JÖRG DRÄGER
RALPH MÜLLER-EISELT

WE HUMANS AND THE INTELLIGENT MACHINES

How algorithms
shape our lives
and how we can
make good use of them

We Humans and the Intelligent Machines

Jörg Dräger, Ralph Müller-Eiselt

We Humans and the Intelligent Machines

How algorithms shape our lives and how we can make good use of them

Verlag Bertelsmann**Stiftung**

Contributors:
Carla Hustedt
Sarah Fischer
Emilie Reichmann
Anita Klingel

Editor: André Zimmermann

Contents

The algorithmic society – a preface

Intelligent machines are part of our lives. They help doctors diagnose cancer and dispatch policemen to find criminals. They preselect suitable candidates for HR departments and suggest the sentences judges should impose. It is not science fiction, it is reality. Algorithms and artificial intelligence increasingly determine our everyday lives.

Only a fine line separates fascination from horror. Many things sound promising: defeating cancer before it develops, stopping crime before it happens, getting the dream job without the right connections, serving justice freed from subconscious prejudices. All of that sounds auspicious, yet the negative narrative is just as impressive: healthcare systems which are no longer based on social solidarity, minority groups which suddenly find themselves disadvantaged, individuals who are completely excluded from the job market. In this scenario, people become playthings, the victims of digitally determined probabilities.

Whether promise or peril – the changes will be radical. We must therefore re-evaluate and readjust the relationship between humans and machines. How does artificial intelligence (AI) affect us, our lives and our society? Where can algorithms enrich us, where must we put an end to their threatening omnipotence? Who wins and who loses through digital disruption? These questions are reminiscent of earlier upheavals of similarly broad scope. The Industrial Revolution also changed economic and social conditions, engendering hope for the future, along with considerable fear and social tensions. In retrospect, technological progress has made most people's lives better and has increased prosperity, life expectancy and social standards.

Who would today seriously long to return to the pre-industrial era of the early 18th century?

It would be naïve, however, to simply trust that again this time everything will turn out for the better. Whether intelligent machines will improve society or make it worse is far from clear. The good news is that it is up to us to shape how things change. Algorithms are created by humans and do what humans tell them to do. We are therefore the ones who can decide which interests and values they should serve.

The purpose of this book is to encourage everyone to get involved. We want to show how intelligent machines can be used to serve society, which is one of the most important policy tasks of our time. The book is full of international examples but written from the perspective of Germany, where politicians have been somewhat slow and negligent in responding to digital change. While the debate in our country has generally been a long lament about insufficient wireless coverage and slow Internet access, other nations have clearly outpaced us. In early 2016 – an eternity in digital times – then US President Barack Obama convened a high-ranking expert commission to develop recommendations on how American society could use AI to its advantage. Immediately after taking office, French President Emmanuel Macron made European cooperation on this issue one of his core concerns. It will indeed be necessary to join forces in Europe, since China is prepared to invest the equivalent of $150 billion in AI projects in the coming decade.

Algorithms are here to stay. The Algorithmic Revolution is not something we will simply be able to sit out. It is not a purely economic phenomenon, social concerns are at least as urgent. Intelligent machines can directly impact the common good – which is why we have written this book. In the first part, *The algorithmic world,* it examines the far-reaching changes transforming our lives and the necessity for humans and machines to find a meaningful way to complement their respective strengths. The second part, *What algorithms can do for us,* provides a structured overview of the broad use of algorithms in society and their opportunities, risks and consequences. The third part, *What we must do now,* develops specific proposals for creating a sound algorithmic society,

followed by a brief outlook. With this mix of wake-up call, analysis and ideas for solutions, we hope to fuel a broader societal debate.

That is why this book is not about technology, but about its social consequences and requirements to shape the future. We are not concerned with business models, but with social models. Many practical examples illustrate how the increasing use of seemingly intelligent machines affects both individuals and society as a whole. *Seemingly* refers to the simple fact that algorithms can imitate human intelligence and, in some areas, even outperform us in cognitive terms. This so-called artificial intelligence, however, is limited to narrowly defined tasks and lacks precisely what continues to make human beings unique: our ability to combine different facts, to evaluate and transfer knowledge, and to weigh conflicting interests and goals. Whenever in this book we speak of "intelligent machines" as synonyms for algorithms – even more correctly, as synonyms for algorithmic (software) systems – we are very aware of this essential limitation of their "intelligence." Even then, however, their impact remains extremely far-reaching.

Our book was originally published in the spring of 2019 in German. Since the topic is global and since we received a lot of interest from abroad, we decided to follow up with this English translation. Consequently, it was carried out by the artificially intelligent translation software DeepL, enriched by some editing. We hope that the outcome of this machine-human collaboration enables a broader community to build upon our thinking.

We Humans and the Intelligent Machines looks at the great challenges caused by the Algorithmic Revolution through the lens of the common good – independently and impartially, but by no means apolitically. Like the Bertelsmann Stiftung's *Ethics of Algorithms* project (www.ethicsofalgorithms.org), we want to raise awareness of upcoming changes, structure the debate, develop solutions and help to initiate their implementation. In doing so, we are guided by a clear precept: The motivation to take action must not be triggered by what is technically possible, but by what is socially meaningful. This book is intended to encourage you to take up this notion and get involved. It remains up to us to ensure algorithms and AI are here to serve humanity.

MEDICINE

Diabetes treatment
Vivienne Ming (USA)
Personalization, p. 59

Telemedicine for at-risk patients
Charité (Germany)
Leeway, p. 97

FINANCE

High-school admissions
New York City (USA)
Always, p. 21

Learning math in school
Bettermarks (Germany)
Personalization, p. 61

Learning math in school
New Classrooms (USA)
Personalization, p. 64

Student recruitment
ITT Institute (USA)
Personalization, p. 66

PUBLIC SECURITY

Predicting recidivism
COMPAS (USA)
Always, p. 17,
Algorithms, p. 45,
Justice, p. 144

Predictive firefighting
Fire Cast (USA)
Distribution, p. 116

Recruitment: discrimination
Kyle Behm (USA)
Algorithms, p. 50

Recruitment: matching
Knack (worldwide)
Access, p. 77

Recruitment: equal opportunity
Pymetrics (worldwide)
Justice, p. 139

VISUALS AND LANGUAGE

Language testing
Immigration assessment (Australia)
Algorithms, p. 42

AI for the blind
AIpoly Vision, Seeing AI (worldwide)
Empowerment, p. 85

Translation
Google Translate, DeepL (worldwide)
Empowerment, p. 88

ONLINE PLATFORMS

Calculating trust
Trooly, AirBnB (worldwide)
Empowerment, p. 90

Online dating and matching
Parship (Germany)
Connection, p. 148

The algorithmic world

1 Always everywhere

*"In short, success in creating effective AI could
be the biggest event in the history of our
civilization, or the worst. We just don't know."*[1]

Stephen Hawking, physicist (1942–2018)

December 11, 2017. It is the day the New York City Council reclaims its right to self-determination.[2] For the 8.6 million residents of the US metropolis, it is an important victory to ensure that the algorithms used there will become more transparent. As a result, New Yorkers are perhaps the world's first citizens to have the right to know where, when, how and according to which criteria they are governed by machines. The man who leads the fight is James Vacca – a Bronx Democrat who heads the Committee on Technology during his third and final term as a member of the City Council. The law to be passed today will become part of his political legacy, and its significance could potentially extend far beyond New York and the United States.

"We are increasingly governed by technology."[3] With this sentence, Vacca begins his speech introducing the bill. By "we" the 62-year-old means the citizens of the city but also himself and his fellow City Council members. New York's public administrators have been using algorithms for some time and in a wide variety of areas: law enforcement, the judiciary, education, fire protection, social transfers – all with very little transparency. Neither the public nor their elected representatives know which data are fed into the algorithms and how they are weighted. In such situations, it is just as difficult for citizens to object to automated decisions taken by the authorities as it is for elected representatives to exercise political control. Vacca fights against this lack of transparency, wanting every office that uses algorithms to be accountable to the City Council and to the public. He wants to shed light on the black box of the algorithmic society.

Much has changed since Vacca first began working nearly 40 years ago. At the beginning of his career, letters were written on typewriters. When they were to be replaced by computers, he thought it was a waste of money. Vacca is anything but a digital native. But he is not a digital naive either. Through his work for the Committee on Technology, he knows to what extent computer-based decisions affect the daily lives of New Yorkers: Police officers patrol on the basis of machine-generated crime forecasts, students are assigned to their secondary schools by computers, social welfare payments are checked by software, and pre-trial detention is imposed on the basis of algorithmically calculated recidivism rates. In principle, Vacca has no objection to that. Yet he wants to understand how these decisions are made.

Vacca was irritated by the lack of openness in administrative procedures as early as the 1980s. At the time, he was annoyed by what he considered a shortage of personnel at the Bronx police station which he oversaw as district manager. When he turned to the relevant government agency, he was told that the crime rate in his district was too low for more policemen. The underlying formula used to calculate the rate, however, was not given to him. Therefore, he could neither understand nor question the quota, nor take action against it.

Vacca wanted more transparency. In August 2017, he presented the first version of the bill to the City Council. It would have required all public authorities to disclose the source code for their algorithms. Yet the experts put the brakes on during the Committee on Technology hearing: The subject area is still too unknown, they said. Too much transparency would endanger public safety, make the systems vulnerable to hackers and violate software manufacturers' intellectual property.

Vacca had to make concessions. A commission of academics and experts was set up to draft rules, due by the end of 2019, on how City Council members and the public will be informed about such automated decisions. Vacca was nevertheless satisfied because the commission has a clearly defined mandate: "If machines, algorithms and data determine us, they must at least be transparent. Thanks to the trans-

parency law, we will have a better overview and understanding of algorithmic decision-making, and we will be able to make agencies accountable."[4] The trend towards more openness and regulation seems unstoppable.

The legislative initiative has already stimulated a number of changes. The use of algorithms is now on New York's public agenda – in the City Council, in the media, among the city's residents. Algorithms are a political issue. A debate is taking place about what they are used for. And they are already used very broadly.

In the service of safety

It is not only 911 emergency calls but also computer messages that send New York police officers out on their next assignment.[5] No crime has occurred at the scene assigned to the police by the software. According to the automated data analysis, however, the selected area is likely to be the site of car theft or burglary in the next few hours – crimes that could be prevented by increased patrols.

Algorithms are managing law enforcement activities. In the 1990s, New York City was notorious for its high crime rate and gangsterism. Within one year, 2,000 murders, 100,000 robberies and 147,000 car thefts took place. New York was viewed as one of the most dangerous cities in the world. Politicians reacted. Under the slogan "zero tolerance," tougher penalties and higher detection rates were meant to make clear: Crime does not pay.

But what if modern technology could be used to prevent crime before it even occurs? The New York police force also considered this, although it initially sounded like science fiction. The Spielberg thriller *Minority Report*, based on the short story by Philip K. Dick, played the idea through in 2002: In a utopian society, serious crimes no longer happen because three mutants have clairvoyant abilities and reliably report every crime – a week before it is committed. Potential offenders are detained. Chief John Anderton, played in the movie by Tom Cruise,

leads the police department and is proud of its results until one day his own name is spat out by the system. He is now considered a murderer-to-be and desperately tries to prove his innocence.

In New York City, algorithms play the same role that the three mutants do for Dick and Spielberg: They provide crime forecasts. Yet with one decisive difference: The computer does not predict *who* will commit a crime in the near future but *where* it will take place. The term for this is "predictive policing."

And it works like this: Software evaluates the history of crime for each district of New York in recent years and compares the identified patterns with daily police reports. Crime may seem random at first glance, but in fact certain crimes such as burglary or theft adhere to patterns that can be worked out. These patterns depend on demographics, the day of the week, the time of day and other conditions. Just as earthquakes occur at the edges of tectonic plates, crime takes place around certain hot spots, such as supermarket parking lots, bars and schools. The predictive policing software marks small quadrants of 100 to 200 meters in length, where thefts, drug trafficking or violent crimes have recently taken place, which – according to the analysis – are often followed by other crimes.

Since law enforcement officers started using predictive policing, their day-to-day work has changed. In the past, they were only called when a crime had already been committed and needed to be solved. Today, the computer tells them where the next crime is most likely to occur. In the past, they often took the same route every day, but now the software determines so-called crime hotspots where they need to be present to monitor what is going on. The police can thus better plan and deploy their resources and work more preventively. "The hope is the holy grail of law enforcement – preventing crime before it happens," says Washington law professor Andrew G. Ferguson.[6] New York Mayor Bill de Blasio sees this in a more pragmatic and less poetic way: Algorithmic systems, he argues, have made police work more effective and more trustworthy. The city is now safer and more livable.[7] In fact, within 20 years the number of murders in New York City has fallen by 80 percent to only about 350 per year. Thefts and robberies also fell by 85 percent. It is not possible to determine

exactly how much predictive policing has contributed to this. In any case, the software enables policemen to be where they are needed most.

The specific functioning of the algorithms, however, remains hidden from the public: How do these programs work? What data do they collect? There are lawsuits pending against the New York police for violating the Freedom of Information Act. People have just as little knowledge about where the algorithms are used, the plaintiffs argue, as they do about how the calculations take place. The first court to hear the case ruled in favor of the plaintiffs. Nevertheless, the police continue to refuse to publish detailed information about their predictive policing.

The New York Fire Department also prefers preventing fires to extinguishing them.[8] But like the police, it struggles with limited resources. Not all of the 330,000 buildings in New York can be inspected every year. The firefighters must therefore set priorities and identify the buildings most at risk. But which ones are they? This selection process alone used to occupy an entire department. For a few years now, the firefighters have been using a computer program that algorithmically calculates the risk of each building catching fire. Taking into account the size, age, building material, pest infestation and inhabitant density as well as the history of fires in the neighborhood, the algorithm creates an inspection list for the next day (see Chapter 10).

In the service of justice

"Smaller, safer, fairer."[9] Using this motto, Mayor de Blasio presented his plan to close New York's largest prison in June 2017.[10] In the 1990s, most of the city's then 20,000 prisoners were incarcerated on Rikers Island, once known as the new Alcatraz. By now, less than 10,000 New Yorkers are imprisoned and Rikers Island, which costs $800 million a year to run, is partly empty. Moreover, the prison has recently been shaken by a scandal about the mistreatment of a juvenile detainee. De Blasio therefore has several reasons for wanting to close the facility. He also would like to further reduce the number of prisoners: to 7,000 in five years and to 5,000 in the long term.

His biggest lever: algorithms. They are supposed to help New York's judges better assess risks, for example, whether pre-trial detention is necessary or whether an early release is adequate. The probabilities to be assessed here are, in the first case, the danger that the alleged person will flee before the trial and, in the second case, the threat of recidivism. These probabilities depend on so many factors that a judge can hardly be expected to evaluate all of them adequately in the time allotted for each case.

COMPAS (Correctional Offender Management Profiling for Alternative Sanctions) is the software that calculates the risk of flight and recidivism. While the company that developed the program refuses to publish the algorithm behind it, research by ProPublica, a non-profit organization for investigative journalism, has shown that such systems collect and analyze a large amount of data, such as age, gender, residential address, and type and severity of previous convictions. They even gather information on the family environment and on existing telephone services. All in all, *COMPAS* collects answers to 137 such questions.

The potential for providing algorithmic support to judges is huge. In a study in New York City, researchers calculated that if prisoners with a low probability of recidivism were released, the total number of detainees could be reduced by 42 percent without increasing the crime rate.[11] In Virginia, several courts tested the use of algorithms. They ordered detainment only in half as many cases as when judges issued a ruling without such software. Despite that, there was no increase in the rate of people who did not show up for their trial or who committed a crime in the interim.

Algorithmically supported decisions improve forecasts even if they do not offer 100-percent accuracy. In addition, they could also reduce variations in the sentences handed down. In New York City, for example, the toughest judge requires bail more than twice as often as the most lenient of his colleagues. The fluctuations may be due to the attitude of the judges but also to their workload, since they only have a few minutes to decide what bail to set.

What promises advantages for society can, however, result in tangible disadvantages for the individual. Hardly anyone knows this better

than Eric Loomis, a resident of the state of Wisconsin. In 2013, he was sentenced to six years in prison for a crime that usually draws a suspended sentence. The *COMPAS* algorithm had predicted a high probability of recidivism, contributing to the judge's decision in favor of a long prison sentence. The discrimination that can result from the use of algorithms will be discussed in more detail in Chapter 4.

In the service of efficiency

Every autumn in New York City, the application phase for high school begins.[12] For many parents this is a time of stress and uncertainty because there are too few places at the popular schools known for getting their students into good colleges and thus providing better career prospects. The teenagers and their parents research secondary schools for months, and some have taken admission tests or gone in for interviews. The right high school should be academically challenging, have good sports facilities and ideally be located in the neighborhood. Naturally, it would also have a high graduation rate and be seen as competitive. Approximately 80,000 young people and their parents have until December 1 to choose 12 schools from over 400 options on the application form. The following March, the Department of Education will tell them which school they can attend.

Until 2003, the department's staff had to allocate slots manually – a complex task that took place under considerable time pressure. The amount of administrative work was immense, and the result was unsatisfactory because 41 percent of the students did not get a place at one of the four schools they could select back then. Dissatisfaction among students and families was correspondingly high. Children with poor grades or from poorer households were seldom given a chance, while highly committed parents always came up with some new way to get their offspring into one of the best schools.

Today, New York's young people have a better chance of going to a school of their choice since neither administrators nor lotteries are selecting the secondary school. That is now the job of an algorithm.

A method derived from game theory allows a much more accurate fit between students' preferences and schools' capacities. Today, 96 percent of the students in America's largest city go to a high school of their choice, and not only because the wish list has been expanded from four to twelve. Half of the students receive a place at their most preferred school, another third at their second choice. The new system prevents instances such as those occurring in the past where some children were accepted at several of their chosen schools and others at none at all. The matching has become much more efficient.

New York City uses algorithms to optimize a standard distribution problem: Too many applicants have to be assigned to too few places. With other high-demand goods, such as tickets for a popular concert, the solution would be simple. Prices would simply be increased until supply and demand are balanced. But access to public goods such as school education needs to be determined by other criteria – which were developed for New York by a Nobel Prize laureate. Alvin E. Roth of Stanford University designed an algorithm that only makes a final allocation after several preliminary rounds of virtual matching, taking into account both students' preferences and the schools' capacities and selection criteria.

Nevertheless, this algorithm does not solve all the problems faced by the city's education system: Social inequalities, for example, are not eliminated by efficient allocation, nor is the fact that pupils from different backgrounds tend to go to different schools. Furthermore, there are still not enough slots at the popular schools and there is a clear gap between the educational opportunities in New York's richer and poorer areas. Children from socially disadvantaged households and with lower grades still tend to end up in underfinanced and poorer schools. Parents in underserved neighborhoods may be happier because their child is given a place at the nearest school, but that does not make the school the best choice for the child. Students from more affluent households, on the other hand, often receive intensive support, including from professional consultants, in drawing up their wish lists.

Algorithms that try to solve complex tasks more efficiently are not only used in New York City's schools. How social welfare benefits are verified has also been automated.[13] In 2009, 48,000 investigations into

welfare fraud were carried out manually, with only $29 million seized as a result. Today, an algorithm recognizes the patterns of fraud much more reliably. The number of investigations has been reduced and, with that, the number of false accusations; at the same time the amount of money recovered has increased. In 2014, $46.5 million was recovered after only 30,000 investigations. However, the lack of transparency remains a problem here as well. Although fraud perpetrated at the expense of the general public can now be detected more efficiently than in the past, the individuals involved are given little insight into the criteria used to investigate them. Yet a high degree of transparency would be desirable, especially when it comes to distributing social benefits, since that would increase credibility and trust that administrative decisions are taken fairly.

Setting the course

New York City is not the only American city where algorithms are omnipresent. Chicago and Los Angeles provide their judges with support in the form of software or use predictive policing as well. Algorithmic systems are also used outside the US, for example in Australia, where they decide on social benefits and even automatically send reminders and warnings when potential fraud is perceived (see Chapter 9). Germany is not there yet but initial applications do exist: In Berlin, places at primary schools are allocated using software (see Chapter 10) and algorithms check tax returns for plausibility. Six of the country's states use different forms of predictive policing (see Chapter 11). Especially in large cities, public administration has become so complex that municipal services from police patrols to waste collection can hardly be managed without technological support – including the use of algorithms. They are part of the daily life of every citizen. But most citizens do not know these algorithms exist, let alone understand how they function. People do not need to understand, you might say. They should be happy if the garbage is picked up on time and no unnecessary costs arise for them as taxpayers.

Yet with decisions about imprisonment, access to the best educational path or governmental support, algorithms intervene deeply in the fundamental rights of individuals. This makes the software and its design highly political. Such seemingly intelligent systems should not only be debated behind closed doors or among academics but also in a broad social and political discourse – especially since even well-designed algorithms can discriminate. In the fight against crime, they can be self-reinforcing: The police find the most crime in the areas they investigate the most. Minor drug offenses, for example, common in most parts of a city, are identified disproportionately frequently in certain neighborhoods, leading to even more police checks there. Or in the case of the courts: When an algorithm sends people to prison for a longer period of time, they are more likely to remain unemployed after their release. They will also have less contact with family and friends and will therefore be more likely to become repeat offenders which confirms the algorithm's predictions. Critics argue that all this reinforces the discrimination against and stigmatization of certain social groups.

As New York City shows, algorithms can solve tasks that are too complex for humans. They can be useful helpers for us and our societies. But whether or not they are successful depends on the goals we set for them. They are neither inherently good nor bad. Ideally, they result in more safety, justice and efficiency. At the same time, however, they can reinforce existing social inequalities or even create new forms of discrimination. It is up to us to set the course so that things develop in the right direction.

James Vacca now teaches at Queens College, City University of New York. His years on the City Council are over since its members can serve a maximum of two consecutive terms. He proudly looks back on December 11, 2017, and his greatest legacy, the algorithmic accountability law, saying: "We were the first to politically concern ourselves with algorithms. Algorithms are helpful, it would be wrong to ban them. But we have to regulate how to deal with them. It is *the* political task of our time."[14]