

Christine T. Wolf (2020): AI Ethics and Customer Care: Some Considerations from the Case of “Intelligent Sales”. In: Proceedings of the 18th European Conference on Computer-Supported Cooperative Work: The International Venue on Practice-centred Computing on the Design of Cooperation Technologies - Exploratory Papers, Reports of the European Society for Socially Embedded Technologies (ISSN 2510-2591), DOI: 10.18420/ecscw2019_n02

Copyright 2020 held by Authors, DOI: 10.18420/ecscw2020_n02

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, contact the Authors.

AI Ethics and Customer Care: Some Considerations from the Case of “Intelligent Sales”

Christine T. Wolf
IBM Research – Almaden
ctwolf@us.ibm.com

Abstract. This note considers the topic of AI ethics as it relates to applied, industrial AI projects. In particular, it examines “intelligent sales,” a strategic management concept that envisions the enhancement of workflows within a sales organization with the use of Big Data and artificial intelligence (AI). This note examines common depictions of intelligent sales campaigns in management literature, identifying key topics in these discourses: *data fusion*, *responsive customer care*, *automation*, and *streamlined delivery*. We focus on ways in which intelligent sales is envisioned to enrich the customer-sales relationship through the use of Big Data and AI, surfacing ethical considerations around: training data and the use of AI outputs in everyday work practices. This paper contributes to discourses on the fairness, accountability, and transparency (FAcCT*) of algorithmic systems by raising a number of emergent concerns in enterprise AI applications and in particular some considerations from the emergent management concept intelligent sales.

Introduction

Contemporary organizations are increasingly scrutinized on their ability to perform “data drivenness,” that is, demonstrate the innovative development and use of Big Data and artificial intelligence (AI) technologies and capabilities to enhance, standardize, and optimize their business operations. Commonly signaled today in popular business and management trade press under monikers such as “intelligent enterprise” (Schoemaker and Tetlock 2017) or “cognitive enterprise” (IBM Institute for Business Value 2018), these trends are not new. Indeed, they follow a long-tailed trajectory of scientific management (Bear 2009; Braverman 1998; Taylor 2003) and statistical governance (Desrosières 2002).

The aspirations of the cognitive enterprise are propelled by two entwined phenomena: the emergence and maturing of the Big Data era; and the renaissance of AI and machine learning (ML) development that has followed in its wake. The impact of Big Data and AI are seen as offering great potential to transform contemporary society, similar to the PC and computerization social movements of prior decades (Kling and Iacono 1988). It is undeniable that these two keywords (“Big Data,” and “AI”) are driving forward contemporary business and strategic management discourses, obvious to even the casual reader of trade publications like *Harvard Business Review*, *MIT Technology Review*, *Forbes*, *CIO*, and *The Economist*, to name a few.

Big Data and AI have changed many aspects of everyday, consumer life, but these trends also pose opportunities for transformation in a wide range of industrial domains and processes. As Chui et al. (2018) discuss in a *Harvard Business Review* article, supply chain management, as well as sales and marketing are the two domain areas heralded as having the most immediate promise for enterprise AI. While specific imagined application areas are diverse, popular management discourses often focus on two overarching visions of how data will transform businesses. One, is the ability of a firm to leverage previously “untapped” data sources within their organization, often unstructured data like text documents (“data re-use” transformation). Another is the incorporation of novel data sources into organizational information practices, such as bringing social media or Internet of Things (IoT) data to bear on internal firm processes (“novel data” transformation). AI capabilities build on these Big Data stores, with novel algorithms and techniques introduced on an almost-daily basis that can provide predictive outputs and enable machine action to execute computational tasks with remarkable accuracy.

In this note, we focus on Big Data and AI transformations in the sales domain, and in particular efforts commonly called “intelligent sales.” Sales is a complex and multi-faceted business process, requiring the coordinated efforts of many departments within a firm (e.g., sales staff, but also: marketing, accounting,

account management, procurement, delivery/fulfillment, to name a few). Considerable effort is required to see a “lead” (a potential customer, also called a “prospect”) all the way through to “cash” (revenue flowing from the customer to the sales organization). This overarching super-process is often referred to as “lead-to-cash,” a simple phrase that encompasses a complex web of interdependent processes and sub-processes that weaves together many sociotechnical actors. As Sinha (2018) describes in a *Future of Customer Engagement and Experience* report: “Lead to cash is arguably the most important customer-centric process in an organization, starting with the customer’s intention to buy, and ending with revenue recognition.” Many leading IT service providers, such as SAP, Salesforce, IBM, or Microsoft, offer software products and services to support the complex information processes of sales organizations. “Intelligent sales” is an emergent management concept that sees sales processes as ones ripe for Big Data/AI transformations.

Although particular implementations may differ, there are several key themes salient across intelligent sales products and services, as discussed in grey literature from the business management field. Grey literature is “the diverse and heterogeneous body of material available outside, and not subject to, traditional academic peer-review processes” (Adams et al. 2017) (p. 433). While care must be taken to contextualize its non-peer-reviewed status, grey literature can productively enhance scholarly knowledge by providing “relevant contemporary material in dynamic and applied topic areas where scholarship lags” (Adams et al. 2017) (p. 433). In this note, we outline four key themes identified in popular business and management trade press discourses on the topic of intelligent sales:

- *data fusion* – pooled and integrated data across departments within the sales organization, providing a “global” view of customer accounts and profiles;
- *responsive customer care* – enhanced sales relationships, where salespeople have broader, data-driven understandings of customer needs;
- *automation* – the reduction of repetitive tasks throughout the sales process via robotic process automation (RPA); and
- *streamlined delivery* – standardized, data-driven processes of offering selection, pricing, contracting, and fulfillment.

In each of these, Big Data and AI play particular roles in changing segments of the lead-to-cash process into various data-driven workflows. These hoped-for changes aim to transform organizational life sociomaterially and at varying scales, with implications at the technological, work practice, as well as strategic level. Readers will note threads which tie intelligent sales to longstanding issues within the CSCW community: data integration and re-use, workplace automation and standardization, and changing service relationships – these have all been topics of concerns since workplace computerization in the 1970s and 1980s. We make note of these ties and also draw attention to the novel considerations raised in the Big Data/AI context and in particular how the case of intelligent sales presents us with emergent questions on the design and study of data-driven

workplace technologies. Most notable of these questions are ones of ethics raised by Big Data/AI. We identify a number of ethical issues that the intelligent sales case implicates, organizing them into two themes: training data and the use of AI outputs in everyday work practices. This note brings discourses on the fairness, accountability, and transparency (FAcCT*) of algorithmic systems in conversation with those in the CSCW community on practice-oriented computing, raising a number of open and pressing questions in need of further interrogation and scrutiny.

The rest of this note is laid out as follows. In Section 2, we provide an in-depth description of the “intelligent sales” concept, its key features, and case studies. In Section 3, we discuss two ethical themes surfaced through the intelligent sales case study. Then in Section 4, we conclude the paper.

Intelligent Sales

Sales is a complex business process that involves many different departments within an organization – for example, marketing, sales/account management, finance, procurement, delivery/fulfillment, and customer service, to name a few. These departments must coordinate efforts and information – towards shared goals of contracts being signed, revenue being realized, and services delivered. Sales organizations are driven by both a commercial logic (closing deals) and also a service ethic (serving customers) (Brokling 2014; McClaren 2000; Oakes 1992) and intelligent sales is seen as a way of meeting both demands together. Many leading IT service providers (e.g., SAP, Salesforce, IBM) offer software products and services to support sales organizations. Intelligent sales is a strategic concept and approach that envisions a “nextgen” sales experience transformed through the infusion of Big Data and AI capabilities to enhance, streamline, and make more efficient the various work processes of a sales organization (Sinha 2018).

We organize these into four key topics that feature prominently in management literature on the concept of intelligent sales. *Data fusion* is a technical antecedent for the intelligent sales concept – data must be pooled and integrated across sales departments before “downstream” processes may be data-infused. But data fusion reflects a desire to transform organizational life both infrastructurally (i.e., data practices) as well as epistemically (i.e., knowledge practices) by imagining new practices of knowledge generation and the surveillance of distributed, diverse organizational actions such pooled data might make possible. *Responsive customer care* and *automation* are both constitutive elements of the intelligent sales concept that focus on the work-practice level, imagining how the everyday work practices of salespeople might change, whether augmented with data (responsive customer care) or off-loaded onto machine agents (automation). *Streamlined delivery* also constitutes the intelligent sales concept – but instead imagines how it might change operations at an

organizational level, by standardizing and rationalizing various sales processes like pricing or fulfillment. We elaborate on each of these four topics below.

Data Fusion

Like all AI endeavors, data are the lifeblood of the intelligent sales concept. Given the complexity of sales organizations, with various departments involved along the lead-to-cash process, a key premise in intelligent sales is a unified and integrated data lake at the heart of the sales organization. This data lake provides a “single source of truth” about a customer – meaning that different departments can gain visibility into information about what is or has happened with the customer. As a first step (even before AI insights) this data integration is seen to enhance coordination and collaboration across departments. For example, Salesforce’s documentation for *Einstein*, their intelligent sales solution, states the benefits of this type of data integration thusly: “All data are connected so all departments can work together seamlessly sharing data to help each other succeed” (“Artificial Intelligence Technology and Resources” n.d.).

This type of cross-functional coordination and collaboration is envisioned as enabling a wholistic or “global” views of lead-to-cash sub-processes, providing executive management with data-driven measurement on their return-on-investment (ROI) and efficacy. For example, according to Salesforce documentation, marketing efforts (such as emails or adverts sent to customers) can be “traced” through the sales pipeline, measuring if and how different tactics influence conversion (turning a prospect or lead into a paying customer): “Companies want to know if they sent emails to prospects and they’re not coming to their web site, if those marketing efforts are having any impact,” one executive shared in a business press article, “Now, if I can see the prospects are having meetings and exchanging emails with salespeople, I can see the impact of my marketing dollars.” (“Accelerating Sales and Marketing Efforts Through Artificial Intelligence” 2019) (p. 6).

Responsive Customer Care

A core value proposition for intelligent sales is its potential to enrich the customer relationship. One envisioned way for the predictive power of AI to transform sales work practices is creating more *pro-active* (rather than re-active) relationships with customers, as Sinha (2018) writes in a *Future of Customer Engagement and Experience* report. For example, in an article for the magazine *CIO*, Thomas (2019) describes AI that can analyze historical data to predict which customers might churn and why, alerting salespeople to accounts in need of additional care. Even before leads or prospects become customers, intelligent sales is imagined to provide insights, for example, in developing customer profiles, matching sales reps to prospects with similarities to their previous successful accounts (Thomas 2019). “To pursue an opportunity, salespeople

spend extensive time searching for customer information that is typically spread throughout a disparate system, and then go with their “gut feel” or personal network to pursue it,” challenges seen as remedied by an AI solution with integrated data and predictive insights (Sinha 2018). Salespeople have to “make decisions on a daily, or even hourly, basis as to where to focus their time when it comes to closing deals to hit their monthly or quarterly quota,” as Antonio (2018) writes in *Harvard Business Review*. In an intelligent sales ecosystem, the decision on where to focus can be made data-driven through the use of AI: “With AI, the algorithm can compile historical information about a client, along with social media postings and the salesperson’s customer interaction history (e.g., emails sent, voicemails left, text messages sent, etc.) and rank the opportunities or leads in the pipeline according to their chances of closing successfully” (Antonio 2018).

Like many enterprise applications of AI, the intelligent sales concept emphasizes *augmenting* workers rather than replacing them, as Pettey (2018) writes in a *Gartner* report. An intelligent sales solution is seen as offering salespeople a broader and more robust picture of their customers – a data-driven portrait – supporting sales staff in their efforts to tweak tactical approaches (and the offerings they package) to customers’ unique situations. Sales is not solely pecuniary work – sales organizations are driven by both a commercial logic (closing deals and realizing revenue), but also a service ethic (working with customers towards shared goals) (Brokling 2014; McClaren 2000; Oakes 1992). The enhancement of sales processes through the infusion of Big Data and AI are aimed as serving both of these logics – intelligent sales is seen as enriching customer experience in the sales process, as well as those in the sales organization. In a recent survey by *Harvard Business Review*, 82% of respondents at large companies believed AI “has the potential to make human work in sales and marketing more meaningful and valuable” (“Accelerating Sales and Marketing Efforts Through Artificial Intelligence” 2019) (p.1).

In this vein, AI is seen as a much-needed innovation for the daily practices of salespeople, as they struggle to remain a customer’s “trusted advisor” amid increasingly complex and multi-faceted market dynamics (Sinha 2018). For example, a Deloitte case study describes the use of AI in the sales organization of a food & beverage company (“Business Analytics and AI Case Studies” n.d.). The project used various Big Data sources and AI to transform a national-level sales strategy into a “hyper-localized” one, enabling sales teams to tailor their approaches for regional segments/markets and provide more customized offerings to customers that leveraged a “locally-focused, fact-based strategy” (“Business Analytics and AI Case Studies” n.d.).

Automation

Another key theme in the intelligent sales depictions is automation, often through the use of robotic process automation (RPA). A common application of RPA, as

described in *Harvard Business Review*, is in automatically generating the documentation that is required of salespeople after client contact and then entering it into customer relationship management (CRM) systems. Many CRM systems fail to flourish, with many salespeople seeing them as a “necessary evil,” rather than a useful or insightful tool (Edinger 2018; “Why Do CRM Projects Fail” n.d.). A recent *Harvard Business Review* survey reported on two case studies which used RPA to address the problem of client contact documentation – one at Lyft Business, a unit of the ridesharing company focused on the enterprise market, and another at Gainsight, a software and services company (“Accelerating Sales and Marketing Efforts Through Artificial Intelligence” 2019). The sales team at Lyft Business grew rapidly from 20 to 200 people across locations, which made it difficult to understand differences in productivity across accounts and reps (“Accelerating Sales and Marketing Efforts Through Artificial Intelligence” 2019). By automatically collecting data on the sales teams’ activities, the case study reported that the organization was able to run subsequent analyses and detect patterns driving productivity numbers. Sales reps were concerned with the heightened surveillance of their activities, worries reported to have “faded quickly when [the reps] realized...the data could help them make more sales” (“Accelerating Sales and Marketing Efforts Through Artificial Intelligence” 2019) (p.3). Automating data collection, then, is seen as beneficial not only in that it frees salespeople up from tedious and repetitive tasks (e.g., creating CRM entries after every phone call or email from a client) but also in that it produces more data which can be analyzed to provide further insights into the sales process. This was a similar finding in the Gainsight case study, where automating client contact documentation via RPA was reported to have helped reveal trends in top closing accounts: “Gainsight learned, for example, that 94% of its top deals correlated with salespeople scheduling a meeting within the last two weeks before closing” (“Accelerating Sales and Marketing Efforts Through Artificial Intelligence” 2019) (p. 3). Other example applications of automation in sales documentation practices include the automatic transcription and summarization of sales calls and providing real-time “to do” items for salespeople to follow up on, as discussed in a recent *Gartner* report (Panette 2019).

Automation is also seen as transforming client contact itself, for example using chatbot and other AI agents to follow up with prospects and leads, as discussed in recent *Harvard Business Review* and *MIT Technology Review* articles (“AI-powered software robots are getting into the sales business” 2018; Kannan and Bernoff 2019; Kardon 2019; Shaner 2018). Some have noted the ethical issues with automating business activities like this in publications such as *TechTarget*, *Strategic Finance*, and *Business Horizons* (Botelho 2017; Castelluccio 2019; Przegalinska et al. 2019). A recognition of these ethical concerns has led some companies to establish guidelines for the ethical use of chatbots. For example, IBM has issued a “code of ethics for chatbots” which establishes the rights of

customers to know they are interacting with a machine and not a human (Reddy 2017). Other issues include intellectual property (IP) and the ownership of information shared with a chatbot, as well as safeguarding privacy of personal data shared with a chatbot (Reddy 2017).

Streamlined Delivery

Streamlined delivery is another thread in intelligent sales rhetoric, which envisions the transformation of various dimensions of offering selection, pricing, contracting, and fulfillment. Through the use of AI, for example, strategic recommendations could be made on upsell opportunities and prices could be optimized for individual customer accounts (Antonio 2018). Furthermore, AI forecasting can help sales managers more effectively leverage and coach sales teams across open accounts (“Accelerating Sales and Marketing Efforts Through Artificial Intelligence” 2019; Antonio 2018). Text-generating AI can also be used to create contract drafts automatically (Lohr 2017; Rich 2018), while other forms of natural language processing (NLP) can help manage changes and discrepancies over the contract lifecycle (Gabbard 2019).

The online retailer Amazon has famously disrupted the eCommerce industry by using AI in every aspect of their business model, including order fulfillment and shipping (Karlinsky 2019; Terdiman 2018). But more traditional retailers also leverage AI to gain real-time predictions on inventory and shipping costs. For example, in a *IBM Blog* article, Geoffroy (2019) writes about an REI engagement with Watson that aimed to implement an AI-driven order management system to provide real-time store inventory data, as well as distribution center inventory data, to customers. Instead of losing sales due to out-of-stock inventory, the article states the AI system helped the company “see an increase of \$100 million in sales the first year, in addition to a 27% lift in ‘buy online, pickup in store’ orders” (Geoffroy 2019). In business-to-business (B2B) settings, delivery and fulfillment are also tied closely with supply chain and logistics. As this note mentioned earlier, these two domains (sales/marketing, and supply chain/logistics) are seen as the most promising enterprise application areas for AI (Chui et al. 2018).

AI Ethics in Intelligent Sales

All workers encounter ethical dilemmas in their everyday work practice. The ethical dilemmas faced by salespeople can be unique, though, due to their role as “boundary-role performers,” meaning their work practices maneuver between the sales organization and that organization’s customers (Weeks and Nantel 1992). Navigating between various stakeholders (e.g., prospects, customers, competitors, managers, executives, order management, and delivery support etc) often creates a dynamic and high-adrenaline daily practice, which can lead to increased

workplace stress and tension, causing a heightened affective experience and perception of ethical dilemmas (Weeks and Nantel 1992). This ambivalent dynamism, along with the twinned logics of commerce and service that drive sales work, can create a fragmented experience for salespeople and the feeling of being pulled in many directions (Brokling 2014; McClaren 2000; Oakes 1992). Introducing emergent technologies (like Big Data and AI) into sales work, then, heightens an already intense – and ethically-laden – workplace experience.

The past few years have seen a rapid growth of communities of scholars investigating the fairness, accountability, and transparency (FAccT*) of AI systems. For example, a number of workshops have been held at leading AI conferences such as AAAI, NeurIPS, and ICML. The growth of these communities is motivated by a recognition that AI systems can and do have great societal power, which requires careful consideration on their sociotechnical considerations and repercussions. Concerns on the ethical, humane, and intuitive use of emergent technologies is of course not new to those in the social computing field. Indeed, the contemporary AI renaissance – and the ethical issues raised in its wake – provides a pressing opportunity for social computing researchers to work together with technical AI community to design just and human AI futures (Loi et al. 2019; Wolf et al. 2018).

The ethical questions around Big Data and AI are sociotechnical (boyd and Crawford 2012; Shin and Choi 2015). This means that what we have come to know as “Big Data” or “AI” is comprised of and steered by both technological and social actors, and importantly the entangled interplay amongst them (boyd and Crawford 2012). The datasets used to train AI models, the algorithms developed to organize them, the compute power and cloud computing that makes such data transformations possible, the interfaces that display model outputs – they cannot be understood or analyzed without also looking to the developers who clean data and write code, the subject-matter experts who label data and shape training sets, the project managers and executives who oversee applied AI projects and re-design business processes, the workers who touch AI outputs, and so on (Wolf 2020). The design of the overarching AI ecosystem – and in particular, the ways in which business processes are re-designed – are also ethical questions, of worthy scrutiny alongside those which are more technologically-focused (Martin 2019; Wolf and Blomberg 2019). How will an AI system be integrated into organizational practice and what roles and responsibilities will it re-configure? How will these changes re-design the nature and experience of work?

Next, we highlight two themes with ethical valence: invisible actors (training data) and explainability/literacy (use of AI outputs in everyday work practices). While we organize these into two themes to provide analytical traction, we recognize the concerns they surface are dynamic and fluid, overlapping in situated practice.

Invisible Actors – Training Data

Recent work in the FAccT* community has called attention to the important role that training data plays in the development of AI systems (Gebru et al. 2020; Holland et al. 2018; Mitchell et al. 2019). While theoretical research and algorithm development in the AI community typically uses standardized “off the shelf” datasets (such as MNIST¹ the dataset of hand-written digits, or LISA² photographs of road signs), in applied AI projects, real-world training sets derived from the domain setting are used (Wolf 2019a). There is considerable data engineering work (commonly called data pre-processing) that AI/ML developers and data scientists must undertake to ready real-world training sets, which are often messy, noisy, and incomplete (Wolf 2019b, 2019a).

As we have outlined above, there are many potential applications of AI to various dimensions of the overarching lead-to-cash sales process – from predictive pricing and offering recommendations, to automated recordkeeping, client interactions, and content creation. Across these various use cases, an important question to ask is: what data are these AI models trained on? Understanding a model’s training set is important in not only assessing whether it encompasses a representative sample of the domain and actions of interest; it is also important in appropriately bounding subsequent outputs and any insights they might offer. AI models are statistical representations of their training set – thus the training set becomes the “universe” of information the model considers and accordingly acts upon.

Who and what is represented in the training sets that underly AI sales applications? Ethical concerns over training data are absent from current management discourses on intelligent sales. While others have pointed to ethical issues with Big Data use in marketing (particularly digital advertising (Kitchin 2014; Nunan and Di Domenico 2013; Pasquale 2015; Zwitter 2014)), the intelligent sales concept sees marketing blending with sales through cross-department data fusion, raising new questions around ethics and transparency that require careful scrutiny.

In B2B settings, whose behaviors and actions are captured and used to train AI models? Take a hypothetical scenario in the retail industry, where an IT services firm (the sales organization in this scenario) is selling a cognitive service to a clothing retailer that monitors social media activity (e.g., clothing worn in photos by people in particular market segments of interest, like women ages 20-27) and predicts regional fashion trends for that segment (insights useful to the retailer’s warehousing and fulfillment). Such a scenario raises issues of consent and privacy around the initial data collection, of course; but it also raises concerns over

¹ <http://yann.lecun.com/exdb/mnist/>

² <http://cvrr.ucsd.edu/LISA/lisa-traffic-sign-dataset.html>

transparency in whose behaviors and actions are depicted in the training set (i.e., representativeness) and also whose subsequent behaviors and actions the trained model is used to predict (i.e., training/target alignment). Do the ethical stakes change if the industry is in the food/beverage sector? Pharmaceuticals? What about tobacco products?

Furthermore, who has visibility into these pieces of the AI pipeline and to whom are discussions of their ethical considerations owed? Many real-world applications of AI leverage pre-trained models, often accessed through “model marketplaces” made possible via application programming interfaces (APIs) (“AI Market Leaders Join Forces to Release Modzy™, an AI Platform and Model Marketplace, Introducing Choice, Scale, and Security” 2019; “Democratizing Data Science in Your Organization” 2018; Pscheid 2018). This creates distance between model creation and model use, rendering important questions on what is made visible at different points in the AI pipeline and what becomes obscured. To what extent does a data science team picking up and bootstrapping off a pre-trained model know the origins and composition of the model’s training set? To what extent can (and should) they be able to describe the model’s underpinning to their subsequent users? As noted above, in creating their “chatbot code of ethics,” IBM listed the right of customers to know they are interacting with a virtual agent (instead of a human being). How might such a right to visibility play out in scenarios where pre-trained models are used? Careful consideration is needed to more fully think-through the ethical questions raised in applied AI projects, such as intelligent sales.

Explainability/Literacy – Use of AI outputs in Everyday Work Practices

The ability for humans to make sense of and reason about algorithmic, AI actions – the “explainability” of AI – is an important topic of research that has flourished in recent years (Adadi and Berrada 2018; Guidotti et al. 2018). Explainability is a particular concern in AI systems which use black-box modelling techniques, like Deep Learning, that model data in extremely high dimensionality. This complexity makes it difficult to explain model actions in ways that humans can comprehend. This has led to a fast-growing area of technical research known as XAI (eXplainable AI). Despite significant technical progress in XAI, a gap exists between explainability of AI from an academic or research point of view and the types of explainability requirements that arise in real-world AI deployments and settings of end-use (Wolf 2019c). A number of concerns arise in business settings, where AI capabilities may be infused in one sub-process – deeply embedded within, and perhaps invisible to downstream actors in an overall workflow (as discussed above).

For example, consider the scenario where AI is embedded into a business decision-making point, a typical use case for enterprise AI. Interacting with the

AI output requires a certain level of technical literacy – the business user must be able to understand and assess the AI model’s output in the context of the business decision at hand. Does this mean all workers must now have a basic understanding of AI in order to appropriately use cognitive services in their everyday work practices? A further consideration is the collaborative nature of many work practices. Is only one worker involved in the decision-making process? If others are involved, do they also interact with the AI output or must they rely on the “user” (the person who actually touches and interacts with the output)? What new forms of translation work are required on the part of the user to adequately convey the technical details of the AI system in others involved in the decision-making process?

Further questions around explainability and literacy also arise in relation to the sales-customer relationship. For example, does a salesperson need to be able to explain to the customer why the system has made a particular recommendation – whether it be the prediction of churn or a customized offering layout? Are sales staff to disclose the use of an AI system to a customer when presenting a bid? Does a customer have a “right to know,” as we saw articulated in the chatbot code of ethics? What about in self-service situations, where the customers might be browsing a catalogue of offerings and experience directly AI-driven customization? Such questions remain ripe for ethical interrogation. Such considerations pose open questions around labor transformations in almost every imagined enterprise application of AI.

Another concern around the intelligibility of AI systems is the “freshness” or currency/relevancy of data streams. Related to the issue of training data (whose data are used in training sets and what activities are being modelled) are questions of when is training data collected and modelled. As AI systems are deployed into settings of end use, their deployment can influence the object of interest – similar to the adage “the simple fact of observation changes behavior,” introducing an AI model into a situation also changes it. Questions of freshness – of the “old” data used to train the model, as well as the ongoing alignment between the model’s predictive targets and the phenomena of interest, require careful attunement.

Further, downstream in an applied AI pipeline, where are input data being pulled from (i.e., the new pieces of data that will be run through the AI model to make some sort of prediction) and are such input streams current? In the sale context, questions of currency and relevancy can be particularly tricky given the cyclical nature of sales – in B2B settings, these cycles are driven often by procurement and budgetary concerns, and in consumer settings, they are often driven by things like seasons and holidays, as well as inflation and labor market conditions. How do these concerns factor into the ways in which input data are collected and predictions are subsequently made? How should they frame the potential insights such predictions might offer?

What makes the explainability and literacy of enterprise AI systems an ethical concern? As practice-oriented CSCW scholars, we are concerned with everyday practices – and the ways in which those practices are aligned, integrated, burdened, displaced, or replaced by technological systems. In our endeavors to understand and influence the entangled nature of practices and technologies, we are beckoned to interrogate the societal impacts of the systems we build and deploy. Workplace technologies can have a “dark side” that we must not ignore (Nauwerck and Cowen Forssell 2018). As Walsham (2012) provocatively asks: are we making a better world with ICTs? Applying this question to the topic of enterprise AI, we might ask ourselves: are we making a better workplace? Narratives around enterprise AI often emphasize time savings and the relief that machine action and automation might provide to workers from burdensome or boring tasks. But workplace stress and burnout are global problems in the contemporary experience of work (Carod-Artal and Vázquez-Cabrera 2013). “Do more with less” is often experienced as an oppressive imperative of organizational life, not simply a motivational aspiration. Does adding complex, statistically-driven AI systems into work processes (and the sensemaking practices it demands of workers) increase the risk of workplace stress or burnout? Is the re-skilling required to effectively work with AI systems adequately scaffolded and supported, through the organizational provisioning of appropriate training materials, as well as the time to actually study and learn them? The explainability and literacy of enterprise AI systems raises important questions on the nature and experience of everyday work, as well as its implications for workplace empowerment and well-being.

Conclusion

In this note, we have outlined an emergent management concept known as “intelligent sales,” which sets out a vision of transforming sales processes through the use of Big Data and AI capabilities. We have surfaced a number of ethical issues that case of intelligent sales raises, particularly: invisible actors (training data) and explainability/literacy (use of AI outputs in everyday work practices).

These issues – and indeed, the broader ethical landscape of intelligent sales – requires further investigation and analysis. This note is intended as a starting place, provoking further attention within the CSCW community on the fairness, accountability, and transparency (FAT*) of algorithmic systems by raising a number emergent in AI applications in enterprise settings.

In addition to further research by social computing scholars and researchers, we also note the need for collaboration around these issues that is both social and technical. Given that intelligent sales is a strategic management vision for the future of organizations, the management community needs to be engaged in shaping conversations around the ethical (and unethical) uses of Big Data and AI

capabilities. Foundational research on ethical behavior in sales work more broadly has found efficacy in the posting clear ethical guidelines for sellers – that is, when an organization makes clear the ethical bounds for sellers, sellers found ethical dilemmas more manageable and less distressing, leading to increased ethical behavior (Weeks and Nantel 1992). While these insights draw on sales work generally (not specific to Big Data or AI) they are instructive in how we can shape the future of intelligent sales today – through clear guidance on the ethical use of such systems and the data-driven insights they might provide. In presenting the case of intelligent sales, and several ethical issues it raises, this note hopes to start a conversation within our community on how to work towards such a goal.

Acknowledgments

Thank you to the reviewers for helping refine this work. Thanks also to Mary Roth for early discussions on this topic. All opinions expressed herein are my own and do not reflect any institutional endorsement.

References

- Accelerating Sales and Marketing Efforts Through Artificial Intelligence. (2019, September 19). *Harvard Business Review*.
<https://hbr.org/sponsored/2019/09/accelerating-sales-and-marketing-efforts-through-artificial-intelligence>. Accessed 6 January 2020
- Adadi, Amina; and Mohammed Berrada. (2018). Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). *IEEE Access*, vol. 6, 2018, pp. 52138–52160.
- Adams, Richard J.; Palie Smart; and Anne Sigismund Huff. (2017). Shades of Grey: Guidelines for Working with the Grey Literature in Systematic Reviews for Management and Organizational Studies. *International Journal of Management Reviews*, vol. 19, no. 4, 2017, pp. 432–454.
- AI Market Leaders Join Forces to Release Modzy™, an AI Platform and Model Marketplace, Introducing Choice, Scale, and Security. (2019, November 4). *Bloomberg.com*. <https://www.bloomberg.com/press-releases/2019-11-04/ai-market-leaders-join-forces-to-release-modzytm-an-ai-platform-and-model-marketplace-introducing-choice-scale-and-security>. Accessed 12 January 2020
- AI-powered software robots are getting into the sales business. (2018, August 21). *MIT Technology Review*. <https://www.technologyreview.com/f/611922/ai-powered-software-robots-are-getting-into-the-sales-business/>. Accessed 6 January 2020

- Antonio, Victor. (2018, July 30). How AI Is Changing Sales. *Harvard Business Review*. <https://hbr.org/2018/07/how-ai-is-changing-sales>. Accessed 6 January 2020
- Artificial Intelligence Technology and Resources: Salesforce Einstein. (n.d.). *Salesforce.com*. <https://www.salesforce.com/products/einstein/overview/>. Accessed 12 January 2020
- Bear, Laura. (2009). *Industrial Work and Life: An Anthropological Reader*. (M. Mollona, G. D. Neve, and J. Parry, Eds.). Oxford England ; New York: Berg Publishers.
- Botelho, Bridget. (2017, April 5). Chatbot technology raises ethical questions. *TechTarget*. <https://searchcustomerexperience.techtarget.com/opinion/Chatbot-technology-raises-ethical-questions>. Accessed 6 January 2020
- boyd, danah; and Kate Crawford. (2012). Critical Questions for Big Data. *Information, Communication & Society*, vol. 15, no. 5, June 2012, pp. 662–679.
- Braverman, Harry. (1998). *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*. New York: Monthly Review Press.
- Brokling, Kimberly. (2014). *Taking the slime out of sales: Discursive resources and identity construction among salespeople* (Master's Thesis). University of Colorado. Retrieved from <https://search.proquest.com/openview/ec245b67338654d74f0045d21fffa2bb/1?pq-origsite=gscholar&cbl=18750&diss=y>
- Business Analytics and AI Case Studies. (n.d.). *Deloitte United States*. <https://www2.deloitte.com/us/en/pages/deloitte-analytics/articles/business-analytics-case-studies.html>. Accessed 12 January 2020
- Carod-Artal, Francisco Javier; and Carolina Vázquez-Cabrera. (2013). Burnout Syndrome in an International Setting. In S. Bährer-Kohler (Ed.), *Burnout for Experts: Prevention in the Context of Living and Working*. Boston, MA: Springer US.
- Castelluccio, Michael. (2019, December). Creating Ethical Chatbots. *Strategic Finance*.
- Chui, Michael; Nicolaus Henke; and Mehdi Miremadi. (2018, July 20). Most of AI's Business Uses Will Be in Two Areas. *Harvard Business Review*. <https://hbr.org/2018/07/most-of-ais-business-uses-will-be-in-two-areas>. Accessed 6 January 2020

- Democratizing Data Science in Your Organization. (2018, April 5). *Harvard Business Review*. <https://hbr.org/sponsored/2019/04/democratizing-data-science-in-your-organization>. Accessed 12 January 2020
- Desrosières. (2002). *The Politics of Large Numbers: A History of Statistical Reasoning*. Cambridge, Mass.: Harvard University Press.
- Edinger, Scott. (2018, December 20). Why CRM Projects Fail and How to Make Them More Successful. *Harvard Business Review*. <https://hbr.org/2018/12/why-crm-projects-fail-and-how-to-make-them-more-successful>. Accessed 6 January 2020
- Gabbard, Jason. (2019, October 24). Insight: AI-Powered Contracts Can Help Legal Departments Avoid Pitfalls. *Bloomberg Law*. <https://news.bloomberglaw.com/tech-and-telecom-law/insight-ai-powered-contracts-can-help-legal-departments-avoid-pitfalls>. Accessed 6 January 2020
- Gebru, Timnit; Jamie Morgenstern; Briana Vecchione; Jennifer Wortman Vaughan; Hanna Wallach; Hal Daumeé III; and Kate Crawford. (2020). Datasheets for Datasets. *arXiv:1803.09010 [cs]*, January 2020. <http://arxiv.org/abs/1803.09010>. Accessed 12 January 2020
- Geoffroy, Jeff. (2019, May 20). Does my fulfillment operation really need AI? *IBM Blog: Watson*. <https://www.ibm.com/blogs/watson-customer-engagement/2019/05/20/does-my-fulfillment-operations-really-need-ai/>. Accessed 6 January 2020
- Guidotti, Riccardo; Anna Monreale; Salvatore Ruggieri; Franco Turini; Fosca Giannotti; and Dino Pedreschi. (2018). A Survey of Methods for Explaining Black Box Models. *ACM Comput. Surv.*, vol. 51, no. 5, August 2018, pp. 93:1–93:42.
- Holland, Sarah; Ahmed Hosny; Sarah Newman; Joshua Joseph; and Kasia Chmielinski. (2018). The Dataset Nutrition Label: A Framework To Drive Higher Data Quality Standards. *arXiv:1805.03677 [cs]*, May 2018. <http://arxiv.org/abs/1805.03677>. Accessed 12 January 2020
- IBM Institute for Business Value. (2018). *The Cognitive Enterprise: Reinventing your company with AI*. <https://www.ibm.com/thought-leadership/institute-business-value/report/cognitive-enterprise>. Accessed 12 January 2020
- Kannan, P. V.; and Josh Bernoff. (2019, May 21). Does Your Company Really Need a Chatbot? *Harvard Business Review*. <https://hbr.org/2019/05/does-your-company-really-need-a-chatbot>. Accessed 6 January 2020
- Kardon, Brian. (2019, February 13). Five AI Solutions Transforming B2B Marketing. *MIT Sloan Management Review*.

<https://sloanreview.mit.edu/article/five-ai-solutions-transforming-b2b-marketing/>. Accessed 6 January 2020

- Karlinsky, Neal. (2019, June 5). How artificial intelligence helps Amazon deliver. *Day One: The Amazon Blog*. <https://blog.aboutamazon.com/innovation/how-artificial-intelligence-helps-amazon-deliver>. Accessed 6 January 2020
- Kitchin, Rob. (2014). *The Data Revolution: Big Data, Open Data, Data Infrastructures and Their Consequences*. Thousand Oaks, CA: SAGE Publications Ltd.
- Kling, Rob; and Suzanne Iacono. (1988). The Mobilization of Support for Computerization: The Role of Computerization Movements. *Social Problems*, vol. 35, no. 3, June 1988, pp. 226–243.
- Lohr, Steve. (2017, March 19). A.I. Is Doing Legal Work. But It Won't Replace Lawyers, Yet. *The New York Times*. <https://www.nytimes.com/2017/03/19/technology/lawyers-artificial-intelligence.html>. Accessed 6 January 2020
- Loi, Daria; Christine T. Wolf; Jeanette L. Blomberg; Raphael Arar; and Margot Brereton. (2019). Co-designing AI Futures: Integrating AI Ethics, Social Computing, and Design. In S. Harrison, S. Bardzell, C. Neustaedter, and D. G. Tatar (Eds.), *Companion Publication of the 2019 on Designing Interactive Systems Conference, DIS 2019, San Diego, CA, USA, June 23-28, 2019*. ACM, pp. 381–384.
- Martin, Kirsten. (2019). Ethical Implications and Accountability of Algorithms. *Journal of Business Ethics*, vol. 160, no. 4, December 2019, pp. 835–850.
- McClaren, Nicholas. (2000). Ethics in Personal Selling and Sales Management: A Review of the Literature Focusing on Empirical Findings and Conceptual Foundations. *Journal of Business Ethics*, vol. 27, no. 3, October 2000, pp. 285–303.
- Mitchell, Margaret; Simone Wu; Andrew Zaldivar; Parker Barnes; Lucy Vasserman; Ben Hutchinson; et al. (2019). Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency, FAT* 2019, Atlanta, GA, USA, January 29-31, 2019*. ACM, pp. 220–229.
- Nauwerck, Gerolf; and Rebecka Cowen Forssell. (2018). The Digital Work Environment—a Challenge and an Opportunity for CSCW, 2018.
- Nunan, Daniel; and MariaLaura Di Domenico. (2013). Market Research and the Ethics of Big Data. *International Journal of Market Research*, vol. 55, no. 4, July 2013, pp. 505–520.

- Oakes, Guy. (1992). *The soul of the salesman: the moral ethos of personal sales*. Atlantic Highlands, N.J.: Humanities Press International.
- Panette, Kasey. (2019, February 5). The CIO's Guide to Artificial Intelligent. *Gartner*. <https://www.gartner.com/smarterwithgartner/the-cios-guide-to-artificial-intelligence/>. Accessed 12 January 2020
- Pasquale, Frank. (2015). *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge: Harvard University Press.
- Pettey, Christy. (2018, February 9). Lessons From Artificial Intelligence Pioneers. *Gartner*. <https://www.gartner.com/smarterwithgartner/lessons-from-artificial-intelligence-pioneers/>. Accessed 12 January 2020
- Przegalinska, Aleksandra; Leon Ciechanowski; Anna Stroz; Peter Gloor; and Grzegorz Mazurek. (2019). In bot we trust: A new methodology of chatbot performance measures. *Business Horizons*, vol. 62, no. 6, November 2019, pp. 785–797.
- Pscheid, Julian. (2018, September 18). API AI Examples: Machine Learning in Enterprise App Development. *EMERGE | Digital Product Agency*. <https://www.emergeinteractive.com/insights/detail/api-ai-examples-enterprise-apps/>. Accessed 12 January 2020
- Reddy, Trips. (2017, October 15). The code of ethics for AI and chatbots that every brand should follow. *IBM Blogs: Watson*. <https://www.ibm.com/blogs/watson/2017/10/the-code-of-ethics-for-ai-and-chatbots-that-every-brand-should-follow/>. Accessed 6 January 2020
- Rich, Beverly. (2018, February 12). How AI Is Changing Contracts. *Harvard Business Review*. <https://hbr.org/2018/02/how-ai-is-changing-contracts>. Accessed 6 January 2020
- Schoemaker, Paul J.H.; and Philip E. Tetlock. (2017, Spring). Building a More Intelligent Enterprise. *MIT Sloan Management Review*. <https://sloanreview.mit.edu/article/building-a-more-intelligent-enterprise/>. Accessed 12 January 2020
- Shaner, Jared. (2018, January 4). Are Bots The Future Of Sales? *Forbes*. <https://www.forbes.com/sites/forbesbusinessdevelopmentcouncil/2018/01/04/are-bots-the-future-of-sales/>. Accessed 6 January 2020
- Shin, Dong-Hee; and Min Jae Choi. (2015). Ecological views of big data: Perspectives and issues. *Telematics and Informatics*, vol. 32, no. 2, May 2015, pp. 311–320.
- Sinha, Swati. (2018, May 18). Intelligent sales lead to cash. *The Future of Customer Engagement and Experience*. <https://www.the-future-of->

- commerce.com/2018/05/18/intelligent-sales-lead-to-cash/. Accessed 5 January 2020
- Taylor, Frederick Winslow. (2003). *Scientific management*. London: Routledge. <http://site.ebrary.com/id/10166472>. Accessed 12 January 2020
- Terdiman, Daniel. (2018, October 5). How AI is helping Amazon become a trillion-dollar company. *Fast Company*. <https://www.fastcompany.com/90246028/how-ai-is-helping-amazon-become-a-trillion-dollar-company>. Accessed 6 January 2020
- Thomas, Brian E. (2019, December 2). AI is already disrupting sales, today. *CIO*. <https://www.cio.com/article/3481741/ai-is-already-disrupting-sales-today.html>. Accessed 6 January 2020
- Walsham, Geoff. (2012). Are We Making a Better World with Icts? Reflections on a Future Agenda for the IS Field. *Journal of Information Technology*, vol. 27, no. 2, June 2012, pp. 87–93.
- Weeks, William A.; and Jacques Nantel. (1992). Corporate codes of ethics and sales force behavior: A case study. *Journal of Business Ethics*, vol. 11, no. 10, October 1992, pp. 753–760.
- Why Do CRM Projects Fail. (n.d.). *Salesforce.com*. <https://www.salesforce.com/hub/crm/why-do-crm-projects-fail/>. Accessed 6 January 2020
- Wolf, Christine; and Jeanette Blomberg. (2019). Evaluating the Promise of Human-Algorithm Collaborations in Everyday Work Practices. *Proc. ACM Hum.-Comput. Interact.*, vol. 3, no. CSCW, November 2019, pp. 143:1–143:23.
- Wolf, Christine T. (2019a). Professional Identity and Information Use: On Becoming a Machine Learning Developer. In N. G. Taylor, C. Christian-Lamb, M. H. Martin, and B. Nardi (Eds.), *Lecture Notes in Computer Science. Information in Contemporary Society*. Cham: Springer International Publishing, pp. 625–636.
- Wolf, Christine T. (2019b). Conceptualizing Care in the Everyday Work Practices of Machine Learning Developers. In *DIS '19 Companion. Companion Publication of the 2019 on Designing Interactive Systems Conference 2019 Companion*. New York, NY, USA: ACM, pp. 331–335.
- Wolf, Christine T. (2019c). Explainability scenarios: towards scenario-based XAI design. In W.-T. Fu, S. Pan, O. Brdiczka, P. Chau, and G. Calvary (Eds.), *Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI 2019, Marina del Ray, CA, USA, March 17-20, 2019*. ACM, pp. 252–257.

Wolf, Christine T. (2020). AI Models and Their Worlds: Investigating Data-Driven, AI/ML Ecosystems Through a Work Practices Lens. In A. Sundqvist, G. Berget, J. Nolin, and K. I. Skjerdingsstad (Eds.), *Lecture Notes in Computer Science. Sustainable Digital Communities*. Cham: Springer International Publishing, pp. 651–664.

Wolf, Christine T.; Haiyi Zhu; Julia Bullard; Min Kyung Lee; and Jed R. Brubaker. (2018). The Changing Contours of “Participation” in Data-driven, Algorithmic Ecosystems: Challenges, Tactics, and an Agenda. In V. Evers, M. Naaman, G. Fitzpatrick, K. Karahalios, A. Lampinen, and A. Monroy-Hernández (Eds.), *Companion of the 2018 ACM Conference on Computer Supported Cooperative Work and Social Computing, CSCW 2018, Jersey City, NJ, USA, November 03-07, 2018*. ACM, pp. 377–384.

Zwitter, Andrej. (2014). Big Data ethics. *Big Data & Society*, vol. 1, no. 2, July 2014, pp. 2053951714559253.

Submitted in February 2020, revised in April 2020, accepted in May 2020.