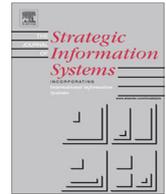




ELSEVIER

Contents lists available at ScienceDirect

Journal of Strategic Information Systems

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

Strategic opportunities (and challenges) of algorithmic decision-making: A call for action on the long-term societal effects of ‘datification’

Sue Newell^{a,*}, Marco Marabelli^b^a School of Business, Management and Economics, University of Sussex, Brighton BN1 9RH, UK^b IPM Department, Bentley University, Waltham, MA 02452, USA

ARTICLE INFO

Article history:

Accepted 6 February 2015

Available online xxxx

Keywords:

Algorithmic decision-making

Digital traces

Sensors

Strategic

Social and ethical issues

ABSTRACT

Today, digital data are captured through a variety of devices that have the ability to monitor the minutiae of an individual's everyday life. These data are often processed by algorithms, which support (or drive) decisions (termed ‘algorithmic decision-making’ in this article). While the strategic value of these data (and subsequent analysis) for businesses is unquestionable, the implications for individuals and wider society are less clear. Therefore, in this Viewpoint article we aim to shed light on the tension between businesses – that increasingly profile customers and personalize products and services – and individuals, who, as McAfee and Brynjolfsson (2012, p. 5) suggest, are ‘walking data generators’ but are often unaware of how the data they produce are being used, and by whom and with what consequences. Issues associated with privacy, control and dependence arise, suggesting that social and ethical concerns related to the way business is strategically exploiting digitized technologies that increasingly support our everyday activities should be brought to the fore and thoughtfully discussed. In this article we aim to lay a foundation for this discussion in the IS community and beyond.

© 2015 Elsevier B.V. All rights reserved.

Introduction

The last decade has witnessed the widespread diffusion of digitized devices that have the ability to monitor the minutiae of our everyday lives (Hedman et al., 2013). Nolan (2012, p. 91) argues that ‘Global IT has enabled information on most everything to flow most everywhere at stealth speed’. The data trail we leave is increasingly used by companies to manage employees and target and personalize products and services for clients and customers, based on developing algorithms that can make predictions about individuals by recognizing complex patterns in huge data sets compiled from multiple sources. In this article we consider some of the observed and potential consequences of this new type of data-driven, algorithmic decision-making, illustrating that while it can offer strategic opportunities for business and sometimes benefits for individuals, there are also costs, hence raising societal issues: as Galliers et al. (2012) indicate, there can be a difference between how *business* is benefiting and how *society* is benefiting – or otherwise.

The IS literature has already raised social and ethical concerns associated with IT (Smith, 2002; Smith and Hasnas, 1999), and in particular those concerns are often associated with privacy issues (e.g., see Belanger and Crossler, 2011; Chan et al., 2005; Coll, 2014; Greenaway and Chan, 2005). However, few IS studies have linked these concerns with the digitization of

* Corresponding author.

E-mail addresses: sue.newell@sussex.ac.uk (S. Newell), mmarabelli@bentley.edu (M. Marabelli).

our everyday life (exceptions include Abbas et al., 2014; Boyd and Crawford, 2014; Lyon, 2014; Slade and Prinsloo, 2013), and fewer still have discussed this phenomenon in relation to algorithmic decision-making (one exception being Schroeder and Cowls, 2014). Here, we focus on the consequences of ‘algorithmic decision-making’, which occurs when data are collected through digitized devices carried by individuals such as smartphones and technologies with inbuilt sensors – and subsequently processed by algorithms, which are then used to make (data-driven) decisions. That is, decisions are based on relationships identified in the data, and the decision maker often ignores why such relationships may be present (Mayer-Schonberger and Cukier, 2013). While these data-driven decisions made by businesses lead to personalized offerings to individuals, they also result in the narrowing of their choices (Newell and Marabelli, 2014).

Given the above, we argue that algorithmic decision-making has societal consequences that may not always be positive and, in this Viewpoint article, we aim to articulate such concerns. In so doing, we bring to the fore the issues related to algorithmic decision-making and highlight the interdisciplinary nature of this topic (Chen et al., 2012; Smith et al., 2011). As we have indicated, some work has been done to shed light on the social implications of the widespread diffusion of digital devices in the IS community, but also in other disciplines such as sociology – as in the work of Lyon (2001, 2003, 2014), Doyle et al. (2013), and Ball (2002, 2005) on impacts of monitoring and surveillance on society, and of Castells et al. (2009) and Campbell and Park (2008) on societal changes determined by the diffusion of digital devices. Here, we call for IS research that examines (and challenges) corporations (and governments) in terms of the strategic decisions that are being made based on data that we are now constantly providing them (see also MacCrory et al., 2014), whether we realize it or not. Next, we define some key concepts and set the boundaries of our analysis.

Big data, little data, and algorithmic decision-making

Data-driven or ‘algorithmic’ decision-making is based on collecting and analyzing large quantities of data that are then used to make strategic decisions. Algorithmic decision-making incorporates two main characteristics: firstly, decision-makers rely on information provided by algorithms that process huge amounts of data (often big data, as we will explain next); secondly, the reasons behind the ‘suggestions’ made by the algorithms are often ignored by decision-makers (Mayer-Schonberger and Cukier, 2013). We expand on both characteristics below.

Digitized technologies and data analytics

Data that originate from digitized devices are increasingly permeating our everyday lives. These digitized devices have the ability to keep track of and record what we do. As a result, somebody else may eventually be able to use the data thus produced – often with purposes different from those originally intended. Thus, we focus on ‘digital traces’ – all data provided by individuals (1) during ‘IT-related’ activities, captured from social networks, online shopping, blogs, but also ATM withdrawals, and other activities that will leave a ‘trace’ (Hedman et al. 2013; Wu and Brynjolfsson, 2009) and (2) that are captured through technologies that we use that have in-built sensors. These technologies include LBS (Location Based Technologies) that are IT artifacts equipped with GPS systems and so have the ability to collect a user’s location such as a smartphone with GPS – see Abbas et al. (2014) and Michael and Michael (2011) for social implications – and other surveillance and monitoring devices – see the previously cited work of Lyon (2001, 2003, 2014) for privacy implications.

It is clear that the huge amount of digital trace data that are collected through the many digitized devices that we now use to support our daily activities fall into the ‘big data’ umbrella. The big data (analytics) concept is very similar to the more familiar (and less sexy) business intelligence that has been studied for the past decade or so (e.g., Negash, 2004; Power, 2002; Rouibah and Ould-ali, 2002; Thomsen, 2003). McAfee and Brynjolfsson (2012). Following Gartner’s (2001) definition, it is the three Vs of big data¹ on which we focus: Volume (the amount of data determines value); Variety (data arise from different sources/databases and are cross-matched to find relationships), and Velocity (data are generated quickly). Big data encompasses much more than this individually generated data trail (see Chen et al., 2012 for a broad discussion of big data analytics) but here we focus just on this everyday digital trail that we each leave. That is, we focus on *those* big data that are generated by individuals during their everyday lives (and are captured as digital traces). In other words, we focus on data that arise as a consequence of each of us now being a ‘walking data generator’ (McAfee and Brynjolfsson, 2012, p. 5). This attention to the digitization of our everyday life allows us to narrow the focus of our inquiry and to expand on concerns regarding the use (and abuse) of one aspect of big data analytics that concerns algorithm-driven decision-making and associated personalization – to which we now turn.

Algorithmic decision-making

(Big) data captured through digitized devices are processed by algorithms aimed at predicting what a person will do, think and like on the basis of their current (or past) behaviors. These algorithms can predict particular outcomes, as with

¹ The definition of big data was updated by Gartner in 2012 as they now describe the concept as ‘high volume, high velocity, and/or high variety information assets that require new forms of processing to enable enhanced decision making, insight discovery and process optimization’ (Gartner, 2012). Moreover, others have added ‘new Vs’ – e.g., veracity, variability, visualization, and value, viewing big data in terms of 5 or even 7 Vs. Here, where we stick with the original definition (Gartner, 2001) as this reflects the essence of big data for the purposes of this article.

the numbers of ‘friends’ on Facebook being used to predict a person’s credit risk (<http://www.google.com/patents/US8560436>) or an individual’s Facebook ‘likes’ on a college Facebook page, used to predict her/his willingness to become a donator (http://www.nytimes.com/2015/01/25/technology/your-college-may-be-banking-on-your-facebook-likes.html?_r=0). Interestingly, these predictions often represent a black-box: while humans must decide what to measure and produce the algorithms to analyze the data being collected, these decisions do not necessarily involve understanding the causes and consequences of particular patterns of behavior that are identified (Mayer-Schonberger and Cukier, 2013). Rather, it is deemed sufficient that connections are discovered. Traditionally, making decisions has been a human-centered, knowledge-based activity with humans discriminating on the basis of an understanding of theory or context (Tsoukas and Vladimirov, 2001). By contrast, algorithmic decision-making means that discriminations are increasingly being made by an algorithm, with few individuals actually understanding what is included in the algorithm or even why. In other words, it is seen as being sufficient that an algorithm is successfully predictive, never mind if the reasons for the associations found in the data from different sources are unknown. We argue that this is likely to create problems when no one in a corporation really understands why some decisions are made. For example, one could argue that the last financial crisis was at least partially a product of this problem, with the algorithms that predicted the pricing for mortgage-backed securities clearly not taking into account all the risks while at the same time not being subject to question because the basis of the algorithm was neither clear nor easily accessible, either to the senior managers in the financial institutions where the algorithms were being used or to the credit rating agencies who were evaluating these products (Clark and Newell, 2013).

In sum, here we focus on data collected through digitized devices that we increasingly use to support our everyday activities. This is ‘big data’, because the three (or more) Vs of Gartner’s (2001, 2012) definition apply. In fact, data coming from digitized technologies are high in volume because of the widespread diffusion of digital devices that allow access to social networks at any time, as well as all other types of technologies that record what we do even if we do not ‘own’ them (e.g., surveillance cameras, or an ATM card machine, where the usage information goes into our bank’s database). Thus, data come from different sources (variety). For instance, data used for making ‘algorithmic decisions’ may come from a combination of contributions on social networks and LBS systems (e.g., a ‘check in’), or spending capacity of consumers associated with personal facts of individuals (e.g., the partner’s birthday). Data velocity is clearly another characteristic of the digitization of our everyday life, because we are ‘walking data generators’ 24/7 and ‘More data cross the Internet every second than were stored in the entire Internet just 20 years ago’ (McAfee and Brynjolfsson, 2012, p. 4). On this point, it is worth noting that most of the digitized devices that collect such individual level activity data fall under the Internet of Things (IoT) umbrella (Miorandi et al., 2012; Xi et al., 2012). However, we do not restrict our analysis to those digitized devices that are connected to the Internet because some devices remain (for now) independent of the Internet (e.g., some OBD devices). One such example is provided by Progressive Insurance in the USA (<http://www.progressive.com>), which provides a memory stick that is plugged into a car’s on-board computer and the data must be uploaded to the insurance company rather than automatically sent via the Internet.

Potential discriminations associated with the (ab)use of algorithmic decision-making: big and little data

The use of algorithmic decision-making associated with data coming from the digitization of our everyday lives improves the capacity of a business to make discriminations. Thus, businesses have always discriminated in terms of to whom they offer products and services, because products and services are targeted to different audiences (we cannot, unfortunately all afford to buy a Bentley car). With algorithmic decision-making they are simply taking this a step further. For example, they can now much more precisely target and personalize offerings to customers and potential customers – those predicted to buy particular products or services. As a more specific example, a car’s computer that monitors speed, usage of brakes, horn, lights, etc. (such as Progressive Insurance’s OnStar OBD technologies mentioned above) has the ability capture all these details that are then sent to data centers. Computers then analyze the (big) data and insurance companies are able to use the results to discriminate (e.g., by charging young men higher premiums because the data indicate that they – generally – drive less safely than other categories of drivers). Such data-driven decision-making has been questioned because it can go against the ethical principle of equal or fair treatment. This is exemplified in the recent case in the EU, where insurers are required to no longer use statistical evidence about gender differences to set premiums. Thus, despite the fact that gender differences are clear from the data (e.g., young male drivers are ten times more likely to be killed or injured than those – of both sexes – over the age of 35; women live, on average, longer than men), it is considered to be discriminatory (following an EU ruling that came into effect in December 2012) to use this trend evidence to differentiate between premiums (e.g., car insurance or actuarial rates) for men and women. The point about this change in the law is that it was considered to be discriminatory because, for example, while young men in general may drive more recklessly and so be more prone to accidents, an individual young man may not and would therefore be discriminated against when insurers set premiums based on group trends observable in collective data.

While using big data and algorithmic decision-making to observe trends and so discriminate between groups of individuals can have social consequences that are potentially unfair, this targeting can now be taken further when data are used not to predict group trends but to predict the behavior of a specific individual. This is sometimes described as ‘little’ data – although it should be noted that little data are actually based on big data but are simply used in a more targeted way. Thus, little data focuses on the everyday minutiae of specific individuals, using computing capacity to collect extremely granular data (Munford, 2014). Drawing on the previous example of a car’s OBD, little data can now allow us to concentrate on a

specific driver, and we can decide whether an individual is a good or bad driver based on the sensor data from his/her car. Sensors have the ability to capture individual's behaviors and are widespread. As an illustration, consider that approximately 85% of handsets now have a GPS system chipset installed (Abbas et al., 2014). By using sensor data, the insurer would not be setting premiums based on the general trends in accident rates between groups, but instead would base their calculations on the actual driving habits of an individual. However, if little data are more 'objective' in terms of discriminations made by corporations, it probably poses more issues for societies given the observed or potential social consequences; for instance, in terms of an individual's privacy (Lyon, 2014) or in terms of the exploitation of the vulnerable – an issue that IS scholars seem not to have fully addressed as yet.

It is then clear that algorithmic decision-making poses two main concerns in terms of big and little data: first, (in terms of big data) this data trail provides the opportunity for organizations to move to algorithmic decision-making, which McAfee and Brynjolfsson (2012) argue, is superior to traditional 'HiPPO' (highest-paid person's opinion) decision-making. Algorithmic decision-making is, they argue, superior to human judgment-based decisions because of all the inherent biases in human judgment (Hodgkinson et al., 2002). However, we question this assumption because making decisions on the basis of big data (and algorithms) might create unfair discriminations. Second, we argue that monitoring an individual's behavior poses societal concerns since "the digital artifacts will be able to remember where they [individuals] were, who used them, the outcomes of interactions, etc." (Yoo, 2010, p. 226) and this often happens without individuals even being aware that they are being monitored. Thus, we posit that undertaking research to assess these societal harms, so that corporations can be held responsible and citizens become more aware, can potentially be very useful.

Below we identify three tradeoffs that involve issues associated with the use by corporations (and governments) of data from digitized devices that support our daily activities, and in particular with the strategy of using data analytics. The first of these considers the privacy of individuals versus security for society – an issue that is preeminent in people's minds following the recent terrorist attacks, particularly in Paris, in January 2015.

Tradeoffs and societal issues associated with big (and little) data

Privacy versus security

Digitized devices can improve security, and examples include the security-tracking systems adopted for prison populations, when prisoners are released but are required to wear a tracking ankle-bracelet. These systems are aimed at improving the overall security of our society, with the sensor acting as a deterrent for prisoners to escape or commit a crime when they are on parole. Other instances where security is enhanced by everyday digitized devices is in the capacity of sensors to trace a stolen device, or a kidnapped child, as in the case that occurred in September 2013 in Texas, where the Houston police were able to trace the whereabouts of a kidnapper by tracing the iPad that he had with him in his car (<http://abc13.com/archive/9242256/>). A similar example relates to police authorities being able to detect a crime because it is all 'caught on tape', for example with sensor-activated security cameras and, potentially, Google Glass or other camera-based devices that are now routinely carried by many.

All these examples of companies, government agencies and private individuals using digitized technologies to increase security come at some costs in terms of individuals' privacy. In terms of locating a lost smartphone, it has to be the user who, deliberately, accepts giving up her/his (right of) privacy by activating the 'find my phone' option (<https://itunes.apple.com/us/app/find-my-iphone/id376101648?mt=8>). The example of Google Glass or digital cameras worn, for example, by cyclists or skiers to record their journey, is more complex since the privacy that a user gives up affects others' privacy, thus representing a shift from the individual to the societal level. In some circumstances one's use of social software applications affects others' privacy, as for example, for people who are tagged in somebody's Facebook profile without them knowing. Perhaps not surprisingly, privacy advocates have argued that in these types of exchanges consumers are justified in expecting that the data they collect and share should remain private among those to whom it was originally disclosed – dependent on users' risk perceptions, as noted by Gerlach et al. (2015) – rather than being shared with third parties who may subsequently behave opportunistically (Beldad et al., 2011; Petronio, 2002).

Thus, it is clear that improving security across society, based on digital devices, can impact on individual's privacy. Some companies are doing something about this. For instance, Facebook no longer allows a user's 'wild tagging' and, instead, an automatic email is sent to a user who is tagged, for approval (or at least this is a configurable option under privacy settings). Yet, the exponential diffusion of tracking software embedded in social networks such as Facebook and the sensors and cameras in many other digital devices lead us to think that it will be hard for organizations (or governments) to regulate how individuals use responsibly technologies that enable tracking (i.e., in a way that balances security and privacy). The societal issue is raised because the move towards using devices and applications to gain increased security comes at the expense of reduced privacy. This points to a question about whether users (and more broadly society) want to give up some security potential to ensure more privacy (Culnan and Williams, 2009; Velasquez, 2006). This is a decision that citizens need to debate with their politicians (Dinev et al., 2008) and that governments in turn need to debate with businesses, since it is businesses that collect and analyze digital traces. This is exemplified by the Lee Rigby case (the British soldier killed on a street in London), where Facebook was accused of not helping to protect security because it did not use its analytical capability to detect and report the fact that the killer was posting that he was intending to commit just such a murder (<http://www.theguardian.com/uk-news/live/2014/nov/25/lee-rigby-woolwich-inquiry-report-published-live-coverage>).

Other privacy/security tensions are reflected in the current debate on whether US police officers should wear cameras following recent cases involving police officers' improper use of force (see for instance the developments over the Michael Brown case <http://www.cnn.com/2014/12/04/us/eric-garner-ferguson-body-cameras-debate/>). Here, a sensor technology would be employed but would not actually generate data that will be processed by algorithms, since the camera records would be reviewed only in particular circumstances. However, this and other types of sensor are pervasive (Lyytinen and Yoo, 2002) (and invasive), and the data (e.g., the camera records) would be stored. In such circumstances, we do not know whether in the future somebody will develop an algorithmic-based decision system to analyze the data (e.g., to assess the performance of police officers). It is thus clear that the widespread diffusion of digitized technologies can be harmful to individuals' privacy while justified by corporations and governments in the name of public security – a tradeoff widely discussed by Lyon (2003, p. 79) in terms of ID cards that are supposed to improve national security in that he raises the issue of whether as citizens we are willing to “pay the price in liberty for security”. This tradeoff, then, raises complex social issues because of the ready availability of these data and because of the capacity of algorithms to discriminate almost in real time – for instance, to determine that particular categories of people (based on race, income, job, etc.) are more likely to commit a crime, and could, therefore, be subjected to higher levels of policing and potentially also face discrimination in other areas (Lyon, *ibid*). This, therefore, pits an individual's privacy against the security of society, but also suggests broader social issues in terms of freedom versus control, as we articulate next.

Freedom versus control

The 'datification' of everything means that we can use devices to constantly track every decision made and place that a person visits (be they an employee, a citizen, or our child), and use these data to monitor and control (some now prefer to use the term 'nudge') behavior (Whitman, 2011). This second tradeoff between freedom and control is more complex than the previous one because, here, individuals can be aware that they are being controlled. This is informed control (e.g., because they are required to carry RFID badges at the work place or maybe even have chips implanted under their skin, another example of invasive technology – <http://www.bbc.co.uk/news/technology-31042477> – or because they decide to use an electronic collection system in their car to drive through tolled roads and bridges). However, they can also be controlled without knowing that they are being monitored. This is uninformed control. Uninformed control happens, for instance, when tracking cookies monitor someone's online activity, or, more generally, when 'second hand' use of data originating from digitized technologies are used.

Freedom versus informed control

Surveillance based on parents tracking their children's every move (e.g., using an application on the child's smart-phone) is clearly informed control and would allow parents to feel that they are in control of their children's movements. However, the loss of freedom (and privacy, as we have already pointed out) of those subjected to this surveillance might have far-reaching effects, for instance in terms of children's feelings of personal responsibility. After all, we know that punishment is not always an effective deterrent because, once the punishment is removed, the individual often resorts to the prior (undesirable) behavior; so, if individuals conform only because they know they are being monitored, will their behavior change once the monitoring ceases? With constant surveillance, like punishment, while we may change behavior, the beliefs about what is appropriate or inappropriate may remain (Podsakoff et al., 1982; Staples, 2013). This tension, then, is between improved control (by business but also government or private citizens) at the expense of individuals feeling that they have some freedom and autonomy – a feeling that we know has a significant influence on motivation in the long-term (Hasan and Subhani, 2011). One such example is Hitachi's new digital identification badge that collects data on individual employees' exact location within an office, records who the person has spoken to, for how long and how energetically (<http://www.cnn.com/2014/02/02/opinion/greene-corporate-surveillance>). Adopting this kind of technology as a strategy for constant monitoring may, however, affect employees' motivation and perhaps also their capacity to produce disruptive innovation. Indeed, productivity might benefit (at least in the short term) from such an aggressive approach to control in the workplace. However, the longer-term consequences of such constant surveillance may be more problematic. For instance, Lyon (2003, p. 20) points out that a “surveillance system obtains personal and group data in order to classify people and populations according to varying criteria, to determine who should be targeted for special treatment, suspicion, eligibility, inclusion, access, and so on”, arguing that such “social sorting” leads to long-term discriminations. He states that “data about transactions is used both to target persons for further advertising and to dismiss consumers who are of little value to companies” (*ibid*,; 1), leading to long-term social differences. Moreover, breakthrough innovation, which is more risky and leads to more frequent 'failures' (O'Reilly and Tushman, 2004), might be jeopardized because individuals who are constantly monitored are less likely to expose themselves to failure in front of their peers and superiors. This suggests that those making strategic decisions about how to use this new tracking technology (whether business, government or private individual) might want to think about reducing the amount of surveillance on employees, customers, family members or citizens since this would be the price they might want to pay for allowing people to feel in control of the decisions they make – in other words, being informed and not automated – to use the language of Zuboff (1984). This supports our argument that a tradeoff emerges between control and freedom in the context of the digitization of our everyday lives.

Freedom versus uninformed control

The feeling of being controlled, as we have discussed, might lead to some unwanted consequences (e.g., loss of a sense of responsibility or lower productivity in the work place). However, probably a more relevant societal issue emerges when control (over an individual's freedom) is made without the individual even knowing that she/he is being controlled (when this is not made explicit or is not requested). To this end, here we provide an example involving individuals' online activities, where the 'free' access to information is increasingly controlled as Internet companies (social networks, news, etc.) now determine (based on algorithms) what we see. For instance, we may see many posts about the newest iPhone (6, at the time of our writing) and assume that many of our Facebook friends are posting articles about this new technology. However, the frequency with which we see these posts may be partially due to us having clicked on an advertisement related to the iPhone 6: Facebook's algorithm decides that we are interested in such products and then shows us others' posts that are related to the iPhone 6. A consequence of such use of algorithms by corporations to decide – for the consumer – the posts, news or advertising that they are exposed to, is that it may lead to a slow and often subtle manipulation of consumers' worldviews as well as to new forms of discrimination. Simply put, what is presented to the reader is decided by an algorithm – tapping into prior searches – and is not based on an explicit personal choice. An example of uninformed control by a corporation that produces worrisome societal issues is found in the account presented by Eli Pariser, who showed that "Facebook was looking at which links I clicked on, and it was noticing that I was clicking more on my liberal friends' links than on my conservative friends' links. And without consulting me about it, it had edited them out. They disappeared." (Pariser, 2011). In the longer term, this manipulation by corporations of what the consuming public is exposed to – exposing us only to things that we like (or the things that an algorithm assumes we like) – may produce societal changes. For instance, our exposure to online diversity will be reduced, as in the example of Eli Pariser. More recently, Greg Marra, a Facebook engineer argued that, "We think that of all the stuff you've connected yourself to, this is the stuff you'd be most interested in reading", explaining further that an algorithm monitors 'thousands and thousands' of metrics to decide what we should see on our Facebook page. These metrics include what device we use, how many comments or 'Likes' a story has received and how long readers spend on each article/post. The assumed goal, as a New York Times article suggests, is that companies are using algorithmic decision-making "to identify what users most enjoy" (http://www.nytimes.com/2014/10/27/business/media/how-facebook-is-changing-the-way-its-users-consume-journalism.html?_r=0). However, this also indicates that this practice of showing us only things that 'fit' with our (little) data profile, limits our possibility to choose, and might inhibit our capacity to make informed decisions (on what we buy and even what we think).

These strategies, then, that are adopted by organizations to allow them to tailor results and personalize offerings to individual consumers are leading to citizens (i.e., all us of who 'surf the web') being exposed to less and less diversity online. A potential consequence is that we may become less tolerant to diversity, meaning that we may as a result become less able to listen to someone who thinks differently (e.g., a Republican, in Pariser's example). Moreover, there may be other, more worrying consequences in the long-term that are associated with race-diversity intolerance and the increased exploitation of the vulnerable. For example, in relation to the latter issue, if algorithms work out who is less capable of making good financial decisions, personalized advertisements can then be sent persuading these people to take out risky loans, or high-rate instant credit options, thereby exploiting their vulnerability. The strategic use of our own data by corporations to personalize our Internet, in other words, is just another and potentially more pernicious way of allowing discrimination; pernicious because the only person who has access to the outcomes of the discrimination is the individual being discriminated against (who is often not aware of the fact that they are exposed to discriminatory information – uninformed control), making it easy for unscrupulous businesses to use personalization in a way that harms the vulnerable.

Another way to illustrate how societal concerns emerge as a consequence of businesses (and governments) using data from the digitization of our everyday life is by articulating the tradeoff between independence and dependence, to which we now turn.

Independence versus dependence

Citizens in many countries increasingly depend on digital devices for many activities. However, here, a tradeoff originates from the tension between the willingness to depend on IT devices and being able to live without them (i.e., to be independent), should the need arise. Think of our decreasing sense of direction due to our dependency on GPS systems or, were we to consider safety issues, think of those sensor-based systems that are able to park our car – or even drive it! These driving systems use onboard cameras and laser rangefinders to identify obstacles (or hazards, if the onboard computer controls the car while it is 'in motion'); then an algorithm is able to scan the surrounding environment and to identify safe zones, avoiding for example other cars (see for instance a 2012 MIT study on these algorithms – <http://newsoffice.mit.edu/2012/mechanical-engineers-develop-intelligent-car-co-pilot-0713>). In the case of car autopilots, algorithmic decision-making has a twofold role: first, data on 'real drivers' are collected and analyzed so that the algorithm can make the appropriate decisions (e.g., reacting as a driver would, but with the difference that the algorithm is never tired or inattentive, thus carrying a safety advantage with respect to humans). Second, sensors embedded in cars (laser rangefinders, in this example) collect environmental data that are analyzed in real time, so the algorithm has the ability to either assist the driver by supporting his/her decisions (with warnings to the driver) or to make decision on its own – when the car is in full autopilot mode. Drivers, thus, are somewhat 'tempted' to benefit from the comforts companies now design into their products using digital technology, but this necessarily takes place at the expense of our future independence. In fact, if our car computer emits a

warning signal while we drive on a highway, suggesting that we should slow down, we might argue that it is because the GPS embedded in the car computer has just received a 'traffic warning' or, because the weather channel is broadcasting heavy rain in minutes, or because we are about to drive through a road work area, but we do not really know the actual reason of the warning, yet we slow down – this (again) illustrating that algorithmic decision-making incorporates advantages (in this context, for users) but at the same time precludes a full understanding of why some decisions are being made. This limits learning through practice (Brown and Duguid, 1991) that in the long term might modify an individual's ability to learn new tasks and, more generally, adapt to the workplace or to society more generally (Dall'Alba and Sandberg, 2010; Nicolini et al., 2003).

While it is certain that there are good reasons for companies designing and for users adopting these automated systems, as we saw, this might also lead to a change in our ability to undertake particular activities without sensors, and learn. In the example of the autopilot, once our car parks itself, will we forget how to park on our own? IT-assisted systems have been around for a while in commercial planes, but pilots are constantly trained on how to pilot a plane in case the autopilot stops working. However, would individuals be trained on how to drive a car once such 'autopilot' systems become common in private motor vehicles? This example brings to the fore the point that digital technologies and devices (and the associated algorithmic decision-making) are increasingly influencing and even managing our lives, leaving unanswered the question on whether these algorithms are just supporting our activities, or whether they are actually in charge (e.g., controlling what we do) – and if they are taking over, does this excess of control occur at the expense of our ability to improvise and respond to emergencies? Thus, in the car example, it is clear that issues associated with safety emerge. In fact, as car drivers who now rely on sensors, we do not have the luxury that airline pilots have, of simulators to ensure that we maintain our skills so that we are prepared for an emergency. Nevertheless, even pilots (who, unlike private citizens, are trained on how to operate aircrafts) are not free of the consequences from technology that 'takes over', as the US NTSB (National Transportation Safety Board) reports in relation to some major plane accidents (see for instance the case of Air France Flight 447 in 2009, <http://spectrum.ieee.org/riskfactor/aerospace/aviation/air-france-flight-447-crash-caused-by-a-combination-of-factors>).

The negative consequences of an excess of IT dependence are associated with the risks we are exposed to when we forget how to do certain things. Progress necessarily involves automation (and the loss of certain manual capacities), and many innovations developed by corporations positively contribute to our quality of life (and to our safety). However, it is digital devices and the associated algorithmic decision-making that pose issues, especially when supervising or undertaking human activities that might involve life-threatening outcomes were the technology to stop working. Moreover, because of the connectivity between sensor devices, there is also the potential of chaos occurring if everything stops working for everyone simultaneously. In particular, we argue, it is the diffusion of such IT automations among common citizens that creates threats were we to become fully dependent on the technology and unable to operate without it. However, the adoption of some of these automations is (or will become) virtually mandatory for many – creating discriminations against those who do not conform. One simple example relates to US residents who, if desiring to use cars equipped with a standard stick shift (instead of an automatic), will have to pay more, just because 'standard' is not a standard in the US. On the other hand, those who can only drive automatic cars will have to pay more if they want to rent a car when they travel overseas, because most cars will have a standard shift and there will be a premium for an automatic car. This point raises the issue of the role of business in promoting such (automated) digitized technologies: does business have a responsibility for thinking about such consequences and building in opportunities for learning to reduce our over-dependence?

In sum, we argue that this tradeoff is affected by the willingness of users to give up some of the comforts that come from dependence on IT, in the interests of preserving their ability to cope when the IT does not work as expected. Yet, digital technologies are extremely tempting, and now widely adopted. For instance past research on mobile technologies has already shed light on users' needs to be 'always on', with the consequence that a feeling of 'dependency' arises (Jarvenpaa and Lang, 2005). However, here we go beyond the psychological feeling of dependency and point to the users' need to be somewhat assisted (if not led or managed) by digital technology (that involves algorithmic decision-making) – with discrimination being the consequence of not conforming to this dependence. Companies too need to include sensor-based technologies in their products and services to remain competitive. For example, a logistics company that does not use GPS-equipment to determine best routing opportunities would experience difficulties in finding partners to develop a supply chain – being thus discriminated against (again, for not conforming). However, we suggest that companies might also usefully start to think about how they can and should, in some circumstances at least, support the development of the ability to cope with situations of technology failure, with the consequence that algorithms assist decision-makers but do not entirely take over from human judgment. In our view, a balance must be struck, which to date seems to favor increasing dependence on IT over being able to cope in the face of IT failure.

We have thus far identified three key tradeoffs: between privacy and security, control and freedom, and dependence and independence, which are obviously inter-related. We do not claim that these are the only tensions that are relevant; however, they do allow us to provide concrete examples of strategic opportunities for businesses as well as societal issues emerging from corporations' (and governments') exploitation of data coming from the widespread use of digital technologies that support – and impact – our everyday lives. In the next section we discuss the more general social issues arising from these tensions.

Social consequences of digital technology and algorithmic decision-making

While in the past knowledge and learning have been recognized as path-dependent (Cohen and Levinthal, 1990; Zahra and George, 2002), in this era of widespread diffusion of digital technologies that capture our everyday activities, our awareness about things appears to be not so much path-dependent as *determined* by our past actions and algorithmic rules. For example, the algorithm EdgeRank is used by Facebook to weight 'likes' and modify an individual's Facebook page as a result, therefore manipulating the 'wisdom of the crowd' (Kittur et al., 2007). While this may make sense from a marketing perspective for businesses (e.g., it is helpful to identify a customer's interests), it poses concerns for society because of the potential broader and longer-term social impacts. More specifically, our examples for each tension suggest that businesses (and at times governments and private individuals) are generally in favor of a more secure society over an individual's privacy, of a more controlled population (employees, customers and citizens) over individual freedom – leaving more and more people increasingly dependent upon technology, at the expense of personal independence.

To consider these issues, we start from the premise that digital trace data *is here to stay*; companies will increasingly include tracking software and sensors in the products and the services they offer, and so collect masses of data on our everyday habits with a view to using these data to develop algorithms that drive decision-making. In fact, whether data are gathered from social networks, an ATM transaction, or from a sensor-based device, there are many aspects associated with companies using such data that many users want, hence it is unlikely to 'go away'. As a result, businesses will keep exploiting big and little data potentials to profile customers, please social network users, and grant (commercial) opportunities to those who, at best, accept being controlled, reducing their need to learn while giving up some privacy. On the one hand, individuals benefit from corporations' use of big/little data analytics – one can save some money on an insurance policy, access a free show if willing to watch commercials, or just be pleased to see that everybody thinks her/his way (see the Facebook experiment by Pariser, above). On the other hand, businesses are aware that improving their knowledge about employees and customers will lead to more control of employees and more (addressed and effective) sales, and therefore more profits. And it is this enticement by the business world that leads people to assume that they have to give up some privacy/freedom/independence, whether this is because it is a way to access a line of credit to buy a house or because they want to use social networks to fulfill their social needs (Kane, 2014).

As we previously pointed out when we provided our definition of algorithmic decision-making, this might lead to very superficial understandings of why things happen, and this will definitely not help managers, as well as 'end users' build cumulative knowledge on phenomena. Since decisions are made following an algorithm, how the algorithm came up with a particular result is unknown; as a result, there will be very little opportunity to learn from mistakes. Ironically, therefore, decision-makers might be losing the capacity to make decisions on their own, thereby making them a good example of (over) dependence on digital technology and algorithmic decision-making (cf. our third tradeoff). However, perhaps more important (from a societal perspective) than the lack of lessons learned, is the need to discuss the creation of new forms of discrimination as a result of algorithmic decision-making and the associated personalization of information. Think again of when algorithms determine that particular categories of people (e.g., based on race, income, job) are more likely to commit a crime and, as a result, those concerned find difficulty in obtaining a loan or changing job, never mind being subjected to tighter police scrutiny. This clearly violates basic privacy rights, but is justified based on the idea that it will increase security in society. Or, think again of the control exercised by algorithms in sensor-equipped cars on teenagers: these data are used by insurance companies to decide whether a driver is good or bad, again on the basis of an algorithm (the tradeoff between control and freedom). Similarly, when we give new technologies the possibility to take over our learning and let our car park and drive for us (or let our car 'suggest' what we should do in order to perform a perfect parallel park), our decision-making is being driven by algorithms (the tradeoff between independence and dependence).

These tradeoffs operate together rather than independently. For instance, if we use an app. that 'knows' what music we like so that, when we start driving, we do not need to search for a particular song, this is because we have enabled functionality on our radio/phone that is able to 'see' our favorite playlists, or that looks into our past purchases. For instance, iTunes Genius works with a 'secret algorithm' created by Apple that compares our library of tracks to all other Genius users' libraries and considers complex 'weight factors' to then come up with the appropriate playlist for a specific user (Mims, 2010). Here, we do not aim to go into technical details on how these algorithms work – as Apple engineer Erik Goldman said, the algorithm is 'secret', jokingly noting that 'if he told you how Genius works, he'd have to kill you' (Mims, *ibid.*) – to highlight the relevance and commercial value of these algorithms. However, this example reflects how, in some circumstance we are literally at the mercy of an algorithm, which makes a decision for us. What if we look for vegetarian food just because we go out for dinner with friends who happen to be vegetarian? Does this mean that, due to the connections between databases of large companies (or because of tracking cookies), we will be denied the opportunity of seeing advertisements for steakhouses on our Facebook webpage? Or will we be classified as good drivers because a sensor detects that most of the time we obey speed limits (even if the reason is that we know that where we drive the speed limits are strictly enforced)? Or will our preferences in terms of the music that we listen to be so reinforced by the automatic selections based on the algorithm that we reduce our exposure to alternative genres?

It is clear that the three tradeoffs showcase interests and needs of individuals on the one hand, and the somewhat opportunistic strategic moves of businesses (and governments) on the other. Moreover, our discussion illustrates the relevant role of algorithms in making decisions about an individual's characteristics/preferences based on trends, and therefore about

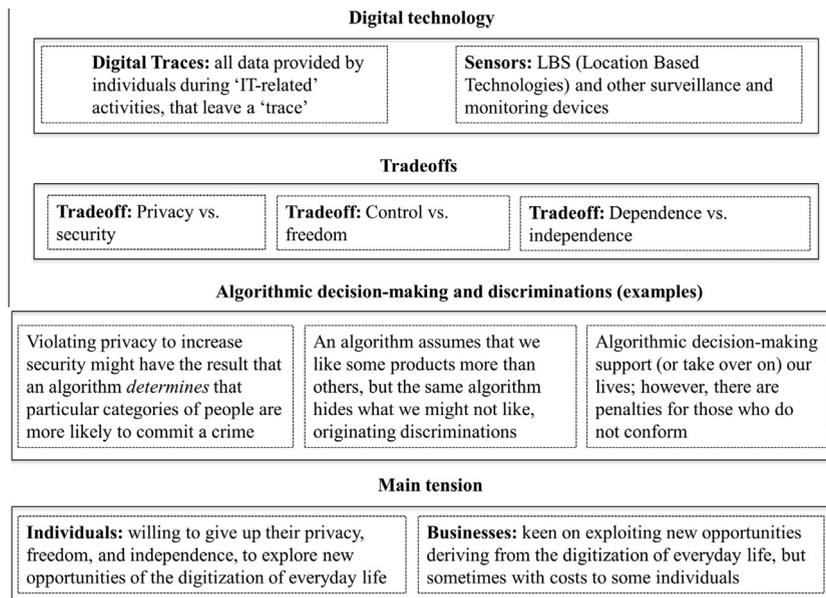


Fig. 1. A summary framework of the consequences of an algorithm-based world.

what individuals should see and are likely to buy. Eric Schmidt (Google) said in 2010 that, "It will be very hard for people to watch or consume something that has not in some sense been tailored for them". This statement involves privacy issues (businesses will know almost everything about consumers), control issues (consumers are literally monitored and then controlled with choices made for them), and dependence issues (loss of independence in making informed decisions, since the information provided about a specific choice will be driven by online feeds – created by algorithms). We posit that IS research is needed to examine the social issues that are emerging in relation to the strategic uses made by corporations of data from the digitization of our everyday lives. With the intention to provide an overall picture of the ethical issues and challenges created by this increasing trend, above we present a framework (Fig. 1) that illustrates how digital technology (first layer) generates tradeoffs (second layer), when this technology is combined with algorithmic decision-making (third layer), leading to tensions (fourth layer). This summary framework has the purpose of showcasing strategic opportunities as well as societal challenges in an era of widespread diffusion of digital technology, and of supporting further interdisciplinary research on this topic, along with our suggested new avenues of research and potential research questions (in the last section that follows).

Research agenda and concluding remarks

We do not know for sure the extent to which digital technology and the associated big/little data analytics are going to impact society in the long term. However, we suggest that individuals seem to be likely to accept the 'dark side' of datification through digital traces (always there), and constant monitoring through sensors because they are persuaded that the benefits outweigh the costs. Thus, businesses (and governments) try to send to citizens the message that security is more important than privacy (to fight terrorism, for instance). And the same businesses make us believe that if we want to quickly find what we are looking for (whether it is a movie that we like, through Netflix, or a specific piece of information, through Google) we need the support of algorithms, that 'know' us and what we want – precluding our exposure to diversity. And finally, businesses develop digital technologies that 'help' us do new things more quickly, but simultaneously make us more reliant on (and so more vulnerable to) these same technologies as well as reducing our ability to learn.

Therefore we suggest that research should be carried out that considers broad social issues associated with businesses' (and government's) strategic use of data, especially so because we currently have very little understanding of what the consequences of corporations' non-responsible use of these data will be for society (Tene and Polonetsky, 2013), albeit we have suggested some negative impacts above. One way of looking at these social issues may be using an ethical dilemma lens, where we consider individuals' right to maintain their privacy, freedom and independence, against businesses' right to discriminate to promote sales – using cutting-edge technology such as big data analytics. We suggest that such dilemmas can be addressed using the teleological or deontological approaches to ethics (or both). The deontological approach (utilitarianism) is the best-known consequentialist theory (Bentham, 1776; Mill, 1863), and suggests that ethical behavior is one that maximizes societal welfare while minimizing social harm (Vallentyne, 1987; Berente et al., 2011). According to this approach, insurance companies may be right in being algorithm driven – thus applying higher premiums to those who, according to the data analytics, are more at risk of having car accidents. However, the (contrasting) deontological approach

bases ethical decisions on broad, universal ethical principals and moral values such as honesty, promise keeping, fairness, loyalty and rights (e.g. to safety, privacy) so that the process or means by which an individual does something, rather than the outcome, is the focus of decision-making (e.g. lying is dishonest as it is one's duty to be honest regardless of whether this might lead to some ultimate good), therefore the end never justifies the means (Mingers and Walsham, 2010). According to this latter approach, discriminations should not take place if a minority is adversely and unfairly treated – never mind if following an algorithm maximizes positive consequences for society. More specifically, here we want to identify research questions that examine the social issues related to each of our tradeoffs, as described next.

Firstly, in terms of the tradeoff between privacy and security, one aspect that deserves attention relates to how far different countries (and regulators) will balance this tradeoff. From an ethical perspective, for instance, they might choose to privilege the maximization of societal welfare (so taking a teleological approach) or to pay attention to 'minorities', who are penalized by the discriminations of algorithmic decision-making (deontological approach). Information privacy laws and regulations – how citizens perceive (the value of) privacy – are country-specific and are related to cultural and historical issues (Milberg et al., 1995). One example is the recent debate about the 'right to be forgotten' (http://ec.europa.eu/justice/data-protection/files/factsheets/factsheet_data_protection_en.pdf), which forced Google to delete some information (and to implement processes to do so in the future, should people ask) from the results of its search engine in Europe, while this issue is not perceived currently as a relevant one in other countries such as in the US. To this end, it would be interesting to dig deeper into research questions such as: 'How far do institutions and governments influence the balance between privacy and security associated with digital technologies that collect data on our everyday lives?' 'What are the historical, cultural and social reasons behind the variety of approaches to digital privacy adopted by different countries?' 'Do these different approaches reflect differences in citizens' ethical and moral values?' 'Do (or will) social networks have the ability, in the long term, to modify ethical and moral values about privacy in different countries?' 'Will the diffusion of digital technology (and the IoT) lead to the standardization of ethical and moral values across countries, in the long-term?'

In terms of the tradeoff between freedom and control, we know very little about how far users are aware that they are being controlled by large Internet companies (especially if we think of 'second hand' data), and if they are, it would be interesting to learn about whether individuals' needs to enact social networks (Kane, 2014) prevail over the potentially uncomfortable feeling of being profiled (little data). Moreover, we do not have specific quantitative data that illustrates the effectiveness of algorithmic decision-making in identifying people's needs – for instance we know that little data has the ability to 'know' (or assume) what people want to purchase on the basis of a number of digital traces, but little is known about the actual revenues that derive from this – and whether the costs of implementing 'smart' algorithms and maintaining expensive hardware that can process big data is covered by the increased sales. We should assume that businesses achieve positive bottom lines from big data, since datification and algorithmic decision making is widely adopted and is expensive (<http://www.forbes.com/sites/ciocentral/2012/04/16/the-big-cost-of-big-data/>), but we do not know, in concrete terms, the extent to which this has improved sales, customer satisfaction, inventory management or other financial, operational and organizational parameters. Knowing this would perhaps indicate a price for a loss of individuals' privacy and freedom. After all, one of the commonly cited examples of the usefulness of algorithmic decision-making was of Google being able to predict the location of a US flu epidemic, based on searches for flu remedies, faster than the Center for Disease Control (CDC). Yet, the story often remains untold, that they have been unable to repeat this success (<http://www.theguardian.com/technology/2014/mar/27/google-flu-trends-predicting-flu>). Thus, it is important that we conduct research that looks at the benefits for citizens of having a 'tailored' Internet, as against the costs of the benefits of living in an 'Internet bubble'. And finally there is a question about, 'what ethical guidelines might businesses usefully adopt to manage big/little data and produce the algorithms from these data'? For instance, Facebook has indicated that they have developed ethical policies for those who design algorithms, but such policies are not disclosed to the public. It is important that we research these issues so that we understand the ways in which businesses are using algorithms for discriminating so that we can enter a debate with business about associated ethical concerns (much as, for example, was the case in relation to the use of child labor in the past). Consider, for example, a monitoring system that profiles utility customers and sets different prices for gas and electricity, based on geographical areas and demand (another example of uninformed control). In this instance, maximizing societal welfare (cheap electricity for the majority) at the expense of minorities may well be unacceptable from an ethical standpoint (since those who end up paying more, are likely ironically also to be the very people who may be the most vulnerable and least able to pay). As a start in this process, we need research that sheds better light on the overall awareness of individuals in terms of how their data are being used by businesses and whether people are happy with this, especially as this exploits the more vulnerable in society. Thus, while big data analytics has the potential to shed light on important human and societal issues (Markus, 2015), this should not happen at the expense of the vulnerable.

In terms of the tradeoff between independency and dependency, we think that major societal issues are associated with the lack of opportunities for individuals to learn – and this poses issues from a knowledge creation and sharing perspective. As we pointed out earlier in this article, knowledge develops cumulatively and, according to the practice perspective (Feldman and Orlikowski, 2011; Schatzki et al., 2001; Sandberg and Tsoukas, 2011; Whittington, 2014) knowledge equates with practice. However, in the context of this tradeoff, it is the algorithm that gains knowledge about the minutiae of individuals – for instance, analyzing how humans drive a car, so that it can then operate as such, while humans may not gain a better understanding from this process. This poses relevant issues that involve both the private and work life of individuals. For example, 'will individuals lose their capacity to learn (even from mistakes)?' 'Will IT-assisted systems reach a point that they will impair an individual's problem solving skills and abilities in her/his everyday life?' This issue can be taken further if

we refer to the potential decreased ability of managers to make decisions on their own, due to the few opportunities to ‘practice’ decision-making processes. For instance, ‘Will an individual’s capacity to adapt to a new organization and a new job be compromised by increased control (made by algorithms, and that leads to living in a ‘bubble’, see our previous discussion of Pariser’s ‘experiment’), which makes people less likely to be flexible and accepting towards diversity?’ Also, there is the issue of the legitimacy of decisions that are based on algorithms and whether they will be accepted by those affected by the decision when the reasons for the decision are not actually known, even by those developing the algorithms. Thus, ‘will businesses face more claims of unfair discrimination in the future when people identify that they have been treated differently to others but when the justification for this is not understood and cannot be clearly articulated by anyone?’ “The computer says ‘no’” (e.g., <https://www.youtube.com/watch?v=AJQ3TM-p2QI>), may come to be an unacceptable justification for being discriminated against as we are increasingly confronted by this ‘rationale’.

The examination of the above social issues demands a multi-disciplinary approach that considers economic, legal, organizational, ethical, cultural and psychological consequences of the digitization of our everyday lives for different populations. In examining these issues, we would do well to remember that the ways new computing technologies (and the associated data) are used is not neutral in terms of the consequences for the human actors who leave digital traces that are then collected and analyzed. Corporations (and governments) have choices about how and what digital traces they collect and measure, and about the algorithms that they develop to make decisions based on this measurement, even if these decisions are increasingly distributed throughout a corporation rather than in the hands of the CIO (Nolan, 2012). These choices raise fundamental social questions as we have seen. As researchers we have an opportunity – and a responsibility (cf. Desouza et al., 2006, 2007) – to expose this empirically and theoretically and so promote an agenda of ‘responsible analytics’ that attempts to reduce the long-term negative social consequences of this new era concerned with the digitization of society.

In conclusion, this paper is an explicit call for action. We argue that researchers as well as practitioners should take these issues into serious consideration and articulate an interdisciplinary debate on how the datification of our everyday lives and the associated algorithmic decision-making (and the IoT) will affect society. This consequent research agenda requires a multi-disciplinary perspective. The issues are extremely relevant, strategic research topics. Whether we are interested in finding ways to increase business value or we are concerned with broader social issues of equality and democracy, they require immediate action. Strategic IS scholars are interested in “the way IT ‘delivers the goods’ by providing business AND social benefits” (Galliers et al., 2012, emphasis added). We would argue that minimizing social harm, even if to a minority, should be added to this agenda.

Acknowledgements

We would like to thank Sirrka Jarvenpaa and Yolande Chan for their extremely useful feedback given to previous versions of this manuscript.

References

- Abbas, R., Katina, M., Michael, M.G., 2014. The regulatory considerations and ethical dilemmas of location-based services (LBS): a literature review. *Inf. Technol. People* 27 (1), 2–20.
- Ball, K., 2002. Elements of surveillance: a new framework and future directions. *Inf., Commun., Soc.* 5 (4), 573–590.
- Ball, K., 2005. Organization, surveillance, and the body: towards a politics of resistance. *Organization* 12 (1), 89–108.
- Belanger, F., Crossler, R.E., 2011. Privacy in the digital age: a review of information privacy research in information systems. *MIS Q.* 35 (4), 1017–1041.
- Beldad, A., de Jong, M., Steehouder, M., 2011. I trust not therefore it must be risky: determinants of the perceived risk of disclosing personal data for e-government transactions. *Comp. Hum. Behav.* 27 (6), 2233–2242.
- Bentham, J., 1776. *A Fragment of Government*. London (Preface (2nd para)).
- Berente, N., Gal, U., Hansen, S., 2011. Ethical implications of social stratification in information systems research. *Inf. Syst. J.* 21 (4), 357–382.
- Boyd, D., Crawford, K., 2014. Critical questions for big data: provocations for a cultural, technological, and scholarly phenomenon. *Inf., Commun. Soc.* 15 (5), 662–679.
- Brown, J.S., Duguid, P., 1991. Organizational learning and communities of practice: towards a unified view of working, learning, and innovation. *Organ. Sci.* 2 (1), 40–57.
- Campbell, S.W., Park, Y.J., 2008. Social implications of mobile telephony: the rise of personal communication society. *Sociol. Compass* 2 (2), 371–387.
- Castells, M., Fernandez-Ardevol, M., Linchuan Qiu, J., Sey, A., 2009. A cross-cultural analysis of available evidence on the social uses of wireless communication technology. In: *Mobile Communication and Society: A Global Perspective*. The MIT Press, Cambridge, MA.
- Chan, Y.E., Culnan, M.J., Greenaway, K., Laden, G., Levin, T., 2005. Information privacy: management, marketplace, and legal challenges. *Commun. AIS* 16, 270–298.
- Chen, H., Chiang, R.H.L., Storey, V.C., 2012. Business intelligence and analytics: from big data to big impact. *MIS Q.* 36 (4), 1165–1188.
- Clark, C., Newell, S., 2013. Institutional work and complicit decoupling across the U.S. capital markets: the case of rating agencies. *Bus. Ethics Q.* 23 (1), 1–30.
- Cohen, W.M., Levinthal, D.A., 1990. Absorptive capacity: a new perspective of learning and innovation. *Adm. Sci. Q.* 35 (1), 128–152.
- Coll, S., 2014. Power, knowledge, and the subjects of privacy: understanding privacy as the ally of surveillance. *Surveillance* 17(10), 1250–1263.
- Culnan, M.J., Clark-Williams, C., 2009. How ethics can enhance organizational privacy: lessons from the choicepoint and TJX data breaches. *MIS Q.* 33 (4), 673–687.
- Dall’Alba, G., Sandberg, J., 2010. Learning through practice: professional and practice-based learning, 1, 104–119.
- Desouza, K.C., El Sawy, O.A., Galliers, R.D., Loebbecke, C., Watson, R.T., 2006. Beyond rigor and relevance towards responsibility and reverberation: information systems research that really matters. *Commun. Assoc. Inf. Syst.* 17 (16).
- Desouza, K.C., Ein-Dor, P., McCubbrey, D.J., Galliers, R.D., Myers, M.D., Watson, R.T., 2007. Social activism in information systems research: making the world a better place. *Commun. Assoc. Inf. Syst.* 19, 261–277.
- Dinev, T., Hart, P., Mullen, M., 2008. Internet privacy concerns and beliefs about government surveillance – an empirical investigation. *J. Strateg. Inf. Syst.* 7 (3), 214–233.
- Doyle, A., Rippert, R., Lyon, D., 2013. *The Global Growth of Camera Surveillance*. Routledge Publisher.

- Feldman, M.S., Orlikowski, W., 2011. Theorizing practice and practicing theory. *Organ. Sci.* 22 (4), 1240–1253.
- Galliers, R.D., Jarvenpaa, S.L., Chan, Y.E., Lyytinen, K.L., 2012. Strategic information systems: reflections and prospectives. *J. Strateg. Inf. Syst.* 21 (2), 85–90.
- Gartner, 2001. 3D Data Management: Controlling Data Volume, Velocity, and Variety, by D. Laney <<http://blogs.gartner.com/doug-laney/files/2012/01/ad949-3D-Data-Management-Controlling-Data-Volume-Velocity-and-Variety.pdf>> (accessed 06.02.15).
- Gartner, 2012. The Importance of 'Big Data': A Definition, by Beyer M.A., Laney D. <<http://www.gartner.com/document/2057415>>.
- Gerlach, J., Widjaja, T., Buxmann, P., 2015. Handle with care: How online social network providers' privacy policies impact users' information sharing behavior. *J. Strat. Inf. Syst.* 24(1), <http://dx.doi.org/10.1016/j.jsis.2014.09.001> (available online 12.10.14).
- Greenaway, K., Chan, Y.E., 2005. Theoretical explanations for firms' information privacy behaviors. *J. AIS* 6 (6), 171–198.
- Hasan, S., Subhani, M., 2011. Top management's snooping: is sneaking over employees' productivity and job commitment a wise approach. *Afr. J. Bus. Manage.* 6 (14), 5034–5043.
- Hedman, J., Srinivasan, N., Lindgren, R., 2013. Digital traces or information systems: sociomateriality made researchable. In: Proceedings of the 34th ICIS. Milan, Italy.
- Hodgkinson, G.P., Maule, A.J., Brown, N.J., Pearman, A.D., Glaister, K.W., 2002. Further reflections on the elimination of framing bias in strategic decision-making. *Strateg. Manag. J.* 23 (11), 1069–1073.
- Jarvenpaa, S.L., Lang, K.R., 2005. Managing the paradoxes of mobile technology. *Inf. Syst. Manage.* 22 (4), 7–23.
- Kane, G.C., 2014. Psychological stages of symbolic action in social media. In: Proceedings of the 35th ICIS December 14–17. Auckland, NZ.
- Kittur, A., Chi, E., Pendleton, B.A., Suh, B., Mytkowicz, T., 2007. Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeois. *WWW: World Wide Web* 1 (2), 19–28.
- Lyon, D., 2001. Surveillance and Society: Monitoring Everyday Life. McGraw-Hill International – Publisher.
- Lyon, D., 2003. Surveillance as Social Sorting: Privacy, Risk, and Digital Discrimination. Routledge, London.
- Lyon, D., 2014. Surveillance, snowden, and big data: capacities, consequences, critique. *Big Data Soc.* 1 (2). <http://dx.doi.org/10.1177/2053951714541861>.
- Lyytinen, K., Yoo, Y., 2002. Issues and challenges in ubiquitous computing. *Commun. ACM* 45 (12), 62–96.
- MacCrory, F., Westerman, G., Alhammedi, Y., Brynjolfsson, E., 2014. Racing with and against the machine: changes in occupational skill composition in an era of rapid technological advance. In: Proceedings of the International Conference of Information Systems (ICIS). Auckland, NZ.
- Markus, M.L., 2015. New games, new rules, new scoreboards: the potential consequences of big data. *J. Inf. Technol.* <http://dx.doi.org/10.1057/jit.2014.28> (available online 20.01.15).
- Mayer-Schonberger, V., Cukier, K., 2013. Big Data: A Revolution That Will Transform How We Live, Work, and Think. Houghton Mifflin Harcourt Publishing Company, New York, NY.
- McAfee, A., Brynjolfsson, E., 2012. Big data: the management revolution. *Harv. Bus. Rev.* 90 (10), 60–68.
- Michael, K., Michael, M.G., 2011. The social and behavioral implications of location-based services. *J. Loc. Based Serv.* 5 (3–4), 121–137.
- Milberg, S.J., Burke, S.J., Smith, H.J., Kallman, E.A., 1995. Values, personal information, privacy, and regulatory approaches. *Commun. ACM* 38 (12), 65–74.
- Mill, J.S., 1863. Utilitarianism, reprinted in 1906 in Chicago by the University of Chicago Press (first appearance in Fraser's Magazine, in 1861, then collected and reprinted as a single book in 1863).
- Mims, C., 2010. How iTunes Genius Really Works, MIT Technology Review <<http://www.technologyreview.com/view/419198/how-itunes-genius-really-works/>> (accessed February 02.02.15).
- Mingers, J., Walsham, G., 2010. Toward ethical information systems: the contribution of discourse ethics. *MIS Q.* 34 (4), 833–854.
- Miorandi, D., Sicari, S., De Pellegrini, F., Chlamatac, I., 2012. Internet of things: vision, applications, and research challenges. *Ad Hoc Netw.* 10, 1497–1516.
- Munford, M., 2014. Rule changes and big data revolutionise Caterham F1 chances. The Telegraph, Technology Section, 23 February 2014 <<http://www.telegraph.co.uk/technology/technology-topics/10654658/Rule-changes-and-big-data-revolutionise-Caterham-F1-chances.html>> (accessed 15.11.14).
- Negash, S., 2004. Business intelligence, communication of the association for information systems, 13 (article 15) <<http://aisel.aisnet.org/cais/vol13/iss1/15>>.
- Newell, S., Marabelli, M., 2014. The crowd and sensors era: opportunities and challenges for individuals, organizations, society, and researchers, in: Proceedings of 35th ICIS, December 14–17. Auckland, NZ.
- Nicolini, D., Gherardi, S., Yanow, D., 2003. Introduction: toward a practice-based view of knowing and learning in organizations. In: *Knowing in Organizations: A Practice-based Approach*. Barnes & Noble.
- Nolan, R., 2012. Ubiquitous IT: the case of the Boeing 787 and implications for strategy research. *J. Strateg. Inf. Syst.* 21 (2), 91–102.
- O'Reilly, C.A., Tushman, M.L., 2004. The ambidextrous organizations. *Harv. Bus. Rev.*, April, 1–10.
- Pariser, E., 2011. Beware Online 'Filter Bubbles' TED-2011 <http://www.ted.com/talks/eli_pariser_beware_online_filter_bubbles> (accessed 01.10.14).
- Petronio, S., 2002. Boundaries of Privacy: Dialectics of Disclosure. State University of New York Press, Albany, NY.
- Podsakoff, P., Todor, W., Skov, R., 1982. Effects of leader contingent and noncontingent reward and punishment behaviors on subordinate performance and satisfaction. *Acad. Manag. J.* 25 (4), 810–821.
- Power, D.J., 2002. Decisions Support Systems: Concepts and Resources for Managers. Quorum Books, Westport, CT.
- Rouibah, K., Ould-ali, S., 2002. Puzzle: a concept and prototype for linking business intelligence to business strategy. *J. Strateg. Inf. Syst.* 11 (2), 133–152.
- Sandberg, J., Tsoukas, H., 2011. Grasping the logic of practice: theorizing through practical rationality. *Acad. Manag. Rev.* 36 (2), 338–360.
- Schatzki, T.R., Knorr-Cetina, K., von Savigny, E., 2001. The Practice Turn in Contemporary Theory. Routledge, London.
- Schroeder, R., Cowls, J., 2014. Big data, ethics, and the social implications of knowledge production. *GeoJournal* <<https://dataethics.github.io/proceedings/BigDataEthicsandtheSocialImplicationsofKnowledgeProduction.pdf>> (accessed 24.01.15).
- Slade, S., Prinsloo, P., 2013. Learning analytics: ethical issues and dilemmas. *Am. Behav. Sci.* 57 (10), 1510–1529.
- Smith, H.J., 2002. Ethics and information systems: resolving the quandaries. *ACM SIGMIS Datab.* 33 (3), 8–22.
- Smith, H.J., Hasnas, J., 1999. Ethics and information systems: the corporate domain. *MIS Q.* 23 (1), 109–127.
- Smith, H.J., Dinev, T., Xu, H., 2011. Information privacy research: an interdisciplinary review. *MIS Q.* 35 (4), 989–1016.
- Staples, W.G., 2013. *Everyday Surveillance: Vigilance and Visibility in the Postmodern Life*. Rowman & Littlefield.
- Tene, O., Polonetsky, J., 2013. Big data for all: privacy and user control in the age of analytics, 11 Nw. J. Technol. Intelect. Prop., 239 <<http://scholarlycommons.law.northwestern.edu/njtip/vol11/iss5/1>> (accessed 02.11.14).
- Thomsen, E., 2003. BI's promised land. *Intell. Enterp.* 6 (4), 21–25.
- Tsoukas, H., Vladimirou, E., 2001. What is organizational knowledge? *J. Manage. Stud.* 38 (7), 973–993.
- Vallentyne, P., 1987. The teleological/deontological distinction. *J. Value Inq.* 21, 21–32.
- Velasquez, M., 2006. *Business Ethics: Concepts and Cases*. Upper Saddle River, NJ, Pearson.
- Whitman, G., 2011. The new paternalism: unraveling 'nudge'. *Econ. Affairs* 31, 4–5.
- Whittington, R., 2014. Information systems strategy and strategy-as-practice: a joint agenda. *J. Strateg. Inf. Syst.* 23 (1), 87–91.
- Wu, L., Brynjolfsson, E., 2009. The future of prediction: how google searches foreshadow housing prices and quantities. In: Proceedings of 31st ICIS, December 15–18. Phoenix, AZ, paper 147.
- Xi, F., Yang, L.T., Wang, L., Vinel, 2012. Internet of things. *Int. J. Commun. Syst.* 25, 1101–1102.
- Yoo, Y., 2010. Computing in everyday life: a call for research on experiential computing. *MIS Q.* 34 (2), 213–231.
- Zahra, S.A., George, G., 2002. Absorptive capacity: a review, reconceptualization, and extension. *Acad. Manag. Rev.* 27 (2), 185–203.
- Zuboff, S., 1984. In the Age of Smart Machine: The Future of Work and Power. Basic Books, NY.