soldo, Dina Beseio, Sarah Scodel, et al. 2023. Large Language Models for User Interest Journeys. *arXiv preprint arXiv:2305.15498* (2023).
- [41] Charles H Christiansen. 1999. Defining lives: Occupation as identity: An essay on competence, coherence, and the creation of meaning. *The American Journal of Occupational Therapy* 53, 6 (1999), 547–558.
- [42] Daniela Barreiro Claro, Marlo Souza, Clarissa Castellã Xavier, and Leandro Oliveira. 2019. Multilingual open information extraction: Challenges and opportunities. *Information* 10, 7 (2019), 228.
- [43] Robert William Clowes. 2020. The internet extended person. *Límité. Revista Interdisciplinaria de Filosofía y Psicología* (2020), 1–23.
- [44] Ronald Harry Coase. 1995. *The nature of the firm*. Springer.
- [45] Charles Horton Cooley. 1902. Looking-glass self. *The production of reality: Essays and readings on social interaction* 6 (1902), 126–128.
- [46] Eric Corbett and Emily Denton. 2023. Interrogating the T in FAacT. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 1624–1634.
- [47] Francis Macdonald Cornford et al. 1976. *The republic of Plato*. CUP Archive.
- [48] Kimberlé W Crenshaw. 2017. *On intersectionality: Essential writings*. The New Press.
- [49] Jennifer Crocker and Brenda Major. 1989. Social stigma and self-esteem: The self-protective properties of stigma. *Psychological review* 96, 4 (1989), 608.
- [50] Sunhao Dai, Chen Xu, Shicheng Xu, Liang Pang, Zhenhua Dong, and Jun Xu. 2024. Bias and Unfairness in Information Retrieval Systems: New Challenges in the LLM Era. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 6437–6447.

- [51] Thomas H Davenport and Laurence Prusak. 1998. *Working knowledge: How organizations manage what they know*. Harvard Business Press.
- [52] Shumin Deng, Yubo Ma, Ningyu Zhang, Yixin Cao, and Bryan Hooi. 2022. Information Extraction in Low-Resource Scenarios: Survey and Perspective. *arXiv preprint arXiv:2202.08063* (2022).
- [53] Michael A DeVito, Jeremy Birnholtz, Jeffery T Hancock, Megan French, and Sunny Liu. 2018. How people form folk theories of social media feeds and what it means for how we study self-presentation. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–12.
- [54] Michael A DeVito, Darren Gergle, and Jeremy Birnholtz. 2017. " Algorithms ruin everything" # RIPTwitter, Folk Theories, and Resistance to Algorithmic Change in Social Media. In *Proceedings of the 2017 CHI conference on human factors in computing systems*. 3163–3174.
- [55] Natalia Díaz-Rodríguez, Javier Del Ser, Mark Coeckelbergh, Marcos López de Prado, Enrique Herrera-Viedma, and Francisco Herrera. 2023. Connecting the dots in trustworthy Artificial Intelligence: From AI principles, ethics, and key requirements to responsible AI systems and regulation. *Information Fusion* 99 (2023), 101896.
- [56] Upol Ehsan, Philipp Wintersberger, Q Vera Liao, Elizabeth Anne Watkins, Carina Manger, Hal Daumé III, Andreas Riener, and Mark O Riedl. 2022. Human-Centered Explainable AI (HCXAI): beyond opening the black-box of AI. In *CHI conference on human factors in computing systems extended abstracts*. 1–7.
- [57] Malin Eiband, Hanna Schneider, Mark Bilandzic, Julian Fazekas-Con, Mareike Haug, and Heinrich Hussmann. 2018. Bringing transparency design into practice. In *Proceedings of the 23rd International Conference on Intelligent User Interfaces*. 211–223.
- [58] Sheena Erete, Yolanda A Rankin, and Jakita O Thomas. 2021. I can't breathe: Reflections from Black women in CSCW and HCI. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW3 (2021), 1–23.
- [59] Thomas Erickson and Wendy A Kellogg. 2000. Social translucence: an approach to designing systems that support social processes. *ACM transactions on computer-human interaction (TOCHI)* 7, 1 (2000), 59–83.
- [60] Motahhare Eslami, Karrie Karahalios, Christian Sandvig, Kristen Vaccaro, Aimee Rickman, Kevin Hamilton, and Alex Kirlik. 2016. First I " like" it, then I hide it: Folk Theories of Social Feeds. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 2371–2382.
- [61] Francesco Fabbri, Maria Luisa Croci, Francesco Bonchi, and Carlos Castillo. 2022. Exposure inequality in people recommender systems: the long-term effects. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 16. 194–204.
- [62] Maryam Fazel-Zarandi and Mark S Fox. 2011. Constructing expert profiles over time for skills management and expert finding. In *Proceedings of the 11th International Conference on Knowledge Management and Knowledge Technologies*. 1–6.
- [63] Antonio Ferrara, Lisette Espin-Noboa, Fariba Karimi, and Claudia Wagner. 2022. Link recommendations: Their impact on network structure and minorities. In *Proceedings of the 14th ACM Web Science Conference 2022*. 228–238.
- [64] Megan Rebecca French. 2018. *Algorithmic mirrors: An examination of how personalized recommendations can shape self-perceptions and reinforce gender stereotypes*. Stanford University.
- [65] Batya Friedman and Peter H Kahn Jr. 2007. Human values, ethics, and design. In *The human-computer interaction handbook*. CRC press, 1267–1292.
- [66] Anna Gausen, Bhaskar Mitra, and Siân Lindley. 2023. A Framework for Exploring the Consequences of AI-Mediated Enterprise Knowledge Access and Identifying Risks to Workers. *arXiv preprint arXiv:2312.10076* (2023).
- [67] Viktor Gecas and Michael L Schwalbe. 1983. Beyond the looking-glass self: Social structure and efficacy-based self-esteem. *Social psychology quarterly* (1983), 77–88.
- [68] A Gilbert. 2023. Reframing Automation: a new model for anticipating risks and impacts. *Institute for the Future of Work* (2023).
- [69] Glean. 2024. *Glean*. Retrieved January 15, 2024 from <https://www.glean.com/product/workplace-search-ai>
- [70] Erving Goffman. 2009. *Stigma: Notes on the management of spoiled identity*. Simon and schuster.
- [71] Erving Goffman et al. 2002. The presentation of self in everyday life. 1959. *Garden City, NY* 259 (2002).
- [72] Rodrigo Gonçalves and Carina Friedrich Dorneles. 2019. Automated expertise retrieval: a taxonomy-based survey and open issues. *ACM Computing Surveys (CSUR)* 52, 5 (2019), 1–30.
- [73] Amy L Gonzales and Jeffrey T Hancock. 2008. Identity shift in computer-mediated environments. *Media Psychology* 11, 2 (2008), 167–185.
- [74] Hanneke Grutterink and Alyson Meister. 2022. Thinking of you thinking of me: An integrative review of meta-perception in the workplace. *Journal of organizational behavior* 43, 2 (2022), 327–341.
- [75] Hanneke Grutterink, Gerben S Van der Vegt, Eric Molleman, and Karen A Jehn. 2013. Reciprocal expertise affirmation and shared expertise perceptions in work teams: Their implications for coordinated action and team performance. *Applied Psychology* 62, 3 (2013), 359–381.
- [76] Ido Guy, Inbal Ronen, and Eric Wilcox. 2009. Do you know? Recommending people to invite into your social network. In *Proceedings of the 14th international conference on Intelligent user interfaces*. 77–86.
- [77] Shuguang Han, Daqing He, Jiepu Jiang, and Zhen Yue. 2013. Supporting exploratory people search: a study of factor transparency and user control. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*. 449–458.
- [78] Michael A Hedderich, Lukas Lange, Heike Adel, Jannik Strötgen, and Dietrich Klakow. 2020. A survey on recent approaches for natural language processing in low-resource scenarios. *arXiv preprint arXiv:2010.12309* (2020).
- [79] Aaron Hess. 2014. You Are What You Compute (and What is Computed For You): Considerations of Digital Rhetorical Identification. *Journal of Contemporary Rhetoric* 4 (2014).
- [80] Frederik Hogenboom, Flavius Frasinca, Uzay Kaymak, Franciska De Jong, and Emiel Caron. 2016. A survey of event extraction methods from text for decision support systems. *Decision Support Systems* 85 (2016), 12–22.



- [81] Muhammad Humayoun, Harald Hammarström, and Aarne Ranta. 2022. Urdu morphology, orthography and lexicon extraction. *arXiv preprint arXiv:2204.03071* (2022).
- [82] Declan Humphreys, Abigail Koay, Dennis Desmond, and Erica Mealy. 2024. AI hype as a cyber security risk: the moral responsibility of implementing generative AI in business. *AI and Ethics* (2024), 1–14.
- [83] Omayma Husain, Naomie Salim, Rose Alinda Alias, Samah Abdelsalam, and Alzubair Hassan. 2019. Expert finding systems: A systematic review. *Applied Sciences* 9, 20 (2019), 4250.
- [84] Sarah Inman and David Ribes. 2019. "Beautiful Seams" Strategic Revelations and Concealments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [85] Claudiu Gabriel Ionescu and Monica Licu. 2023. Are TikTok Algorithms Influencing Users' Self-Perceived Identities and Personal Values? A Mini Review. *Social Sciences* 12, 8 (2023), 465.
- [86] William James. 1984. *Psychology; briefer course*. Vol. 14. Harvard University Press.
- [87] Mohammad Hossein Jarrahi, David Askay, Ali Eshraghi, and Preston Smith. 2023. Artificial intelligence and knowledge management: A partnership between human and AI. *Business Horizons* 66, 1 (2023), 87–99.
- [88] Heng Ji and Ralph Grishman. 2011. Knowledge base population: Successful approaches and challenges. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*. 1148–1158.
- [89] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *Comput. Surveys* 55, 12 (2023), 1–38.
- [90] Anna Jobin, Marcello Ienca, and Effy Vayena. 2019. The global landscape of AI ethics guidelines. *Nature machine intelligence* 1, 9 (2019), 389–399.
- [91] Adam N Joinson, Ulf-Dietrich Reips, Tom Buchanan, and Carina B Paine Schofield. 2010. Privacy, trust, and self-disclosure online. *Human-Computer Interaction* 25, 1 (2010), 1–24.
- [92] Alexandros Karatzoglou, Linas Baltrunas, and Yue Shi. 2013. Learning to rank for recommender systems. In *Proceedings of the 7th ACM Conference on Recommender Systems*. 493–494.
- [93] Krishnaram Kenthapadi, Himabindu Lakkaraju, and Nazneen Rajani. 2023. Generative ai meets responsible ai: Practical challenges and opportunities. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 5805–5806.
- [94] Pranav Khadpe, Ranjay Krishna, Li Fei-Fei, Jeffrey T Hancock, and Michael S Bernstein. 2020. Conceptual metaphors impact perceptions of human-ai collaboration. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–26.
- [95] George Lakoff and Mark Johnson. 2008. *Metaphors we live by*. University of Chicago press.
- [96] Ida Larsen-Ledet and Siân Lindley. 2022. Ways of seeing and being seen: People in the algorithmic knowledge base. In *Workshop at the 20th Eur. Conf. Comput.-Supported Cooperative Work*.
- [97] 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. *arXiv preprint arXiv:2209.03819* (2022).
- [98] Stefan Larsson and Fredrik Heintz. 2020. Transparency in artificial intelligence. *Internet Policy Review* 9, 2 (2020).
- [99] Angela Y Lee, Hannah Mieczkowski, Nicole B Ellison, and Jeffrey T Hancock. 2022. The algorithmic crystal: Conceptualizing the self through algorithmic personalization on TikTok. *Proceedings of the ACM on Human-computer Interaction* 6, CSCW2 (2022), 1–22.
- [100] Min Kyung Lee, Daniel Kusbit, Evan Metsky, and Laura Dabbish. 2015. Working with machines: The impact of algorithmic and data-driven management on human workers. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 1603–1612.
- [101] Nancy H Leonard, Laura L Beauvais, and Richard W Scholl. 1999. Work motivation: The incorporation of self-concept-based processes. *Human relations* 52 (1999), 969–998.
- [102] Jing Li, Aixin Sun, Jianglei Han, and Chenliang Li. 2020. A survey on deep learning for named entity recognition. *IEEE transactions on knowledge and data engineering* 34, 1 (2020), 50–70.
- [103] Zixuan Li, Yutao Zeng, Yuxin Zuo, Weicheng Ren, Wenxuan Liu, Miao Su, Yucan Guo, Yantao Liu, Xiang Li, Zhilei Hu, et al. 2024. KnowCoder: Coding Structured Knowledge into LLMs for Universal Information Extraction. *arXiv preprint arXiv:2403.07969* (2024).
- [104] Q Vera Liao and Jennifer Wortman Vaughan. 2023. AI Transparency in the Age of LLMs: A Human-Centered Research Roadmap. *arXiv preprint arXiv:2306.01941* (2023).
- [105] Jay Liebowitz. 2001. Knowledge management and its link to artificial intelligence. *Expert systems with applications* 20, 1 (2001), 1–6.
- [106] Shuyi Lin, Wenxing Hong, Dingding Wang, and Tao Li. 2017. A survey on expert finding techniques. *Journal of Intelligent Information Systems* 49 (2017), 255–279.
- [107] Siân E Lindley and Denise J Wilkins. 2023. Building Knowledge through Action: Considerations for Machine Learning in the Workplace. *ACM Transactions on Computer-Human Interaction* (2023).
- [108] Nick Logler, Daisy Yoo, and Batya Friedman. 2018. Metaphor cards: A how-to-guide for making and using a generative metaphorical design toolkit. In *Proceedings of the 2018 designing interactive systems conference*. 1373–1386.
- [109] Brenda Major and Jennifer Crocker. 1993. Social stigma: The consequences of attributional ambiguity. In *Affect, cognition and stereotyping*. Elsevier, 345–370.
- [110] Márcio José Mantau and Fabiane Barreto Vavassori Benitti. 2022. Awareness support in collaborative system: Reviewing last 10 years of cscw research. In *2022 IEEE 25th International Conference on Computer Supported Cooperative Work in Design (CSCWD)*. IEEE, 564–569.

- [111] Catherine C Marshall and Siân E Lindley. 2014. Searching for myself: motivations and strategies for self-search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 3675–3684.
- [112] Alice E Marwick and Danah Boyd. 2011. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New media & society* 13, 1 (2011), 114–133.
- [113] Stefanie Mauksch, Heiko A von der Gracht, and Theodore J Gordon. 2020. Who is an expert for foresight? A review of identification methods. *Technological Forecasting and Social Change* 154 (2020), 119982.
- [114] George Herbert Mead et al. 1934. *Mind, self, and society*. Vol. 111. University of Chicago press Chicago.
- [115] Alyson Meister, Emma Zhao, Carol Gill, Karen Jehn, and Amanda Sinclair. 2023. How feeling misidentified can drive negative attitudes yet increase performance: The role of appraisals. *Journal of Applied Social Psychology* (2023).
- [116] Jacob Metcalf, Emanuel Moss, Elizabeth Anne Watkins, Ranjit Singh, and Madeleine Clare Elish. 2021. Algorithmic impact assessments and accountability: The co-construction of impacts. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. 735–746.
- [117] Milagros Miceli, Julian Posada, and Tianling Yang. 2022. Studying up machine learning data: Why talk about bias when we mean power? *Proceedings of the ACM on Human-Computer Interaction* 6, GROUP (2022), 1–14.
- [118] Microsoft. 2023. *Microsoft Viva*. Retrieved October 18, 2023 from <https://www.microsoft.com/en-us/microsoft-viva>
- [119] Microsoft. 2023. *Topics in Viva Connections*. Retrieved January 15, 2024 from <https://learn.microsoft.com/en-us/microsoft-365/topics/topics-card-viva-connections?view=o365-worldwide>
- [120] Microsoft. 2023. *Viva Topics*. Retrieved October 18, 2023 from <https://learn.microsoft.com/en-us/microsoft-365/topics/topic-experiences-overview?view=o365-worldwide>
- [121] Microsoft. 2024. *Ask questions and get answers using Microsoft 365 Copilot Business Chat*. Retrieved September 19, 2024 from <https://support.microsoft.com/en-us/topic/ask-questions-and-get-answers-using-microsoft-365-copilot-business-chat-fd8d88af-9492-48cd-8385-7e8615b42d80>
- [122] Microsoft. 2024. *Microsoft Copilot 365*. Retrieved January 15, 2024 from [https://www.microsoft.com/en-us/microsoft-365/enterprise/copilot-for-microsoft-365#tabs-pill-bar-oca455\\_tab1](https://www.microsoft.com/en-us/microsoft-365/enterprise/copilot-for-microsoft-365#tabs-pill-bar-oca455_tab1)
- [123] Microsoft. 2024. *Microsoft Copilot for Work*. Retrieved January 15, 2024 from [https://www.microsoft.com/en-us/microsoft-365/copilot-for-work#tabs-pill-bar-oc58d2\\_tab5](https://www.microsoft.com/en-us/microsoft-365/copilot-for-work#tabs-pill-bar-oc58d2_tab5)
- [124] Dave Murray-Rust, Iohanna Nicenboim, and Dan Lockton. 2022. Metaphors for designers working with AI. (2022).
- [125] Devesh Narayanan. 2023. Welfarist Moral Grounding for Transparent AI. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 64–76.
- [126] Ikujiro Nonaka. 1994. A dynamic theory of organizational knowledge creation. *Organization science* 5, 1 (1994), 14–37.
- [127] Judith S Olson, Jonathan Grudin, and Eric Horvitz. 2005. A study of preferences for sharing and privacy. In *CHI'05 extended abstracts on Human factors in computing systems*. 1985–1988.
- [128] Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Emre Kiciman. 2019. Social data: Biases, methodological pitfalls, and ethical boundaries. *Frontiers in big data* 2 (2019), 13.
- [129] Alexandra Olteanu, Michael Ekstrand, Carlos Castillo, and Jina Suh. 2023. Responsible AI Research Needs Impact Statements Too. *arXiv preprint arXiv:2311.11776* (2023).
- [130] Wanda J Orlikowski. 2002. Knowing in practice: Enacting a collective capability in distributed organizing. *Organization science* 13, 3 (2002), 249–273.
- [131] Leysia Palen and Paul Dourish. 2003. Unpacking "privacy" for a networked world. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 129–136.
- [132] Hyanghee Park, Daehwan Ahn, Kartik Hosanagar, and Joonhwan Lee. 2022. Designing fair AI in human resource management: Understanding tensions surrounding algorithmic evaluation and envisioning stakeholder-centered solutions. In *Proceedings of the 2022 CHI conference on human factors in computing systems*. 1–22.
- [133] Sachin Pawar, Girish K Palshikar, and Pushpak Bhattacharyya. 2017. Relation extraction: A survey. *arXiv preprint arXiv:1712.05191* (2017).
- [134] Marta Pérez-Verdugo. 2022. Situating Transparency: An Extended Cognition Approach. *Teorema* 41, 3 (2022).
- [135] Filip Radlinski, Krisztian Balog, Fernando Diaz, Lucas Dixon, and Ben Wedin. 2022. On Natural Language User Profiles for Transparent and Scrutable Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2863–2874.
- [136] Dillon Reisman, Jason Schultz, Kate Crawford, and Meredith Whittaker. 2018. Algorithmic impact assessments: a practical Framework for Public Agency. *AI Now* 9 (2018).
- [137] Advait Sarkar. 2023. Enough with "human-AI collaboration". In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–8.
- [138] Advait Sarkar, XT Xu, N Toronto, I Drosos, and C Poelitz. 2024. When Copilot Becomes Autopilot: Generative AI's Critical Risk to Knowledge Work and a Critical Solution. *EuSpRIG Proceedings* (2024).
- [139] Laura Sartori and Giulia Bocca. 2023. Minding the gap (s): public perceptions of AI and socio-technical imaginaries. *AI & society* 38, 2 (2023), 443–458.
- [140] Andrew D Selbst. 2021. An institutional view of algorithmic impact assessments. *Harv. JL & Tech.* 35 (2021), 117.
- [141] Andrew D Selbst and Solon Barocas. 2018. The intuitive appeal of explainable machines. *Fordham L. Rev.* 87 (2018), 1085.

- [142] Abigail Sellen and Eric Horvitz. 2024. The rise of the AI Co-Pilot: Lessons for design from aviation and beyond. *Commun. ACM* 67, 7 (2024), 18–23.
- [143] Khaled Shaalan, Sanjeera Siddiqui, Manar Alkhatib, and Azza Abdel Monem. 2019. Challenges in Arabic natural language processing. In *Computational linguistics, speech and image processing for Arabic language*. World Scientific, 59–83.
- [144] Renee Shelby, Shalaleh Rismani, Kathryn Henne, AJung Moon, Negar Rostamzadeh, Paul Nicholas, N'Mah Yilla-Akbari, Jess Gallegos, Andrew Smart, Emilio Garcia, et al. 2023. Sociotechnical harms of algorithmic systems: Scoping a taxonomy for harm reduction. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*. 723–741.
- [145] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of exposure in rankings. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 2219–2228.
- [146] Irene Solaiman, Zeerak Talat, William Agnew, Lama Ahmad, Dylan Baker, Su Lin Blodgett, Canyu Chen, Hal Daumé III, Jesse Dodge, Isabella Duan, et al. 2023. Evaluating the social impact of generative ai systems in systems and society. *arXiv preprint arXiv:2306.05949* (2023).
- [147] Joseph E Stiglitz. 2000. The contributions of the economics of information to twentieth century economics. *The quarterly journal of economics* 115, 4 (2000), 1441–1478.
- [148] H Colleen Stuart, Laura Dabbish, Sara Kiesler, Peter Kinnaird, and Ruogu Kang. 2012. Social transparency in networked information exchange: a theoretical framework. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*. 451–460.
- [149] Christopher A Summers, Robert W Smith, and Rebecca Walker Reczek. 2016. An audience of one: Behaviorally targeted ads as implied social labels. *Journal of Consumer Research* 43, 1 (2016), 156–178.
- [150] Henri Tajfel. 1978. Social categorization, social identity and social comparison. *Differentiation between social group* (1978), 61–76.
- [151] Zephyr Teachout. 2022. The Boss Will See You Now. *The New York Review of Books* 69, 13 (2022).
- [152] John C. Thomas, Wendy A. Kellogg, and Thomas Erickson. 2001. The knowledge management puzzle: Human and social factors in knowledge management. *IBM systems journal* 40, 4 (2001), 863–884.
- [153] Dianne M Tice. 1992. Self-concept change and self-presentation: the looking glass self is also a magnifying glass. *Journal of personality and social psychology* 63, 3 (1992), 435.
- [154] SM Tonmoy, SM Zaman, Vinija Jain, Anku Rani, Vipula Rawte, Aman Chadha, and Amitava Das. 2024. A comprehensive survey of hallucination mitigation techniques in large language models. *arXiv preprint arXiv:2401.01313* (2024).
- [155] Sarah J Tracy and Angela Trethewey. 2005. Fracturing the real-self fake-self dichotomy: Moving toward “crystallized” organizational discourses and identities. *Communication theory* 15, 2 (2005), 168–195.
- [156] Chun-Hua Tsai and Peter Brusilovsky. 2018. Beyond the ranked list: User-driven exploration and diversification of social recommendation. In *23rd international conference on intelligent user interfaces*. 239–250.
- [157] Chun-Hua Tsai and Peter Brusilovsky. 2021. The effects of controllability and explainability in a social recommender system. *User Modeling and User-Adapted Interaction* 31 (2021), 591–627.
- [158] Elizabeth Anne Watkins, Emanuel Moss, Jacob Metcalf, Ranjit Singh, and Madeleine Clare Elish. 2021. Governing algorithmic systems with impact assessments: Six observations. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 1010–1022.
- [159] Denise Wilkins, Siân Lindley, Max Meijer, Richard Banks, and Britta Burlin. 2020. Designing AI systems that make organizational knowledge actionable. *Interactions* 27, 6 (2020), 72–75.
- [160] Genta Indra Winata, Alham Fikri Aji, Zheng-Xin Yong, and Thamar Solorio. 2022. The decades progress on code-switching research in nlp: A systematic survey on trends and challenges. *arXiv preprint arXiv:2212.09660* (2022).
- [161] John Winn, Matteo Venanzi, Tom Minka, Ivan Korostelev, John Guiver, Elena Pochernina, Pavel Mishkov, Alex Spengler, Denise Wilkins, Siân Lindley, et al. 2021. Enterprise Alexandria: Online high-precision enterprise knowledge base construction with typed entities. In *3rd Conference on Automated Knowledge Base Construction*.
- [162] Christine T Wolf. 2016. Seeing work: constructing visions of work in and through data. In *Proceedings of the 2016 ACM International Conference on Supporting Group Work*. 509–512.
- [163] Christine T Wolf. 2020. From knowledge graphs to knowledge practices: On the need for transparency and explainability in enterprise knowledge graph applications. In *Proceedings of the KG-BIAS Workshop 2020 at AKBC 2020*.
- [164] Mingman Xu, Qiong Chang, and Jun Miyazaki. 2023. How Does the System Perceive Me?—A Transparent and Tunable Recommender System. In *International Conference on Database and Expert Systems Applications*. Springer, 33–48.
- [165] Dawit Yimam-Seid and Alfred Kobsa. 2003. Expert-finding systems for organizations: Problem and domain analysis and the DEMOIR approach. *Journal of Organizational Computing and Electronic Commerce* 13, 1 (2003), 1–24.
- [166] Kenji Yoshino. 2006. The pressure to cover. *New York Times Magazine* 32 (2006).
- [167] Bowen Zhang and Harold Soh. 2024. Extract, Define, Canonicalize: An LLM-based Framework for Knowledge Graph Construction. *arXiv preprint arXiv:2404.03868* (2024).
- [168] Yongfeng Zhang, Xu Chen, et al. 2020. Explainable recommendation: A survey and new perspectives. *Foundations and Trends® in Information Retrieval* 14, 1 (2020), 1–101.
- [169] Lingfeng Zhong, Jia Wu, Qian Li, Hao Peng, and Xindong Wu. 2023. A comprehensive survey on automatic knowledge graph construction. *Comput. Surveys* 56, 4 (2023), 1–62.
- [170] Valentin Zieglmeier and Alexander Pretschner. 2023. Rethinking people analytics with inverse transparency by design. *arXiv preprint arXiv:2305.09813* (2023).

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009