# Rathenau Instituut

# This is how we put AI into practice based on European values



**How much freedom should we allow the technology and the companies that are behind artificial intelligence (AI)? That question is drawing increasing attention. In recent months, a whole range of organisations published their own ethics codes for AI. It's now time for shared rules and legislation. That demands technological, social, and legal innovation, with respect for European values.**

*By Roos de Jong, Linda Kool, and Rinie van Est*

## Contents

## Introduction

This article was written in the run-up to Netherlands Digital Day on 21 March 2019. During this digital summit, parties from the business community, the scientific sector, government, and civil-society organisations will discuss further development of the National Digitalisation Strategy and the question of how we can give shape to the digital transition.

At the beginning of 2017, the Rathenau Instituut's *Urgent Upgrade* report drew attention to the wide range of social and ethical issues with which digitalisation is confronting us. Awareness of these issues is growing, partly due to the current interest in AI and ethics. AI is also the overarching theme for the digital summit.

This article positions AI within the broader digital transition. We prefer to speak of "AI systems". These are automated decision systems and therefore raise countless political questions, for example: who possesses the knowledge needed to take decisions? And who determines who possesses that knowledge?

It is now widely acknowledged that AI can have a profound impact on us as people. The many ethics codes developed by various parties over the last two years have contributed to this. After this crucial process of "evaluating" – in which goals, boundaries, and preconditions have been clarified – it is now time to move towards *shared* rules and legislation for the development and use of AI.

We make three suggestions in that regard:
- take standards and legislation as the basis;
- seek the link between innovation policy and values more emphatically, including within the national digitalisation strategy; and
- make the values perspective central to experimentation in actual practice.

## 1. What is AI?

The media show us new possibilities of AI every day, for example precision detection of cancer, large-scale tests of self-driving cars, and drones that may one day deliver our parcels. But we are confronted with the negative aspects too: elusive algorithms play a role in the distribution of disinformation, smart software systems turn out to discriminate, and smart devices in our homes prove to be easy to hack.

We often associate AI with futuristic applications or science fiction. But we've already been using many AI applications for a long time. Just take the examples of spam filters in your mailbox, online search engines, and the recommendations we get from Amazon or Netflix. Those applications are fully integrated into our daily lives and have been made possible by AI.

### 1.1. A short history of AI

AI isn't something new. As far back as the 1950s, scientists, mathematicians, and philosophers were exploring the concept of artificial intelligence. And the basic AI techniques used in **machine learning** and **deep learning** (see the box "Not all AI is the same") have also been known for a long time. Machine learning and deep learning have really gained momentum over the past two decades due to the increased processing power of computers and "big data". Where some very specific skills are concerned, AI systems can now outclass people. In 1997, IBM's chess-playing computer Deep Blue

defeated world champion Gary Kasparov, and in 2017 Google's AlphaGo beat the world's top Go player Ke Jie.

**AI is part of a system**

AI is not just a single technology but is better understood as a "cybernetic system" that observes, analyses (thinks) and acts, and can learn from doing so. Various technologies (sensors, big data, robot bodies, or an app) together ensure that in the right environment a system or machine can exhibit a certain degree of intelligent behaviour. It is the *combination* of ubiquitous devices and sensors connected to the Internet, that is capable, to a greater or lesser extent, of carrying out actions independently.

AI can therefore not be separated from other digital technologies, such as robotics, the Internet of Things, digital platforms, biometrics, virtual and augmented reality, persuasive technology, and big data. AI can be seen as the "brain" behind various "smart" applications, for example:

- a self-driving car is a combination of the Internet of Things and AI;
- Netflix recommendations are a combination of big data, AI, and a platform.

In the Rathenau Instituut's reports on the digital society, we therefore talk about intelligent machines and digitalisation, so as to emphasise that the issues relate to this broad cluster of technologies. In various reports we discuss the significance of applications such as self-driving cars, drones, smart healthcare applications and platforms.

**Not all AI is the same**

Discussions of AI in the media often start with visions of the future involving situations in which AI transcends human intelligence. That is also referred to as "**technological singularity**". But it's still a long way off. Google's AlphaGo has proven to be better at the game of Go than humans, but it can't suddenly play a different game. The same applies to the AI that can recognise skin cancer better than experienced dermatologists – it can't be used for any other task. This illustrates how most AI specialises in one specific skill.

In 2018, AlphaZero came along, a system that can play several games extremely well. Unlike AlphaGo, AlphaZero can win at chess, Go, and Shogi against professional human opponents. A more essential difference between AlphaGo and AlphaZero is in *how* they learn. AlphaGo was trained based on games played by humans, whereas AlphaZero learned by playing against itself.

Nowadays, a distinction is often made between "**rule-based AI**" and "**machine learning**". Rule-based AI is based on programmed "if this, then that" instructions. Such a system doesn't learn from itself but it does exhibit intelligent behaviour by analysing the environment and taking action – with a certain degree of autonomy – to achieve specific goals (see the European Commission's definition of AI).

Machine learning is based on detecting and learning from patterns in data. It involves developing software that improves its own performance, and it relies heavily on statistics**. Deep learning** is a type of machine learning based on neural networks (inspired by the biology of the brain). It involves combining weightings with input so as to classify and cluster that input.

Although public and political debate currently focuses mainly on forms of machine learning, it is important to understand that machine learning is only one type of AI. Machine learning and deep learning do not replace other types of AI; an AI system often involves a combination of multiple AI technologies, such as decision trees or logical reasoning.

## 1.2.  Smart or stupid AI: automated decision-making

AI systems are often good at a specific skill. In certain situations, however, the software does not deliver trustworthy results and is easy to fool. Changing just a few pixels or pasting a visibly distorting pattern over an image can cause a smart system to come up with a silly answer. Due to a bit of "noise", automatic image recognition can suddenly classify a panda as a gibbon (for a nice overview see this article (in Dutch) in the *Volkskrant* newspaper).

What seems simple for a human can be very difficult for an AI system. That a six-year-old child won't regard a panda as a gibbon and understands depth and shadows in 2D drawings is the result of a complex brain that has evolved over millions of years.

Various smart systems are dependent on the data they receive as input. If that data contains errors, then it's a matter of "garbage in, garbage out". Microsoft, for example, had to suspend its smart Taj chatbot within 24 hours when people uploaded racist remarks and Taj itself posted racist and inflammatory tweets. Biased data turns out to be a problem that is difficult to tackle. For example, Amazon had to suspend a system for selecting applicants automatically because it put women's applications at the bottom of the pile. The system was based on historical data (more men than women work in the IT industry), but it was not intended that it would utilise that fact as a selection criterion.

## 1.3.  Who is responsible for decisions taken by AI?

However smart or stupid an AI system may be, governments, companies, and other organisations are now deploying such systems for all kinds of purposes: driving cars, assisting doctors, checking creditworthiness, detecting fraud, and waging war. An important part of the discussion therefore focuses on the decisions these systems make. AI systems are decision systems, and that fact raises questions.

How do these systems arrive at a decision or a recommendation? Can that be checked, for example by regulators or by the person whom the decision concerns? To what extent is a human being involved in the decision? Who is responsible for the decision taken? As a doctor, will you soon have to provide an explanation if you haven't followed the recommendations generated by the system? And what about a motorist who doesn't follow the instructions of his smart car?

We are happy to take advantage of the convenience that digital services (using AI) can offer. But there is more behind those services. Collection and utilisation of data about our behaviour are omnipresent, and they are used to influence our behaviour. Harvard professor Shoshana Zuboff studies the business models and strategies of major technology companies. In her book *The Age of Surveillance Capitalism* (2019), she poses three crucial questions:

1. Who has the knowledge?
2. Who decides who has the knowledge?
3. Who decides who can decide who has the knowledge?

In the next section we discuss the social and ethical issues involved in AI and automatic decision-making, and in Section 3 we deal in greater detail with the above questions of power.

## 2.     From individual ethics codes to European frameworks

In the *Urgent Upgrade* report, we show that digitalisation involves a wide range of social and ethical issues (see Table 1). Besides privacy and security, issues such as control of technology, justice, human dignity, and unequal power relationships also play a role. In the *Human Rights in the Robot Age* report we describe how human rights such as non-discrimination and the right to a fair trial can come under pressure as a result of digitalisation (among other things by AI).

| Theme | Social and ethical issues |
|-------|---------------------------|
| Privacy | Data protection, digital inviolability of the home, mental privacy, surveillance, goal displacement |
| Autonomy | Freedom of choice, freedom of expression, manipulation (spreading disinformation, micro-targeting), protection of democracy, paternalism, skills, limits of self-reliance |
| Security | Information security, identity fraud, physical safety |
| Control over technology | Control over and insight into algorithms, responsibility, predictability |
| Human dignity | Dehumanisation, instrumentalisation, de-skilling, de-socialisation, unemployment |
| Justice | Discrimination, exclusion, equal treatment, stigmatisation |
| Power relationships | Unfair competition, exploitation, relationship consumer/company, relationship company/platform |

Over the past two years, public awareness of these issues has increased enormously. In the *Directed Digitalisation* report, we provide a detailed overview of initiatives by various parties. The discussion around AI plays a driving role in this.

There have been many initiatives within science, business, and government to look at AI from an ethical perspective. This has led to a wide range of ethics codes. We discuss some of them below.

- In January 2017, the Future of Life Institute (which is made up of researchers and technology companies and aims to prevent or reduce the risks posed by AI) organised the Asilomar Conference on Beneficial AI. The conference produced a

list of principles to which researchers in both academia and industry wish to commit themselves. Among other things, the 23 principles concern research strategies, for example, that the purpose of AI research is to develop beneficial and "not undirected" intelligence. As regards the purposes and behaviour of AI systems, it is stated that they must align with ideals of human dignity, data rights, freedoms, and cultural diversity. In addition, an arms race in lethal autonomous weapons should be avoided.

- In 2018, professional associations, for example the IEEE, and technology companies, such as Google and Telefónica, published statements and ethical guidelines on AI. Google refers to important issues such as interpretability, fairness, security, human-machine cooperation, and liability. According to these codes, AI should promote human well-being. The codes also state ambitions with regard to inclusiveness, fairness, reliability, explainability, security, responsibility, and accountability.

- In the Netherlands, the Platform for the Information Society (ECP) – made up of companies, government and civil-society organisations – released its *AI Impact Assessment tool* at the end of 2018. This takes nine values as the basic principle, based on the European advisory group's recommendations on AI, robotics, and autonomous systems (see below).

- There has also been attention to ethics and AI within government. In Europe, the European Group on Ethics in Science and New Technologies (EGE) published a *Statement on Artificial Intelligence, Robotics and Autonomous Systems*. In it, this ethics advisory organisation formulates a number of fundamental ethical principles and democratic requirements. These are based on the values laid down in European conventions and in the European Union's Charter of Fundamental Rights: human dignity, autonomy, responsibility, justice, equity and solidarity, democracy, the rule of law and accountability, security, safety, bodily and mental integrity, data protection and privacy, and sustainability.

- In June 2018, the European High-Level Expert Group on AI ("AI HLEG") – consisting of academics and representatives of companies and civil society – was set up to assist the European Commission in drawing up an AI strategy. The EU wishes to become the world leader in development and application of responsible, "human-centric" AI. In December 2018, the AI HLEG published its *Draft Ethics Guidelines for Trustworthy AI*; the definitive version is expected in April.

  According to the AI HLEG, two conditions must be met if AI is to be trustworthy:

  Firstly, it should respect fundamental rights, applicable regulations and core principles and values, ensuring an "ethical purpose". To ensure that people are central, the potential effects of AI on the individual and society must be assessed against social standards and values and five ethical principles: beneficence, non-

maleficence, autonomy, justice, and explicability.

Secondly, the technology must be robust and trustworthy. This concerns issues such as accountability, data governance, design for all, governance of AI autonomy (human oversight), non-discrimination, respect for human autonomy, respect for privacy, robustness, safety, and transparency.

## 2.1.  Consensus on AI and European values

The ethics codes that we have examined show an overlap as regards issues and values that can be affected by AI, with at the core central European values, as enshrined in human rights conventions and the *Charter of Fundamental Rights of the European Union*.

We also see that the broad spectrum of issues in the AI discussion is very similar to the issues identified in the *Urgent Upgrade* report. This shows that discussion of AI has indeed contributed to broadening public debate about digitalisation.

Ethics professors Floridi, Dignum et al. argue that many of the principles that have been formulated regarding the use of AI are very much in line with important bioethical principles, such as beneficence, non-maleficence, autonomy, and justice. Bioethics concerns technologies that intervene in the human body and brain. In its *Intimate Technology* and *Human Rights in the Robot Age* reports, the Rathenau Instituut has shown how digital technology also intervenes strongly in us as people: our thinking, social life, and action. Discussions on AI and ethics show that many parties now acknowledge that AI can intervene in our lives in a fundamental way.

Recognition of the fact that medical technology brings with it all kinds of bioethical issues has had a major impact on the way society deals with it, and on the role of regulation. Similarly, broad awareness that AI – and digitalisation in general – raise numerous bioethical issues will have a major impact on how we shape policy and monitor the development and use of IT systems. Ethics codes alone will not be enough; legislation will also be needed in order to make AI truly "human-centric" and trustworthy.

## 3.   The emergence of geopolitical debate

Our investigation of the various ethics codes reveals that one major social issue is often not mentioned, namely unequal power relationships. As we noted above, AI is about making decisions: who takes the decision, and who decides who may take that decision? In the past eighteen months, various questions of power have emerged as a new dimension of the AI debate.

### 3.1.   The AI race: techno-economic, socio-political, and military

Many people see AI as a general purpose technology that makes possible many new products and services, as a technology that will affect domains such as mobility, care, investigation, services, energy, education, labour, agriculture, and justice. AI is therefore also seen as the driver of economic growth and technological innovation, and the key to military technology.

On the economic level, dilemmas of economic and technological dependence and independence are involved. Winner-takes-all dynamics mean that companies in China and the US have access to huge quantities of data, which they utilise to improve their algorithms and their products and services. In this way, these companies are gaining an ever-increasing share of the market. Some large technology companies are repositioning themselves as AI companies, for example Facebook, Apple, Amazon, Netflix and Google in the US, and Baidu, Alibaba and Tencent in China. They dominate the AI market and often have a larger R&D budget than entire countries (total of public and private investment in R&D). According to Bloomberg, Amazon, for example, spent 22.6 billion dollars on R&D in 2017, while Alphabet spent 16.6 billion.

According to the United Nations, socio-political questions and implications for the international order are also involved. In China, AI contributes to the assessment and modification of citizens' behaviour. In liberal democracies, discussion has arisen about the role of AI and the influence on public and political debate, for example through political micro-targeting and the dissemination of disinformation. The international community is therefore talking about the protection of fundamental human rights and the sovereignty of the Internet: what role does the state have in protecting the Internet from foreign influences?

Finally, attention is being paid to the implications for defence and military power relationships. Some experts warn that the international consequences of AI are comparable to the impact of the development of nuclear weapons in the last century. At the Asilomar Conference on Beneficial AI, various scientists and companies agreed to prevent an arms race in autonomous weapon systems, but there is as yet no international ban. Should a particular state decide to include lethal autonomous weapons in its arsenal, other states may well feel compelled to follow. International

relations can also be set on edge with video or audio fragments faked using AI ("deep fakes", see this example).

## 3.2.   National AI strategies

The importance of these questions of power is now widely recognised. A number of countries have therefore published a national AI strategy. These strategies often express the ambition to remain, or become, a world leader in AI, for example those of the US and China (see the boxes below). Several European countries have also introduced a similar strategy, including France in March 2018. In his strategy, the French President, Emmanuel Macron, identifies European values as the starting point.

The European Commission (EC) followed in April 2018, with the EU prioritising European values as regards innovation and use of AI, thus adopting its own approach. The EC encourages each Member State to develop its own AI strategy. In the Netherlands, the "AI strategic action plan" is being prepared. This is expected before the summer.

### China – AI used for administrative purposes

With its *New Generation Artificial Intelligence Development Plan*, (July 2017), China expressed its aim to be the leading AI power by 2030. In this national AI strategy, the country designated development of the AI sector as a national priority. Pursuant to the strategy, AI technologies and applications "made in China" must be comparable to AI from the rest of the world by 2020. The same applies to Chinese companies and research facilities. By five years after that, breakthroughs in specific AI disciplines must have placed China in a leading position. In the final phase, China expects to become the world's leading centre for AI innovation.

**United States – AI driven by the private sector**

The White House has also made American leadership in AI a priority. For example, the Obama administration published a number of influential AI reports, including *Preparing for the Future of Artificial Intelligence*, the *National Artificial Intelligence Research and Development Strategic Plan* and *Artificial Intelligence, Automation, and the Economy*. The Trump administration has made AI an R&D priority, with a great deal of scope for American companies. In May 2018, representatives of industry, academia, and the government met for a summit on AI. This produced four central objectives: (1) strengthen the national R&D ecosystem, (2) support US workers in taking full advantage of the benefits of AI, (3) remove barriers to AI innovations, and (4) make possible high-quality, sector-specific AI applications. In February 2019, President Trump issued an executive order aimed at retaining America's position as world leader.

**Europe – trustworthy and human-centric AI**

Various EU countries have a national AI strategy. In April 2018, twenty five EU Member States signed a declaration that they would cooperate on the development of AI. This was followed by the European Commission's *Communication on Artificial Intelligence for Europe*, setting out a tripartite EU approach to AI: (1) Boost the EU's technological and industrial capacity and AI uptake across the economy, (2) prepare for socio-economic changes brought about by AI by encouraging the modernisation of education and training systems, nurturing talent, anticipating changes in the labour market, supporting labour market transitions and adaptation of social protection systems, and (3) ensure that there is an appropriate ethical and legal framework.

This third point distinguishes Europe from the USA and China. December 2018 saw the appearance of the *Coordinated Plan on Artificial Intelligence*, aimed at the EU becoming the world leader in the responsible development and application of AI. The European Commission is committed to "human-centric" AI, based on European standards and values. In addition to a finalised version of the *Ethics Guidelines for Trustworthy AI* (see above), a directive on product liability will also be published this year. Various websites provide a useful overview of national AI strategies; see for example the webpages of the OECD and the *Future of Life Institute*, and this article on Medium.

### 3.3.    Rule of law versus powerful companies

The AI race seems to mainly involve countries such as the United States and China. The EU positions itself on the world stage by focusing on a values-driven approach. At the same time, there are concerns about the power of countries, given the growing power of large technology companies. A unique concentration of power has arisen among the technology giants. According to Paul Nemitz, a strategic advisor to the European Commission, this is because these can make large-scale investments and they have increasing control of the infrastructures of public discourse. They can also collect personal data, create profiles, and dominate the development of AI services.

By formulating codes of ethics, technology companies are responding strategically to the many ethical and social issues that exist within society. That is important, but it does not solve all the problems. Firstly, parties cooperate within a chain, so whose code then in fact applies? Who oversees the various codes and coordinates them in the event of disparities between them? Joint action is therefore required. Parties such as the United Nations and the ISO are therefore working on collective standards.

Secondly, ethics codes *supplement* existing legislation. Social and ethical issues cannot be resolved only by drawing up codes; companies must also comply with existing legislation.

In the next section we look in more detail at how parties can innovate with AI based on European values.

## 4.     Towards values-driven innovation

Since AI touches on so many public values, the challenge is to shape innovation in a targeted way based on shared public values. In recent years the Rathenau Instituut has carried out a great deal of research into various aspects of values-driven innovation (see the reports *Valuable Digitalisation*, *Industry seeking University*, *Living Labs in the Netherlands*, and *Gezondheid Centraal* [Focus on Health]). In those reports, we make it clear that innovation policy is not only about developing new technological applications but also about the purpose which those applications serve, namely addressing the challenges facing society.

That is why values-driven innovation includes a focus on the development of suitable revenue models, appropriate legislation and regulations, and the social embedding of new applications. This approach to innovation recognises the complex nature of innovation at an early stage. We will discuss three ways in which innovation can be given shape on the basis of values.

### Take legislation as the starting point

The avalanche of ethics codes may give the impression that the development of AI is taking place in a legal vacuum. That is not the case, of course. There is a lot of existing legislation with which the development and use of AI must also comply. This involves not only fundamental rights, including constitutional rights, but also specific legislation, such as the EU's *General Data Protection Regulation*, or sector-specific legislation in fields such as healthcare or the transport market.

The recent history of digitalisation reveals that various IT platforms show little respect for the law. By dismissing existing legislation as obsolete, technology companies are attempting to evade various statutory responsibilities. This creates uncertainty regarding rights, obligations, and responsibilities (see also our report *Eerlijk delen* [A Fair Share]).

Courts are now clarifying matters in various legal cases, think of the ruling by the European Court of Justice on the responsibilities that Uber has. The Court ruled that Uber offers a transport service within the meaning of EU law. This means that the Member States are free to determine, at national level, the conditions subject to which that service may be provided. Regulatory bodies also play an important role in clarifying legal uncertainties.

Another factor is that the platforms often do not fit precisely into the existing legal categories. This leads to conceptual and policy uncertainty: is Facebook a social media platform, or a news company with the associated responsibilities? That is why it is often necessary to update existing legislation. The European Commission is currently preparing for revision of a large number of legal frameworks, including consumer law, copyright, audio-visual media, privacy, digital security, and competition law.

**Innovation policy: make public values central**

Greater attention has been paid in recent years to ethics in innovation policy. In June 2018, for example, the Dutch government published the national digitalisation strategy. This addresses numerous issues, including privacy, cybersecurity, and a fair data economy. The final section of the document concerns constitutional rights and ethics. The government is currently developing two "visions" on AI, a strategic AI action plan and a vision on constitutional rights and AI. The latter is a key component of the action plan.

It is important to see ethics not as a separate or final element of innovation programmes but as an integral part of them. The challenge is to make shared public values the starting point. Examples of this can be found in other European countries. In the area of mobility, for example, the United Kingdom drew up cybersecurity standards for self-driving cars at an early stage, as a basis for their development. Germany has drawn up ethical guidelines for the development of self-driving cars. The Netherlands, too, can shape and direct innovation by imposing preconditions in the fields of privacy, cybersecurity, transparency and other basic principles, for example in areas of experimentation or when regulators grant (temporary) permits ("regulatory sandboxes").

## 4.1. Actual practice: commit to technological and legal innovation at the same time

In actual practice, it turns out that technological innovation cannot be viewed separately from revenue models and regulations; these develop in tandem. For example, innovative cities such as Eindhoven and Amsterdam found themselves confronted by issues regarding the collection and use of sensor data within public space. Who has control of that data? What purposes can it be used for? How can a data monopoly be prevented? Amsterdam and Eindhoven therefore called for the development of national ground rules. A guide has since been produced.

In the healthcare context, too, a development can be identified in which innovation is embedded in a local care context involving doctors, patients, researchers, and developers. This benefits the quality of the new applications. The focus is no longer on the quantity of data but on its quality, and the higher purpose, namely improved health (see also our report *Gezondheid Centraal* [Focus on Health]).

## 5.    Conclusion

The increased attention for AI and ethics has broadened the public debate on digitalisation. Various actors have committed themselves to formulating and endorsing ethics codes and principles for the development and deployment of AI. It is good to see that the debate on values is being taken up internationally and that a great deal of attention is being paid to ethics. At the same time, however, this enormous attention paid to ethics threatens to stand in the way of follow-up steps. It's now time to deal with the practical aspects.

Individual ethics codes are just part of the story. Alignment and collective standards are also important. The movement to introduce such standards is now getting underway. In addition, individual codes of ethics do not constitute enforceable rules that have been legitimised by a democratic process. In a democracy based on the rule of law, legislation must play a central role in the development of AI. The role of legislation and the democratic process in achieving this is essential, certainly in view of changing geopolitical circumstances and the far-reaching power of large technology companies.

This means that the connection between innovation and public values must be pursued far more emphatically. In practice, that is happening more and more. Innovation means not only technological innovation but also social, economic, and legal innovation. In this article, we have presented various examples that show that such innovation must be – and can be – a matter of collaboration.

### Further reading
You can also read our publications on this subject:

- A Fair Share (2017)
- Urgent Upgrade (2017)
- Living labs in the Netherlands (2017)
- Directed digitalisation (2018)
- *Wethouders en raadsleden, durf te vragen*
  [City Executives and Councillors, Dare to Ask] (2018)
- Decent Digitisation (2018)
- Artificial Intelligence, what's New? (2018)
- Industry seeking University (2018)
- *Gezondheid Centraal* [Focus on Health] (2019)
- Overview of ethics codes and principles for AI (2019)

## Authors

Roos de Jong, Linda Kool, and Rinie van Est

## Preferred citation form:

Jong, R. de, L. Kool, and R. van Est (2019). *This is how we put AI into practice based on European values*. Rathenau Instituut: The Hague

## Open Access

The Rathenau Instituut has an Open Access policy. Reports, background studies, scientific articles, and software are published publicly and are available free of charge. Research data are made freely available, while respecting laws and ethical research norms regarding copyright, privacy and the rights of third parties.

## Contact details

Anna van Saksenlaan 51
PO Box 95366
NL-2509 CV The Hague
The Netherlands
+31 (0)70 342 1542
info@rathenau.nl
www.rathenau.nl

The Rathenau Instituut supports the formation of public and political opinion on the socially relevant aspects of science and technology. It conducts research and organises discussion of science, innovation, and new technologies.