

Lifting the curtain: Strategic visibility of human labour in AI-as-a-Service

Big Data & Society
January–June: 1–14
© The Author(s) 2021
DOI: 10.1177/20539517211016026
journals.sagepub.com/home/bds



Gemma Newlands 

Abstract

Artificial Intelligence-as-a-Service (AIaaS) empowers individuals and organisations to access AI on-demand, in either tailored or ‘off-the-shelf’ forms. However, institutional separation between development, training and deployment can lead to critical opacities, such as obscuring the level of human effort necessary to produce and train AI services. Information about how, where, and for whom AI services have been produced are valuable secrets, which vendors strategically disclose to clients depending on commercial interests. This article provides a critical analysis of how AIaaS vendors manipulate the visibility of human labour in AI production based on whether the vendor relies on paid or unpaid labour to fill interstitial gaps. Where vendors are able to occlude human labour in the organisational ‘backstage,’ such as in data preparation, validation or impersonation, they do so regularly, further contributing to ongoing techno-utopian narratives of AI hype. Yet, when vendors must co-produce the AI service with the client, such as through localised AI training, they must ‘lift the curtain’, resulting in a paradoxical situation of needing to both perpetuate dominant AI hype narratives while emphasising AI’s mundane limitations.

Keywords

AI-as-a-Service, artificial intelligence, commercial logics, hidden work, machine learning, transparency

Introduction

In the current zeitgeist of research into artificial intelligence (AI), academic interest has largely focused on AI as an abstract sociotechnical phenomenon, or as embedded in major platforms such as Facebook, Google or Twitter. However, this focus has overlooked the widespread growth of commercial AI-as-a-Service offerings (AIaaS), where organisations access specific AI capabilities via cloud computing (Casati et al., 2019; Parsaeefard et al., 2019; Tubaro et al., 2020). Enabling organisations to scale their AI usage based on strategic requirements, AIaaS can range from conversational bots to knowledge mapping, computer vision and speech recognition. While AIaaS adoption in organisations continues apace (Vesa and Tienari, 2020; von Krogh, 2018), a complex AI production ecosystem has emerged which connects vendors, clients, developers and data workers through both informal and formal ties.

Yet, AI production remains shrouded in what Costas and Grey (2014, 2016) refer to as organisational secrecy. Information about how, where, and for whom AI services have been produced are valuable secrets,

which vendors strategically disclose to clients depending on commercial interests. The societal imagination of AI is fuelled by fantasy more than reality, driven by media depictions of AI that can do everything (Giuliano, 2020; Selbst, 2017). As such, corporate sociotechnical imaginaries (Hockenhull and Cohn, 2021) construct understandings of how AI services ought to function effortlessly in order to trigger investment, foster innovation, and sell services (Pettersen, 2019).

Observing an industry-wide strategy to de-emphasise the human factor in AI production, Tubaro (2021) claims that, ‘corporate communication highlights the role of technology, not human contribution, especially in the AI industry’ (Tubaro, 2021: 13). However, accurate identification of where human subjectivity enters AI production is essential since engineers and AI production workers infuse AI systems

BI Norwegian Business School, Oslo, Norway

Corresponding author:

Gemma Newlands, BI Norwegian Business School, 0484 Oslo, Norway.
Email: gemma.e.newlands@bi.no



with values and normativities (Elish and Boyd, 2018; Grosman and Reigeluth, 2019). Moreover, such identification is necessary to safeguard the environmental and social sustainability of AI production (Villena and Gioia, 2018).

Drawing on Daniels' (1987) concept of invisible work as that which 'disappears from our observations and reckonings' (403), I offer a critical analysis on the human labour involved in AI production. Specifically, I explore how this labour is often rendered either invisible or hypervisible through a vendor's organisational practices (Rabelo and Mahalingam, 2019). I argue that where vendors are able to occlude human labour in the organisational 'backstage' (Goffman, 1959), such as in data preparation, validation or impersonation, they do so regularly, further contributing to ongoing technoutopian narratives of AI hype. Yet, when vendors must co-produce the AI service with the client, such as through localised AI training, they must 'lift the curtain', resulting in a paradoxical situation of needing to both perpetuate dominant AI hype narratives while emphasising AI's mundane limitations.

This vendor-centric approach thus advances ongoing discussions on labour in AI production, which have so far focused on micro-work or social media platforms (Bechmann and Bowker, 2019; Grosman and Reigeluth, 2019; Tubaro et al., 2020). Where possible, I have also drawn on examples from the Nordic AIaaS ecosystem, in order to highlight the efforts of regional AI development and to counterbalance the gravitational orbit of North American companies in academic research. Centering on the Global North also highlights ongoing global inequalities in AI production, where localised AI production is highlighted while remote production is often occluded.

The rest of the paper is structured as follows. After a contextualising review of both organisational secrecy and AIaaS, I will address three critical processes of strategic visibilisation. For the first process, I will detail how the AIaaS vendor performs strategic secrecy in the 'organisational backstage', obscuring the AI production labour upstream in the AI supply chain (Villena and Gioia, 2018). Indeed, the AI supply chain can be conceptualised as the chain of collection, curation and custody of data from source to model, passing through the hands of potentially infinite numbers of data workers, data brokers and data scientists on the way. Yet, mirroring opacities encountered even in supply chains for quotidian products such as bread, socks or pain medication, uncovering even the identity of upstream (data) suppliers can be impossible due to commercial secrecy and the complexity of global networks. Because this labour is purchased by the vendor or performed by intra-organisational members, the

vendor is able to obscure the details and render much of it invisible to the client.

For the second process, I explore how the necessity for AI co-production with organisational clients means that vendors must 'lift the curtain' and explain in clear terms how much mundane human effort is involved in AI production. Because this labour must be co-opted from the client organisation to make the AI function as promised, the vendor must sacrifice techno-hype for practical reality. Nevertheless, in this section I discuss how vendors are reluctant to dismiss techno-hype entirely, resulting in an often conflicting and paradoxical situation where AI services are presented to clients as simultaneously effortless and arduous.

For the third process, I briefly explore how clients can also engage in strategic visibilisation when using the AI service to interface with external customers. In AIaaS interview services or chatbots, the vendor makes the human involvement in training and operating the AIaaS visible to the client. However, the client can then render their own human labour invisible to the end customer, thus perpetuating and benefiting from techno-hype.

Organisational secrecy and invisible work

Organisational secrecy, defined by Costas and Grey (2014, 2016), is the deliberate concealment of information by actors within organisations. As both a process and a universal sociological form, organisational secrecy is neither good nor bad, even though most research to date on organisational secrecy has tended to view it negatively (Birchall, 2011). When considering how organisations keep certain information secret, Erving Goffman's (1959) dramaturgic theory of the 'frontstage' and 'backstage' offers a useful lens through which to view social actions and processes. In *The Presentation of the Self in Everyday Life*, Goffman (1959) describes how actors seek to give a favourable impression in line with current societal values. Within this framework, Goffman (1959) offers the concept of *strategic secrecy*, which refers to the protection of valuable information, as well as *dark secrecy*, which refers to the protection of shameful or dirty secrets.

Although Goffman's (1959) theory initially referred only to face-to-face interactions with co-present actors, research has extended it to online (Kilvington, 2020) and organisational settings (Manning, 2008; Whittle et al., 2020). As such, we can observe how organisations conceal processes and information in their 'backstage', while performing impression management to shape idealised self-presentations on the 'frontstage' (Ringel, 2019). The backstage, in this case, refers to the activities and processes conducted by the organisation, while the 'frontstage' refers to the self-impression

of the organisation to external stakeholders. As an example of this, Laguecir and Leca (2019) explain how the organisational backstage can be used to conceal misconduct. To quote Whittle et al. (2020), the ‘frontstage is conceptualised as the place where a more anodyne and sanitised version of the organisation is presented, a version that is detached from the messy and sometimes ‘dirty’ reality’ (5).

In the production of AI services, one element often kept strategically secret on the ‘backstage’ is the amount and type of labour required. Human labour, as described by Daniels (1987), is often rendered invisible by being overlooked, marginalised and under rewarded. Although Daniels’ (1987) research on invisible work is centred on domestic work, numerous scholars have built upon her foundations to offer a robust discourse around invisible and hidden work in an organisational context (Otis and Zhao, 2016). Hatton (2017), for instance, argues that work is made invisible through cultural, legal, and spatial mechanisms of invisibility, whereby sociospatial mechanisms devalue labour through physical separation between the worker and the observer. As with organisational secrecy, visibilisation is also a process which is imposed and normalised through social interaction and organisational practices (Rabelo and Mahalingam, 2019). In contrast to invisible labour, it is also important to acknowledge the potential for hypervisible labour, primarily referring to work that has an aesthetic component of its observation (Crain et al., 2016). Although hypervisible labour has so far focused only on aesthetically pleasing labour, such as expert cooking, it provides a valuable way of also describing labour, which is emphasised and highlighted for strategic purposes.

While visibility is complex concept, engaging multiple multi-directional stakeholders (Kaplan et al., 2020), (in)visible work must have a potential observer for whom the work is rendered invisible or visible. In almost every conceivable case, the worker will observe and acknowledge that they are working, that they are engaging in an act of production, even if they cannot see the end results of their labour. There will also always be epistemological gaps between the lived and observed experience of one’s own work with the observation of others (Newlands, 2020). However, making one’s work visible to external stakeholders offers opportunities for recognition, esteem subjectivity and motivation (Angella, 2016; Dashtipour and Vidaillet, 2017). By contrast, keeping work invisible to direct supervisors can offer workers discretion in their own labour process (Brighenti, 2007; Star and Strauss, 1999). Visibility is thus neither inherently good nor bad, with workers always facing trade-offs in making their work visible to different audiences (Timonen and Vuori, 2018).

Within a digital setting, there have been a number of discussions about how human labour has become less visible when mediated. For instance, critical attention has been paid to unpaid ‘digital labor’ (Scholz, 2013) and the ‘digi-housekeeping’ necessary to maintain flexible working arrangements (Whiting and Symon, 2020). Recent research has also discussed how much work becomes invisible in data intensive science, particularly referring to cleaning and archiving datasets (Plantin, 2019; Scroggins and Pasquetto, 2020). In this article, however, I will focus specifically on the labour which is made invisible or hypervisible in the production of AI services.

AI-as-a-Service

Drawing on the technical heritage of Software-as-a-Service, Platform-as-a-Service and Infrastructure-as-a-Service, many organisations are adopting AIaaS, accessing specific AI capabilities through the cloud (Casati et al., 2019; Parsaeefard et al., 2019; Tubaro et al., 2020). In the same way that cloud technologies have democratised access to software, data storage and data processing (Mosco, 2015), AIaaS empowers organisations to access AI on-demand, in either tailored or ‘off-the-shelf’ forms. In-house AI development is costly in terms of time and money, particularly given the scarcity of AI talent on the job market. By using AIaaS, individuals and organisations can scale their usage based on strategic requirements, fostering rapid experimentation and overall reduction of return-on-investment. Hailed as the next big wave in computing, AIaaS is currently offered by major vendors including Amazon Web Services, Microsoft Azure, Google Cloud, IBM Cloud, Oracle and Salesforce. AIaaS is also the key business model for a host of start-ups and small-medium enterprises (Metelskaia et al., 2018).

However, most AIaaS vendors operate within competitive innovation environments shaped by technoutopian narratives. Discourse around AI, as explained by Giuliano (2020), imbues AI with magical, quasi-religious notions. Promising to make life easier, to mimic human intelligence and surpass it, aspirational rhetorics around AI silence critics by promising that any shortfalls will be solved in the future (Florida, 2020). The sociotechnical imaginaries surrounding AI (Jasanoff and Kim, 2009) are multiple and powerful, though increasingly commodified by vendors to generate what has recently been termed the ‘corporate socio-technical imaginary’ (Hockenull and Cohn, 2021; Mager and Katzenbach, 2021). From this perspective, vendors and other industry actors generate visions of the future full of effective and ubiquitous AI services as a mechanism to generate increased funding (Mützel, 2021). Indeed, company valuations are

critically aligned with the technicality of their AI offering, specifically its scalability and ability to reduce labour costs.

Nevertheless, in order to identify and acknowledge the human labour involved in AI production, it is important to counterbalance discourse with materiality. Despite being accessed primarily through the cloud, it would be a fallacy to consider AI services as ‘immaterial media’ (Cubitt, 2017). Rather, data and technology have an extensive material footprint (Prainsack, 2019). The metaphor of the ‘cloud’ in cloud computing is often merely a fog that obscures the social and technical reality behind it (Wyatt, 2021).

Drawing on the ‘biography of artifacts’ approach, we can observe that AI services, as with all technical entities, have biographies, and different human hands are involved in their production (Williams and Pollock, 2012). Indeed, if we briefly unpack the concept of AI from an etymological standpoint, it is important to acknowledge the human craft inherent in the term. Artificial, derived from ‘artificium’ (a work of art) and ‘artifex’ (craftsman), directly acknowledges the human skill involved in the production. By treating AI as a form of digital artefact crafted by human hand, it is possible to consider its production through long-standing discourses around production, craft and labour; to position its creation as a critical topic of study.

Ongoing discussions regarding the future of work centre around whether jobs will be replaced by automation or AI, leading to widespread un- and underemployment (Acemoglu and Restrepo, 2018; Lloyd and Payne, 2019). Indeed, automation and AI technologies are regularly implemented in order to reduce the labour force, as discussed below with regard to chatbots and ‘digital employees’. However, an important counterargument is the continued reliance on hidden human work in the production, training and execution of AI (Bechmann and Bowker, 2019). Dataset development, for instance, has complex chains of human labour and, as Raisch and Krakowski (2021) describe, there is a close intertwinement of humans and machines in the production of AI. However, institutional separation between development, training, and deployment can lead to critical opacities, such as obscuring the types of human effort necessary to produce and train AI services. Separation between the AIaaS vendor and the client, particularly when the AI service is constructed by multiple parties with expertise gaps, means that important information can be distributed across organisations. For strategic reasons, the vendor may render this human involvement either invisible or hypervisible. Building on this, I will now address the three critical processes of strategic visibilisation.

The invisible backstage of AI production

AIaaS, particularly machine-learning based AIaaS, is data intensive and thus labour intensive (Casati et al., 2019; Tubaro and Casilli, 2019). Machine-learning models must be trained with accurate, relevant, and high-quality training datasets (von Krogh, 2018). Cognilytica (2019), for instance, estimates that data preparation tasks form over 80% of the time consumed in most AI and ML projects, with the market for third-party data labelling solutions reaching \$150 million in 2018, and expected to surpass \$1 billion by 2023. As Bechmann and Bowker (2019) explain, AI production requires layers of ‘data collection; data cleaning; data partitioning; model selection; model training (including tuning and assessment) and model deployment’ (1). Types of data preparation can include semantic segmentation, object detection, and image classification. While some of these tasks can be relatively basic, Schmidt (2019) has highlighted the increasing complexity of data tasks over time. However, datasets for AI do not emerge from thin air, nor from a vacuum untouched by power asymmetries and commercial interests. Since an AI model is only as good as the data that feeds it, a nascent body of research has studied the provenance of AI datasets, defined here as ‘information about the creation, chain of custody, modifications or influences’ (Cheney et al., 2009: 960).

Despite the crucial nature of this work for AI production, AIaaS vendors implement strategic secrecy to render the form and scope of data work invisible to clients. Such ‘invisibility’ means that it goes unacknowledged through lack of attribution or the blurring between human and technical achievement. Because this data is often generated externally from the AIaaS vendor, the work to collect and prepare it becomes even more removed from the client’s potential view. Researchers have identified the often haphazard methods of generating and using datasets for AI training, referring to a ‘laissez-faire’ attitude towards data collection which trades rigour for speed and accessibility (Jo and Gebru, 2020).

However, there is further lack of visibility about these workers because the vendor may not know the source of the data due to data brokerage (Crain, 2018). Data brokerage, referring to the business model of gathering and distributing data, is usually conceptualised as part of the advertising and customer-targeting ecosystem (Yeh, 2018). Although data brokers operate as a global infrastructure supporting the development of AI services, the complex nature of data marketplaces, where data points are repackaged and resold, make transparency about their practices difficult to ascertain. This is in addition to the understandable inclination of

data brokers to maintain trade secrecy about their data sources.

AI production operates in a global network, where AIaaS providers, the workforce and their customers are distributed worldwide. Yet, as with global outsourcing elsewhere, the labour-intensive activity of cleaning, curating and annotating datasets is outsourced to lower labour cost countries and thus outside of the vendor's direct observation. Data work in AI production is usually performed in low income countries and poorly compensated (Tubaro et al., 2020). Indeed, such offshoring, or 'botshoring', is deeply entwined in labour inequalities. Miceli et al. (2020), conducting fieldwork with data annotation companies in Argentina and Bulgaria, demonstrate how the work of data annotators is often fulfilled by poorly paid workers. Researchers have questioned whether these interstitial tasks, filling the gaps in AI production, will be left to low-skilled and low-income workers (Celentano, 2019; Tubaro and Casilli, 2019). Platform interfaces, between vendors and data workers, in particular enable 'entrepreneurs to imagine workers in a better place than they actually are' (Gruszka and Böhm, 2020: 4).

Artificial intelligence preparation

Needing increasing precision in data quality, AIaaS providers and users are moving to specialised platforms for AI data preparation such as Appen (Schmidt, 2017, 2019). Tubaro and Casilli (2019), for instance, use the automotive industry as a case study for hidden work in AI, pointing out how the production of data is labour intensive but often hidden through platforms. In this case, data preparation is conducted by individuals on dedicated AI micro-work sites in a distributed fashion.

Yet, considerable research to date has focused on the role of generalised crowdworking sites as hubs for AI production (Altenreid, 2020). Gray and Suri (2019), for example, highlight in *Ghost Work* how humans are required to label mass amounts of data on Amazon Mechanical Turk. They argue that the platform design is engineered to anonymise workers, making their individual contributions invisible within the total community. Moreover, research has showed that crowdworkers face considerable inequalities in accessing such work in the first place (Newlands and Lutz, 2020). Yet, this research and similar research by Irani and Silberman (2013; Irani, 2015) predominantly focus on the worker angle, without connecting the crowdwork to specific vendors in a larger AI production ecosystem. This additional link between data preparation and the vendors who are purchasing such labour warrants further research, since the vendors set the standards and requirements for the tasks, establishing

certain parameters of the working conditions (Miceli et al., 2020). However, due to corporate secrecy, information about how the datasets are gathered remains on the 'backstage.'

Data preparation work also increasingly occurs through prison labour, a situation compounding questions of visibility and inequality. Examples of inmates performing data entry work can be found in the US, a country highly reliant on prison labour (Cao, 2019). However, prison labour for data preparation can also be observed in Finland. Niche languages, such as Finnish, make data collection and preparation difficult to outsource to lower-income countries such as the Philippines, Venezuela or Brazil where micro-work sites are more common. In Finland, a small group of prisoners were hired to prepare data for an AIaaS startup called Vainu (Chen, 2019; Kaun and Stierstedt, 2020). Vainu focuses on using AI to process business related articles to identify and classify contractors and companies by industry. While the prison workers are referred to as 'AI Trainers' in the limited media coverage, it is evident that their tasks are basic data preparation work.

Artificial intelligence verification

While no client likely believes that the AI emerges *ex nihilo*, techno-utopian rhetoric surrounding AI and particularly machine-learning AI can lead the client to vastly underestimate the amount of labour involved in AI production. One facet of AI production, which is kept particularly secret, is AI verification/validation (Tubaro et al., 2020). Based on his recent research into AI data workers in South America, Schmidt (2019) has pointed out that most of the value is in validation data, where 'human cognition is needed to evaluate the decisions that machine-learning systems have made' (9). Moreover, research has pointed out how Facebook obscures its human content moderators (Gillespie, 2020), and how Google uses human raters to perform quality assessments (Bilić, 2016). The type of validation work is also found in the use of human labor in home assistants, such as Amazon Echo and Google Home (Day et al., 2019; Verheyden et al., 2019). In this case, contracted human workers listen to audio recordings of conversations in order to review the effectiveness of the AI. Users have genuine privacy concerns regarding third-parties listening to their data (Lutz and Newlands, 2021). As demonstrated by the media backlash towards such data work, whereby users felt deceived that humans were involved in the production, such labour is often rendered invisible for strategic purposes.

Artificial intelligence impersonation

From the perspective of the AIaaS vendor, occluded human work is not limited to the data and the training. Sometimes the AI can be, in fact, humans all along. Fauxtimation, a term coined by Astra Taylor (2018) and also referred to as ‘pseudo-AI,’ ‘forged labor’ or ‘AI impersonation’ (Tubaro et al., 2020), involves a process of ontological obfuscation whereby technological deficiencies are bootstrapped through the use of human workers. Fauxtimation goes beyond the standard use of human workers to oversee AI training. Instead, human workers secretly fulfil the tasks explicitly described as being AI-powered. Shestakofsky (2017) notes how human ‘computational labor’ takes on the function of software algorithms as one part of a broader human-software configuration. Gulfs between imagination and technological realities are resolved by ‘computational labor’, using secret workers in the Philippines to perform repetitive information-processing tasks. Yet, this displacement of computer work onto human workers was ‘not because full automation was impossible, but because developers believed that achieving it would be inefficient’ (Shestakofsky, 2017: 391).

Such impersonations can be viewed as a form of ‘dark secrecy’, since they represent problematic actions firmly contained within the organisational backstage (Costas and Grey, 2016; Goffman, 1959). We can also understand these actions as Ringel’s (2019) ‘secrets of imperfection,’ since they refer to inconsistencies in private actions in contrast to expected norms. Tubaro et al. (2020), for instance, refer to Julie Desk, a French start up scheduling assistant that required human workers to masquerade as the AI assistants. In this case, the AI was supported by human work to temporarily fill the gaps between promises and reality, though the existence of human effort was not revealed to the clients. A similar instance of deception can be seen in X.ai, an ‘AI’ personal assistant service which promised to automatically scan emails, schedule meetings and email about appointments. Here, however, human workers had to pretend to be the AI during interactions with clients (Richardson, 2018). A former Xai employee recounts that the company employed people to manually receive and reply to emails, masquerading as the AI assistant. In other examples, human workers bootstrap the cognitive tasks that AI is supposed to perform. Expensify, for instance, used human workers to manually process receipt data instead of using their AI-driven ‘SmartScan technology’ (Newman, 2017; Peterson, 2017). According to the CEO of Expensify, the use of humans was merely ‘a technical detail’. Yet, the insertion of humans into this process has significant

implications for data privacy since clients’ un-redacted information was made available to third party contractors. In these few examples, the human worker is entirely hidden from the client’s view.

In current discussions, fauxtimation has been frequently conflated with the longstanding ‘Wizard of Oz’ (WOZ) HCI simulation technique (Thomason and Litman, 2013). This technique involves a human ‘wizard’ masquerading as a putative computer system to assess whether it has sufficient usability. However, Fraser and Gilbert (1991), in their review of the WOZ technique, underline that one of the pre-conditions for a WOZ simulation is that ‘it must be possible to specify the future system’s behaviour’ (82). Continuous beta-stage bootstrapping of a future system, based on undeveloped technology, is therefore not WOZ because the purpose of using human ‘wizards’ in fauxtimation is not to test usability, but to replace the system altogether. Gregory Koberger, the CEO of ReadMe, conflating fauxtimation and WOZ, is quoted as saying that fauxtimation is ‘essentially prototyping the AI with human beings’ (Solon, 2018). However, this statement is disingenuous, helping to mask a deceptive practice which has serious consequences for users and for the workers who must masquerade as an AI system.

The commercial incentives for hiding this labour are evident, namely increased venture capital funding, customer growth, and business development. As with ongoing debates around whether certain platforms are ‘technology platforms’ or ‘labour platforms’ (c.f. Uber), high valuation and funding demands minimal reliance on employees to perform vital functions. By rushing under-developed AIaaS offerings to the market and by making the human effort invisible, companies can also test out user demand without having to wait and invest in a finished service offering. Reflective of how such labour is often deeply hidden in the organisational backstage, evidence of its existence comes to light primarily through workers speaking out, from investigative reporting and research, or from the work of NGOs.

AI co-production: Lifting the curtain

As co-creators of value, or ‘prosumers’ (Toffler, 1980), clients are often expected to make their products or services work (Dujarier, 2016). This can be observed with the growth of self-service technologies (Ritzer, 2015), giving rise to concerns about the ‘overworked consumer’ (Andrews, 2019). Active consumer participation can, however, lead to role stress and role ambiguity, resulting in service failure if the service does not meet their expectations (Blut et al., 2020; Plé, 2016). In a recent study, Castillo et al. (2020) have even proposed

the possibility of co-destruction in AI-powered service interactions due to misalignments between expectations and reality. Any additional effort on behalf of a client necessary to make AIaaS function properly can also hinder uptake and cause frustration.

In many cases, AIaaS vendors must co-opt the additional labour of clients to make certain AI services function. Often this takes the form of AI training (Grønsund and Aanestad, 2020). Since this extra labour is required to be provided for free, the vendor must expose some of the reality of AI production. This offers an interesting example of the breach between the organisational frontstage and the backstage, since vendors are limited to what extent they can draw on techno-utopian narratives. As Costas and Grey (2016) explain, such organisational barriers ‘are never fully effective and may be subverted’ (83). Vendors must therefore balance hype discourses around the described benefits of the service with the reality of human intervention.

Designed to be just intelligent enough to perform specific tasks, a common manifestation of AIaaS is in the form of ‘AI colleagues,’ ‘robot workers’ or ‘digital employees’ (Huang and Rust, 2018). In an organisational setting, these AIaaS applications can be used both to replace and augment human workers (Davenport and Kirby, 2015), blurring the lines between the human workforce and what Dholakia and Firat (2019) term the ‘mechanized working class’ (1506). We can observe such ‘digital employees’ as a form of AI co-production, since significant human labour is still required to embed the AIaaS applications into the enterprise and to keep them running (Lyytinen et al., 2020). In this instance, there is a fascinating discourse between the idea of effortless AI and the reality of continuous training by human workers.

One example is offered by a Norwegian start-up, Simplifai. With offices in Norway, India and Ukraine, Simplifai offers different ‘bots’ as a service, which they term ‘digital employees’ (Simplifai, 2020a). These bots can be ‘hired’ and are given names and pictures such as Embla the E-mail processor, Sigve the Tax Advisor, or Liv in Customer Support. Once operational, these AI applications run with minimal human intervention. However, these ‘digital employees’ require training by staff members to get running. Simplifai’s e-mail bot, for instance, requires human workers to ‘tell the bot the type and category of the e-mail, what intents it has and what kinds of entities are specified’ (Simplifai, 2020c). We can observe a thought-provoking paradox in how Simplifai simultaneously highlights the need for human training, while emphasising digital employees’ inhuman capacities. As the

CCO of Simplifai, Daniel Kohn, explains: a digital employee:

never shows up late, never goes home, gets sick, skips work, or leaves early to go on a trip. Instead, this loyal, hard-working employee is always doing its job and has infinite resources. It has up to five times greater work capacity and works 5 to 10 times faster than a human being. In many ways, this is the employee of your dreams. (Simplifai, 2020b)

It is important to briefly underline here the terminological significance of AIaaS providers utilising the terms ‘employee’ to describe a technological process. Specifically, an employee is usually considered to be a human in a contractual relationship with an employer, where labour is performed for monetary compensation. By co-opting the terminology of employment, AIaaS providers are potentially destabilising the value and dignity of the human work required to produce and train the AI service. This rhetoric frames human workers as somehow defective for wanting to go home at the end of the working day. Somehow humans are not loyal, or hard-working enough because they might get sick (a particularly poignant phrase during a global pandemic).

A similar example of co-opting client labour can be observed with Memory.AI, a Norwegian AIaaS provider specialising in time-tracking. The company advertises its product, Timely, as ‘Automatic time tracking. Get the complete picture of your work day with zero effort’ (Memory AI, 2020a). As the CEO claims ‘the only thing our customers have to do is to hit an Accept button and you’re done with your timesheet’ (O’Hear, 2018). However, looking closer, we can observe greater customer effort required than merely hit ‘Accept’:

Memory AI needs input from you to get started. That input comes in the form of the hours you log from Memory, which tells the AI what projects you want those memories to be associated with. The results of that input are suggestions made by Memory AI that you can adjust then approve or reject. What really helps train the AI is not only the input you provide from Memory but also adjusting and approving the suggestions the AI makes so it can improve and make more accurate suggestions going forward. (Memory AI, 2020b)

In addition to manually annotating the data collected by Memory AI, each user has to continuously train the AI through approving or rejecting suggestions. What is particularly interesting is the refusal by MemoryAI to set a time limit on how much training is

required: ‘There is no set amount of time before Memory AI starts working. Memory AI will start learning your logging behavior as soon as you’ve logged your memories for the first time. The more you log, the faster you will receive suggestions’ (Memory AI, 2020b).

Further exemplifying this notion of co-production is Iris.ai, a Norwegian start-up offering an AIaaS scientific researcher. As advertised on Iris.ai’s website, ‘Iris.ai helps you to map out an overall landscape around your research challenge as well as build a specific reading list covering the papers that you should actually spend time on reading’ (Iris.ai, 2020a). However, making Iris.ai function requires a significant outlay of time on behalf of the client.

Step 1: ‘Write out the problem you are trying to solve in your own words., Use the visual maps to gain an overview of your problem., Use the hierarchy editor to fine-tune your map results., Bookmark papers and maps for later visits.’

Step 2: ‘Pull together your corpus, Select include/exclude concepts, Select include/exclude topics, Manually verify the results.’ (Iris.ai, 2020b)

In these two steps, it is evident that Iris.ai demands extensive human intervention to generate useful results from the AI. Yet, what is particularly interesting is how Iris.ai adopts anthropomorphic language to create a notion that the AI is an independent agent being ‘helped’ by those involved in its production. For example, Iris.ai is trained using an evaluated annotation set from a community of volunteer AI trainers. As explained by the company, ‘Iris.ai, the young AI scientist, is ready to start her first grade . . . She can’t figure out the world by herself. She needs your help!’ (Ritola, 2016).

Extending the backstage

When AIaaS vendors must involve the human labour of clients to co-produce the AI service, the role of human intervention is thus made hypervisible. However, in cases where clients use the AIaaS as an intermediary between themselves and customers, such as through interview services or chatbots, we can observe an interesting extension of the ‘backstage’ outwards. In these cases, clients themselves can engage in a form of strategic secrecy to obfuscate the level of human effort involved in the AI. Without the requirement to co-opt customer labour, the clients can similarly draw on techno-utopian concepts of AI as being effective and unbiased.

Interview services

Currently, there are a number of AI-based interview services on the market. One of these, HireVue, is a US-based video platform which uses algorithms to assess answers and facial expressions. Seedlink, based in Amsterdam, also asks candidates to answer questions on a smartphone and uses AI to analyse language for cultural fit.

TengAI, a Swedish HR start-up founded in 2019, uses machine learning to power an interview robot and recruitment assistant (TengAI, 2020a). TengAI’s main market offering is the promise of an unbiased interview experience, as well as a socially-distanced interview option suitable for the Covid-19 pandemic. However, the company simultaneously emphasises the commercial benefits of reducing labour costs. As they emphasise, it is ‘possible to interview more candidates since TengAI can perform 4 times faster and 7 times more robot interviews than a human recruiter’ (TengAI, 2020b). In terms of the process, TengAI requires their enterprise clients to tailor the AI parameters, the ‘TengAI Performance Indicator’, to the specific hiring requirements (TengAI, 2020c). These parameters are then used to assess candidates and provide a shortlist. However, the interview process retains critical gaps where human effort is made internally hypervisible.

Firstly, TengAI gives the client the opportunity to combine the AI’s performance scores with a manual scoring of interview audio. In addition, TengAI emphasises that the HR team can ‘conduct personal interviews with final candidates to ensure motivation and drive’ (TengAI, 2020c). Although the TengAI experience is emphasised as being unbiased and AI-driven, particularly towards interviewees, the process retains a considerable level of human subjectivity. In company documentation, TengAI (2020d) emphasise that their process shifts ‘the subjectivity along in the process (where it is less damaging). The in-depth assessment will still be done by an experienced recruiter trained in unbiased recruitment’ (TengAI, 2020d). This balance of human and AI can be understood through a power lens, since organisations want to retain control rather than surrender full autonomy over hiring decisions to an AI.

Yet, we can observe the extension of the ‘backstage’ by the client, through the contrast between the interviewee and interviewer understanding of the AI service. While the client organisation is fully aware that the interview process is full of human labour and subjective choices, the interviewee is given the impression that the process is more automated. Interview processes are critical for initial socialisation and to lay the psychological contract between employers and potential

employees. By hiding the human subjectivity in the process behind a robotic face, organisations are potentially risking the interview's success as a socialisation mechanism, while only gaining a small reduction in their bias (depending on how the assessment parameters are coded). Moreover, as with the use of AI-based CV screening, AIaaS interview experiences may alienate and dehumanise potential applicants (Buranyi, 2018). What organisations do gain, however, in such an AIaaS adoption and the occlusion of the human subjectivity, is a reduction in labour costs and the reputational benefits of being AI-driven and unbiased.

Chatbots

One of the most prominent channels of AIaaS adoption in enterprise is the use of chatbots to replace and augment front line employees (FLEs), such as customer service agents (Xiao and Kumar, 2019). Chatbots enable rich interactions with people, triggering the view that they are social entities (Jörling and Böhm Paluch, 2019). Yet, although advances in chatbot technology have continued apace, chatbots offer limited functionality to users without manual training by dedicated workers (Følstad and Brandtzaeg, 2017). Chatbots otherwise remain incapable of making sense of nuance, meaning or social relationships (Chakrabarti and Luger, 2015; Pantano and Pizzi, 2020).

There are two key types of human workers necessary to make chatbots function: chatbot trainers and chatbot operators. Although the terminology used varies between vendors, chatbot trainers can be understood as those workers who manually train the chatbot, determining parameters and validating responses. Chatbot training is key for implementation (Kvale et al., 2019) and for best results, chatbot training is a continuous process which must be undertaken within the client organisation. Chatbot operators, by contrast, are the FLEs who work alongside the chatbot in answering customer queries. In the majority of cases, organisations are using chatbot AIaaS as a form of complementary augmentation (Davenport et al., 2020). Rather than the chatbot taking full control, FLEs and chatbots work together and 'handover' customers between them. In both these cases, since the additional training must be co-opted, the necessity for human labour is emphasised by the vendor.

As an example of the value of making AI training hypervisible, Norwegian chatbot vendor Boost.ai boasts of more than 1500 certified AI trainers. While offering a 'self-learning AI,' they explain that:

To achieve this unparalleled consistency, it requires not just the very best natural language technology but also the expertise of human customer service staff. Existing

employees can translate their extensive knowledge of your company's products and services into a new role called the 'AI Trainer' that is responsible for training and maintaining the virtual agent. Imagine the potential of a virtual agent that never sleeps, combined with the wealth of experience of your most competent customer support staff. (Boost.ai, 2020)

In this case, the AI Trainer is essential in writing dialogue for the chatbot, but also in shaping the chatbot's personality and validating its decisions. Rather than hide this labour away, they highlight it, offering an AI Trainer conference, an AI Trainer knowledge base, and an 'Advanced AI Trainer course' (Boost.ai, 2018). Yet, Boost.ai does not hide the key market value of their service: the ability to reduce labour costs by having a virtual agent that never sleeps.

This hypervisibility of the chatbot trainer labour to the client, however, is contrasted to the often obfuscated nature of the chatbot operator labour to the end customer. Frequently, end customers are unable to tell whether they are communicating with a chatbot or a human operator. Here, the client is able to use the AI service to create their own 'backstage', which blurs the line between AI and human labour for their own commercial benefit. However, during critical incidents when chatbot users want immediate or sensitive responses, chatbots are unsuitable and their use can cause frustration and anger among users. As Castillo et al., (2020) demonstrate, customers feel deceived if they interact with a chatbot rather than a human agent, feeling like the enterprise is unwilling to provide adequate resources to deal with them.

Addressing this widespread practice, BotXO, a Danish chatbot provider, emphasises the importance of 'Bot Operators'. Rather than merging the chatbots with the human operators, BotXO emphasises the importance of building trust between customers and business by advertising the ontological status of the chatbot:

It is also essential to have your bot admits its own limits and shortcomings, when an error or understanding occurs . . . You can easily prevent this by always providing the customer with the choice of getting transferred from chatbot support to human customer support, in case a chat goes south. Make your bot admit that it is, indeed, just an AI bot that is still learning. This will give a more personal touch and will reduce the risk of the user feeling alienated or angered. (BotXO.ai, 2020)

By encouraging client disclosure of a bot's shortcomings and the option for human customer support, BotXO pursues strategic commercial incentives of reducing customer frustration. By downplaying the bot's capabilities

and making the human labour hypervisible, clients can thus retain the benefits (labour cost reduction) while avoiding the negatives (customer frustration).

Conclusion

The commercial growth of AIaaS in recent years continues to shape many individuals' direct experience with AI, whether in terms of enterprise solutions, customer service or time-saving technologies. Yet, while human effort remains essential in AIaaS at every stage of its production and deployment, not all human effort is rendered visible in the same way. Within this article, I have therefore used Goffman's (1959) dramaturgic lens of the 'backstage' to understand how human labour is rendered visible or hypervisible in AI production. In essence, AIaaS vendors manipulate the visibility of human involvement in AI based on whether the AIaaS vendor must co-opt paid or unpaid labour to fill the remaining gaps.

When vendors must co-opt unpaid labour from clients to make AI services function, such as through localised chatbot training or user validation of output, AIaaS vendors emphasise the limitations of the AI service and highlight the importance of human-involvement. For instance, the valorisation of chatbot trainers highlights how the AIaaS vendor and user must co-produce the eventual service. An emphasis on interviewer subjectivity and control in AI interview applications enables the client to benefit from the notion of AI, while retaining full human control over the process.

By contrast, where the AIaaS vendors make use of paid workers 'behind the curtain', such as in data preparation or vendor-situated fauxtimation, they are able to commercially benefit by emphasising the technical wizardry and rendering the workforce invisible. By paying for the labour that feeds the AI, vendors are thus able to make that very labour invisible. A Norwegian AIaaS company, for instance, can thus market itself as a local exemplar of Nordic innovation, while outsourcing the entire technological development to India or Eastern Europe and its data supply chain to South America.

While this article has focused on the dynamic visibility between vendors and clients, it cannot be overstated how complex the formal and informal networks have become among different stakeholders in the AI production ecosystem. Going forward, it is pertinent to continue to research AI production from multiple angles, perspectives, and contexts. This strategic rendering of AI production raises several key issues which deserve further research.

Firstly, there are ethical concerns regarding how to ensure transparency and accountability when using

externally developed AI applications, where the client cannot access, query, nor understand how AI-derived decisions are made (Felzmann et al., 2019). Intellectual property rights and technical skill gaps can restrict even powerful clients from understanding how their AIaaS operates, a particularly worrying prospect if off-the-shelf AI solutions are rolled out in essential public services.

Secondly, there are commercial concerns which only intensify as the gap between hype and reality widen. If AIaaS is advertised as effortless and quasi-magical, but relies on the flow of human sweat more than the flow of electric currents, then companies may have to be careful with dissatisfied customers or consumer protection agencies.

Thirdly, there are labour concerns. AI production operates in a global network, where AIaaS providers, the workforce and their customers are distributed worldwide. While identifying how work is hidden in AI production, it is essential to identify where the workers are, and whose work is being hidden. Itemised acknowledgement of the role of data-workers is not feasible, nor is it probably desirable by the workers themselves to have their names attached to the infinite variety of AI solutions deriving from their data-work. However, general acknowledgement, from an industry and societal level, of the role of data-workers and the reliance of AIaaS on such data work, could enhance the prestige, salary, and recognition of the work. As an emerging industry, AIaaS vendors can set the benchmarks for adequate conditions, compensation and recognition commensurate with the actual work contributed.


Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

This work was supported by the Research Council of Norway within the FRIPRO TOPPFORSK (275347) project 'Future Ways of Working in the Digital Economy'.

ORCID iD

Gemma Newlands  <https://orcid.org/0000-0003-0851-384X>

References

- Acemoglu D and Restrepo P (2018) The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review* 108(6): 1488–1542.
- Altenreid M (2020) The platform as factory: Crowdwork and the hidden labour behind artificial intelligence. *Capital & Class* 44(2): 145–158.

- Andrews CK (2019) *The Overworked Consumer: Self-Checkouts, Supermarkets, and the Do-It-Yourself Economy*. London: Lexington Books.
- Angella M (2016) Work, recognition and subjectivity: Relocating the connection between work and social pathologies. *European Journal of Social Theory* 19(3): 340–354.
- Bechmann A and Bowker G (2019) Unsupervised by any other name: Hidden layers of knowledge production in artificial intelligence on social media. *Big Data & Society* 6(1): 1–11.
- Bilić P (2016) Search algorithms, hidden labour and information control. *Big Data & Society* 3(1): 1–9.
- Blut M, Heirati N and Schoefer K (2020) The dark side of customer participation: When customer participation in service co-development leads to role stress. *Journal of Service Research* 23(2): 156–173.
- Birchall C (2011) Introduction to ‘secrecy and transparency’ the politics of opacity and openness. *Theory, Culture & Society* 28(7–8): 7–25.
- Boost.ai (2018) AI Trainer: a job of the future. Available at: <https://www.boost.ai/articles/ai-trainer-a-job-of-the-future> (accessed 17 November 2020).
- Boost.ai (2020) AI for customer service: 5 unique benefits. Available at: <https://www.boost.ai/ai-for-customer-service> (accessed 17 November 2020).
- BotXO.ai (2020) Are you building a chatbot? Stop making these 10 UX mistakes. Available at: <https://www.botxo.ai/en/blog/top-ux-mistakes-chatbot/> (accessed 17 November 2020).
- Brighenti A (2007) Visibility: a category for the social sciences. *Current Sociology* 55(3): 323–342.
- Buranyi S (2018) ‘Dehumanising, impenetrable, frustrating’: The grim reality of job hunting in the age of AI. *The Guardian*, 4 March. Available at: <https://www.theguardian.com/inequality/2018/mar/04/dehumanising-impenetrable-frustrating-the-grim-reality-of-job-hunting-in-the-age-of-ai> (accessed 6 May 2021).
- Cao L (2019) Made in the USA: Race, trade and prison labor. *NYU Review of Law and Social Change* 43: 1–58.
- Casati F, Govindarajan K, Jayaraman B, et al. (2019) Operating enterprise AI as a service. In: Yanguis S, Bouassida Rodriguez I, Drira K and Tari Z (eds) *ICSOC: Service-oriented Computing*. Cham, Switzerland: Springer, pp.331–344.
- Castillo D, Canhoto AI and Said E (2020) The dark side of AI-powered service interactions: Exploring the process of co-destruction from the customer perspective. *The Service Industries Journal*. Epub ahead of print 30 June 2020. DOI: 10.1080/02642069.2020.1787993.
- Celentano D (2019) Automation, labour justice, and equality. *Ethics and Social Welfare* 13(1): 33–50.
- Chakrabarti C and Luger GF (2015) Artificial conversations for customer service chatter bots: Architecture, algorithms, and evaluation metrics. *Expert Systems with Applications* 42(20): 6878–6897.
- Chen A (2019) Inmates in Finland are training AI as part of prison labor. *The Verge*, March 28. Available at: <https://www.theverge.com/2019/3/28/18285572/prison-labor-finland-artificial-intelligence-data-tagging-vainu> (accessed 6 May 2021).
- Cheney J, Chong S, Foster N, et al. (2009) Provenance: a future history. In: *Proceedings of the 24th ACM SIGPLAN conference companion on object oriented programming systems languages and applications*, Orlando, Florida, 25–29 October 2009, pp. 957–964.
- Cognilytica (2019) *Data Engineering, Preparation, and Labeling for AI 2019*. Report n. cgr-de100. Washington DC: Cognilytica.
- Costas J and Grey C (2014) Bringing secrecy into the open: Towards a theorization of the social processes of organizational secrecy. *Organization Studies* 35(10): 1423–1447.
- Costas J and Grey C (2016) *Secrecy at Work: The Hidden Architecture of Organizational Life*. Stanford, CA: Stanford University Press.
- Crain M (2018) The limits of transparency: Data brokers and commodification. *New Media & Society* 20(1): 88–104.
- Crain M, Poster WR and Cherry MA (2016) *Invisible Labor: Hidden Work in the Contemporary World*. Oakland, CA: University of California Press.
- Cubitt S (2017) *Finite Media*. Durham, NC: Duke University Press.
- Daniels A (1987) Invisible work. *Social Problems* 34(5): 403–415.
- Dashtipour P and Vidaillet B (2017) Work as affective experience: The contribution of Christophe Dejours’ ‘psychodynamics of Work’. *Organization* 24(1): 18–35.
- Davenport T, Guha A, Grewal D, et al. (2020) How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science* 48(1): 24–42.
- Davenport TH and Kirby J (2015) Beyond automation. *Harvard Business Review* 93(6): 59–65.
- Day M, Turner G and Drozdak N (2019) Amazon workers are listening to what you tell Alexa. *Bloomberg*. 11 April. Available at: <https://www.bloomberg.com/news/articles/2019-04-10/is-anyone-listening-to-you-on-alexa-a-global-team-reviews-audio> (accessed 6 May 2021).
- Dholakia N and Firat AF (2019) Markets, consumers and society in the age of heteromation. *European Journal of Marketing* 53(8): 1504–1520.
- Dujarier MA (2016) The three sociological types of consumer work. *Journal of Consumer Culture* 16(2): 555–571.
- Elish MC and Boyd D (2018) Situating methods in the magic of Big Data and AI. *Communication Monographs* 85(1): 57–80.
- Felzmann H, Fosch Villaronga E, Lutz C, et al. (2019) Transparency you can trust: Transparency requirements for artificial intelligence between legal norms and contextual concerns. *Big Data & Society* 6(1): 1–14.
- Floridi L (2020) AI and its new winter: from myths to realities. *Philosophy & Technology* 33(1): 3.
- Følstad A and Brandtzaeg PB (2017) Chatbots and the new world of HCI. *Interactions* 24(4): 38–42.
- Fraser NM and Gilbert GN (1991) Simulating speech systems. *Computer Speech & Language* 5(1): 81–99.
- Gillespie T (2020) Content moderation, AI, and the question of scale. *Big Data & Society* 7(2): 1–5.

- Giuliano RM (2020) Echoes of myth and magic in the language of artificial intelligence. *AI & Society* 35(4): 1009–1024.
- Goffman E (1959) *The Presentation of Self in Everyday Life*. New York: Doubleday.
- Gray ML and Suri S (2019) *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Eamon Dolan Books.
- Grønund T and Aanestad M (2020) Augmenting the algorithm: Emerging human-in-the-loop work configurations. *The Journal of Strategic Information Systems* 29(2): 101614.
- Grosman J and Reigeluth T (2019) Perspectives on algorithmic normativities: Engineers, objects, activities. *Big Data & Society* 6(2): 1–12.
- Gruszka K and Böhm M (2020) Out of sight, out of mind?(in) visibility of/in platform-mediated work. *New Media & Society*. Epub ahead of print 25 December 2020. DOI: 10.1177/1461444820977209
- Hatton E (2017) Mechanisms of invisibility: Rethinking the concept of invisible work. *Work, Employment and Society* 31(2): 336–351.
- Hockenhull M and Cohn ML (2021) Hot air and corporate sociotechnical imaginaries: Performing and translating digital futures in the Danish tech scene. *New Media & Society* 23(2): 302–321.
- Huang MH and Rust RT (2018) Artificial intelligence in service. *Journal of Service Research* 21(2): 155–172.
- Irani L (2015) The cultural work of microwork. *New Media & Society* 17(5): 720–739.
- Irani LC and Silberman MS (2013) April) Turkopticon: Interrupting worker invisibility in Amazon Mechanical Turk. In: *Proceedings of the SIGCHI conference on human factors in computing systems*, Paris, France. 27 April–2 May 2013, pp. 611–620.
- Iris.ai (2020a) What is Iris.ai?. Available at: <https://help.iris.ai/basics-of-the-iris-ai-science-assistant/what-is-iris-ai/> (accessed 17 November 2020).
- Iris.ai (2020b) Features. Available at: <https://iris.ai/features> (accessed 17 November 2020).
- Jasanoff S and Kim SH (2009) Containing the atom: Sociotechnical imaginaries and nuclear power in the United States and South Korea. *Minerva* 47(2): 119.
- Jo ES and Gebru T (2020) Lessons from archives: Strategies for collecting sociocultural data in machine learning. In: *Proceedings of the 2020 conference on fairness, accountability, and transparency*, Barcelona, Spain, 27–30 January 2020, pp.306–316.
- Jörling M and Böhm Paluch S (2019) Service robots: Drivers of perceived responsibility for service outcomes. *Journal of Service Research* 22(4): 404–420.
- Kaplan A, Sabbah-Karkabi M and Herzog H (2020) “When I iron my son’s shirt, I feel my maternal role”: Making women’s invisible work visible. *Journal of Family Issues* 41(9): 1525–1545.
- Kaun A and Stierstedt F (2020) Prison media work: From manual labor to the work of being tracked. *Media, Culture & Society* 42(7–8): 1277–1292.
- Kilvington D (2020) The virtual stages of hate: Using Goffman’s work to conceptualise the motivations for online hate. *Media, Culture & Society* 43(2): 256–272.
- Kvale K, Sell OA, Hodnebrog S, et al. (2019) Improving conversations: Lessons learnt from manual analysis of chatbot dialogues. In: *International workshop on chatbot research and design*, pp.187–200.
- Laguecir A and Leca B (2019) Strategies of visibility in contemporary surveillance settings: Insights from misconduct concealment in financial markets. *Critical Perspectives on Accounting* 62: 39–58.
- Lloyd C and Payne J (2019) Rethinking country effects: Robotics, AI and work futures in Norway and the UK. *New Technology, Work and Employment* 34(3): 208–225.
- Lutz C and Newlands G (2021) Privacy and smart speakers: A multi-dimensional approach. *The Information Society*. Epub ahead of print 27 March 2021. DOI: 10.1080/01972243.2021.1897914.
- Lyytinen K, Nickerson JV and King JL (2020) Metahuman systems= humans+ machines that learn. *Journal of Information Technology*. Epub ahead of print 25 May 2020. DOI: 10.1177/0268396220915917.
- Mager A and Katzenbach C (2021) Future imaginaries in the making and governing of digital technology: Multiple, contested, commodified. *New Media & Society* 23(3): 223–236.
- Manning PK (2008) Goffman on organizations. *Organization Studies* 29(5): 677–699.
- Memory AI (2020a) Automatic time tracking. Available at: <https://memory.ai/timely/features/automatic-time-tracking> (accessed 17 November 2020).
- Memory AI (2020b) Ask the experts: All about memory AI. Available at: <https://support.memory.ai/en/articles/3405984-memory-ai-faqs> (accessed 17 November 2020).
- Metelskaia I, Ignatyeva O, Deneff S, et al. (2018) A business model template for AI solutions. In: *Proceedings of the international conference on intelligent science and technology (ICIST ’18)*. New York, NY: Association for Computing Machinery, pp.35–41.
- Miceli M, Schuessler M and Yang T (2020) Between subjectivity and imposition: Power dynamics in data annotation for computer vision. arXiv:2007.886.
- Mosco V (2015) *To the Cloud: Big Data in a Turbulent World*. London: Routledge.
- Mützel S (2021) Unlocking the payment experience: Future imaginaries in the case of digital payments. *New Media & Society* 23(2): 284–301.
- Newlands G (2020) Algorithmic surveillance in the gig economy: The organization of work through Lefebvrian conceived space. *Organization Studies*. Epub ahead of print 9 July 2020. DOI: 10.1177/0170840620937900
- Newlands G and Lutz C (2020) Crowdsourcing and the mobile underclass: Barriers to participation in India and the United States. *New Media and Society*. Epub ahead of print 27 January 2020. DOI: 10.1177/1461444820901847
- Newman L (2017) It’s not always AI that sifts through your sensitive info. *Wired*, 29 November. Available at: <https://www.wired.com/story/not-always-ai-that-sifts-through-sensitive-info-crowdsourced-labor/> (accessed 6 May 2021).
- O’Hear S (2018) Memory raises \$5M to bring AI to time tracking. *Tech Crunch*. Available at: <https://techcrunch.com/2018/10/17/memory-timely/> (accessed 23 October 2020).

- Otis EM and Zhao Z (2016) Producing invisibility: surveillance, hunger, and work in the produce aisles of Wal-Mart, China. In: Crain M, Poster W and Cherry M (eds) *Invisible Labor*. Berkeley, CA: University of California Press, pp.148–168.
- Pantano E and Pizzi G (2020) Forecasting artificial intelligence on online customer assistance: Evidence from chatbot patents analysis. *Journal of Retailing and Consumer Services* 55: 1–9.
- Parsaeefard S, Tabrizian I and Leon-Garcia A (2019) Artificial Intelligence as a Service (AI-aaS) on software-defined infrastructure. In *2019 IEEE conference on standards for communications and networking (CSCN)*, Granada, Spain, pp.1–7.
- Peterson B (2017) Expensify CEO says its ‘automated’ services ‘proud’ to use humans to process receipts. *Business Insider*, 28 November. Available at: <https://www.businessinsider.nl/expensify-is-proud-to-use-humans-in-its-automated-service-2017-11/?international=true&r=US> (accessed 6 May 2021).
- Pettersen L (2019) Why artificial intelligence will not outsmart complex knowledge work. *Work, Employment and Society* 33(6): 1058–1067.
- Plantin JC (2019) Data cleaners for pristine datasets: Visibility and invisibility of data processors in social science. *Science, Technology, & Human Values* 44(1): 52–73.
- Plé L (2016) Studying customers’ resource integration by service employees in interactional value co-creation. *Journal of Services Marketing* 30(2): 152–164.
- Prainsack B (2019) Logged out: Ownership, exclusion and public value in the digital data and information commons. *Big Data & Society* 6(1): 2053951719829773.
- Rabelo VC and Mahalingam R (2019) “They really don’t want to see us”: How cleaners experience invisible ‘dirty’ work. *Journal of Vocational Behavior* 113: 103–114.
- Raisch S and Krakowski S (2021) Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review* 46(1): 192–210.
- Richardson J (2018) AI chatbots try to schedule meetings without enraging us. *Wired*, 24 May. Available at: <https://www.wired.com/story/xai-meeting-ai-chatbot/> (accessed 6 May 2021).
- Ringel L (2019) Unpacking the transparency–secrecy nexus: Frontstage and backstage behaviour in a political party. *Organization Studies* 40(5): 705–723.
- Ritola M (2016) Ready to train Iris.ai?. *Iris.ai Blog*, 10 May. Available at: <https://iris.ai/launch/aitraining/> (accessed 6 May 2021).
- Ritzer G (2015) Prosumer Capitalism. *The Sociological Quarterly* 56: 413–445.
- Schmidt C (2017) *Digital Labour Markets in the Platform Economy: Mapping the Political Challenges of Crowdwork and Gig Work*. Friedrich-Ebert-Stiftung, Bonn, Germany.
- Schmidt F (2019) *Crowdproduktion von Trainingsdaten: Zur Rolle von Online-Arbeit Beim Trainieren Autonomer Fahrzeuge*. Hans-Böckler-Stiftung, Dusseldorf, Germany.
- Scholz T (2013) *Digital Labor: The Internet as Playground and Factory*. London: Routledge.
- Scroggins MJ and Pasquetto IV (2020) Labor out of place: On the varieties and valences of (in) visible labor in data-intensive science. *Engaging Science, Technology, and Society* 6: 111–132.
- Selbst A (2017) A mild defense of our new machine overlords. *Vanderbilt Law Review* 70: 87–104.
- Shestakofsky B (2017) Working algorithms: Software automation and the future of work. *Work and Occupations* 44(4): 376–423.
- Simplifai (2020a) Digital Employees. Available at: <https://www.simplifai.ai/en/digital-employee/> (accessed 23 October 2020).
- Simplifai (2020b) The digital employee who loves the tedious tasks. Available at: <https://www.simplifai.ai/en/claims-link/> (accessed 23 October 2020).
- Solon O (2018) The rise of ‘pseudo-AI’: How tech firms quietly use humans to do bots work. *The Guardian* 6 July. Available at: <https://www.theguardian.com/technology/2018/jul/06/artificial-intelligence-ai-humans-bots-tech-companies>
- Star SL and Strauss A (1999) Layers of Silence, Arenas of Voice: The Ecology of Visible and Invisible Work. *Computer Supported Cooperative Work (CSCW)* 8: 9–30.
- Taylor A (2018) *The automation charade*. Available at: <https://logicmag.io/failure/the-automation-charade/> (accessed 6 May 2021).
- TengAI (2020a) The leading interview robot and recruitment assistant. Available at: <https://www.tengai-unbiased.com/> (accessed 17 November 2020).
- TengAI (2020b) Unique offer: Corona safe recruitment with interview robot. Available at: <https://www.tengai-unbiased.com/corona-safe-recruitment-with-interview-robot/> (accessed 17 November 2020).
- TengAI (2020c) Recruitment solutions. Available at: <https://www.tengai-unbiased.com/recruitment-solutions/> (accessed 17 November 2020).
- TengAI (2020d) Interview robot FAQ-the most frequently asked questions about Tengai. Available at: <https://www.tengai-unbiased.com/interview-robot-faq/> (accessed 17 November).
- Thomason J, and Litman D (2013) Differences in user responses to a wizard-of-oz versus automated system. In: *Proceedings of the 2013 conference of the North American chapter of the association for computational linguistics: human language technologies*, Atlanta, Georgia, 9–14 June 2013, pp.796–801.
- Timonen H and Vuori J (2018) Visibility of work: How digitalization changes the workplace. In: *Proceedings of the 51st Hawaii international conference on system sciences*, Honolulu, Hawaii, 3 January 2018, pp.5075–5084.
- Toffler A (1980) *The Third Wave*. New York: William Morrow.
- Tubaro P and Casilli A (2019) Micro-work, artificial intelligence and the automotive industry. *Journal of Industrial and Business Economics* 46(3): 333–345.
- Tubaro P, Casilli AA and Coville M (2020) The trainer, the verifier, the imitator: Three ways in which human

- platform workers support artificial intelligence. *Big Data & Society* 7(1): 1–12.
- Tubaro P (2021) Disembedded or deeply embedded? A multi-level network analysis of online labour platforms. *Sociology*. Epub ahead of print 31 January 2021. DOI: 10.1177/0038038520986082.
- Verheyden T, Baert D, van Hee L, et al. (2019) Google employees are eavesdropping, even in your living room, VRT NNWS has discovered. *VRT* 10 July. Available at: <https://www.vrt.be/vrtnws/en/2019/07/10/google-employees-are-eavesdropping-even-in-flemish-living-rooms/> (accessed 6 May 2021).
- Vesa M and Tienari J (2020) Artificial intelligence and rationalized unaccountability: Ideology of the elites? *Organization*. Epub ahead of print 22 October 2020. DOI: 10.1177/1350508420963872
- Villena VH and Gioia DA (2018) On the riskiness of lower-tier suppliers: Managing sustainability in supply networks. *Journal of Operations Management* 64: 65–87.
- von Krogh G (2018) Artificial intelligence in organizations: New opportunities for phenomenon-based theorizing. *Academy of Management Discoveries* 4(4): 404–409.
- Whittle A, Gilchrist A, Mueller F, et al. (2020) The art of stage-craft: A dramaturgical perspective on strategic change. *Strategic Organization*. Epub ahead of print 27 April 2020. DOI: 10.1177/1476127020914225.
- Whiting R and Symon G (2020) Digi-housekeeping: The invisible work of flexibility. *Work, Employment and Society*. Epub ahead of print 19 June 2020. DOI: 10.1177/0950017020916192.
- Williams R and Pollock N (2012) Research commentary—moving beyond the single site implementation study: How (and why) we should study the biography of packaged enterprise solutions. *Information Systems Research* 23(1): 1–22.
- Wyatt S (2021) Metaphors in critical Internet and digital media studies. *New Media & Society* 23(2): 406–416.
- Xiao L and Kumar V (2019) Robotics for customer service: A useful complement or an ultimate substitute? *Journal of Service Research*. Epub ahead of print 30 September 2019. DOI: 10.1177/1094670519878881.
- Yeh CL (2018) Pursuing consumer empowerment in the age of big data: A comprehensive regulatory framework for data brokers. *Telecommunications Policy* 42(4): 282–292.