

Privacy in an age of the internet, social networks and big Data



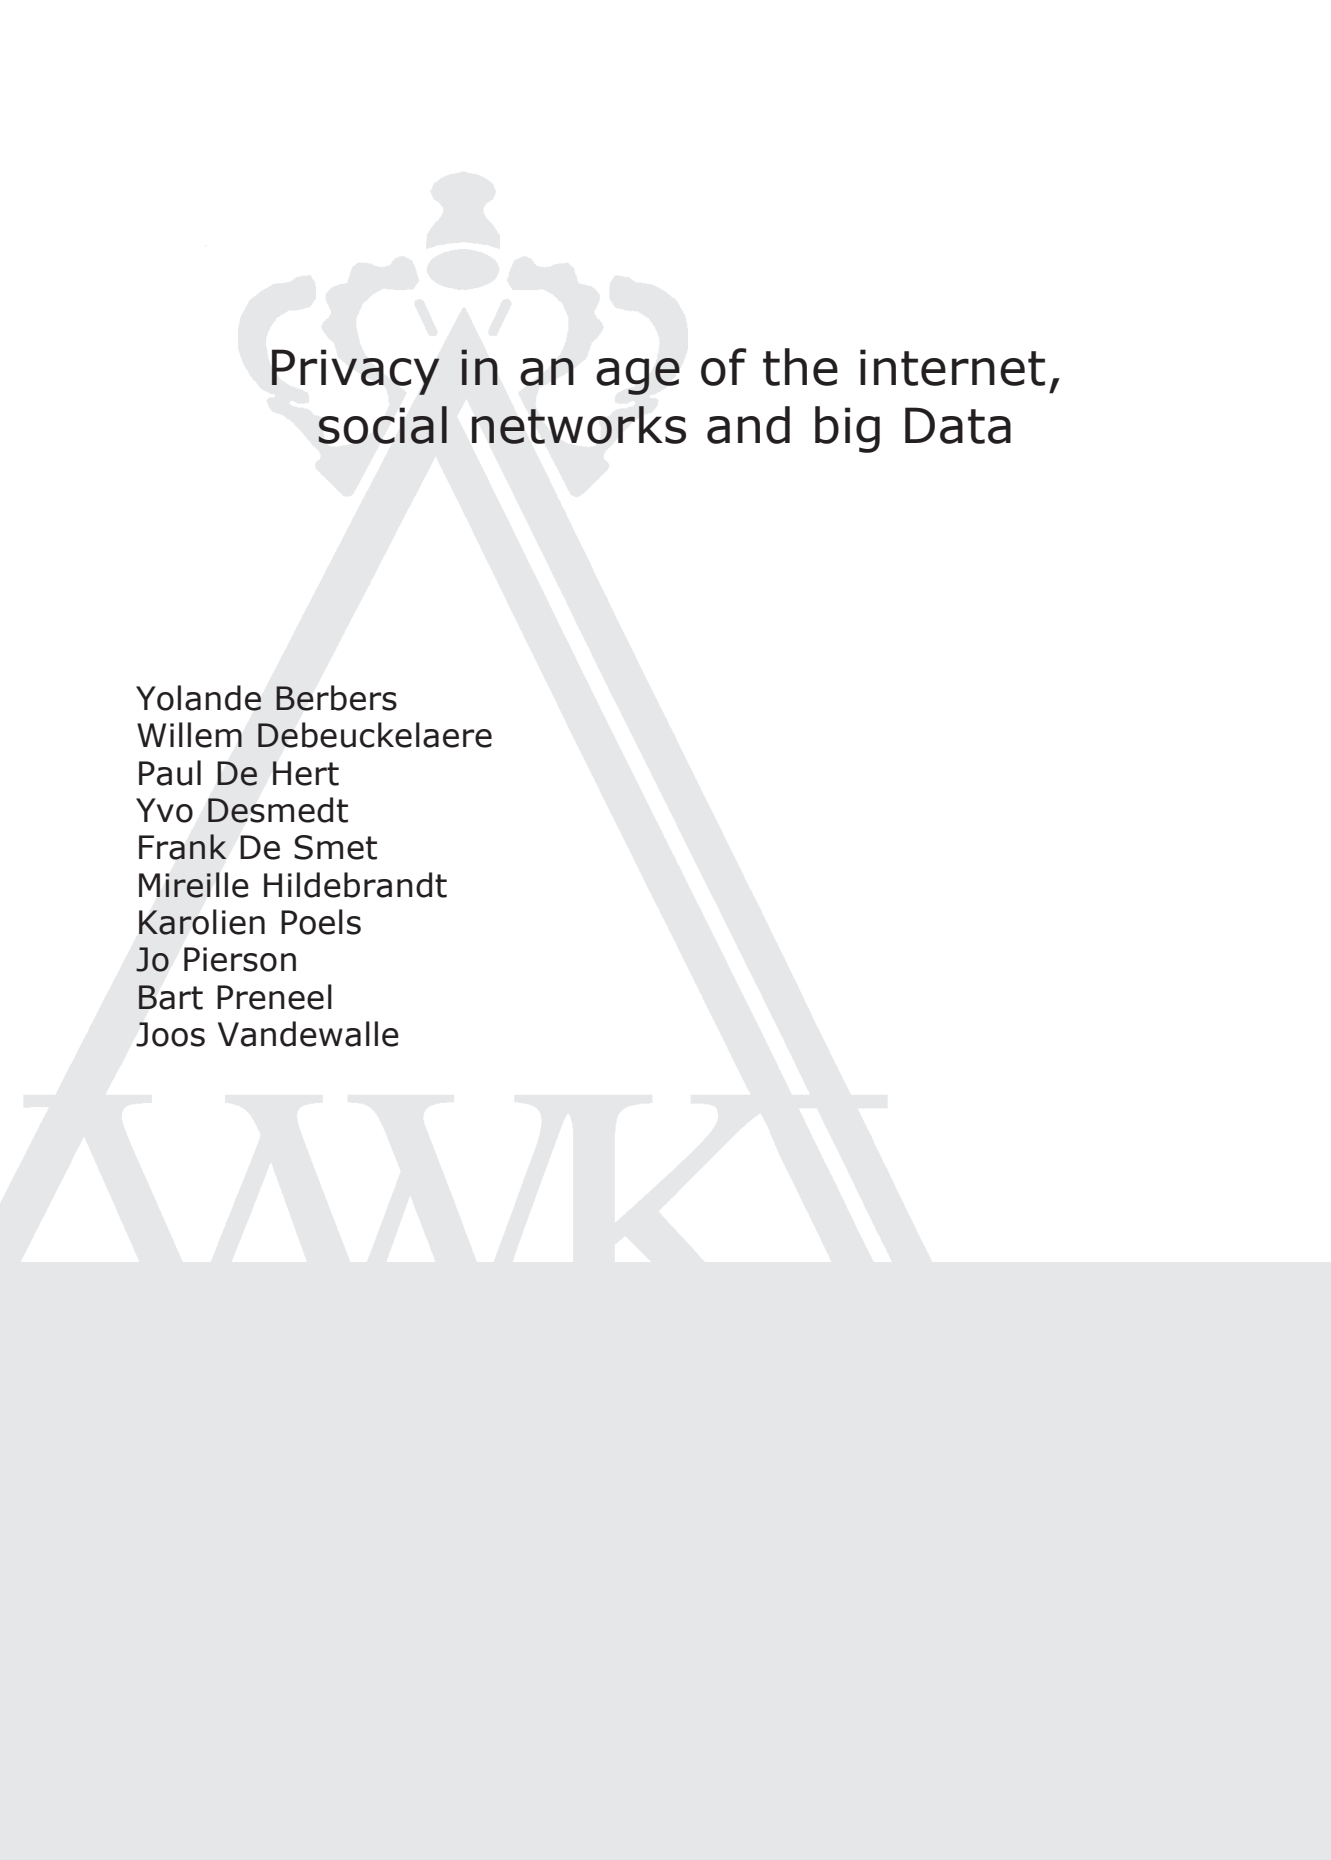
Published
by
the Royal
Flemish Academy
of Belgium
for
Science
and the Arts

Standpunten nr. 49 b



KVAB Press

Hertogsstraat 1
1000 Brussel
Tel. 02 550 23 23
www.kvab.be
info@kvab.be



Privacy in an age of the internet, social networks and big Data

Yolande Berbers
Willem Debeuckelaere
Paul De Hert
Yvo Desmedt
Frank De Smet
Mireille Hildebrandt
Karolien Poels
Jo Pierson
Bart Preneel
Joos Vandewalle

Partial reproduction is permitted provided the source is mentioned.

Suggested citation: Yolande Berbers, Mireille Hildebrandt, Joos Vandewalle (et al), *Privacy in an age of the internet, social networks and big Data* KVAB Standpunt 49 b, 2018 (Original text: Dutch).

© Copyright 2018 KVAB
D/2018/0455/02
ISBN 978 90 6569 918 28

Translation Dutch manuscript: An Baeyens

Photo and cover design: Anne-Mie Van Kerckhoven

Privacy in an age of the internet, social networks and big Data

TABLE OF CONTENTS

Executive summary	2
Preface	5
1. Introduction and context.	6
2. Description and context of the key terms	7
2.1 Exploring the terminology	8
2.2 Scoping the problem.	11
2.3 A survey of potential solutions	20
3. Analysis of privacy from some relevant case studies.	32
4. Conclusions and recommendations tailored to particular groups.	46
4.1. Conclusions	46
4.2. Recommendations.	48
References	52
Members of working group	57

Executive summary

The current use of the Internet, social media and big data severely affects the privacy of ordinary users. This positioning paper is primarily aimed at the private user young and old who did not have special education or training regarding ICT but still uses these services intensively and who, whether or not, rightly worries about the hazards to which his or her privacy is exposed. This requires not only a better and deeper understanding of the technological possibilities and limitations, but also the commercial interests, and their relation to the constraints and threats of our personal privacy when using the many often valuable services. The specific aspects of privacy as patients, or the privacy regulations for companies and institutions that track and process files with data from individuals, employees, students, or customers, is not dealt with but is referred to other reports. This positioning paper has been conceived by a working group of members of KVAB and external experts covering the different aspects of this interdisciplinary subject, that have met regularly over a period of one year.

Since the ICT world is often overwhelmed with “jargon” words, the scope of which does not penetrate or because the newspapers sometimes describe very frightening low-backed situations, we first discuss the main concepts both at the level of the machine learning, data extraction and the big data, as well as the privacy issues that arise, and finally the ways in which a better privacy can be acquired.

In order to make this more concrete for the modal reader, we discuss important privacy hazards in a number of concrete situations, such as the digital life of a family, the big data police in passenger profiles, the internet of things, the context of smart cities, distributed information versus central collection, autonomous vehicles, and location information. Although this digital revolution is not over yet, the modal user can already modify his behavior.

There is extensive scientific literature on this subject, but there are also many widely accessible texts available recently, including websites, to which the interested reader is referred to in the bibliography.

The ten recommendations mainly focus on various target groups and situations.

Recommendation 1: Responsibilities. Privacy in the big data is an issue for citizens, engineers, consumers, companies, institutions, media and governments. This calls for the provision of sufficient resources to the supervisors, especially with regard to companies that derive their earnings model from big data analysis.

Recommendation 2: Alert citizens. Citizens, whose data are being processed, should try to maximize their rights under the GDPR. The verification of personal

data requires that the individual gains insight into the use and misuse of the data, as a precondition for genuine freedom of choice. Precisely because it is extremely difficult for individuals, we recommend that those concerned use the opportunity to exercise their claims through mandating to consumer or privacy organizations (Article 80 GDPR).

Recommendation 3: Providence, Profile Transparency, and Goal Binding. Although the profiles themselves are not related to a particular person and thus are not personal data themselves, the application applies to a person who fits within the "validation" of the profile, under the fundamental right to data protection (GDPR). The right to profile transparency implies the obligation to inform stakeholders and explain how they are profiled and this beyond a correlation or statistical relationship.

Recommendation 4: Power Unbalance. If the person responsible for an ICT service relies on the consent for the use of personal data, then it must be easy to withdraw, with a limitation of permission in time. They will not apply a manifest power imbalance between the data subject and the controller or processor, e.g. because the responsible person provides the dominant (or only) service in the market. The controller must demonstrate that there is no power imbalance or that this imbalance cannot affect the consent of the person concerned.

Recommendation 5: The builders of ICT and IoT devices must make use of technologies that maintain privacy and allow transparency for the end user. They need to work on 'privacy by design', taking privacy from the start of the design as an important requirement, and not being "stuck" afterwards. The service providers must allow users to assemble services of different origins. The designers of algorithms must write their algorithms to ensure users' privacy. Application designers need to allow transparency, work on efficient and effective technologies that allow users to authorize their data usage. Additionally, one must make certification of applications so users are sure that the applications are safe. Typically privacy must be default.

Recommendation 6: Role of government and companies. It is the duty of government and companies to check for each big data solution whether the risks for the protection of personal data and the risks to society as a whole outweigh the benefits. In doing so, one should always check if it is not possible to achieve the same goal by using less data or aggregating data.

Recommendation 7: Preventing unwanted data bias. The responsible designers and service providers must always check whether inaccurate or unfair 'data bias', 'algorithm bias' or 'output bias' is hidden in the data sets with which algorithms are being trained, either in mathematical models themselves or in the output (indirect discrimination).

Recommendation 8: Limits to the use of big data by the government. The use of public sector big data, both in the field of detection of tax and social security fraud and in the context of national security, crime and law enforcement, should always be subject to a review by the relevant supervisors. In addition, the legitimacy and the related proportionality must be paramount, which also requires a marginal efficiency test. It is essential that legislation be provided that determines how and when the result of data mining and statistical analyzes (correlations) by the government may or may not be used as legal evidence to make decisions in individual cases (e.g. In dealing with fraud, law enforcement ...).

Recommendation 9: Establishing a digital clearing house. It is advisable to set up a Digital Clearing House (DCH) that monitors the quality of the various digital market regulators.

Recommendation 10: Task of education. Specific to young people, education has a task of bringing awareness, attitudes, skills and behavior from the actual life spheres such as home, school and friends (eg youth associations). It is important to point out to young people the "pitfalls" of their own behavior, as expressed, for example, in the privacy paradox.

Preface

The Academy's *Standpunten* series (Position Papers) contributes to a scientifically validated debate on current social and artistic topics. The authors, members and workgroups of the Academy write under their own name, independently and in full intellectual freedom. The quality of the published studies is guaranteed by the approval of one or several of the Academy's classes. This position paper was approved for publication by the meeting of the Class of Technical Sciences on May 18 2017.

1. Introduction and context

Information and Communication Technology (ICT) has evolved at lightning speed over the past fifty years. It was in 1969 that a couple of brief messages were first exchanged between two computers at a distance of 500 miles (800 kilometres). The Internet was born in 1983 with the launch of TCP/IP, a protocol that allows messages to be split up into packets for sending. The World Wide Web (WWW) was first used at CERN in 1990, with Flanders' own Robert Calliau playing a part. In 2008, social networks went mainstream, with users young and old now using these very user-friendly, and even tempting, technologies all over the world. The evolution of ICT continues at an ever increasing pace.

While ICT brings many benefits with it for us and society, many of its social side-effects were unforeseen and could not have been predicted while these tools were being designed and developed. For instance, it could hardly have been anticipated that these technologies would lead to the concentration of personal data in the hands of a few big players. The harvesting and use of personal data and Big Data is opening up potentially very compelling opportunities for companies to offer new, improved products and services, and one of the things it is bringing about is an expectation among citizens of greater security. However, it could also pose a threat to the personal lives of users and might catalyse the undesirable use and collection of data, unwanted advertising, blackmail and indeed computer crimes. Tim Berners-Lee, the inventor of the Web, has even been warning that privacy might be on its deathbed.

Users are often totally unaware of the risks. Moreover, ICT companies have modest interest at best in these problems, and society is not properly equipped, either in organisational, legal or technical terms, to tackle them. In late 2015, KVAB (the Royal Flemish Academy of Belgium for Science and the Arts) decided to launch a working group to document the most pressing aspects of the issue from the technical, sociological and legal perspectives. The present document is one of the results of that decision. At its close, it enunciates evidenced recommendations for the various parties involved, including government, business, education and the general public. This is done in view of the desirability of having the present and next generation of technological services (collectively known as the Fourth Industrial Revolution) respect and sustain basic human rights worldwide. The individual must be able to have confidence that his personal data are being processed lawfully, and must have sufficient wherewithal to prevent any illegitimate data processing. For this, effective transparency will be needed, even penetrating the walls of trade secrets and/or national security.

While historical comparisons are only ever partial, they can alert us to the grave impact which might ensue from using collections containing great quantities of personal data (nowadays called Big Data). As early as 1933, Nazi Germany held a

national census of its population of 41 million, keeping a record of personal data including ethnic origin. For the purposes of the Holocaust, a central archive of punch cards (Hollerith cards) was maintained, which could be read mechanically at speed, containing information on 17.5 million people in Germany and the occupied countries. The archive documents took up over 27 kilometres' worth of shelving: lists, inventories, personal descriptions, medical experiment reports, decrees, etc. They demonstrate the Civil Service direction of the murder machine and its scale. The personal data at the disposal of today's Big Data ICT companies are many times greater in size. Moreover, today's computers can automatically search, retrieve and mechanically process them in the blink of an eye.

Readers unfamiliar with the terminology or context of this topic would be best advised to begin with Chapter 2, which explains concepts such as Big Data, machine learning and data mining, sets out the various (and often as yet unknown) problems involved, and discusses the technological, organisational, social and legal solutions. Chapter 3 considers specific case studies, placing the issues raised in Chapter 2 in realistic contexts. Finally, Chapter 4 draws conclusions and makes recommendations for the various players in this issue. Readers who would prefer to gain a concrete impression rapidly of the range of situations which jeopardise their privacy can skip straight to Chapter 3 and flick back where necessary to Chapter 2 for definitions of any unfamiliar terms they come across. A third category of readers, those seeking to know above all how the various actors in the issue ought to deal with present developments, might prefer to go straight to Chapter 4. The evidentiary basis of the conclusions and recommendations made there will be found in the foregoing chapters.

This position paper addresses the generic use of the Internet, social networks and Big Data. Their use in healthcare is the topic of another position paper issued by KVAB [Verdonck, Van Hulle et al. 2017]. Recently, a more specialist report [CBPL 2017], largely of use to those responsible for processing personal data, containing 33 recommendations, was issued by the Belgian Commission for the Protection of Privacy (CBPL).

2. Description and context of the key terms

Since the key terms in this issue are recent coinings and tend to be understood only vaguely or incompletely, they are first described concretely, together with the context and developments that have already occurred in the lives of these terms. We shall first consider several of the basic concepts, including Big Data, machine learning and data mining; then, a number of privacy-related problems will be outlined; and finally, we shall consider possible avenues for solutions.

2.1 Exploring the terminology

Big Data and data mining

Big Data is said to be 'different' from previous types of data collection because of the factors known as the three V's: the volume (an exponentially-increased quantity), the velocity (processing speed can even be real-time) and the variety (many distinct types of data can now be integrated for processing) all differ from what has been known in the past. At the least, we have to do with great quantities of both structured and unstructured data. The latter refers to data in many different formats (text, images, sound) and from a plethora of sources, including e-mails, videos, letters, reports, blogs, postings, figures, archives, sensors, cameras and more. Some data are associated with identifiable people; with a well-defined aspect of their identity, for instance. This is the specific meaning of 'personal data', to which particular legal provisions apply. Personal data might be freely volunteered data (e.g. online forms which people submit), observed data (e.g. behavioural observations made by software or sensors) or derivatives (e.g. a creditworthiness score calculated from these data). Other data types relate to the management of product life cycles, transport, or critical infrastructure (e.g. metal fatigue, water levels or climate change). As a term, Big Data has come to mean, and now consistently means, machine-readable digital information that can be processed by computer systems and that is directly linked to technologies that enable the querying and analysis of large quantities of data, not necessarily pre-sorted. This is what is known as 'analytics': analyses made of electronic datasets using digital analysis techniques, calculation programmes or algorithms. In this context, we also speak of 'data mining', by analogy from the extraction of material resources from the earth.

It is sometimes suggested that Big Data contains such quantities of data that analyses of it are now flawless. Any errors in it would supposedly be ironed out because — statistically speaking — we have all the data that matter. To claim this is wrong-headed. For example, it is often a challenge to obtain germane data at all. Acquiring them may be expensive, or technically impossible, or thwarted by privacy objections, trade secrets or intellectual property law. Therefore the temptation is considerable to work with more easily-obtained data (often called 'low-hanging fruit'), yet such data might be irrelevant or incomplete, or might contain prejudices (distortion or bias), which will quickly throw up erroneous or irrelevant predictions in the model. Sometimes, these data are mere traces of or references to events, not measurements as such of the phenomenon being studied. In addition, we can by definition only work with data from the past or (where there is data streaming) from the present; future data points will invariably be extrapolations. The central question, then, is always to what extent the available data are representative of new or future data. This question can only be answered by considering the purpose to which machine learning is being applied. There is

no such thing as 'the right way to represent reality using data'. Those setting out to predict tax fraud patterns ought to collect data that allow for a distinction to be made between future defrauders and honest taxpayers; the only problem is that it is not known to us what data would make such a distinction possible to draw, and this might well tempt some to strike out and make a start with the kinds of data (the low-hanging fruit) available to the tax authority in question. A supplementary point here is that it is not even known which tax frauds have evaded detection heretofore. We see, then, that it is a perilous venture to collect ever more data points and to seek patterns in these: the infringements of privacy that this causes can only increase, and success is not guaranteed. There might well be a self-fulfilling prophecy at play here: if we focus on those known at any given time to be committing fraud, then we are wrongly assuming that the patterns relevant to our task will be found among this set [Harcourt 2007].

Machine learning

Machine learning is a sub-discipline of computer science that has revolutionised artificial intelligence. To grasp the impact of Big Data, it is crucial to have a sound understanding of machine learning. The simplest yet clearest definition [Tom Mitchell 1997] is: '*A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P , if its performance at tasks in T , as measured by P , improves with experience E .*'

Importantly, if computer systems are to have the ability to learn, they must be capable of accumulating 'experiences', so that they can measure the effect of their own conduct and adjust that conduct accordingly. This is an entirely automated process, achieved by what are known as 'algorithms'. An algorithm is a step-by-step plan, a kind of recipe or set of instructions that computers can use to undertake a task. In machine learning, these plans are partly made up of mathematical functions that go looking for relationships between different data points, or that seek to build up profiles of people, groups of people or organisations. For instance, data held by a credit provider might indicate that people who place an order after midnight are much more likely to fail to pay the bill. A web shop that uses that credit provider might respond to that fact by deciding that people exhibiting this behaviour are not allowed to pay in arrears. In order to discover this correlation, the credit provider will have had to feed the algorithm with data, known as the 'training set'. In a properly-run system, the patterns discovered should regularly be tested (and the algorithm might have to be retrained if the system performance deteriorates) using new data; these form the 'test set'. A worthwhile algorithm, then, will be well capable of 'generalising', i.e. correctly classifying previously unseen data. The reliability of machine learning largely depends on the relevance and completeness of the training set and test set employed, and moreover on the algorithms trained to make the predictions, and on the speed with which results are required.

There are two pitfalls that we can point out here. One is that algorithms have been known to identify extremely detailed relationships which document the training set fairly accurately, yet are not so adequate at predicting relationships that occur in the next dataset. This is known as 'overfitting': the fit between the data and the mathematical model is 'too good', preventing decent generalisation. The other is that the linkages might be highly universal but still do not apply to many individual cases; this is known as 'overgeneralisation'. There are always trade-offs to be made between the scope, completeness and relevance of training sets, the speed with which they can generate a result, and the level of detail in the relationships. It is impossible to score highly on all these criteria simultaneously. In practice, all sorts of choices are therefore made, which may have a negative impact on the results. Consequently, the wise course of action is to keep scrutinising the intended use of machine learning in any application. It affects matters greatly whether predictions have been made on the basis of a sufficiently rich dataset or rather are just a very rough indication of potentially relevant relationships, particularly where we have to make decisions about people. In our credit assessment example, one can well imagine that when extra data become available, the correlation between night-time ordering and failure to pay has to be contextualised; it might, for example, turn out that the correlation applies to men only. In our tax authority example, it might be that keeping an eye open for as yet unknown fraud cases, and setting priorities on which kinds of fraud to tackle hardest, will keep throwing up new insights.

The need for quality data

Controlled versus uncontrolled machine learning. Certainly, machine learning comes in a variety of forms, just as there are various types of algorithms available. It is not always obvious which kind of machine learning and which types of algorithms are most fit for purpose. Currently, most machine learning systems employ mathematical formulations which seek to describe mathematically, and to optimise, the relationships between the available input and the desired output of a system. Any kind of machine learning will involve what is known as a 'hypothetical space', although it may be that the hypotheses involved were partly developed by the software itself. It is important to embed data-driven practices within empirically- and theoretically-driven practices, such that data-driven research does not determine the entire methodology. This will obviate the risk of being impelled by data and algorithms, when one ought to be steered by facts and insights. There is a tendency abroad to confuse the two, as though data constituted facts and as though algorithms were a kind of magic potion. Nothing could be further from the truth. Data are traces of, references to, or representations of facts; nothing more and nothing less than that. As stated previously, we have no inherent guarantees that data are up to date, complete and/or relevant to our tasks. Computer scientists are more aware of this than anyone else. Sadly, it is often far from straightforward for policymakers, advertisers and the various

kinds of vendors of services to avoid the methodological pitfalls that mitigate the usefulness of data-driven practices.

It will be evident that a great quantity and high quality of data is needed to make good decisions, ones based on reliable associations. This brings us to three points of attention:

- Decisions regarding people will often require the processing of many personal details, potentially from widely differing contexts. This might in turn accrue detailed profiles of individuals, greatly infringing their privacy;
- Even where these profiles remain at an abstract level, such as because they are derived from aggregated data, they might in application actually have a major impact on people's personal lives;
- In instances where individuals' rights and liberties might be infringed, the significant, sustained involvement of a human being might remain necessary in the making of decisions based on data (although even then, the chance of erroneous decisions remains). For instance, machine learning might be suitable for use to conduct a daily preliminary screening of millions of credit card operations in order to pick out the potentially fraudulent among them, and these might best be further examined by human expert eyes. Machines are not the only entities that learn: the people who work with them, too, must learn lessons when machine decisions lead to unjust infringements of privacy.

2.2 Scoping the problem

This section will discuss problems and mechanisms that often fail to register with users but which may significantly influence or subconsciously direct them, including: bias, choice architecture, nudging, AB testing, consumer influencing, tracking, and search engine advertising.

Bias

One of the problems that Big Data and machine learning give rise to is the possible prejudices inherent in (1) the data, (2) the analytical techniques and/or (3) the findings of the analysis. The literature typically refers to this problem as 'bias': presumptions or prejudices which lead one in a certain direction, thereby potentially both creating and limiting the boundaries of one's hypothetical space. An unbiased dataset (with random distribution) does not exist — unless one manufactures it to be such, and even that is hard to achieve. We thus always face the question of whether the bias in our model is relevant and reliable, and, by extension, whether our bias might arise from a problematic social distribution, such as ones determined by income, education, criminal record, state of health and the like. Hypothetical space is a sum of the parts of mathematical functions able to detect

patterns in datasets; without a hypothetical space, no value can be added to raw data. One might work with a very simple hypothetical space, by looking for just a couple of obvious relationships (for example, all linear correlations among the data); conversely, it might be a highly complex space (in instances where non-linear correlations or underlying causal relationships are also being elicited). There is no causal relationship between children's shoe size and their level of general development, but there is a correlation, the causes of which have to be sought in other factors. This is an example that anyone can grasp, but when we turn to consider other relationships, it is tempting to treat correlations as though we had to do with cause and effect, even though in many cases it takes follow-up studies to establish whether there are such relationships and how they might be explained. Examples would include correlations between eating habits and obesity, or between a given genetic profile and susceptibility to certain illnesses. Such correlations depend upon a complex interplay of causes; taking decisions on the basis of a presumed simple correlation is therefore potentially both dangerous (if it overlooks the actual cause) and futile (and may lead to a waste of resources).

As set out earlier in our definition of machine learning-related terms, many decisions have to be taken when deriving new insights from datasets, and these decisions boil down to making trade-offs, such as between the extent, relevance, completeness and correctness of datasets, the types of data points, and/or data format. Every one of these decisions entails financial implications (more and improved data might simply be too expensive or not feasible to acquire), but no less than this, they also have consequences for the reliability of the findings (for instance, where relevant data points were not included in the dataset used, or where the hypothetical models developed failed to pick them up). Trade-offs between cost, speed of obtaining results, and reliability of results are inevitable when the applications are specific. If one is able to compare all the possibly relevant data and to query them using very complex algorithms (the preconditions for a large hypothetical space), the description that one obtains of the data will be highly detailed. As discussed above, that description might even be so precise that the patterns detected cannot be generalised to other data. If one then attempts to retool the algorithm into a more generic description, it may well be that better predictions are achieved, yet even then they will remain statistical in nature. This means, for example, that behavioural profiles obtained by this tweaking will apply to the average person covered by the data profile, but most likely not to any specific individual. If someone fits in a data profile for which the average risk of contracting stomach cancer is 70%, this does not mean that they have a 70% risk of getting stomach cancer. The reason for this has to do with the distribution of patterns across datasets, which become compelling only if they are discrepant from the average. If our given person has relatives with stomach cancer, then his own risk will be above 70%; if these relatives are already old without stomach cancer ever having manifested, then his risk may well be below 70%. Discrepancy from random distribution is the productive bias that makes it possible at all

to discern patterns. As philosophers [e.g. Gadamer 2010] and scientists [e.g. Wolpert 2013] have been noting for years, bias is the precondition for the making of discriminations that enable thought and understanding. Therefore, bias is not only inevitable but in a sense is the foundation for all perception and cognition. This does not mean that 'anything goes', however. Sloppily-prepared, carelessly-tested bias will produce untenable results that yield a distorted impression of reality. In some cases, this is hazardous (one need only think of critical infrastructure); in others, it may lead to unjustifiable discrimination (an example being US judges' use of software to determine the severity of the sentence; software in which the same statistical signifier raises the sentence for black perpetrators and lowers it for white perpetrators [Angwin 2016]).

Illegal or morally reprehensible bias. The fact that bias is inevitable and productive does not exclude the possibility that a given bias might lead to illegal or morally reprehensible discrimination. If the distribution of income between men and women is unequal because women are unjustly paid less for the same work, the dataset on which an algorithm is trained will display a bias that will be reflected in the results. Now, imagine that a researcher is setting out to establish whether women are just as competent as men, and takes mean income as the measure for the study. An algorithm trained on the correct data will produce the finding that women are less competent. For reasons such as this, it is essential to monitor whether datasets themselves display any biases; any such bias must be made recognisable before policymakers base anything on incorrect presumptions. Techniques have now been developed to allow this, such as 'discrimination-aware data mining' [Berendt and Preibusch 2014]. In this example, we have to do with the presumption that higher earners are more competent people. This kind of presumption continually crops up, and it necessitates that we remain watchful for underlying connections which might give rise to illegal indirect discrimination or which might be morally indefensible.¹

Choice architecture, nudging, AB testing

Advertising, marketing and policymaking are all domains where thinking is increasingly revolving around 'choice architecture' and 'nudging'. The idea here is that it is possible as it were to pre-sort the choices made by people, thereby boosting the chances that the 'desired' choice is made. This might, for example, take the form of making the desired choice the default option offered; while people can make other choices instead, experience teaches that most do not bother to do so. It therefore makes a great deal of difference whether default settings on laptops, smartphones, smart meters or social networks minimise the sharing of personal data (with an opt-*in* to onward processing) or to maximise them (with

¹ On the consequences of relying upon machine-learning systems in medical diagnostics, see [Cabitza 2016].

an opt-out of onward processing). Having standard settings that minimise data processing is one way of protecting data 'by default'. Companies whose business model is dependent on the processing of great quantities of behavioural data will nevertheless be predisposed to offer a choice architecture which by default allows the harvesting of all data. Government agencies labouring under the impression that Big Data is going to solve countless problems for them will likewise tend to draft legislation allowing as broad competences as possible for data interception, data requisitioning and/or data reuse.


When machine learning is combined with 'nudging' (giving people a slight push), then it is easy to construct a choice architecture which will tempt unaware citizens and consumers to share unprecedentedly huge quantities of data. As a term, 'nudging' is derived from social psychology and behavioural economics and is based on the premise that people frequently behave *irrationally* in a *predictable* manner. Once you have understood the irrational tendencies that steer our conduct, you have a useful tool at your disposal. By making this step, nudging has left behind the untenable assumptions of 'rational choice theory', which was long in vogue in the economic and policymaking sciences, but today's thinking continues to centre around utility maximisation and rationality.

Another way of influencing as much behaviour as possible as effectively as possible is 'AB testing'. It has become a widely used method of determining which website designs obtain the best yields. In this, we first make up a version of the website that needs 'optimising' and call it version A. We then adapt one aspect of the website (whether by the addition of a new selection button, a different colour scheme, more rapid click-through programming or changed use of language), and call it version B. With these two versions in place, we send half the website's visitors to version A and the other half to version B and record their click behaviour to see which version produces the desired behaviour (whether higher purchase rates, more thorough reading or more penetrating clicks through the website). We then choose the version which better matches the desired output as our new version. This process can be iterated repeatedly, and because visitor behaviour is machine-readable, it can quickly be calculated how to influence visitors successfully. AB testing thus helps design the choice architecture that the client requires.

However, it is pertinent that the legal framework of data protection sets requirements regarding choice architecture and limits the scope for tempting people subtly, as it were, to share their behavioural data. The legal framework requires data protection 'by default' and 'by design'. The question facing us, then, is not *whether* we wish to design a choice architecture, but *which*.

Influencing the consumer

The Internet and social media are inextricably linked to today's consumers and the ways in which they are influenced. People give away an inordinate amount



<http://geek-and-poke.com>

of personal data about themselves online. They do so by their surfing, search, like and click behaviour, and by creating social media profiles, but they do so equally by means of the great variety of peripheral data emitted by devices (such as location and sound captures). Moreover, e-commerce is a booming industry. Increasingly, people are searching for, comparing and buying their goods and services in an online environment. Companies are highly interested in harnessing these various types of personal data to approach actual or potential customers with personalised adverts and promotions, often without the consumers realising what is going on. Again, this consumer influencing is mostly carried out through online channels such as websites, mobile device apps and social media. Many such websites and platforms are provided for free — consumers 'pay' with their marketable behaviour data — and are largely funded by advertising. These datasets are also used by marketers to track trends and emerging markets. Thanks to Big Data and the concomitant data mining techniques, corporations assess that they now know better than ever before what 'their' consumers are feeling, thinking, wishing for and buying, and they can come up with new or improved products and services based on those insights. No wonder, then, that the mass take-up of

the Internet and social media and of the associated Big Data has caused a whole series of transformations in the world of advertising and marketing.

Tracking

'Tracking' is the collection by websites and companies of Internet users' personal data. We have to distinguish between 'first-party' and 'third-party' tracking. First-party tracking is when a company/organisation itself collects information on its users; it might be an e-commerce site, but equally it might be some other kind of website, such as a news site or the website of a public body. The kinds of things that are tracked might include who visits a website when; which parts of the website are visited or clicked on; which products or services are purchased; and the like. An organisation might then use those data to send users targeted advertising, to make recommendations which they see on subsequent visits, or to improve the experience of its own website (possibly using AB testing). An organisation might also sell this information to third parties, in which case it is possible that 'third-party trackers' are active on its website. These tend to be advertising networks and data brokers, who make it their business to follow people across every website they can and consequently have very complete details of someone's online behaviour. This makes these data brokers very powerful, especially if they go on to link online data to offline behaviour, such as supermarket purchases or financial transactions, or when they link the data to location, health and lifestyle data collected from what are known as 'wearable tech devices' or 'wearables' (such as smartphones, smartwatches, activity trackers and mobile applications). By amalgamating all these data flows, data brokers can document a remarkable amount about people's habits and preferences, including sensitive and ultra-personal details about their health, sexuality, politics, finances and more. These collated data streams are particularly valuable to advertisers and marketers, and in view of this importance, they are sold at high premiums. In conclusion [Wearable]:

'They know all about you. They know who you are and where you live, where you work and how you worship, what magazines you read and what websites you visit, what books you love and bands you loathe, what you earn and what you save, what you like to eat and do and say and see and buy. They're the data brokers, and your business is their business.'

'Behaviour-driven advertising' is, then, enabled by tracking. Which advert (including banner advert) someone ends up seeing on a website is something increasingly determined by the website user's personal profile and online behaviour. 'Retargeting' is a major advertising strategy which usually involves banner adverts for products or services which a user previously looked at or purchased. In other words, this type of advertising is based on previous surfing habits. Retargeting crops up on all sorts of websites and is an extremely popular way of advertising on social media platforms such as Facebook. Although recent studies have shown

that it is a highly effective means of advertising (since with it, people are more inclined to click and to carry through to making the purchase), and although it is often regarded as more purposeful than 'random advertising', it is also evident that there is little knowledge of just how retargeting works. When people are briefed on how their prior surfing habits are harnessed and on the concomitant online monitoring that enables this, they assume a more critical and concerned stance on privacy as regards this form of advertising.

Popular websites, such as online newspapers, sell advertising space to advertisers who hope thereby to reach their intended target audience. This might be done directly, as in the classic media: in this model, an advertiser purchases advertising space on a given news website because the typical readership is a good match for its intended customer type. Increasingly, however, online advertising space is being sold on the principle of 'real-time bidding': at the very moment a user logs on to a website, the advertising space is put up for automated auction to the highest bidder. While the advertiser does not get to see the user's personal details, it can employ the intermediary services of an 'ad network' to be able to place the banner in the view of a user who has a particular profile; thereafter, it has to pay repeat fees per 'impression' (the number of times that the advert is shown to a website user) or per 'click' (the number of times that the user follows through to the advert). See, for example, [Google], [Facebook] or [Coursera]. Some mobile apps even take this a stage further, by feeding through users' personal details, such as phone numbers, locational and health data, to advertisers; this is enabled by app developers' addition of libraries of advertising software to their apps as a revenue generator. In addition, it has been demonstrated that this information is often being sent in the clear and is readily interceptable by third parties, such as telcos and government agencies [Demetriou 2016], [Vanrykel 2016].

This approach is an attempt to personalise the entirety of advertising on a website or in an app to the online behavioural profile of each individual user. All this takes just fractions of a second to accomplish, without the user being aware of it. It is uncertain to what extent this kind of 'behavioural advertising' exceeds in effectiveness other types of advertising, such as 'contextual adverts' (those dictated by the type of website on which they are placed).

Search engine advertising

A comparable, much-used type of consumer influencing online is 'search engine advertising' (SEA), where companies pay the search engines, such as Google, to promote their hits in search results or to display them prominently alongside the results, in order to boost the chances of people clicking on them and thus coming across their company. Advertising based on the user's previous searches can also crop up on other websites via this method. Search engine advertising is regarded as a form of 'native advertising', a term for online advertising which is made

to resemble the inherent, non-commercial content of the website one is looking at — such as a listing of search results, a supposed editorial (in online newspapers) or comments masquerading as having been left by fellow users on, for example, a social network website news feed. However, native advertising encompasses 'sponsored' (a euphemism for 'paid') content by a given company, and cannot readily be distinguished from the original content of a website. Its growth is partly a response to the plummeting advertising revenue being earned from traditional ad banners and the mass take-up by surfers of ad blocking software. The native advertising that a given user gets to see is more and more adapted dynamically to the indications of their personal online profile, interests, searches, purchasing conduct and other individual, social and contextual factors [Working party 2010].

A key caveat here is that the data gathered by data brokers and the personalised use of these data by ad networks is bound up with problematic prejudices and can even lead to undesired discrimination. Advertisers have a particular notion of their target audience and of what makes people tick. Hence, women who enter search terms related to pregnancy or trying to get pregnant might immediately find themselves put down as belonging to the 'mothers-to-be' category and then be overwhelmed with unsought-for advertising targeting expectant mothers. Research has also indicated that if one searches using a personal name more typical of black people, one will obtain types of personalised advertising different from those targeted at names more typical of white people. For example, names generally thought of as typically black attract significantly more search engine advertising about negative issues, such as arrest, than 'white-sounding' names do [Sweeney 2013]. This is certainly problematic when someone's name is entered to search for their profile (including professional profile) in the context of a job interview or other important social interactions.

Mighty commercial players; enfeebled users

In this summary of online consumer influencing techniques, it is important to consider a few extremely major, powerful players who have a number of key roles. Google is one of the most prominent examples of such: as the world's number one search engine by far, Google has gradually grown to become the king of the online advertising jungle. This one company has substantially determined how consumer influencing works online. Google has a sophisticated, highly lucrative model of search engine advertising based on paid search terms and per-click payments: frequently-searched terms are sold at a premium to advertisers who pay per user click. Google also runs real-time auctions to respond to search term entry: if, for example, a user types in 'summer holiday in Spain', various travel agencies will compete to have their advert — which takes the form of a search result — appear at the top of the results [Rathenau Inst. 2010]. An additional factor is that since 2008, Google has owned the Web's largest ad network, DoubleClick. DoubleClick is one of the largest third-party trackers and manages the placement and sale of

online advertising space via linkages to personal data and browser histories on countless websites. As of 2016, Google has amended its privacy policy to entitle itself to link DoubleClick data to the data held by its search engine proper and also to Google users' personal accounts (e.g. via its Gmail e-mail service). By this stroke, Google now combines first-party tracking, via its own website and services, with a major third-party tracking system by means of the integration with DoubleClick data (as a result of which, DoubleClick has become a de-facto first-party tracker, too). Other key players, too, such as Facebook, Amazon and Apple, have massive amounts of data at their disposal, and thus huge power when it comes to consumer influencing. To take one example, Facebook can track its users not just on its own platform but also across all kinds of websites which enable means of sharing and following entities via users' Facebook accounts (this is known as a 'social plugin' and takes the form of a Facebook Like button on other websites, etc.).

The question arises of why users are sharing, and continuing to share, valuable personal data so massively, thus apparently volunteering their lives to commercial entities. The prime explanation for this is that they tend not to be aware of what they are doing. It is often argued that when people share their personal details, they are indulging in a spot of 'privacy calculus'. In other words, they have to weigh up whether disclosing their personal details is worth it in view of the quid pro quo — more personalised offers of enhanced relevance to them, free news or other information, or entertainment. Because people have a tendency to focus on the quick wins and to be less aware of the fact that their surrendered data will continue to have a life of their own, they opt (or may opt) to disclose their data. The lack of transparency and difficulty that there is in grasping the extent and nature of the tracking that occurs, makes it a rational impossibility for people to make a proper calculation in this privacy calculus. The calculus tends to be done subconsciously or incompletely. What is more, writers often refer to the 'privacy paradox': even when people speak of their concerns about their privacy online, their minds are not on this at the crucial decision-making moments, such as when they create a personal profile on social media, when making purchases online, when making highly personal and extensive Internet searches, when taking part in attractive online promotions and competitions, when playing amusing games on offer, and so on and so forth. Fun, social drive or the immediately-felt benefit of the information obtained from the deal dominates the calculation. People do not weigh up the safeguarding of their privacy against that side of the bargain; at least, not at these moments of truth.

It would seem, then, that people — not least due to the manipulation of their cognitive biases and emotional finiteness (their lack of total knowledge and transparency about the quantity and nature of data harvesting operations, the immediate appeal of the online environment and the goods and services offered there) — are not properly equipped to deal with these data collection practices

and the associated applications of consumer influencing which loom on every hand. This is concomitant with the previously-discussed nudge practices which, besides exploiting the human limitations just mentioned, also very much shape the environment in which these vulnerabilities feature.

2.3 A survey of potential solutions

What can be done to tackle all these problems? The best approach seems to be a multi-stakeholder approach, in which all parties take responsibility. First and foremost, there are the ICT-driven methods of cryptography, data protection and anonymisation. There is also the educational angle, with data sense and advertising literacy being taught from the youngest years [website ik beslis] and also the older generations being briefed. Finally, there are legal strands, with four other components besides the EU's General Data Protection Regulation (GDPR): transparency of profiling; purpose limitation and the data controller's legitimate interest; the presumption of innocence by police and judicial authorities; and objectivity of private-law liability in cases of illegitimate processing.

Cryptography, back doors, mass surveillance, security

'Cryptography' is the branch of science that deals with protecting information, including digital information. Seen historically, the emphasis in cryptography was on protecting *communications in transit*, such that only the intended recipient could read the information. When information becomes digital, the importance of protecting data against alteration (safeguarding data integrity) grows, as does the need to identify sender and recipient properly. For the purposes of protecting people's private lives, there is an increasing drive to protect communications metadata, too: concealing the sender's and recipient's identities and locations from third parties. The increasing saturation of computer usage has caused more attention to be paid to protection of information *at rest* on PCs, tablet devices, smartphones or cloud storage. A recent development is the desire to protect information even while it undergoes *processing*. Picture the situation where a user wants to use his own healthcare data to calculate his risk of contracting certain conditions. A user in such a situation will be keen to keep his data confidential, while the service provider to which he consigns his data will most likely want to protect its calculation methods or algorithms. At first sight, this might seem an irresolvable conflict, but dedicated cryptographic algorithms do permit the uploading of sensitive data to the cloud in enciphered form so that the service provider can run its calculations on them. The enciphered result can then be downloaded and deciphered.

Since we are inexorably becoming a world of Big Data, the flow of information being sent, stored and processed is only growing and growing. This being so, personal information is being spread increasingly widely, with enhanced risks of

misuse by other users, companies and government bodies ('data pollution'). The best way of protecting information effectively and keeping data pollution in check is to use cryptography. Yet there are limits to the capabilities of cryptography. If the analysis is done on enciphered data, data will remain well protected but a chance remains that the analysis will be inherently discriminatory or will infringe privacy. To cover this risk, other solutions are called for, ones discussed earlier in this paper. Besides, in certain applications, such as on social networks, sharing information with peers is the very essence; here, cryptography can help by restricting personal information to only those peers selected by the sharer and by screening them off from being exploited by major service providers such as the social network platform (e.g. Facebook) or the network operator (e.g. Vodafone). With cryptography, data protection becomes a matter of guarding digital keys. If Alice wishes to send Bob a secret, she will first have to agree a secret key with him. Using the key, Alice will then encipher her message into a cipher text, which she then sends to Bob. Availing himself of the same key, Bob will work the cipher text back into the plain message and will also seek to assure himself that this is a genuine message from Alice. For anyone lacking the key, the message is unreadable. In massive networks, however, such as the Internet, it is not feasible for every service provider and service user to pre-agree between sender and recipient a secret key of this type before every transmission. Here, 'public key cryptography' provides a solution: the encryption key is public, and the only element that has to remain secret is the decryption key that makes the information legible again. One could compare this arrangement to a postbox: anyone can deposit a letter through the slot, but only the keyholder can empty the box.

Until the 1980s, cryptography remained the preserve of the military, government and banks. Nowadays, though, it is a mass technology: more than 30 billion devices are equipped with cryptography. The largest application is still in bank cards, but cryptography is now also found in all mobile telephones and WiFi networks, browsers, smartphones, digital ID cards, passports, entry passes, car keys, DVD and BluRay players, and chat apps such as WhatsApp, iMessage, etc. Since cryptography also has military applications, it falls under dual use legislation, such that the use, import and export of cryptography are government-regulated in many countries. Until the late 1980s, cryptography was found almost exclusively in expensive, bulky hardware sets, making it more straightforward to check for its use. Yet the arrival of the World Wide Web enabled a whole gamut of new applications that called for cryptography: online shopping, online banking, music downloads, app downloads, etc. Moreover, because computers' number-crunching capacity grew relentlessly, it became feasible to embed cryptographic features within software, making the regulation of cryptography a much taller order.

There is a multitude of reasons why government bodies would much prefer to keep regulating cryptography. The first is military: cryptography is strategic knowledge and ought not to fall into the wrong hands. The second is the interests

of police forces, which heretofore have always relied upon their ability to open letters and tap telephone calls, with judicial warrant. Now, if both parties choose to encrypt their communications, that option is denied them. Third, the intelligence agencies have an interest in not seeing cryptography become generalised, since their business is collecting information on foreign countries and certain domestic groups.

In present-day society, it has become hard to conceive that governments might outlaw cryptography. For one thing, every country has a strategic interest in its citizens', companies' and public bodies' ability to protect themselves adequately against criminals and state actors. Moreover, a growing number of sectors critical to the economy are continuing their digitisation agenda. Industries such as power generation, transport and healthcare have a very pronounced need for secure IT and communications; nothing but cryptography makes that possible.

Government bodies have come up with a range of strategies to manage this internal conflict of interest. The first solution thought up was either to strictly regulate or entirely to forbid cryptography, but — as noted above — this proved to be an illusion, because the mass transition to encrypted software was a fait accompli. Next, governments sought to cripple cryptography (such as by setting short maximum lengths for encryption keys), so that government agencies could break into communications but supposedly no-one else could. This is a problematic solution to adopt, as it exposes citizens and companies to foreign state attack and indeed to organised crime. What is more, technology develops at such speed that what is today a halfway secure system might well be hopelessly outdated within a decade. Despite the problems with this solution, governments have carried on with it regardless. EU law stipulates that cryptographic hardware may only be exported if its keys consist of 56 bits or fewer. This decision was made over twenty years ago; nowadays, it takes a central signals intelligence agency a couple of seconds to retrieve the cleartext, and a group of academic researchers a couple of hours.

The next strategy which governments adopted was putting backdoors in cryptographic systems for government agencies to use. One example was the early-1990s Clipper Chip in the United States, which reportedly provided the US Government with access to enciphered communications. This proposal was withdrawn after lobbying by industry and academia, but that is not to say that government agencies relinquished their systemic control. For instance, we know from the Snowden papers that the US signals intelligence agency, the National Security Agency (NSA), built a backdoor into an industry-standard software package for generating encryption keys. Governments can also oblige communications service providers to yield the keys they hold to government bodies, using what is known as a 'security letter'. The Investigatory Powers Act was passed by the British Parliament in late 2016, extending previous statutes. Section 217 of the

new Act empowers the UK Government to force the providers of goods or services to build a backdoor access into their systems.

Experts are in agreement that building in such backdoors is a highly dangerous affair, for there is a risk that another state or a criminal organisation will discover and use it, rendering the whole system in question insecure. In the case of Juniper routers, this risk has already become reality [Juniper routers trapdoor]. Moreover, where such backdoors exist, mass interception — the storing and analysing of everyone's data on the system — remains a possibility, and that is in contravention of the European Convention on Human Rights and is a gateway to large-scale abuses. Experts even cast doubt on whether mass interception is an effective way of combating organised crime and terrorism. The lack of transparency on these issues has prevented public debate.

Finally, another possible solution is using malware, whereby a target's computer or telephone is infected so as to gain access to the data required for investigations. This method is often called 'remote hacking', and it enables the capture of data either before encryption or after decryption. The competences for its use are set out by law. In principle, this is a superior solution, although it brings with it the problem that oversight of how it is used is very difficult to accomplish. In addition, there are proliferation risks: if others come across the malware, they can quite easily adapt it and launch it at other targets, including government bodies. We conclude that regulating or deliberately hamstringing cryptography is not a good solution. Robust cryptographic solutions are a necessity to protect society effectively. The alternative approach — remote hacking of IT systems — requires legal safeguards that must be vigorously enforced.

Legal approaches to data protection with Big Data

The General Data Protection Regulation (GDPR) was published as Regulation (EU) 2016/679 in the Official Journal of the EU (Document 32016R0679) on 4 May 2016 [GDPR 2016]. Its provisions entered into force on 24 May 2016 and it becomes applicable law on 25 May 2018. The Regulation reconciles two aims: better protection of individuals' personal data and greater scope for business to exploit the digital single market through the direct effect of a single regulatory framework for all member states. Supplementary national legislation is permitted only insofar as discretionary room has been allowed and/or reserved for that purpose.

Most distinctive about the GDPR is that it is more tailored to the age of Big Data than the EU Directive it is replacing. This has everything to do with the uniformity of its text, which applies throughout EU territory, allowing transnational companies to incorporate the provisions of the new Regulation more easily into their business processes. Moreover, the maximum fines are now very high (comparable with

those levied for breaches of competition law), namely 4% of global turnover, thus creating the right incentive structure for stakeholders to conduct their business legally. Finally, private-law liability for illegal data processing is now better regulated, not least by the stipulation that data subjects can bring class actions by mandating their infringement claims to an NGO.

In sum, while the Regulation provides substantially the same protection as the current Directive [EU Data Protection Directive 1995], which continues to be the applicable law until May 2018, it is more effective and more attuned to the mass scale and complexity of Big Data and machine learning. Improvements include:

1. citizens' greater ease of access to their own personal data;
2. more exacting requirements for the obtaining of consent;
3. the requirement for withdrawal of consent to be as easy as giving it;
4. the requirement for easily-comprehensible information to be provided on what will be done with people's data;
5. the right to have one's personal data deleted, in certain circumstances;
6. the right to user-friendly transfer ('portability') from one responsible party to another or to oneself;
7. the obligation to build protection into the design of the IT systems in question, also known as 'data protection by design';
8. the obligation to set all the standard settings of these systems such that only the necessary processing occurs ('data protection by default');
9. the setting of limits to 'profiling' (defined as the automated processing of personal data for the intended purpose of evaluating personal characteristics).

We now go on to discuss a few of the particularly relevant domains of European data protection law, to the extent that they concern Big Data and machine learning.

1. Transparency of profiling

Classically, in regards to privacy, the emphasis has lain upon the safeguards required in the collection of personal data. With Big Data, this emphasis has shifted to the analysis and use of personal data [WRR 2016]. This is particularly the case where automated decisions are made which are of major significance to the person in question. In such cases, it is crucial — both vis-à-vis the data subject and vis-à-vis the regulatory authority (Belgium: *Commission de la protection de la vie privée/Commissie voor de bescherming van de persoonlijke levenssfeer*; the Netherlands: *Autoriteit Persoonsgegevens*; UK: the Information Commissioner's Office) — that transparency be provided on three matters. First, there must be demonstrable clarity that the decision was taken on the grounds of these kinds of analysis (often called 'profiling'), what the exact objectives are of the processing, and who bears responsibility. Second, the underlying reasoning must be set out in plain language. Third, the processing party must state what consequences

it foresees that the profiling will have. As regards the underlying reasoning, it is crucial that information be provided on the methodological choices made in the analysis. This will include, for example, details of the choice of algorithms (including training algorithms) and model structure used, any parameters applied, the variables taken into account, and the data type used to train the model. This is intended to make the processing of data as reproducible as possible (this is the obligation of transparency and the obligation to inform). In best practice, the decision algorithms which the process gives rise to will (together with the specific input-output relationship arising from the methodology used) also be published and set out in comprehensible terms, so that the data subjects have sufficient information to understand upon what logic decisions about them are based. It is also important that the accuracy and performance of any models employed (on independent test data) are notified as exactly as possible, thus disclosing the profiling's margin of error. These requirements are intended to enable data subjects to mount an appeal, should it prove necessary, against decisions taken about them where Big Data analyses played a role. A key consideration here is that one cannot reasonably expect individuals to have to dig all this information out for themselves and to enter into dialogue with the entity that has been making decisions about them. There is an obvious task here for oversight authorities, both those which cover data processing and the consumer watchdogs, or for civil society players.

This transparency must take account of the basic premise that there ought not to be any needless infringements of the rights or freedoms of others, including trade secrets and intellectual property, and particularly not of the copyright protecting the software [para. 63 of Preamble, GDPR 2016]. In fact, this is less of a problem than might be imagined, because what the party affected is entitled to is *significant* information on the profiles and algorithms used. It is precisely not the intention of this framework to encumber the data subject with technical gobbledegook that is of no use to him. This does not detract, however, from the consideration that the aforementioned debate on the methodological integrity of profiling and research into possibly illegal bias is one that the individual concerned must be able to participate in. This might, not infrequently, lead to difficult trade-offs on trade secrets and intellectual property rights. At any rate, under the new GDPR, it will be the national Data Protection Authorities which receive competence to require documentation of the working of these IT systems and, by way of verification measure, also to demand access to the servers and sight of the method used.

2. Purpose limitation and the data controller's legitimate interest

A key prerequisite for personal data processing, including in cases of Big Data, is that it be made clear in advance (1) for what specific purpose the data are going to be processed and (2) that the processing really will be restricted to that aim. This is known in European law as the 'finality principle', alias the 'requirement of

purpose limitation'. Given the complexity of data streams, this requirement is one that many responsible parties find highly tricky to meet. With Big Data, the motto often seems to be: collect as much data as possible from all over the place first, and work out what we can do with it later. As described above, this can easily degenerate into a habit of working with low-hanging fruit, which is not conducive to the reliability of findings. Purpose limitation, then, is as much a requirement for methodologically-sound data processing as it is for legitimate processing. An example to illustrate this would be the data streams which are generated when using a smartphone: they go to the hardware provider, the operating system vendor, the firmware owners, the browser programmer, and to countless app owners. The Internet of Things (IoT) adds many more data streams to the mix, as do smart energy infrastructure and robotics; among these extra streams, there is a preponderance of machine-readable behavioural data. Nevertheless, the purpose must — very much including when it involves Big Data — remain lawful at all times and must always be properly specific and, indeed, explicit (i.e. known). Processing for any other purpose than stated may only occur where the purpose is congruent with the original purpose, so that it remains reasonably foreseeable to the data subjects what their data might be used for.

Big Data often entails the reuse of data first possessed by other entities and/or held for other purposes. For instance, energy consumption data are harnessed to detect benefit (welfare) fraud, or locational data required to provide some service are rehashed to calculate health risks. In such cases, the overriding principle is that fresh legal justification is needed to proceed. This certainly need not always involve seeking permission; a legal obligation might suffice (for the detection of fraud), or the legitimate interest of the responsible party (such as economic interest). For instance, many companies provide free services (search engines, social networks, etc.) which they are only able to maintain and improve due to advertising revenue which is contingent upon the processing of users' behavioural data.

Consequently, where Big Data is used, it seems that seeking permission would often be unfeasible and unfair. For individual consumers, it is hardly feasible at all to deduce which patterns and linkages are being derived from their data and what the knock-on effect might be (the implications could involve refusal of credit lines, increased insurance premiums and denial of job or training opportunities). For this reason, indicating mere agreement will usually not meet the requirement that there should be unambiguous informed consent. Accordingly, the jurisprudence of the Court of Justice of the European Union, which is supreme in such matters, presumes that where a company's business model is predicated upon the processing of personal data, the applicable basis will be 'the legitimate interest of the responsible party'. This is a legal test which always entails a balancing of that interest against the data subjects' rights and liberties, which certainly must be honoured. In this balancing, much will depend on the technical and

organisational measures which the responsible parties are taking — to protect the data from unauthorised usage and from hacking; to prevent unwanted targeting; and to achieve the transparency of profiling discussed above. These may involve anonymisation and simple ways of reviewing data or of stopping their processing (the GDPR lays down that withdrawal of consent must be just as straightforward as providing consent in the first place). As noted already, under the GDPR, these kinds of measures are obligatory as a means of data protection by design and by default.

3. Presumption of innocence by police and judicial authorities

It is not only in the private sector that Big Data and machine learning are used. Police forces and judicial authorities have begun employing techniques such as 'crime mapping' to detect the hot spots where particular problems can be expected (be these public order offences, criminal acts, disasters or other). Authorities are also investing in software applications that award convicts a score which represents their estimated risk of reoffending; these scores may be taken into account when determining the length and type of sentence that will be sought. Here, too, we have to do with the processing of personal data, but in this case the legal framework is different and transparency is not organised along the same lines, because here confidentiality is often essential for the task at hand. All these elements call the presumption of innocence into question, precisely because people's general tendency to become systematic in their monitoring of groups of people based on what Big Data analysis and machine learning is telling them does not cease to apply in this domain. Such monitoring could then, for instance, be based on a potential suspicion of potential future criminal acts or — even more broadly — on suspicions of future undesirable behaviour [Hildebrandt 2016].

Since the Snowden 'revelations', it has been evident that intelligence and security services are also among those who are intensively involved in collecting and analysing a huge range of communication and behavioural data to be able to predict intended terrorist attacks accurately. As things currently stand, it is unclear whether or not the multitude of watch lists of potentially dangerous individuals has actually helped prevent attacks. One of the reasons for this lack of clarity is that the phenomena in question are of insufficiently regular nature to enable the composition of a reliable profile.

This throws up at least three controversial issues: (1) the public-private partnerships that facilitate the batch-forwarding of Big Data from the private sector to prosecutorial authorities and security agencies; (2) the crypto-war, discussed above, that is being waged around the question of whether national and international security is really served by the 'backdoor' accesses that government bodies have to their own and other countries' citizens' data in storage and in transit; (3) the decreasing ability of citizens to know in advance what kind of

conduct on their part might trigger further investigation. In addition, it can prove a real problem to challenge suspicions, clues and presumptions against one, because the profiles involved remain secret.

As regards the security services, it will be clear that they are a national competence and thus exempt from the operation of European Union data protection law, although their work is covered by the European Convention on Human Rights. This being so, the European Court of Human Rights regularly makes pronouncements on the strict conditions under which the security services may undertake profiling.

4. Private-law liability for illicit data processing

In addition to a public-law liability of citizens towards government bodies, there is a private-law liability of citizens towards each other. If EU member states' individual data protection authorities (perhaps obliged to take such action by the Court of Justice of the European Union) are able to devote sufficient budgeting, technical expertise and staff to the matter, we may expect that the high amounts of fines and the Union-wide scope of application of the Regulation will provide reasonably effective protection against abuses. However, the Regulation also calls for effective private-law liability for infringements of the legal provisions. These can be understood as failures to meet obligations, such as by omitting to report a data leak, but also cover failures to meet security requirements, or simply illicit processing of data (in the absence of a legitimate reason to do so or overstepping the bounds of the stated purpose of processing). For such private-law liability, it is required that there be demonstrable material or immaterial damages, which in turn must — likewise demonstrably — have been caused by the pertinent infringement(s) of rights or of legal obligations. Yet the Regulation speaks of an 'effective judicial remedy', one which must offer real protection when the rights of data subjects are infringed. Judges around the world seem to be more prepared to award damages in cases of immaterial damage, such as uncertainty as to possible identity fraud or reputational damage. As noted above, the Regulation additionally stipulates that EU member states enact a means for affected individuals to hand their claims over to NGOs, allowing for class-action lawsuits.

We may expect that it will at some point be considered whether 'strict liability' is a good solution to this issue, so that the affected party is not saddled with the impossible evidentiary standard of having to prove that the unjust data processing really did cause the damages. This might, for example, allow for a reversal of the burden of proof. It is improbable that any infringement will in and of itself be found to be tantamount to damages. Thus, the protection offered by this type of liability will remain restricted to issues where it can be shown plausibly that material or immaterial damages have been incurred. In addition, liability law is a national competence and so reserved to member states' parliaments and courts. As things stand, 'European private law' is still an admixture of comparative law

and germane EU directives, such as that on product liability. We can expect, however, that it will continue to develop as a field of law, now that it has become crucial to the smooth working of the single market and to the effective protection of consumers who use transnational companies' services.

Anonymity and anonymisation

Anonymity plays a central role in issues of privacy and Big Data. It is a lively term, particularly at law. Personal data are protected for privacy reasons, but if they are anonymised, then they cease to be personal data and thus lose their legal protection.

How so? The explanation is straightforward: anonymised data no longer refer to individuals. Nobody's privacy is affected by the statement, 'There's somebody with a beard.' Nobody will be able to work out who is being talked about, and so saying such a thing is regarded as privacy-irrelevant. This ceases to be the case if the set of people potentially being talked about consists of few bearded men and many clean-shaven men, for then it will quickly be apparent which bearded man is meant. So anonymisation of personal data appears highly conducive to privacy, as long as anonymised data cannot be used in breach of anyone's privacy. Our knowledge economy practices anonymisation for the purposes of further use and sharing of available data. Often, companies and public bodies can achieve their purposes with anonymised data, and personal data prove not to be really needed. Big Data operations are, then, often achievable using anonymised data. The operator of a bike share scheme needs to know where in his city, at which bike racks, there are too many or too few bikes. Knowing who is using each bike, or where a given bike is, is not necessary to that process. Working with personal data is always interesting, of course, and in some cases necessary. Incontrovertibly, the operator of a bike scheme who knows which user is cycling where, on which bike, has a knowledge advantage; the question is whether he has a need to know. By law, personal data may be processed, but only where it is a real necessity and the task cannot be done in any other way. Naturally, this is a hotly-contested issue.

A separate and more fundamental debate is that about anonymisation itself. In an increasing number of cases, technology enables us to penetrate the anonymity of data and uncover the identity of the person signified by such a statement as 'There's somebody with a beard'. Moreover, data protection legislation lays down that data which carry a reasonable probability of (re-)identification are by definition personal data. As soon as it is possible, such as by combining data in order to associate a data point with an identifiable individual, we have to do with personal data. Thus, both dynamic and static IP addresses are regarded as personal data insofar as they can reasonably be linked to the person of the user. Legally speaking, then, using most anonymisation techniques does not constitute anonymisation proper but rather what the GDPR calls 'pseudonymisation', for reasons such as the

availability of a decryption key. (Even if such a key is managed by other entities, the law will often not brook the notion of the data being anonymised, as long as the responsible party can reasonably be expected to be in a position to obtain access to the key, such as by court order.) Legally, even pseudonymised data are by definition personal data. Nevertheless, pseudonymisation can in some cases be a good way of meeting the requirements of the GDPR [GDPR 2016]. Where this is so, pseudonymisation is a form of data protection by design. In such cases, the supplementary data that would reveal the identity of the pseudonymous data subjects must be stored separately, in both technical and organisational terms.

It has also now become possible in principle to furnish every data user with an individual key so that the data obtained by various parties cannot be accrued. This is of particularly great importance in medical and study-related data, where professionals have their own responsibility to ensure the confidentiality of data but where there is also a pressing need for access to Big Data for medical research or to develop data-driven learning environments [Verheul et al. 2016]. However, the privacy protection offered by pseudonymisation is limited [de Montjoye 2013].

Online platforms and shared responsibility

The Internet, and by extension online platforms, are increasingly penetrating the very fibres of society and the economy. 'Online platforms' have been called the 'technological, economic and socio-cultural infrastructure for facilitating and organising online social and economic traffic between users and suppliers, with data (including user data) being the fuel' [van Dijck et al. 2016, translation]. This may involve information and communication via social media (e.g. Facebook), but equally encompasses a great range of types of service provision using the sharing economy, such as transport (e.g. Uber), hospitality (e.g. Airbnb), education (e.g. Coursera), healthcare (e.g. PatientsLikeMe) and many other industries besides. User data and user metadata are a crucial element of these platforms' earning model, so that increasingly major interests are at stake. This issue boils down to striking a balance between the opportunities that data offer society and the corporate world, and the concomitant risks to citizens' fundamental rights (privacy, data protection, etc.) and to (European) public values (equal treatment, inclusiveness, diversity, etc.).

To achieve that balance, it is necessary that all relevant public and private parties be brought together and enter into dialogue. This will entail arranging a cooperative or collaborative scheme of responsibilities between policymakers, the industry, users and civil society organisations. Current legislation has a tendency to accord responsibility to a single central player (e.g. the data controller, editor or service provider), in order to keep oversight and liability efficiently-organised. However, that does not resolve all the issues. We thus see the GDPR using the language of 'joint controllers', and this term also provides for processors of personal data

acting on commission to be held liable for unlawful processing. Comparable issues of shared responsibility are in fact also at play in the domain of freedom of expression, and the problems that have been cropping up in that domain are not necessarily resolvable via the data-protection route. In the context of the Internet and online platforms, we see that various parties can and must shoulder their responsibilities, not only as regards privacy and data protection but also on matters such as controversial content, hate speech, diversity and transparency. Platform owners provide the infrastructure through which users encounter each other and share information; it is the users who choose — sometimes in huge numbers — to share particular content and data; it is government that sets the policy and legal framework for data collection, storage, processing and protection.

In order for this kind of collaborative responsibility to be well organised, society needs a species of multi-stakeholder consultative organ, one that will factor in ethical, legal and social aspects (referred to by the acronym ELSA). First and foremost, this requires that all stakeholders (including social media and online platforms) accept that they have a responsibility in the matter, one which will vary according to context. Thereafter, the actors involved need to arrive at a common view of how privacy is defined and treated. In this process, they could learn valuable lessons from the long tradition of technology assessment, particularly from the variant known as 'constructive technology assessment'. Finally, each stakeholder ought to resolve to make concrete efforts in practice. For social media and online platforms, for instance, this would entail building functions into the system to give users proper transparency and to allow them control over their data. Only in this manner can solutions regarding future digital media and data technologies be found which enjoy widespread social traction. It will be apparent that this approach foresees there being an incentivisation structure within which stakeholders are obliged to shoulder their responsibilities; this is particularly so given the transnational and global nature of the major players. High fines and private-law liability, as laid down by the GDPR, are a first effort towards creating a level playing field that will enable companies and government bodies to assume their responsibilities without being elbowed out of the market.

User data-literacy

Often, legal and technical studies on privacy, data and the Internet discuss the user from an expert perspective, and have too little to say to users (or do not let users speak) from the perspective of users' everyday lives. This sometimes gives rise to a one-sided and sketchy concept of 'the' user, whereas real interaction with the digital environment is far richer and more complex as regards the spectrum of awareness, attitudes, skills and behaviour. Account must be taken, particularly where young people are concerned, of the various domains of their lives: home, school, friendship groups (e.g. youth associations). Only by researching users empirically in such a nuanced manner is it possible to achieve user empowerment

and avoid user disempowerment. 'Empowerment' can be defined as the process of fortifying users/citizens/consumers so that they can have a good grip on their own situation and on their environment, by obtaining control, honing their critical awareness and fostering participation.

Given the importance of empowerment to socio-technological evolutions in the Internet and Big Data world, the well-established notion of media-savvy now needs to be supplemented with a new concept, data literacy. This term connotes the degree of insight, control and confidence that people have in how their data (including personal data) are collected, stored, processed and used (or reused). The data referred to by this term are those shared of one's own volition, which encompasses both data shared deliberately (photos, interests, address details, etc.) and shared unawares (surfing history, location, cookies, etc.), but also — no less so — derived data (credit scores, profiling, emotional states, etc.). To boost data literacy, one ought to bear in mind not only digital media (such as the Web and social media) but also a burgeoning number of other technologies that generate and process data (wearable tech, drones, the Internet of Things, smart grids, etc.). One of the greatest challenges in this regard is how to render comprehensible and meaningful the complexity of technological aspects that impact upon data, metadata and privacy, such as algorithms, APIs, machine learning, default privacy settings, and AB testing.

3. Analysis of privacy from some relevant case studies

The following case studies are considered with an eye on the future and the whole system as the average user experiences it. They each go into detail on the aspects of Big Data and learning-enabled systems that are fundamentally changing the privacy scene. We will in each case refer to the considerations that are weighed up between user friendliness, functionality, privacy protection, cost, quality and comprehensiveness.

Case study 1: A family's digital life

Saturation of the Internet and of web-enabled devices is extremely high in Flanders [Digimeter 2016]. Here, we sketch a fictional but realistic situation from the daily life of a family consisting of a mother, father and two daughters aged nine and thirteen. Many aspects of the terminological and problem scoping that this paper deals with are expressed in the case study.

It's the weekend. Everyone has got up and breakfast is on the table. The family has a no tablets or smartphones rule during meals. After clearing away breakfast, everyone in the household gets busy with their smartphone or tablet. Mum is planning the summer holiday in France. She's keen to have a look at a few more



<http://ShoeBoxBlog.com>

campsite websites. Dad is looking forward to a cycling trip with some mates, and is curious what route they have planned — and he hopes to sneak a peek at where his mates have been training this week. The younger daughter is a YouTube addict. Tired by her week at school, she just feels like watching some films on her tablet to relax. The elder daughter, who has her own smartphone, retires to her bedroom, because her friends are all awake and the messages are in full flow on Instagram.

Mum starts surfing the Web, using Google to find campsites. She types 'small campsite southern France'. The first thing to appear is a privacy reminder by Google. She scans it quickly; it is a lengthy text. She registers that it has to do with 'retaining data, YouTube, other services, improving user experience, personalised search results...'. She also reads that all settings can be changed, but decides that she doesn't have time for all this. Clicking 'I accept', she moves on: the privacy paradox in practice. Mum is immediately interested by the first campsites in her

search results. She also likes the fact that a map pops up immediately showing the campsite locations and reviews by other tourists. She clicks on the first few hits — which, as we have seen, are search engine advertising. She notices that on one of the campsite websites she opens, she sees an advert for the shoes she was browsing online earlier in the week. She's not sure whether she'll order them; she is creeped out by the fact that these shoes keep popping up wherever she goes. She resolves that such is life and that at least she's seeing things she's actually interested in. This is the effect of retargeting. Unfortunately, it often means that she keeps getting this kind of advert after having ordered the item, which is a bit annoying.

Dad is relishing the prospect of his weekly spin on the bike. He's bought a new heart rate monitor and wants to install it quickly before he sets off. He downloads the app, which asks him to link his heart rate data to Strava, the app he uses to keep track of his cycling runs and achievements. 'Nice,' he thinks to himself, 'that'll let me compare my heart rate with my mates. Which of us is fittest?' He receives a prompt: will he also allow the data to be shared publicly? 'Hmm, maybe not just yet. But if I find that my heart is in great shape, then I'm all up for that.' He has a look at his friends' profiles on the app. It's obvious that they've been training hard this week; one even went cycling in the Flemish Ardennes. He hops across to the website where he buys his cycling gear and sees that he's getting a recommendation to buy a new type of tyre. It seems they're on special offer. He also looks at the cycling gloves he's been wanting to get for a while: 'Only a few left in stock!' He decides to order the tyres and gloves without delay; nudging at work.

The younger daughter installs herself in the living room with her tablet. She goes straight to YouTube to see whether the people she follows have uploaded any new videos. One of them is a vlogger, a funny lad who has a massive followership and who comes out with a new film about his life twice a week. Today, he's talking about a trip he made to a new theme park. 'That looks super cool!', she thinks, and decides to ask her parents later whether they can go there together this summer. So she has been charmed by a social influencer, who might have been offered a free trip to the theme park in question. After that video, she plays a free game (Mum and Dad are not keen on paying for tablet games). It's a race; you have to knock over as many soft drink bottles as you can. 'Really fun!' She shares her score with her friends and sends an invitation to a few other friends to join in the game. 'Where do they sell that drink?', she wonders. She has just played an advergame, made and paid for by a trading company. It is not evident to younger children that this is advertising.

Withdrawing to her bedroom, the elder daughter opens Instagram. Her friends are posting funny photos and film clips thick and fast, and leaving comments on what they have watched. Scrolling through the news reports, she sees one for a

trainers chain which she popped into with her friends yesterday. While they were in the store, they had pointed out to each other which pairs of trainers they'd like to buy. The Instagram message has a link to take surfers straight to their web shop. She decides she'll show Mum this shortly, as Mum promised that she'd be getting a new pair of trainers. The message placed by the trainers chain is an instance of geolocation-based advertising. Because Instagram, a social media platform, knows from smartphones who has been where, Instagram can tout that information to advertisers to target potential customers. All of a sudden, the elder daughter sees a photo of herself flash past, one from several years ago. It's a photo she's deeply ashamed of. Evidently, someone once uploaded that photo to a website, and this means you can find it by Googling her name. She has no idea how the photo could be deleted, but hopes that this photo will not keep dogging her. 'Everyone does silly things sometimes, don't they?' In other words, the right to be forgotten is quite fundamental. It's now noon and the family sits down to eat again. They chat about summer holiday plans, the forthcoming bike trip, a new theme park and cool trainers. Not a word about a certain photograph.

Case study 2: Big Data in passenger profiling

After the 9/11 attacks in New York, it only stood to reason for authorities to take a fresh look at how air traffic was organised. The United States, more particularly the Department of Homeland Security (DHS), came up with the idea of blacklisting passengers who ought never again to board an aircraft. All airlines use the same Passenger Name Records (PNR) system in ticketing and flight manifests. Airlines using US airspace are obliged to e-mail these data before departure: no data, no clearance. The DHS seeks out suspicious passengers among this mountain of data. To draw up the blacklists, the DHS used passenger profiling, not just the names of already-known suspects. A suspicious flight purchase (bought using cash, ticket bought last-minute, etc.) will trigger an addition to the list of potential suspects. This is Big Data through and through.

This system is now used all over the world, including in Europe. Sadly, there are no real overviews in figures available, so whether it works is unknown. We do, however, know that people have found themselves on the blacklist who have no connection to terrorism but who were profiled by chance. One of those to whom this has happened is the singer Cat Stevens. The government agencies responsible simply call this tough luck for those caught up in it. There are no proper arrangements for getting oneself extricated from passenger blacklists. Moreover, the system is not what one could call friendly; one is only told that one is on the no-fly list when one turns up at the airport with bags packed.

How should this Big Data application be thought of? In the absence of any impact reporting, we can have no debate on the issue, but quite apart from that, it is

striking how little effort is being made to make this Big Data project appealing to citizens. A first step towards that end would be admitting honestly that the system is to some extent based on guesstimates, and that consequently errors are quite possible. A second step would be having a good, long think in specific terms about passengers unjustly prevented from boarding, and to set up an all but automatic compensation scheme for erroneous cases. Such a system already exists for flight cancellations or major delays. This would be enough to alleviate some of the frustration. A third step would be making the system truly user-friendly. Let blacklisted people be sent an e-mail when they have bought their ticket to inform them that they will not be allowed to fly. Why not provide that transparency? Is this a matter of national security? Blacklisted citizens will find out anyway that they are on the list, but in more trying circumstances (at the airport). It would be a different matter if they were to be told when purchasing their ticket that they are on the no-fly list. That would enable the subjects to challenge the decision and hire a lawyer. It would at least allow affected people to cancel appointments abroad. By treating Big Data properly, passenger rights can be honoured, even in the domain of counter-terrorism. The inevitable mistakes made by algorithms will then be more acceptable [DH 2011].

Case study 3: The Internet of Things in the context of smart cities

Two tendencies made the Internet of Things (IoT) possible: modern electronics, which enables cheap, tiny sensors and microprocessors, and the ubiquitous networking that allows everything to be connected to everything else. Through their sensors, items can collect data and interact with each other and the wider world. This offers extensive new types of applications, ranging from health (where personal parameters such as blood pressure, temperature and other vital signs can be tracked) via smart houses and buildings (where everything from lights to heating and doors can be monitored and controlled) through to smart cities (where mobility and transport but also water consumption, waste generation and energy usage can be monitored and controlled).

The sensors typically used in the IoT supply a continuous stream of data (a great volume and huge variety of data at high frequency) that can be analysed. These data can also be combined with other data, such as social media data (Facebook, Twitter, etc.). This means that the IoT constitutes a fantastically good source for Big Data applications. Various algorithms will seek to detect patterns and correlations in the flood of data, in order to produce contextually-enriched information that can be used to draw conclusions. Even supposedly anonymised data can be harnessed to pick a unique profile out of the mass of data, thereby enabling the individuals to whom the data appertain to be identified after all.

This case study focuses on the IoT in 'smart cities'. This term is often used to denote cities which foster innovation and creativity, whether using ICT or otherwise, as well as sustainability, entrepreneurship and ICT education. The term is also used

to refer to cities which have a citizen-centred model of urban development, with consideration given to social innovation and cohesion, equality of opportunity and participatory city governance.

Why this interest in smart cities? Globally, we see that a growing number of people live in cities: in 2014, over 50% of the world's population was already urban. The United Nations forecasts that this percentage will have grown to 66% in 2050 [UN 2014]. To keep cities habitable, great efforts are being made to obtain better services and living conditions. Smart cities are intended to be a contribution to this end, by such means as enabling better, cheaper, more efficient and more sustainable services [Ballon 2016]. A few examples of how the IoT works will serve to illustrate this.

- Smart mobility: Sensors measure traffic streams in the city, of pedestrians, cyclists and cars as well as of public transport. This allows people to know, for instance, where there are jams or which bus lines need extra drivers. It also assists motorists looking for a place to park. Some Dutch cities are considering programming the traffic lights to give cyclists extra priority in rainy weather. Car-sharing and bike-sharing is made easier thanks to smart apps linked to the network to which the vehicles to be pooled are connected. The city's entire transport system is monitored, controlled and where necessary adjusted as an overall whole.
- Heightened security with quicker, more targeted emergency interventions: A network of smart cameras and other sensors, combined with coordinated assistance services and manned control centres, can enable quicker detection of and more efficient interventions on problems, natural disasters and acts of terrorism.
- Smart waste management: By placing sensors in rubbish containers and bins, waste removal services can adapt their rounds to real-time needs.
- Environmental monitoring: A network of sensors can monitor air pollution, noise nuisances, water purity and the water level of streams and rivers, in order to intervene more rapidly or take preventive measures.
- Smart lighting: Instead of having lighting on all night, sensor measurements will enable lighting to be on (at least, on at full power) only when there is movement.
- Electricity is used everywhere in cities and might also be generated within cities by solar panels or other small-scale green energy undertakings. Equipping every building with a smart meter is one aspect of 'smart grids', referring to energy networks. Since smart meters have intelligence, they can also take decisions [Belmans et al. 2016].
 - With a smart meter, any building can — besides consuming electricity from the mains — also deliver electricity generated in the building to the grid. This renders the entire energy network dynamic, in contrast to what is typically seen today.

- With a smart meter, residents can more easily optimise their energy usage, such as by setting heavy equipment like washing machines to work only during off-peak hours, or only charging their electric cars at those times.
- Electricity usage often sees surges in demand (such as at the start of the evening, when many people arrive home, cook and watch TV). At these peak times, setting the air conditioning or the washing machine to take a short break might be a smart decision that a smart meter could take.
- Overall, the expectations cherished for smart grids are that they will deliver a more reliable system (as it will be decentralised, obviating reliance on a single source, because some buildings in the city will also be able to generate electricity) and a more efficient system. In addition, smart grids promote renewable energy sources, and hence sustainability, because renewable energy can be more easily fed into the mains in this system.
- Smart cities contain smart houses, where the heating can be set individually for each room, depending on what activities are or are not happening in each; the heating in each room can also be adjusted remotely. What applies to heating applies equally to lighting or air conditioning. All manner of devices 'plugged into the Internet' can be remote-controlled.

Used as described in the above examples, sensors yield great quantities of information, some of which will be highly sensitive. This is because the habits of residents in a smart city are captured in various ways, with an unprecedented intensity. Their data are used to provide 'smart' services, and many data might be linked to individuals, even though pre-anonymised. The combination of data from various sources might enable individuals' life patterns to be documented in fairly minute detail! This has prompted an adjustment to the regulation around smart electricity meters in the Netherlands. Initially, it had been envisaged to make their installation obligatory in new buildings and during renovations and whenever an old meter was replaced. For privacy reasons, however, this obligation was rescinded; users are still free now to choose a traditional, 'dumb' meter. There have been increasing concerns about smart meter safety in the United Kingdom, with fears that smart meters, which tend to communicate wirelessly, could lead to unwelcome exposure to pulsed radio-frequency radiation in the home. The UK, too, has decided not to make smart meters compulsory. The EU [Smart metering in EU 2014] and the Flemish Government [Digitale meters 2017], however, continue to urge the roll-out of smart meters.

The privacy challenges associated with applications for smart cities are huge.

- Many of these challenges have to do with transparency and openness.
 - Every individual must be in a position to give consent efficiently and effectively for the use of his or her data. This applies not just to direct usage alone; data intended for one application often end up being reused

for other purposes. The same obtains with combined data: data ensuing from various smart city applications might be combined to yield fresh information.

- Every individual has the right to know what is done with data collated on him or her. There is nothing self-evident about this; the algorithms employed are often lacking in transparency. It increasingly seems as though all our information is disappearing into a black hole, from which decisions are unaccountably spewed out. When we speak of the use of data, we also mean reuse and the combination of multi-source data.
- The user must be free to choose devices and apps and be at liberty to combine them. Currently, device vendors, service providers and app developers are not particularly transparent and seek a monopoly position in order to have a hold over users.
- Security: Even when people consent to, for example, data from a smart-metered house being used by apps which ensure aims such as efficient energy consumption, they do not want their data to fall into the hands of thieves, who would be able to work out the household's routines exactly from them and thus be able to pick their moment to break in.
- The legal framework: The above norms are now laid down more rigorously than previously by the GDPR, which will have an immediate EU-wide impact and will enable effective enforcement. This is of crucial importance now that many ICT applications, including very varied apps offered in the interests of the smart cities agenda, can be downloaded from the Internet, a technology which knows no national borders.

Much remains to be done:

- *Makers of IoT devices* have to come up with technologies that respect privacy and allow transparency for the end user. They must make privacy by design a priority, with privacy being considered as one of the key requirements from the beginning of the design procedure and not 'tacked on' as an afterthought;
- *Service providers* must allow users to put together services of disparate provenance;
- *Algorithm programmers* must write their algorithms such as to safeguard user privacy;
- *App developers* must permit transparency and make efforts to achieve efficient, effective technologies, ones in which users can provide consent for the use of their data. They also need to become serious about app certification, so that users can be sure that their apps are secure. As regards both devices and apps, there is a great need for better, more granular ways of storing settings and preferences, and privacy must be the presumed choice in these settings (privacy by default);
- Finally, there is much *regulation* yet to be enacted, which must bring about standardisation (without putting a dampener on innovation) and must provide a basis for certification [Perera 2015], [DG Int. Pol. 2015], [UN 2014].

Case study 4: Distributed information versus central collection

There are many reasons why companies and government bodies are pushing to centralise more and more information in a data cloud — which might either be a public cloud offered by a major player such as Amazon, Google or Microsoft or be hosted on an organisation's own infrastructure. First, it offers a noticeable cost saving, largely due to economies of scale. It also allows great flexibility of response to fluctuations in data volume, without major outlay of funds. A second reason for using the cloud is that it makes data available to those who need them at any time and in any place. Examples include photos that people want to be able to view on smartphones and tablets, but which can also easily be shared. A third reason is the scope that the cloud gives for analysing data, known as data mining, which enables the discovery of new information and new unmet needs, but also enables targeted advertising — yielding mountains of income for Google and Facebook.

This approach is not without its risks. Month by month, the personal data of millions of users are leaked after hacking attacks or website breaches occasioned by negligent maintenance [website-breaches]. Clearly, we are not always managing to protect properly the huge quantities of sensitive data stored in the cloud, including the publicly-accessible cloud. In addition, this information also ends up in government hands: Edward Snowden revealed in June 2013 that the PRISM program gives the NSA access to data stored in the cloud of online players such as Microsoft, Facebook, Google and Apple [PRISM]. Finally, the centralised concentration of information entails a growing risk of secondary usage. Even if the user might have given consent for it, he will find it very difficult to understand the full implications.

In principle, it would also be possible to offer a great wealth of services based on local information ('fog computing' and 'edge computing'). For instance, a device should be able merely on the basis of data that the device itself holds to decide which adverts would interest the user, without sensitive information having to be uploaded to the cloud. Online social networks could also be developed using that methodology. Users have now grown to expect automatic syncing of data between their smartphones, tablets and laptops, and an automated back-up provision; this would be possible using a local server shared within a household or by a group of friends, which would keep all the information local. Should it nevertheless have to be stored in the cloud, this could be done in encrypted form; this would render the data unreadable to the cloud provider, thwarting any nefarious use.

For some applications, however, central processing is a must; an example would be complex analyses carried out on massive medical datasets. For such cases, we require strict rules on data security, access, logging, and auditing. Also, this is another domain where cryptography is a help: the past few years have seen

the skill of crunching enciphered data come on in leaps and bounds. This entails the server being able to search through enciphered data on condition that the search terms were known at the time of encipherment. It is also possible now to calculate statistical functions (averages, correlations, simple polynomial functions) efficiently on enciphered data. What results from this is a 'cipher text', decipherable only by keyholders. In its blandest form, fully homomorphic encipherment, this allows for any possible data operation to be carried out, but unfortunately it is still a very sluggish and costly business. Finally, cryptography allows various parties to collaborate to perform calculations across a network on all their pooled data, yet without leaking any information about the data. These solutions, something of a magic bullet, allow data to be protected robustly but at the same time to be milked usefully. They do still cost a great deal more than current conventional solutions in which all data end up in a single undifferentiated database, but there are some applications and datasets that for societal reasons oblige a distributed approach to data processing.

Case study 5: Connected and Autonomous Driving (CAD)²

A leaked European Commission Communication [Commission EU Parliament] on *Building a European Data Economy* proposed that Connected and Autonomous Driving (CAD) would be a suitable test case for investigating how data exchange within the EU's single market could be made as efficient, effective and responsible as possible. The main thing that makes CAD interesting is that it serves as an example of the aforementioned Internet of Things, a cyber-physical infrastructure connecting 'things' (in this case, vehicles) online and letting them talk to each other ('machine-to-machine communication') via the cloud. CAD also serves as a fine example of data streams that consist partly of personal and partly of non-personal data (the latter being road conditions, tyre tread and traffic volumes), although in the current state of the art, that distinction will not mean much, because even non-personal data can be coupled to an identifiable person at any time, causing them to qualify as personal data after all.

In Belgium as in other countries, there are calls to allow the police access after serious traffic accidents to data from the black box that is obligatory to be fitted in all new vehicles in the EU from 2018 onwards in order to allow automatic calls to the emergency services (eCall). The black box, officially an 'Event Data Recorder' (EDR), can in principle send and/or save all kinds of data. For the time being, the only functionality obliged by law is that the box should be able to place an emergency 112 (999) call, but a black box of this kind could also be used to record the data from the last five seconds prior to a crash, giving an indication of speed and braking, which traditionally have been worked out from skid marks

² This section is a reworking of [Hildebrandt 2017], a piece commissioned by the Dutch periodical *Ars Aequi*.³

on the road surface. With the introduction of Anti-lock Brake Systems (ABS) and Electronic Stability Programmes (ESP), such skid marks are visible less and less often, so an EDR would provide insight into the precursors to a crash [Post-Crash Voertuig Diagnose]. Without access to such data, reckless drivers might escape their just punishment, or careful ones might unjustly remain suspects.

Besides an EDR, new cars nowadays have a gamut of other systems that log and in some cases save and/or forward data. These include the on-board computer (the car's operating system, together with the aforementioned braking and stabilisation system); the navigation system; any driving style detectors (sensors); and even the mobile telephone that drivers bring on board with them. These constitute an almost inexhaustible supply of data. Some of the streams concern the state of the vehicle (fuel tank nearly empty, oil level below minimum, tyres too bald); others have more immediately to do with the conduct of driver and/or passengers — seatbelts done up or not done up; car left unlocked; courtesy lamps turned off. Yet others give a direct reflection of an individual's behaviour: speed, braking, signs of tiredness, eating or drinking while driving, or even having a row with the passengers. Locational data now include the movement patterns of both the car itself and the driver and any passengers.

The composite of data available in this domain allows very specific profiles to be put together on individual people, from travel behaviour and driving style right up to an estimate of personality traits. In addition, aggregated driving conduct data can be used to make a great many predictions about individual driving behaviour and the driving behaviour of particular types of road users: older drivers, smokers, young men, immigrants, vegetarians, the unemployed, glasses-wearers, etc. It is important to distinguish between two types of profiles: (1) an individual profile consisting of *historic personal data*, and (2) an individual or group profile consisting of *predictions* of future behaviour, based on aggregated datasets of a multitude of vehicles, drivers, passengers and the smart environment that now surrounds road users. These kinds of predictions may give rise to risk analyses, and to decisions, some of which will be automated. With the transition from 'connected driving' to 'autonomous driving', we will find ourselves increasingly having to do with automated decisions based on the braking behaviour of a self-driving car that is 'in conversation' with other vehicles to avoid a collision. Decisions of this type will not be directly relevant to the privacy of drivers or passengers, as long as they do not concern evaluating a driver or passenger or forming a judgement on driving capability, risk perception or risky conduct. However, as soon as we have to do with evaluations of people, this bears upon several fundamental rights: privacy, but also non-discrimination and possibly even the presumption of innocence. An example would be the fluctuation of insurance premiums based on predicted risky driving conduct, or estimates of the personal character of the person involved (algorithms declaring a tendency to impulsiveness) which will permeate across to data brokers, who in turn sell them on to employers or investigative services.

This is not science fiction, although we are still in the very early days of these and similar developments [Viereckl 2016].

Predictive profiles of this kind, continually adapted to continuous data streams, are giving rise to a new sort of 'transparency'. People are becoming 'transparent' in the sense that the profiles which exist on them allow a look, penetrating through them as individuals, at (statistically) similar people and allow numerous automated decisions to be taken on that basis. Here, too, a case in point would be supplements or rebates on insurance premiums, cancellation or suspension of driving licences, overriding of drivers' speed or braking, and even the monitoring of individuals on the basis of predicted hazards. The latter infringes more than just individuals' autonomy; it creates an entirely new choice architecture, with citizens' and consumers' behaviour increasingly being consciously steered, limited and tweaked by others. Citizens and consumers are as yet blissfully unaware of this development; they have little idea of this 'steer' [Sunstein 2016], [Yeung 2017].

The privacy conundrum, then, is thrown into sharp relief by the fact that, while people's conduct is becoming increasingly machine-predictable, the consequences that their conduct will entail for themselves are increasingly hard to foresee. This is giving rise to a very essential kind of uncertainty and a problematic kind of potential manipulability. In particular, were CAD data — perhaps via data brokers — to find their way into other contexts (the employment relationship, medical records or police files), privacy would seem altogether a thing of the past.

As to types of solutions: *Predictability and transparency of profiling*. As stated above, predictions can be derived either from a person's machine-read driving style or from aggregated data on many drivers. Although in the latter case profiles are *not* inherently related to a particular person and therefore *not* personal data in and of themselves, the application of them to a person who 'fits' the profile does fall under the fundamental right to data protection. It is of great importance to be able to predict what automated decisions might be made on the basis of this kind of profiling, and to obtain an impression of the logic used to spin off profiles. As seen in Part I of this position paper, we really ought not to be expecting Big Data or machine learning to work miracles for us. Consequently, it is essential that applied artificial intelligence is testable and, insofar as decisions based upon it have a major impact upon the lives of the data subjects, that those decisions can also be challenged at law. Both the current EU privacy legislation and the forthcoming GDPR provide a right to transparency of profiling, with top priority accorded to the obligation to inform data subjects of profiling and to explain how they are being profiled. Now that the GDPR, in a departure from current legislation, provides for an enforcement regime with teeth, it seems that this transparency of profiling is a crucial type of solution to the issue.


Predictability and purpose limitation. It is also important that citizens and consumers have a good idea of the purposes for which their CAD data are being

processed. Some are of the opinion that purpose limitation is an outdated concept because, supposedly, it hampers the rise of Big Data and machine learning. Only the boundless collection and analysis of CAD data would unlock the potential of such data to add value, we are told, precisely because that lack of limitation would allow an n=all scenario, with all relevant data to hand. Yet this is far from the case; as we saw earlier in our definition of Big Data and machine learning terminology, effective data analysis requires the careful consideration of all kinds of trade-offs, with the purpose itself of the analysis being the decisive factor in methodological choices. Moreover, that purpose is what enables us to keep the pressure on those wishing to profit from the data and, where necessary, to hold them accountable, given the foreseeable likely use of those data.

Case study 6: Information collation across locations

A significant trend in service provision is contextual adaptation of services. Factors at play here include users' general profiles, the devices they are using, and also their locations. Where fixed devices are being used, locational data can lead to price discrimination (higher prices are calculated for users in countries or neighbourhoods with a more expensive cost of living); in more extreme cases, they can even determine whether or not a service is available (e.g. Netflix or BBC iPlayer). Mobile devices allow for more exact geolocation, enabling the display of restaurants, shops or fuel stations in the vicinity. More sophisticated services are also enabled by this geolocation, such as the provision of traffic information, advertisements or discount codes that depend on the user's location, or a notification of which of the user's contacts are in the area. There is no doubting the high commercial premium that locational information provides, when coupled with the user's response to these services. By the same token, however, locational data can be very sensitive: they enable people, for example, to work out the user's home and workplace address, but also social class, religious affiliation and information on state of health. Research has found that locational information is not only being collected at large scale by the major Internet players but is also being processed by mobile app developers and by companies that offer advertising software to those developers. Even where the username is not saved, locations form a unique identifier; a 2013 study, for example, demonstrated that just four locational points and times were sufficient to identify 95% of users [de Montjoye 2013].

In addition to legal protection, one of the requirements is that users expressly give consent for the processing of their data, it is possible to protect locational data by technical means when providing location-dependent services [Shokri 2011]. In the most extreme form, a user's location can be suppressed; using techniques such as private information retrieval, information can be elicited from a database without the service provider knowing what information has been requested. Users can obscure their location by giving a location that is out by a couple of hundred



It's Google. They say you left your keys in the left-hand pocket of your other pants.

yards or by a mile or two: the greater the degree of obscurity, the more privacy, but the less useful the services obtained. Another solution is to add to one's profile some locations where one has never been; this is fine if one is looking for restaurants, but less apt if one is looking to see which of one's friends are in the vicinity. Finally, users can reduce the accuracy of the locational data they are giving off by telling the service provider that one is in a particular area (e.g. within a two-kilometre radius of downtown Brussels) without giving any further details.

Which of these approaches is best will depend on the particular service being used. When using a wired connection, one can use Tor to hide one's IP address,

and thereby also one's location. Tor is a network developed to obfuscate analysis of a user's Internet traffic by others to determine the origin and destination of data packets. While Tor offers protection from website owners and telcos, it does not hide one from government agencies that can avail themselves of a large number of interception points. Orbot offers similar functionality for Android devices, but this does not mean that one's location is really concealed. Turning to practical cases, some services offer an option to withhold one's location from others, but it is all but impossible to keep one's location truly secret. A company like Apple or Google will know the physical location of a huge number of Bluetooth and WiFi networks, and can use that to determine the user's location — even if WiFi or GPS services are disabled on the device. The only way to avoid this way of being geolocated is to jailbreak an iPhone, or root an Android-running phone, and then to install tools for the purpose. This does, however, make the user more vulnerable to hacking. Similarly, network operators are able to work out which 2G, 3G or 4G cell a user is located in, particularly in urban areas. In this case, the only³ solution is to put one's phone on flight mode or switch it off, but of course by doing so one is going without most of the functions the device offers.

Our conclusion is that locational data are particularly sensitive data, even though in practice it is nearly impossible to mask one's location from the major Internet and telecom players while using a mobile device [de Montjoye 2013], [Shokri 2011].

4. Conclusions and recommendations tailored to particular groups

The growing use of social media, cloud computing, the Internet and smartphones, and the power of machine learning when trained on ICT companies' and major public bodies' Big Data, is giving rise to an entirely new paradigm for the privacy of all ages of the population. Although we are still seeing rapid developments, this evolution calls for a coherent body of measures to be taken by those responsible for data processing, and for better awareness-raising among the general public.

4.1 Conclusions

What can be done to tackle all these problems? The control of personal data requires first and foremost that the individual has a grasp of the use and misuse of those data, and a genuine freedom of choice. These are prerequisites for the protection of fundamental rights, more particularly the fundamental right to protection of personal data. Nevertheless, we must face up to the fact that the ramifications of Big Data and machine learning are complex and often hard to retain in view, so

³ There are also legal guarantees limiting the use of locational data by operators [Wet op de Elektronische Communicatie].

that equipping individuals to make well-informed choices might seem an illusion. The entities that make hay from crunching Big Data are, therefore, going to have to shoulder their responsibility and will need to be obliged by law to do so. After all, one cannot expect an entity to conduct itself ethically if this were to mean that they would be squeezed out of the market; we need a level playing field, one providing data protection by default.

Certainly, we should be looking to a wide range of ICT methods to offer solutions, such as cryptography, and to other forms of data security, such as anonymisation and pseudonymisation. There is also a task for education: inculcating data and advertising sense/literacy from a tender age, and indeed for the mature generations. Furthermore, an approach of empowerment and resistance is called for. As a final consideration, all of the foregoing needs to be robustly embedded in legislation: both the previously-mentioned data protection regulation at EU level (regulating transparency, purpose limitation, criminal law and competences of the security services, data minimisation and the approach taken to data separation) and private-law enforcement. Laws must be attuned to the social and technological context, and must take due account of the paucity of knowledge among individuals [Council of Europe 2017].

The concentration of data in the hands of the big players of the ICT domain of the Internet and social media, and the powerful profiling techniques and business models that they practise, necessitate international regulation and active watchdogs with sufficient resources to enforce their oversight.

The responsible processing entities play a key role in weighing up the probable impact of their intended processing against the basic rights and fundamental freedoms of those involved. EU legislation stipulates in this regard that they identify and evaluate the risks of each processing operation involving Big Data and the potential adverse effect upon the rights of the individual and upon fundamental freedoms, in the interests of the right to protection of personal data, having a care for the right to non-discrimination, and bearing in mind social and ethical consequences. Risks can be contained by means of fitting measures, such as privacy by design and privacy by default. Finally, responsible parties will have to ascertain and approve the effectiveness of the solutions, under the watchful eye of the national data protection authorities.

The process of fortifying users/citizens/consumers so that they can have a good grip on their own situation and on their environment, by obtaining control, honing their critical awareness and fostering participation, is known as empowerment. One of the greatest challenges in this regard is how to render comprehensible and meaningful the complexity of technological aspects that impact upon data, metadata and privacy, such as algorithms, APIs, machine learning, default privacy settings, and AB testing. Education and training can assist people of all ages to

get to grips with the implications of the use of their personal data in Big Data contexts. For this reason, schools and educational institutions must start regarding information literacy and digital literacy as a core educational skill.

More profound public awareness is needed particularly as regards what is known as metadata and the very intimate profiling that they enable. Locational data call most prominently for attention here. There are next to no services on the market for users to mask their location from suppliers of services, the big players of the Internet world and the mobile operators without losing access to information relevant to where they are.

Finally, since responsibilities are divided across a range of actors, we require public debate and multi-stakeholder consultation involving all affected bodies. A key responsibility is borne by government and the industry, but always in close consultation with regulatory bodies, scientists, civil society bodies and citizens. Only in this way can a kind of 'collaborative responsibility' be attained in order to reach effective, broadly-supported solutions that will maintain and reinforce public values and fundamental rights in the digital age.

4.2 Recommendations

The following recommendations have been drawn up with a view to the roles of the citizen, government, ICT companies and the designers and researchers of ICT services. The recently-published report of the Belgian Privacy Commission [CBPL 2017] discusses Big Data privacy issues at length and contains 33 worthy recommendations, focused largely on those in organisations and companies who are responsible for processing data and on government bodies in light of the GDPR. Below, we incorporate and amplify recommendations 2, 8, 11, 12, 17 and 23 of that report. We also provide four recommendations of our own: for citizens, education, ICT service designers and government.

Recommendation 1: Responsibilities. Privacy as regards Big Data is an issue for citizens, engineers, consumers, companies, public bodies, media organisations and government institutions. This does not detract from the greater responsibility borne by big players who profit from Big Data analysis, nor from the fact that ultimate responsibility for ensuring respect for human rights — and thus for ensuring that the right efforts are made at the right levels — rests with government. What is called for, then, is for sufficient resources to be placed at the disposal of regulators, particularly to watch companies whose business model draws upon the analysis of Big Data (an area in which, incidentally, government bodies themselves are expected to lead by example). What is also called for is privacy curricula to be written for education, with encouragement of and attention paid to best practices and products. Consumer bodies and other civil society organisations have an important role to play in this matter: not just in making available comparisons of

Big Data and Internet service providers' privacy terms and conditions but also, for example, in exercising mandated rights of subject access requests and of resistance.

Recommendation 2: Alert citizens. Citizens whose data are processed must seek to exercise their GDPR-given rights to the full. Control of personal data is predicated upon individuals having sight of the use and misuse of those data, for only then can there truly be said to be freedom of choice. As this is a particularly tall order for individuals on their own, we recommend that those affected avail themselves of the opportunity to exercise their claims by means of mandating consumer or privacy bodies (Art. 80, GDPR). This approach will prove to be an effective way of contacting the party responsible for data processing (for purposes such as subject access requests or the right to have a copy or impose limitations on one's data) when the available information on the use of people's data is insufficiently transparent, incomplete or excessively vague, or when people seek to overturn a particular conclusion drawn from their data (e.g. fraud profiling).

Recommendation 3: Predictability, transparency of profiling, purpose limitation. Although a profile itself is *not* inherently related to a particular person and therefore does *not* constitute personal data, its application to a person who 'fits' the profile is covered by the fundamental right to data protection (GDPR). The right to transparency on profiling implies the obligation to inform the data subjects of profiling and to explain to them how they are being profiled. This is more far-reaching than correlations or statistical connections. In addition, it is vital that citizens and consumers have a sound impression of the aims for which their personal data are processed.

Recommendation 4: Addressing the imbalance of power. If the party responsible for an ICT service makes use of consent for the use of personal data, that consent must be straightforward to withdraw again and must always be time-limited. Moreover, such consent will not be valid where there is a manifest imbalance of power between the data subject and the responsible party or data processor, such as when the responsible party is the dominant (or monopoly) service in its market. The responsible party should have to demonstrate the absence of an imbalance of power, or that any imbalance cannot affect the data subject's ability to consent.

Recommendation 5: Designers of ICT and IoT devices must become serious about developing technologies that safeguard privacy and that offer transparency to the end user, such as by building in privacy by design, whereby privacy is reckoned with as a key requirement from the commencement of the design process onwards rather than being 'tacked on' as an afterthought. Service providers must allow users to compose services of mixed origin. *Algorithm programmers* must write their algorithms such as to guarantee user privacy. *App programmers* must permit transparency and make efforts to deliver efficient, effective technologies, with

users able to give their consent for the use of their data. App certification also needs to become a priority, so that users can be assured that the apps they use are secure. For both devices and apps, there is a great need for better, more granular ways of entering settings and preferences. Maximum privacy 'by default' must be the rule here. There is also a great deal of further *regulation* to be enacted, to ensure standardisation (without putting a dampener on innovation) and to form a basis for certification.

Recommendation 6: Role of government and business. It is the duty of government and business to make a fundamental assessment for each Big Data solution of whether the benefits outweigh the risks for the protection of personal data and for society as a whole (what if the data are leaked?). In each such assessment, it must be pondered whether or not the same objective could be accomplished using fewer data or less data aggregation. This ensues from the basic principles of EU legislation: data protection by design and data protection by default. In light of the plethora of leaks of data libraries, it is time for the Data Protection Authorities to invoke their new powers to impose effective solutions that feature data protection by design, with data remaining local as much as possible and stored centrally as little as possible.

Recommendation 7: Avoiding undesirable data bias. Designers and vendors responsible for services must continually consider whether improper or unfair data bias, algorithm bias or output bias are lurking in the datasets used to train their algorithms. Such bias might reside either in the mathematical models themselves or in the output (i.e. indirect discrimination). Questions needing to be addressed will include: Why are particular sections of the population excluded? Which data points are less well visible in training or tests of algorithms? Discrimination-aware data mining will help here.

Recommendation 8: Setting limits to government use of Big Data. The use of Big Data in the public sector — both in detection of tax and benefit fraud and in the realm of national security, the fight against crime and the maintenance of law and order — must continually be subject to the scrutiny of appropriate oversight bodies. This scrutiny must major on lawfulness and the concomitant concept of proportionality, which in turn will require a continual testing of the margin of effectiveness. It is crucial that a legal framework be put in place to lay down how and when the outcome of data mining and statistical analyses (correlations) made by government can be used as evidence in court proceedings in individual cases (e.g. fraud or public order offence trials), and when not.

Recommendation 9: Setting up a digital clearing house. It is recommendable that a Digital Clearing House (DCH) be set up to watch over the quality of the various regulators active in the digital market.

Recommendation 10: The role of education. With a particular eye to minors, education has a duty to discharge in terms of awareness-raising, attitude-setting, skills training and modelling conduct in the specific domains in which children and young people live: home, school, friendship groups (e.g. youth associations), etc. It is important to brief young people on the pitfalls of their digital behaviour, such as those expressed by the privacy paradox.

References

[Angwin 2016] J. Angwin e.a., 'Machine Bias: There's Software Used Across the Country to Predict Future Criminals. And It's Biased Against Blacks', *ProPublica*, 23 May 2016, zie www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

[AVG 2016] Algemene Verordening Gegevensbescherming, Regulation EU 2016/679. Aangenomen op 27 april 2016 en treedt in werking op 25 mei 2018. <http://eur-lex.europa.eu/legal-content/NL/TXT/?uri=CELEX%3A32016R0679>. Engelse versie zie [GDPR, 2016].

[Ballon 2016] P. Ballon, *Smart Cities: Hoe technologie onze steden leefbaar houdt en slimmer maakt*, Tiel: Lannoo, 221.

[Berendt and Preibusch 2014] B. Berendt and S. Preibusch, 'Better Decision Support through Exploratory Discrimination-Aware Data Mining: Foundations and Empirical Evidence', *Artificial Intelligence and Law* 22, no. 2 (1 June 2014): 175–209, doi:10.1007/s10506-013-9152-0.

[Belmans et al 2016] R. Belmans, P. Vingerhoets, I. Van Vaerenbergh et al, *The Central Role of End Users in the Energy Transition*, KVAB Standpunt 44 b, 2017.

[Cabitza 2016] F. Cabitza, 'The Unintended Consequences of Chasing Electric Zebras', IEEE SMC Interdisciplinary Workshop HUML 2016, The Human Use of Machine Learning, 12/16/ 2016, Venice, Italy https://www.researchgate.net/publication/311702431_The_Unintended_Consequences_of_Chasing_Electric_Zebras

[CBPL 2017] Commissie voor de bescherming van de persoonlijke levenssfeer, Big Data Rapport AH-2016-0154, <https://www.privacycommission.be/nl/publieke-consultatie-big-data-rapport>

[Commission EU Parliament] Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions: 'Building a European Data Economy', www.euractiv.com/wp-content/uploads/sites/2/2016/12/data-communication.pdf.

[Council of Europe 2017] Council of Europe - Guidelines on the Protection of Individuals with Regard to the Processing of Personal Data in a World of Big Data <http://rm.coe.int/CoERMPublicCommonSearchServices/DisplayDCTMContent?documentId=09000016806e7a>

[Coursera] <https://www.coursera.org/learn/friends-money-bytes/lecture/CKluM/selling-ad-spaces-through-auctions>

[DCS] Digital clearinghouse DCS https://secure.edps.europa.eu/EDPSWEB/edps/site/mySite/Big_data_rights_Lets_get_together

[De Hert 2011] P. De Hert, & R. Bellanova, R., 'Mobility should be fun. A consumer (law) perspective on border check technology', *The Scientific World JOURNAL*, 2011, vol. 11, 490-502.

[Demetriou 2016] S. Demetriou, W. Merrill, W. Yang, A. Zhang, C. A. Gunter: *Free for All! Assessing User Data Exposure to Advertising Libraries on Android*, NDSS, 2016.

[deMontjoye2013] Y.-A. de Montjoye, C.A. Hidalgo, M. Verleysen, V. D. Blondel, *Unique in the Crowd: The privacy bounds of human mobility*. Nature Scientific Reports 3, 2013.

[DG Int.Pol. 2015] Big Data and Smart Devices and Their Impact on Privacy. Study for the LIBE Committee. Directorate General for Internal Policies, Policy Department Citizen's Rights and Constitutional Affairs. [http://www.europarl.europa.eu/RegData/etudes/STUD/2015/536455/IPOL_STU\(2015\)536455_EN.pdf](http://www.europarl.europa.eu/RegData/etudes/STUD/2015/536455/IPOL_STU(2015)536455_EN.pdf)

[Digimeter 2016] Digimeter report 2016 <http://www.imec-int.com/en/digimeter>

[Digitale meters 2017] <https://www.vlaanderen.be/nl/nbwa-news-message-document/document/09013557801c194c>

[Dual use] <http://ec.europa.eu/trade/import-and-export-rules/export-from-eu/dual-use-controls/>

[EU Data protection directive 1995] http://ec.europa.eu/justice/policies/privacy/docs/95-46-ce/dir1995-46_part1_en.pdf

[Facebook] <https://www.facebook.com/business/help/430291176997542>

[Gadamer] Gadamer in *Wahrheit und Methode*, Bateson over 'the difference that makes a difference', Wolpert's 'no free lunch theorem', Gigerenzer over heuristieken, Schauer's work on stereotyping.

[Gadamer 2010] H. G. Gadamer, *Gesammelte Werke: Band 1: Hermeneutik I: Wahrheit und Methode: Grundzüge einer philosophischen Hermeneutik, 7.*, durchges. A. edition, Tübingen: Mohr Siebeck, 2010.

[GDPR, 2016] General Data Protection Regulation, (GDPR) Regulation EU 2016/679. Aangenomen op 27 april 2016 en treedt in werking op 25 mei 2018. <http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32016R0679>. Nederlandse versie zie [AVG 2016].

[Google] <https://support.google.com/adwords/answer/2996564?hl=nl>

[Hildebrandt 2016] M. Hildebrandt, 'Data-gestuurde intelligentie in het strafrecht, Preadvies Nederlandse Juristen Vereniging 2016', in: E.M.L. Moerel e.a., *Homo Digitalis* (Handelingen Nederlandse Juristen-Vereniging 2016-I), Den Haag: Wolters Kluwer 2016, p. 137-240, zie ook <http://njv.nl/preadviezen/preadviezen-2016/>

[Hildebrandt 2017] M. Hildebrandt, 'Wat weet mijn auto nog meer? Juridische bescherming *by design* in tijden van Internet van de Dingen', *Ars Aequi*, februari 2017, 97-102.

[Juniper routers trapdoor] https://en.wikipedia.org/wiki/Dual_EC_DRBG

[Mitchell, 1997] T. Mitchell, *Machine Learning*, McGraw Hill, 1997.

[NAP 2016], E. Grumbling, *Privacy Research and Best Practices: Summary of a Workshop for the Intelligence Community*, The National Academies Press, 2016.

[Perera 2015] C. Perera, R. Ranjan, L. Wang, S. U. Khan, A. Y. Zomaya, 'Big Data Privacy in the Internet of Things Era', *IEEE IT Professional Magazine: Special Issue Internet of Anything 2015*, Issue No.03 - May-June, 2015 vol.17.

[Post-Crash Voertuig Diagnose] <http://www.p-crashvd.nl>. Licentiehouders van een aantal merkspecifieke diagnosesystemen, zoals die van ODIS (Volkswagen, Audi, Skoda en Seat), VCDS (VAG-COM), BMW ISTA (BMW en Mini) en BMW Keyreader, Mercedes Xentry (Mercedes personenauto's en lichte bedrijfswagens), Volvo VIDA (Volvo personenauto's).

[PRISM] [https://en.wikipedia.org/wiki/PRISM_\(surveillance_program\)](https://en.wikipedia.org/wiki/PRISM_(surveillance_program))

[Rathenau Inst. 2010] Rathenau Instituut, 'Databases. Over ICT-beloofes, informatiehonger en digitale autonomie' http://www.cs.ru.nl/B.Jacobs/PAPERS/Rapport_Databases_Rathenau_Instituut_nov_2010.pdf

[Royal Society] The Royal Society, 'Progress and research in cybersecurity; supporting a resilient and trustworthy system for the UK?', 2016. <http://royalsociety.org/cybersecurity>

[Shokri 2011] R. Shokri, G. Theodorakopoulos, J.-Y. Le Boudec, J.-P. Hubaux, 'Quantifying Location Privacy', *IEEE Symposium on Security and Privacy*, 2011: 247-262.

[Smart metering in EU 2014] Smart metering deployment in the European Union <http://ses.jrc.ec.europa.eu/smart-metering-deployment-european-union>

[Sunstein 2016] C.R. Sunstein, *The ethics of Influence. Government in the Age of Behavioral Science*, Cambridge University Press, 2016.

[Sweeny 2013] <https://arxiv.org/ftp/arxiv/papers/1301/1301.6822.pdf>

[Tene, Polonetsky 2015] O. Tene, J. Polonetsky, 2015, 'A theory of creepy: Technology, privacy, and shifting social norms', *Yale Journal of Law and Technology* 16.1:2 (2015), <http://digitalcommons.law.yale.edu/yjolt/vol16/iss1/2/>

[UK Information Commissioner 2017] UK Information Commissioner, 'Big Data, Artificial Intelligence, Machine Learning and Data Protection', <http://iconewsblog.wordpress.com/2017/03/03/ai-machine-learning-and-personal-data/>

[van Dijck 2016] J. van Dijck, T. Poell, & M. de Waal, 2016, *De platformsamenleving: strijd om publieke waarden in een online wereld*, Amsterdam: Amsterdam University Press, 180.

[Vanrykel 2016] E. Vanrykel, G. Acar, M. Herrmann, C. Diaz: 'Leaky Birds: Exploiting Mobile Application Traffic for Surveillance', in: *Financial Cryptography and Data Security – 20th International Conference, FC 2016, Lecture Notes in Computer Science 9603*, Springer-Verlag.

[Verdonck, Van Hulle 2017] M. Van Hulle, P. Verdonck et al, 'Data Science and Healthcare', KVAB Standpunt 48 b, 2018.

[Verheul et al. 2016] E. Verheul et al., 'Polymorphic Encryption and Pseudonymisation for Personalised Healthcare', 2016, <https://eprint.iacr.org/2016/411>

[Viereckl 2016] R. Viereckl et al., 'Connected Car Report 2016: Opportunities, Risk, and Turmoil on the Road to Autonomous Vehicles' (PWC), <http://www.strategyand.pwc.com/reports/connected-car-2016-study>.

[VN2014] UN Department of Economic and Social Affairs, 'World Urbanisation Prospects'. United Nations, New York, 2014 revision, p. 1, <http://esa.un.org/unpd/wup/Highlights/WUP2014-Highlights.pdf> .

[Wearable] <https://www.wearable.com/internet-of-things/whos-watching-your-smartwatch>

[website-breaches] <http://www.informationisbeautiful.net/visualizations/worlds-biggest-data-breaches-hacks>

[website ik beslis] <https://www.ikbeslis.be>

[Wet Elektronische Communicatie] Wet Elektronische Communicatie, Art 123. http://www.ejustice.just.fgov.be/cgi_loi/change_lg.pl?language=nl&la=N&cn=2005061332&table_name=wet

[Wolpert 2013] D. H. Wolpert, 'Ubiquity Symposium: Evolutionary Computation and the Processes of Life: What the No Free Lunch Theorems Really Mean: How to Improve Search Algorithms', *Ubiquity 2013*, no. December: 2:1–2:15, doi:10.1145/2555235.2555237.

[Working Party] Working Party on data protection and privacy, European advisory body, 'Opinion 2/2010 on online behavioural advertising', 2010 http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2010/wp171_en.pdf

[WRR, 2016] Wetenschappelijke Raad voor het Regeringsbeleid WRR in Nederland 2016, WRR-rapport 95: *Big Data in een vrije en veilige samenleving* <http://www.wrr.nl/publicaties/publicatie/article/big-data-in-een-vrije-en-veilige-samenleving/>

[WRR 2017] Wetenschappelijke Raad voor het Regeringsbeleid WRR in Nederland WRR-Policy Brief 6: *Big Data and Security Policies: Serving Security, Protecting*

Freedom 2017 <https://www.wrr.nl/publicaties/publicaties/2017/01/31/big-data-and-security-policies-serving-security-protecting-freedom>

[Yeung 2017] K. Yeung, "Hypernudge": Big Data as a Mode of Regulation by Design', *Information, Communication & Society* (20) 2017, afl. 1, p. 118-136.

Members of working group

Yolande BERBERS (KTW, KU Leuven)

Willem DEBEUCKELAERE (Commissie voor de bescherming van de persoonlijke levenssfeer)

Paul DE HERT (VUB, Tilburg University TILT)

Yvo DESMEDT (KTW, University of Texas, University College London)

Frank DE SMET (Commissie voor de bescherming van de persoonlijke levenssfeer)

Jos DUMORTIER (Time.Lex)

Mireille HILDEBRANDT (University of Nijmegen, VUB)

Eleni KOSTA (Tilburg University TILT)

Karolien POELS (JA, UAntwerp)

Jo PIERSON (VUB)

Yves POULLET (Centre de Recherche Information, Droit et Société CRID, FUNDP Namur)

Bart PRENEEL (ESAT COSIC, KU Leuven)

Joos VANDEWALLE (KTW, KU Leuven)

Karel VELLE (KMW, State Archives, Ghent University)

KTW = Class of Technical Sciences

KMW = Class of Humanities

JA = Jonge Academie (The Young Academy)

RECENT POSITION PAPERS (from 2014)

27. Giovanni Samaey, Jacques Van Remortel e.a. – *Informaticawetenschappen in het leerplichtonderwijs*, KVAB/Klasse Technische wetenschappen en Jonge Academie, 2014.
28. Paul Van Rompuy – *Leidt fiscale autonomie van deelgebieden in een federale staat tot budgettaire discipline?* KVAB/Klasse Menswetenschappen, 2014.
29. Luc Bonte, Paul Verstraeten e.a. – *Maatschappelijk verantwoord ondernemen. Meedoen omdat het moet, of echt engagement?* KVAB/Klasse Technische wetenschappen, 2014.
30. Piet Van Avermaet, Stef Slembrouck, Anne-Marie Simon-Vandenberghe – *Talige diversiteit in het Vlaams onderwijs: problematiek en oplossingen*, KVAB/Klasse Menswetenschappen, 2015.
31. Jo Tollebeek – *Metamorfoses van het Europese historisch besef, 1800-2000*, KVAB/Klasse Menswetenschappen, 2015.
32. Charles Hirsch, Erik Tambuyzer e.a. – *Innovative Entrepreneurship via Spin-offs of Knowledge Centers*, KVAB/Klassen Natuurwetenschappen en Technische wetenschappen, 2015.
33. Georges Van der Perre en Jan Van Campenhout (eds.) – *Higher education in the digital era. A thinking exercise in Flanders*, Denkersprogramma KVAB/Klasse Technische wetenschappen, 2015.
34. Georges Van der Perre, Jan Van Campenhout e.a. – *Hoger onderwijs voor de digitale eeuw*, KVAB/Klasse Technische wetenschappen, 2015.
35. Hugo Hens e.a. – *Energiezuinig (ver)bouwen: geen rechttoe rechtaan verhaal*, KVAB/Klasse Technische wetenschappen, 2015.
36. Marnix Van Damme – *Financiële vorming*, KVAB/Klasse Menswetenschappen, 2015.
37. Els Witte – *Het debat rond de federale culturele en wetenschappelijke instellingen (2010-2015)*, KVAB/Klasse Menswetenschappen, 2015.
38. Irina Veretennicoff, Joos Vandewalle e.a. – *De STEM-leerkracht*, KVAB/Klasse Natuurwetenschappen en Klasse Technische wetenschappen, 2015.
39. Johan Martens e.a. – *De chemische weg naar een CO₂-neutrale wereld*, KVAB/Klasse Natuurwetenschappen, 2015.
40. Herman De Dijn, Irina Veretennicoff, Dominique Willems e.a. – *Het professoraat anno 2016*, KVAB/Klasse Natuurwetenschappen, Klasse Menswetenschappen, Klasse Kunsten en Klasse Technische wetenschappen, 2016.
41. Anne-Mie Van Kerckhoven, Francis Strauven – *Een bloementapijt voor Antwerpen*, KVAB/Klasse Kunsten, 2016.
42. Erik Mathijs, Willy Verstraete (e.a.), *Vlaanderen wijs met water: waterbeleid in transitie*, KVAB/Klasse Technische wetenschappen, 2016.
43. Erik Schokkaert - *De gezondheidszorg in evolutie: uitdagingen en keuzes*, KVAB/Klasse Menswetenschappen, 2016.
44. Ronnie Belmans, Pieter Vingerhoets, Ivo Van Vaerenbergh e.a. – *De eindgebruiker centraal in de energietransitie*, KVAB/Klasse Technische Wetenschappen, 2016.
45. Willem Elias, Tom De Mette – *Doctoraat in de kunsten*, KVAB/Klasse Kunsten, 2016.
46. Hendrik Van Brussel, Joris De Schutter e.a., *Naar een inclusieve robotsamenleving*, KVAB/Klasse Technische Wetenschappen, 2016.
47. Bart Verschaffel, Marc Ruyters e.a., *Elementen van een duurzaam kunstenbeleid*, KVAB/Klasse Kunsten, 2016.
48. Pascal Verdonck, Marc Van Hulle (e.a.) - *Datawetenschappen en gezondheidszorg*, KVAB/Klasse Technische wetenschappen, 2017.
- 48 b. Pascal Verdonck, Marc Van Hulle (e.a.) – *Data Science and Healthcare*, KVAB/Technical Science Class, 2018.

The complete list of position papers and all PDFs can be viewed at
www.kvab.be/en/position-papers