Journal of Strategic Information Systems 29 (2020) 101614

Contents lists available at ScienceDirect

Journal of Strategic Information Systems

journal homepage: www.elsevier.com/locate/jsis

Augmenting the algorithm: Emerging human-in-the-loop work configurations

Tor Grønsund[a,*], Margunn Aanestad[b,a]

[a] Department of Informatics, University of Oslo, Oslo, Norway
[b] Department of Information Systems, University of Agder, Kristiansand, Norway

ARTICLE INFO

Keywords:
Digital work
Automation
Augmentation
Human–machine configurations
Human-in-the-loop
Artificial intelligence (AI)

ABSTRACT

How do configurations of humans and algorithms evolve as firms adopt artificial intelligence (AI) capabilities, and what are the implications for work and organization? We explored these questions through a two-year long case study of an organization in the international maritime trade that introduced automated algorithmic support for data analysis and prediction work. Drawing on a human–machine configuration perspective, we found that humans and the algorithm were configured and reconfigured in multiple ways over time as the organization dealt with the introduction of algorithmic analysis. In contrast to replacing human work, the emergent configurations required new roles and redistribution of extant expertise to augment and improve the accuracy of the algorithm. Our analysis suggests that the new configuration resembled a human-in-the-loop pattern, comprised of both the augmentation work of *auditing* (i.e. the generation of a ground truth and assessment of the algorithmic output against this) as well as the work of *altering* the algorithm and the data acquisition architecture. Our research points to the strategic importance of a human-in-the-loop pattern for organizational reflexivity to ensure that the performance of the algorithm meets the organization's requirements and changes in the environment.

Introduction

Advances in 'data-driven intelligence engines' (Baptista et al., 2017a) based on availability of data and artificial intelligence (AI) have the potential to transform the nature of work and organizations (Newell and Marabelli, 2015). Information Systems (IS) researchers have raised fundamental questions as to the ramifications of both datafication processes (Galliers et al., 2017, Günther et al., 2017; Markus, 2017) and intelligent technologies (Bailey and Barley, 2019; Faraj et al., 2018; von Krogh, 2018; Rai et al., 2019), including the implications for organizational strategy (Constantiou and Kallinikos, 2015; Woerner and Wixom, 2015; Bhimani, 2015) and work in settings ranging from education (Marjanovic and Cecez-Kecmanovic, 2017), policing (Waardenburg et al., 2018), recruitment (Broek et al., 2019), and medical diagnosing (Lebovitz et al., 2019).

Automation and augmentation figure centrally yet often at either end of the human–machine spectrum in current debates on organizational uptake of AI (Raisch and Krakowski, 2020). Data-powered algorithms allows automation of cognitive, discretionary, and decision-making tasks that humans used to perform. We currently see a significant deployment of simpler types of automation such as Robotic Process Automation tools (based on rule-based engines), and speech recognition and conversational agents used for first line customer interaction (based on Natural Language Processing techniques). More advanced algorithmic automation may

* Corresponding author.
E-mail addresses: torrg@ifi.uio.no (T. Grønsund), margunn.aanestad@uia.no (M. Aanestad).

https://doi.org/10.1016/j.jsis.2020.101614

Available online 24 June 2020

include machine learning, computation, and statistical techniques that "rely on large data sets to generate responses, classifications, or dynamic predictions that resemble those of a knowledge worker" (Faraj et al., 2018, p. 62). These are still relatively rarer (Brynjolfsson and Mitchell, 2017), nevertheless, popular debates tend to profess a "replacement" trope, where human jobs are taken over by intelligent machines (McAfee and Brynjolfsson, 2017).

More nuanced accounts instead emphasize interrelation and augmentation (Bailey and Barley, 2019; von Krogh, 2018; Markus, 2017; Davenport and Kirby, 2016), where humans and algorithms interact to perform a task. While automation of a human worker's tasks may result in replacement of the human worker altogether, more often there is a partial automation of specific tasks, resulting in a division of labour between the human and the technology, where novel tasks also emerge and ensure a continued need for the human worker. In addition, one may see indirect effects where not only does automation substitute labour, it also "complements labor, raises output in ways that lead to higher demand for labor, and interacts with adjustments in labor supply" (Autor, 2015, p. 5). For instance, by automating routine manual and cognitive tasks, human capacity is freed to take on non-routine tasks and capacity and productivity can increase, which again drives an increase in demand for non-routine problem-solving (Nedelkoska and Quintini, 2018).

In this paper we address these questions with a specific focus on automating/augmenting organizational data analysis and prediction work. Algorithmic decision-making is based on data that are captured through digitized devices and processed by algorithms. The aim is generally towards predicting objects' (future) behavior based on their current or past behavior (Newell and Marabelli, 2015, p. 4). As a core task underlying such work, prediction refers to taking data or information you do have to generate information you did not previously have (Agrawal et al., 2019, p. 1), including assignment of categorical class labels (classification) and future events (prediction). The capabilities of utilizing data has in itself become a significant strategic asset and an object of strategy formulation (DalleMule and Davenport, 2017; Iansiti and Lakhani, 2020), and data-driven decision-making is increasingly implemented (Brynjolfsson and McElheran, 2016). The strategic value of data may lead organizations to automate and scale up data-driven analysis using algorithmic technologies, as the amounts and complexity of novel data surpass the organization's existing analysis capacity. At the same time, algorithmic automation requires quality data input to be useful to organizations and thus might require human judgement and verification (von Krogh, 2018). With the possibility of a significant change of work, organizing and strategizing, it is pertinent that IS researchers examine what this entails for work. Such studies could provide research-based knowledge to organizations that consider how to respond strategically to the new influx of data-intensive and algorithmic technologies.

We have conducted a qualitative case study of an organization that introduced algorithmic support for its analysis of geospatial satellite data and contextual data sources. The site of our field work was a global ship brokering company that intermediates exchanges between ship owners and charterers and/or cargo owners. This business relies on vast up-to-date data and analysis on maritime activity and market developments to predict ship behavior and, in turn, trade flows. Here, a central data source is the global Automatic Identification System (AIS), which contains information about ships' identity, location, and cargo. In recent years, AIS data has become more openly available and accessible, and thus the information asymmetries that the brokers traditionally could capitalize on were eroded. Facing a risk of disintermediation, the company responded with a strategic agenda to generate value-added analyses, through ramping up its data acquisition and introducing algorithmic support for data analysis in order to improve its accuracy and timeliness. We have followed the process of introducing the algorithm to the fabric of the organization over two years and are interested to understand the distribution of work tasks and roles between the algorithm and humans, its organizing and strategic significance.

Our guiding research question is thus: *how do configurations of humans and algorithms evolve as a firm develops AI capabilities, and how do these configurations impact work and the organization?* To answer this, we describe the changes in the organization as it sought to integrate algorithmic data processing and classification into its value creation activities, both in individual worker's tasks, the emergence of new tasks, roles, and capabilities, and the ongoing configuration of the work organization.

The article is structured as follows. We start with a review of the literature on automation and augmentation by algorithmic technologies. Here we articulate the notion of human–machine configuration (Suchman, 2007) as our main analytical frame. We develop the specific configuration of "human-in-the-loop" as particularly relevant to unpack the emergence and nature of augmentation work. We then describe our methodological approach and present our case study of the ship brokering firm. The empirical material is presented according to a chronological structure representing shifting configurations of human–machine work. In the discussion, we elaborate on the novelty and nature of human–machine augmentation work and point to human-in-the-loop configuration as a new form of organizing. Following this, we adopt the notion of reflexivity to frame this configuration as a strategic, dynamic capability.

Related research and analytic framework

Algorithmic automation, augmentation, and configurations

Automation is not an all-or-nothing phenomenon. Firstly, automation can be applied to different functions of an organization including information acquisition, analysis, decision, and action. Secondly, automation may vary along a continuum from fully manual (i.e., human) performance to fully automatic performance (Parasuraman et al., 2000). Thirdly, across this continuum and these functions, automation may augment human performance of a particular task through increasing the capability of humans to approach a problem situation, gain comprehension, and derive solutions (Engelbart, 1962, p. 1; Licklider, 1960). AI-enabled automation, then, "may range from task substitution (AI substitutes humans) to task augmentation (AI and humans complement one and another) to task assemblage (AI and humans are dynamically brought together to function as an integrated unit)" (Rai et al., 2019, p. iv). von Krogh (2018) asserts that it is in situations of problem-solving (algorithms providing alternative courses of action to resolve a problem), rather than decision-making (conclusions reached by the algorithms based on the data available), that algorithms appear to augment rather than substitute humans in their performance of tasks. Especially in complex situations where solutions may benefit from the relative strengths of humans and algorithms, an "assemblage" of team problem-solving and automated solution generation seems advantageous (von Krogh, 2018, p. 407). Indeed, a widespread assumption is that "computers and human beings have complementary strengths and problem-solving capabilities […] and human beings generally outperform machines when dealing with ambiguity, vagueness and incomplete information" (Pavlou, 2018, p. 44-45). It has long been suggested that algorithms are more appropriate for structured problem-solving and where outcomes can be clearly described; and human processing is more apt for ill-structured problem-solving, to which solutions tend to require human interpretation and judgment (e.g., Simon, 1973; Shrestha et al., 2019).

While these distinctions may be useful heuristics, we do not base the study on a given categorization of humans' and machines' attributes, similarities and differences. Rather, we seek to study their interplay and association, prompted by the notions of augmentation and assemblage. To do this, we build on Lucy Suchman's seminal study at Xerox PARC (Suchman, 1987) which challenged fundamental ideas about humans' and machines' cognition and communication; ideas that had guided the design of the intelligent interface of the advanced copying machines used by the scientists she observed. Her studies revealed that the interface designed based on a knowledge representation in a computationally encoded control structure did not work, since human interaction with each other and with objects in the world hinges on a more complex process of mutually interactive constitution of intelligibility (Suchman, 2007, p. 10, 1993). Therefore, she argued we need to move "from categorical debates to empirical investigations of concrete practices" (Suchman, 2007, p.1) in which humans and machines act together. We seek to build on this and take the notion of *human–machine configurations* as our analytic frame for studying this interplay. A configuration is a specific set of relations between human(s) and machine(s) with a certain division of tasks and responsibilities between them. There can be several possible configurations with different specific work distributions, degrees of automation and types of task sharing over time. Reconfiguration, then, is the process by which new configurations of humans and machines and related objects of work practice emerge (Mazmanian et al., 2014, p. 832; Suchman, 2012). The notion of "configuration" draws our attention to how sociotechnical phenomena have been shaped through being joined together and reminds us that this is an enacted and ongoing accomplishment (Suchman, 2012).

Furthermore, our analysis will focus on a particular configuration otherwise known as human-in-the-loop. This is a term that originated within the fields of modelling and simulation, indicating human supervisory control in computational systems' operations. Typically, this implied human feedback and responsibility for performance management, exception handling and improvement (Sheridan, 1995). The notion has been revived within current advances in AI, especially the interactive forms of machine learning, such as reinforcement and active learning (Rahwan, 2018; Brynjolfsson and Mitchell, 2017). Here humans participate in different roles such as the creation of data sets (e.g., generating interaction data) and labelling of events or data unknown to the algorithm (i.e., edge cases) that are used as training examples for the algorithmic system to learn from. The notion of human-in-the-loop is also used when discussing regulation and control of ethically problematic autonomous systems, e.g., human oversight over AI-enabled autonomous weapons (Citron and Pasquale, 2014). Whereas AI-based automation might take the human out of the loop and replace them with algorithms, augmentation might involve a human-in-the-loop (Norman, 1990; Bailey and Barley, 2019; Markus, 2017). The concept of human-in-the-loop thus helps us examine the tasks assigned to and performed by humans in the human–machine configurations. We find it useful to identify novel digital work that emerges as humans and algorithms are configured to transform data into value-added output. In the next section we introduce our case organization and the strategic context for its activities, before we describe the design and execution of the research.

Research approach and methodology

Case background: Brokering activities and open, global data

About 90 percent of the world's trade is carried by sea, "with ships and ports acting as the arteries of the global economy" (Saul, 2017). Due to the complexity involved in global maritime trade, brokering firms have traditionally earned their right to serve as intermediators of logistics in this industry. Our case is a financial services and brokering firm which has operated in this industry since the mid-1800s. The firm (dubbed ShipCo for anonymity) today employs several hundred employees distributed across >10 offices worldwide. ShipCo's brokers intermediate and negotiate exchange between ship owners and cargo owners and/or charterers through identifying opportunities for transport and arbitrage, i.e., profit from a temporary difference in prices on a commodity. To intermediate such exchange, ShipCo relies upon up-to-date data on market developments and maritime activity. In their everyday work, shipbrokers need to know not only current supply and demand for a given commodity, but also which ships are located where and when, and whether or not these ships are available.

The core system and data source in our study is the Automatic Identification System (AIS), a global, standardized communication system enabling exchange of data on ships' navigational status and voyages. Since 2004, all ships of 300 or more gross tonnage engaged on international voyages are mandated by the International Maritime Organization (IMO), a UN agency, to carry AIS transponders. The AIS transponder (the on-board device that sends and receives signals) uses a very high frequency (VHF) radio transmitter and the Global Positioning System (GPS) to broadcast information on its activity to nearby receiver devices on other ships, buoys, and ground stations. The default transmission rate is every few seconds, and the message includes automatically updated (i.e., dynamic) data on ships' identity, position coordinates, course, and speed. In addition, the message carries manually entered (and more static) data such as cargo, point of departure, destination and estimated time of arrival.

There are several issues contributing to varying data quality in the AIS system. For instance, AIS signals can interfere with one another in crowded waters, or the manually entered information, such as a ship's destination is prone to omission, misspellings, and distortions. The destination port can be spelled in multiple ways (e.g. "Singapore" or "Sng"). Information about the destination can be withheld or misleading in order to hide the intended destination, or it can also be used to flag security measures, e.g., in pirate waters where "Armed guards on board" is a common value in the destination field. Furthermore, the transponder can be turned off by the ship crew to avoid being detected by pirates or to conceal ship identity, location, or destination for economic or political gain. This is a common occurrence, e.g., in the waters around the Korean peninsula to evade the sanctions on trading with North Korea.

While the AIS was designed to prevent ships from colliding (by allowing ships to know nearby ships' location), other usages have emerged. This is due to more actors entering the infrastructure as low-earth orbit satellites became increasingly available during the last decade. Uses of AIS data today range from search and rescue operations, piracy monitoring, smuggling prevention, disease mitigation, environmental protection, economic forecasting, commodity trading and intelligence, with the latter being the focus of this article.

In ShipCo, AIS data is used to make aggregate assumptions about trade flows and to predict supply and demand for commodities such as oil and gas. Here, however, competition has picked up. There are also other, novel actors in the markets who offer AIS data and analysis of maritime activity to industry professionals such as analysts, traders, and policy makers. They include data service providers, usually upstream satellite operators, who sell raw AIS data and/or satellite images; downstream online ship tracking services; as well as analytics and intelligence firms which acquire, refine, and resell the data. Increasingly, data service providers combine AIS data with contextual data and machine learning algorithms to generate value-added information about trade activity. Contextual data sets include, but are not limited to, ship registers, weather reports, financial projections, customs and port authority statistics, governmental databases, and satellite imagery. This description of the moving frontier of data analytics in a global, competitive sector provides the backdrop for our case study, which we now will describe in more detail.

Research approach

We have conducted an in-depth, qualitative case study based on the interpretive research paradigm (Walsham, 1995). The fieldwork was conducted between September 2017 and October 2019 and covered the process of the organization's design, implementation, and use of algorithm-supported data analysis (Bailey and Barley, 2019). The selection of the site and negotiation of access happened through an extant development project between the shipbroking firm and an enterprise software company. The company delivered cloud and business intelligence (BI) solutions to the maritime sector, and the first author was an employee of the company while doing his PhD. Access was negotiated in spring 2017 on behalf of the software company and with guarantee of confidentiality regarding business-critical information. During the fieldwork, the researcher role and intent was clearly communicated, and observations and interviews were based on informed consent from the informants. Our initial access was negotiated through meetings with the firms' top management in early fall 2017. As we were granted access to the firm's headquarters and the research unit, our role of involvement gradually evolved from participating in strategy planning meetings to also taking part in implementation and evaluation meetings. This allowed us to participate in separate and joint meetings with the IT development team, analysts, and brokers from various 'desks'. In our account we have anonymized the informants.

Data collection

The main method was participant observation, including informal interviews and conversation with organizational members. The first author participated in 39 meetings, which, with some exceptions, were on-site. The meetings lasted between thirty minutes and five hours, with an average duration of 104 min. In addition, a two-week-long design workshop was used for data collection in January 2018. The first author was engaged in organizing the workshop, which involved the wider organization in the design of new digital solutions for data acquisition. The Chief Digital Officer (CDO) and brokers were involved throughout the two-week workshop period and engaged particularly in five half-day sessions giving feedback on the evolving designs. The outcome of this process was presented to the executive team. During our observations, we avoided recording so as to ensure openness and trust. Notes were taken during and written out to more comprehensive research notes after the observations.

During the fieldwork period, informal interviews and conversation were regularly conducted to follow up and verify our observations, including one-to-one meetings, lunch meetings, face-to-face conversations in the field as well as over Skype and phone. Documentary sources were reviewed including company PowerPoints and reports, spreadsheets, contract excerpts, whitepapers, newspaper articles, web sites, software, and internal databases. We were also given a company chat and email account on which we regularly conversed with our key informants; this contains > 200 emails.

Data analysis

Our data analysis strategy follows the four-stage approach of Jarzabkowski et al. (2016) to analyzing qualitative process data and everyday activities as they unfold. This approach contrasts to studying formal processes and scripted patterns in that there is no defined sequence to be enacted. Rather, we pay attention to situations where the informants themselves are trying to set forth a direction or order within their activities (p. 244). As such, we assume that it is in the change of practice that new task relations emerge and, eventually, form a pattern of human and machine interplay (Suchman, 2007). This approach aligns with other studies concerned with the relationship between technology and organizational work (e.g., Barley, 1990).

The first of the four stages involved using initial observations from meetings and conversations to confirm that our setting allowed access to phenomena of interest. This was done mainly through participation in the initial meetings with the executive team members at ShipCo and the software provider. For example, from one of the meetings we produced and circulated a summary of our understanding of the strategic conditions and problem context, which served as a means for confirmation and contextualization (Klein and Myers, 1999). Using a sparse form of open coding (Walsham, 2006) of our field notes, we sought out patterns and selected empirically informed themes such as “secrecy” and “obscuring”, all relating to episodes of questioning the validity of external data. These initial codes allowed us to consider our data in context of candidate constructs and to sharpen the focus of subsequent data collection.

For the second stage of analysis, we set boundaries for our analysis by selecting a sub-sample of our data that centered on meetings revolving around data quality and evaluation of algorithmic outcome. This meant that we concentrated our analysis on the research unit including analysts as well as the IT development team. The gradual deepening of access and exposure to evaluation and implementation issues, helped narrow down our focus and orient our analysis towards the interplay between extant and novel data-intensive work. This aligns with, e.g., Leonardi (2011) taking as his unit of analysis the process by which human and technology are intertwined. Our unit of analysis was the emerging human-algorithm configuration with a focus on how existing tasks were (re-) distributed and which new tasks emerged.

The third stage of analysis involved a deeper production of case narratives (Langley, 1999) and vignettes from our empirical material, while still engaging in the field and collecting data. Here quotes, memos, and early analytical notes were documented to permit a better understanding of what was going on, what the key conditions were, who the key entities were, and what constituted an event of interest (Jarzabkowski et al., 2016). Visual maps served as intermediary data bases for the identification of relevant themes and sequences of work (Feldman, 2016). Drawing on these empirical data bases together with key informants allowed us to revisit our narratives in light of new material and gradually add new insights as patterns of problematic situations (such as ships “going dark”) and human-algorithmic configurations started to emerge.

For the fourth phase of analysis, we explored data-theory links (Walsham, 2006) across related research to develop a conceptual basis for our analysis. Emergent theoretical concepts such as augmentation and assemblage provided a “repertoire of lenses” (Venters et al., 2014) which helped us illuminate and validate observations related to regularities such as ‘training of the robot’ (the CDO's terms when referring to the algorithm). The relevance of human-in-the-loop as a key concept surfaced in this phase as we engaged with our data and discovered the emergent patterns and relations between algorithmic processing and human intervention.

The following section describes our empirical study of the process through which an algorithm was introduced into the organizational unit. The account focuses on the impact this had on the organizing of the research and analysis function. Following that, we then present our case analysis focused on emergent human–machine configurations with a particular focus on the notion of human-in-the-loop.

Case findings

We will now present our empirical material using a chronological structure representing shifting configurations of human–machine interplay. First, we describe the nature of the research unit’s extant work practices (Pre-introduction of the algorithm) before presenting an account of the introduction of an algorithm leading to three different configurations (Introduction of the algorithm). Finally, we describe the process by which the emergent work configuration achieved satisfactory accuracy for the algorithm and expanded into prediction for other types of ships and commodities (Post-introduction of the algorithm).

Pre-introduction of the algorithm

Manually producing tradetables to support analysis

Since the mid-1900′s the research unit in ShipCo’s headquarters has monitored world seaborne trade and produced ‘intel’. Here, a researcher analyzes data on ships’ voyages and activities in certain ports to produce a current image of trade flows. The main recipient for this information is the analyst, who construct aggregates to predict fleet trade flows and supply-and-demand for certain commodities. The primary function of an analyst is to provide shipbrokers, clients, and other stakeholders, with up-to-date aggregates, projections on trade developments, and recommendations for arbitrage opportunities. Their mode of work is illustrated in Fig. 1.

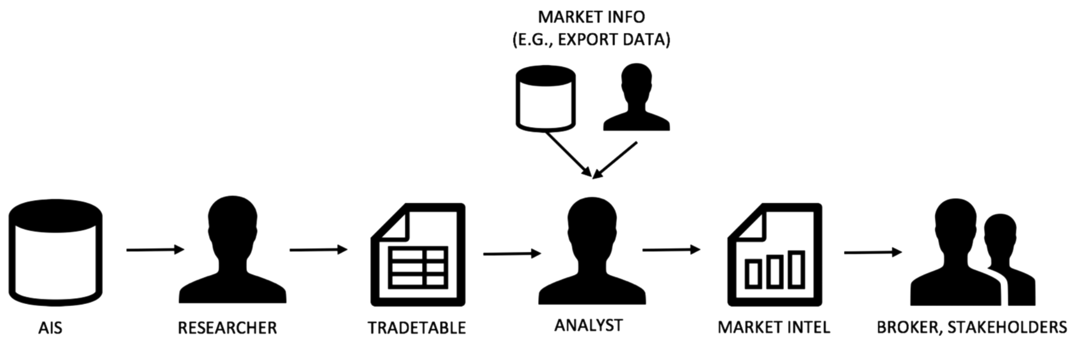

Fig. 1. The roles and process of transforming AIS data into tradetables and market intelligence.

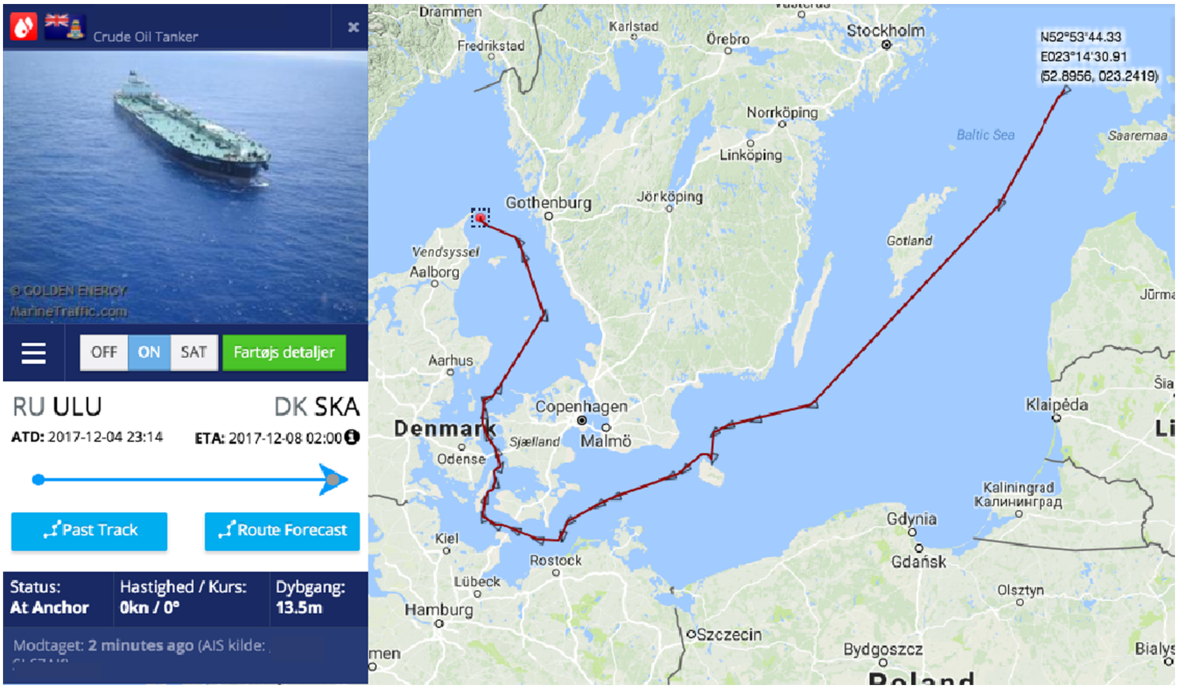

Fig. 2. Tracking of ship using AIS data and the online tracking system MarineTraffic.com.

The current paper focuses on the researcher working with AIS data (the leftmost person in the figure). The researcher's work involves the considerable manual work of gathering and analyzing historical and current data on fleet and market developments. Both researchers and analysts would typically specialize on certain geographical regions and/or commodity markets (called "segments") such as liquefied petroleum gas (LPG) and oil.

The researcher worked to identify maritime activity and provide information that supported the analysts in intelligence gathering used to predict trade flows (see also Fig. 1). The main output from the researcher's activity was *tradetables* – spreadsheet documents which contain information about activities of certain ships, including timestamps of departures and arrival, destinations, and in which ports they loaded or discharged cargo.

To generate the tradetables, the researcher gathered, cross-referenced, and classified AIS data using an interactive ship tracking system (Fig. 2 shows an example of one of the commonly used shiptracking tools). The tradetables were then provided to the analysts where they were combined with contextual information. The generation of tradetables was time consuming and might involve non-trivial decisions, as the following fieldwork vignette illustrates:

On one morning in early 2018, the researcher arrives at her desk at the shipbroking company's headquarters, immediately opens the ship tracking tool and begins digging into the activity of a 53 000-gross tonnage VLCC (Very Large Crude Carrier) between certain Saudi and African port terminals. She pulls a range of temporal observations on this ship from the ship tracking tool and tediously dives into its recent movements, including wherefrom and at what time it departed; when and at which port it arrived; what type of cargo and how much it, if at all, loaded or discharged. Yet, with this voyage, like many others, there are abnormalities, and it is hard to tell what the ship actually did at certain ports at certain points in time. Consequently, the researcher pulls out a second data source, another third-party, online ship tracking tool, for cross-referencing and making sense of the ship's actual activity. Documenting the ship's activity in her custom Excel spreadsheet system, she manages to clarify some but not all of the abnormalities. Subsequently, the researcher reaches out to a colleague at the operations desk for more information and knowledge. Her colleague quickly relates the inconsistences to a sudden, major oil transaction by the Horn of Africa. Given this new insight, the researcher returns to her desk and more confidently classifies this ship's activity along with an estimation of how much cargo was loaded or discharged on a scale of 0–100. Before forwarding the updated tradetable to the tanker analyst, she moves on to classify the next voyage on her list.

While the manual classification of ship activities and preparation of tradetables was a tedious process, the researcher was assisted by her deep and long-term understanding of vessel movement patterns, routes, and terminals across the world's ports. As the CDO explained, referring to the researcher, "*she knows almost all ships and ports inside out*". The researcher was drawing on more than three decades of experience of sorting through, revising, and verifying detailed data on particular ships and their voyages: to rule out abnormalities, properly document their port activity, and produce as accurate as possible tradetables. This process of extracting and interrogating AIS data was critical for the researchers' judgment on whether the ship was loading or discharging in a specific harbor, or whether it did neither, but just bunkered fuel.

Using tradetables to predict ship behavior and trade flows

Analysts at ShipCo's research unit used the tradetables to construct aggregates on trade flows and supply-and-demand for certain commodities such as crude oil. These were then used for predicting where and when opportunities for business might occur. Having received the tradetables electronically from the researcher, the analyst combined these files with additional, contextual information. This included internal sources, such as databases and time series supplied from the research unit, as well as external sources such as interviews, fleet lists, tracking software, and market reports. From these inputs the analyst produced up-to-date reports, projections and visualizations, which were used to provide the brokers, clients, and other stakeholders with actionable advice. Here, one senior analyst describes the importance of AIS data along with information on availability:

> '*We want to predict the number of vessels, their availability, and whether they are fixed or not. The key problem to solve is: how likely is it that they* [i.e.*, the clients] are going to pay X? Answers to this come down to ship availability, where and when the ship will go, where it was fixed, its ballast and speed.'*

Not only accurate, but also timely information was crucial. For example, as soon as available ships and cargo were negotiated and 'fixed' between other actors in the marketplace, ShipCo's window of opportunity would terminate. An unforeseen congestion of ships awaiting to enter the Panama or Suez channels might influence ships' routes and thus ship availability. Prices on oil might drop due to political conditions, causing the supply and demand balance to shift. As such, there was a need for analysts to constantly stay on top of a variety of parameters which might influence the market conditions and rates. The ability to foresee 'where ships actually will be going' is considered to be crucial for estimating where and when opportunities for arbitrage will occur. Perceiving the strategic importance of this, the CDO commented that "*there are others out there who are not direct competitors today who can enter and take the market from us… to maintain our existence, it is 'alpha and omega' that we do this well.*" However, keeping up with diverse data sources and a continuous stream of vast AIS data required a lot of capacity; more than was available with the extant practice of producing tradetables.

Detecting human bottleneck and setting strategic direction

Guided by the CDO, the IT team had initiated some work on improving data analysis capacity in the autumn of 2017, however, towards the end of 2017 the importance of the challenges became an urgent concern across the upper echelon in the firm. During a client meeting, one of the senior analysts was surprised to discover that a client kept more up-to-date data than he did. Looking back on this event, the senior analyst expressed his concerns:

> *'Ever since I have been here, the trend is: brokers have scaled down on research, while the customers have scaled up on research. Tracking data [on vessel activity from the AIS] is what made me go from zero to hero. Today, however, my customers have the same or even better data than what I do. This is a huge challenge.'*

The analyst's competitive edge was lost. Returning from the client meeting, the analyst took his concerns straight to the senior management. They traced the issue back to the limited capacity of the researcher in keeping up with the incoming amount of data on vessel activity, and thus producing tradetables on time. Subsequently, the management decided that improving the production of up-to-date tradetables for this analyst's segment (oil freight) should receive the highest priority. This event boosted the process of adding automation to the data acquisition and analysis process and guided the focus towards this specific segment.

Realizing that they needed to obtain control over data quality, rather than rely solely on intermediary providers' representation of the data, ShipCo approached a satellite operator and data supplier that collected and offered a stream of AIS data as a service. To achieve a more fine-grained and accurate picture of activity in the oil freight segment the company subscribed to a higher frequency AIS data stream, going from daily to hourly updates. This would also allow ShipCo to bypass downstream data providers as well as enable more timely production of tradetables. Also, the predictive power of the tradetables was expected to increase, since the algorithmic tables would be based upon more frequent data, spanning larger chunks of time than the researcher currently used, and pull in additional data sets. The new data acquisition strategy meant an increase in AIS data input which necessitated installation of additional data warehouse capacity. It also generated more work on preparing and processing this data; the increased uptake of data required more capacity for classification. However, the algorithm was expected to deliver this increased capacity for analysis of the incoming data.

Introduction of the algorithm

Automating tradetable production

The algorithm-supported analysis system was designed to automate both data acquisition and the processing of data for subsequent analysis. Acquisition of data was automated by the system pulling streams of data on ship activity from the satellite-AIS data provider, along with additional data such as vessel descriptions and geospatial data, into a Hadoop-based data warehouse repository. Here the data were extracted and consolidated, then classified using rule-based NLP (Natural Language Processing) classification, and finally presented in BI tools that allowed human interpretation of the output. The intention was to have the algorithm clean, prepare and classify 'raw' AIS data, similarly to what the researcher manually did to generate the tradetables. As the team started to look into the AIS data, they realized that the quality of data was insufficient to fully automate the data processing directly. There were anomalies in the data sets, such as variations in how names of port destinations had been spelled as this had been entered as free text by the various ship crews. For example, on several occasion members of the team referred to "*a thousand spelling variations of Singapore*". According to the data scientist, the system initially was based on "*hard-coded heuristics or simple pieces of logic that were* *configured to capture certain events and errors in the data sets*". When an issue with the data was detected, such as the port names, the data scientist would set out to collect all known portmanteaus, abbreviations and acronyms related to the port in question. He would then instruct the algorithm to manipulate these various strings of port spellings and geographical data, so that it could automatically normalize cases of data anomalies.

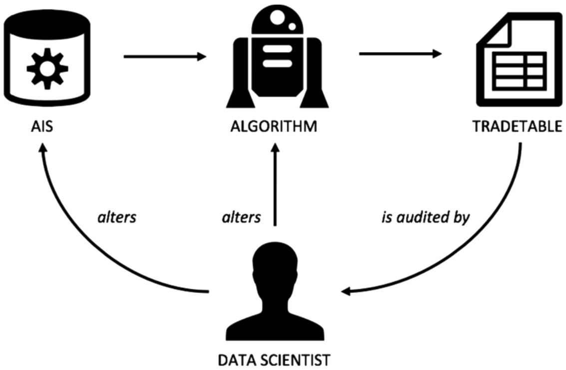

Fig. 3. Configuration upon the introduction of the algorithm: auditing and altering conducted by the data scientist.

Let us consider this initial version of a human–machine configuration (see Fig. 3): The novel work tasks in the initial phase fell to a large degree on the data scientist and consisted of evaluating or *auditing* the tradetables generated by the algorithm and manually altering the algorithm to improve its performance. In addition, the experiences made during the first phase also effected change in the data acquisition. This included work of ensuring availability, negotiating access with the AIS data supplier, connecting to new streams of data, and adjusting the data structures. This work of *altering* was not one-off but would be repeated regularly. For instance, when novel variations in port names occurred, they would be added to the algorithm's rule set. This initial configuration saw a human-in-the-loop – the data scientist who worked to adjust both the algorithm and the data acquisition process.

Enrolling domain expertise to audit algorithmic output

While the team's emphasis initially was on data sourcing and processing, their focus shifted toward evaluation of the algorithmic output as the implementation of the algorithmic system progressed. During this phase, the CDO called for meetings with the data scientist and the researcher to evaluate output from the algorithm. The data scientist would present algorithmic results in the form of complete tradetables and map plots of ship voyages. During the discussions the team would zoom into specific voyages to resolve data disparities. The data scientist would examine the underlying code, switching between an Emacs and a SQL editor as well as an interactive maritime map. The data scientist explained how he targeted abnormal or unconfident output from the algorithm:

> *'I passed [edge cases] on to the researcher and one broker for assurance and feedback. For example, in one case, one vessel reported the same port, Sungai Linggi in Malaysia, for both departure and arrival. What we are interested in is understanding what really happened here so that we can have the algorithm to automatically deal with such cases in the future.'*

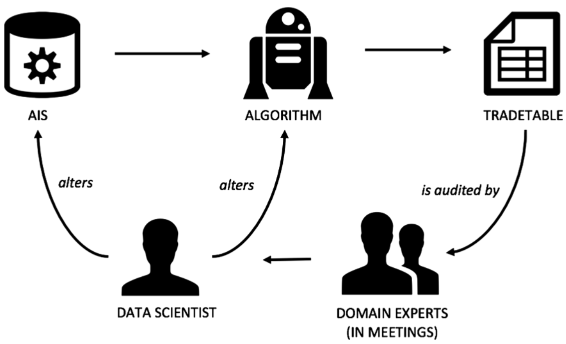

Fig. 4. Reconfiguration involving domain experts in auditing and data scientist in altering of the algorithm.

Collaborative evaluative events, like the one described above, frequently took place in meetings during the initial implementation of the algorithm. As expressed by the CDO: "*Every week we have to tune on the edge cases"*. In these meetings the researcher who carried out the manual data processing and preparation, was a crucial participant. On an as-needed basis, various participants with detailed knowledge on specific segments, including regions, ports, and ships would also join the meetings and contribute their interpretation of ambiguous output. Based on the participants' interpretation and justification, the data scientist and the CDO would adjust the algorithmic set-up as well as the data structures serving the algorithm. These face-to-face meetings where characterized by cycles of interpretation and refinement of the algorithmic system through an evaluative practice.

The second version of the human–machine configuration (see Fig. 4) was different from the first. Now more humans were added into the loop - some of them permanent members of the evaluative practice, and others were occasional participants. They did not all have the same tasks, as some (the domain experts and the researcher) provided an evaluation of the algorithms output and others (the data scientist and the CDO) also adjusted the algorithm and data acquisition accordingly.

Hiring data analyst and repurposing the role of the researcher

The work of evaluation evolved further when the team started to compare the algorithmic outputs (i.e., tradetables) with the tradetables that were produced manually by a researcher. Each tradetable contained timestamps for departure, regions and ports for loading, timestamps for arrival, regions and ports for discharge, draught changes, and number of discharges, where the latter item (number of discharges) had been estimated based on the other information in the table. The two tradetables were displayed side by side in a spreadsheet tool so that they could easily be compared. When going through the two outcomes, the data scientist would continue to bring in the researcher and domain expertise for interpretation of discrepancies.

Throughout this work, the researcher's manually generated tradetable continued to serve as a 'ground truth' or 'gold standard' reference against which the algorithmic tradetable was evaluated. The accuracy of the algorithm was estimated by comparing and calibrating the number of estimated discharges. For instance, with the manual tradetable identifying 70 and the algorithmic tradetable 77 discharges, the discrepancy between the two were estimated to be ~10%. Based on the identified gaps, the data scientist and the CDO adjusted the algorithm and altered the data acquisition strategy. However, both the data scientist and the researcher were now facing increasing workloads. As the CDO recognized the need to free up their time, the management decided by early 2019 to hire a new resource within the team. The role of this new resource—called "data analyst"— was largely to evaluate the algorithmic output and provide the data scientist with information about its performance (see Fig. 5). As the work had progressed, the researcher's concerns for her future job had been raised, and the CDO tried to assure her that she would not be replaced: *"[The] data scientist's analyses and predictions will be fed back into your system…. [It] will help you save time on mundane tasks, they will allow you to work more on recent data, and to focus on quality assurance … you will train the robot.*"

In this third version of the configuration (see Fig. 5), tradetables were generated by both the researcher and the algorithm, and then compared by the data analyst. Only on some occasions where the two reference outputs deviated, the researcher, data scientist or other domain experts were called in. The CDO and the data scientist still adjusted the algorithm and allowed it to be released to the intended users, the company's analysts. The extant, manual work of tradetable generation continued, as it generated the reference against which the algorithm would continually be audited. The data analyst was feeding back information both on individual edge cases, but also on overall performance. He compared the manual and algorithmic output and published the result from this work to a Business Intelligence (BI)-based dashboard that was consulted by the CDO and data scientist who in turn made decisions for altering the algorithmic system.

Post-introduction of the algorithm

Achieving satisfactory accuracy and expanding into new segments

By October 2019, the algorithm had successfully achieved an accuracy level equal to or better than that of the researcher. The oil tanker analyst had accepted the automatically generated tradetables for use in his analysis and prediction of trade flows. This was a milestone and allowed the team to start adapting the algorithmic system and configuration to analysis of adjacent segments. The

Fig. 5. Reconfiguration repurposing the role of the researcher and enrolling data analyst to the auditing work.

following vignette illustrates how the system had to be adjusted to allow the team to expand into a new segment.

> *"We have just started on a new segment", says our informant. He explains that they were happy with the performance for oil tankers and now started with LPG (Liquefied Petroleum Gas) tankers. Then they had to do a new round of dealing with port names, since sometimes the LPG ships go to different ports. "See here, for instance, there are a lot of ships a bit north of Chennai, and we haven't registered any port terminal here". He points to what he later explains is a heatmap of ship locations which he generated himself from the BI tool. He zooms in on the map and true enough, there is a port terminal there. He explains that because of this expansion in the algorithm's scope, the data scientist has now provided him with a list of port names that he has to go through. These are port names in use in the AIS system, but where the spelling may vary (as with translating Arabic or Chinese into Latin letters), where the names is misspelled, or where an area within the larger port area is not specified. The list is sorted so that the most frequent occurrences are on top. Our informant then goes through the list and identifies the port's coordinates and enters it into the 'mapping list' used by the algorithm. "Here it says "Ocoa Bay", and if I Google it, I find that it is in the Dominican Republic. I then go to the map and zoom in, from the heatmapping I see where the ships are, and then I find their coordinates and can register these. Over here I have made myself a Python script that helps me to do the registration more efficiently. Yesterday I covered around 1800 entries in the port file.". … "When I started here, the mapping list had around 8000 entries, now this has around 29,000 entries (*i.e. *where variations in port names are mapped onto standardized forms)."*

While the team was achieving their goals, expanding into a new segment required adaption of both the algorithm and the data acquisition. As the CDO explained: "*the different ship types have different behavior patterns, the algorithms for each segment must be switched on*." The work contained similar tasks that had been covered earlier, such as normalization of port names and resolution of ambiguous ship movements, however, in a new context with new entities.

Embracing a culture of learning and experimentation

One of the intentions of the data analyst's work was to create an easy-to-use interface to the evaluative task for other organizational members. The idea was to facilitate a way for them to continue to "train" the algorithmic system on a more permanent basis and through a less costly mechanism than the meetings. According to the CDO: *"Human-in-the-loop is important; it is next in our pipeline…. We can also put the senior analyst and a broker in the loop…. We need to develop an interface that is user-friendly so that non-data scientists can understand and efficiently use it to train the algorithm."* This is in line with the CDO's statement towards the future vision: *"'Today we are using data-supported human analysis. I want the opposite. I want human-supported data analysis…. This means that we can apply our team's knowledge in a new way to IT, rather than viewing their knowledge as obsolete."* His vision was exemplified by the repurposing of the role of the researcher, from manual classification work toward 'training the robot', which was achieved mainly through providing the auditing process with a ground truth reference. Importantly, this role and reference allowed the feedback mechanism through which the system could adapt.

Looking ahead, the executive team aims to do more than simply implement the algorithmically supported data analysis of trade flows in everyday organizational work. They envision a multifaceted digital strategy involving algorithmic prediction of global trade rates and identification of the market drivers influencing these rates. Core to achieving this is building a culture of learning and experimentation. For instance, previously, while the team grappled with making sense of questionable AIS data, the CDO had initiated experimentation with various machine learning algorithms such as neural networks. The goal was to identify key drivers of shifts in supply and demand in maritime trade. Here, internal and external data such as time series on regional weather conditions, corn production, and petroleum consumption were combined and fed into the various algorithms. The algorithmic outcomes were then evaluated and compared by the team, including the CDO and a manager from the financial department, and the parameters were in turn adjusted throughout multiple iterations. When evaluated against historical data, however, the team discovered that the simpler algorithms, linear regression in particular, returned the most accurate results. Following the experiment, the team reoriented their efforts toward acquisition and transformation of data while keeping humans in the loop for design and implementation of the algorithm. While still a strategic vision, these directions indicate an approach to strategizing that emphasize experimentation and learning. Some obvious questions are how much can be automated with a reasonable cost-to-benefit ratio, and what the nature of the augmentation work is going to be.

Analysis

Our case study covers the initial phases of ShipCo's trajectory of co-evolving human and algorithmic intelligence. The strategic starting point for the initiative was that the organization sought to maintain its competitive advantage through turning vast data into value-added analysis and "intel". The expectation was that the algorithm would resolve an analysis capacity bottleneck, with the goal of achieving greater accuracy and speed in the delivery of analysis. When the algorithm was introduced, the team discovered that it needed adaptations and optimizations. While this is by no means surprising, we think that investigating this early phase allowed us to map how the division of work tasks between humans and algorithm, and among different organizational members, evolved. In other words, we could study in detail how tasks, work and organizing were impacted by the introduction of the algorithm. In the following we analyze the novel tasks, work processes and the configurations that emerged in relation to the algorithm.

Ongoing configuration and new work of augmenting the algorithm

The introduction of the algorithm has strengthened core organizational activities, more specifically the organization's capacity to perform analysis of incoming data. More data can be analyzed (i.e., more high-frequency data for larger time periods) and a higher number of analyses can be conducted (i.e., tradetables for a higher number of ships can be produced) in a more timely manner (without the time lag). The new algorithmic analysis could also generate other information products that had previously been manually generated, such as dynamically updated graphs of trading activity. As a result, the algorithm strengthened the organization's data analysis activities.

However, the introduction of the algorithm has also required the development of novel capabilities. Employees with data science skill sets were hired, and organizational members gained new work tasks. The new work tasks fell into two broad categories. Firstly, the performance of the algorithm needed to be monitored and evaluated – or *audited*. Secondly, the algorithm, the data input or the presentation mode had to be continuously improved – or *altered.* As a collective, the team's work of auditing and altering the algorithm comprise what we call "augmenting work". The nature of this augmenting work, to whom it was assigned, and the delegation of auditing and altering tasks, varied throughout this early phase of ShipCo's trajectory.

The first configuration saw the enrolment of a data scientist who, alongside the CDO, built an initial version of the algorithm and ensured its data input stream. Once the system was in place, the data scientist both audited the algorithmic output and altered the algorithm: tasks which were novel for him as well for the organization.

In the second configuration, the algorithm had improved enough to generate actual output. Now domain expertise was required to assess the output, and the data scientist enrolled the researcher and sometimes the analyst for the auditing work. The domain experts' auditing generated crucial resources for the data scientist. First, they pointed out errors in the algorithm's performance – this indicated where there was a need to alter either the algorithm or the data input. Secondly, they attempted to explain the errors and to provide a correct interpretation of data irregularities. This interpretation indicated what kind of altering was required, e.g. whether to add additional rules or to adjust the existing rules so that they performed more in line with the domain experts' judgments. The initial form for this collective augmentation work was the ad hoc meetings, which over time evolved into more structured meetings, where collective auditing of algorithmic output was followed by updating of the algorithm by the data scientist. However, this configuration of multiple humans-in-the-loop was costly, as this meeting activity required allocation of significant resources.

In the third configuration we see that the role of domain experts as "auditors" became less central again; now a new dedicated resource who was neither a data scientist nor a domain expert was delegated the task of auditing. This was only doable because a set of known errors (including explanations) and strategies for dealing with them had been established. The newly hired data analyst was tasked with auditing the algorithm's output, while the data scientist was responsible for altering the algorithm. This particular configuration for augmenting the algorithm held up also after the first version of the algorithm was set in production for the particular market segment (oil tankers). As the team moved onto a new segment of LPG tankers (Liquid Petroleum Gas), there were novel issues to deal with, e.g., new destination ports for LPG tankers what had not been used by oil tankers. The process of mapping the variation and normalizing the spelling had to be repeated for these ports. However, the process of auditing and altering the algorithm did not require the organization to go back to the meeting-based mode of working; it was still done with the same configuration of the data analyst and the data scientist as central actors.

In the company rhetoric, the expenditure of resources (including the new employee) was seen as part of building long-term algorithmic capabilities in the organization. This included building company culture, mandate, and talent. The strategic agenda was also visible in the future-oriented justification of other investments, such as preparation of high-quality data sets for future, more advanced machine learning. This would necessitate building other algorithms for predictive modelling of market developments, which would require utilization of other data sets, both data such as ship registers, market data such as export/import statistics, weather and climate data (to predict crop productivity, and freight conditions). These undertakings would come with new needs for augmentation, most likely with more demanding processes for auditing the outputs and altering the algorithms employed. The initial steps of algorithmic classification and production of tradetables served to build capacity and establish data resources, data practices, and work culture that would be a basis for further employment of algorithms.

Human-in-the-loop: An emergent work pattern for algorithmic decision-making

We have argued that the organization's analysis capacity increased, and that the researcher as well as the domain experts gained new work tasks. In addition, the organization allocated resources and hired new staff to perform the novel work with the algorithm. What about the manual work process of the researcher, will she be replaced by the algorithm? While we do not know the ultimate answer to this question, we wish to emphasize that the work of auditing that underlies the augmentation of the algorithm and the overall learning, depend crucially on having a ground truth.

At the point in time of the study, the researcher's manual work process represented this ground truth and was still necessary. Our argument is not that there is an inherent complexity of the task that makes it "un-automatable". In one sense, the problem domain was not very complex as there was a limited number of ports, ships, routes and actions. The algorithm used was also relatively simple as it was based on a rule engine that was gradually extended and improved, based on explication and encoding of knowledge from the researcher and analysts. Rather our argument centers on the observation that the algorithm required human auditing when being built, and that this auditing work required a reference measure against which to compare the algorithmic output. Where can this reference measure be found? There was no easy access to any objective "ground truth" from the outer world in this case. On the contrary, the researcher's central challenge was to interpret messy, erroneous, missing, or obfuscated data from the external sources. The problem was not resolvable, as the data sources were not under ShipCo's control. The AIS provider did not have a solution for the data problem, and the introduction of the algorithm did not change this situation. This is a significant challenge for a learning process that builds on feedback: spotting and acting on the gap between the actual and reference level output is core to how algorithms learns. Without any ground truth, there is a missing feedback component in the human–machine configuration.

As we saw in the vignette describing the researcher's manual work practice, the researcher chose to consult a colleague to ask for assistance in resolving the anomalies she saw in the incoming data. Later, when the algorithm was being built, a similar mechanism for collective deliberation emerged in the joint evaluation meetings. However, as both domain experts and data scientists were scarce resources, it was crucial to organize this auditing work efficiently. In our case, a data analyst who was neither a domain expert nor data scientist was tasked with the auditing work. This configuration emerged after a certain plateau of learning was reached, where a set of common errors had been identified and dealt with. We argue that the continued production of manual tradefiles was crucial for this configuration to work. The availability of a (human-generated) approximation of a ground truth made algorithmic auditing possible to delegate to a non-expert. Having this reference available, the data analyst was able to identify and investigate errors because he could compare the manual and algorithmic output. Without the provision of such a reference, the auditing task could not have been delegated to the new resource. Later the work of algorithmic auditing might become even more automated, but again not without a reference measure to compare against. We thus argue that a core reason for a human-in-the-loop configuration is to set the standards for auditing.

Towards the assemblage: Reciprocal human–machine augmentation

The two types of augmentation work—auditing and altering—are dependent on one another: the auditing requires input from the algorithm being altered; and the altering requires input from the auditing process. As the human reference (i.e., the ground truth) and algorithmic reference output become integrated into this cycle, the auditing work of identifying and assessing the gap between the two references can increasingly be delegated to non-experts. In this case, both the tasks of generating the ground truth and assessing the gap between the algorithmic and human reference were new. Generating a ground truth required redistribution and repurposing of extant work as the researcher's outcome now also served as a performance metric. The latter task required a new dedicated resource – the data analyst. Fig. 6 illustrates the overall task dependencies and feedback dynamics involved in the augmenting work.

From the left in the illustration, *Data Acquisition* denotes the work of data sourcing, transformation, and storage, a process which in this case was increasingly automated. The *Algorithmic Classification* relies on this data acquisition as it automatically takes input data, processes it, and produces a prediction outcome. Traditionally, the organization have relied on manual *Human Classification* (illustrated by the dotted arrows and outer feedback loop), which was identified as a bottleneck in the phase before the introduction of the algorithm. However, the extant human work of classification persisted and became central to the *Human Auditing* process as it served as a ground truth against which the algorithmic outcome was compared. The outcome of this auditing work was in turn used in the *Human Altering* of the algorithm as well as the data acquisition architecture. The bold arrow feeding back from auditing to the algorithm denotes altering through training: as the algorithm evolved toward a fully-fledged learning algorithm (i.e., one that learns from humans' learning examples), it would learn mainly through the auditing-algorithm loop, rather than the full auditing-altering-algorithm loop. From the right in the illustration, *Application* indicates acts of applying or putting in operation the prediction

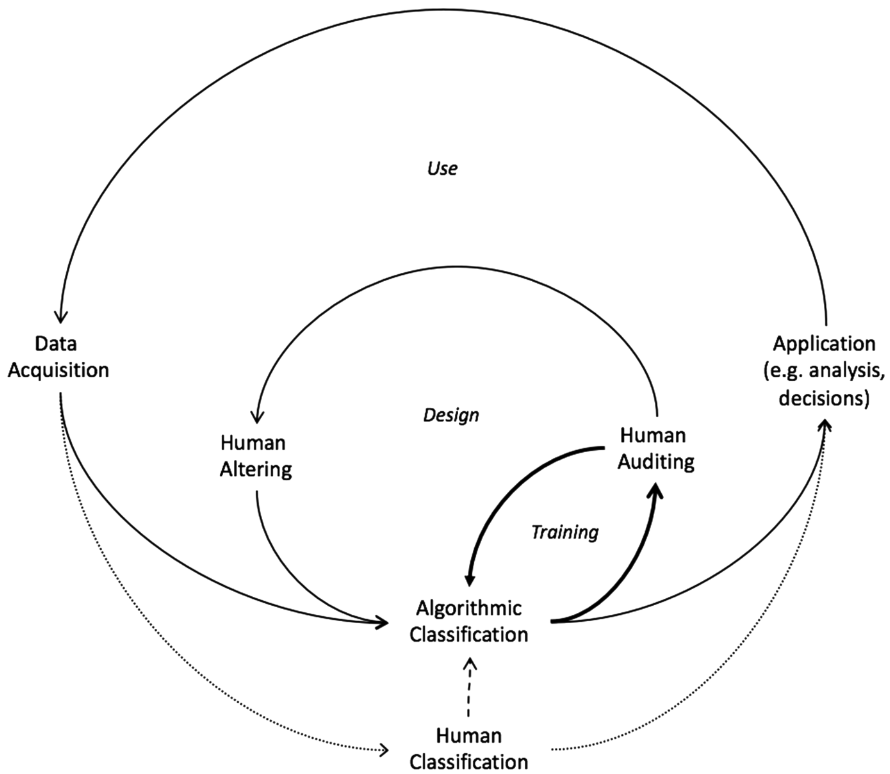

Fig. 6. Emergent human-in-the-loop configuration including augmentation work, and cycles of design and use.

outcomes resulting from the algorithmic (or human) classification work. Actions taken on the analysis at time *t0* would potentially influence the field-level dynamics and, thus, data that is used by the algorithm at *t1*. For example, in a scenario where ShipCo would use their analyses to advise a client to transfer hundreds of thousands of tons of oil, the client's action on this 'intel' would feed-back on and influence the data that ShipCo again acquires.

Following Fig. 6, our observations indicate that the (outer) feedback loop of manual human work shifted to the (mid) loop of auditing and altering the algorithm, which in turn, following the vision of the company, was shifting toward the (inner) loop of training the algorithm by examples (i.e., supervised machine learning). While shedding light on the work configurations which AI springs from, this dynamic indicates that human classification work becomes automated and that the new augmentation work of auditing might prevail whereas the work of altering might exhaust. As the organization expands its scope into new segments and contextual data sources, however, work of altering might be required to further adjust and fit the algorithm to the changing environment. In either case, a ground truth is likely to be needed in the auditing work and to maintain a reflexive operating mode.

Discussion

Our account and articulation of augmentation work contributes to recent IS and organizational literature concerned with understanding whether and under what conditions intelligent technologies will automate or augment human work practices and tasks (Bailey and Barley, 2019; Faraj et al., 2018; von Krogh, 2018; Markus, 2017). In particular, the human-in-the-loop configuration identified and illustrated in this case is one answer to Markus' call that "we should focus on the organizational and environmental conditions that promote or inhibit the effective performance of humans in the control loop" (Markus, 2017, p. 237). Addressing our research question and following Baptista et al.'s (2017a) call for strategic perspectives on digital work and organizational transformation, we organize our discussion by three themes: new patterns of work, new forms of organizing, and their strategic significance (see also Günther et al., 2017).

Augmenting the algorithm: Auditing and altering

Overall, our study confirms earlier research on the non-trivial nature of organizational value creation from data-driven and algorithmic technologies (see e.g., Watson, 2017; Scheepers et al., 2018). Work and resources are required for developing and training the algorithm, as well as for acquiring and preparing the data sets both for the training phase and for the algorithms' regular working. Tarafdar et al. (2019), for example, argues that an organization needs five crucial capabilities: data science competence, business domain proficiency, enterprise architecture expertise, operational IT backbone expertise, and digital inquisitiveness.

Our empirical account sheds additional light on exactly how certain work tasks and roles were changed upon the introduction of an algorithm. In particular, our study contributes with an empirically based understanding of the nature of "augmentation work", comprised of auditing and altering the algorithm. Table 1 further describes the two roles of augmentation work, along with their requirements and task association. Following the task flow, we see that the work of auditing and altering are mutually dependent on each another, eventually forming a feedback loop (see, Fig. 6). As such, both roles are crucial parts of the human-in-the-loop configuration of improving the accuracy of the algorithm and transforming data into value.

In our analysis, we also asked whether the researcher ultimately would be replaced by algorithms. Traditionally, human classification supported or augmented the work of the analyst. We saw that this human work of classification was being automated through the introduction of the algorithm. However, human work was not discontinued, but repurposed to provide a ground truth and still operated in parallel to the algorithmic classification. This prompts another question: whether the researcher and role of human classification will be kept in the loop as the algorithmic accuracy is improved. It has for a long time been known that continual feedback and interaction is crucial for automation to work (Norman, 1990). Our study underscores this and highlights the roles of augmentation work in providing this feedback. In particular, we saw that human classification played a critical role to provide a ground truth – a component which was necessary not only to enable automation but also to obtain and maintain the feedback loop and overall learning of the system. This suggests that the processes of automation and augmentation influence one another in a reciprocal manner (Raisch and Krakowski, 2020), and that their organizing emerges as a core capability.

Table 1
Synthesis of augmentation work (roles, requirements, and task association).

	ROLES	REQUIREMENTS AND TASK ASSOCIATION
AUGMENTATION WORK	AUDITING THE ALGORITHM	Reference input from algorithmic classification or prediction outcome (contingent on data acquisition) Reference input from human classifier or trusted source, a performance metric (e.g., ground truth or gold standard) Comparison, identification, and representation of gap based on discrepancy between the reference inputs (algorithmic and human), e.g., by data analyst
	ALTERING THE ALGORITHM	Input is information related to the gap identified in the auditing work Decision and action to be taken on gap: data science skills to asses algorithmic performance, alter and adjust the algorithm Decision and action to be taken on gap: decision-maker and data science skills to negotiate and change the data acquisition strategy and architecture

Human-in-the-loop configuration for augmenting the algorithm

Often, research on autonomous systems and learning algorithms propose a human-in-the-loop configuration either to avoid harmful consequences or to install accountability for the outcomes: "learning algorithms require humans to ensure accountability" (Faraj et al, 2018, p. 66). Our case is not about a learning algorithm or AI system per se, but the work and process from which such systems spring (Bailey and Barley, 2019). Here the rationale for a human-in-the-loop is more connected to the original training purpose. A classification algorithm is trained through providing it with learning examples and evaluating the proportion of correct classifications. In this case, the ground truth against which the algorithm's classifications were evaluated, was the researcher's manual evaluations of the same data. This human-in-the-loop configuration emerged as necessary during the design and instruction of the algorithm. The question of future configurations is still open, however, we hypothesize that the need for human auditing may re-surface as the algorithm needs to evolve, either because of new tasks or because of changes in the external environment. Indeed, extant literature seem to point to this being a permanent need. Scheepers et al. (2018) point out the need to treat algorithmic technologies as "lifelong learners" (p. 102). Similarly, Agrawal et al. (2019) assert that what limits algorithms is the need for a "willing human trainer" (p. 5). A reference and a feedback component, then, is a crucial resource necessary to ensure continued learning in the overall system. Seidel et al. (2019) argues that an organization that employs algorithmic technologies will need to continuously frame, evaluate and adjust the role of the algorithmic tool in the workflow. Humans will have a role either as "coaches who guide" or "laboratory scientists who experiment" (Seidel et al., 2019, p. 57). While trained data scientists may be in short supply, "citizen data scientists" (Watson, 2017), subject matter experts (Scheepers et al., 2018) or human labelers (Jaton, 2017) could play crucial roles.

That the human-in-the-loop configuration continued to be relevant in our case had to do with the un-resolvable issues with data quality and a long tail of edge cases in this environment. Issues with data veracity was a primary reason for a less than complete automation of the analysis process, and we believe this may point to a situation that can be found in other cases and areas, since "AI places great demand on the quality of input data (correct labelling, complete data, and detectable noise) [and] the requirement for flawless data input to AI still strongly influences when and where it can be put to use in organizations (e.g.., where tasks are repetitive and quality data are generated)" (von Krogh, 2018, p. 405). In the absence of quality data input, the evolution of the algorithm required support from humans, including a ground truth. The human-in-the-loop pattern thus appears to have an augmentative rather than a controlling (cf. Markus, 2017) purpose, where the former configurations serve as a basis for subsequent configurations (see, Fig. 6).

Considering the design and use of intelligent technologies (Bailey and Barley, 2019) as a constitutive and generative process (Suchman, 2012), the feedback loop comprised by acquisition-algorithm-application appears to capture the *use* of the algorithmic classification, while the loop of algorithm-auditing-altering captures the *design* of the algorithm. In the center of Fig. 6, the loop of algorithm-auditing figures design by human *training* as congruent with interactive machine learning (Rahwan, 2018). Thus, our case suggests that the use of algorithmic outcomes feeds back into the design cycle of the algorithm, where humans continue to play a crucial role. The human-in-the-loop configuration emerges as a new form of organizing that helps bring together at once humans and algorithms. It allows the organization to incorporate both augmentation and automation of tasks (Raisch and Krakowski, 2020), and design and use of the algorithm (Bailey and Barley, 2019), in a reciprocal manner. This may also be a relevant pattern in other cases that are contingent on commoditization of data and competitive pressures.

Strategic implications of algorithmic analysis of external data

We adopt the notion of reflexivity to frame our articulation of human-in-the-loop configuration as a strategic capability. Taking configurations as constitutive, generative, and at once action and effect (Suchman, 2012, p. 49), we follow an interactional perspective on reflexivity, which "involves a circular, recursive process or pattern involving feedback loops" (Lynch, 2000, p. 27). Congruent with our conceptualization of augmentation work, feedback is defined as "information about the gap between the actual level and the reference level of a system parameter which is used to alter the gap in some way" (Ramaprasad, 1983, p. 4; Brohman et al., 2019). This approach allows us to pay attention not only to the intertwining of human and machine work, augmentation and automation, but also their underlying feedback mechanisms as a possible source of competitive advantage (Davenport and Kirby, 2016; Iansiti and Lakhani, 2020).

Externally sourced data has become increasingly significant for both operations and strategy and has required development of novel strategic capabilities. For instance, Baptista et al. (2017b) showed the need for the organization to develop "reflexiveness" to adequately deal with social media data in strategizing. In the current case we charted the development of a similar organizational capability of algorithmic reflexiveness: the human-in-the-loop configuration allowed the incorporation of feedback through auditing and altering the algorithm. In our case the data were not necessarily trustworthy, and strategies were devised to resolve the challenges this creates. Here the problem was not just faulty transmitters, noise and errors, but also strategic gaming (cf. Espeland and Sauder, 2007; Marjanovic and Cecez-Kecmanovic, 2017). For instance, mandated data sharing can undermine competitiveness as competitors get hold of strategically sensitive information (Hautz et al., 2017, p. 302), and drive, e.g., ship crews to obfuscate or manipulate data about their true destination. Thus, there was a continued need for human judgment on the output of the algorithmic data analysis. Our study thus indicates that a human-in-the-loop configuration would be a strategic capability in situations where data are sourced from environments characterized by increased transparency and availability of strategic and competitive information. The strategic leveraging of such commodified data requires the work of auditing and altering both for data acquisition and algorithm refinement. Through the evolving configurations of intertwined human and algorithmic work the organization developed a capability to dynamically organize these core tasks.

Our approach to 'reflexive capability' as a core characteristic of the human-in-the-loop configuration can add to the strategy literature on dynamic capabilities, defined as a firm's ability to integrate, build, and reconfigure internal and external resources and competences to address changing environments (Teece et al., 1997, p. 516). While the terms 'reconfigure' and 'configurations' are central to extant definitions of dynamic capabilities (Teece et al., 1997; Eisenhardt and Martin, 2000; Winter, 2003), the terms are not unpacked in detail, neither generally nor specifically relating to human–machine interplay. In this case, resources and competencies comprised an assemblage of humans and algorithms (and data) that was (re)configured in new ways to address changes in the competitive environment (as illustrated by the risk of disintermediation, the need to catch up with clients' analysis and the varying quality and trustworthiness of data). We learned that the human–machine configuration was designed to be reflexive. As Eisenhardt and Martin pointed out, "competitive advantage lies in the resource configurations that managers build using dynamic capabilities, not in the capabilities themselves" (Eisenhardt and Martin, 2000, p. 1107). In this view, human-in-the-loop configuration and the augmentation work that it entails might lead organizations toward dynamic organizing and competitive advantage upon the introduction of intelligent technologies.

Conclusion

Research on algorithmic and intelligent technologies has generated important insights about their potential to replace human work and their relative strengths compared with human capabilities. However, it has largely not investigated the emergent configurations by which humans and algorithmic interplay emerge. Despite the recognition that automating and replacing human work may be more challenging than first expected, detailed research on its processes is still limited. Our two-year long study of an organization that introduced algorithmic support for analysis, is therefore a timely contribution to the building of improved understanding. Upon the introduction of the algorithm, we observed how humans and the algorithm came to augment one another. A set of human work tasks emerged around the automated algorithm. Some tasks, namely auditing and altering the algorithm, were new and constituted the core of augmentation work. Another extant role, the manual work of classifying data, was not replaced, but repurposed to provide the auditing process with a ground truth against which the algorithmic outcome was compared. This emergent human-in-the-loop configuration enabled a feedback mechanism which allowed reflexiveness and scale in the sense of increased accuracy and timeliness. We suggest that the processes of the design and use of the algorithm are tightly coupled as the use feeds into the design mode. Therefore, the organizational choices regarding the division of labour between the algorithm and the humans are crucial. The human-in-the-loop configuration emerges as a strategic capability. It ensures that the performance of the algorithm meets the organization's requirements, because it provides a feedback loop that supports learning. Therefore, it enables ongoing adaptations, extension and improvement of the algorithm's performance. This configuration of human and algorithmic work thus hinges on the organizational reflexivity which is required to successfully generate value out of algorithm-supported data analysis.

Acknowledgements

This work was financially supported by The Research Council of Norway and Iserv Computing, a Norwegian software company. This paper presents independent research and the views expressed are those of the authors and not necessarily those of the Research Council or the software company.

References

Agrawal, A., Gans, J.S., Goldfarb, A., 2019. Exploring the impact of artificial intelligence: Prediction versus judgment. Inf. Econ. Policy 47 (1), 1–6.
Autor, D.H., 2015. Why are there still so many jobs? The history and future of workplace automation. J. Econ. Perspect. 29 (3), 3–30.
Bailey, D.E., Barley, S.R., 2019. Beyond design and use: How scholars should study intelligent technologies. Information and Organization. Published online 26 December 2019. doi: 10.1016/j.infoandorg.2019.100286.
Baptista, J., Stein, M.-K., Lee, J., Watson-Manheim, M.B., Klein, S., 2017a. Call for papers: Strategic perspectives on digital work and organizational transformation. J. Strateg. Inf. Syst. 26 (4), i–iii.
Baptista, J., Wilson, A.D., Galliers, R.D., Bynghall, S., 2017b. Social media and the emergence of reflexiveness as a new capability for open strategy. Long Range Plan. 50 (3), 322–336.
Barley, S.R., 1990. The alignment of technology and structure through roles and networks. Adm. Sci. Q. 35 (1), 61–103.
Bhimani, A., 2015. Exploring big data's strategic consequences. J. Inf. Technol. 30 (1), 66–69.
Broek, E.V., Sergeeva, A., Huysman, M., 2019. Hiring algorithms: an ethnography of fairness in practice. In: Proceedings of ICIS 2019, December 15-18th 2019, Munich, Germany.
Brohman, K., Addas, S., Dixon, J., Pinsonneault, A., 2019. Cascading feedback: A longitudinal study of a feedback ecosystem for telemonitoring patients with chronic disease. MIS Quart. 43 (4), 1–23.
Brynjolfsson, E., McElheran, K., 2016. The rapid adoption of data-driven decision-making. Am. Econ. Rev. 106 (5), 133–139.
Brynjolfsson, E., Mitchell, T., 2017. What can machine learning do? Workforce implications. Science 358 (6370), 1530–1534.
Citron, D., Pasquale, F., 2014. The scored society: Due process for automated predictions. Washington Law Rev. 89 (1), 1–34.
Constantiou, I.D., Kallinikos, J., 2015. New games, new rules: Big data and the changing context of strategy. J. Inf. Technol. 30 (1), 44–57.
DalleMule, L., Davenport, T.H., 2017. What's your data strategy. Harvard Bus. Rev. 95 (3), 112–121.
Davenport, T.H., Kirby, J., 2016. Just how smart are smart machines? MIT Sloan Manage. Rev. 57 (3), 21–25.
Eisenhardt, K.M., Martin, J.A., 2000. Dynamic capabilities: what are they? Strateg. Manag. J. 21 (10–11), 1105–1121.
Engelbart, D.C., 1962. Augmenting human intellect: A conceptual framework. Summary report AFOSR3223 under Contract AF 49(638)-1024, SRI Project 3578 for Air Force Office of Scientific Research. Stanford Research Institute: Menlo Park, CA. http://www.dougengelbart.org/pubs/augment3906.html (accessed 30 October 2019).
Espeland, W.N., Sauder, M., 2007. Rankings and reactivity: How public measures recreate social worlds. Am. J. Sociol. 113 (1), 1–40.
Faraj, S., Pachidi, S., Sayegh, K., 2018. Working and organizing in the age of the learning algorithm. Inf. Organ. 28 (1), 62–70.
Feldman, M.S., 2016. Making process visible: Beyond boxes and arrows. In: Langley, A., Tsoukas, H. (Eds.), The SAGE Handbook of Process Organization Studies.

SAGE, London, pp. 625–635.
Galliers, R.D., Newell, S., Shanks, G., Topi, H., 2017. Datification and its human, organizational and societal effects. J. Strateg. Inf. Syst. 26 (3), 185–190.
Günther, W.A., Mehrizi, M.H.R., Huysman, M., Feldberg, F., 2017. Debating big data: A literature review on realizing value from big data. J. Strateg. Inf. Syst. 26 (3), 191–209.
Hautz, J., Seidl, D., Whittington, R., 2017. Open strategy: Dimensions, dilemmas, dynamics. Long Range Plan. 50 (3), 298–309.
Iansiti, M., Lakhani, K.R., 2020. Competing in the Age of AI: Strategy and Leadership When Algorithms and Networks Run the World. Harvard Business Review Press.
Jarzabkowski, P., Lê, J., Spee, P., 2016. Taking a strong process approach to analyzing qualitative process data. In: Langley, A., Tsoukas, H. (Eds.), The SAGE Handbook of Process Organization Studies. SAGE, London, pp. 237–253.
Jaton, F., 2017. We get the algorithms of our ground truths. Soc. Stud. Sci. 47 (6), 811–840.
Klein, H.K., Myers, M.D., 1999. A set of principles for conducting and evaluating interpretive field studies in information systems. MIS Quart. 23 (1), 67–94.
Langley, A., 1999. Strategies for theorizing from process data. Acad. Manag. Rev. 24 (4), 691–710.
Leonardi, P.L., 2011. When flexible routines meet flexible technologies: affordance, constraint, and the imbrication of human and material agencies. MIS Quart. 35 (1), 147–167.
Lebovitz, S., Levina, N., Lifshitz-Assaf, H., 2019. Doubting the Diagnosis: How Artificial Intelligence Increases Ambiguity during Professional Decision Making. Available at SSRN 3480593.
Licklider, J.C.R., 1960. Man-computer symbiosis. IRE Trans. Hum. Factors Electron. 1, 4–11.
Lynch, M., 2000. Against reflexivity as an academic virtue and source of privileged knowledge. Theory Cult. Soc. 17 (3), 26–54.
Marjanovic, O., Cecez-Kecmanovic, D., 2017. Exploring the tension between transparency and datification effects of open government IS through the lens of Complex Adaptive Systems. J. Strateg. Inf. Syst. 26 (3), 210–232.
Markus, M.L., 2017. Datification, organizational strategy, and IS research: what's the score? J. Strateg. Inf. Syst. 26 (3), 233–241.
Mazmanian, M., Cohn, M., Dourish, P., 2014. Dynamic reconfiguration in planetary exploration. MIS Quart. 38 (3), 831–848.
McAfee, A., Brynjolfsson, E., 2017. Machine, Platform, Crowd: Harnessing Our Digital Future. WW Norton & Company.
Nedelkoska, L., Quintini, G., 2018. Automation, skills use and training. OECD social, employment and migration working papers, No. 202 OECD Publishing, Paris.
Newell, S., Marabelli, M., 2015. Strategic opportunities (and challenges) of algorithmic decision-making: A call for action on the long-term societal effects of 'datification' J. Strateg. Inf. Syst. 24 (1), 3–14.
Norman, D.A., 1990. The 'problem' with automation: inappropriate feedback and interaction, not 'over-automation' Philos. Trans. R. Soc. Lond. B Biol. Sci. 327 (1241), 585–593.
Parasuraman, R., Sheridan, T.B., Wickens, C.D., 2000. A model for types and levels of human interaction with automation. IEEE Trans. Syst. Man Cybern.-Part A: Syst. Hum. 30 (3), 286–297.
Pavlou, P.A., 2018. Internet of things-will humans be replaced or augmented? GfK Mark. Intell. Revi. 10 (2), 43–48.
Rai, A., Constantinides, P., Sarker, S., 2019. Editor's comments: next-generation digital platforms: toward human-AI hybrids. MIS Quart. 43 (1), iii–ix.
Rahwan, I., 2018. Society-in-the-loop: programming the algorithmic social contract. Ethics Inf. Technol. 20 (1), 5–14.
Ramaprasad, A., 1983. On the definition of feedback. Behav. Sci. 28 (1), 4–13.
Raisch, S., Krakowski, S., 2020. Artificial Intelligence and Management: The Automation-Augmentation Paradox. Academy of Management Review. Published online 11 February 2020. doi: 10.5465/2018.0072.
Saul, J., 2017. Global shipping feels fallout from Maersk cyber attack. Retrieved from https://www.reuters.com/news/picture/global-shipping-feels-fallout-from-maers-idUSKBN19K2LE.
Scheepers, R., Lacity, M.C., Willcocks, L.P., 2018. Cognitive automation as part of Deakin University's digital strategy. MIS Quart. Executive 17 (2), 89–107.
Seidel, S., Berente, N., Lindberg, A., Lyytinen, K., Nickerson, J.V., 2019. Autonomous tools and design: a triple-loop approach to human-machine learning. Commun. ACM 62 (1), 50–57.
Sheridan, T.B., 1995. Human centered automation: oxymoron or common sense?. In: 1995 IEEE International Conference on Systems, Man and Cybernetics. Intelligent Systems for the 21st Century, vol. 1, pp. 823–828. IEEE.
Shrestha, Y.R., Ben-Menahem, S.M., von Krogh, G., 2019. Organizational decision-making structures in the age of artificial intelligence. California Manage. Rev. 61 (4), 66–83.
Simon, H.A., 1973. The structure of ill structured problems. Artif. Intell. 4 (3–4), 181–201.
Suchman, L.A., 1987. Plans and Situated Actions: The Problem of Human-Machine Communication. Cambridge University Press, New York.
Suchman, L., 1993. Do categories have politics? The language/action perspective reconsidered. In: Proceedings of the Third European Conference on Computer-Supported Cooperative Work 13–17 September 1993, Milan, Italy ECSCW'93. Springer, Dordrecht, pp. 1–14.
Suchman, L., 2007. Human-Machine Reconfigurations: Plans and Situated Actions, second ed. Cambridge University Press, New York.
Suchman, L., 2012. Configuration. In: Lury, C., Wakeford, N. (Eds.), Inventive Methods: The Happening of the Social. Routledge, New York, pp. 48–60.
Tarafdar, M., Beath, C.M., Ross, J.W., 2019. Using AI to enhance business operations. MIT Sloan Manage. Rev. 60 (4), 37–44.
Teece, D.J., Pisano, G., Shuen, A., 1997. Dynamic capabilities and strategic management. Strateg. Manag. J. 18 (7), 509–533.
Venters, W., Oborn, E., Barrett, M., 2014. A trichordal temporal approach to digital coordination: the sociomaterial mangling of the CERN grid. MIS Quart. 38, 927–949.
von Krogh, G., 2018. Artificial intelligence in organizations: New opportunities for phenomenon-based theorizing. Acad. Manage. Discov. 4 (4), 404–409.
Walsham, G., 1995. Interpretive case studies in IS research: nature and method. Eur. J. Inf. Syst. 4 (2), 74–81.
Walsham, G., 2006. Doing interpretive research. Eur. J. Inf. Syst. 15 (3), 320–330.
Watson, H., 2017. Preparing for the cognitive generation of decision support. MIS Quart. Executive 16 (3), 153–169.
Waardenburg, L., Sergeeva, A., Huysman, M., 2018. Hotspots and blind spots. A case of predictive policing in practice. In: Schultze et al. (Eds.), Living with Monsters? Social Implications of Algorithmic Phenomena, Hybrid Agency, and the Performativity of Technology. Working Conference on Information Systems and Organizations, IFIP WG 8.2. Springer, Cham, pp. 96–109.
Winter, S.G., 2003. Understanding dynamic capabilities. Strat. Manage. J. 24 (10), 991–995.
Woerner, S.L., Wixom, B.H., 2015. Big data: extending the business strategy toolbox. J. Inf. Technol. 30 (1), 60–62.