The Management of Data Analytics Services: New Challenges and Future Directions

Citation: de Véricourt, Francis and Georgia Perakis. "Frontiers in service science: The management of data analytics services: New challenges and future directions." Service Science 12, 4 (November 2020): 121-129. © 2020 INFORMS

As Published: http://dx.doi.org/10.1287/serv.2020.0262

Publisher: Institute for Operations Research and the Management Sciences (INFORMS)

Persistent URL: https://hdl.handle.net/1721.1/130071

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

SERVICE SCIENCE
Vol. 00, No. 0, Xxxxx 0000, pp. 000–000
ISSN 2164-3962 | EISSN 2164-3970 | 00 | 0000 | 0001

INFORMS
DOI 10.1287/xxxx.0000.0000

The Management of Data Analytics Services: New Challenges and Future Directions

Francis de Véricourt
European School of Management a Technology, Berlin devericourt@esmt.org

Georgia Perakis
Sloan School of Management, MIT, Cambridge, MA 02139, georgiap@mit.edu

In this short paper, we discuss the impact of data analytics in services and delinieate future research directions for the field. After illustrating how data analytics are transforming different service sectors, we consider the provision of data analysis as a service in its own right. We discuss how the very nature of data and certain features of the machine learning method give rise to new issues and pitfalls for the management of these services, which delineates as many future research directions. We also discuss the co-production of services by humans and machines, and call for more research on responsible data analytics services to tackle some of the most pressing ethical issues in our societies.

Key words: Service, Big Data, Machine Learning
History:

1. Introduction

Data and analytics have become essential resources for the design and management of services. The growing availability of data coupled with advances in analytics and machine learning in particular, have enabled service organizations to improve their performance and create better value for their services to consumers. Yet, the wealth and very nature of data poses new challenges for the management of service systems. In this short paper, we present what we think are the most novel and important questions that the rise of data analytics pose in service science.

The field often referred to as data analytics concerns the whole process by which starting with using data one creates models leading to decisions that create value to different organizations. Data analytics have allowed researchers and practitioners alike in recent years to generate new insights from different services based on actual data. In the first section of this paper, we illustrate how these approaches are fundamentally transforming several service industries, including healthcare services, online services and transportation services to cite a few. We believe that the growing availability of data will continue to provide ample opportunities for researchers to develop and apply approaches ranging from machine learning and data mining to econometrics and field experiments.

Yet, the field of data analytics does not solely concern itself with analyzing data to gain knowledge and produce insights for specific service systems. The production of data analysis constitutes a service in itself, with specific workflows and managerial challenges.

We explore these issues in the second section of the paper. In particular, we argue that the so-called unreasonable effectiveness of data (Halevy et al. 2009) makes the acquisition of large amount of data the most important step (as opposed, for instance, to the improvement of ML algorithms) in the development of data analytics services. This gives rise to new but under-studied business models and approaches that facilitate the acquisition of data, such as the AI flywheel effect (Sarikaya 2019) or crowdsourcing for supervised learning. We also discuss how the nature of data and certain algorithms may give rise to novel incentive problems that could hamper the adoption of data analytics services.

We further believe that the future of service science lies in the co-production of services by humans and machine. How to optimally integrate machine-based recommendations to the decision-making process of a human service provider remains largely an open question. We explore this aspect in the fourth section of the paper.

We conclude the paper by raising some ethical issues associated with the collection of data and usage of ML algorithms to deliver services.

2. Applications of ML and Big Data in Services

In this section, we build and expand on Misic and Perakis (2019) to briefly show how the availability of a large amount of data and the application of advanced analytics techniques are transforming different service sectors. We illustrate this point for data rich service industries, including healthcare, supply chain and retail services. These applications of data analytics approaches have not only improved these sectors and revealed new insights but also enabled the development of new techniques and methodologies,

Healthcare services is an important sector where the availability of data such as electronic health records in the recent years has allowed analytics to have an impact. Nowadays the data from patients available to healthcare providers, allows them to improve patient services. Papers in the field of service and management science have already been addressing many issues in the healthcare space in the recent years. Healthcare policy is one of the areas in this space. Examples include the work of Aswani et al. (2018) that studies the Medicare Shared Savings Program (MSSP). Another important application relates to kidney allocations (with live as well as deceased organ donors). For deceased organ donors, the topic has been examined in several papers such as Bertsimas et al. (2013) as well as by Baris et al. (2017) and the references therein. For live donors, the work of Ashlagi and Roth (2020) and the references therein discuss this topic in detail. Both kidney exchange programs (whether it is for live or deceased donors) have become topics of national visibility where our community is having an impact.

There is also a growing literature that examines analytics in services at the hospital-level. For example, Rath et al. (2017), Rath and Rajaram (2018), Ang et al. (2015) and Zenteno et al. (2016) are examples of papers that develop optimization and predictive methods to study the problem of staffing, scheduling of surgeries and ED wait times at different hospitals in the United States. The methods used in these papers, are not only optimization methods but also predictive methods. Some of the approaches proposed in the literature to predict ED wait times include rolling average estimators, metrics based on fluid models and quantile regression as well as the LASSO method of statistical learning (Tibshirani 1996) with predictor variables that are derived from a generalized fluid model of the queue. Papers such as Ang et al. (2015) show, using real hospital data, that such methods outperform other approaches to predicting ED wait times, leading to reductions in mean squared error of up to 30% relative to the rolling average method that is commonly used in hospitals today.

Finally, the increasing availability of electronic health record data has led to the application of analytics methods to patient-level problems. Papers, such as Bastani and Bayati (2015) and Bertsimas et al. (2017) consider personalized approaches taking into account patient features (e.g., patient's gender, age, whether the patient has other pre-existing conditions, whether the patient is taking other drugs). There are many other exciting examples where analytics are impacting healthcare services. The papers by Mahmoudi et al. (2017) and Farias and Li (2017) are such examples. These papers devise methods for dealing with missing data in order to design testing for cancer such as for example blood tests for cancer.

Apart from healthcare, another important sector where analytics is having an impact includes supply chain management. In fact, many important problems in supply chain management are being re-examined under the lens of analytics. Predicting demand and optimizing the supply chain with this information is a critical problem. The newsvendor problem in inventory management has been a cornerstone problem. When one has data on past demand, together with other features that may be predictive of demand (these features could include weather forecasts and economic indicators, as an example), one can consider building a demand distribution that is feature-dependent, and then find the optimal order quantity for the distribution corresponding to a given realization of the features. Ban and Rudin (2019) were perhaps the first in the field of service and management science to consider a data-driven approach that incorporates feature-dependent demand in the newsvendor problem for inventory management.

Connected to supply chain lies the area of retail services. For example, how to manage the omnichannel operations of a large online retailer is an important problem where analytics are nowadays playing an important role. "Omnichannel" refers to the integration of an e-commerce channel and a network of brick-and-mortar stores. Acimovic and Graves (2014) consider the problem of how to optimally fulfill each customer's order to minimize average outbound shipping costs. Using data from a large online retailer, they propose a method that reduces outbound shipping costs on the order of 1%. Harsha et al. (2019) enhance this problem by considering the omnichannel price optimization problem, that is, pricing in the presence of cross-channel interactions in demand and supply, with exogenous cross-channel fulfillment. They implemented their approach commercially as part of the IBM Commerce Markdown Price Solution. In addition, they did a causal analysis based on data from a pilot implementation at a large U.S. retailer to suggest their approach could yield an estimated 13.7% increase in clearance period revenue.

Related to retail services, analytics is also having an impact on how firms can sell their products to customers. A first order problem connects with distribution. An example includes Avrahami et al. (2014) who improved how Yedioth in Israel distributes print magazines and newspapers using real-time information on the newspaper sales at the different retail outlets to enable pooling of inventory in the distribution network. The methods developed in this paper generated substantial cost savings at Yedioth due to both a reduction in magazine production levels and a reduction in the return levels.

Connected to this problem the service and management science community has also developed analytics approaches for assortment optimization. The goal in this problem is to optimize which products to offer to which customers, in order to maximize the revenue earned when customers make purchase decisions from the offered products. A major complicating factor in this problem is the presence of substitution behavior, that is, customers who may arrive with a particular product in mind, but purchase a different product if that first product is not available. Numerous discrete choice models, such as the multinomial logit (MNL) model, have been proposed for modeling how customers make purchase decisions (Train 2009). Both the estimation and the optimization are important problems in this area. Among the many papers in this space, examples include the paper of Farias et al. (2013) who proposed a nonparametric model for representing customer choice behavior, by modeling the customer population as a probability distribution over rankings of the products. On the other hand, the paper of Blanchet et al. (2016) introduced a Markov chain-based choice model where substitution from one product to another is modeled as a state transition in a Markov chain.

Subsequently, the community has also recently considered how analytics can be helpful with pricing. For example, the paper of Ferreira et al. (2015), were among the first to study the pricing problem of Rue La La's weekly sales events (a large online retailer that does flash sales). The paper combined machine learning and optimization to predict demand and prescribe optimal pricing recommendations. Related to the topic of pricing is also promotion planning. Price promotion planning is concerned with when and how deeply to promote each item over a time horizon. Cohen et al. (2017a), Cohen et al. (2017b) as well as Cohen-Hillel et al. (2019a), Cohen-Hillel et al. (2019b) consider how analytics can help for price promotion planning first for a single item and subsequently, for multiple items. These are examples of papers that develop methods for optimally deciding price promotions by combining machine learning and optimization.

The increasing availability of data has also led to many important methodological advances. Several researchers have developed new related methods such as regularized regression methods as well as tree ensemble-based methods among many others (see for example, Tibshirani 1996, Breiman et al. 1984, Perakis et al. 2019 among others). The development of such methods has in fact seen unprecedented growth. In addition, the rise of analytics has also allowed researchers to develop further optimization methods such as for example, methods for integer optimization (see for example, Vielma 2015) and robust optimization (see for example, Bertsimas et al. 2011) among many others.

These advances are leading researchers and practitioners alike to answer many important questions in several different service sectors such as how to streamline and unify online and offline services, how to price and promote products in retail services and how to improve healthcare delivery services among many others. The review paper by Misic and Perakis (2019) provides more details in terms of methodology and applications in this space.

Despite the huge growth in the recent years, we believe that the application of analytics in the service science field is still at its infancy. As a result, in what follows we will take a different approach. In particular, we will discuss some of our thoughts and vision in terms of different issues associated with analytics in the area of services. We will discuss what could be some of the underlying associated issues such as pitfalls related to data and very importantly the role of humans in data analytics. Overall our goal in what follows, is to discuss how analytics can be viewed as a new kind of service on its own.

3. Data analytics as a service

Producing the right data and the right analysis through the right methodology can be viewed as a service by itself. This raises new and specific managerial challenges due to the very nature of data as well certain aspects of analytics techniques and machine learning in particular. In this section, we first briefly describe the process of creating and delivering a data analytics based process. We then identify what we see as the main novel challenge associated with each step of the process.

3.1. Data analysis, processes and workflows

Practitioners and academics have proposed several workflows (see for instance Bruha and Famili 2000, Xin et al. 2018b, GSMA 2018), which consist of different steps that are needed to successfully carry out a data science project. These processes commonly consist of iterating between three main components: 1. Data acquisition and pre-processing, 2. Analysis and learning and 3. Post-processing and deployment. While the AI community has developed tools to optimize many aspects of these processes (see Xin et al. 2018c, Sparks et al. 2017, for instance), the delivery of data science products and insights remains driven by human analysts.

As such, the production and delivery of data analysis is subject to managerial challenges that permeates any service systems, such as problems of resource allocation, outsourcing decisions or incentives issues, to cite a few. The management and economic literature has developed successful frameworks to address these issues in other service contexts. Yet, the very nature of data and the specificities of a data analysis workflow often preclude the application of these frameworks.

In fact, most challenges in a data analysis workflow are related to data and deployment issues, as opposed to the most technical aspects of the process. Based on a survey from the Machine Learning literature, Xin et al. (2018b) present descriptive statistics on how data scientists allocate their effort in a data analysis process. The study reveals that for ML based projects, data scientists allocate most of their effort to data preprocessing (step 1.) and exert more effort in post-processing activities (step 3.) than in the learning tasks (step 2.). However, much remains to be done to address the specific challenges associated with these steps, which we describe next.

3.2. The unreasonable effectiveness of data

Among the different components of the data science workflow, data acquisition and pre-processing (step 1.) emerge as the most critical steps. One reason for this is the so-called “unreasonable effectiveness of data”, an expression first coined by Halevy et al. (2009). This term asserts that simple models trained on large datasets perform better than elaborated models that try to discover general rules. For instance, Banko and Brill (2001) show that simple models and a lot of data trump more elaborate models based on less data for a language disambiguation task. Similarly, Sun et al. (2017) study the impact of the size of datasets on the accuracy of deep learning in vision tasks. The authors conclude by encouraging their community to exert more efforts in building larger datasets. Some data scientists also argue that collecting data is more important than improving algorithms to develop machines that can help cure cancer (Warren 2016).

The unreasonable effectiveness of data does not only rely on the sheer volume of data, but also on its variety. Indeed, combining heterogeneous data sources is another key driver of the recent successes of data analytics and machine learning. In health care, for instance, Mayer-Schönberger and Ingelsson (2018) suggest that a variety of data of limited quality may outperform small homogenous datasets of high quality. Some of the most insightful results from the analysis of Big Data in medicine actually comes from the combination of heterogenous large datasets (see for example Frei et al. 2011 who merge three large heterogenous datasets to study the link between mobile phone usage and brain tumors). Similarly, scholars and practitioners also argue that the value of big data in business often comes from the combination of internal and external large datasets (Thomas and Leiponen 2016, Gopalkrishnan et al. 2013, Davenport 2013).

This data effectiveness also enables a virtuous cycle (sometimes referred to as "the AI flywheel effect" - see, for instance, Sarikaya 2019, Arslanian and Fischer 2019), by which more heterogenous data improve the performance of machine learning algorithms, which fosters the diffusion of the service and, in turn, generate more data. For instance, Amazon's Alexa detects whether a customer may rephrase her request as a signal that the first request should have elicited the same response, and thus improve the accuracy of the recommendation algorithm (Sarikaya 2019).

Despite the importance of the unreasonable effectiveness of data, the managerial literature has barely studied how this effect affects the development and delivery of data analysis services. A notable exception is Gurkan and de Véricourt (2020), which study how data effectiveness and "the AI flywheel effect" interact with the design and outsourcing of an AI based service. We believe, however, that exploring this effect further is a very fruitful research direction for the management of data analytics services .

3.3. The role of users in producing machine learning models

Machines are sometimes automatically trained by solely relying on data, without resorting to external human expertise other than data scientists. This approach has shown great success for certain tasks, such as image and speech recognition, translations or recommendation systems. Yet, automatically training an algorithm has also reached its limit in many service applications. This is especially the case when large enough training sets are missing, or when the service task pertains to low probability events such as diagnosing a rare disease (Holzinger 2016).

For situations such as these, a so-called human-in-the-loop approach has been advanced (Amershi et al. 2014). The approach consists of developing a tight coupling between the user and the machine, by which human experts directly guide the learning algorithm (step 2.). The problem becomes then to minimize the time and costs associated with the interaction cycles between humans and machine, as well as having users with no machine-learning expertise efficiently influences the learning algorithms.

The development of tools and techniques to mitigate these issues is the topic of current research in the machine learning community (Wilson and Daugherty 2018, Xin et al. 2018a, for example). Yet, letting humans directly guide a learning algorithm can introduce new types of biases that a pure technical solution cannot fully address. For instance, Thomaz and Breazeal (2008) and Knox and Stone (2012) find that when guiding the learning process of a machine, people tend to overemphasize correct decisions (by providing more positive rewards) than wrong ones (by giving fewer negative rewards). Furthermore, integrating users with no machine-learning expertise directly into the training process requires rendering algorithms understandable to the users. This, however, may incentivize the data scientists to develop less accurate but more explainable algorithms for a specific audience, which Herman (2017) identifies as an implicit human cognitive bias.

How to efficiently manage the dynamic process of integrating domain expertise and users? Knowledge into the training of an algorithm remains an open question, that the management science literature has not addressed thus far. In this respect, one promising research direction concerns the use of crowdsourcing for supervised learning, in which the crowds label inputs, or evaluate the outputs of an algorithm (see Good and Su 2013, for instance). Crowdfunding has recently received a lot of attention in the finance and economics literature (Strausz 2017) in the context of project funding. How to efficiently design and manage crowdsourcing for machine learning remains poorly understood.

3.4. Machine learning interpretability and the problem of adoption

In addition, the deployment of these machine learning based products (Step 3) is also limited by the machine's inability to explain its recommendations and predictions. These machines are often black-boxes (Ribeiro et al. 2016), not only to the end users, but also to the very data scientists who created them. Data scientists certainly know how to develop datasets and training algorithms, but the final outputs of the system remain an emerging phenomenon. In other words, while the statistical and optimization principles for training a machine is understood, the underlying explanatory structure which eventually determines the machine's outputs is not (Holzinger et al. 2017).

This lack of interpretability has managerial consequences and, in particular, may hinder the adoption of machine learning models. Indeed, end users need to be confident that an algorithm behaves reasonably well before adopting it. The algorithm's evaluation on test data, as is practiced by the machine learning community, does not fully reflect the actual performance in the real world. To make this assessment, users also need to leverage their domain knowledge (see Ribeiro et al. 2016, for instance), which requires understanding the rationale behind a prediction. However, the difficulty to explain a machine's output obfuscates this assessment and may reduce the user's trust in the system (Kim 2015, Ridgeway et al. 1998).

An approach to mitigate this issue consists in relying on more interpretable models such as regression trees or Bayesian neural networks (see Rani et al. 2006, for instance). These methods arguably provide better cues about the underlying reasons for the predictions of an algorithm, enabling users to make an assessment based on their expertise. Yet, this approach restricts the sets of methods available to the data scientist, and thus may result in lower levels of performance. This yields a tradeoff between performance and interpretability, by which building trust promotes the deployment of an algorithm at the cost of sub-optimal performance. How this tradeoff affects managerial decisions concerning the design and deployment of machine-learning based services remains largely an open question.

4. The co-production of services by machines and humans

The increasing adoption of smart machines and data-base technologies have questioned the role of humans in the delivery of services. The introduction of these new technologies has even met intense resistance by human operators in certain service sectors (Katz 2017). While new technologies can substitute for labor, a wealth of evidences suggests that they also complement human skills (see for instance Felten et al. 2019 and references therein). This is because machines often substitute for only a subset of the different tasks required to perform an activity (Autor 2015, European Economic and Social Committee 2017).

For instance, AI-driven tools can generate basic memos for lawyers, who then add deeper analysis and improve the language. The same is true for human translators, who use translation machines to produce a first draft, that human translators improve to adjust the tone to the context (Katz 2017). The potential for synergies such as these has led The Defense Advanced Research Projects Agency to invest $2 Billion with the goal of improving the collaborations between humans and machines (DARPA 2018).

This complementarity stems from the fact that, despite their impressive achievements, machine learning technologies sill lack the impressive cognitive flexibility of humans (Diamond 2013, Laureiro-Martínez and Brusoni 2018). This flexibility enables integrating various information from a context and other background knowledge, which go beyond the data on which machines are trained (Marcus 2018). In the financial service industry, for instance, mutual funds managers better detect aggregate shock in the market than machine-based quantitative funds do (Abis 2017). Similarly, humans make hiring recommendations based on information that ML algorithms cannot access (Hoffman et al. 2017). More generally, the purpose of many ML applications is to provide useful information to human service providers, who makes decisions that do not solely rely on the algorithm's outputs (Lipton 2016).

To the extent that data-based technologies can improve part of a service delivery process, the coproduction of decisions by humans and machines should result in overall higher performance

(Mims 2017). Yet, not all performance measures systematically improve. For example, Stoffel et al. (2018) found that when radiologists took into account the deep-learning analysis of ultrasound images, the diagnoses of breast tumors significantly improved. This improvement, however, mainly stemmed from a radical decrease in false negative, while false positives rate did not significantly change.

Psychology research has long explored the question of how people integrate machine-based recommendations into their decision-making process. In particular, it has been suggested that humans are hesitant to rely on these recommendations (Meehl 1954, Dawes 1979), a notion that Dietvorst et al. (2015) recently “Algorithm Aversion” (and that the business community has started to explore - see Frick 2015, Harrell 2016). By contrast, Logg et al. (2019) demonstrate that people can appreciate the value of algorithmic recommendations. Several studies have further explored under which conditions and to which extent humans account for machine-based judgments to make a decision (see for instance Dietvorst et al. 2016, Yeomans et al. 2017, for recent work on the topic).

We see this coproduction of services by machines and humans, and its implications for quality and costs as the most pressing research area for the field service sciences.

5. Responsible Service Analytics

In closing this paper, we want to emphasize the importance of ethics for data analytics. In the early days of the field, ethical concerns related to data as well as the methods used, had been at times either ignored or at the very best considered “after the fact”. With the increasing success of machine learning at making predictions many problematic - if not shocking- cases have since emerged. For instance, Google photos’ image labeling algorithm has labeled dark-skinned people as gorillas (Pachal 2015). Certain automatic soap dispensers also do not recognize black hands (Mills 2017) and Apple’s Iphone X face recognition algorithm could not differentiate between two different Asian people (Papenfuss 2017). More generally, sociologist Benjamin makes the point that algorithms and data reenforce racial inequity in the US, and proposes a framework grounded in the science and technology studies to uncover these less apparent discriminations (Benjamin 2019). These failures have increasingly pressured the field of service science to address issues of fairness and discrimination.

Biases such as these do not come so much from the training algorithms and statistical models (although issues exist there as well - see Obermeyer et al. 2019 for instance), but are typically due to the dataset on which these methods are used. This data often reflects and directly encodes some of the social, gender and racial biases of societies. For instance, the US criminal justice increasingly relies on artificial intelligence in its decision processes (Isaac 2017). These algorithms, however, are trained on criminal datasets that show patterns of discriminatory policing (Angwin et al. 2017, Richardson et al. 2019). Even when the data does not directly encode the information that may lead to discrimination, the data may be rich enough to enable algorithms to reveal certain aspects of a customer identity. For instance, ML algorithm are better than humans at predicting the sexual orientation of individuals through facial recognition (Wang and Kosinski 2018). This, in turn, relates to consumer privacy and data protection issues, which have been the topic of intense political debates and constitutes an important research direction.

Data analytics services, therefore, run the risk of perpetuating and even reinforcing existing societal biases at the expense of certain communities. By the same token, however, we believe the management of data analytics services provides a unique opportunity to tackle these societal issues. It is indeed much easier to debias a dataset than directly change systematic discriminations embedded in societies. A promising example of this is Google Translate. Many languages such as Hungarian or Chinese are gender neutral. Yet Google Translate exhibited a strong bias towards male when translating certain sentences in English (such as "he is a doctor", Prates et al. 2020). Google has since successfully updated its system to provide both feminine and masculine translations (Johnson 2017).

These issues and the potential of the field to reduce discriminations in societies call for robust research on *responsible analytics services*. In this direction, we believe that the main problem lies in incentivizing analytics service providers to properly debias their database and algorithms. More generally, as Ronald Coase mentioned (Nobel Prize laureate) "if you torture the data long enough, it will confess". And because the results of ML algorithms are typically difficult to interpret, verifying to which extent the data has been "tortured" can be hard. How to properly incentive providers of data analytics services and verify their work remain largely an open question.

6. Conclusion

In summary, this paper only scraped the surface regarding the role of data analytics in service science. We believe, however, that framing the provision of data analytics as a service in its own right, as we do in this paper, gives rise to new and fruitful research directions for the field of service science. Focusing on this approach is also crucial for the field to have a meaningful societal impact. Indeed, managing the delivery of analytical services responsibly is one of the key to address the discriminations and fairness issues that permeate our societies.

References

Abis S (2017) Man vs. machine: Quantitative and discretionary equity management, working paper, Columbia University.

Acimovic J, Graves S (2014) Making better fulfillment decisions on the fly in an online retail environment. *Manufacturing & Service Operations Management* 17(1):34–51.

Amershi S, Cakmak M, Knox WB, Kulesza T (2014) Power to the people: The role of humans in interactive machine learning. *AI Magazine* 35(4):105–120.

Ang E, Kwasnick S, Bayati M, Plambeck EL, Aratow M (2015) Accurate emergency department wait time prediction. *Manufacturing & Service Operations Management* 18(1):141–156.

Angwin J, Larson J, Mattu S, Kirchner L (2017) Machine bias: There's software used across the country to predict future criminals and it's biased against blacks. URL `www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing`, accessed: 2020-08-10.

Arslanian H, Fischer F (2019) Future Trends in Artificial Intelligence. Arslanian H, Fischer F, eds., *The Future of Finance: The Impact of FinTech, AI, and Crypto on Financial Services*, 231–247 (Cham: Springer International Publishing).

Ashlagi I, Roth A (2020) Kidney exchange: an operations perspective. *Working paper* .

Aswani A, Shen ZJM, Siddiq A (2018) Data-driven incentive design in the medicare shared savings program. *Operations Research* Forthcoming; available at SSRN 2808211.

Autor DH (2015) Why are there still so many jobs? the history and future of workplace automation. *The Journal of Economic Perspectives* 29(3):3–30.

Avrahami A, Herer YT, Levi R (2014) Matching supply and demand: delayed two-phase distribution at yedioth group – models, algorithms, and information technology. *Interfaces* 44(5):445–460.

Ban GY, Rudin C (2019) The big data newsvendor: Practical insights from machine learning. *Operations Research* 67(1):90–108.

Banko M, Brill E (2001) Scaling to very very large corpora for natural language disambiguation. *Proceedings of the 39th annual meeting on association for computational linguistics*, 26–33 (Association for Computational Linguistics).

Baris A, Freidewald J, Randa AC (2017) Organ transplantation. *Handbook of Healthcare Analytics: Theoretical Minimum for Conducting 21st Century Research on Healthcare Operations. John Wiley & Sons* .

Bastani H, Bayati M (2015) Online decision-making with high-dimensional covariates. *Working paper* Available at SSRN 2661896.

Benjamin R (2019) *Race After Technology: Abolitionist Tools for the New Jim Code* (New Jersey, USA: John Wiley & Sons).

Bertsimas D, Brown DB, Caramanis C (2011) Theory and applications of robust optimization. *SIAM review* 53(3):464–501.

Bertsimas D, Farias VF, Trichakis N (2013) Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research* 61(1):73–87.

Bertsimas D, Kallus N, Weinstein AM, Zhuo YD (2017) Personalized diabetes management using electronic medical records. *Diabetes Care* 40(2):210–217.

Blanchet J, Gallego G, Goyal V (2016) A Markov Chain Approximation to Choice Modeling. *Operations Research* 64(4):886–905.

Breiman L, Friedman J, Stone CJ, Olshen RA (1984) *Classification and regression trees* (CRC press).

Bruha I, Famili F (2000) Postprocessing in machine learning and data mining. *SIGKDD Explorations* 2(2):110–114.

Cohen MC, Kalas J, Perakis G (2017a) Optimizing promotions for multiple items in supermarkets. *Working paper* Available at SSRN 3061451.

Cohen MC, Leung NHZ, Panchamgam K, Perakis G, Smith A (2017b) The impact of linear optimization on promotion planning. *Operations Research* 65(2):446–468.

Cohen-Hillel T, Panchamgam K, Perakis G (2019a) Bounded memory peak-end models can be surprisingly good. *Working paper* Available at SSRN 3360643.

Cohen-Hillel T, Panchamgam K, Perakis G (2019b) High-low promotion policies for peak end demand models. *Working paper* Available at SSRN 3360680.

DARPA M (2018) Darpa announces $2 billion campaign to develop next wave of ai technologies. URL `www.darpa.mil/news-events/2018-09-07`, accessed: 2019-09-20.

Davenport TH (2013) Analytics 3.0. *Harvard Business Review* 91(12):64–+.

Dawes RM (1979) The robust beauty of improper linear models in decision making. *American psychologist* 34(7):571.

Diamond A (2013) Executive functions. *Annual review of psychology* 64:135–168.

Dietvorst BJ, Simmons JP, Massey C (2015) Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General* 144(1):114.

Dietvorst BJ, Simmons JP, Massey C (2016) Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them. *Management Science* 64(3):1155–1170, URL `http://dx.doi.org/10.1287/mnsc.2016.2643`.

European Economic and Social Committee (2017) Impact of digitalisation and the on-demand economy on labour markets and the consequences for employment and industrial relations. Technical report, European Economic and Social Committee.

Farias VF, Jagabathula S, Shah D (2013) A nonparametric approach to modeling choice with limited data. *Management Science* 59(2):305–322.

Farias VF, Li AA (2017) Learning preferences with side information. *Management Science* Forthcoming; available at `http://www.mit.edu/~aali/files/Tensor.pdf`.

Ferreira KJ, Lee BHA, Simchi-Levi D (2015) Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management* 18(1):69–88.

Frick W (2015) Here's why people trust human judgment over algorithms. *Harvard Business Review* .

Good BM, Su AI (2013) Crowdsourcing for bioinformatics. *Bioinformatics* 29(16):1925–1933.

Gopalkrishnan V, Steier D, Lewis H, Guszcza J, Lucker J (2013) Big data 2.0: New business strategies from big data. *Deloitte Review* 54–69.

GSMA (2018) The data value chain. URL `www.gsma.com/publicpolicy/wp-content/uploads/2018/06/GSMA_Data_Value_Chain_June_2018.pdf`, accessed: 2019-09-20.

Gurkan H, de Véricourt F (2020) Contracting, pricing, and data collection under the ai flywheel effect. *ESMT Berlin Working Paper* .

Halevy A, Norvig P, Pereira F (2009) The Unreasonable Effectiveness of Data. *IEEE Intelligent Systems* 24(2):8–12, ISSN 1541-1672, URL `http://dx.doi.org/10.1109/MIS.2009.36`.

Harrell E (2016) Managers shouldn't fear algorithm-based decision making. URL `hbr.org/2016/09/managers-shouldnt-fear-algorithm-based-decision-making`, accessed: 2020-09-05.

Harsha P, Subramanian S, Uichanco J (2019) Dynamic pricing of omnichannel inventories. *Manufacturing & Service Operations Management* 21:47–65.

Herman B (2017) The promise and peril of human evaluation for model interpretability. *arXiv preprint arXiv:1711.07414* .

Hoffman M, Kahn LB, Li D (2017) Discretion in hiring. *The Quarterly Journal of Economics* 133(2):765–800.

Holzinger A (2016) Interactive machine learning for health informatics: when do we need the human-in-the-loop? *Brain Informatics* 3(2):119–131.

Holzinger A, Biemann C, Pattichis CS, Kell DB (2017) What do we need to build explainable ai systems for the medical domain? *arXiv preprint arXiv:1712.09923* .

Isaac WS (2017) Hope, hype, and fear: the promise and potential pitfalls of artificial intelligence in criminal justice. *Ohio St. J. Crim. L.* 15:543.

Johnson M (2017) Providing gender-specific translations in google translate. URL `ai.googleblog.com/2018/12/providing-gender-specific-translations.html`, accessed: 2020-08-10.

Katz M (2017) Welcome to the era of the ai coworkera. URL `www.wired.com/story/welcome-to-the-era-of-the-ai-coworker/`, accessed: 2019-09-20.

Kim B (2015) *Interactive and interpretable machine learning models for human machine collaboration.* Ph.D. thesis, Massachusetts Institute of Technology.

Knox WB, Stone P (2012) Reinforcement learning from human reward: Discounting in episodic tasks. *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication*, 878–885.

Laureiro-Martínez D, Brusoni S (2018) Cognitive flexibility and adaptive decision-making: Evidence from a laboratory study of expert decision makers. *Strategic Management Journal* 39(4):1031–1058.

Lipton ZC (2016) The mythos of model interpretability. *arXiv preprint arXiv:1606.03490* .

Logg JM, Minson JA, Moore DA (2019) Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes* 151:90–103.

Mahmoudi M, Caracciolo G, Safavi-Sohi R, Poustchi H, Li AA, Vasighi M, Chiozzi RZ, Capriotti AL, Lagana A, Hajipou M, Di Domenico M, Di Carlo A, Caputo D, Aghaverdi H, Xu X, Papi M, Palmieri V, Santoni A, Shi J, Jon S, Palchetti S, Pozzi D, Malekzadeh R, Suslick KS, Farias VF, Farokhzad OC (2017) Multi-nanoparticle protein corona characterization of human plasma and machine learning enable accurate identification and discrimination of cancers at early stages. *Working paper* .

Marcus G (2018) Deep learning: A critical appraisal. *arXiv preprint arXiv:1801.00631* .

Mayer-Schönberger V, Ingelsson E (2018) Big Data and medicine: A big deal? *Journal of Internal Medicine* 283(5), ISSN 0954-6820, 1365-2796, URL `http://dx.doi.org/10.1111/joim.12721`.

Meehl PE (1954) Clinical versus statistical prediction: A theoretical analysis and a review of the evidence, university of Minnesota Press.

Mills K (2017) 'Racist' soap dispenser refuses to help dark-skinned man wash his hands - but twitter blames 'technology'. URL `www.mirror.co.uk/news/world-news/racist-soap-dispenser-refuses-help-11004385`, accessed: 2020-08-10.

Mims C (2017) Without Humans, Artificial Intelligence Is Still Pretty Stupid. *Wall Street Journal* ISSN 0099-9660.

Misic V, Perakis G (2019) Data analytics in operations management: A review. *MSOM* .

Obermeyer Z, Powers B, Vogeli C, Mullainathan S (2019) Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366(6464):447–453.

Pachal P (2015) Google photos identified two black people as 'gorillas'. URL `https://mashable.com/2015/07/01/google-photos-black-people-gorillas/`, accessed: 2020-08-10.

Papenfuss M (2017) Woman in china says colleague's face was able to un- lock her iphone x. URL `www.huffpostbrasil.com/entry/iphone-face-recognition-double_us_5a332cbce4b0ff955ad17d50`, accessed: 2020-08-10.

Perakis G, Singhvi D, Skali-Lami O (2019) Xstrees: A tree sampling framework for interpretable tree ensembles. *https://papers.ssrn.com/sol3/3505431* .

Prates MO, Avelar PH, Lamb LC (2020) Assessing gender bias in machine translation: a case study with google translate. *Neural Computing and Applications* 32:6363–638.

Rani P, Liu C, Sarkar N, Vanman E (2006) An empirical study of machine learning techniques for affect recognition in human–robot interaction. *Pattern Analysis and Applications* 9(1):58–69.

Rath S, Rajaram K (2018) Staff planning for hospitals with cost estimation and optimization. *Working paper* Available at SSRN 3241063.

Rath S, Rajaram K, Mahajan A (2017) Integrated anesthesiologist and room scheduling for surgeries: Methodology and application. *Operations Research* 65(6):1460–1478.

Ribeiro MT, Singh S, Guestrin C (2016) Why should i trust you?: Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1135–1144 (ACM).

Richardson R, Schultz J, K C (2019) Dirty data, bad predictions: How civil rights violations impact police data, predictive policing systems, and justice. *New York University Law Review Online* Forthcoming.

Ridgeway G, Madigan D, Richardson T, O'Kane J (1998) Interpretable boosted naïve bayes classification. *Proceedings of the fourth international conference on knowledge discovery and data mining*, 101–104.

Sarikaya R (2019) How alexa learns. URL `blogs.scientificamerican.com/observations/how-alexa-learns/`, accessed: 2019-09-20.

Sparks ER, Venkataraman S, Kaftan T, Franklin MJ, Recht B (2017) Keystoneml: Optimizing pipelines for large-scale advanced analytics. *2017 IEEE 33rd international conference on data engineering (ICDE)*, 535–546 (IEEE).

Stoffel E, Becker AS, Wurnig MC, Marcon M, Ghafoor S, Berger N, Boss A (2018) Distinction between phyllodes tumor and fibroadenoma in breast ultrasound using deep learning image analysis. *European Journal of Radiology Open* 5:165–170, URL `http://dx.doi.org/10.1016/j.ejro.2018.09.002`.

Strausz R (2017) A theory of crowdfunding: A mechanism design approach with demand uncertainty and moral hazard. *American Economic Review* 107(6):1430–76.

Sun C, Shrivastava A, Singh S, Gupta A (2017) Revisiting Unreasonable Effectiveness of Data in Deep Learning Era. *arXiv:1707.02968 [cs]* .

Thomas LD, Leiponen A (2016) Big data commercialization. *IEEE Engineering Management Review* 44(2):74–90.

Thomaz AL, Breazeal C (2008) Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence* 172(6-7):716–737.

Tibshirani R (1996) Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)* 267–288.

Train KE (2009) *Discrete choice methods with simulation* (Cambridge university press).

Vielma JP (2015) Mixed integer linear programming formulation techniques. *SIAM Review* 57(1):3–57.

Wang Y, Kosinski M (2018) Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. *Journal of personality and social psychology* 114(2):246.

Warren M (2016) The cure for cancer is data, mountains of data. URL `https://www.wired.com/2016/10/eric-schadt-biodata-genomics-medical-research/`, accessed: 2019-09-20.

Wilson HJ, Daugherty PR (2018) Collaborative intelligence: humans and ai are joining forces. *Harvard Business Review* 96(4):114–123.

Xin D, Ma L, Liu J, Macke S, Song S, Parameswaran A (2018a) Accelerating human-in-the-loop machine learning: Challenges and opportunities. *Proceedings of the Second Workshop on Data Management for End-To-End Machine Learning*, 9 (ACM).

Xin D, Ma L, Song S, Parameswaran A (2018b) How developers iterate on machine learning workflows–a survey of the applied machine learning literature. *arXiv preprint arXiv:1803.10311* .

Xin D, Macke S, Ma L, Liu J, Song S, Parameswaran A (2018c) Helix: Holistic optimization for accelerating iterative machine learning. *Proceedings of the VLDB Endowment* 12(4):446–460.

Yeomans M, Shah A, Mullainathan S, Kleinberg J (2017) Making sense of recommendations. *Journal of Behavioral Decision Making* .

Zenteno AC, Carnes T, Levi R, Daily BJ, Dunn PF (2016) Systematic OR Block Allocation at a Large Academic Medical Center. *Annals of surgery* 264(6):973–981.